Fourier or Wavelet bases as counterpart self-attention in spikformer for efficient visual classification

Qingyu Wang^{1,2*}, Duzhen Zhang^{1,2*}, Tilelin Zhang^{1,2†}, and Bo Xu^{1,2}

¹ School of Artificial Intelligence, University of Chinese Academy of Sciences ² Institute of Automation, Chinese Academy of Sciences

{wangqingyu2022, zhangduzhen2019, tielin.zhang, xubo}@ia.ac.cn

Abstract. Energy-efficient spikformer has been proposed by integrating the biologically plausible spiking neural network (SNN) and artificial Transformer, whereby the Spiking Self-Attention (SSA) is used to achieve both higher accuracy and lower computational cost. However, it seems that self-attention is not always necessary, especially in sparse spike-form calculation manners. In this paper, we innovatively replace vanilla SSA (using dynamic bases calculating from Query and Key) with spike-form Fourier Transform, Wavelet Transform, and their combinations (using fixed triangular or wavelets bases), based on a key hypothesis that both of them use a set of basis functions for information transformation. Hence, the Fourier-or-Wavelet-based spikformer (FWformer) is proposed and verified in visual classification tasks, including both static image and event-based video datasets. The FWformer can achieve comparable or even higher accuracies (0.4%-1.5%), higher running speed (9%-51% for training and 19%-70% for inference), reduced theoretical energy consumption (20%-25%), and reduced GPU memory usage (4%-26%), compared to the standard spikformer. Our result indicates the continuous refinement of new Transformers, that are inspired either by biological discovery (spike-form), or information theory (Fourier or Wavelet Transform), is promising.

Keywords: spiking neural network · Transformer · visual classification

1 Introduction

Spiking neural network (SNN) is considered the third generation of artificial neural networks [27] for its biological plausibility of event-driven characteristics. It has also received extensive attention in the computation area of neuromorphic hardware [6], exhibiting a remarked lower computational cost on various machine learning tasks, including but not limited to, visual classification [50], temporal auditory recognition [36], and reinforcement learning [34]. The progress in SNN is contributed initially by some key computational modules inspired by the biological brain, e.g., the receptive-field-like convolutional circuits, self-organized plasticity propagation [47], and other multi-scale inspiration from the single neuron or synapse to the network or cognitive functions. Simultaneously, the SNN

also learns from the artificial neural network (ANN) by borrowing some mathematical optimization algorithms, e.g., the approximate gradients in backpropagation (BP), various types of loss definitions, and regression configurations.

Even though various advanced architectures have been proposed and contributed ANN to a powerful framework, the efforts to promote its training speed and computational consumption have never been stopped. As the well-known Transformer for example, it contains a rich information representation formed by multi-head self-attention, which calculates Query, Key, and Value from the inputs to connect each token in a sequence with every other token. Although having achieved rapid and widespread application, the $\mathcal{O}(N^2)$ complexity (with N representing the sequence length) results in a huge training cost in Transformer that can not be neglected. Many works have tried to solve this problem, including but not limited to, replacing self-attention with unparameterized transform formats, for example, using Fourier Transform (FNet [23]) or Gaussian Transform (Gaussian attention [45]). Another attempt is to integrate some key features of ANNs and SNNs to exhibit their advantages, such as the higher accuracy performance in ANNs and the lower computational cost in SNNs.

The spikformer [50] explores self-attention in SNN for more advanced deep learning. It introduces a spike-form self-attention called Spiking Self-Attention (SSA). In SSA, the floating Query, Key, and Value signals are sent to leakyintegrated and fire (LIF) neurons to generate spike sequences that only contain binary and sparse 0 and 1 vision information, which results in non-negativeness spiking attention map. This special map doesn't require the complex softmax operation anymore for further normalization, which means a lower computational consumption is needed compared to that in vanilla Self-Attention. However, even though many efforts have been made, it seems that the SSA still exhibits an $\mathcal{O}(N^2)$ complexity, whereby further refinement is necessary. Given binary and sparse spikes for information representation, we here question whether it is still necessary to retain the original complex structure of Self-Attention in spikformer. Here, we give a hypothesis that although Self-Attention with learning parameters has been generally considered more flexible, it is still not suitable in the spike stream context, since the correlation between sparse spike trains is too weak to form closed similarity. Hence, an intuitive approach is to convert these sparse spike trains in spatial domains to the equivalent frequency domains with the help of Fourier transformation.

Here we propose a new hypothesis: Just like the Fourier Transform, Self-Attention can also be thought of as using a set of basis frequency functions for information representation. The main difference between these two methods is that the Fourier Transform uses fixed triangular basis functions to transform signals into the frequency domain, while on the contrary, the Self-Attention calculates higher-order signal representation from compositions of the input to produce more complex basis functions ($Query \times Key$). This understanding may explain why FNet [23] performs well, since fixed basis functions may also work in some cases by offering structured prior information. Following this perspective, an intuitive plan is to integrate all these key features together, towards a reduced

computational cost and accelerated running speed, including unparameterized transforms (e.g., Fourier Transform and Wavelet Transform), and spike-form sparse representation. Our main contributions can be summarized as follows:

We propose a key hypothesis that the Self-Attention in Transformer works by using a set of basis functions to transform information from Query, Key, and Value sequences, which is very similar to the Fourier Transform. Hence, after jointly considering the shortcomings of spikformer, we replaced SSA with spike-form Fourier Transform and Wavelet Transform. Mathematical analysis indicates a reduced time complexity from $\mathcal{O}(Nd^2)$ or $\mathcal{O}(N^2d)$, to $\mathcal{O}(N \log N)$ or $\mathcal{O}(D \log D) + \mathcal{O}(N \log N)$, under the same accuracy performance.

The results validate that our method achieves superior accuracy on event-based video datasets (improved by 0.3%-1.2%) and comparable performance on spatial image datasets, compared to spikformer with SSA. Furthermore, it exhibits significantly enhanced computational efficiency, reducing memory usage by 4%-26%, reducing theoretical energy consumption by 20%-25%, and achieving approximately 9%-51% and 19%-70% improvements in training and inference speeds, respectively.

We further analyze the orthogonality of self-attention as a set of basis functions. We find during training, that the orthogonality is continuously decreasing, which inspires us to use combined different wavelet bases with nonlinear, learnable parameters as coefficients to form structured nonorthogonal basis functions. In the second round of experiments, the experiments show even better accuracy performance on event-based video datasets (improved by 0.4%-1.5% compared to spikformer).

2 Related Work

Vision Transformers The vanilla Transformer architecture, initially designed for natural language processing [35], has then demonstrated remarkable success in various other computer-vision tasks, including image classification [9], semantic segmentation [37], object detection [1], and low-level image processing [3]. The critical component that contributes to the success of the Transformer is the selfattention mechanism. In Vision Transformer (ViT), self-attention can capture global dependencies between image patches and generate meaningful representations by weighting the features of these patches, using the dot-product operation between Query and Key, followed by the softmax normalization [19]. The structure of ViT also fits for conventional SNNs, offering potential Transformer-type architectures for achieving higher accuracy performance.

Spiking Neural Networks In contrast to traditional ANNs that employ continuous floating-point values to convey information, SNNs utilize discrete spike sequences for communication, offering a promising energy-efficient and biologically plausible alternative for computation. The critical components of SNNs encompass spiking neuron models, optimization algorithms, and network architectures. Spiking neurons serve as the fundamental non-linear spatial and temporal information processing units in SNNs, responsible for receiving from continuous inputs and converting them to spike sequences. Leaky Integrate-and-Fire (LIF) [7], PLIF [11], Izhikevich [17] neurons are commonly used dynamic neuron models in SNNs for their efficiency and simplicity. There are primarily two optimization algorithms employed in deep SNNs: ANN-to-SNN conversion and direct training. In ANN-to-SNN conversion [32], a high-performance pre-trained ANN is converted into an SNN by replacing Rectified Linear Unit (ReLU) activation functions with spiking neurons. However, the converted SNN requires significant time steps to accurately approximate the ReLU activation, leading to substantial latency [14]. In direct training, SNNs are unfolded over discrete simulation time steps and trained using backpropagation through time [33]. Since the eventtriggered mechanism in spiking neurons is non-differentiable, surrogate gradients are employed to approximate the non-differentiable parts during backpropagation by using some predefined gradient values to replace infinite gradients [22].

With the advancements in ANNs, SNNs have improved their performance by incorporating advanced architectures from ANNs. These architectures include Spiking Recurrent Neural Networks [26], ResNet-like SNNs [16], and Spiking Graph Neural Networks [42]. Recently, exploring Transformer in the context of SNNs has received a lot of attention. For example, temporal attention has been proposed to reduce redundant simulation time steps [43]. Additionally, an ANN-SNN conversion Transformer has been introduced, but it still retains vanilla self-attention that does not align with the inherent properties of SNNs [29]. Furthermore, spikformer [50] investigates the feasibility of implementing selfattention and Transformer in SNNs using a direct training manner.

In this paper, we argue that the artificial Transformer can be well integrated into SNNs for higher performance, while at the same time, the utilization of SSA in spiking Transformer (spikformer) can be further replaced by a special module based on Fourier Transform or Wavelet Transform, which to some extent, indicating an alternative more efficient effort to achieve fast, efficient computation without affecting the accuracy.

3 Preliminaries

3.1 Spiking Neuron Model

The spiking neuron serves as the fundamental unit in SNNs. It receives the current sequence and accumulates membrane potential, which is subsequently compared to a threshold to determine whether a spike should be generated. In this paper, we consistently employ LIF at all Spiking Neuron Layers.

The dynamic model of the LIF neuron is described as follows:

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

$$S[t] = \mathcal{G} \left(H[t] - V_{th} \right), \tag{2}$$

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

where τ represents the membrane time constant, and X[t] denotes the input current at time step t. When the membrane potential H[t] exceeds the firing threshold V_{th} , the spiking neuron generates a spike S[t]. The Heaviside step function $\mathcal{G}(v)$ is defined as 1 when $v \geq 0$ and 0 otherwise. The membrane potential V[t] will transition to the reset potential V_{reset} if there is a spike event, or otherwise it remains unchanged as H[t].

3.2 Spiking Self-Attention

The spikformer utilizes the SSA as its primary module for extracting sparse visual features and mixing spike sequences. Given input spike sequences denoted as $\boldsymbol{X} \in \mathbb{R}^{T \times N \times D}$, where T, N, and D represent the time steps, sequence length, and feature dimension, respectively, SSA incorporates three key components: Query (\boldsymbol{Q}) , Key (\boldsymbol{K}) , and Value (\boldsymbol{V}) . These components are initially obtained by applying learnable matrices $\boldsymbol{W}_Q, \boldsymbol{W}_K, \boldsymbol{W}_V \in \mathbb{R}^{D \times D}$ to the input sequences \boldsymbol{X} . Subsequently, they are transformed into spike sequences through Spiking Neuron Layers, formulated as:

$$\boldsymbol{Q} = \mathcal{SN}(BN(\boldsymbol{X}\boldsymbol{W}_Q)), \boldsymbol{K} = \mathcal{SN}(BN(\boldsymbol{X}\boldsymbol{W}_K)), \boldsymbol{V} = \mathcal{SN}(BN(\boldsymbol{X}\boldsymbol{W}_V)), \quad (4)$$

where \mathcal{SN} denotes the Spiking Neuron Layer, BN denotes Batch Normalization and $\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V} \in \mathbb{R}^{T \times N \times D}$. Inspired by vanilla Self-Attention [35], SSA adds a scaling factor s to control the large value of the matrix multiplication result, defined as:

$$SSA(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V}) = S\mathcal{N} \left(\boldsymbol{Q} \ \boldsymbol{K}^{\mathrm{T}} \ \boldsymbol{V} * s \right),$$

$$\boldsymbol{X}' = S\mathcal{N}(BN(Dense(SSA(\boldsymbol{Q}, \boldsymbol{K}, \boldsymbol{V})))),$$
(5)

where $\mathbf{X}' \in \mathbb{R}^{T \times N \times D}$ are the updated spike sequences. It should be noted that SSA operates independently at each time step. In practice, T represents an independent dimension for the SN layer. In other layers, it is merged with the batch size. Based on Equation (4), the spike sequences \mathbf{Q} and \mathbf{K} produced by the SN layers SN_Q and SN_K , respectively, naturally have non-negative values (0 or 1). Consequently, the resulting attention map is also non-negative. Therefore, according to Equation (5), there is no need for softmax normalization to ensure the non-negativity of the attention map, and direct multiplication of \mathbf{Q} , \mathbf{K} and \mathbf{V} can be performed. This approach significantly improves computational efficiency compared to vanilla Self-Attention.

However, it is essential to note that SSA remains an operation with a computational complexity of $\mathcal{O}(N^2)$.³ Within the spike-form frameworks, we firmly believe that SSA is not essential, and there exist simpler sequence mixing mechanisms that can efficiently extract sparse visual features as alternatives.

³ Although SSA can be decomposed with an $\mathcal{O}(N)$ attention scaling, this complexity hides large constants, causing limited scalability in practical applications. For more detailed analysis, refer to **Time Complexity Analysis of FW vs. SSA** Section.

3.3 Fourier Transform

The Fourier Transform (FT) decomposes a function into its constituent frequencies. For an input spike sequence $\boldsymbol{x} \in \mathbb{R}^{N \times D}$ at a specific time step in \boldsymbol{X} , we utilize the FT to transform information from different dimensions, including 1D-FT and 2D-FT.

The discrete 1D-FT along the sequence dimension \mathcal{F}_{seq} to extract sparse visual features is defined by the equation:

$$\boldsymbol{x}_{n}' = \mathcal{F}_{\text{seq}}(\boldsymbol{x}_{n}) = \sum_{k=0}^{N-1} \boldsymbol{x}_{k} e^{-\frac{2\pi i}{N}kn}, n = 0, ..., N-1,$$
(6)

where *i* represents the imaginary unit. For each value of *n*, the discrete 1D-FT generates a new representation $\boldsymbol{x}'_n \in \mathbb{R}^D$ as a sum of all of the original input spike features $\boldsymbol{x}_n \in \mathbb{R}^{D}$.⁴

Similarly, the discrete 2D-FT along the feature and sequence dimensions $\mathcal{F}_{seq}(\mathcal{F}_{f})$ is defined by the equation:

$$\boldsymbol{x}_{n}' = \mathcal{F}_{\text{seq}}(\mathcal{F}_{\text{f}}(\boldsymbol{x}_{n})), n = 0, ..., N - 1.$$
(7)

Notably, equations (6) and (7) only consider the real part of the result. Therefore, there is no need to modify the subsequent MLP sub-layer or output layer to handle complex numbers.

3.4 Wavelet Transform

Wavelet Transform (WT) is developed based on Fourier Transform to overcome the limitation of Fourier Transform in capturing local features in the spatial domain.

The discrete 1D-WT along the sequence dimension W_{seq} to extract sparse visual features is defined by the equation:

$$\boldsymbol{x}_{n}' = \mathcal{W}_{\text{seq}}(\boldsymbol{x}_{n}) = \frac{1}{\sqrt{N}} \Big[\boldsymbol{T}_{\varphi}(0,0) * \varphi(\boldsymbol{x}_{n}) + \sum_{j=0}^{J-1} \sum_{k=0}^{2^{j}-1} \boldsymbol{T}_{\psi}(j,k) * \psi_{j,k}(\boldsymbol{x}_{n}) \Big], \quad (8)$$

$$T_{\varphi}(0,0) = \frac{1}{\sqrt{N}} \sum_{k=0}^{N-1} \boldsymbol{x}_{k} * \varphi(\boldsymbol{x}_{k}), \quad T_{\psi}(j,k) = \frac{1}{\sqrt{N}} \sum_{k'=0}^{N-1} \boldsymbol{x}_{k'} * \psi_{j,k}(\boldsymbol{x}_{k'}), \quad (9)$$

where n = 0, ..., N - 1, $N = 2^J$ (N is typically a power of 2), * denotes elementwise multiplication, $T_{\varphi}(0,0)$ are the approximation coefficients, $T_{\psi}(j,k)$ are the detail coefficients, $\varphi(x)$ is the scaling function, and $\psi_{j,k}(x) = 2^{j/2}\psi(2^jx - k)$ is

⁴ It is important to note that the weights in Equation (6) are fixed constant and can be pre-calculated for all spike sequences.

the wavelet function. Here, we use the Haar scaling function and Haar wavelet function for example, which is defined by the equation:

$$\varphi(x) = \begin{cases} 1 & 0 \le x < 1\\ 0 & \text{otherwise} \end{cases}, \quad \psi(x) = \begin{cases} 1 & 0 \le x < 0.5\\ -1 & 0.5 \le x < 1\\ 0 & \text{otherwise} \end{cases}$$
(10)

Similarly, the discrete 2D-WT along the feature and sequence dimensions $\mathcal{W}_{seq}(\mathcal{W}_{f})$ is defined by the equation:

$$\boldsymbol{x}_{n}' = \mathcal{W}_{\text{seq}}(\mathcal{W}_{\text{f}}(\boldsymbol{x}_{n})), n = 0, ..., N - 1,$$
(11)

In the subsequent experimental section, we also delve into the exploration of different basis functions as well as their potential combinations.

4 Method

Following a standard Vision Transformer architecture, the vanilla spikformer incorporates several key components, including the Spiking Patch Splitting (SPS) module, Spikformer Encoder Layers, and a Classification head for visual classification tasks. Here, we directly replace vanilla SSA head with the FW head to efficiently manage spike-form features.

In the following sections, we provide an overview of our proposed FW former in Figure 1, followed by a detailed explanation of the FW head. Finally, we compare the time complexity of both of these two heads.

4.1 Overall Architecture

We provide Figure 1 for an overview of our FWformer. First, for a given 2D image sequence $I \in \mathbb{R}^{T \times C \times H \times W}$,⁵ the goal of the Spiking Patch Splitting (SPS) module is to linearly project it into a *D*-dimensional spike-form feature and split this feature into a sequence of *N* flattened spike-form patches $P \in \mathbb{R}^{T \times N \times D}$. Following the approach of the vanilla spikformer, the SPS module employs convolution operations to introduce inductive bias [41].

Second, to generate spike-form Relative Position Embedding (RPE), the Conditional Position Embedding (CPE) generator [4] is utilized in the same manner as the spikformer. The RPE is then added to the patch sequence \boldsymbol{P} , resulting in $\boldsymbol{X}_0 \in \mathbb{R}^{T \times N \times D}$.

Third, the *L*-layer FW Encoder is designed to manage X_0 . Different from spikformer encoder layer with SSA head, our FW Encoder Layer consists of an FW sub-layer and an MLP sub-layer, both with batch normalization and Spiking

⁵ In the event-based video datasets, the data shape is $I \in \mathbb{R}^{T \times C \times H \times W}$, where T, C, H, and W denote the time step, channel, height, and width, respectively. In static datasets, a 2D image $I_s \in \mathbb{R}^{C \times H \times W}$ needs to be repeated T times to form an image sequence.



Fig. 1: The overall architecture of our proposed FW former. It mainly consists of three components: (1) Spiking Patch Splitting (SPS) module, (2) FW former Encoder Layer, and (3) Classification Layer. Additionally, we highlight the similarities between the FW head and SSA head at a single time step, which inspires us to choose the former as an exploration for more efficient calculations within the spike-form framework.

Neuron Layer. Residual connections are also applied to both the modules. The FW head in FW sub-layer serves as a critical component in our encoder layer, providing an efficient method for spike-form sparse representation. We have provided two implementations for FW head, including Fourier Transform (FT) and Wavelet Transform (WT). Many works in the past have used FT and WT to alternate between the spatial and frequency domains, allowing for efficient analysis of signals. While in this paper we treat them as structured basis functions with prior knowledge for information transformation. These implementations will be thoroughly analyzed in the next section.

Finally, following the processing in spikformer, a Global Average-Pooling (GAP) operation is applied to the resulting spike features, generating a D-dimensional feature. The feature is then fed into the Classification module consisting of a spiking fully-connected (SFC) layer, which produces the prediction

Y. The formulation of our FW former can be expressed as follows:

$$\boldsymbol{P} = \mathrm{SPS}\left(\boldsymbol{I}\right),\tag{12}$$

$$RPE = CPE(\boldsymbol{P}), \tag{13}$$

$$\boldsymbol{X}_0 = \boldsymbol{P} + \mathrm{RPE},\tag{14}$$

$$\boldsymbol{X}_{l}^{\prime} = \mathcal{SN}(\mathrm{BN}(\mathrm{FW}(\boldsymbol{X}_{l-1}))) + \boldsymbol{X}_{l-1},$$
(15)

$$\boldsymbol{X}_{l} = \mathcal{SN}(\mathrm{BN}(\mathrm{MLP}(\boldsymbol{X}_{l}'))) + \boldsymbol{X}_{l}', \qquad (16)$$

$$Y = SFC(GAP(X_L)), \tag{17}$$

where $\boldsymbol{I} \in \mathbb{R}^{T \times C \times H \times W}$, $\boldsymbol{P} \in \mathbb{R}^{T \times N \times D}$, RPE $\in \mathbb{R}^{T \times N \times D}$, $\boldsymbol{X}_0 \in \mathbb{R}^{T \times N \times D}$, $\boldsymbol{X}'_l \in \mathbb{R}^{T \times N \times D}$, $\boldsymbol{X}_l \in \mathbb{R}^{T \times N \times D}$ and l = 1, ..., L.

Moreover, the Membrane Shortcut (MS), which has been applied in many existing works [2, 44], is also utilized in our model for comparison. It establishes a shortcut between the membrane potential of spiking neurons in various layers to enhance performance and increase biological plausibility [44].

4.2 The FW head

Given input spike sequences $\boldsymbol{X} \in \mathbb{R}^{T \times N \times D}$, these features are then transformed into spiking sequences $\boldsymbol{X}' \in \mathbb{R}^{T \times N \times D}$ through a \mathcal{SN} layer. The formulation can be expressed as:

$$FW(\boldsymbol{X}) = FT(\boldsymbol{X}) \text{ or } WT(\boldsymbol{X}),$$

$$\boldsymbol{X}' = \mathcal{SN}(BN((FW(\boldsymbol{X})))),$$

(18)

In contrast to the SSA head in Equation (5), the FW head does not involve any learnable parameters or Self-Attention calculations. Here, we can choose from Fourier Transform (FT) and Wavelet Transform (WT) with fixed basis functions. We can also combine different wavelet bases to form a superior function, which is defined as follows:

$$Base = a \cdot Base1 + b \cdot Base2 + c \cdot Base3,$$

FW(**X**) = Base(**X**), (19)

where a, b, c are learnable parameters, and *Base1*, *Base2*, *Base3* are selected bases. Since Wavelet Transform is a linear transformation, Equation 19 can also be written as:

$$FW(\boldsymbol{X}) = a \cdot Base1(\boldsymbol{X}) + b \cdot Base2(\boldsymbol{X}) + c \cdot Base3(\boldsymbol{X})$$
(20)

Further analysis and experiments will be conducted based on the proposed FW head in the following sections.

4.3 Time Complexity Analysis of FW vs. SSA

We make a time complexity analysis between SSA, Fourier Transform (FT), and Wavelet Transform (WT). The results are presented in Table 1. In the subsequent experimental section, we also conduct a more specific comparison of the training and inference speeds between FW and SSA under the same conditions.

In SSA (Equation (5)), since there is no softmax operation, the order of calculation between Q, K, and V can be changed: either QK^{T} followed by V, or $K^{T}V$ followed by Q. The former has a time complexity of $\mathcal{O}(N^{2}d)$, while the latter has $\mathcal{O}(Nd^{2})$, where d is the feature dimension per head.⁶ The second complexity, $\mathcal{O}(Nd^{2})$, cannot be simply considered as $\mathcal{O}(N)$ due to the large constant d^{2} involved. Only when the sequence length N is significantly larger than the feature dimension per

Table 1: The time complexity for different methods. We have N = 64, D = 384 or 256, and d = 32.

,	,
Methods	Time Complexity
SSA [50]	$\mathcal{O}(N^2d)$ or $\mathcal{O}(Nd^2)$
1D-FFT	$\mathcal{O}(N \log N)$
2D-FFT	$\mathcal{O}(D\log D) + \mathcal{O}(N\log N)$
2D-WT	$\mathcal{O}(D\log D) + \mathcal{O}(N\log N)$

head d does it demonstrate a significant computational efficiency advantage over the first complexity, $\mathcal{O}(N^2 d)$.

In our implementation, we utilize the Fast Fourier Transform (FFT) algorithm to compute the discrete FT. Specifically, we employ the Cooley-Tukey algorithm [5], which recursively expresses the discrete FT of a sequence of length $N = N_1N_2$ in terms of N_1 smaller discrete FTs of size N_2 , reducing the time complexity to $\mathcal{O}(N \log N)$ for discrete 1D-FT along the sequence dimension. Similarly, for discrete 2D-FT first along the feature dimension and then along the sequence dimension, the time complexity is $\mathcal{O}(D \log D) + \mathcal{O}(N \log N)$. In general, the complexity of WT is comparable to that of FFT [13].

5 Experiments

We conduct experiments on event-based video datasets (CIFAR10-DVS and Dvs-Gesture), as well as static image datasets (CIFAR10 and CIFAR100). The FW-former is trained from scratch and compared with existing methods, including spikformer with SSA and its variant. More analyses are also given about the effects of different wavelet bases and their combinations.

5.1 Experiment Settings

To ensure a fair comparison, we ensure the same configurations of spikformer with SSA for datasets, implementation details, and evaluation metrics. To conduct the experiments, we implement the models using PyTorch and Spiking-Jelly $[10]^7$. All experiments are conducted on NVIDIA A100 GPU.

Event-based Video Datasets For the CIFAR10-DVS and DvsGesture datasets, which have an image size of 128×128 , we employ the Spiking Patch Splitting (SPS) module with a patch size of 16×16 . This configuration splits each image into a sequence with a length N of 64 and a feature dimension D of 256. We utilize 2 FWformer encoder layers and set the time step of the spiking neuron to

⁶ In practice, the SSA in Equation (5) can be extended to multi-head SSA. In this case, d = D/H, where H is the number of heads.

⁷ https://github.com/fangwei123456/spikingjelly

16. The training process consists of 106 epochs for CIFAR10-DVS and 200 epochs for DvsGesture. We employ the AdamW optimizer with a batch size of 16. The learning rate is initialized to 0.1 and reduced using cosine decay. Additionally, data augmentation techniques, as described in [25], are applied specifically to the CIFAR10-DVS dataset.

Static Image Datasets For the CIFAR10/100 datasets featuring an image size of 32×32 , we employ the SPS module with a patch size of 4×4 , which splits each image into a sequence of length N = 64 and a feature dimension of D = 384. For the FW former Encoder, we use 4 layers, and the time-step of the spiking neuron is set to 4. During training, we utilize the AdamW optimizer with a batch size of 128. The training process spans 400 epochs, with a cosine-decay learning rate starting at 0.0005. Following the approach outlined in [46], we apply standard data augmentation techniques such as random augmentation, mixup, and cutmix during training.

5.2 Accuracy Performance

We evaluate the accuracy performance on visual classification tasks, utilizing Top-1 accuracy (Top-1 acc.) as the performance metric. The results of our FW-former, spikformer with SSA, and other existing methods (both SNNs and ANNs, including the spikformer variant [44]) on event-based video datasets as well as static image datasets are presented in Table 2.

Table 2: Accuracy performance comparison of our method with existing methods on CIFAR10-DVS (DVS10), DvsGesture (DVS128), CIFAR10, and CIFAR100. Our FWformer (* means we replace vanilla residual connection with Membrane Shortcut (MS)) outperforms spikformer with SSA on event-based video datasets in terms of Top-1 acc. and achieves comparable accuracy on static datasets (the text in **bold** indicates the best results). It is necessary to mention that 2-256 signifies a configuration with 2 encoder layers and a feature dimension of 256.

Methods	Architecture	Time Step (DVS10/128)	Top-1 acc. (DVS10/128)	Methods	Architecture	Time Step	Top-1 acc. (CIFAR10/100)
LIAF [40]	LIAE-Net	10/60	70 4/97 6	Hybrid training [31]	VGG-11	125	92.22/67.87
TA-SNN [43]	TA-SNN	10/60	72.0/98.6	Diet-SNN [30]	ResNet-20	10/5	92.54/64.07
Rollout [21]	_	48/240	66 8/97 2	STBP [38]	CIFARNet	12	89.83/-
DECOLLE [18]	_	- /500	- /95.5	STBP NeuNorm [39]	CIFARNet	12	90.53/-
tdBN [49]	ResNet-19	10/40	67.8/96.9	Dspike [24]	_	6	94.3/74.2
PLIF [11]	_	20/20	74.8/97.6	TSSL-BP [48]	CIFARNet	5	91.41/-
D-ResNet [12]	Wide-7B-Net	16/16	74.4/97.9	STBP-tdBN [49]	ResNet-19	4	92.92/70.86
Dspike [24]	_	10/-	75.4/-	TET [8]	ResNet-19	4	94.44/74.47
SALT [20]	_	20/-	67.1/-		RosNot 10	1	04 07 /75 35
DSR [28]	-	10/ -	77.3'/-	ANN Methods	Transformer-4-384	1	96.73/81.02
SDSA [44]	Spikformer-2-256	16/16	80.0/99.3	SDSA [44]	Spikformer-4-384	4	95.6/78.4
SSA [50]	Spikformer-2-256	16/16	80.9/98.3	SSA [50]	Spikformer-4-384	4	95.51/78.21
1D-FFT	FW former - 2 - 256	16/16	80.5/99.0	1D-FFT	FWformer-4-384	4	94.9/77.3
2D-FFT	FWformer-2-256	16/16	80.6/98.4	2D-FFT	FWformer-4-384	4	95.1/77.9
2D-WT-Haar	FWformer-2-256	16/16	81.0/98.5	2D-WT-Haar	FWformer-4-384	4	95.2/78.1
1D-FFT*	FWformer-2-256	16/16	80.8/ 99.5	$1D-FFT^*$	FWformer-4-384	4	95.5/78.0
$2D-FFT^*$	FWformer-2-256	16/16	80.7/98.2	2D-FFT*	FWformer-4-384	4	95.0/78.3
2D-WT-Haar*	FWformer-2-256	16/16	81.2/99.1	$2D-WT-Haar^*$	FWformer-4-384	4	95.6 /78.2

Our FWformer achieves remarkable accuracy, reaching 81.2% on CIFAR10-DVS with 2D-WT-Haar, and an impressive 99.5% on DvsGesture with 1D-FFT. The performances surpass the spikformer with SSA by 0.3% and 1.2%, respectively. While on static datasets, our FWformer variants demonstrate comparable Top-1 accuracy. The results demonstrate the advantage of our methods, particularly on event-based video datasets.

5.3 Computational Costs and Speed Performance

Furthermore, we conduct a comprehensive comparison between existing works and our FW former in terms of GPU memory usage, training speed, and inference speed, ensuring identical operating conditions. The training speed represents the time taken for the forward and back-propagation of a batch of data, while the inference speed denotes the time taken for the forward-propagation of a batch of data in milliseconds (ms). To minimize variance, we calculate the average time spent over 100 batches. The results are presented in Table 3 (DWT-C means 2D-WT combination with learnable parameters, which will be discussed in the next subsection). We also provide an analysis of energy efficiency in **Appendix**.

In the case of event-based video datasets, our FWformer achieves a significant reduction in the number of parameters, approximately 20%, under identical hyperparameter configurations and operating conditions. This reduction, attributed to the absence of learnable parameters, translates to around 4%-5% memory savings. Moreover, our FWformer demonstrates remarkable improvements in both training and inference speeds, showing increases of approximately 9%-51% and 33%-70%, respectively, compared to SSA. While in the case of static datasets, our FWformer also shows several advantages under identical hyperparameter configurations and operating conditions. It achieves a notable reduction in the number of parameters, approximately 25%, leading to memory savings of around 26%. Furthermore, our FWformer enhances both training and inference speeds by approximately 18%-29% and 19%-61%, respectively.

Table 3: Memory usage and speed performance comparison of our method with existing methods on CIFAR10-DVS (DVS10), DvsGesture (DVS128) and CIFAR-static (CIFAR10 and CIFAR100). Our FWformer outperforms spikformer with SSA [50] and its variant [44] when comparing GPU memory usage, training speed and inference speed under identical operating conditions.

Methods	Param (M)	Memory (DVS10/128) (GB)	Training Speed (DVS10/128) (ms/batch)	Inference Speed (DVS10/128) (ms/batch)
STBP-tdBN [49]	12.63	25.86/25.87	65/194	27/98
TET [8]	12.63	36.13/36.17	71/203	22/77
SDSA [44]	2.59	9.02/9.03	73/245	29/101
SSA [50]	2.59	9.02/9.03	76/246	30/105
1D-FFT	2.06	8.67/ 8.71	51/121	11/32
2D-FFT	2.06	8.54/8.74	55/135	21/37
2D-WT-Haar	2.06	8.70/8.73	62/139	21/46
DWT-C	2.06	8.55/8.74	69/158	23/48

Methods	Param (M)	Memory (CIFAR-static) (GB)	Training Speed (CIFAR-static) (ms/batch)	Inference Speed (CIFAR-static) (ms/batch)
STBP-tdBN [49]	12.63	8.02	155	20
TET [8]	12.63	8.19	148	23
SDSA [44]	9.32	11.69	162	33
SSA [50]	9.32	11.69	166	31
1D-FFT	6.96	8.61	118	12
2D-FFT	6.96	8.75	122	13
2D-WT-Haar	6.96	9.33	121	19
DWT-C	6.96	9.86	136	25



Fig. 2: (A) We treat spiking self-attention as a set of basis functions and proceed to measure the changes in their orthogonality throughout the training process. (B) A diagram visualizing how the basis functions, spanning a feature space, are transformed from orthogonal to non-orthogonal, with only two axes used for simplification. (C) A diagram visualizing our endeavor to employ fixed non-orthogonal bases.

5.4 From orthogonal to non-orthogonal bases

In the previous experiments, the Haar base was used as the default choice for Wavelet Transform. We have also compared the performance of some other wavelet bases including Db1, Bior1.1, and Rbio1.1, each having different functions for deconstructing the spike-form feature while maintaining orthogonality. The results on CIFAR10-DVS and DvsGesture are presented in Table 4. Interestingly, most alternative basis functions yield similar Top-1 accuracy. Their performance is comparable to or even better than that of spikformer. It is essential to highlight that Wavelet Transform offers numerous different basis function options, and our exploration has not been exhaustive. Investigating the influence of more basis functions, is an avenue for future research.

However, a more interesting question arises: Is it always necessary to pursue orthogonality? Although in many cases, orthogonality signifies sparse and efficient information representation, neural networks may show the opposite phenomenon in the actual training process, that is, parameters naturally tend toward overlapping representations, and SSA is no exception. To illustrate this, we treat SSA as basis functions as proposed in the previous sections, and then quantitatively measure changes in their orthogonality during training. We calculate the inner product of each row vector (one base) in $Q \times K$ with others (other bases) and sum them in each training step. The variation trend is shown in Figure 2 (A). Initially, network parameters are nearly orthogonal at initialization, but their orthogonality is continuously decreasing during training. A diagram visualizing how the basis functions change during training is provided in Figure 2 (B). Inspired by this phenomenon, we further explore the combination of different wavelet bases to form fixed non-orthogonal basis functions, as depicted in Figure 2 (C). We assume that the pre-trained dynamic bases may serve an equivalent function to the fixed non-orthogonal bases.

Here we choose Bior1.1, Haar, and Db1 for further exploration. To search for the proper coefficients of their combinations, we set an initial set of learnable coefficients first and then conduct training. The new FW head will have three parameters to learn, which exert minimal influence on the overall computation of the network but play a crucial role in finding suitable combinations. The results are presented in Table 4. Consistent with our hypothesis, the use of fixed nonorthogonal basis functions further improves accuracy performance (0.4%-1.5% improvements on event-based video datasets, compared to vanilla spikformer).

We attempt to conduct a preliminary analvsis of the situations in which our FWformer is applicable: In contrast to conventional signal processing, complex tasks such as NLP and ASR need the designed models to learn diverse syntactic and semantic relationships. which can hardly be represented simply by fixed basis functions such as Fourier bases. For this reason, the network has to form dynamic higher-order basis functions, which are adjusted by not only the changing inputs but also the parameter learning of the network itself. We can regard these basis functions in networks as hyper-parameters that need continuous adjustment. However, this also means each time the basis functions change, the rest of the network has to adapt accordingly, which is understandable in complex tasks but not

Table 4: Accuracy performance of different wavelet bases as well as their combination with learnable parameters (DWT-C) on DVS10 and DVS128 datasets (* means replacing vanilla residual connection with Membrane Shortcut (MS)).

Methods	Architecture	Top-1 acc. $DVS10/128$
Db1	FWformer-2-256	81.0/98.7
Bior1.1	FWformer-2-256	80.9/98.2
Rbio1.1	FWformer-2-256	80.4/98.1
DWT-C	FWformer-2-256	81.3 /99.1
DWT-C*	FW former - 2 - 256	81.2/ 99.8

necessary in some other cases (e.g. event-based video tasks). Moreover, within spike-form frameworks, the features are represented by such sparse spiking signals that the correlation between them is too weak to form closed similarity, so it is more suitable to use structured fixed basis functions (e.g. Fourier bases and Wavelet bases) containing prior knowledge to get a simplified network.

6 Conclusion

We present the FW former that replaces SSA with spike-form FW head, based on the hypothesis that both of them use dynamic or fixed bases to transform information. The proposed model achieves comparable or better accuracy, higher training and inference speed, and reduced computational cost, on both eventbased video datasets and static datasets. We analyze the orthogonality in SSA during training and assume that the pre-trained dynamic bases serve an equivalent function to the fixed bases, which inspires us to explore non-orthogonal combined bases and get even higher accuracy. Additionally, we provide an analysis of why and under what scenarios our FW former is effective, indicating the promising refinement of new Transformers in the future, which is inspired by biological discovery and information theory.

7 Appendix

We estimate the theoretical energy consumption of FW former mainly according to [15, 44, 50]. It is calculated by the following two equations:

$$SOPs(l) = Rate \times T \times FLOPs(l),$$
 (21)

$$E_{FWformer} = E_{MAC} \times \text{EL}_{Conv}^{1} + E_{AC} \times (\sum_{k=2}^{K} \text{SOP}_{Conv}^{k} + \sum_{m=1}^{M} \text{SOP}_{FC}^{m} + \sum_{n=1}^{N} \text{SOP}_{FW}^{n}),$$
(22)

SOPs(l) means synaptic operations (the number of spike-based accumulate (AC) operations) of layer l, Rate is the average firing rate of input spike train to layer l, T is the time window of LIF neurons, and FLOPs(l) refers to the floating point operations (the number of multiply-and-accumulate (MAC) operations) of layer l. We assume that the MAC and AC operations are implemented on the 45nm hardware [15], with $E_{MAC} = 4.6pJ$ and $E_{AC} = 0.9pJ$.

 EL_{Conv}^{1} represents the FLOPs of convolution module in ANNs. It is used for the first layer to convert static images into spike trains, which can also be written as SOP_{Conv}^{1} for event-based video datasets, and SOP_{Net}^{l} (SOP_{Conv}^{k} , SOP_{FC}^{m} , SOP_{FW}^{n}) is for the rest of FWformer.

The experimental settings are the same as in the main text. Each cell in the table below contains results presented in the form of OPs(G)/Power(mJ), where OPs refers to the total SOPs in a SNN model, and Power refers to the average theoretical energy consumption when predicting one sample from the datasets.

The results indicate that our methods can achieve a reduction in energy consumption of approximately 20%-25% compared to SSA [50], and 4%-9% compared to its variant [44]. This is primarily due to lower computational complexity of the FW head, as reflected in fewer total SOPs (OPs). Our FW former demonstrates enhanced energy efficiency.

Table 5: The theoretical energy consumption on DVS10 (CIFAR10-DVS), DVS128 (DvsGesture), CIFAR10 and CIFAR100 dataset. DWT-C means 2D-WT combination with learnable parameters.

Methods	DVS10 OPs(G)/Power(mJ)	DVS128 OPs(G)/Power(mJ)	Methods	CIFAR10 OPs(G)/Power(mJ	CIFAR100) OPs(G)/Power(mJ)
SDSA [44] SSA [50]	$\frac{1.561}{0.816}$ $\frac{1.852}{0.943}$	$\frac{1.620}{0.713}$ $\frac{1.914}{0.822}$	SDSA [44 SSA [50]	$\begin{array}{c} 0.951/0.415 \\ 1.186/0.523 \end{array}$	1.446/0.609 1.737/0.748
1D-FFT 2D-FFT 2D-WT DWT-C	1.547/0.752 1.548/0.752 1.549/0.753 1.553/0.753	$\begin{array}{c} 1.608/0.650\\ 1.609/0.653\\ 1.609/0.651\\ 1.613/0.652\end{array}$	1D-FFT 2D-FFT 2D-WT DWT-C	$\begin{array}{c} 0.942/0.392\\ 0.943/0.393\\ 0.944/0.393\\ 0.944/0.393\\ 0.947/0.394\end{array}$	$\begin{array}{c} 1.438/0.578\\ 1.438/0.584\\ 1.439/0.584\\ 1.442/0.586\end{array}$

References

- Carion, N., et al.: End-to-end object detection with transformers. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 213–229. Springer (2020) 3
- Chen, G., Peng, P., Li, G., Tian, Y.: Training full spike neural networks via auxiliary accumulation pathway. arXiv preprint arXiv:2301.11929 (2023)
- Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12299–12310 (2021)
 3
- Chu, X., Tian, Z., Wang, Y., Zhang, B., et al.: Twins: Revisiting the design of spatial attention in vision transformers. Advances in Neural Information Processing Systems 34, 9355–9366 (2021) 7
- Cooley, J.W., Tukey, J.W.: An algorithm for the machine calculation of complex Fourier series. Mathematics of computation 19(90), 297–301 (1965) 10
- Davies, M., Srinivasa, N., Lin, T.H., Chinya, G., Cao, Y., Choday, S.H., Dimou, G., Joshi, P., Imam, N., Jain, S., et al.: Loihi: A neuromorphic manycore processor with on-chip learning. Ieee Micro 38(1), 82–99 (2018) 1
- 7. Dayan, P., Abbott, L.F.: Theoretical neuroscience: computational and mathematical modeling of neural systems. MIT press (2005) 4
- Deng, S., Li, Y., Zhang, S., Gu, S.: Temporal Efficient Training of Spiking Neural Network via Gradient Re-weighting. In: International Conference on Learning Representations 11, 12
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In: International Conference on Learning Representations (2020) 3
- Fang, W., Chen, Y., Ding, J., Chen, D., Yu, Z., Zhou, H., Masquelier, T., Tian, Y., other contributors: SpikingJelly. https://github.com/fangwei123456/ spikingjelly (2020), accessed: YYYY-MM-DD 10
- Fang, W., Yu, Z., Chen, Y., et al.: Incorporating learnable membrane time constant to enhance learning of spiking neural networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2661–2671 (2021) 4, 11
- Fang, W., Yu, Z., Chen, Y., Huang, T., Masquelier, T., Tian, Y.: Deep residual learning in spiking neural networks. Advances in Neural Information Processing Systems 34, 21056–21069 (2021) 11
- Gonzales, R.C., Wintz, P.: Digital image processing. Addison-Wesley Longman Publishing Co., Inc. (1987) 10
- Han, B., Srinivasan, G., Roy, K.: Rmp-snn: Residual membrane potential neuron for enabling deeper high-accuracy and low-latency spiking neural network. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13558–13567 (2020) 4
- Horowitz, M.: 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). pp. 10–14. IEEE (2014) 15
- 16. Hu, Y., Tang, H., Pan, G.: Spiking deep residual networks. IEEE Transactions on Neural Networks and Learning Systems (2021) 4
- Izhikevich, E.M., Gally, J.A., Edelman, G.M.: Spike-timing dynamics of neuronal groups. Cerebral cortex 14(8), 933–944 (2004) 4

- Kaiser, J., Mostafa, H., Neftci, E.: Synaptic plasticity dynamics for deep continuous local learning (DECOLLE). Frontiers in Neuroscience 14, 424 (2020) 11
- Katharopoulos, A., Vyas, A., et al.: Transformers are rnns: Fast autoregressive transformers with linear attention. In: International Conference on Machine Learning. pp. 5156–5165. PMLR (2020) 3
- Kim, Y., Panda, P.: Optimizing deeper spiking neural networks for dynamic vision sensing. Neural Networks 144, 686–698 (2021) 11
- Kugele, A., Pfeil, T., Pfeiffer, M., Chicca, E.: Efficient processing of spatio-temporal data streams with spiking neural networks. Frontiers in Neuroscience 14, 439 (2020) 11
- Lee, C., Sarwar, S.S., et al.: Enabling spike-based backpropagation for training deep neural network architectures. Frontiers in neuroscience p. 119 (2020) 4
- 23. Lee-Thorp, J., et al.: Fnet: Mixing tokens with fourier transforms. arXiv preprint arXiv:2105.03824 (2021) 2
- Li, Y., Guo, Y., Zhang, S., Deng, S., Hai, Y., Gu, S.: Differentiable spike: Rethinking gradient-descent for training spiking neural networks. Advances in Neural Information Processing Systems 34, 23426–23439 (2021) 11
- Li, Y., Kim, Y., Park, H., Geller, T., Panda, P.: Neuromorphic data augmentation for training spiking neural networks. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII. pp. 631–649. Springer (2022) 11
- Lotfi Rezaabad, A., Vishwanath, S.: Long short-term memory spiking networks and their applications. In: International Conference on Neuromorphic Systems 2020. pp. 1–9 (2020) 4
- Maass, W.: Networks of spiking neurons: the third generation of neural network models. Neural networks 10(9), 1659–1671 (1997) 1
- Meng, Q., Xiao, M., Yan, S., Wang, Y., et al.: Training high-performance lowlatency spiking neural networks by differentiation on spike representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12444–12453 (2022) 11
- Mueller, E., Studenyak, V., et al.: Spiking transformer networks: A rate coded approach for processing sequential data. In: 2021 7th International Conference on Systems and Informatics (ICSAI). pp. 1–5. IEEE (2021) 4
- Rathi, N., Roy, K.: Diet-snn: Direct input encoding with leakage and threshold optimization in deep spiking neural networks. arXiv preprint arXiv:2008.03658 (2020) 11
- Rathi, N., Srinivasan, G., et al.: Enabling Deep Spiking Neural Networks with Hybrid Conversion and Spike Timing Dependent Backpropagation. In: International Conference on Learning Representations 11
- Rueckauer, B., Lungu, I.A., Hu, Y., et al.: Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. Frontiers in neuroscience 11, 682 (2017) 4
- Shrestha, S.B., Orchard, G.: Slayer: Spike layer error reassignment in time. Advances in neural information processing systems 31 (2018) 4
- Tang, et al.: Deep reinforcement learning with population-coded spiking neural network for continuous control. In: Conference on Robot Learning. pp. 2016–2029. PMLR (2021) 1
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017) 3, 5

- Wang, Q., Zhang, T., Han, M., Wang, Y., Zhang, D., Xu, B.: Complex dynamic neurons improved spiking transformer network for efficient automatic speech recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 102–109 (2023) 1
- 37. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., et al.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 568–578 (2021) 3
- Wu, Y., Deng, L., et al.: Spatio-temporal backpropagation for training highperformance spiking neural networks. Frontiers in neuroscience 12, 331 (2018) 11
- Wu, Y., Deng, L., Li, G., Zhu, J., Xie, Y., Shi, L.: Direct training for spiking neural networks: Faster, larger, better. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 1311–1318 (2019) 11
- Wu, Z., Zhang, H., Lin, Y., Li, G., Wang, M., Tang, Y.: Liaf-net: Leaky integrate and analog fire network for lightweight and efficient spatiotemporal information processing. IEEE Transactions on Neural Networks and Learning Systems 33(11), 6249–6262 (2021) 11
- Xiao, T., Singh, M., Mintun, E., et al.: Early convolutions help transformers see better. Advances in Neural Information Processing Systems 34, 30392–30400 (2021) 7
- Xu, M., Wu, Y., Deng, L., Liu, F., Li, G., Pei, J.: Exploiting Spiking Dynamics with Spatial-temporal Feature Normalization in Graph Learning. In: Zhou, Z. (ed.) Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021. pp. 3207-3213. ijcai.org (2021). https://doi.org/10.24963/ijcai.2021/441, https://doi.org/10.24963/ijcai.2021/441 4
- Yao, M., Gao, H., et al.: Temporal-wise attention spiking neural networks for event streams classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10221–10230 (2021) 4, 11
- Yao, M., Hu, J., Zhou, Z., Yuan, L., Tian, Y., Xu, B., Li, G.: Spike-driven transformer. Advances in Neural Information Processing Systems 36 (2024) 9, 11, 12, 15
- You, W., Sun, S., Iyyer, M.: Hard-coded gaussian attention for neural machine translation. arXiv preprint arXiv:2005.00742 (2020) 2
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., et al.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 558–567 (2021) 11
- 47. Zhang, T., Cheng, X., Jia, S., Poo, M.M., Zeng, Y., Xu, B.: Self-backpropagation of synaptic modifications elevates the efficiency of spiking and artificial neural networks. Sci Adv 7(43), eabh0146 (2021). https://doi.org/10.1126/sciadv. abh0146, https://www.ncbi.nlm.nih.gov/pubmed/34669481 1
- Zhang, W., Li, P.: Temporal spike sequence learning via backpropagation for deep spiking neural networks. Advances in Neural Information Processing Systems 33, 12022–12033 (2020) 11
- Zheng, H., Wu, Y., Deng, L., Hu, Y., Li, G.: Going deeper with directly-trained larger spiking neural networks. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 11062–11070 (2021) 11, 12
- Zhou, Z., et al.: Spikformer: When Spiking Neural Network Meets Transformer. In: International Conference on Learning Representations (2023) 1, 2, 4, 10, 11, 12, 15