

Binary Event-Driven Spiking Transformer

Honglin Cao¹, Zijian Zhou¹, Wenjie Wei¹, Yu Liang¹, Ammar Belatreche², Dehao Zhang¹, Malu Zhang^{1*}, Yang Yang¹ and Haizhou Li^{3,4}

¹University of Electronic Science and Technology of China

²Northumbria University

³The Chinese University of Hong Kong, Shenzhen

⁴National University of Singapore

{honglincao, zijianzhou}@std.uestc.edu.cn, maluzhang@uestc.edu.cn

Abstract

Transformer-based Spiking Neural Networks (SNNs) introduce a novel event-driven self-attention paradigm that combines the high performance of Transformers with the energy efficiency of SNNs. However, the larger model size and increased computational demands of the Transformer structure limit their practicality in resource-constrained scenarios. In this paper, we integrate binarization techniques into Transformer-based SNNs and propose the Binary Event-Driven Spiking Transformer, i.e. BESTformer. The proposed BESTformer can significantly reduce storage and computational demands by representing weights and attention maps with a mere 1-bit. However, BESTformer suffers from a severe performance drop from its full-precision counterpart due to the limited representation capability of binarization. To address this issue, we propose a Coupled Information Enhancement (CIE) method, which consists of a reversible framework and information enhancement distillation. By maximizing the mutual information between the binary model and its full-precision counterpart, the CIE method effectively mitigates the performance degradation of the BESTformer. Extensive experiments on static and neuromorphic datasets demonstrate that our method achieves superior performance to other binary SNNs, showcasing its potential as a compact yet high-performance model for resource-limited edge devices. The repository of this paper is available at <https://github.com/CaoHLin/BESTFormer>.

1 Introduction

Spiking Neural Networks (SNNs) have attracted significant attention as third-generation artificial neural networks, known for their high biological plausibility and low power consumption [Maass, 1997]. The spiking neuron utilizes binary spikes as the fundamental units for information transmission and works in a sparse spike-driven manner [Zhang *et al.*

*Corresponding author.




Figure 1: Accuracy vs. NS-ACE & Model Size. Our method achieves superior computational and storage efficiency while outperforming other quantized SNNs on ImageNet. Neuromorphic Synaptic Arithmetic Computation Effort (NS-ACE) assesses SNN resource use in neuromorphic computing environments [Shen *et al.*, 2024].

et al., 2021]. This sparse synaptic transmission in SNNs simplifies multiply-accumulate (MAC) operations into accumulate (AC) operations, thereby significantly enhancing computational efficiency [Li *et al.*, 2023; Xu *et al.*, 2024a]. Furthermore, the energy-efficient feature of SNNs has driven the development of neuromorphic hardware, such as TrueNorth [Akopyan *et al.*, 2015], Loihi [Davies *et al.*, 2018], and Tianjic [Pei *et al.*, 2019]. However, despite the notable energy efficiency of SNNs, their performance in complex tasks still requires improvement.

In recent years, there have been several studies integrating Transformers into SNNs, leading to a series of high-performance models, such as Spikformer [Zhou *et al.*, 2023b], Spikingformer [Zhou *et al.*, 2023a], Spike-Driven Transformer v1 and v2 [Yao *et al.*, 2024a; Yao *et al.*, 2024b], and SNN-ViT [Wang *et al.*, 2025]. Compared to convolutional architectures in SNNs, these Transformer-based models have demonstrated significant performance improvements [Zhang *et al.*, 2022]. However, their advancements typically rely on large model size, which is accompanied by substantial memory storage and computational overhead, limit-

ing their deployment on resource-constrained edge devices. Therefore, there is an urgent need for a compact yet high-performance Transformer-based SNN.

Quantization is a highly effective method for compressing large-scale models, which reduces model parameters from 32-bit to a low bit-width representation [Wei *et al.*, 2025]. As an extreme form of quantization, binarization maximizes model size compression and accelerates computational speed by employing bitwise operations [Qin *et al.*, 2020a]. Therefore, incorporating binarization with Transformer-based SNNs is promising for achieving an efficient and high-performance model. It is worth noting that current research on binarization in SNN domains is primarily focused on convolutional structures, while Transformer-based structures remain unexplored [Yin *et al.*, 2024; Wei *et al.*, 2024; Liang *et al.*, 2025].

In this paper, we explore the application of binarization technique in Transformer-based SNNs and propose a Binary Event-Driven Spiking Transformer (BESTformer) to minimize the model size and computational cost. Despite its high efficiency, the proposed BESTformer suffers from a significant performance drop due to limited information representation capability of binarization. To address this issue, we propose the Coupled Information Enhancement (CIE) method to maximize the mutual information between the binary model and its full precision counterpart as much as possible. By utilizing the CIE method in BESTformer, we improve its performance significantly while maintaining its efficiency advantage, as shown in Figure 1. The main contributions of this paper are summarized as follows:

- We explore the combination of binarization with high-performance, low-power event-driven self-attention paradigm, proposing the Binary Event-Driven Spiking Transformer (BESTformer). The proposed BESTformer compresses both the weight parameters and attention map into mere 1-bit representations, aiming to reduce the model size and the excessive computational burden of Transformer-based SNNs.
- We identify and analyse the performance degradation in BESTformer, which we attribute to the constrained information representation capability caused by binarization. Inspired by information theory, we propose the CIE method. This method utilizes a reversible framework and information enhancement distillation to maximize the mutual information between BESTformer and its full precision counterpart, leading to enhanced performance.
- We conduct extensive experiments on static and neuromorphic datasets and demonstrate that the proposed BESTformer with the CIE method outperforms other binary SNNs. It’s important to note that our method achieves a 7.85% performance improvement on ImageNet-1k datasets compared to other models of similar scale at a time step of 1.

2 Related Works

2.1 Transformer-based SNNs

In recent years, there have been several studies integrating Transformer architectures into SNNs, leading to a series of high-performance SNN models. Spikeformer [Li *et al.*, 2022b] is the first to integrate the Transformer architecture with SNNs, however, it retains numerous floating-point operations, making it unsuitable for neuromorphic computation. Spikformer [Zhou *et al.*, 2023b] introduces the Spiking Self Attention (SSA) mechanism, which enhances both energy efficiency and performance of Transformer-based SNNs. Based on this, Spikingformer [Zhou *et al.*, 2023a] modifies the residual connection within it to achieve a purely spike-driven Vision Transformer, further enhancing model’s efficiency. Spike-driven Transformer [Yao *et al.*, 2024b] proposes a spike-driven self-attention mechanism with linear complexity, significantly reducing energy consumption. Then, to ensure versatility and high performance across various vision tasks, [Yao *et al.*, 2024a] expand the original architecture into Meta-SpikeFormer, also known as Spike-driven Transformer v2. Moreover, SNN-ViT [Wang *et al.*, 2025] achieves state-of-the-art performance in spiking vision tasks by introducing a saccadic self-attention mechanism specifically designed for spatio-temporal spike trains, maintaining linear computational complexity for edge applications. Despite much progress, these models are limited by substantial memory and computational overheads, underscoring the need for further compression to reach their full potential.

2.2 Quantization techniques in SNNs

Various approaches have been proposed to quantize SNNs to low-bits. [Deng *et al.*, 2021] employs spatiotemporal backpropagation (STBP) to directly train quantized SNNs and introduces the alternating direction method of multipliers (ADMM) to solve the performance degradation caused by quantization. Then, to further enhance the performance, [Yoo and Jeong, 2023] uses constrained backpropagation (CBP) with the Lagrangian function as an objective function to quantize SNNs in training. As an extreme form of quantization, binarization has also been widely studied. [Qiao *et al.*, 2021] presents a weight-binarized SNN to efficiently process event-based data, addressing the training demand of neuromorphic hardware for event data. Moreover, [Pei *et al.*, 2023] proposes the accuracy loss estimator and binary weight optimization to achieve ultra-low latency adaptive local binary SNNs, which reduce memory storage by over 20% while still maintaining high recognition accuracy. Recently, [Wei *et al.*, 2024] introduces a quantized SNN (Q-SNN) that reduces both weight and membrane potential representation, they also propose a weight-spike dual regulation (WS-DR) method to enhance the performance of Q-SNN. Despite the effectiveness of these binarization methods in SNNs, they mainly focus on spiking convolutional architectures, while Transformer-based SNNs have not been explored.

3 Method

In this section, we first introduce the construction of BESTformer, including the weight binarization and the attention bi-

Attn	Boolean	HR-LIF	SR-LIF
[4,0,0,0]	[1,0,0,0]	[1,0,0,0]	[1,1,0,0]
[1,5,0,0]	[1,1,0,0]	[1,1,0,0]	[1,1,1,0]
[0,3,1,0]	[0,1,1,0]	[0,1,1,0]	[0,1,1,0]

Table 1: A simple example of binarizing Attn with boolean function, HR-LIF, and SR-LIF. The number of split patches N , time step T , threshold V_{th} , and time constant τ are set to 1, 4, 1, 0.5, respectively.

narization. Subsequently, we analyze the challenge of limited information representation capability in the binary model. In order to address this challenge, we take inspiration from the information theory and propose the CIE method, which encompasses a reversible framework and information enhancement distillation.

3.1 Binary Event-Driven Spiking Transformer

Weight binarization

Existing Transformer-based SNNs typically utilize the Leaky Integrate-and-Fire (LIF) model with the hard reset mechanism (HR-LIF), its dynamics can be described as the following discrete form:

$$\tilde{U}_l[t] = \tau U_l[t-1] + X_l[t], \quad (1)$$

where τ is time constant factor, $U_l[t]$ is the membrane potential of neurons in layer l at time $t-1$, $\tilde{U}_l[t]$ is its intermediate representation, and $X_l[t] \in \mathbb{R}^{C \times H \times W}$ is the input current. $X_l[t]$ is integrated by presynaptic neurons, described as:

$$X_l[t] = \mathcal{BN}(W_l S_{l-1}[t]), \quad (2)$$

where \mathcal{BN} represents batch normalization, W_l is the 32-bit weight matrix, and $S_{l-1}[t]$ is binary spike activities. Once membrane potential $U_l[t]$ reaches its firing threshold V_{th} , the neurons will generate a spike which can be described as:

$$S_l[t] = \begin{cases} 1, & \text{if } \tilde{U}_l[t] \geq V_{th}, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Neurons reset their membrane potential after emitting a spike. Typically, we set the reset potential to 0:

$$U_l[t] = (1 - S_l[t]) \cdot \tilde{U}_l[t]. \quad (4)$$

To further reduce storage and computation demands, we quantize W_l into 1-bit representation through the following formulas [Qin *et al.*, 2020b]:

$$\hat{W}_l = W_l - \bar{W}_l, \quad \hat{W}_l^{\text{std}} = \frac{\hat{W}_l}{\sigma(\hat{W}_l)}, \quad (5)$$

$$B_w = \begin{cases} +1, & \text{if } w \geq 0, \\ -1, & \text{otherwise,} \end{cases} \quad w \in \hat{W}_l^{\text{std}}, \quad (6)$$

where \bar{W}_l is the mean value of W_l , $\sigma(\hat{W}_l)$ is the standard deviation of \hat{W}_l . According to these two formulas, it can be seen that B_w has two features: zero mean and normalization, where zero mean maximizes the information entropy of

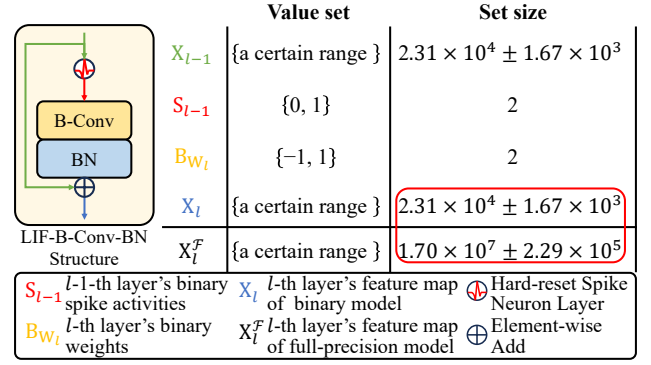


Figure 2: The 'LIF-B-Conv-BN' structure of BESTformer and representation capability of variables in the structure. Value set indicates the collection of all values present in a variable. Set size indicates the size of a value set.

weight and normalization can accelerate the convergence process [Salimans and Kingma, 2016; Qin *et al.*, 2020b]. Therefore, Equation 1 can be rewritten as:

$$U_l[t] = \tau U_l[t-1] + \mathcal{BN}(B_{W_l} \otimes S_{l-1}[t]), \quad (7)$$

where \otimes is efficient bitwise operations, theoretically offering a $32 \times$ memory saving and speedup compared with 32-bit operations [Rastegari *et al.*, 2016].

Attention binarization

Aside from 1-bit weights and 1-bit spike activities, BESTformer also binarizes another crucial component in the Transformer-based SNNs, i.e., the attention map Attn. Typically, Attn is obtained through the matrix multiplication of two 1-bit spike vectors, Query Q and Key K, yielding a non-negative integer result, described as:

$$\text{Attn} = \text{QK}^\top, \quad \text{Attn} \in \mathbb{N}^{(T \times N \times N)}, \quad (8)$$

where N is the number of split patches, T is the time step of BESTformer. To ensure full bitwise operations in BESTformer, we further binarize this attention map Attn. However, directly binarizing it using the boolean or sign function will lead to limited information retention. As the network depth increases, this limited information retention will result in severe performance degradation. Fortunately, this issue can be alleviated by leveraging the additional temporal dimension of spiking neurons while maintaining the binary nature of BESTformer. Given that each non-zero item in Attn is an integer at least 1, HR-LIF with a threshold of 1 will degenerate into the boolean function. Therefore, in this paper, we use LIF with the soft reset mechanism (SR-LIF) to binarize Attn, mathematically defined as:

$$B_{\text{Attn}} = \lambda \cdot \text{SR-LIF}(\text{Attn}), \quad (9)$$

where $\lambda \in \mathbb{R}^{(T \times 1 \times 1)}$ is the layer-wise learnable factors used to minimize binarization errors, which can be incorporated into the firing threshold during inference without requiring extra computations. SR-LIF calculates the membrane potential after spike emission by subtracting the threshold, i.e., $U_l[t] = \tilde{U}_l[t] - \theta S_l[t]$. As shown in Table 1, when compared with boolean function or HR-LIF neurons, binarizing Attn using SR-LIF can preserve more information.

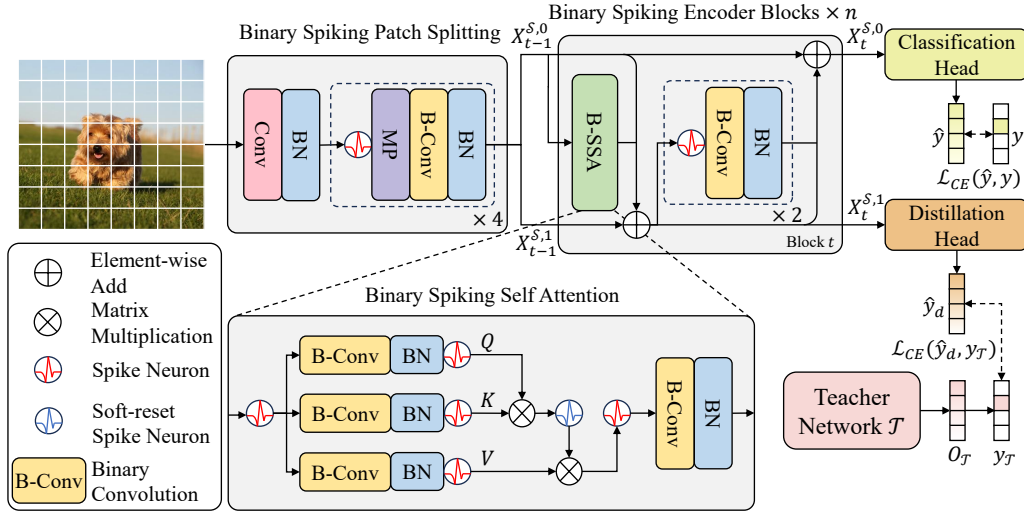


Figure 3: Overview of our BESTformer with the Coupled Information Enhancement method, which consists of a Binary Spiking Patch Splitting Module(BSPS), Reversible Binary Spiking Transformer Encoder Blocks, Classification and Distillation Heads.

3.2 Challenge analysis

BESTformer follows the event-driven ‘LIF-Conv-BN’ design as commonly used in [Zhou *et al.*, 2023a; Yao *et al.*, 2024a], but replaces vanilla convolution with binary convolution. As shown in Figure 2, X is obtained by convolution and \mathcal{B}/\mathcal{N} operations on binary B_{W_l} and S_{l-1} .

Previous research indicates that the poor performance of binarized neural networks is attributable to their low representational capability [Liu *et al.*, 2018; Guo *et al.*, 2022; Guo *et al.*, 2024]. In our research, we experimentally analyzed the average representation capability of the full-precision network and BESTformer on ImageNet-1k, respectively. We use the value set size of feature maps as the measure of representation capability, which refers to the maximum number of distinct values in a feature map. The specific results are shown in Figure 2.

According to the huge difference in set size values of X_l^F and X_l , we find that the representation capability of full-precision model is significantly greater than binary model. This indicates that the information carried by BESTformer is severely constrained than its full-precision counterpart. This results in BESTformer losing a significant amount of useful information during the forward process, leading to unsatisfactory performance, especially in deep networks.

3.3 Coupled information enhancement BESTformer

Due to the constrained information representation capability of the binarized model, the ideal binary model should aim to retain the information representation of their full-precision counterparts as much as possible, thus the mutual information between the binarized and full-precision models’ representations should be maximized [Li *et al.*, 2022a; Xu *et al.*, 2023; Xu *et al.*, 2024b]. Therefore, we propose a knowledge distillation framework, optimizing the mutual information \mathcal{I} between the student model \mathcal{S} and the full-precision teacher

model \mathcal{T} , which can be formalized as:

$$\max_{\theta^S} \mathcal{I}(X_n^S; X_m^T), \quad (10)$$

where θ^S represents the parameters of the student model, and X_n^S and X_m^T correspond to the final (n -th and m -th) encoder blocks’ outputs of the student and teacher models, respectively. It’s challenging to solve the maximization problem directly. Hence, we decompose this optimization objective into the difference between two entropy terms:

$$\mathcal{I}(X_n^S; X_m^T) = \mathcal{H}(X_n^S) - \mathcal{H}(X_n^S | X_m^T), \quad (11)$$

where $\mathcal{H}(X) = -\int p(x) \log p(x) dx$ and $p(x)$ denotes the probability density function of the random variable. We employed coupled information enhancement methods to optimize the objective: (1) maximizing $\mathcal{H}(X_n^S)$ to its upper bound by reversible framework [Gomez *et al.*, 2017], and (2) minimizing $\mathcal{H}(X_n^S | X_m^T)$ by information-enhanced distillation. The overall architecture diagram of applying the CIE method to BESTformer is shown in Figure 3.

Reversible framework

The information entropy of the model’s feature maps exhibits a decreasing trend as the number of layers increases. This can be shown explicitly by the inequality in Proposition 1.

Proposition 1. *In the context of deep neural networks, the information entropy of feature maps exhibits a non-increasing trend with respect to the depth of the network. That is, for X_l where $l \in \{1, \dots, n\}$, $\mathcal{H}(X_{l-1}) \geq \mathcal{H}(X_l)$ always holds true. Furthermore, we have $\mathcal{H}(X_0) \geq \mathcal{H}(X_1) \geq \dots \geq \mathcal{H}(X_n)$.*

Proposition 1 shows that information retained by the model tends to decrease as the neural network deepens, which is contrary to our objective of maximizing $\mathcal{H}(X_n^S)$. To address this issue, we utilize a reversible forward mapping in the encoder, as the reversible connection shown in Figure 3. We




Figure 4: Illustration of the forward and inverse process of the proposed reversible framework. The inverse process indicates that the inputs can be reconstructed from the outputs, i.e., this framework is reversible and no information is lost.

define forward mapping $\Phi_l(X_{l-1}^{S,0}, X_{l-1}^{S,1}) = (X_l^{S,0}, X_l^{S,1})$ as:

$$\begin{aligned} X_l^{S,0} &= \text{BSSA}(X_{l-1}^{S,1}) + \frac{1}{2}(X_{l-1}^{S,0} + X_{l-1}^{S,1}), \\ X_l^{S,1} &= \text{BMLP}(X_{l-1}^{S,0}) + \frac{1}{2}(X_{l-1}^{S,1} + X_{l-1}^{S,0}), \end{aligned} \quad (12)$$

where BSSA denotes the Binary Spiking Self Attention and BMLP represents the Binary MLP. As shown in Figure 4, in the reversible encoder, the input can be accurately reconstructed from the output, ensuring that no information is lost in the process. To verify this, we explicitly formulate the hidden inverse mapping that corresponds to the forward mapping as $\Phi_l^{-1}(X_l^{S,0}, X_l^{S,1}) = (X_{l-1}^{S,0}, X_{l-1}^{S,1})$ as:

$$\begin{aligned} X_{l-1}^{S,1} &= 2(X_l^{S,1} - \text{BMLP}(X_{l-1}^{S,0})) - X_l^{S,0}, \\ X_{l-1}^{S,0} &= 2(X_l^{S,0} - \text{BSSA}(X_{l-1}^{S,1})) - X_{l-1}^{S,1}. \end{aligned} \quad (13)$$

Inverse mapping does not directly engage in network’s operations, but it indirectly confirms that the information entropy within the reversible framework remains constant across all encoder blocks. Therefore, $\mathcal{H}(X_n^S)$ reaches its upper bound $\mathcal{H}(X_0^S)$ as demonstrated by Proposition 2, which aligns with the objective of maximizing $\mathcal{H}(X_n^S)$. The detailed proofs of Proposition 1 and 2 are given in **Supplementary Materials**.

Proposition 2. *In the context of reversible deep neural networks, the information entropy of the feature map remains invariant with respect to the depth of the network, i.e. $\mathcal{H}(X_0) = \mathcal{H}(X_1) = \dots = \mathcal{H}(X_n)$.*

Information enhanced distillation

After modifying the Transformer Encoder Blocks of our binarized model to a reversible connection form, we maximize $\mathcal{H}(X_n^S)$ to its information entropy upper bound $\mathcal{H}(X_0^S)$. Maximizing $\mathcal{H}(X_0^S)$ can be achieved through standard network training and weight standardization. Consequently, the core optimization objective becomes:

$$\min_{\theta^S} \mathcal{H}(X_n^S | X_m^T). \quad (14)$$

Given the challenges associated with directly minimizing $\mathcal{H}(X_n^S | X_m^T)$, we propose a knowledge distillation approach to implicitly achieve this minimization. Inspired by [Touvron *et al.*, 2021], we employ a dual-head architecture to fully leverage the information contained in the output features of the reversible network. Specifically, $X_n^{S,0}$ and $X_n^{S,1}$ are fed into the classification and distillation heads, respectively, generating the model’s outputs \hat{y} and \hat{y}_d . Let O_T denote the teacher model’s output, and $y_T = \arg \max O_T$ represent the teacher model’s ‘hard decision’. The global loss function incorporating distillation is formulated as:

$$\mathcal{L}_{global} = (\mathcal{L}_{CE}(\hat{y}, y) + \mathcal{L}_{CE}(\hat{y}_d, y_T))/2. \quad (15)$$

This dual-head design offers a significant advantage over traditional distillation methods by decoupling the gradient backpropagation for distillation and classification at the head level. This decoupling reduces mutual interference between the two gradients during backpropagation, facilitating more effective parameter updates. It is worth noting that $X_n^{S,1}$, being the output of a deeper network compared to $X_n^{S,0}$, is more susceptible to overfitting when used for classification. Therefore, we strategically feed $X_n^{S,1}$ into the distillation head. This approach leverages the teacher model’s output, which encapsulates ‘error information’, thereby alleviating the overfitting problem and enhancing model’s generalization capabilities.

4 Experiments

In this section, we first assess the classification performance of the proposed BESTformer with the CIE method on small-scale datasets, including CIFAR [Krizhevsky *et al.*, 2009], CIFAR10-DVS [Li *et al.*, 2017]. Following this, we evaluate the method’s performance on large-scale image dataset, ImageNet-1K [Deng *et al.*, 2009] to verify the scalability of our approach. Finally, we perform a series of ablation studies to validate the effectiveness of our method. The implementation details are provided in **Supplementary Materials**.

4.1 Comparison with Related Work

Results on small-scale datasets classification

We evaluate our BESTformer with the CIE method on small-scale datasets, including static datasets, CIFAR-10 and CIFAR-100, and neuromorphic datasets CIFAR10-DVS. The experimental results in Table 2 demonstrate the superior performance and efficiency of our proposed method across multiple benchmark datasets.

On the CIFAR-10 dataset, our 1-bit Bestformer-4-384 achieves a remarkable accuracy of 95.73%, significantly outperforming previous state-of-the-art Transformer-based methods (93.91%) and conventional Conv-based approaches (91.66%). Similarly, on CIFAR-100, our 1-bit model attains 79.80% accuracy, representing a substantial improvement of 3.89% over the previous best result of 75.91% achieved by 2-bit. Notably, this performance enhancement is achieved while maintaining exceptional model efficiency. Our 1-bit Bestformer-4-384 requires only 1.18MB of storage for CIFAR-10 and 1.31MB for CIFAR-100, representing more than a 93% reduction in model size compared to

Dataset	Method	Architecture	Weight Bits	Time Step	Model Size (MB)(↓)	Accuracy(↑)
CIFAR-10	[Yoo and Jeong, 2023]	VGG16	1	32	1.89	91.51%
		VGG16	2	32	3.73	91.66%
	[Hu <i>et al.</i> , 2024]	ResNet18	1	1	1.42	93.74%
	[Shen <i>et al.</i> , 2024]	Spikformer-4-384	1	4	1.17	93.91%
		Spikformer-4-384	2	2	2.28	93.56%
	Ours	Bestformer-2-384	1	4	0.73	94.98%
		Bestformer-4-384	1	2	1.18	95.19%
		Bestformer-4-384	1	4	1.18	95.73%
CIFAR-100	[Yoo and Jeong, 2023]	VGG16	1	32	2.08	66.53%
		VGG16	2	32	3.90	66.46%
	[Wei <i>et al.</i> , 2024]	ResNet19	1	2	1.56	78.77%
	[Shen <i>et al.</i> , 2024]	Spikformer-4-384	1	4	1.24	74.13%
		Spikformer-4-384	2	2	2.34	75.91%
	Ours	Bestformer-2-384	1	4	0.86	78.23%
		Bestformer-4-384	1	2	1.31	79.23%
		Bestformer-4-384	1	4	1.31	79.80%
ImageNet	[Yoo and Jeong, 2023]	SEW-ResNet18	1	4	3.36	54.34%
		SEW-ResNet18	2	4	4.88	58.04%
		SEW-ResNet34	1	4	4.59	60.10%
		SEW-ResNet34	2	4	7.13	62.98%
	[Shen <i>et al.</i> , 2024]	SEW-ResNet34	1	1	4.59	52.17%
		SEW-ResNet34	2	2	7.13	60.15%
		Spikformer-8-512	1	1	4.60	54.54%
		Spikformer-8-512	2	2	8.07	61.37%
CIFAR10-DVS	[Yoo and Jeong, 2023]	Wide-7B-Net	1	16	0.17	74.70%
		Wide-7B-Net	2	16	0.32	75.30%
	[Shen <i>et al.</i> , 2024]	Spikformer-2-256	1	16	0.33	79.80%
	Ours	Bestformer-2-256	1	10	0.34	78.70%
		Bestformer-2-256	1	16	0.34	80.80%

Table 2: Classification performance comparison on CIFAR-10, CIFAR-100, ImageNet and CIFAR10-DVS.

full-precision counterparts. This demonstrates that our quantization approach effectively preserves model accuracy while significantly reducing memory requirements.

The robustness of our method is further validated on neuromorphic datasets. On CIFAR10-DVS, our 1-bit Bestformer-2-256 achieves 80.80% accuracy, a state-of-the-art result. This consistent performance across both static and neuromorphic datasets underscores the versatility and effectiveness of our approach in different application scenarios.

Results on ImageNet-1k classification

On the challenging ImageNet dataset, Our model consistently shows superior performance compared to other quantization methods. According to Table 2, our 1-bit Bestformer-8-512 achieves an impressive accuracy of 63.46% with 4 time steps, outperforming other quantized methods such as CBP-QSNN

(60.10% with SEW-ResNet34) and Quantized Spikformer (61.37% with 2-bit weights). Moreover, with just 1 time step, the model attains an accuracy of 62.39%, improving performance by 7.85% (54.54% with 1-bit weight and 1 time step by Quantized Spikformer) while maintaining a similar model size and comparable computational complexity.

Compared to full-precision models, our method significantly reduces model size and computational efficiency. According to Table 3, our 1-bit Bestformer-8-512 model is 5.57MB, approximately 10 times smaller than its full-precision counterpart (59.36MB). In terms of computational efficiency, our 1-bit Bestformer-8-512 model with 4 time steps requires only 5.67G Neuromorphic Synaptic Arithmetic Computation Effort (NS-ACE [Shen *et al.*, 2024], more details are shown in **Supplementary Materials**), a 94.6% reduction from the 104.32G NS-ACE of the full-

Model	Model Size (MB)(↓)	SOPs (G)(↓)	NS-ACE (G)(↓)
Full-precision†	59.36	6.52	104.32
Q-SEW-ResNet‡	7.13	2.06	4.11
Q-Spikformer‡	8.07	3.93	7.86
Ours	5.57	2.13	2.13

Table 3: Comparison of resource consumption on ImageNet-1k. † The results for the full-precision model are taken from [Zhou *et al.*, 2023a]. ‡ The quantized results are taken from [Shen *et al.*, 2024].

precision model. This substantial decrease in NS-ACE indicates markedly improved energy efficiency of our model on neuromorphic hardware.

4.2 Ablation study

Impact of components of CIE on model accuracy

To validate the effectiveness of our proposed components, we conduct comprehensive ablation studies on CIFAR-100. As shown in Table 4, we progressively incorporate different components into our framework and observe their individual and combined effects. The baseline model achieves 77.77% accuracy. Adding the reversible architecture (RF) brings a 0.50% improvement, while incorporating information enhancement distillation (IED) alone leads to a more substantial 1.34% gain. When combining both components, we explore two variants: using $X_l^{S,0}$ for distillation and $X_l^{S,1}$ for classification yields a 1.90% improvement, while the reverse configuration achieves the best performance with a 2.03% accuracy gain. These results demonstrate that both the reversible architecture and information enhancement distillation contribute positively to the model’s performance, with their combination producing synergistic effects.

Impact of CIE on different architectures

Our ablation experiments on CIFAR100 demonstrate the effectiveness of the CIE method across different Bestformer architectures. The results of these ablation experiments are presented in Figure 5a. The Bestformer-2-384 architecture achieved 76.75% accuracy, which was significantly enhanced to 78.23% with CIE integration. Similarly, Bestformer-4-384 exhibited improved performance from 77.77% to 79.80%, while Bestformer-6-384 achieved the highest overall accuracy of 79.98% after applying CIE, compared to its base-

Method	Accuracy (%)	Increment (%)
baseline	77.77	-
w/ RF	78.27	+ 0.50
w/ IED	79.14	+ 1.34
w/ RF & IED†	79.67	+ 1.90
w/ RF & IED‡	79.80	+ 2.03

Table 4: Ablation study of CIE design on CIFAR-100. RF stands for reversible architecture, and IED stands for information enhancement distillation. Symbol † means that $X_l^{S,0}$ in Figure 3 is used for distillation and $X_l^{S,1}$ is used for classification. Symbol ‡ means that $X_l^{S,0}$ is used for classification and $X_l^{S,1}$ is used for distillation.

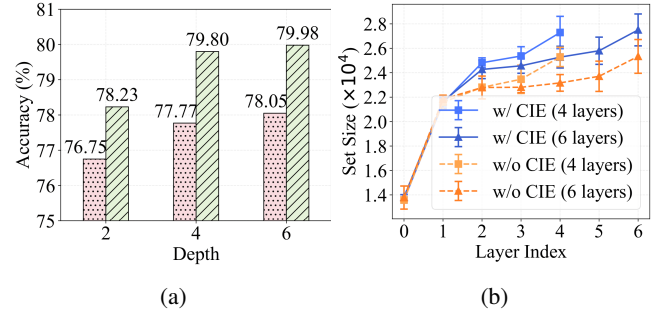


Figure 5: Ablation study for CIE method on CIFAR-100. (a) Impact of CIE on model accuracy across different architectures. (b) A comparative analysis of information representation capability: evaluating models with and without CIE method across variable architecture depths.

line of 78.05%. The consistent performance improvements across all architectures suggest that CIE effectively enhances feature representation regardless of model depth. Additionally, while deeper architectures demonstrated higher baseline performance, Bestformer-4-384 achieved the most substantial improvement (+2.03%) with CIE integration, indicating an optimal balance between model capacity and enhancement effectiveness.

Impact of CIE on information representation

To further validate the effectiveness of CIE, we conducted an in-depth analysis of information representation capability across different encoder blocks on CIFAR-100 datasets as shown in Figure 5b. We conducted multiple experiments to obtain the average values and corresponding error margins for all architectures to assess the overall trend. The results demonstrate that BESTformer with CIE method consistently maintains higher representation capability than original baseline (without CIE), especially in the later encoder blocks. The improved information representation ability helped the model with 4 layer encoder blocks achieve a 2.03% performance improvement. The representation gap between them remains relatively stable across different architectures, underlining the consistent superiority of our approach. These findings support the effectiveness of our CIE method, which contributes to improved performance in downstream tasks.

5 Conclusion

This work introduces a Binary Event-Driven Spiking Transformer that significantly reduces the storage and computational demands of Transformer-based Spiking Neural Networks. To address the constrained information representation capability caused by binarization, we propose the Coupled Information Enhancement (CIE) method, which combines a reversible framework and information enhancement distillation. Extensive experiments on both static and neuromorphic datasets demonstrate that BESTformer with CIE achieves superior performance compared to other binary SNNs while maintaining high efficiency. Our work provides a promising direction for developing compact yet high-performance models for resource-constrained edge devices.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China (No. U2333211, U20B2063 and 62220106008), in part by the Project of Sichuan Engineering Technology Research Center for Civil Aviation Flight Technology and Flight Safety (No. GY2024-27D), in part by the Open Research Fund of the State Key Laboratory of Brain-Machine Intelligence, Zhejiang University (Grant No. BMI2400020), in part by the Shenzhen Science and Technology Program (Shenzhen Key Laboratory, Grant No. ZDSYS20230626091302006), in part by the Shenzhen Science and Technology Research Fund (Fundamental Research Key Project, Grant No. JCYJ20220818103001002) and in part by the Program for Guangdong Introducing Innovative and Entrepreneurial Teams, Grant No. 2023ZT10X044.

Contribution Statement

Honglin Cao and Zijian Zhou contribute equally to this work.

References

- [Akopyan *et al.*, 2015] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE transactions on computer-aided design of integrated circuits and systems*, 34(10):1537–1557, 2015.
- [Davies *et al.*, 2018] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro*, 38(1):82–99, 2018.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [Deng *et al.*, 2021] Lei Deng, Yujie Wu, Yifan Hu, Ling Liang, Guoqi Li, Xing Hu, Yufei Ding, Peng Li, and Yuan Xie. Comprehensive snn compression using admm optimization and activity regularization. *IEEE transactions on neural networks and learning systems*, 34(6):2791–2805, 2021.
- [Gomez *et al.*, 2017] Aidan N Gomez, Mengye Ren, Raquel Urtasun, and Roger B Grosse. The reversible residual network: Backpropagation without storing activations. *Advances in neural information processing systems*, 30, 2017.
- [Guo *et al.*, 2022] Yufei Guo, Yuanpei Chen, Liwen Zhang, Xiaode Liu, Yinglei Wang, Xuhui Huang, and Zhe Ma. Im-loss: information maximization loss for spiking neural networks. *Advances in Neural Information Processing Systems*, 35:156–166, 2022.
- [Guo *et al.*, 2024] Yufei Guo, Yuanpei Chen, Xiaode Liu, Weihang Peng, Yuhang Zhang, Xuhui Huang, and Zhe Ma. Ternary spike: Learning ternary spikes for spiking neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 12244–12252, 2024.
- [Hu *et al.*, 2024] Yangfan Hu, Qian Zheng, and Gang Pan. Bitsnns: Revisiting energy-efficient spiking neural networks. *IEEE Transactions on Cognitive and Developmental Systems*, 2024.
- [Krizhevsky *et al.*, 2009] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [Li *et al.*, 2017] Hongmin Li, Hanchao Liu, Xiangyang Ji, Guoqi Li, and Luping Shi. Cifar10-dvs: an event-stream dataset for object classification. *Frontiers in neuroscience*, 11:309, 2017.
- [Li *et al.*, 2022a] Yanjing Li, Sheng Xu, Baochang Zhang, Xianbin Cao, Peng Gao, and Guodong Guo. Q-vit: Accurate and fully quantized low-bit vision transformer. *Advances in neural information processing systems*, 35:34451–34463, 2022.
- [Li *et al.*, 2022b] Yudong Li, Yunlin Lei, and Xu Yang. Spikeformer: a novel architecture for training high-performance low-latency spiking neural network. *arXiv preprint arXiv:2211.10686*, 2022.
- [Li *et al.*, 2023] Guoqi Li, Lei Deng, Huajing Tang, Gang Pan, Yonghong Tian, Kaushik Roy, and Wolfgang Maass. Brain inspired computing: A systematic survey and future trends. *Authorea Preprints*, 2023.
- [Liang *et al.*, 2025] Yu Liang, Wenjie Wei, Ammar Belatreche, Honglin Cao, Zijian Zhou, Shuai Wang, Malu Zhang, and Yang Yang. Towards accurate binary spiking neural networks: Learning with adaptive gradient modulation mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 1402–1410, 2025.
- [Liu *et al.*, 2018] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737, 2018.
- [Maass, 1997] Wolfgang Maass. Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671, 1997.
- [Pei *et al.*, 2019] Jing Pei, Lei Deng, Sen Song, Mingguo Zhao, Youhui Zhang, Shuang Wu, Guanrui Wang, Zhe Zou, Zhenzhi Wu, Wei He, et al. Towards artificial general intelligence with hybrid tianjic chip architecture. *Nature*, 572(7767):106–111, 2019.
- [Pei *et al.*, 2023] Yijian Pei, Changqing Xu, Zili Wu, Yi Liu, and Yintang Yang. Albsnn: ultra-low latency adaptive local binary spiking neural network with accuracy loss estimator. *Frontiers in Neuroscience*, 17, 2023.

- [Qiao *et al.*, 2021] GC Qiao, Ning Ning, Yue Zuo, SG Hu, Qi Yu, and Yecheng Liu. Direct training of hardware-friendly weight binarized spiking neural network with surrogate gradient learning towards spatio-temporal event-based dynamic data recognition. *Neurocomputing*, 457:203–213, 2021.
- [Qin *et al.*, 2020a] Haotong Qin, Ruihao Gong, Xianglong Liu, Xiao Bai, Jingkuan Song, and Nicu Sebe. Binary neural networks: A survey. *Pattern Recognition*, 105:107281, 2020.
- [Qin *et al.*, 2020b] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. Forward and backward information retention for accurate binary neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2250–2259, 2020.
- [Rastegari *et al.*, 2016] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: ImageNet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016.
- [Salimans and Kingma, 2016] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016.
- [Shen *et al.*, 2024] Guobin Shen, Dongcheng Zhao, Tenglong Li, Jindong Li, and Yi Zeng. Are conventional snns really efficient? a perspective from network quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27538–27547, 2024.
- [Touvron *et al.*, 2021] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021.
- [Wang *et al.*, 2025] Shuai Wang, Malu Zhang, Dehao Zhang, Ammar Belatreche, Yichen Xiao, Yu Liang, Yimeng Shan, Qian Sun, Enqi Zhang, and Yang Yang. Spiking vision transformer with saccadic attention. In *International Conference on Learning Representations (ICLR)*, 2025.
- [Wei *et al.*, 2024] Wenjie Wei, Yu Liang, Ammar Belatreche, Yichen Xiao, Honglin Cao, Zhenbang Ren, Guoqing Wang, Malu Zhang, and Yang Yang. Q-snns: Quantized spiking neural networks. *arXiv preprint arXiv:2406.13672*, 2024.
- [Wei *et al.*, 2025] Wenjie Wei, Malu Zhang, Zijian Zhou, Ammar Belatreche, Yimeng Shan, Yu Liang, Honglin Cao, Jieyuan Zhang, and Yang Yang. Qp-snn: Quantized and pruned spiking neural networks. In *International Conference on Learning Representations (ICLR)*, 2025.
- [Xu *et al.*, 2023] Sheng Xu, Yanjing Li, Mingbao Lin, Peng Gao, Guodong Guo, Jinhua Lü, and Baochang Zhang. Q-detr: An efficient low-bit quantized detection transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3842–3851, 2023.
- [Xu *et al.*, 2024a] Qi Xu, Xuanye Fang, Yaxin Li, Jiangrong Shen, De Ma, Yi Xu, and Gang Pan. Rsnn: Recurrent spiking neural networks for dynamic spatial-temporal information processing. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 10602–10610, 2024.
- [Xu *et al.*, 2024b] Qi Xu, Yaxin Li, Xuanye Fang, Jiangrong Shen, Qiang Zhang, and Gang Pan. Reversing structural pattern learning with biologically inspired knowledge distillation for spiking neural networks. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3431–3439, 2024.
- [Yao *et al.*, 2024a] Man Yao, JiaKui Hu, Tianxiang Hu, Yifan Xu, Zhaokun Zhou, Yonghong Tian, Bo XU, and Guoqi Li. Spike-driven transformer v2: Meta spiking neural network architecture inspiring the design of next-generation neuromorphic chips. In *The Twelfth International Conference on Learning Representations*, 2024.
- [Yao *et al.*, 2024b] Man Yao, Jiakui Hu, Zhaokun Zhou, Li Yuan, Yonghong Tian, Bo Xu, and Guoqi Li. Spike-driven transformer. *Advances in neural information processing systems*, 36, 2024.
- [Yin *et al.*, 2024] Ruokai Yin, Yuhang Li, Abhishek Moitra, and Priyadarshini Panda. Mint: Multiplier-less integer quantization for energy efficient spiking neural networks. In *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 830–835. IEEE, 2024.
- [Yoo and Jeong, 2023] Donghyung Yoo and Doo Seok Jeong. Cbp-qsnns: Spiking neural networks quantized using constrained backpropagation. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2023.
- [Zhang *et al.*, 2021] Malu Zhang, Jiadong Wang, Jibin Wu, Ammar Belatreche, Burin Amornpaisannon, Zhixuan Zhang, Venkata Pavan Kumar Miriyala, Hong Qu, Yansong Chua, Trevor E Carlson, et al. Rectified linear post-synaptic potential function for backpropagation in deep spiking neural networks. *IEEE transactions on neural networks and learning systems*, 33(5):1947–1958, 2021.
- [Zhang *et al.*, 2022] Jiqing Zhang, Bo Dong, Haiwei Zhang, Jianchuan Ding, Felix Heide, Baocai Yin, and Xin Yang. Spiking transformers for event-based single object tracking. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 8801–8810, 2022.
- [Zhou *et al.*, 2023a] Chenlin Zhou, Liutao Yu, Zhaokun Zhou, Zhengyu Ma, Han Zhang, Huihui Zhou, and Yonghong Tian. Spikingformer: Spike-driven residual learning for transformer-based spiking neural network. *arXiv preprint arXiv:2304.11954*, 2023.
- [Zhou *et al.*, 2023b] Zhaokun Zhou, Yuesheng Zhu, Chao He, Yaowei Wang, Shuicheng YAN, Yonghong Tian, and Li Yuan. Spikformer: When spiking neural network meets transformer. In *The Eleventh International Conference on Learning Representations*, 2023.