

Spike2Former: Efficient Spiking Transformer for High-performance Image Segmentation

Zhenxin Lei^{1, 2*}, Man Yao^{2†}, Jiakui Hu^{3, 4*}, Xinhao Luo², Yanye Lu^{3, 4}, Bo Xu², Guoqi Li^{2†}

¹University of Chinese Academy of Sciences

²Institute of Automation, Key Laboratory of Brain Cognition and Brain-inspired Intelligence Technology, Chinese Academy of Sciences

³Institute of Medical Technology, Peking University Health Science Center, Peking University

⁴National Biomedical Imaging Center, Peking University

Abstract

Spiking Neural Networks (SNNs) have a low-power advantage but perform poorly in image segmentation tasks. The reason is that directly converting neural networks with complex architectural designs for segmentation tasks into spiking versions leads to performance degradation and non-convergence. To address this challenge, we first identify the modules in the architecture design that lead to the severe reduction in spike firing, make targeted improvements, and propose Spike2Former architecture. Second, we propose normalized integer spiking neurons to solve the training stability problem of SNNs with complex architectures. We set a new state-of-the-art for SNNs in various semantic segmentation datasets, with a significant improvement of +12.7% mIoU and 5.0× efficiency on ADE20K, +14.3% mIoU and 5.2× efficiency on VOC2012, and +9.1% mIoU and 6.6× efficiency on CityScapes.

Introduction

Spiking Neural Networks (SNNs) emulate the spatiotemporal dynamics and spike-based communication of biological neurons. The former ensures the network’s representation (Maass 1997), while the spike-driven paradigm introduced by the latter allows SNNs to perform sparse computing when deployed on neuromorphic chips (Merolla et al. 2014; Davies et al. 2018), thereby benefiting from low power consumption (Roy, Jaiswal, and Panda 2019; Schuman et al. 2022). A notable example is the sensing-computing neuromorphic chip Speck, which operates at a power level as low as 0.7mW in typical visual scenes (Yao et al. 2024d).

The complex neuronal dynamics and binary activations make it challenging to train large-scale SNNs. It took a long time for the SNN field to effectively address this issue through surrogate gradient training (Wu et al. 2018; Neftci, Mostafa, and Zenke 2019a) and residual learning design (Fang et al. 2021; Hu et al. 2024b). Currently, SNNs have achieved commendable performance on simple image classification tasks (Yao et al. 2024b,a). Unfortunately, when it comes to complex visual tasks that require the use of sophisticated neural network architectures, SNNs fall short.

For instance, in image segmentation, an additional segmentation head is required alongside the backbone used for image classification, resulting in a network structure that is significantly more complex than that for classification. Simply converting complex Artificial Neural Networks (ANNs) into spiking versions or directly applying residual designs from classification tasks often leads to a notable drop in performance. Consequently, the few existing SNN models (Kim, Chough, and Panda 2022; Zhang, Fan, and Zhang 2023; Yao et al. 2024a; Su et al. 2024; Patel et al. 2021) that tackle image segmentation tasks tend to perform poorly.

This work explores the application of SNNs with more complex architectures to image segmentation tasks. Specifically, Mask2Former (Cheng et al. 2022) is a classic Transformer-based per-mask classification architecture consisting of three parts: a backbone network; a Feature Pyramid Network (FPN) pixel decoder with a multi-scale deformable transformer encoder block; and a transformer decoder block. Directly converting the Mask2Former architecture into a spiking version results in obvious performance degradation and non-convergence. To address this, we investigated which modules in the spiking Mask2Former contribute to the significant loss of information, where the spiking neurons in these modules nearly cease to fire.

The first module with severe spike degradation is the deformable attention transformer encoder block in the FPN decoder. In the vanilla Mask2Former, the query operation in the deformable attention block is inherently sparse; if the queried information consists of sparse spikes, this could result in significant information loss. To address this, we incorporate convolution blocks in encoder blocks and redesign the deformable attention blocks to preserve more effective information and energy efficiency. The second key module that leads to information loss is the final mask embedding layer. This layer is crucial as it outputs the final segmentation results; however, being at the deepest part of the network, the semantic information that reaches this point is already greatly diminished. To end this, we build an auxiliary information branch to enhance the representation of mask embedding.

Another challenge is that, beyond architectural design, spiking neurons inherently introduce information loss by converting continuous values into binary spikes. This issue has long plagued the SNN field, leading to the development

*These authors contributed equally.

†Corresponding author.

of various methods aimed at refining the membrane potential distribution, such as attention mechanism (Yao et al. 2023) and information maximization loss (Guo et al. 2022). Recently, Luo et al. proposed an integer training and spike inference method in spiking CNNs to address this challenge. However, this approach does not extend to more complex architectures like Mask2Former, which require numerical stability and precise representation in the interaction of cross-modal features (Cheng, Schwing, and Kirillov 2021; Carion et al. 2020; Li et al. 2023). To address this, we propose a novel normalization method that normalizes integers during training to facilitate training, while ensuring that the spike-driven nature during inference remains unaffected.

The proposed methods are validated on the popular datasets ADE20k, CityScapes, Pascal VOC2012 for semantic segmentation. Spike2Former far exceeds the best models in SNNs in both performance and energy consumption and is comparable to the ANNs. The main contributions are:

1) **Spike2Former**. We analysis the challenge of applying SNNs to complex architecture and propose a transformer-based image segmentation method, Spike2Former, that incorporates the Spike-driven Deformable Transformer Encoder (SDTE) and Spike-Driven Mask Embedding (SDME) module to resolve the information deficiency and improve model performance.

2) **NI-LIF Spiking Neuron**. We design a novel spiking neuron, NI-LIF, to increase the training stability and reduce the information deficiency in complex segmentation method. The NI-LIF normalizes the integer activation during training and can be equivalent to spike-driven in inference, capitalizing on the low power nature of SNNs.

3) **Performance**: The proposed Spike2Former obtains a remarkable performance improvement in popular segmentation tasks with low power consumption, demonstrating the potential of SNNs in complex scenarios. Spike2Former achieves a new state-of-the-art ADE20k (+12.7% mIoU and 5.0 \times), Pascal VOC2012 (+14.3% mIoU and 5.2 \times), CityScapes (+9.1% mIoU and 6.6 \times).

Related Works

SNN Training

Training methods in Spiking Neural Networks (SNNs) mainly fall into two categories: ANN-to-SNN conversion and direct training. While ANN-to-SNN conversion inherits from traditional neural networks, it suffers from longer time steps and limited real-time processing capabilities (Li et al. 2022; Bu et al. 2022; Wu et al. 2021). Direct training with surrogate gradients (Neftci, Mostafa, and Zenke 2019b) offers better flexibility but typically yields lower performance. Inspired by recent Integer-based Leaky Integrate-and-Fire (I-LIF) neurons, which bridge the gap between training and inference through virtual timesteps, we adopt the direct training approach for its architectural flexibility (Hu et al. 2024b; Fang et al. 2021; Wu et al. 2018). Recently, Luo et al. introduced a novel spiking neuron that trains the model with Integer activation as a virtual timesteps and converts the Integer into binary spikes by expanding virtual timesteps during inference. We take this inspiration and further improve it for

more complex application scenes.

SNN Backbone Design

SNN backbone designs can be broadly categorized into CNN-based and Transformer-based approaches. CNN-based SNNs primarily focus on improving the Spike-ResNet architecture (Zheng et al. 2021) through diverse design of residual connections such as SEW-ResNet (Fang et al. 2021) and MS-ResNet (Hu et al. 2024b). Recently, Transformer-based methods have gained prominence in SNN backbone design. Approaches such as (Wang et al. 2023; Zhou et al. 2023; Leroux, Finkbeiner, and Neftci 2023; Yao et al. 2024b,a; Hu et al. 2024a; Yao et al. 2024c) incorporate spiking neurons into self-attention mechanisms to enhance model performance. These methods are relatively straightforward and lack complex interactions, making them easier to train within the SNN framework. Additionally, Luo et al. recently introduced a spike-driven SpikeYOLO model, adapted from YOLOv8 for object detection, which achieves competitive performance. However, to mitigate information loss during feature interactions, this method simplifies the model design, failing to fully address the challenges of applying SNNs to complex architectures.

Image Segmentation

In ANNs, image segmentation can be categorized into per-pixel and per-mask classification. Per-pixel classification, often using CNNs, employs Feature Pyramid Networks (FPN) to generate segmentation masks. MaskFormer (Cheng, Schwing, and Kirillov 2021) redefined this as per-mask classification, generating binary masks and assigning them semantic classes. Mask2Former (Cheng et al. 2022) and MaskDINO (Li et al. 2023) further improve the MaskFormer with more precise refinement of learnable query. In SNNs, image segmentation remains challenging (Li et al. 2022; Kirkland et al. 2020; Patel et al. 2021). In 2022, (Kim, Chough, and Panda 2022) introduced spike-driven decoders (spike-FCN & spike-Deeplab) by converting ANN architectures to SNNs, while Spiking-CGNet (Zhang, Fan, and Zhang 2023) in 2023 developed a segmentation-specific backbone. More recently, Meta-SpikeFormer (Yao et al. 2024a) directly trained on FPN achieved competitive results on ADE20k (Zhou et al. 2017), highlighting the potential of SNN. However, the performance and energy efficiency of these methods are poor and limit their application to various scenarios.

Method

Spike2Former adopts the architecture of Mask2Former, which includes a pyramid backbone, a pixel decoder with a deformable transformer encoder, a transformer decoder, and a mask embedding module. In this section, we present the improvement in the deformable transformer encoder and mask embedding module, followed by a discussion of the newly proposed spiking neuron NI-LIF.

Information Deficiency in Query

Query features are essential in transformer-based methods (Ding et al. 2023; Wang et al. 2024; Jin et al. 2023). They are primarily used to interact with image features in the Multi-Scale Deformable Attention (MSDeformAttn) transformer

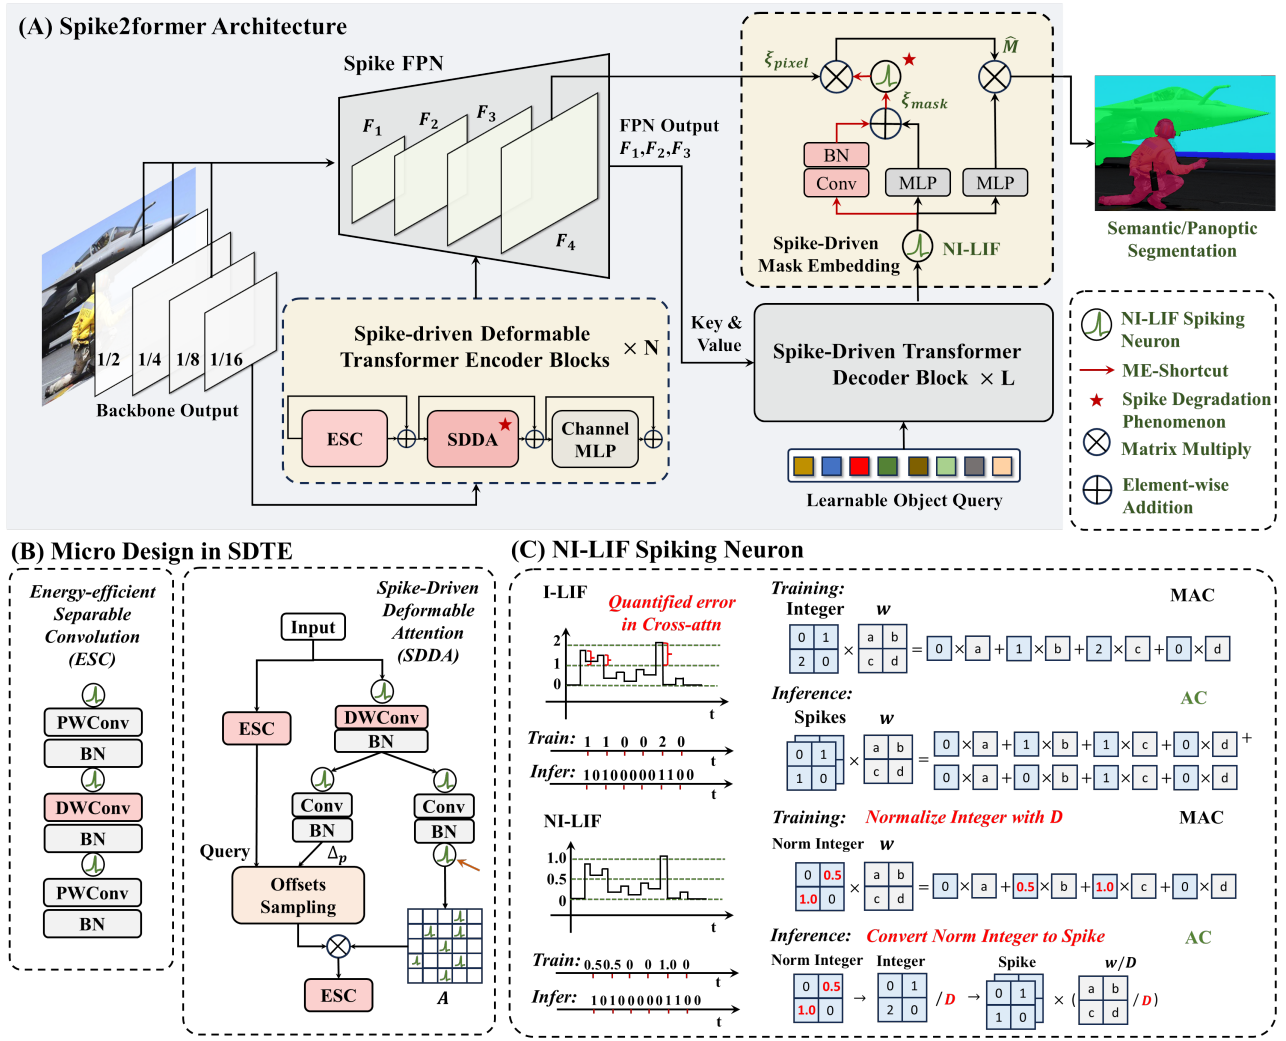


Figure 1: (A) The architecture of Spike2Former, we propose Spike-Driven Mask Embedding(SDME) and Spike-driven Deformable Transformer Encoder to introduce the transformer-based method to SNNs. (B) Micro Design in Spike-Driven Deformable Transformer Encoder(SDTE) including Spike-driven Deformable attention(SDDA) and Spike Separable Convolution(ESC). (C) Comparison between NI-LIF and I-LIF. Integer activation results in information loss. NI-LIF normalizes the integer during training to preserve information and converts them to spikes during inference with only sparse addition.

encoder and transformer decoder within Mask2Former. However, SNNs face challenges in refining queries due to the significant loss of rich semantic representations when converting query into binary spikes. This raise the question: *How can queries preserve information effectively in SNNs?* To address this, we propose modifications in the architecture design, including deformable attention blocks and the mask embedding module, which are particularly prone to bias in SNNs. The details of these modifications will be discussed in the following of this section.

Spike-driven Deformable Transformer Encoder Spike-driven Deformable Transformer Encoder(SDTE) consists of a stack of successive blocks, including an Energy-efficient Separable Convolution(ESC) module to enhance local connections, a Spike-Driven Deformable Attention(SDDA)

module that conducts deformable attention within queries, and a Channel-MLP layer to learn non-linear representations.

Energy-efficient Separable Convolution Blocks Yao et al. utilize a separable convolution to enhance the inductive bias. However, this design significantly increases energy consumption due to the direct connection of depthwise and pointwise convolutions. Thus, we propose adding spiking neuron before the second pointwise convolution for energy efficiency, denoted as $ESC(\cdot)$, and formulated as follows:

$$\mathbf{U}_{pw1} = \text{Conv}_{pw1}(\text{SN}(\mathbf{U})) \quad (1)$$

$$\mathbf{U}_{dw} = \text{Conv}_{dw}(\text{SN}(\mathbf{U}_{pw1})) \quad (2)$$

$$\mathbf{U}_{pw2} = \text{Conv}_{pw2}(\text{SN}(\mathbf{U}_{dw})) \quad (3)$$

where $\mathbf{U} \in \mathbb{R}^{\frac{H}{16} \times \frac{W}{16} \times dim}$ means the input membrane potential (dim denotes the embedding dimension), $\text{Conv}_{dw}(\cdot)$ de-

notes depthwise convolutions, and $\text{Conv}_{pw}(\cdot)$ denote point-wise convolutions.

Spike-Driven Deformable Attention The MSDeformAttn transformer encoder in Mask2Former refines features by attending to the global context, dynamically computing weights from the inputs, and utilizing deformability to adapt the receptive field size (Cavagnero et al. 2024). Although effective in ANNs, the sparse sampling strategy in deformable attention leads to information loss and unstable gradients (Fig. 2). Thus, we propose converting the attention weights into spikes instead of spiking the feature queries to better preserve the semantic information contained within the queries. Furthermore, using multi-scale features as queries significantly increases energy consumption, especially with high-resolution inputs. We recommend using single-scale deep image features with convolution blocks to enhance local connectivity, thereby improving energy efficiency while maintaining performance.

Specifically, for the calculation of attention weight and sampling offsets, we propose adding depthwise convolution (DWConv) to enhance the understanding of scene context. As shown in Fig. 1(B), given an input feature map x_g , the Spike-Driven Deformable Attention can be formulated as:

$$\text{DeformableSDSA}(\mathbf{p}_q, \mathbf{x}_g) = \sum_{g=1}^G \sum_{k=1}^K \mathbf{W}_g \mathbf{A}_{gk} \mathbf{W}'_g \cdot \mathbf{x}_g(p_0 + p_k + \Delta p_{gk}) \quad (4)$$

where G represents the total number of aggregation groups. For the g -th group, \mathbf{W}_g and \mathbf{W}'_g represent the location-irrelevant projection weights. \mathbf{A}_{gk} denotes the attention weight corresponding to the k -th sampling point within the g -th group. Δp_{gk} represents the offset for the k -th sampling location p_k within the g -th group. Subsequently, \mathbf{W}_g , \mathbf{A}_{gk} , and Δp_{gk} are calculated as follows:

$$\mathbf{W}_g = \text{ESC}(x_g), \quad (5)$$

$$x'_g = \text{BN}(\text{DWConv}(\text{SN}(x_g))), \quad (6)$$

$$\mathbf{A}_{gk} = \text{SN}(\text{BN}(\text{Conv}(\text{SN}(x'_g)))), \quad (7)$$

$$\Delta p_{gk} = \text{BN}(\text{Conv}(\text{SN}(x'_g))). \quad (8)$$

Where the SN represents the spiking neuron. We add spiking neuron for \mathbf{A}_{gk} to convert the attention weight into spike and maintain the effective information in feature query. Finally, we apply an ESC block to the input that has been sampled to the embedding dimension. The output from the SDTE is then fed into the SpikeFPN (Yao et al. 2024a) to generate per-pixel embedding.

Spike-Driven Transformer Decoder The Spike-Driven Transformer Decoder (SDTD) contains a Spike-Driven Cross-Attention (SDCA) layer, a Spike-Driven Self-Attention (SDSA) layer, and a Channel-MLP layer. The formulation of SDTD can be written as:

$$\mathbf{Q}' = \mathbf{Q} + \text{SDCA}(\mathbf{Q}, \mathbf{F}_i), i = 1, 2, 3 \quad (9)$$

$$\mathbf{Q}'' = \mathbf{Q}' + \text{SDSA}(\mathbf{Q}'), \quad (10)$$

$$\mathbf{Q}''' = \mathbf{Q}'' + \text{ChannelMLP}(\mathbf{Q}''). \quad (11)$$

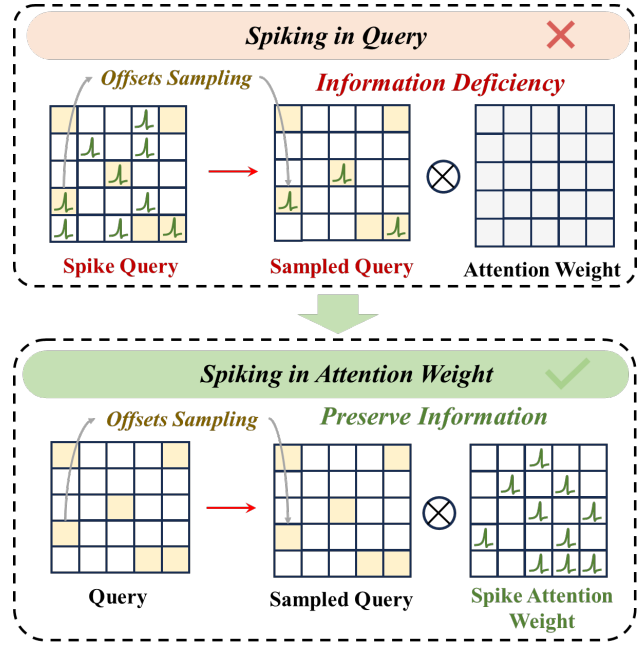


Figure 2: Illustration of offset-sampling operation in Spike-Driven Deformable Attention (SDDA). While directly spiking the query (Sampled Spike Query) leads to information loss, spiking the attention weights (Spike Attention Weight) effectively preserves the crucial query information.

where $\mathbf{Q} \in \mathbb{R}^{N \times C}$ are the N learnable query with learnable positional embedding and $\mathbf{F}_i \in \mathbb{R}^{H_i \times W_i \times B_i}$ indicate the multi-scale feature maps obtained from the pixel decoder, with $i \in \{1, 2, 3\}$.

Additionally, we replace the re-parameterization convolution (Yao et al. 2024a) with the combination of Linear and BatchNorm for energy efficiency. The formulation of SDSA can be written as:

$$\mathbf{Q}_s, \mathbf{K}_s, \mathbf{V}_s = \text{SN}(\text{BN}(\text{Conv}(\text{SN}(\mathbf{Q}_{1-1})))), \quad (12)$$

$$\mathbf{Q}_1 = \text{BN}(\text{Conv}(\text{SN}(\mathbf{Q}_s \mathbf{K}_s^T \mathbf{V}_s * \text{Scale}))) \quad (13)$$

where the $scale$ of SDSA can be re-parameterized into the spiking neuron's threshold. Note that the SDCA obtains the K_s and V_s from multi-scale feature maps \mathbf{F}_i .

Spike-Driven Mask Embedding The mask embedding module in Mask2Former (Cheng et al. 2022) uses a Multi-Layer Perceptron (MLP) to convert the per-segment embedding \mathbf{Q} to N mask embedding $\zeta_{mask} \in \mathbb{R}^{N \times C}$, where C denotes the object class and N denotes the number of query and obtain the binary mask prediction $\hat{\mathbf{M}} \in [0, 1]$ through dot product between per-segment embedding M and per-pixel embedding ζ_{pixel} . However, the rich diversity semantic information corresponding to the segment object within N mask embedding ζ_{mask} will suffer a significant reduction when converting to binary spikes and further affect the binary mask predictions. Furthermore, as shown in Fig. 3 we also found the spike degradation phenomenon in mask embedding that the spiking firing rate before the embedding operation suffers a significant loss and infect the performance.

Method	Model	mIoU (%)	Params (M)	$T \times D$	Power (mJ)	
Pascal VOC2012						
ANN	FCN(Long, Shelhamer, and Darrell 2015)	62.2	49	-	909.6	
	DeepLabv3(Chen et al. 2017)	66.7	68	-	1240.6	
	DeepLabv3+(Chen et al. 2018)	77.2	41	-	326.6	
ANN2SNN	Spike Calibration(Li et al. 2022)	59.6	-	256×1	-	
SNN	SpikeFCN(Kim, Chough, and Panda 2022)	9.9	50	20×1	383.5	
	SpikeDeepLab(Kim, Chough, and Panda 2022)	22.3	68	20×1	523.2	
	SpikeFPN(Yao et al. 2024a)		58.1	17	1×1	81.4
			61.1	59	4×1	179.4
		61.8	34	1×2	50.6	
	Spike2Former (Ours)		62.1	34	2×2	98.3
			75.1	34	1×4	63.0
		75.4	34	4×4	221.9	
ADE20k						
ANN	DeepLabv3+(Chen et al. 2018)	42.7	41	-	818.8	
	MaskFormer(Cheng, Schwing, and Kirillov 2021)	44.5	41	-	243.8	
		46.7	42	-	253	
	Mask2Former(Cheng et al. 2022)	47.2	44	-	326.6	
47.7		47	-	340.6		
SNN	SpikeFPN(Yao et al. 2024a)	33.6	17	4×1	88.1	
	Spike2Former (Ours)	46.3	34	1×4	68.0	
CityScapes						
ANN	DeepLabv3+(Chen et al. 2018)	79.1	65	-	1241.3	
	MaskFormer(Cheng, Schwing, and Kirillov 2021)	78.5	41	-	463.5	
		-	-	-	-	
	Mask2Former(Cheng et al. 2022)	79.6	44	-	601.6	
80.6		47	-	623.4		
SNN	SCGNet-L(Zhang, Fan, and Zhang 2023)	66.5	1.9	4×1	24.6	
	Spike2Former (Ours)	75.2	34	1×4	93.8	

Table 1: The semantic segmentation results on ADE20k, CityScapes, and Pascal VOC2012 are presented. We obtain +14.3%, +12.7%, and +9.1% mIoU on Pascal VOC, ADE20k, and CityScapes separately compared with the state-of-the-art SNN method. The results of MaskFormer and Mask2Former are presented with R50 (Above) and Swin-t (Below) as backbone, respectively. Power is the estimation of energy consumption as formulated in (Yao et al. 2024b).

Thus, we proposed a Membrane Embedding Shortcut (ME-Shortcut) to connect the membrane potential from the output of the transformer decoder to mask embedding, which significantly ensures the training stability and enhances the effective representation of N mask embedding ζ_{mask} .

As shown in Fig. 1(A), we build a ME-Shortcut with a channel convolution (Howard et al. 2019) parallel with the MLP layer to enhance the precise representation of ζ_{mask} and reduce the spike degradation phenomenon. In the following, we convert the ζ_{mask} to spikes before multiplying to pixel embedding ζ_{pixel} for spike-driven. Finally, we obtain the binary mask predictions by:

$$\zeta_{mask} = SN(MLP(SN(\mathbf{Q})) + w_s * BN(Conv(SN(\mathbf{Q})))), \quad (14)$$

$$\hat{\mathbf{M}} = \zeta_{mask} \cdot \zeta_{pixel}. \quad (15)$$

where the $SN(\cdot)$ is the spiking neuron layer, w_s is a learnable weight initialized with 1.0 to select key feature.

NI-LIF Spiking Neuron

The spiking neuron layer integrates spatio-temporal information into the membrane potential and then converts it into binary spikes for spike-driven computing in the following layer. Different from the image classification task, dense prediction necessitates a higher demand for numerical stability. Recent work (Luo et al. 2024) shows a performance increase in Object Detection with Integer Leaky Integrate-and-Fire (I-LIF). However, when we attempt to extend the I-LIF to a more complex architecture like Mask2Former, we found that gradient instability occurs, especially in the spike-driven transformer decoder layer. We argue that the potential reason is that the interaction between object queries and multi-scale feature maps requires numerical stability. However, the large integer value in multi-scale feature maps obtained from I-LIF failed to quantify the precise value and brought quantization error and information loss. Therefore, we intend to tackle this issue by normalizing the $S[t]$ with the virtual timesteps D to

maintain the gradient stable and preserve the information.

We improved I-LIF through normalizing integer-value with virtual timesteps to enhance the numerical and gradient stability, as depicted in Fig. 1(C), whose dynamics are:

$$U[t] = H[t - 1] + X[t], \quad (16)$$

$$S[t] = \text{Clip}(\text{round}(U[t]), 0, D)/D, \quad (17)$$

$$H[t] = \beta(U[t] - S[t] \times D). \quad (18)$$

where $X[t]$ is the spatial input current at timestep t , $U[t]$ denotes the membrane potential that integrates the temporal information $H[t - 1]$ and spatial $X[t]$. $\text{round}(\cdot)$ is a round operation, $\text{Clip}(x, \min, \max)$ implies clipping the input x to $[\min, \max]$, and D is a hyper-parameter to emit the maximum integer value.

Inference As shown in Fig. 1(C), I-LIF (Luo et al. 2024) converts the integer spikes into binary spikes during inference stages, and the input to the spiking neuron at $l + 1$ layer can be described as:

$$X^{l+1}[t] = W^l S^l[t] \quad (19)$$

We also extend the T time step to $T \times D$, and convert the normalized value $S^l[t]$ to a spike sequence $\{S^l[t, d]\}_{d=1}^D$, which can be written as:

$$\sum_{d=1}^D S^l[t, d] = S^l[t] \times D \quad (20)$$

Therefore, the input of the neuron at $l + 1$ layer can be calculated by:

$$X^l[t] = \sum_{d=1}^D (W^l (S^l[t, d] \times D)), \quad (21)$$

$$X^l[t] = \sum_{d=1}^D (W_D^l (S^l[t, d])).$$

where the $W_D^l = W^l \times D$. The spike sequence $S^l[t, d]$ only contains 0/1 after being multiplied by the quantization step, and the MAC (Multiply-Accumulation operations) can be converted to sparse AC (Accumulation operations), which enables sparse addition during inference.

Experiment

Dataset We conduct semantic segmentation on ADE20k (Zhou et al. 2017), CityScapes (Cordts et al. 2016), and Pascal VOC2012 (Everingham et al. 2010) datasets. The details of the training strategy are shown in Tab. 2.

Training Settings We present all our main semantic segmentation results in mean Intersection over Union (mIoU) under single scale inference setting. We use Meta-Spikeformer (Params:15M) (Yao et al. 2024a) as our backbone, which is pre-trained on ImageNet-1k for 200 epochs. More training details can be found in the Appendix.

Setting	ADE20k	CityScapes	Pasval VOC2012
Input Size	512×512	512×1024	512×512
Learning Rate	2e-4	2e-3	2e-3
Optimizer	AdamW	AdamW	AdamW
Training Steps	160k	90k	80k

Table 2: Hyper-parameters setting in Spike2Former.

Experiment Results

In Tab. 1, we comprehensively compare Spike2Former with other ANN and SNN methods in mIoU, parameters, and power. The proposed Spike2Former significantly improves the performance upper bound of SNNs on three public datasets. We obtain **46.3% mIoU**, **75.2% mIoU**, and **75.4% mIoU** in ADE20k, CityScapes, and Pascal VOC2012 which is **+14.3%**, **+9.1%**, and **+14.3%** higher than the previous state-of-the-art SNN method (Yao et al. 2024a; Zhang, Fan, and Zhang 2023), respectively. Spike2Former also demonstrates significant advantages over the existing SNNs in terms of energy consumption: Spike2Former vs. SpikeFPN (Yao et al. 2024a): mIoU **44.6%** vs. 33.6%; Power: **68mJ** vs. 88.1mJ in ADE20k dataset and Spike2Former vs. SpikeFPN: mIoU **75.4%** vs. 61.1%; Power **63.0mJ** vs. 179.4mJ in Pascal VOC2012 dataset. Moreover, the performance gap between SNNs and ANNs is significantly narrowed. Spike2Former got **+1.2%** mIoU compared with MaskFormer(R50) and is competitive with the current classical ANN architecture Mask2Former(R50) in ADE20k dataset, which is 46.3% vs. 44.5% vs. 47.2% mIoU in mIoU while the energy consumption is much lower for **3.58×** and **4.80×** energy efficiency.

Ablation Study

We conduct ablation studies on the various components of Spike2Former to evaluate the contribution of each parts.

Spike-Driven Deformable Transformer Encoder Tab. 3 highlights the performance of our spike-driven Deformable Transformer Encoder (SDTE). As shown in Fig. 2, directly spiking query features reduces mIoU by 3.2%, likely due to excess retention of attention weights, diminishing effective information of query. Our SDTE improves mIoU by 2.5% over SpikeFPN (Yao et al. 2024a) without Transformer encoder and by 1.1% over a vanilla Transformer. While using multi-scale features as queries, as in Mask2Former, yields 46.7% mIoU, it doubles energy consumption (68.0mJ vs. 136.5mJ), prompting us to prioritize single-scale queries for energy efficiency.

Information Deficiency in Query Preserving query information is crucial for effectively applying SNNs to Mask2Former. As demonstrated in Tab. 3, removing the ME-Shortcut connection leads to a significant performance drop (approximately $-4.5%$ mIoU). Fig. 3 visualizes the average binary mask prediction for each query over the validation set of ADE20k, revealing that incorporating shortcut connections (ME-Shortcut) results in more sophisticated and distinct patterns within the mask embedding, indicating higher-

Ablation	Method	Power (mJ)	mIoU (%)
Spike2Former(Baseline)		68.0	46.3
ANN Conversion	Mask2Former(Spiking Version)	-	5.2
Spike-Driven Mask Embedding	Vanilla Spike Mask Embedding	67.9	41.8(-4.5)
	w/o. Conv in ME-Shortcut	67.9	45.7(-0.6)
	ME-Shortcut \rightarrow Membrane Shortcut	67.8	45.0(-1.3)
Spike-driven Deformable Transformer Encoder	w/o. SD Deformable Transformer Encoder	65.2	43.8(-2.5)
	w/o. ESC in SD Deformable attention	64.5	45.6(-0.7)
	w/o. DWConv in SD Deformable attention	67.8	45.7(-0.6)
	Spiking in Query	67.8	43.1(-3.2)
	Multi-Scale feature as query	136.5	46.7(+0.4)
	Vanilla Transformer Encoder	66.9	44.9(-1.4)
Normalize Method in NI-LIF	I-LIF	68.1	37.9(-8.4)
	Norm Only in Cross-Attn	68.1	43.2(-3.1)
	Norm with Const 8	67.4	44.7(-1.3)
	Norm with Sigmoid&L2	-	*

Table 3: We conduct ablation on the proposed Spike-Driven Mask Embedding (SDME), Spike-driven Deformable Transformer Encoder (SDTE) and Normalized Integer-LIF (NI-LIF) in ADE20k dataset. We set $T \times D = 1 \times 4$ and modify just one modification of the baseline to test how the power and performance vary. * Does not converge. The ablation on ANN indicates the direct conversion of Mask2Former with SNNs.

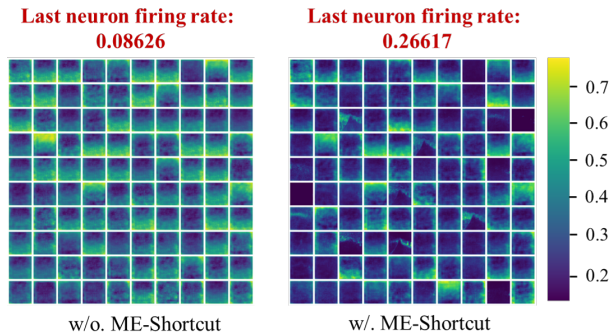


Figure 3: Segment Patterns Represented by Queries. We Average its corresponding segment predictions over the whole validation set of ADE20k. All the predictions are resized to the resolution of 200×200 for illustration purposes.

quality binary mask predictions. Moreover, ME-Shortcut increases the firing rate of the final spiking neuron threefold (from 0.08626 to 0.26617). This suggests that the performance drop without ME-Shortcut is attributable to spike degradation, which our proposed shortcut effectively mitigates. Despite its conceptual simplicity, ME-Shortcut is a key design for the successful application of SNNs to complex architecture.

NI-LIF Spiking Neuron Tab. 3 shows that directly applying I-LIF (Luo et al. 2024) to Spike2Former results in sub-optimal performance (37.9% mIoU compared to 46.3%), primarily due to information deficiency in the cross-attention layer of transformer decoder. Normalizing activations within the cross-attention layer alone yields a substantial +5.3% mIoU improvement compared with I-LIF, highlighting the need for numerical stability and precise representation of

Transformer-based model. Furthermore, when applying the NI-LIF to the whole network, the performance increases 8.4%. This improvement demonstrates the adaptability of NI-LIF in complex architecture and the effectiveness of NI-LIF in reducing the quantization error and enhancing the precise representation of features.

Further analysis in Tab. 1 (Pascal VOC2012) reveals the impact of varying timesteps (T) and quantization steps (D). Increasing timesteps from T=1 to T=2 (with D=2) raises energy consumption from 50.6 mJ to 98.3 mJ with only 0.3% mIoU. Additionally, increasing quantization steps from D=2 to D=4 improves performance from 61.8% to 75.1% mIoU, with a moderate power increase from 50.6 mJ to 63.0 mJ. This suggests that increasing quantization steps can enhance the performance while maintaining reasonable energy consumption.

Conclusion

This work reduces the performance gap between Spiking Neural Networks (SNNs) and Artificial Neural Networks (ANNs) in image segmentation. The proposed Spike2Former introduces architectural innovations and spiking neuron optimizations, including two key modifications to address information deficiency. The Normalized Integer LIF (NI-LIF) mitigates information loss and enhances training stability by converting normalized integers into binary spikes. Spike2Former achieves state-of-the-art performance on three benchmark datasets, highlighting the potential of SNNs for complex segmentation tasks. Our analysis of spike degradation and information deficiency emphasizes the need to reduce information loss in SNNs for advanced architectures. This work lays a foundation for extending SNNs to dense prediction tasks with sophisticated designs.

Acknowledgments

This work was partially supported by CAS Project for Young Scientists in Basic Research (YSBR116), National Distinguished Young Scholars (62325603), National Natural Science Foundation of China (62236009, U22A20103, 62441606, 62406322), Beijing Science and Technology Plan (Z241100004224011), Beijing Natural Science Foundation for Distinguished Young Scholars (JQ21015), China Postdoctoral Science Foundation (GZB20240824, 2024M753497), and Natural Science Foundation of China under Grant (62394314, 62394311).

References

- Bu, T.; Fang, W.; Ding, J.; DAI, P.; Yu, Z.; and Huang, T. 2022. Optimal ANN-SNN Conversion for High-accuracy and Ultra-low-latency Spiking Neural Networks. In *International Conference on Learning Representations*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Cavagnero, N.; Rosi, G.; Cuttano, C.; Pistilli, F.; Ciccone, M.; Averta, G.; and Cermelli, F. 2024. Pem: Prototype-based efficient maskformer for image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15804–15813.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4): 834–848.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 801–818.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshik, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1290–1299.
- Cheng, B.; Schwing, A.; and Kirillov, A. 2021. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34: 17864–17875.
- Cordts, M.; Omran, M.; Ramos, S.; Rehfeld, T.; Enzweiler, M.; Benenson, R.; Franke, U.; Roth, S.; and Schiele, B. 2016. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3213–3223.
- Davies, M.; Srinivasa, N.; Lin, T.-H.; Chinya, G.; Cao, Y.; Choday, S. H.; Dimou, G.; Joshi, P.; Imam, N.; Jain, S.; et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1): 82–99.
- Ding, H.; Wang, B.; Kang, G.; Li, W.; He, C.; Zhao, Y.; and Wei, Y. 2023. DropQueries: A Simple Way to Discover Comprehensive Segment Representations. *IEEE Transactions on Multimedia*.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88: 303–338.
- Fang, W.; Yu, Z.; Chen, Y.; Huang, T.; Masquelier, T.; and Tian, Y. 2021. Deep residual learning in spiking neural networks. *Advances in Neural Information Processing Systems*, 34: 21056–21069.
- Guo, Y.; Chen, Y.; Zhang, L.; Liu, X.; Wang, Y.; Huang, X.; and Ma, Z. 2022. IM-loss: information maximization loss for spiking neural networks. *Advances in Neural Information Processing Systems*, 35: 156–166.
- Howard, A.; Sandler, M.; Chu, G.; Chen, L.-C.; Chen, B.; Tan, M.; Wang, W.; Zhu, Y.; Pang, R.; Vasudevan, V.; et al. 2019. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1314–1324.
- Hu, J.; Yao, M.; Qiu, X.; Chou, Y.; Cai, Y.; Qiao, N.; Tian, Y.; XU, B.; and Li, G. 2024a. High-Performance Temporal Reversible Spiking Neural Networks with $\mathcal{O}(L)$ Training Memory and $\mathcal{O}(1)$ Inference Cost. In *Forty-first International Conference on Machine Learning*.
- Hu, Y.; Deng, L.; Wu, Y.; Yao, M.; and Li, G. 2024b. Advancing Spiking Neural Networks Toward Deep Residual Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 1–15.
- Jin, S.; Li, S.; Li, T.; Liu, W.; Qian, C.; and Luo, P. 2023. You Only Learn One Query: Learning Unified Human Query for Single-Stage Multi-Person Multi-Task Human-Centric Perception. *arXiv preprint arXiv:2312.05525*.
- Kim, Y.; Chough, J.; and Panda, P. 2022. Beyond classification: Directly training spiking neural networks for semantic segmentation. *Neuromorphic Computing and Engineering*, 2(4): 044015.
- Kirkland, P.; Di Caterina, G.; Soraghan, J.; and Matich, G. 2020. SpikeSEG: Spiking segmentation via STDP saliency mapping. In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Leroux, N.; Finkbeiner, J.; and Neftci, E. 2023. Online transformers with spiking neurons for fast prosthetic hand control. In *2023 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 1–6. IEEE.
- Li, F.; Zhang, H.; Xu, H.; Liu, S.; Zhang, L.; Ni, L. M.; and Shum, H.-Y. 2023. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3041–3050.
- Li, Y.; He, X.; Dong, Y.; Kong, Q.; and Zeng, Y. 2022. Spike calibration: Fast and accurate conversion of spiking neural network for object detection and segmentation. *arXiv preprint arXiv:2207.02702*.

- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440.
- Luo, X.; Yao, M.; Chou, Y.; Xu, B.; and Li, G. 2024. Integer-Valued Training and Spike-Driven Inference Spiking Neural Network for High-performance and Energy-efficient Object Detection. *arXiv preprint arXiv:2407.20708*.
- Maass, W. 1997. Networks of spiking neurons: the third generation of neural network models. *Neural Networks*, 10(9): 1659–1671.
- Merolla, P. A.; Arthur, J. V.; Alvarez-Icaza, R.; Cassidy, A. S.; Sawada, J.; Akopyan, F.; Jackson, B. L.; Imam, N.; Guo, C.; Nakamura, Y.; et al. 2014. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345(6197): 668–673.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019a. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6): 51–63.
- Neftci, E. O.; Mostafa, H.; and Zenke, F. 2019b. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6): 51–63.
- Patel, K.; Hunsberger, E.; Batir, S.; and Elias-Smith, C. 2021. A spiking neural network for image segmentation. *arXiv preprint arXiv:2106.08921*.
- Roy, K.; Jaiswal, A.; and Panda, P. 2019. Towards spike-based machine intelligence with neuromorphic computing. *Nature*, 575(7784): 607–617.
- Schuman, C. D.; Kulkarni, S. R.; Parsa, M.; Mitchell, J. P.; Kay, B.; et al. 2022. Opportunities for neuromorphic computing algorithms and applications. *Nature Computational Science*, 2(1): 10–19.
- Su, Q.; He, W.; Wei, X.; Xu, B.; and Li, G. 2024. Multi-scale full spike pattern for semantic segmentation. *Neural Networks*, 176: 106330.
- Wang, P.; Cai, Z.; Yang, H.; Swaminathan, A.; Manmatha, R.; and Soatto, S. 2024. Mixed-Query Transformer: A Unified Image Segmentation Architecture. *arXiv preprint arXiv:2404.04469*.
- Wang, Q.; Zhang, T.; Han, M.; Wang, Y.; Zhang, D.; and Xu, B. 2023. Complex dynamic neurons improved spiking transformer network for efficient automatic speech recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 102–109.
- Wu, J.; Xu, C.; Han, X.; Zhou, D.; Zhang, M.; Li, H.; and Tan, K. C. 2021. Progressive tandem learning for pattern recognition with deep spiking neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 7824–7840.
- Wu, Y.; Deng, L.; Li, G.; Zhu, J.; and Shi, L. 2018. Spatio-temporal backpropagation for training high-performance spiking neural networks. *Frontiers in Neuroscience*, 12: 331.
- Yao, M.; Hu, J.; Hu, T.; Xu, Y.; Zhou, Z.; Tian, Y.; XU, B.; and Li, G. 2024a. 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*.
- Yao, M.; Hu, J.; Zhou, Z.; Yuan, L.; Tian, Y.; Xu, B.; and Li, G. 2024b. Spike-driven transformer. *Advances in Neural Information Processing Systems*, 36.
- Yao, M.; Qiu, X.; Hu, T.; Hu, J.; Chou, Y.; Tian, K.; Liao, J.; Leng, L.; Xu, B.; and Li, G. 2024c. Scaling Spike-driven Transformer with Efficient Spike Firing Approximation Training. *arXiv preprint arXiv:2411.16061*.
- Yao, M.; Richter, O.; Zhao, G.; Qiao, N.; Xing, Y.; Wang, D.; Hu, T.; Fang, W.; Demirci, T.; De Marchi, M.; Deng, L.; Yan, T.; Nielsen, C.; Sheik, S.; Wu, C.; Tian, Y.; Xu, B.; and Li, G. 2024d. Spike-based dynamic computing with asynchronous sensing-computing neuromorphic chip. *Nature Communications*, 15(1): 4464.
- Yao, M.; Zhao, G.; Zhang, H.; Hu, Y.; Deng, L.; Tian, Y.; Xu, B.; and Li, G. 2023. Attention spiking neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8): 9393–9410.
- Zhang, H.; Fan, X.; and Zhang, Y. 2023. Energy-efficient spiking segmenter for frame and event-based images. *Biomimetics*, 8(4): 356.
- Zheng, H.; Wu, Y.; Deng, L.; Hu, Y.; and Li, G. 2021. Going deeper with directly-trained larger spiking neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 11062–11070.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 633–641.
- Zhou, Z.; Zhu, Y.; He, C.; Wang, Y.; YAN, S.; Tian, Y.; and Yuan, L. 2023. Spikformer: When Spiking Neural Network Meets Transformer. In *The Eleventh International Conference on Learning Representations*.