

Rethinking High-speed Image Reconstruction Framework with Spike Camera

Kang Chen^{1,2}, Yajing Zheng^{1,2*}, Tiejun Huang^{1,2,3}, Zhaofei Yu^{1,3*}

¹School of Computer Science, Peking University

²State Key Laboratory for Multimedia Information Processing, Peking University

³Institute for Artificial Intelligence, Peking University

mrchenkang@stu.pku.edu.cn, {yj.zheng,tjhuang,yuzf12}@pku.edu.cn

Abstract

Spike cameras, as innovative neuromorphic devices, generate continuous spike streams to capture high-speed scenes with lower bandwidth and higher dynamic range than traditional RGB cameras. However, reconstructing high-quality images from the spike input under low-light conditions remains challenging. Conventional learning-based methods often rely on the synthetic dataset as the supervision for training. Still, these approaches falter when dealing with noisy spikes fired under the low-light environment, leading to further performance degradation in the real-world dataset. This phenomenon is primarily due to inadequate noise modelling and the domain gap between synthetic and real datasets, resulting in recovered images with unclear textures, excessive noise, and diminished brightness. To address these challenges, we introduce a novel spike-to-image reconstruction framework SpikeCLIP that goes beyond traditional training paradigms. Leveraging the CLIP model’s powerful capability to align text and images, we incorporate the textual description of the captured scene and unpaired high-quality datasets as the supervision. Our experiments on real-world low-light datasets U-CALTECH and U-CIFAR demonstrate that SpikeCLIP significantly enhances texture details and the luminance balance of recovered images. Furthermore, the reconstructed images are well-aligned with the broader visual features needed for downstream tasks, ensuring more robust and versatile performance in challenging environments.

Code — <https://github.com/chenkang455/SpikeCLIP>

Introduction

Spike cameras (Huang et al. 2023), drawing inspiration from the retina, represent an innovative method for capturing visual information by continuously converting light intensity into the binary spike stream. Unlike traditional cameras, which have a fixed exposure period and low capture frequency, spike cameras operate at an ultra-high speed with each pixel continuously monitoring the incoming light. Once the accumulated light reaches a certain threshold, a spike is emitted. This sampling mechanism enables the spike

camera to record visual features at a remarkable rate of 40kHz, making it exceptionally suited for the high-speed imaging task. Additionally, spike cameras offer advantages such as lower bandwidth usage and higher dynamic range compared to traditional high-speed cameras like the Phantom. Nevertheless, the critical challenge of spike-to-image conversion remains to be solved since the binary quantized neuromorphic bitstream produced by the spike camera is not perceivable by human beings.

In recent years, various methods have been developed to map spike inputs to high-quality images. Early model-based approaches, designed based on the spike sampling principle (Zhu et al. 2019), synaptic plasticity (Zheng et al. 2021), and retina model (Zhu et al. 2020), struggled with generalization and consistent reconstruction quality due to numerous adjustable parameters. To address these issues, learning-based methods like WGSE (Zhang et al. 2023) trained the convolutional network on the synthetic dataset with spike-image pairs, but the performance of these approaches degrades under the real-world low-light environment due to the dataset domain gap problem as illustrated in Fig. 1(a). Self-supervised methods have been explored as a solution for this challenge, but existing frameworks like SSML (Chen et al. 2022) and SJRE (Chen, Yu, and Huang 2023) heavily rely on pseudo-labels from TFI (Zhu et al. 2019), which significantly degrade in quality under the low-light scene, leading to poor reconstruction results, as shown in Fig. 1(b).

To sum up, existing methods for spike-based image reconstruction fail to accurately recover sharp texture features under the low-light scene, highlighting the need for a novel paradigm that goes beyond previous approaches. This leads to a critical question: *what other signals can be used to supervise the spike-based image reconstruction network training besides the sharp image?*

To this end, we propose the SpikeCLIP that leverages the text description as the supervision signal for the spike-based image reconstruction task. We utilize the Contrastive Language-Image Pre-training (CLIP) (Radford et al. 2021) model to extract the textual category information of the captured object and the feature representation of well-recovered images from our constructed high-quality image dataset. Benefiting from these designs, our proposed SpikeCLIP can reconstruct texture-rich, noise-free, and brightness-balanced images as shown in Fig. 1(c). Our proposed framework con-

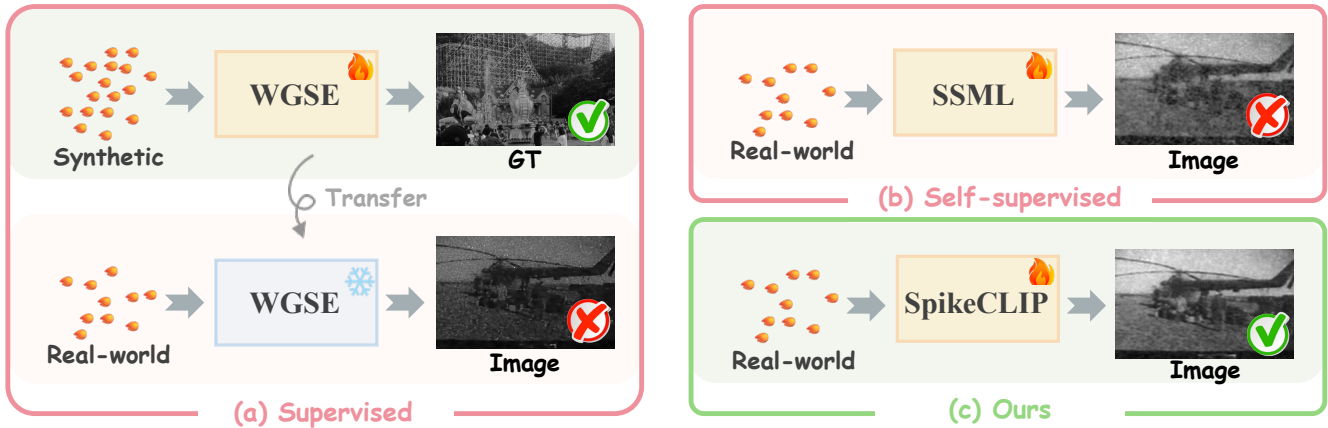


Figure 1: Illustration of the advantages of our method. While previous learning-based approaches struggle with real-world data under extreme conditions, such as low-light scenarios, our proposed SpikeCLIP successfully reconstructs high-quality images.

sists of three stages: (1) *Coarse Reconstruction*, we utilize a lightweight network to map the spike input to an initial reconstruction output; (2) *Prompt Learning*, learnable prompts are optimized to capture the distribution of high-quality and low-quality recovered images; (3) *Refinement*, we leverage class-label features and learned prompts of high-quality images to enhance the network restoration performance further. We conduct quantitative and qualitative analyses on the U-CALTECH and U-CIFAR datasets, with experiments showing that our proposed SpikeCLIP achieves remarkable restoration performance under low-light conditions. The main contributions of this paper are as follows:

- We introduce a novel spike-based image reconstruction framework tailored to the low-light environment, which leverages the CLIP model to supervise the network training by the class label of the captured object and the features of high-quality images.
- We design a high-quality image generation pipeline and demonstrate that a lightweight reconstruction network is sufficient for the spike-to-image task when the supervision signal is weak.
- Our method surpasses previous methods in reconstruction quality on the U-CALTECH and U-CIFAR datasets.

Related Works

Spike Camera Image Reconstruction

Spike-based image reconstruction methods can be roughly categorized into three types: model-based methods, supervised methods, and self-supervised methods.

Model-based methods. Model-based methods aim to build the relationship between the spike stream input and the sharp image output based on the spike camera sampling mechanism and biological principles. Zhu et al. (2019) proposed Texture from Play Back (TFP) and Texture from ISI (TFI) two methods to reconstruct the sharp frame from spike streams based on virtual exposure imaging principles and inter-spike interval analysis. Zheng et al. (2021) introduced

a spike reconstruction framework based on the principles of short-term synaptic plasticity by exploring the temporal regularity of the spike stream.

Supervised methods. Supervised learning methods are designed to train networks on paired synthetic datasets, which are subsequently used for inference in real-world scenarios. She and Qing (2022) introduced SpikeFormer, a hierarchical architecture encoder that progressively exploits multi-scale temporal and spatial features to reconstruct dynamic scenes from binary spike streams. Zhu et al. (2023) proposed a Recurrent Spike-based Image Restoration framework designed to reconstruct sharp frames from spike inputs under general illumination conditions. Zhao et al. (2024b) boost image reconstruction from the perspective of the spike fluctuations and design a spike representation module for fluctuation suppression.

Self-supervised methods. Self-supervised methods are proposed to overcome the domain-gap problem in supervised learning, enabling direct training in real-world scenarios without ground truth. Chen et al. (2022) constructed a self-supervised framework based on distillation learning, where the student model learns from the pseudo-labels provided by the blind-spot self-supervised denoising network (BSN). Chen, Yu, and Huang (2023) proposed a joint optimization framework for the BSN-based spike reconstruction network and the optical flow network, thus providing the mutual constraint for training in the real world.

CLIP in Low-level Vision

CLIP (Radford et al. 2021) has been widely applied in various downstream recognition tasks owing to its robust capability to build the relationship between textual and visual features. However, there remains significant potential for exploring its application in low-level vision. Yang et al. (2024) leverage the CLIP model to estimate accurate blur maps from dual-pixel pairs in an unsupervised manner to aid in enhancing the quality of deblurred images. Liang et al. (2023) introduces a novel unsupervised low-light image enhance-

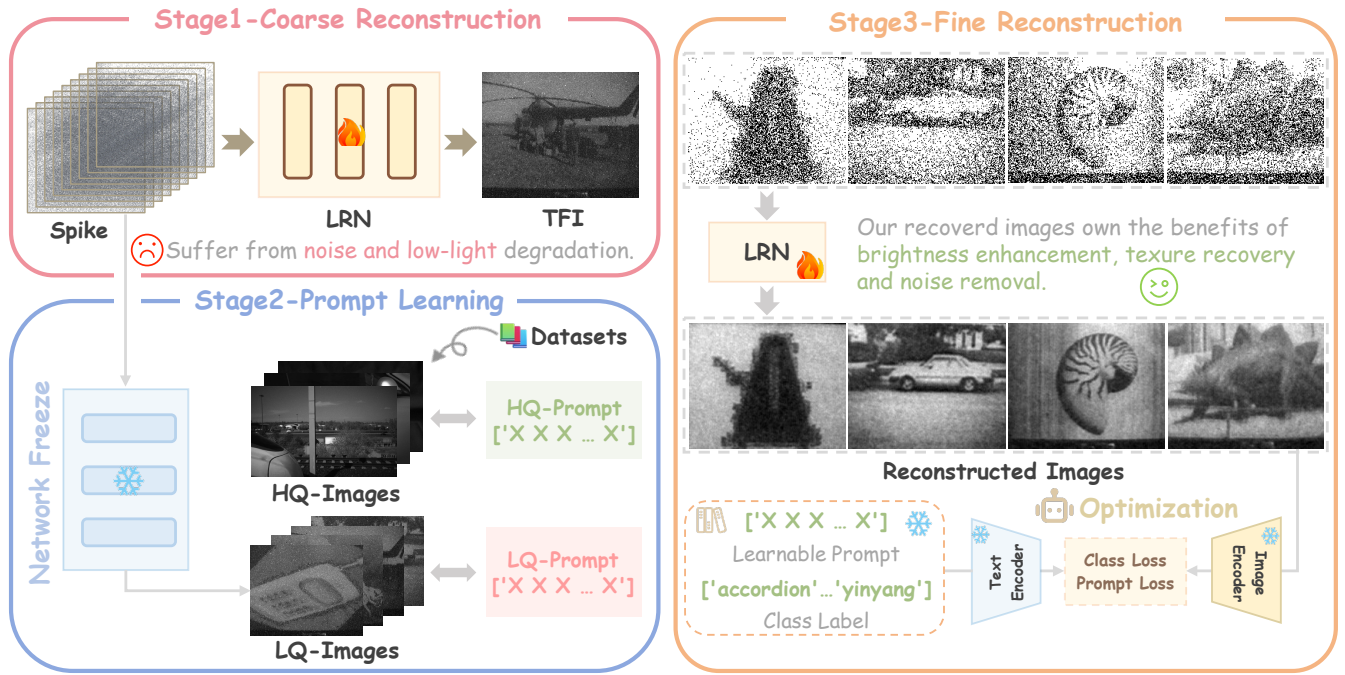


Figure 2: The overall framework of our three-stage spike-based image reconstruction method.

ment framework CLIP-LIT, which utilizes a prompt learning framework to enhance backlit images by imposing the network to learn the bright image features represented by the learnable prompt. Cheng, Liang, and Tan (2024) proposed an asymmetrical encoder-decoder denoising network, focusing on transfer CLIP for generalizable image denoising. Cho et al. (2023) designed a CLIP-based framework that leverages class prompt features to perform event reconstruction and classification simultaneously without the requirement of the labels.

Preliminaries

Spike Camera Principle

In this subsection, we detail the working principle of the spike camera. Initially, the camera photon receptor captures incoming photons and converts the light intensity into a voltage signal, which is continuously accumulated by an integrator. Whenever the voltage of the integrator exceeds the predefined threshold Θ , the signal is reset and the spike camera emits a spike. Mathematically, the spike generation process can be formulated as follows:

$$\int_0^t I(s)ds \geq \Theta, \quad (1)$$

where t represents the spike firing time and $I(s)$ denotes the latent sharp frame at time s . The spike camera employs an asynchronous triggering mechanism with synchronous read-out. Therefore, capturing a scene for T , we can obtain a spike stream with K frames, denoted as $\mathcal{S} \in \{0, 1\}^{K \times H \times W}$ with H and W represent the height and width respectively.

Motivation Clarification

Given the spike stream \mathcal{S} , our objective is to train the spike-based image **Reconstruction Network (Recon-Net)** under the low-light scenario, *i.e.*, establishing the mapping from the sparse spike input to the sharp image I :

$$I = \text{Recon-Net}(\mathcal{S}; \theta), \quad (2)$$

where θ represents the network parameters to be learned.

To optimize the parameters of the Recon-Net, previous supervised methods (Zhao et al. 2021; Zhang et al. 2023; Zhao et al. 2024b) utilize spike-sharp pairs for training, which can be mathematically represented as follows:

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\text{Recon-Net}(\mathcal{S}; \theta), I_{\text{gt}}), \quad (3)$$

with \mathcal{L} denoting the loss function designed to measure the difference between the reconstructed image and the ground truth sharp image I_{gt} , and θ^* represents the optimized network parameters. While SSML (Chen et al. 2022) is the self-supervised method, its framework remains largely the same, except that it substitutes the ground-truth image I_{gt} with the TFI reconstruction pseudo-label and the Recon-Net is the self-supervised denoising network BSN.

In this paper, we propose a new learning-based framework where the supervision signal is not the ground-truth sharp image but rather the class label and unpaired high-quality datasets encoded by the CLIP text encoder, which can be formulated as follows:

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\text{Recon-Net}(\mathcal{S}; \theta), cls, D), \quad (4)$$

where cls represents the class label of the spike stream \mathcal{S} and D represents the unpaired high-quality datasets.

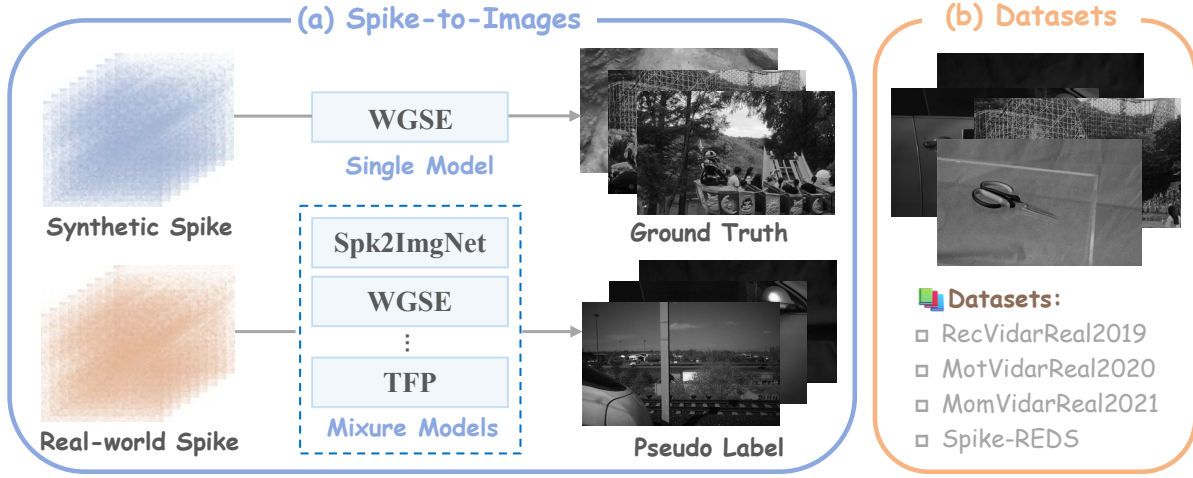


Figure 3: The framework of our designed HQ images generation pipeline.

Methods

Framework Overview

Our framework consists of coarse reconstruction, prompt learning, and fine reconstruction three stages as shown in Fig. 2. In the coarse reconstruction, our designed network is trained to map the spike input to the TFI result. In prompt learning, we design a high-quality image generation pipeline and utilize learnable prompts to capture the distributions of high-quality (HQ) and low-quality (LQ) images. In the fine reconstruction, the network is fine-tuned to enhance the final output, guided by class-label features and the prompts learned from HQ images.

Coarse Reconstruction

TFP and TFI (Zhu et al. 2019) are both effective and straightforward spike reconstruction algorithms, which can serve as the supervision signal for coarse reconstruction as in Eq. (3). However, TFP relies on a virtual exposure window, which requires a relatively long window in low-light scenes to capture sufficient spike information, resulting in severe motion blur. Conversely, TFI is not constrained by this window, making it suitable for low-light image reconstruction.

Given the spike stream \mathcal{S} over the duration T , TFI initially calculates the Inter-Spike Interval (ISI) by searching for the spike firing time t^- and t^+ before and after the middle time of each pixel, formulated as:

$$\text{ISI}_t = \begin{cases} t^+ - t^-, & \text{if } \exists \{t^+, t^-\} \in [0, T], \\ \infty, & \text{otherwise.} \end{cases} \quad (5)$$

After calculating the ISI, the TFI image reconstruction result can be obtained based on the spike sampling principle depicted in Eq. (1), *i.e.*, $I_{\text{tfi}} = \Theta / \text{ISI}$.

Regarding the architecture of the Recon-Net, since our method does not have ground truth as a strong constraint, we design a **Lightweight Reconstruction Network (LRN)** composed of several convolutional layers as in (Chen et al. 2024) to accomplish this task, which is efficient and sufficient as demonstrated by our further experiments. While the

spike stream remains sparse under low-light conditions, we employ the voxel technique widely used in previous event-based vision tasks (Song, Huang, and Bajaj 2022; Chen and Yu 2024) to squeeze the input length of the spike sequence.

To sum up, the coarse reconstruction stage serves two main purposes. First, it enables the network to learn the fundamental mapping from the spike stream to a preliminary image representation. Second, the initial reconstruction result helps to minimize the sample distance between HQ and LQ images, aiding the learnable prompts to capture more distinctive features.

Prompt Learning

We first construct a processing pipeline to generate HQ images, which are rich in texture features and balanced in brightness. Subsequently, we train the learnable prompts to capture and distinguish the features of HQ and LQ images.

HQ Images Generation Pipeline. While we aim to train the LRN based on unpaired HQ datasets, the process of constructing such a dataset remains unsolved. Although synthetic datasets like Spike-REDS (Zhao et al. 2021) contain sufficiently sharp images, there is a disparity in the brightness distribution pattern between these synthetic images and real-world images recovered by the spike camera. To overcome it, we construct the HQ dataset from synthetic and real two parts with the pipeline illustrated in Fig. 3.

For the synthetic dataset Spike-REDS, we reconstruct HQ images based on the best-performing method WGSE (Zhang et al. 2023) on this dataset. For real-world datasets captured under normal light, such as RecVidarReal2019 (Zhu et al. 2020), MotVidarReal2020 (Zheng et al. 2022), and MomVidarReal2021 (Zheng et al. 2023), we design a mixture model that combines algorithms like TFP, TFI, Spk2ImgNet, SSML and WGSE to generate high-quality images together. Among them, image I_{hq} with the best non-reference metric NIQE (Mittal, Soundararajan, and Bovik 2012) score is

added to the HQ-Images dataset, *i.e.*:

$$I_{hq} = \arg \min_{m \in \mathcal{M}} \text{NIQE}(I_m) \quad (6)$$

where \mathcal{M} represents the set of methods and I_m denotes the image reconstructed by method m .

Prompt Optimization. Once the high-quality dataset is constructed, we use the coarse reconstruction images by the LRN under the low light as the low-quality dataset, forming a pair of positive and negative sample datasets. Subsequently, similar to CLIP-LIT (Liang et al. 2023), we employ the learnable prompt framework based on CoOp (Zhou et al. 2022) to obtain the learnable prompt representations of the positive and negative datasets.

Given the high-quality image I_{hq} with its corresponding learnable prompt T_{hq} , and the low-quality image I_{lq} reconstructed by the LRN with its learnable prompt T_{lq} . We utilize the CLIP image encoder Φ_{image} to process each image I in the HQ and LQ datasets, obtaining the encoded image features aligned with the CLIP text encoder Φ_{text} output. Next, we formulate a binary classification problem to obtain the effective representation of the learnable prompts, *i.e.*, the high-quality image should be aligned with the high-quality prompt and vice versa. We feed the learnable prompts into the text encoder and optimize them based on the cross-entropy loss function, formulated as follows:

$$\mathcal{L}_{\text{initial}} = \text{CrossEntropy}(y, \hat{y}), \quad (7)$$

$$\hat{y} = \frac{e^{\Phi_{\text{image}}(I) \cdot \Phi_{\text{text}}(T_{hq})}}{\sum_{i \in \{hq, lq\}} e^{\Phi_{\text{image}}(I) \cdot \Phi_{\text{text}}(T_i)}}, \quad (8)$$

where $I \in \{I_{hq}, I_{lq}\}$ and y is the label of the current image, 0 for the low-quality sample I_{lq} and 1 for the high-quality sample I_{hq} .

Fine Reconstruction

Previous research (Cho et al. 2023) utilized the prompt ‘Image of a [cls]’ to optimize the event reconstruction network. While the CLIP aligns text and images at a high-level semantic layer, directly applying it to the spike reconstruction low-level task introduces significant noise. To overcome this challenge, we utilize the learnable high-quality prompts along with class labels ‘X X...X [cls]’ as the supervision.

Nevertheless, employing the coupled prompts for network optimization directly poses two main problems:

- The class label is not coupled with the HQ prompt during the CLIP model optimization process, which can lead to misalignment when the coupled text is fed into the text encoder.
- The contrastive loss function designed for CLIP training cannot be applied to learnable prompts with two samples.

To the end, we decouple the prompts and design the prompt loss and the class loss to optimize the LRN reconstruction network separately, as illustrated in Fig. 4.




Figure 4: Prompt and class loss illustration.

Prompt Loss. With the learnable HQ and LQ prompts learned from the initialization stage, we can further optimize the LRN network like the previous study (Liang et al. 2023). Specifically, we design the prompt loss to measure the alignment between the reconstructed images and the corresponding HQ prompts in the image-text space. The prompt loss $\mathcal{L}_{\text{prompt}}$ is formulated as:

$$\mathcal{L}_{\text{prompt}} = - \frac{e^{\Phi_{\text{image}}(I) \cdot \Phi_{\text{text}}(T_{hq})}}{\sum_{i \in \{hq, lq\}} e^{\Phi_{\text{image}}(I) \cdot \Phi_{\text{text}}(T_i)}}, \quad (9)$$

where I represents the reconstructed image.

Class Loss. The input spike stream contains limited information under the low light condition. To this end, we utilize the class label ‘[cls]’ of the captured object to supervise the training of the LRN, which aims to guide the network in learning high-level semantic features of the captured scenes from sparse inputs. Given that CLIP utilizes the InfoNCE loss (Oord, Li, and Vinyals 2018) for aligning image-text pairs during training, we adopt a similar strategy by employing InfoNCE to optimize the LRN. Given a batch of spike stream $\mathcal{S} \in \{0, 1\}^{B \times K \times H \times W}$, the loss function is formulated as follows:

$$\mathcal{L}_{\text{class}} = - \sum_{i=1}^B \log \frac{e^{(\Phi_{\text{image}}(I_i) \cdot \Phi_{\text{text}}(T_{c_i})) / \tau}}{\sum_{j=1}^B e^{(\Phi_{\text{image}}(I_i) \cdot \Phi_{\text{text}}(T_{c_j})) / \tau}}, \quad (10)$$

where I_i represents the reconstructed image from the i -th spike input, T_{c_i} is the text feature corresponding to the predicted class c_i , and τ is a temperature parameter that controls the sharpness of the contrast.

Total Loss. The total loss $\mathcal{L}_{\text{total}}$ combines the prompt loss $\mathcal{L}_{\text{prompt}}$ and the class loss $\mathcal{L}_{\text{class}}$ as a weighted sum, formulated as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{class}} + \lambda \mathcal{L}_{\text{prompt}}, \quad (11)$$

where λ is the hyper-parameter controlling the contribution of the prompt loss set to 100 in this task.




Figure 5: Visual comparison of our method with previous methods on the U-CALTEHC dataset.

Experiments

Datasets

We conduct both quantitative and qualitative experiments on the UHSR dataset (Zhao et al. 2024a), which includes real-world spikes captured by the spike camera under low light from ultra-high-speed moving objects of the CALTECH (Fei-Fei, Fergus, and Perona 2004) and CIFAR (Krizhevsky, Hinton et al. 2009) datasets, referred to as U-CALTECH and U-CIFAR respectively. Each dataset contains 5,000 spike-label training pairs and 1,000 test pairs, with 250×250 spatial resolution. We train the LRN network on the training sets and evaluate the image-text quality and classification performance on the test sets.

Training Details

We conduct the comparative experiment based on the Spike-Zoo (Chen et al. 2025) repository, while the overall experimental framework is built on the PyTorch platform and trained based on an NVIDIA 4090 GPU. We employ the Adam optimizer with a learning rate of $4e-5$ on the U-CALTECH dataset and $1e-4$ on the U-CIFAR dataset. The training process consists of 25 total epochs: 5 epochs for coarse reconstruction, 1 epoch for prompt optimization, and 19 epochs for fine reconstruction. We employ the voxelization technique to convert the spike stream from an initial

length of 200 to 50.

Since paired spike-sharp datasets are not available in real-world scenarios, we use non-reference metrics, NIQE (Mittal, Soundararajan, and Bovik 2012), BRISQUE (Mittal, Moorthy, and Bovik 2012), and PIQE (Venkatanath et al. 2015), to evaluate the image reconstruction quality. In addition, we conduct a more comprehensive comparison by measuring the model parameter size (M), FLOPs (G), and time delay (ms) of each method. The FLOPs was calculated based on the standard spike stream with an input size of 250×400 and the time delay was estimated by averaging 100 runs for each method on an NVIDIA 4090 GPU.

Experimental Results

We compare our method with previous state-of-the-art (SOTA) spike-based image reconstruction methods, including model-based approaches TFP (Zhu et al. 2019), TFI (Zhu et al. 2019), and TFSTP (Zheng et al. 2021), supervised learning methods designed for normal lighting conditions including Spk2ImgNet (Zhao et al. 2021), SpikeFormer (She and Qing 2022), and WGSE (Zhang et al. 2023) as well as methods tailored for low-light scenarios RSIR (Zhu et al. 2023), and the self-supervised learning approach SSML (Chen et al. 2022).

We conduct qualitative experiments on the U-CALTECH

Methods	U-CALTECH			U-CIFAR			Params (M)	FLOPs (G)	Latency (ms)
	NIQE ↓	BRISQUE ↓	PIQE ↓	NIQE ↓	BRISQUE ↓	PIQE ↓			
TFP	13.276	27.268	11.201	12.889	32.751	13.464	/	/	/
TFI	11.846	25.198	<u>4.762</u>	12.166	36.661	<u>7.133</u>	/	/	/
TFSTP	11.221	<u>23.446</u>	12.460	11.872	<u>27.215</u>	13.556	/	/	/
SSML	9.887	42.916	27.704	9.474	44.576	29.659	1.192	386	184
SpikeFormer	14.750	72.366	30.880	31.681	81.694	35.944	7.581	67.8	28.7
RSIR	28.302	111.673	15.047	45.640	122.826	18.013	<u>0.255</u>	6.461	<u>7.18</u>
Spk2ImgNet	11.120	74.275	76.348	11.114	81.411	92.902	3.904	1000	70.1
WGSE	<u>7.449</u>	24.683	21.053	<u>9.458</u>	48.438	36.027	3.806	415	40.8
Ours	4.657	14.955	3.265	5.383	25.831	4.298	0.186	<u>18.562</u>	0.540

Table 1: Comparison of our method with SOTA methods on U-CALTECH and U-CIFAR datasets. We use the **bold** and underline to distinguish the best and second-best results.

and U-CIFAR datasets as detailed in Tab. 1, which demonstrates that our method outperforms SOTA on both datasets, delivering superior performance with lower parameters, computational complexity, and latency. Specifically, compared to the second-best method for each metric, our method reduces the NIQE score by 37.9% (compared to WGSE), the BRISQUE score by 36.2% (compared to TFSTP), and the PIQE score by 31.4% (compared to TFI) and by 43.1%, 5.1%, and 39.7% respectively on the U-CIFAR dataset.

We further conduct a visual comparison of our method with SOTA on the U-CALTECH dataset, as shown in Fig. 5. In particular, our method demonstrates superior noise reduction and detail preservation across recovered images. Unlike previous methods producing results with significant noise and blurry details, our approach consistently yields sharper and more accurate reconstructions with brightness balance, which complements the quantitative improvements as discussed earlier.

Ablation Study

In the ablation study, we quantitatively analyze the impact of each component and the model structure from the image quality and classification accuracy two aspects. Since we have aligned the output of the LRN with the text description through the $\mathcal{L}_{\text{class}}$ loss during training, we can leverage the CLIP model to directly classify the class of the captured scene, which assists in evaluating the overall quality of the reconstructed images.

Effectiveness of Each Component. We conducted ablation experiments to evaluate the effectiveness of the first-stage TFI coarse reconstruction, the class loss $\mathcal{L}_{\text{class}}$, and the prompt loss $\mathcal{L}_{\text{prompt}}$ as shown in Tab. 2. While the TFI image feature is not aligned with the CLIP text space, its classification accuracy is relatively low and the image quality suffers due to significant noise and unbalanced luminance. After introducing the classification loss function $\mathcal{L}_{\text{class}}$, we observe substantial improvements across all metrics, benefiting from the high-level features provided by the class features of the captured scene and the alignment with the CLIP text encoder. Additionally, we observe that the initial coarse reconstruction is beneficial for network performance,

TFI	$\mathcal{L}_{\text{class}}$	$\mathcal{L}_{\text{prompt}}$	NIQE ↓	BRISQUE ↓	ACC ↑
✓			15.389	29.957	1.60
	✓		6.043	15.821	58.90
✓	✓		4.799	18.748	63.00
✓	✓	✓	4.656	14.955	64.00

Table 2: Ablation on the contribution of each component.

Model	NIQE ↓	BRISQUE ↓	ACC ↑	Params	Flops
UNet	25.94	69.64	61.70	19.2	46.46
LRN	4.657	14.95	64.00	0.186	18.56

Table 3: Ablation on the network architecture.

likely due to the improved initialization of network parameters. Finally, the model achieves the best performance when all components are active, further demonstrating the benefits of incorporating learnable HQ prompts.

Effectiveness of Network Architecture. While we claim that our designed LRN is well-suited for spike-based image reconstruction under this setting, we further designed a spike-based reconstruction model based on UNet (Ronneberger, Fischer, and Brox 2015) and compare it with the LRN as shown in Tab. 3. Comparison demonstrates the effectiveness of our LRN model, confirming that a lightweight network architecture is not only sufficient but also efficient while the supervision signal is weak.

Conclusion

In summary, we propose a novel spike camera image reconstruction framework tailored for extreme conditions like low light. We design a CLIP-based three-stage coarse to refine reconstruction pipeline, leveraging label textual cues and unpaired high-quality datasets to guide network training. Experimental results on U-CALTECH and U-CIFAR datasets demonstrate that our method surpasses SOTA in recovering images with details and balanced brightness.

Acknowledgments

We sincerely appreciate Yuyan Chen (HUST) for her valuable suggestions and for polishing the figures. This work was supported by the National Natural Science Foundation of China (62422601, 62176003, 62088102, 62306015), the China Postdoctoral Science Foundation (2023T160015), the Young Elite Scientists Sponsorship Program by CAST (2023QNRC001), and the Beijing Nova Program (20230484362).

References

- Chen, K.; Chen, S.; Zhang, J.; Zhang, B.; Zheng, Y.; Huang, T.; and Yu, Z. 2024. SpikeReveal: Unlocking Temporal Sequences from Real Blurry Inputs with Spike Streams. *arXiv preprint arXiv:2403.09486*.
- Chen, K.; Ye, Z.; Huang, T.; and Yu, Z. 2025. Spike-Zoo: A Toolbox for Spike-to-Image Reconstruction. <https://github.com/chenkang455/Spike-Zoo>.
- Chen, K.; and Yu, L. 2024. Motion Deblur by Learning Residual from Events. *IEEE TMM*, 26: 6632–6647.
- Chen, S.; Duan, C.; Yu, Z.; Xiong, R.; and Huang, T. 2022. Self-supervised mutual learning for dynamic scene reconstruction of spiking camera. *IJCAI*.
- Chen, S.; Yu, Z.; and Huang, T. 2023. Self-supervised joint dynamic scene reconstruction and optical flow estimation for spiking camera. In *AAAI*, volume 37, 350–358.
- Cheng, J.; Liang, D.; and Tan, S. 2024. Transfer CLIP for Generalizable Image Denoising. In *CVPR*, 25974–25984.
- Cho, H.; Kim, H.; Chae, Y.; and Yoon, K.-J. 2023. Label-free event-based object recognition via joint learning with image reconstruction from events. In *ICCV*, 19866–19877.
- Fei-Fei, L.; Fergus, R.; and Perona, P. 2004. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPR*, 178–178. IEEE.
- Huang, T.; Zheng, Y.; Yu, Z.; Chen, R.; Li, Y.; Xiong, R.; Ma, L.; Zhao, J.; Dong, S.; Zhu, L.; et al. 2023. 1000× faster camera and machine vision with ordinary devices. *Engineering*, 25: 110–119.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Liang, Z.; Li, C.; Zhou, S.; Feng, R.; and Loy, C. C. 2023. Iterative prompt learning for unsupervised backlit image enhancement. In *ICCV*, 8094–8103.
- Mittal, A.; Moorthy, A. K.; and Bovik, A. C. 2012. No-reference image quality assessment in the spatial domain. *TIP*, 21(12): 4695–4708.
- Mittal, A.; Soundararajan, R.; and Bovik, A. C. 2012. Making a “completely blind” image quality analyzer. *SPL*, 20(3): 209–212.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PMLR.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18, 234–241. Springer.
- She, C.; and Qing, L. 2022. SpikeFormer: Image Reconstruction from the Sequence of Spike Camera Based on Transformer. In *Proceedings of the 2022 5th International Conference on Image and Graphics Processing*, 72–78.
- Song, C.; Huang, Q.; and Bajaj, C. 2022. E-cir: Event-enhanced continuous intensity recovery. In *CVPR*, 7803–7812.
- Venkatanath, N.; Praneeth, D.; Bh, M. C.; Channappayya, S. S.; and Medasani, S. S. 2015. Blind image quality evaluation using perception based features. In *2015 twenty first national conference on communications (NCC)*, 1–6. IEEE.
- Yang, H.; Pan, L.; Yang, Y.; Hartley, R.; and Liu, M. 2024. LDP: Language-driven Dual-Pixel Image Defocus Deblurring Network. In *CVPR*, 24078–24087.
- Zhang, J.; Jia, S.; Yu, Z.; and Huang, T. 2023. Learning temporal-ordered representation for spike streams based on discrete wavelet transforms. In *AAAI*, volume 37, 137–147.
- Zhao, J.; Xiong, R.; Liu, H.; Zhang, J.; and Huang, T. 2021. Spk2imgnet: Learning to reconstruct dynamic scene from continuous spike stream. In *CVPR*, 11996–12005.
- Zhao, J.; Zhang, S.; Yu, Z.; and Huang, T. 2024a. Recognizing Ultra-High-Speed Moving Objects with Bio-Inspired Spike Camera. In *AAAI*, volume 38, 7478–7486.
- Zhao, R.; Xiong, R.; Zhao, J.; Zhang, J.; Fan, X.; Yu, Z.; and Huang, T. 2024b. Boosting Spike Camera Image Reconstruction from a Perspective of Dealing with Spike Fluctuations. In *CVPR*, 24955–24965.
- Zheng, Y.; Yu, Z.; Wang, S.; and Huang, T. 2022. Spike-based motion estimation for object tracking through bio-inspired unsupervised learning. *TIP*, 32: 335–349.
- Zheng, Y.; Zhang, J.; Zhao, R.; Ding, J.; Chen, S.; Xiong, R.; Yu, Z.; and Huang, T. 2023. SpikeCV: Open a Continuous Computer Vision Era. *arXiv preprint arXiv:2303.11684*.
- Zheng, Y.; Zheng, L.; Yu, Z.; Shi, B.; Tian, Y.; and Huang, T. 2021. High-speed image reconstruction through short-term plasticity for spiking cameras. In *CVPR*, 6358–6367.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *IJCV*, 130(9): 2337–2348.
- Zhu, L.; Dong, S.; Huang, T.; and Tian, Y. 2019. A retina-inspired sampling method for visual texture reconstruction. In *ICME*, 1432–1437. IEEE.
- Zhu, L.; Dong, S.; Li, J.; Huang, T.; and Tian, Y. 2020. Retina-like visual image reconstruction via spiking neural model. In *CVPR*, 1438–1446.
- Zhu, L.; Zheng, Y.; Geng, M.; Wang, L.; and Huang, H. 2023. Recurrent spike-based image restoration under general illumination. In *ACMMM*, 8251–8260.