

Spikingformer: A Key Foundation Model for Spiking Neural Networks

Chenlin Zhou^{1, 2}, Liutao Yu², Zhaokun Zhou¹, Han Zhang^{2,3}, Jiaqi Wang^{2,3},
Huihui Zhou², Zhengyu Ma^{2*}, Yonghong Tian^{1,2,4}

¹School of Electronic and Computer Engineering, Shenzhen Graduate School, Peking University

²Pengcheng Laboratory

³Harbin Institute of Technology

⁴School of Computer Science, Peking University

chenlinzhou25@stu.pku.edu.cn, mazhy@pcl.ac.cn, yhtian@pku.edu.cn

Abstract

Spiking neural networks (SNNs) offer a promising energy-efficient alternative to artificial neural networks, due to their event-driven spiking computation. However, some foundation SNN backbones (including Spikformer and SEW ResNet) suffer from non-spike computations (integer-float multiplications) caused by the structure of their residual connections. These non-spike computations increase SNNs' power consumption and make them unsuitable for deployment on mainstream neuromorphic hardware. In this paper, we analyze the spike-driven behavior of the residual connection methods in SNNs. We then present Spikingformer, a novel spiking transformer backbone that merges the MS Residual connection with Self-Attention in a biologically plausible way to address the non-spike computation challenge in Spikformer while maintaining global modeling capabilities. We evaluate Spikingformer across 13 datasets spanning large static images, neuromorphic data, and natural language tasks, and demonstrate the effectiveness and universality of Spikingformer, setting a vital benchmark for spiking neural networks. In addition, with the spike-driven features and global modeling capabilities, Spikingformer is expected to become a more efficient general-purpose SNN backbone towards energy-efficient artificial intelligence.

Code — <https://github.com/TheBrainLab/Spikingformer>

Introduction

Being regarded as the third generation of neural networks (Maass 1997), the brain-inspired Spiking Neural Networks (SNNs) are potential competitors to Artificial Neural Networks (ANNs) due to their high biological plausibility, event-driven property, and low power consumption on neuromorphic hardware (Roy, Jaiswal, and Panda 2019). In particular, the utilization of binary spike signals allows SNNs to adopt low-power accumulation (AC) instead of the traditional high-power multiply-accumulation (MAC), leading to significant energy efficiency gains and making SNNs increasingly popular (Chen et al. 2023).

As SNNs go deeper, their performance has improved significantly (Hu, Tang, and Pan 2021a; Fang et al. 2022;

Zheng et al. 2021; Hu et al. 2021). ResNet with a skip-connection has been extensively studied to extend the depth of SNNs (Fang et al. 2022; Zhou et al. 2023b). SEW ResNet (Fang et al. 2022), a representative Convolutional Neural Network (CNN) based SNN, easily implements identity mapping by Spike-Element-Wise (SEW) Residual connection and overcomes the problems of vanishing/exploding gradients of Spiking ResNet (Hu, Tang, and Pan 2021b). SEW ResNet is the first deep SNN directly trained with more than 100 layers. MS-ResNet (Hu et al. 2021) proposed the Membrane-based Shortcut (MS) Residual connection for spiking neural networks, effectively addressing gradient explosion and vanishing problems without performance degradation in CNN-based SNNs. Spikformer (Zhou et al. 2023b), a directly trained representative transformer-based SNN with SEW Residual connection, is proposed by leveraging both self-attention capability and biological properties of SNNs. It is the first successful exploration of applying flourishing transformer architecture into SNN design, and shows powerful performance. However, Spikformer faces the challenge of non-spike computations (integer-float multiplications) caused by the SEW Residual connection. This not only limits their ability to fully leverage the benefits of event-driven processing in energy efficiency, but also makes it difficult to deploy and optimize their performance on neuromorphic hardware (Chen et al. 2023).

We carry out an in-depth comparison between the SEW Residual connection and the MS Residual connection. Then, we creatively combine MS Residual and Self-Attention in a spike-driven way to address the challenge of non-spike computation in Spikformer while achieving superior performance. We name this model Spikingformer, in contrast to Spike-Driven Transformer (SD-Transformer) (Yao et al. 2023), which combines MS Residual with linear attention yet lacks global modeling capability, limiting its potential as a general-purpose SNN backbone. The main foundation SNN backbones are shown in Tab. 1. Spikingformer is the backbone that contains both spike-driven features and global modeling capabilities. Our contributions are as follows:

1) We conduct a systematic comparison between MS Residual connection and SEW Residual connection, analyzing their spike-driven behavior, firing rate, and performance.

2) We develop a novel spiking transformer model, named Spikingformer, which innovatively integrates MS Resid-

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Settings	SEW-ResNet	MS-ResNet	Spikformer	SD-Transformer	Spikingformer [†]
Test Resolution	224×224	224×224 / 288×288	224×224	288×288	224×224
Model Size	77.28M	78.37M	66.34M	66.34M	66.34M
Model Type	CNN	CNN	Transformer	Transformer	Transformer
Attention	No	No	SSA	SDSA	PSSA
Global Attention	✗	✗	✓	✗	✓
Residual Connection	SEW Residual	MS Residual	SEW Residual	MS Residual	MS Residual
Spike-Driven	✗	✓	✗	✓	✓
ImageNet-1K Acc (%)	69.26	74.21 / 76.02	74.81	77.07	77.64

Table 1: Comparison with other foundation backbones in the SNN domain, including SEW-ResNet (Fang et al. 2022), MS-ResNet (Hu et al. 2021), Spikformer (Zhou et al. 2023b), SD-Transformer (Yao et al. 2023).

ual and Self-Attention in a spike-driven way. Furthermore, leveraging both spike-driven computation and global attention modeling capabilities, Spikingformer is expected to serve as a foundation model for general-purpose energy-efficient artificial intelligence.

3) We validate Spikingformer on 13 datasets across static image classification, neuromorphic classification, and natural language understanding, establishing a vital experimental benchmark for the SNN community.

Related Work

The non-differentiability of spiking neurons in SNNs makes it difficult to train SNNs directly, but a common solution is to employ the surrogate gradient for backpropagation (Neftci, Mostafa, and Zenke 2019). Existing direct training SNNs can be roughly divided into two categories: convolution-based SNNs and transformer-based SNNs.

Convolution-based Spiking Neural Network

In the field of direct training, SNNs are unfolded over simulation time steps and trained with backpropagation through time (Lee, Delbruck, and Pfeiffer 2016; Shrestha and Orchard 2018). Due to the non-differentiability of spiking neurons, surrogate gradient method is employed for backpropagation (Lee et al. 2020; Neftci, Mostafa, and Zenke 2019). SEW ResNet (Fang et al. 2022) is a representative convolution-based SNN model by direct training, and is the first to increase the number of layers in SNNs to be larger than 100. However, the ADD gate in residual connections of SEW ResNet produces non-spike computations of integer-float multiplications in deep convolution layers. (Chen et al. 2023) has identified the problem of non-spike computations in SEW ResNet and Spikformer, and attempts to solve it through adding an auxiliary accumulation pathway during training and removing it during inference. This strategy needs tedious extra operations and results in a significant performance degradation compared with the original models.

Transformer-based Spiking Neural Network

Most existing SNNs borrow architectures from convolutional neural networks (CNNs), so their performance is lim-

ited by the performance of CNNs. The transformer architecture, originally designed for natural language processing (Vaswani et al. 2017), has achieved great success in many computer vision tasks, including image classification (Dosovitskiy et al. 2020; Yuan et al. 2021a), object detection (Carion et al. 2020; Zhu et al. 2020), and semantic segmentation (Wang et al. 2021; Yuan et al. 2021b). The structure of the transformer leads to a novel kind of SNN, with great potential to break through the bottleneck of SNNs’ performance. So far, two main related works: Spikformer (Zhou et al. 2023b) and Spikeformer (Li, Lei, and Yang 2022), have proposed spiking neural networks based on transformer architecture. Although Spikeformer replaces the activation function used in the feedforward layers with a spiking activation function, there are still a lot of non-spike operations remaining, including floating point multiplication, division, and exponential operation. Spikformer proposes a novel Spiking Self Attention (SSA) module by using spike-form Query, Key, and Value without softmax, and achieves state-of-the-art performances on many datasets. However, the structure of Spikformer with residual connection still contains non-spike computation. In our study, we innovatively integrates MS Residual and Self-Attention in a spike-driven way to address this problem.

Methods

Spiking Neuron Model

Spiking neuron is the fundamental unit of SNNs, we choose Leaky Integrate-and-Fire (LIF) model as the spike neuron in our work. The dynamics can be formulated as follows:

$$H[t] = V[t - 1] + \frac{1}{\tau} (X[t] - (V[t - 1] - V_{\text{reset}})), \quad (1)$$

$$S[t] = \Theta(H[t] - V_{th}), \quad (2)$$

$$V[t] = H[t](1 - S[t]) + V_{\text{reset}} S[t], \quad (3)$$

where τ is the membrane time constant, and $X[t]$ is the input current at time step t . When the membrane potential $H[t]$ exceeds the firing threshold V_{th} , the spiking neuron will trigger a spike $S[t]$. $\Theta(v)$ is the Heaviside step function, which equals to 1 when $v \geq 0$ and 0 otherwise. $V[t]$ represents the membrane potential after the triggered event, which equals to $H[t]$ if no spike is generated and otherwise equals to the reset potential V_{reset} .

Spike-Driven Behavior in SNN Residual Learning

There are mainly three types of residual connections in SNNs: Vanilla Residual (Zheng et al. 2021), SEW Residual (Fang et al. 2022), MS residual (Hu et al. 2021). However, Vanilla Residual connection suffers from performance degradation and gradient vanishing/exploding. Therefore, we mainly discuss the latter two residual connections.

The SEW Residual adopted in Spikformer. At present, Spikformer (Zhou et al. 2023b) is the representative work combining deep SNNs with transformer architecture. The residual learning plays an extremely important role in Spikformer, but the SEW Residual connections in Spikformer and SEW ResNet lead to non-spike computation (integer-float multiplications), which are not event-driven computations. As shown in Fig.1(a), the residual learning of Spikformer and SEW ResNet could be formulated as follows:

$$\begin{cases} \mathbf{O}_l = \text{SN}_l(\text{ConvBN}_l(\mathbf{O}_{l-1})) + \mathbf{O}_{l-1}, \\ \mathbf{O}_{l+1} = \text{SN}_{l+1}(\text{ConvBN}_{l+1}(\mathbf{O}_l)) + \mathbf{O}_l, \end{cases} \quad (4)$$

we denotes $\text{SN}_l(\text{ConvBN}_l(\mathbf{O}_{l-1}))$ as \mathbf{S}_l . This residual design inevitably brings in non-spike data and thus MAC operations in the next layer/block. In particular, \mathbf{S}_l and \mathbf{O}_{l-1} are spike signals, and their output \mathbf{O}_l are non-spike signal whose range is $\{0, 1, 2\}$. Non-spike data destructs event-driven computation in the next convolution layer when computing \mathbf{S}_{l+1} of \mathbf{O}_{l+1} . As the depth of the network increases, the range of non-spike data values transmitted to the deeper layer of the network will also expand. In our implementations of Spikformer, the range of the non-spike data could increase to $\{0, 1, 2, \dots, 16\}$ when testing Spikormer-8-512 on ImageNet 2012. Obviously, the range of non-spike data is approximately proportional to the number of residual blocks in Spikformer and SEW ResNet. In fact, integer-float multiplications are usually implemented in the same way as floating-point multiplication in hardware. In this case, the network will incur high energy consumption, approaching to the energy consumption of ANNs with the same structure, which is unacceptable for SNNs.

The MS Residual adopted in Spikingformer. Fig.1(b) shows the MS Residual connection adopted in Spikingformer. It could effectively avoid floating-point multiplications and integer-float multiplications, following the spike-driven principle. The spike-driven residual learning could be easily formulated as follows:

$$\begin{cases} \mathbf{O}_l = \text{ConvBN}_l(\text{SN}_l(\mathbf{O}_{l-1})) + \mathbf{O}_{l-1}, \\ \mathbf{O}_{l+1} = \text{ConvBN}_{l+1}(\text{SN}_{l+1}(\mathbf{O}_l)) + \mathbf{O}_l, \end{cases} \quad (5)$$

In this structure, we denote $\text{ConvBN}_l(\text{SN}_l(\mathbf{O}_{l-1}))$ as \mathbf{S}_l . $\mathbf{S}_l + \mathbf{O}_{l-1}$ belongs to the floating point addition operation, which is the same as the addition operation in SN layer. The floating point addition operation is the most essential operation of SNN. Obviously, the output \mathbf{O}_l is also floating point and will pass through an SN layer before participating in the next ConvBN computation. Therefore, the pure spike-form feature will be generated after the processing of SN layer, and the computation of ConvBN layer will become the floating point addition operation, following the spike-driven principle and reducing energy consumption vastly. In addition, we show the impact of SEW Residual and MS Residual

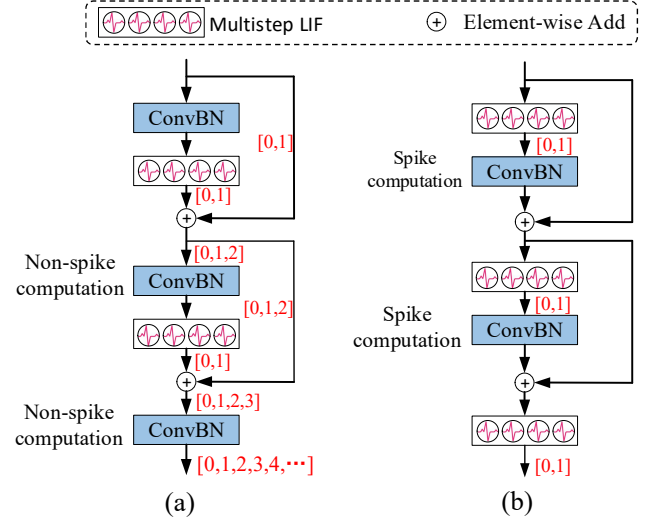


Figure 1: The residual learning in Spikformer and Spikingformer. (a) shows the SEW Residual learning of Spikformer, which contains non-spike computation (integer-float multiplications) in ConvBN layer. (b) shows the MS Residual connection, which is adopted in Spikingformer. MS Residual could effectively avoid integer-float multiplications, following the spike-driven principle.

on the firing behavior and performance of transformer-based SNNs, respectively, in the experimental section.

Spikingformer

Spikingformer family contains two foundation models: Spikingformer and Spikingformer[†]. Spikingformer is a novel and pure transformer-based spiking neural network through integrating spike-driven MS Residual blocks. Spikingformer[†] is a variant of Spikingformer. The pipeline of Spikingformer is shown in Fig.2. In this section, the details of Spikingformer are discussed.

Overall Architecture. Our proposed Spikingformer contains a Spiking Tokenizer (ST), several Spiking Transformer Blocks, and a Classification Head. Given a 2D image sequence $\mathbf{I} \in \mathbb{R}^{T \times C \times H \times W}$ (Note that $C=3$ in static datasets like ImageNet 2012, $C=2$ in neuromorphic datasets like DVS-Gesture), we use the Spiking Tokenizer block for downsampling and patch embedding, where the inputs can be projected as spike-form patches $\mathbf{X} \in \mathbb{R}^{T \times N \times D}$. Obviously, the first layer of Spiking Tokenizer also play a spike encoder role when taking static images as input. After Spiking Tokenizer, the spiking patches \mathbf{X}_0 will pass to the L Spiking Transformer Blocks. Similar to the standard ViT encoder block, a Spiking Transformer Block contains a Spiking Self Attention (SSA) (Zhou et al. 2023b) and a Spiking MLP block. In the last, a fully-connected-layer (FC) is used for the Classification Head. Note that we use a global average-pooling (GAP) before the fully-connected layer to reduce the parameters of FC and improve the classification

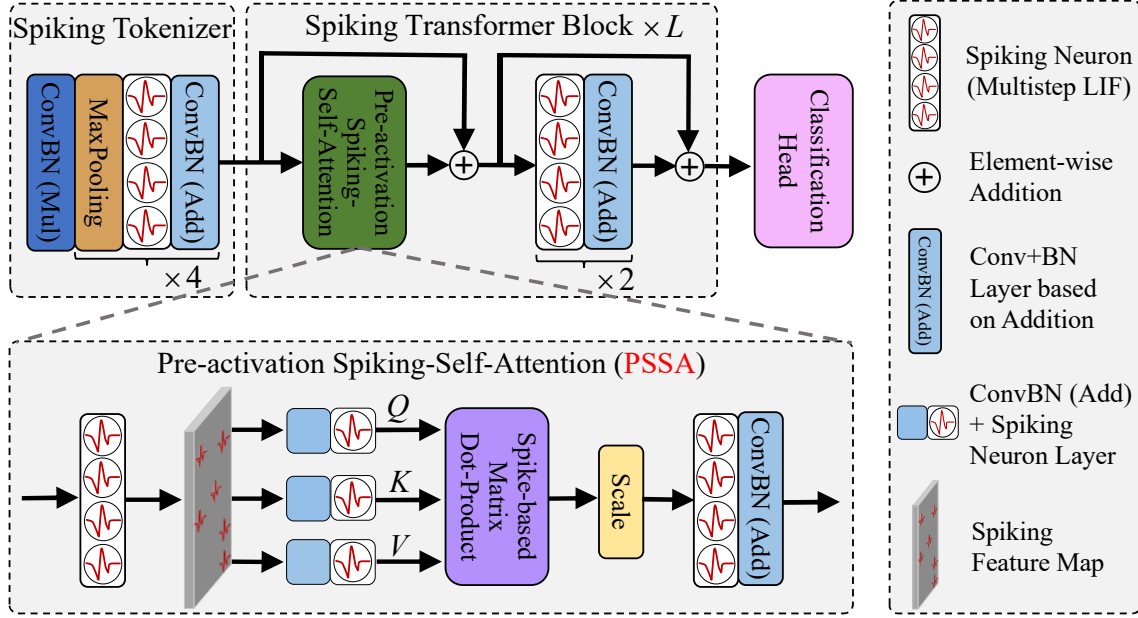


Figure 2: The overview of Spikingformer, which consists of a Spiking Tokenizer, several Spiking Transformer Blocks, and a Classification Head. Note that Multistep LIF is the Leaky Integrate-and-Fire (LIF) neuron model (Fang et al. 2022; Zhou et al. 2023b) with time steps $T > 1$. Same with Spikformer, T is an independent dimension for the spike neuron layer. In other layers, it is merged with the batch size. We use ConvBN to represent a convolution layer and its subsequent BN layer in this work.

capability of Spikingformer.

$$\mathbf{X} = \text{ST}(\mathbf{I}), \quad \mathbf{I} \in \mathbb{R}^{T \times C \times H \times W}, \mathbf{X} \in \mathbb{R}^{T \times N \times D} \quad (6)$$

$$\mathbf{X}'_l = \text{SSA}(\mathbf{X}_{l-1}) + \mathbf{X}_{l-1}, \mathbf{X}'_l \in \mathbb{R}^{T \times N \times D}, \quad (7)$$

$$\mathbf{X}_l = \text{SMLP}(\mathbf{X}'_l) + \mathbf{X}'_l, \mathbf{X}_l \in \mathbb{R}^{T \times N \times D}, \quad (8)$$

$$\mathbf{Y} = \text{FC}(\text{GAP}(\mathbf{X}_L)). \quad (9)$$

Spiking Tokenizer. As shown in Fig.2, Spiking Tokenizer mainly contains two functions: 1) convolutional spiking patch embedding, and 2) downsampling to project the feature map into a smaller fixed size. The spiking patch embedding is similar to the convolutional stream in Vision Transformer (Xiao et al. 2021; Hassani et al. 2021), where the dimension of spike-form feature channels gradually increases in each convolution layer and finally matches the embedding dimension of patches. In addition, the first layer of Spiking Tokenizer is utilized as a spike encoder when using static images as input. As shown in Eq. 10 and Eq. 11, the convolution part of ConvBN represents the 2D convolution layer (stride-1, 3×3 kernel size). MP and SN represent maxpooling (stride-2) and multistep spiking neuron, respectively. Eq. 10 is used for Spiking Patch Embedding without Downsampling (SPE), Eq. 11 is Spiking Patch Embedding with Downsampling (SPED). We could use multiple SPEs or SPEDs for specific classification tasks with different downsampling requirements. For example, we use 4 SPEDs for ImageNet 2012 dataset classification with input size as 224×224 (using 16 times downsampling). we use 2 SPEs and 2 SPEDs for CIFAR dataset classification with input size as 32×32 (using 4 times downsampling). After the

processing of the Spiking Tokenizer block, the input \mathbf{I} is split into an image patch sequence $\mathbf{X} \in \mathbb{R}^{T \times N \times D}$.

$$\mathbf{I}_i = \text{ConvBN}(\text{SN}(\mathbf{I})), \quad (10)$$

$$\mathbf{I}_i = \text{ConvBN}(\text{MP}(\text{SN}(\mathbf{I}))). \quad (11)$$

Spiking Transformer Block. A Spiking Transformer Block contains a Pre-activation Spiking Self-Attention (PSSA) block and a Spiking MLP block. PSSA retains Spiking Self-Attention (SSA)'s global modeling capabilities while making some modifications to be spike-driven and more generalized. Thus, PSSA can be seen as an important variant of SSA. The modifications include: 1) We change the spiking neuron layer position according to our proposed spike-driven residual mechanism, avoiding the multiplication of integers and floating-point weights. 2) To enhance generalization across diverse tasks, we choose ConvBN in place of LinearBN (linear layer and batch normalization) in Spikformer. The PSSA can be formulated as follows:

$$\mathbf{X}' = \text{SN}(\mathbf{X}), \quad (12)$$

$$\begin{cases} \mathbf{Q} = \text{SN}_Q(\text{ConvBN}_Q(\mathbf{X}')), \\ \mathbf{K} = \text{SN}_K(\text{ConvBN}_K(\mathbf{X}')), \\ \mathbf{V} = \text{SN}_V(\text{ConvBN}_V(\mathbf{X}')), \end{cases} \quad (13)$$

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{ConvBN}(\text{SN}(\mathbf{Q}\mathbf{K}^T\mathbf{V} * s)), \quad (14)$$

where $\mathbf{Q}, \mathbf{K}, \mathbf{V} \in \mathbb{R}^{T \times N \times D}$ are pure spike data (only containing 0 and 1). s is the scaling factor as in (Zhou et al. 2023b), controlling the large value of the matrix multiplication result. The Spiking MLP block consists of two SPEs,

Methods	Architecture	Param (M)	Train Size	Test Size	Time Step	Energy (mJ)	Top-1 Acc (%)
SEW ResNet	SEW ResNet-34	21.79	224 ²	224 ²	4	4.04	67.04
SEW ResNet	SEW ResNet-50	25.56	224 ²	224 ²	4	4.89	67.78
SEW ResNet	SEW ResNet-101	44.55	224 ²	224 ²	4	8.91	68.76
SEW ResNet	SEW ResNet-152	60.19	224 ²	224 ²	4	12.89	69.26
MS-ResNet	ResNet-104	78.37	224 ²	224 ²	5	-	74.21
MS-ResNet	ResNet-104	78.37	224 ²	288 ²	5	-	76.02
Spikformer	Spikformer-8-384	16.81	224 ²	224 ²	4	12.43	70.24
Spikformer	Spikformer-8-512	29.68	224 ²	224 ²	4	18.82	73.38
Spikformer	Spikformer-8-768	66.34	224 ²	224 ²	4	32.07	74.81
SD-Transformer	S-Transformer-8-384	16.81	224 ²	224 ²	4	3.90	72.28
SD-Transformer	S-Transformer-8-512	29.68	224 ²	224 ²	4	4.50	74.57
SD-Transformer	S-Transformer-8-768	66.34	224 ²	288 ²	4	6.09	77.07
ANN	Transformer-8-512	29.68	224 ²	224 ²	1	38.34	80.80
Spikingformer	Spikingformer-8-384	16.81	224 ²	224 ²	4	4.69	72.45
	Spikingformer-8-512	29.68	224 ²	224 ²	4	7.46	74.79
	Spikingformer-8-768	66.34	224 ²	224 ²	4	13.68	75.85
Spikingformer[†]	Spikingformer-8-384	16.81	224 ²	224 ²	4	5.61	74.35
	Spikingformer-8-512	29.68	224 ²	224 ²	4	8.68	76.54
	Spikingformer-8-768	66.34	224 ²	224 ²	4	16.30	77.64

Table 2: Results on ImageNet-1k classification. Power is calculated as the average theoretical energy consumption of an image inference on ImageNet, whose detail is shown in Eq.20. Same as Spikformer, our Spikingformer- L - D represents a Spikingformer model with L spiking transformer blocks and D feature embedding dimensions.

which are formulated in Eq.10. Spiking Transformer Block is shown in Fig.2, and it is the main component of Spikingformer.

Classification Head. We use a fully-connected-layer as the classifier behind the last Spiking Transformer Block. In detail, the classifier could be realized in four forms: AvgPooling-FC, SN-AvgPooling-FC, FC-AvgPooling, SN-FC-AvgPooling :

$$\mathbf{Y} = \text{FC}(\text{AvgPooling}(\mathbf{X}_L)), \quad (15)$$

$$\mathbf{Y} = \text{FC}(\text{AvgPooling}(\text{SN}(\mathbf{X}_L))), \quad (16)$$

$$\mathbf{Y} = \text{AvgPooling}(\text{FC}(\mathbf{X}_L)), \quad (17)$$

$$\mathbf{Y} = \text{AvgPooling}(\text{FC}(\text{SN}(\mathbf{X}_L))). \quad (18)$$

AvgPooling after FC (like SN-FC-AvgPooling, FC-AvgPooling) could be considered as computing the average of neuron firing, a post-processing of network, but in this way FC usually requires numerous parameters. AvgPooling before FC (like AvgPooling-FC, SN-AvgPooling-FC) could effectively reduce parameters compared with the previous ways. Only SN-FC-AvgPooling could avoid floating-point multiplication operation, but it needs more FC parameters than AvgPooling-FC or SN-AvgPooling-FC. In addition, it also hinders the classification ability of the network. In this work, we mainly adopt the way of AvgPooling ahead of FC, and choose AvgPooling-FC as the classifier of Spikingformer by default. Some experimental analysis on the classification head will be discussed in Appendix.

Spikingformer[†]. As a key variant, Spikingformer[†] is the

Spikingformer with ConvBN-MaxPool-LIF (CML) down-sampling (Zhou et al. 2023a). CML can overcome the imprecision problem of gradient backpropagation in Spikformer or Spikingformer, which improves network performance while reducing computational cost at the same time.

Theoretical Energy Consumption Calculation

The homogeneity of convolution allows the following BN and linear scaling transformation to be fused into the convolutional layer with an added bias when deployment (Ding et al. 2019, 2021; Hu, Tang, and Pan 2021b; Chen et al. 2023). Therefore, when calculating the theoretical energy consumption, the consumption of BN layers could be ignored. We calculate the number of synaptic operations of spikes before calculating the theoretical energy consumption.

$$SOP^l = fr \times T \times FLOPs^l, \quad (19)$$

where l is a block/layer in Spikingformer, fr is the firing rate of the block/layer, and T is the simulation time step of the spike neuron. $FLOPs^l$ refers to floating point operations of block/layer l , which is the number of multiply-and-accumulate (MAC) operations. And SOP^l is the number of spike-based accumulate (AC) operations. We estimate the theoretical energy consumption of Spikingformer according to (Kundu, Pedram, and Beerel 2021; Hu et al. 2021; Horowitz 2014; Kundu et al. 2021; Panda, Aketi, and Roy 2020). We assume that the MAC and AC operations are implemented on the 45nm hardware (Horowitz 2014), where $E_{MAC} = 4.6pJ$ and $E_{AC} = 0.9pJ$. The theoretical energy

Method	CIFAR10			CIFAR100			DVS128			CIFAR10-DVS		
	Param	T	Acc	Param	T	Acc	Param	T	Acc	Param	T	Acc
MS-ResNet	—	—	91.72	—	—	66.83	—	—	75.6	—	—	—
Spikformer	5.76	4	94.80	5.76	4	76.95	2.57	10	95.8	2.57	10	78.6
Spikformer	9.32	4	95.51	9.32	4	78.21	2.57	16	97.9	2.57	16	80.9
SD-Transformer	10.28	4	95.60	10.28	4	78.4	2.57	16	99.3	2.57	16	80.0
QKFormer	6.74	4	96.18	6.74	4	81.15	1.50	16	98.6	1.50	16	84.0
Transformer (ANN)	9.32	1	96.73	9.32	1	81.02	—	—	—	—	—	—
Spikingformer	5.76	4	95.22	5.76	4	78.34	2.57	10	96.2	2.57	10	79.9
	9.32	4	95.81	9.32	4	79.21	2.57	16	98.3	2.57	16	81.3
Spikingformer[†]	5.76	4	95.54	5.76	4	78.87	2.57	10	97.2	2.57	10	80.5
	9.32	4	95.95	9.32	4	80.37	2.57	16	98.6	2.57	16	81.4

Table 3: Comparison on CIFAR10, CIFAR100, DVS128 and CIFAR10-DVS. "Param" denotes "Parameter (M)", "Acc" denotes "Top-1 Accuracy (%)", and "T" denotes "Time Step".

consumption of Spikingformer can be calculated as follows:

$$E_{Snn} = E_{AC} \times \left(\sum_{i=2}^N SOP_{Conv}^i + \sum_{j=1}^M SOP_{SSA}^j \right) + E_{MAC} \times (FLOP_{Conv}^1), \quad (20)$$

Eq.20 shows the energy consumption of Spikingformer. $FLOP_{Conv}^1$ is the first layer encoding the non-spike input into spike-form. Then the SOPs of N SNN Conv layers and M SSA layers are added together and multiplied by E_{AC} .

Experiments

In this section, we carry out experiments on the static dataset ImageNet (Deng et al. 2009), CIFAR (Krizhevsky 2009) (including CIFAR10 and CIFAR100), the neuromorphic datasets (including CIFAR10-DVS and DVS128-Gesture (Amir et al. 2017)), and the natural language understanding tasks (GLUE) to evaluate Spikingformer.

ImageNet-1k Classification

We compared Spikingformer with Spiking ResNet (Hu, Tang, and Pan 2021a), SEW ResNet (Fang et al. 2022), MS-ResNet (Hu et al. 2021), Spikformer (Zhou et al. 2023b), SD-Transformer (Yao et al. 2023) on ImageNet-1k, which is shown in Tab. 2. Note that we recalculate the energy consumption of Spikformer in Appendix because the non-spike computation of Spikformer can not be directly calculated by Eq.20. Spikingformer-8-512 achieves 74.79% top-1 classification accuracy on ImageNet using 4 time steps, significantly outperforms Spikformer-8-512 by 1.41%, outperforms MS-ResNet model by 0.58% and outperforms SEW ResNet-152 model by 5.53%. Spikingformer-8-512 is with 7.463 mJ theoretical energy consumption, which reduces energy consumption by 60.36%, compared with 18.819 mJ of Spikformer-8-512. Spikingformer-8-768 achieves 75.85% top-1 classification accuracy on ImageNet using 4 time steps, significantly outperforms Spikformer-8-768 by 1.04%, outperforms the MS-ResNet model by

1.64% and outperforms SEW ResNet-152 model by 6.59%. Spikingformer-8-768 is with 13.678 mJ theoretical energy consumption, which reduces energy consumption by 57.34%, compared with 32.074 mJ of Spikformer-8-768. Spikingformer[†]-8-512 achieves 76.54% accuracy, which outperforms Spikformer-8-512 by 3.16% and outperforms SD-Transformer-8-512 by 1.97%. Spikingformer[†]-8-768 achieves 77.64% Top-1 classification accuracy, which outperforms Spikformer-8-768 by 2.83%. Compared with other foundation SNN backbones (SEW-ResNet, MS-ResNet, Spikformer, SD-Transformer), Spikingformer achieves the best performance in ImageNet-1k due to its spike-driven features and global modeling capabilities.

CIFAR and Neuromorphic Tasks

The results are shown in Tab. 3. We compared Spikingformer with MS-ResNet (Hu et al. 2021), Spikformer (Zhou et al. 2023b), SD-Transformer (Yao et al. 2023), QKFormer (Zhou et al. 2024).

CIFAR Classification. From the results, We find that the performance of Spikingformer models surpass all the models of Spikformer with the same number of parameters. In CIFAR10, our Spikingformer achieves 95.81% classification accuracy, significantly outperforms Spikformer by 0.30% and outperforms MS-ResNet-482 by 3.91%. Spikingformer[†] achieves 95.95% accuracy and outperforms SD-Transformer by 0.35%. In CIFAR100, Spikingformer achieves 79.21% classification accuracy, significantly outperforms Spikformer by 1.00% and outperforms MS-ResNet by 12.38%. Spikingformer[†] achieves 80.37% accuracy and outperforms SD-Transformer by 1.97%. Transformer (ANN) is only 0.69% and 1.00% higher than Spikingformer[†] on CIFAR10 and CIFAR100.

Neuromorphic Classification. We compare our method with SOTA methods on DVS-Gesture. In detail, we adopt four SPEDs in Spiking Tokenizer due to the 128*128 image size of CIFAR10-DVS and adopt 2 spiking transformer blocks with 256 patch embedding dimension. The number of time steps of the spiking neuron is 10 or 16. The number

Model	Energy (mj)	Time	MNLI _{m/mm}	QQP	QNLI	SST-2	CoLA	STS-B	MRPC _{F1}	RTE	Avg. Acc (%)
BERT _{base}	51.41	–	84.6/83.4	71.2	90.5	93.5	52.1	85.8	88.9	66.4	79.6
Q2BERT	–	–	47.2/47.3	67.0	61.3	80.6	0.0	4.7	81.2	52.7	49.1
ELMo	–	–	68.6/–	86.2	71.1	91.5	44.1	70.4	76.6	53.4	70.2
SpikeBERT	14.30	4	71.4/71.0	68.2	66.4	85.4	16.9	18.7	82.0	57.5	59.7
Spikingformer	6.76	4	71.9/72.5	84.7	76.0	87.2	24.4	54.5	79.7	55.6	66.8

Table 4: The results on the Natural Language Understanding task (GLUE datasets). "Avg. Acc" denotes "Average Accuracy".

of training epochs is 106, which is the same as Spikformer. The learning rate is initialized to 0.1 and decayed with a cosine schedule. The results of CIFAR10-DVS are shown in Tab.3. Spikingformer achieves 81.3% top-1 accuracy with 16 time steps and 79.9% accuracy with 10 time steps, significantly outperforms Spikformer by 0.4% and 1.3% respectively. Spikingformer[†] achieves 81.4% accuracy with 16 time steps and outperforms SD-Transformer by 1.4%. We compare our method with SOTA methods on CIFAR10-DVS in Tab.3. Spikingformer achieves 98.3% accuracy with 16 time steps and 96.2% accuracy with 10 time steps, outperforms Spikformer by 0.4% and 0.4% respectively.

Natural Language Understanding

We evaluate Spikingformer on the standard GLUE (General Language Understanding Evaluation) benchmark (Wang et al. 2018), which is a widely adopted collection of datasets designed to evaluate and advance natural language understanding (NLU) capabilities in machine learning models. GLUE contains 8 subsets for classification and regression, including single-sentence classification (CoLA, SST-2), pairwise sentence comparison (MRPC, QQP, RTE), and natural language inference (STS-B, MNLI, QNLI, WNLI). We pretrain Spikingformer on Wikipedia-English (Devlin et al. 2019) by masked language pretraining (Devlin et al. 2019), then finetune on GLUE dev set. The experimental results are shown in Tab. 4. We compare Spikingformer with BERT_{base} (Devlin et al. 2019), Q2BERT (Zhang et al. 2020), ELMo (Sarzynska-Wawer et al. 2021), SpikeBERT (Lv et al. 2023). Our Spikingformer achieves 66.8% average accuracy, which outperforms SpikeBERT by 7.1%.

Discussion

In this part, we conduct a systematic comparison between MS Residual connection and SEW Residual connection on Transformer-based SNN (corresponding to Spikingformer and Spikformer, respectively).

Spike Behavior Visualization. We visualize the spike behavior of spikingformer and spikformer in Fig. 3. The results show Spikingformer could effectively avoid integer-float multiplications common in Spikformer. In addition, fr in Fig.3(b) represents the firing rate of Spikingformer. We observe that Spikingformer have a lower firing rate on ImageNet compared with Spikformer (Fig.3(a)), which further reduces synaptic operations and thus energy consumption.

Energy Consumption and Performance Impact. The results can be seen in Tab 2. Compared to Spikformer, Spikingformer achieves higher performance while significantly reducing energy consumption. The primary reason for this

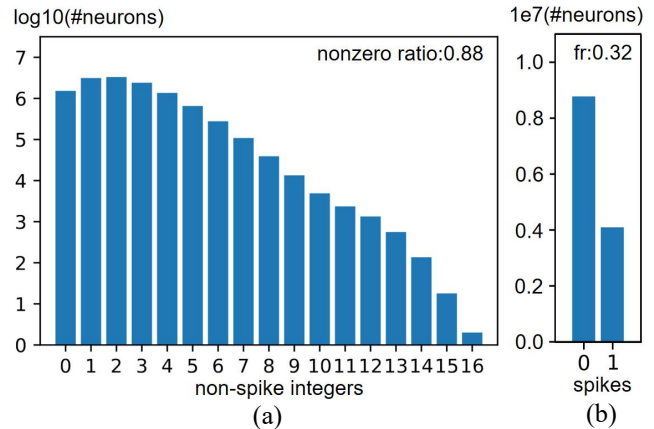


Figure 3: The spike-driven behavior of Spikingformer and Spikformer. (a) Histogram of the input data of block 7 in Spikformer-8-512. The abscissa means non-spike data range with $\{0, 1, 2, \dots, 16\}$ before Conv layer in the transformer block of Spikformer. The nonzero ratio indicates the ratio of non-zero input units. (b) Histogram of the input data of block 7 in Spikingformer-8-512. The abscissa means binary spike data with $\{0, 1\}$ before Conv layer in the transformer block of Spikingformer. The ordinate means of the neuron numbers of $\{0, 1\}$.

energy reduction is Spikingformer’s ability to effectively eliminate integer-float multiplications. Additionally, a lower firing rate on ImageNet further contributes to the improved energy efficiency.

Conclusion

In this work, we propose Spikingformer, a spike-driven transformer-based spiking neural network that innovatively integrates MS Residual connections with Self-Attention in a biologically plausible, spike-driven manner. This design effectively overcomes the non-spike computation limitations of existing spiking transformers, such as Spikformer, while maintaining global modeling capabilities. Extensive experiments across 13 diverse datasets, including large-scale image classification, neuromorphic data classification, and natural language understanding tasks, demonstrate that Spikingformer consistently outperforms previous foundation SNN backbones. With its efficient spike-driven computation and global modeling capacity, Spikingformer establishes itself as a robust, general-purpose SNN backbone and a key benchmark for the SNN community.

Acknowledgements

The study was funded by the National Natural Science Foundation of China under contracts No. 62425101, No. 62332002, No. 62027804, No.62088102, and the major key project of the Pengcheng Laboratory (PCL2025A02).

References

- Amir, A.; Taba, B.; Berg, D.; Melano, T.; McKinstry, J.; Di Nolfo, C.; Nayak, T.; Andreopoulos, A.; Garreau, G.; Mendoza, M.; Kusnitz, J.; Debole, M.; Esser, S.; Delbruck, T.; Flickner, M.; and Modha, D. 2017. A Low Power, Fully Event-Based Gesture Recognition System. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7243–7252.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 213–229. Springer.
- Chen, G.; Peng, P.; Li, G.; and Tian, Y. 2023. Training Full Spike Neural Networks via Auxiliary Accumulation Pathway. *arXiv preprint arXiv:2301.11929*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 248–255.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, 4171–4186.
- Ding, X.; Guo, Y.; Ding, G.; and Han, J. 2019. Acnet: Strengthening the kernel skeletons for powerful cnn via asymmetric convolution blocks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1911–1920.
- Ding, X.; Zhang, X.; Ma, N.; Han, J.; Ding, G.; and Sun, J. 2021. Repvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13733–13742.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*.
- Fang, W.; Yu, Z.; Chen, Y.; Huang, T.; Masquelier, T.; and Tian, Y. 2022. Deep Residual Learning in Spiking Neural Networks. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, volume 34, 21056–21069.
- Hassani, A.; Walton, S.; Shah, N.; Abuduweili, A.; Li, J.; and Shi, H. 2021. Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*.
- Horowitz, M. 2014. 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)*, 10–14. IEEE.
- Hu, Y.; Tang, H.; and Pan, G. 2021a. Spiking Deep Residual Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 1–6.
- Hu, Y.; Tang, H.; and Pan, G. 2021b. Spiking deep residual networks. *IEEE Transactions on Neural Networks and Learning Systems*.
- Hu, Y.; Wu, Y.; Deng, L.; and Li, G. 2021. Advancing residual learning towards powerful deep spiking neural networks. *arXiv preprint arXiv:2112.08954*.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images.
- Kundu, S.; Datta, G.; Pedram, M.; and Beerel, P. A. 2021. Spike-thrift: Towards energy-efficient deep spiking neural networks by limiting spiking activity via attention-guided compression. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 3953–3962.
- Kundu, S.; Pedram, M.; and Beerel, P. A. 2021. Hire-snn: Harnessing the inherent robustness of energy-efficient deep spiking neural networks by training with crafted input noise. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5209–5218.
- Lee, C.; Sarwar, S. S.; Panda, P.; Srinivasan, G.; and Roy, K. 2020. Enabling spike-based backpropagation for training deep neural network architectures. *Frontiers in neuroscience*, 14: 119.
- Lee, J. H.; Delbruck, T.; and Pfeiffer, M. 2016. Training deep spiking neural networks using backpropagation. *Frontiers in neuroscience*, 10: 508.
- Li, Y.; Lei, Y.; and Yang, X. 2022. Spikeformer: A Novel Architecture for Training High-Performance Low-Latency Spiking Neural Network. *arXiv preprint arXiv:2211.10686*.
- Lv, C.; Li, T.; Xu, J.; Gu, C.; Ling, Z.; Zhang, C.; Zheng, X.; and Huang, X. 2023. SpikeBERT: A Language Spikeformer Learned from BERT with Knowledge Distillation. *arXiv preprint arXiv:2308.15122*.
- Maass, W. 1997. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9): 1659–1671.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6): 51–63.
- Panda, P.; Aketi, S. A.; and Roy, K. 2020. Toward scalable, efficient, and accurate deep spiking neural networks with backward residual connections, stochastic softmax, and hybridization. *Frontiers in Neuroscience*, 14: 653.
- Roy, K.; Jaiswal, A.; and Panda, P. 2019. Towards Spike-based Machine Intelligence With Neuromorphic Computing. *Nature*, 575(7784): 607–617.
- Sarzynska-Wawer, J.; Wawer, A.; Pawlak, A.; Szymanowska, J.; Stefaniak, I.; Jarkiewicz, M.; and Okruszek, L. 2021. Detecting formal thought disorder by deep contextualized word representations. *Psychiatry Research*, 304: 114135.

Shrestha, S. B.; and Orchard, G. 2018. Slayer: Spike layer error reassignment in time. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, volume 31.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, volume 30.

Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.

Wang, W.; Xie, E.; Li, X.; Fan, D.-P.; Song, K.; Liang, D.; Lu, T.; Luo, P.; and Shao, L. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 568–578.

Xiao, T.; Singh, M.; Mintun, E.; Darrell, T.; Dollár, P.; and Girshick, R. 2021. Early convolutions help transformers see better. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, volume 34, 30392–30400.

Yao, M.; Hu, J.; Zhou, Z.; Yuan, L.; Tian, Y.; Xu, B.; and Li, G. 2023. Spike-driven transformer. *Advances in neural information processing systems*, 36: 64043–64058.

Yuan, L.; Chen, Y.; Wang, T.; Yu, W.; Shi, Y.; Jiang, Z.-H.; Tay, F. E.; Feng, J.; and Yan, S. 2021a. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 558–567.

Yuan, L.; Hou, Q.; Jiang, Z.; Feng, J.; and Yan, S. 2021b. Volo: Vision outlooker for visual recognition. *arXiv preprint arXiv:2106.13112*.

Zhang, W.; Hou, L.; Yin, Y.; Shang, L.; Chen, X.; Jiang, X.; and Liu, Q. 2020. Ternarybert: Distillation-aware ultra-low bit bert. *arXiv preprint arXiv:2009.12812*.

Zheng, H.; Wu, Y.; Deng, L.; Hu, Y.; and Li, G. 2021. Going Deeper With Directly-Trained Larger Spiking Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 11062–11070.

Zhou, C.; Zhang, H.; Zhou, Z.; Yu, L.; Huang, L.; Fan, X.; Yuan, L.; Ma, Z.; Zhou, H.; and Tian, Y. 2024. Qkformer: Hierarchical spiking transformer using qk attention. *arXiv preprint arXiv:2403.16552*.

Zhou, C.; Zhang, H.; Zhou, Z.; Yu, L.; Ma, Z.; Zhou, H.; Fan, X.; and Tian, Y. 2023a. Enhancing the performance of transformer-based spiking neural networks by SNN-optimized downsampling with precise gradient back-propagation. *arXiv preprint arXiv:2305.05954*.

Zhou, Z.; Zhu, Y.; He, C.; Wang, Y.; YAN, S.; Tian, Y.; and Yuan, L. 2023b. Spikformer: When Spiking Neural Network Meets Transformer. In *The Eleventh International Conference on Learning Representations*.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable DETR: Deformable Transformers for End-to-End Object Detection. *arXiv preprint arXiv:2010.04159*.