

# SpikeNeRF: Learning Neural Radiance Fields from Continuous Spike Stream

Lin Zhu<sup>1</sup> Kangmin Jia<sup>1</sup> Yifan Zhao<sup>2</sup> Yunshan Qi<sup>2</sup> Lizhi Wang<sup>1</sup> Hua Huang<sup>3,1,\*</sup>  
<sup>1</sup>Beijing Institute of Technology <sup>2</sup>Beihang University <sup>3</sup>Beijing Normal University

## Abstract

*Spike cameras, leveraging spike-based integration sampling and high temporal resolution, offer distinct advantages over standard cameras. However, existing approaches reliant on spike cameras often assume optimal illumination, a condition frequently unmet in real-world scenarios. To address this, we introduce SpikeNeRF, the first work that derives a NeRF-based volumetric scene representation from spike camera data. Our approach leverages NeRF’s multi-view consistency to establish robust self-supervision, effectively eliminating erroneous measurements and uncovering coherent structures within exceedingly noisy input amidst diverse real-world illumination scenarios. The framework comprises two core elements: a spike generation model incorporating an integrate-and-fire neuron layer and parameters accounting for non-idealities, such as threshold variation, and a spike rendering loss capable of generalizing across varying illumination conditions. We describe how to effectively optimize neural radiance fields to render photorealistic novel views from the novel continuous spike stream, demonstrating advantages over other vision sensors in certain scenes. Empirical evaluations conducted on both real and novel realistically simulated sequences affirm the efficacy of our methodology. The dataset and source code are released at <https://github.com/BIT-Vision/SpikeNeRF>.*

## 1. Introduction

In recent years, there has been significant progress in the development of neuromorphic cameras, with notable advancements such as event cameras [2, 4] and spike cameras [15, 46]. These innovative devices excel in capturing light intensity changes in high-speed scenes, presenting a breakthrough in the realm of visual perception [9, 11, 13, 31, 32, 38, 49]. Spike camera, wherein each pixel responds independently to the accumulation of photons by generating asynchronous spikes [46]. This distinctive feature enables the spike camera to record full visual details with an ultra-high temporal resolution of up to 40 kHz (or 20 kHz for the

portable version). These advantages include the capability to capture rapid changes in scenes and efficiently represent dynamic environments. With these unique features, spike cameras have demonstrated superiority in handling multiple computer vision tasks [12, 14, 33, 37, 45, 47–49, 51].

In parallel, there has been a growing trend in the computer vision community to explore neural radiance fields (NeRFs) as a solution for scene representation and novel view synthesis [5, 8, 17, 18, 20, 22, 27, 29, 39]. NeRFs employ a multilayer perceptron (MLP) combined with differentiable rendering to represent scenes, allowing for the synthesis of novel views from unseen perspectives. So far, NeRFs have predominantly undergone examination using simulated data and high-quality real-world images acquired under optimal conditions. Some studies have explored the application of NeRF to a distinct type of neuromorphic camera known as an event camera, characterized by differential sampling. Noteworthy among these endeavors are Ev-NeRF [17] and EventNeRF [29], both of which introduced neural radiance fields derived exclusively from event streams. However, the inherent lack of texture details in event data has constrained the effectiveness of these approaches, resulting in limited outcomes. Instead, a spike camera can offer fine texture details through a high temporal resolution spike stream.

However, their application to spike camera data poses a unique challenge due to the distinctive data modality. The output of a spike camera is a spike stream, fundamentally different from traditional images. Each pixel in a spike camera generates a spike, and the accumulation is reset when the photon accumulation exceeds a predefined threshold. At each timestamp, the spike camera outputs a binary matrix, referred to as a spike frame, indicating the presence of spikes at all pixels. Currently, there is no established spike-based method for addressing the challenge of generating dense, photorealistic 3D representations of scenes. Therefore, the challenge tackled in this paper is: *Can we reconstruct a dense volumetric 3D representation of the scene from a spike stream captured by a moving spike camera?*

Addressing the aforementioned challenges, this paper introduces SpikeNeRF, the first approach for deriving a volumetric scene representation from spike camera data using

\*Corresponding author: Hua Huang.

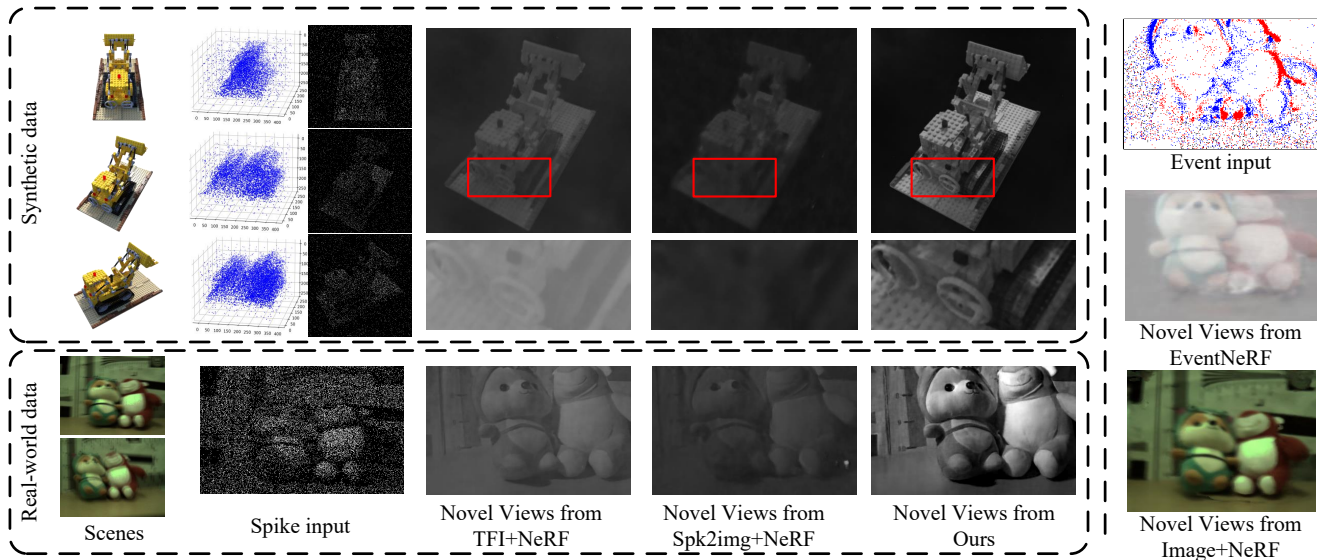


Figure 1. Comparing novel views across various vision sensors, our SpikeNeRF stands out as the first method to learn neural radiance fields from a continuous spike stream. The Spike camera, operating at 20,000 Hz, eliminates motion blur, distinguishing it from traditional cameras. Notably, when contrasted with other methods (e.g., TFI+NeRF [46], Spk2img+NeRF [45], and EventNeRF [29]) and sensors (e.g., event camera), the rendered views of objects or scenes exhibit significantly enhanced sharpness.

Neural Radiance Fields (NeRF). Our method is designed for purely spike-based supervision, preserving texture and motion details in high temporal resolution. We evaluate the efficacy of SpikeNeRF for the task of novel view synthesis on a new dataset of synthetic and real spike sequences. SpikeNeRF leverages NeRF’s inherent multi-view consistency to establish robust self-supervision, mitigating the impact of erroneous measurements in the presence of high noise levels and diverse illumination conditions. Specifically, SpikeNeRF incorporates a spike generation model based on a spiking neuron layer, which considers intrinsic parameters and non-idealities associated with spike cameras. This model enhances the fidelity of spike data, providing a more accurate representation of the scene under varying lighting conditions (see Fig. 1). To enable generalization across diverse illumination scenarios, SpikeNeRF introduces a long-term spike rendering loss. This tandem of loss functions ensures that the model can effectively capture and represent scene geometry, even in challenging lighting conditions. The main contributions of this paper are:

1) We present the first approach for inferring a NeRF volume from only a spike stream. SpikeNeRF is highly robust to various illumination conditions and builds a coherent 3D structure that can provide high-quality observations.

2) We develop a long-term spike rendering loss based on a spiking neuron layer and threshold variation simulation, which are effective in enhancing neural volumetric representation learning.

3) We build both synthetic and real-world datasets for training and testing our model. Experimental results

demonstrate that our method outperforms existing methods. We release the newly recorded spike dataset and source code to facilitate the research field.

## 2. Related Work

### 2.1. Neural Radiance Fields

In recent years, Neural Radiance Fields have garnered significant attention for their remarkable performance in tasks involving neural implicit 3D representation and novel view synthesis. Various enhancements, such as FastNeRF [8] and Depth-supervised NeRF [5], aim to expedite NeRF’s learning speed. Neural scene flow fields [18] delve into 3D Scene Representation Learning for dynamic scenes. Approaches like PixelNeRF [39] and RegNeRF [27] leverage a minimal number of input images for high-quality novel view synthesis. Mip-NeRF [1] introduces a frustum-based sampling strategy for NeRF-based anti-aliasing, addressing artifacts while enhancing training speed. Beyond speed and data augmentation, some works focus on enhancing NeRF with low-quality input images. For instance, NeRF in the wild [21] utilizes low-quality images captured by tourists, tackling occlusions and inconsistent lighting conditions. NeRF in the dark [23] and HDR-NeRF [16] enable the synthesis of high dynamic range new view images from noisy low dynamic images. Additionally, Deblur-NeRF [20] introduces the DSK module for simulating the blurring process, facilitating novel view synthesis from blurry to sharp images. However, Deblur-NeRF may face challenges in scenarios where the camera coincidentally shakes uniformly across

all views or the input images exhibit strong blur.

## 2.2. Scene Reconstruction based on Neuromorphic Camera

Generally speaking, there are two types of bio-inspired visual sampling manner: temporal contrast sampling and integral sampling. Dynamic Vision Sensor (DVS), also known as an event camera, generates events when pixel brightness changes exceed a threshold. Recently, Ev-NeRF [17] and EventNeRF [29] proposed neural radiance fields derived exclusively from event streams. However, due to the absence of texture details in event data, these approaches yield limited results. Ev-NeRF [17] is confined to learning grayscale NeRF, while EventNeRF [29] exhibits noticeable artifacts and chromatic aberration, lacking RGB data for supervision. [28] first introduces event data and blurry RGB frames to achieve high quality clear scene reconstruction. [19] proposes Robust e-NeRF to robustly reconstruct NeRFs from moving event cameras under various real-world conditions,

Our work is based on spike camera, where each pixel responds independently to the accumulation of photons by generating spikes. It records full visual details with an ultra-high temporal resolution (up to 40kHz). Image reconstruction based on spike cameras has emerged by analyzing spike intervals and counts [46]. Leveraging the spiking neuron model, a fovea-like texture reconstruction framework reconstructs images [48]. Spike camera methods have been developed for tone mapping [43], motion deblurring [12], and super resolution [37, 44]. Recent deep learning-based image reconstruction methods [3, 45, 49, 51] for spike cameras outperform traditional methods significantly. The noise characteristics of spike cameras differ substantially from those of traditional cameras. A preliminary analysis in [49] demonstrates that the noise distribution varies under different lighting conditions.

However, there is no established spike-based method for addressing the challenge of generating dense, photorealistic 3D representations of scenes. The challenge is to explore the texture details hidden in the spike stream, understanding the spike distribution is crucial for spike-based applications.

## 3. Methods

### 3.1. Neural Radiance Fields

The fundamental idea of NeRF is to employ a Multilayer Perceptron (MLP) to learn a 3D volume representation [22]. The MLP takes 3D position  $\mathbf{o}$  and 2D ray direction  $\mathbf{d}$  as inputs and produces color  $\mathbf{c}$  and density  $\sigma$  as outputs. In Eq. (1),  $F_\theta$  represents the MLP network with  $\theta$  as its parameters:

$$(\mathbf{c}, \sigma) = F_\theta(\gamma_o(\mathbf{o}), \gamma_d(\mathbf{d})). \quad (1)$$

The functions  $\gamma_o(\cdot)$  and  $\gamma_d(\cdot)$ , defined in Eq. (1), map the input 5D coordinates to a higher-dimensional space, en-

abling the neural network to better capture color and geometry information in the scene. To generate views from the scene representation, NeRF employs a classical volume rendering method described in Eq. 2. For a given ray  $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$  emitted from the camera, the expected color projected on the pixel is denoted as  $C(\mathbf{r})$ . The near and far bounds of the ray are represented by  $l_n$  and  $l_f$ , respectively. NeRF divides the interval  $[l_n, l_f]$  into  $N$  discrete bins, where  $c_i$  and  $\sigma_i$  are the outputs of  $F_\theta$ , indicating the color and density of each bin through which the ray passes. The distance between adjacent samples is  $\delta_i = l_{i+1} - l_i$ , and  $T_i$  represents the transparency of particles between  $l_n$  and bin  $i$ .

$$C(\mathbf{r}) = \sum_{i=1}^N T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (2)$$

where  $T(i) = \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j)$ . To ensure reasonable sampling in the final model, NeRF utilizes a Hierarchical volume sampling strategy, simultaneously optimizing the coarse and fine models and using the density obtained by the coarse model to determine the sampling weight of the fine model.

### 3.2. Spike Sampling Mechanism

In the field of frame-free neuromorphic vision, the spike camera operates by converting the intensity of light into voltage through photoreceptors [7, 46]. Upon reaching a predefined threshold, a one-bit spike is generated, accompanied by a signal to reset the integrator. This mechanism resembles that of an integrate-and-fire neuron. Luminance stimuli  $I$  lead to varying spike firing rates, with asynchronous triggering of output and reset across pixels. Brighter light results in faster firing, governed by the inequality  $\int I dt \geq \phi$ .

The raw data from the spike camera is a three-dimensional spike stream, focusing solely on luminance intensity integration and firing ultra-high frequency spikes. At each sampling moment, a digital signal of “1” (spike) or “0” is output based on whether a spike was fired. The spike firing status of pixel  $(i, j)$  at moment  $t$  is represented as  $S_{i,j}(t) \in \{0, 1\}$ . The spike signal model is defined as:

$$S(x, y, t) = S(x, y, nt_0) = \epsilon \left( \int_{t'}^{nt_0} \frac{C \Delta V}{I_{ph}(x, y, \tau)} d\tau - 1 \right), \quad (3)$$

where  $t'$  is the start time of the current integration phase, reset at the phase’s end.  $\Delta V$  denotes the voltage,  $nt_0$  denotes the  $n$ -th clock signal, with  $t_0$  as the unit time interval (50 $\mu$ s), and  $I_{ph}$  represents the magnitude of the current for photoelectric conversion.  $C$  is the capacitor in parallel with the photodiode, which is used to store electrons for photo-conversion.  $\epsilon(x)$  is a step function where  $\epsilon(x) = 1$  when  $x \geq 0$  otherwise  $\epsilon(x) = 0$ .

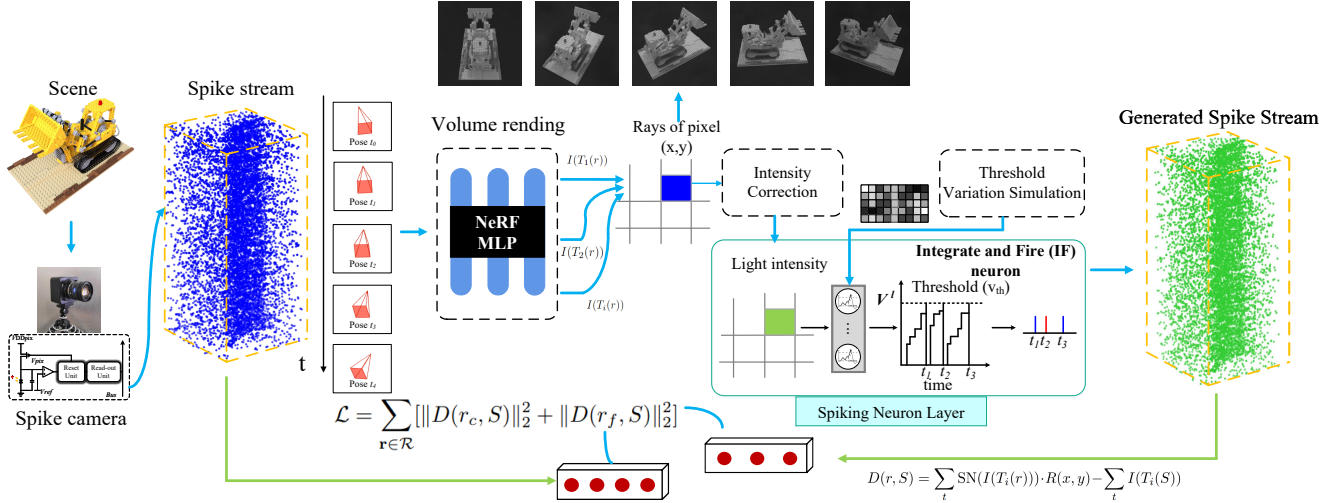


Figure 2. The architecture of SpikeNeRF. Motivated by the objective of learning NeRFs from a continuous spike stream, we establish the connection between the pixel ray  $r$  and the real-world spike stream  $S$ . To quantify the rendering loss in the spike domain, we integrate a spiking neuron layer following the NeRF MLP. The nonuniformity is captured through pixel-to-pixel threshold variation, simulated by the spiking neuron layer. This tandem of loss functions ensures that the model can effectively capture and represent scene geometry.

The noise of the spike camera differs significantly from traditional cameras due to circuit variations. The noise components encompass *Shot Noise*, wherein the photosensor converts photons into photo-electrons, resulting in a photo-current subject to shot noise  $N_p$  even in uniform light scenarios. *Dark Current Noise* arises from thermal diffusion and defects, causing the spike circuit to produce signals even in the absence of scene light, with dark current noise denoted as  $N_d$ . *Response Nonuniformity* stems from chip manufacturing limitations, causing pixel units to exhibit varying sensitivity to light intensity, leading to photo response nonuniformity noise  $N_{nr.u}$ . Lastly, *Quantization Noise* arises from the delay in the release time of the spike signal relative to the generation time of the reset signal, introducing quantization noise  $N_q$ .

### 3.3. Self-Supervision with Volumetric Rendering

An overview of the training pipeline is illustrated in Fig. 2. In line with the spike generation model, we propose self-supervision with volumetric rendering in this section. Different from traditional RGB frames, a binary spike frame can not accurately represent the light intensity at the current timestamp. Let us denote the intensity values of the pixel ray  $r$  at time  $T_i$  as  $\hat{I}(T_i(r))$ . Since the high temporal resolution of the spike camera, the estimated short-term light intensity  $\hat{L}_{T_i}$  can record motion details hidden in the spike stream. A straightforward approach is to calculate the light intensity  $\hat{L}_{T_i}$  at time  $T_i$  based on the original spike stream:

$$\hat{L}_{T_i} = \frac{\theta}{n - m}, \quad m < n, \quad (4)$$

where  $m$  and  $n$  are the timestamps corresponding to two adjacent spikes,  $\theta$  controls the gray-scale level.

However,  $\hat{L}_{T_i}$  is influenced by various sources of noise, particularly the temporal sampling noise [49]. The firing intervals of spikes fluctuate even in constant illumination conditions due to the inherent spike sampling mechanism. For instance, in the integrate-and-fire mechanism, if the membrane potential of the neuron (pixel) is not reset, an object with a light intensity corresponding to an inter-spike interval of 2 might be recorded as 1 or 3. Consequently,  $\hat{L}_{T_i}$  is suitable for capturing motion details in the intensity values of the pixel ray  $r$ , but texture details are often affected.  $\hat{L}_{T_i}$  is susceptible to non-uniform pixel response, akin to standard image sensors. This is manifested by pixel-to-pixel variations in the contrast threshold, analogous to the spike camera’s version of Fixed-Pattern Noise (FPN). Hence, relying solely on the rendered and degraded estimated light intensity values for NeRF training supervision can yield suboptimal results.

#### 3.3.1 Measuring Rendering Loss in Spike Domain

In order to capture the precise distribution of the spike stream, we advocate for employing a long-term spike rate rendering loss rather than directly assessing short-term light intensity. To account for the various noise sources, inspired by the spike simulator [51], we introduce a layer of spiking neurons that generate spike streams based on the intensity values of the pixel ray  $r$  emulating the spike sampling process. We then compute the long-term spike rate rendering loss directly on both the generated spike stream and the




Figure 3. The backpropagation process of spiking neurons. The spiking neuron layer, comprising 256 time steps, follows the NeRF MLP. The weight of the last layer of MLP can be updated through Backpropagation Through Time (BPTT) using Eq. 10.

original spike stream.

Our findings indicate that by combining NeRF’s inherent multi-view consistency with our long-term spike rendering loss, our framework can establish robust self-supervision. This approach proves effective in mitigating the impact of erroneous measurements, particularly in the presence of high noise levels and diverse illumination conditions.

**Relationship between the Pixel Ray  $r$  and Real-world Spike Stream  $S$ .** The objective is to estimate intensity values  $\hat{I}(T_i(r))$  corresponding to the real light of the scene. Assuming that  $\hat{I}(T_i(r))$  is clear, to supervise the  $\hat{I}(T_i(r))$  with the real noisy spike stream, we need to consider multiple noises. Inspired by [51], we can define the relationship between  $\hat{I}(T_i(r))$  and spike noise as:

$$\hat{I}(T_i(r)) = I(T_i(S)) - I(T_i(N)), \quad (5)$$

where  $I(T_i(S)) = \frac{1}{L + N_p + N_d + N_{rnu} + N_q} + N_c$ ,  $I(T_i(N))$  denotes the intensity changes caused by noise,  $L$  represents the scene light intensity,  $Q_r$  is relative quantity matrix of electric charge.  $N_p$ ,  $N_d$ ,  $N_{rnu}$ ,  $N_q$ , and  $N_c$  represent shot noise, dark current noise, response nonuniformity noise, quantization noise, and truncation noise, respectively.

**Spiking Neuron Layer.** To accurately represent the spike distribution, we introduce an additional layer of spiking neurons following the NeRF MLP to generate the spike stream. Given that the spike camera employs an integrate-and-fire mechanism to convert the light intensity of the scene into spike streams, our framework utilizes the Integrate-and-Fire (IF) neuron model. The Integrate-and-Fire (IF) model [10] is a widely employed neuron model in Spiking Neural Networks (SNN) known for its biological realism [36]. For better representation, we represent the function of IF neuron as the discrete form:

$$V_t = V_{t-1} + X_t, \quad S_t = H(V_t - V_{th}), \quad (6)$$

where  $V(t)$  denotes the membrane potential after neuronal dynamics at  $t$ ,  $X(t)$  represents the input to neuron. A spike fires if  $V(t)$  exceeds the threshold  $V_{th}$ ,  $V_{rest}$  is the resting potential after firing,  $S_t$  denotes the spike output at  $t$ ,

$H(\cdot)$  denotes the Heaviside step function which is defined as  $H(x) = 1$  for  $x \geq 0$  and  $H(x) = 0$  for  $x < 0$ .

**Threshold Variation Simulation.** The response nonuniformity noise can be characterized by the pixel-to-pixel variation in the firing threshold. In fact, the deviation matrix corresponding to the response nonuniformity noise can be obtained by capturing a uniform light scene and recording the intensity. Thus, the threshold variation of each spiking neuron can be modeled by Eq. 7, and the spike stream can be generated by

$$\hat{S}(x, y) = \text{SN}(I(T_i(r))) \cdot R(x, y), \quad (7)$$

where  $\text{SN}(\cdot)$  denotes the spiking neuron,  $R(x, y)$  is the nonuniformity matrix and can be obtained by capturing a uniform light scene and recording the intensity<sup>1</sup>.

**Long-term Spike Rendering Loss.** The spike stream can be mapped into an RKHS (reproducing kernel Hilbert space) as continuous-time functions which incorporate a statistical description of spike trains. Based on the RKHS, the kernel method [6] is employed for measuring spike train distances. Additionally, [47] introduces an intensity-based distance metric to quantify spike distances. Nevertheless, the complexity of these two methods is excessively high, rendering NeRF training inefficient.

In order to measure the distance between the generated spike stream  $\hat{S}$  and the input spike stream  $S$ , we propose a long-term spike rendering loss, which can be expressed by

$$D(r, S) = \sum_t \text{SN}(I(T_i(r))) \cdot R(x, y) - \sum_t I(T_i(S)). \quad (8)$$

In theory, employing a longer spike stream for measurement has the potential to mitigate excessive truncation noise. In our experiment, the parameter  $t$  is configured to 256, aligning with the time steps of the spiking neuron. The final loss function can be expressed as follows:

$$\mathcal{L} = \sum_{r \in \mathcal{R}} [\|D(r_c, S)\|_2^2 + \|D(r_f, S)\|_2^2], \quad (9)$$

where  $r_c$  and  $r_f$  are the pixel rays of coarse and fine models, respectively. In this way, the loss function of NeRF with spike stream as input is converted into Eq. 9. We adopt the design of the joint optimization of NeRF’s coarse model and fine model, which is still beneficial in our framework.

### 3.3.2 Backpropagation Training

The spike generation function of an IF neuron is a hard threshold function that emits a spike when the membrane

<sup>1</sup>Choose the pixel  $(x_m, y_m)$  which is closest to the average response value as the reference pixel.  $R(x, y)$  is then obtained by calculating the ratio of the reference pixel’s response value to the response values of other pixels:  $R(x, y) = \frac{(L_2 + L_d(x_m, y_m))T_2(x_m, y_m)}{(L_2 + L_d(x, y))T_2}$ , where  $L$  and  $T$  are variables to be calibrated. See our supplementary materials for details.



Figure 4. Quantitative results on synthetic spike data.

potential exceeds a firing threshold. Due to this discontinuous and non-differentiable neuron model, standard backpropagation algorithms cannot be applied to SNNs in their native form [40, 41]. As shown in Fig. 3, the backpropagated errors pass through the spiking neuron layer using BackPropagation Through Time (BPTT) [35]. Since the spiking neurons are one-to-one connected to the neurons of the last layer of NeRF MLP, the weight of the last layer of NeRF MLP  $w^l$  is updated by BPTT. In BPTT, the network is unrolled for all discrete time steps, and  $w^l$  update is computed as the sum of gradients from each time step as follows:

$$\Delta w^l = \sum_{t=1}^T \frac{\partial \mathcal{L}_{total}}{\partial o_t^l} \frac{\partial o_t^l}{\partial V_t^l} \frac{\partial V_t^l}{\partial w^l}, \text{ where } \frac{\partial o_t^l}{\partial V_t^l} = H_1'(V_t - V_{th}), \quad (10)$$

where  $o_t^l$  is the output of the neuron at time  $t$ ,  $\frac{\partial o_t^l}{\partial V_t^l}$  denotes the derivative of spike with respect to the membrane potential after charging at time step  $t$ . Since  $\frac{\partial o_t^l}{\partial V_t^l}$  is not differentiable, we adopt surrogate gradient method [26] to calculate it. In our work, the approximate IF gradient is computed as  $H_1(x) = \frac{1}{V_{th}}$ , where the threshold value accounts for the change of the spiking output with respect to the input.

Based on the above, a learnable realistic noise-embedded spike stream generation pipeline is proposed. Each time step of the spiking neuron represents a readout circuit process that verifies the reset signal in a clock signal corresponding to a spike camera. The dark current noise,

Table 1. Quantitative evaluation on synthetic dataset.

Method / Metric	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
TFP (32) + NeRF	16.21/16.63	0.142/0.648	0.662/0.058
TFP (256) + NeRF	15.07/19.24	0.094/0.864	0.741/0.085
TFI + NeRF	14.93/19.34	0.108/0.862	0.752/0.090
Spk2img + NeRF	13.41/14.05	0.064/0.749	0.729/0.133
Ours	<b>20.78/22.07</b>	<b>0.209/0.916</b>	<b>0.617/0.053</b>

nonuniformity noise, and quantization noise are collectively treated as nonuniformity noise. These components are incorporated into the generated spike stream through a simulation of threshold variation.

## 4. Experiment

### 4.1. Experimental Setup

**Synthetic Spike Data.** We utilize six synthetic scenes (chair, ficus, hotdog, lego, materials, and mic) to generate synthetic spike data in NeRF. Using the spike generator provided by [51], we simulate spike streams for each viewpoint under different illumination conditions. The original image size is resized to  $400 \times 400$ , resulting in generated spike streams with a resolution of  $400 \times 400 \times 256$ , where 256 represents the time length of each spike stream. In addition to generating spike streams from static viewpoint images,

Table 2. Ablation study on synthetic dataset.

Method	Loss	Static			Dynamic			Average		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
TFP(32)+NeRF	MSE	15.26/15.57	0.154/0.785	0.588/0.065	14.93/15.22	0.112/0.745	0.626/0.109	15.10/15.40	0.133/0.765	0.607/0.087
TFP(256)+NeRF	MSE	14.44/16.41	0.127/0.794	0.652/0.088	14.34/16.29	0.099/0.770	0.693/0.114	14.39/16.35	0.113/0.782	0.672/0.101
TFI+NeRF	MSE	13.77/16.63	0.120/0.793	0.670/0.093	13.76/16.11	0.096/0.769	0.694/0.110	13.76/16.37	0.108/0.781	0.682/0.101
Spk2img+NeRF	MSE	13.42/14.05	0.066/0.726	0.724/0.129	13.39/14.04	0.062/0.772	0.735/0.137	13.41/14.05	0.064/0.075	0.729/0.133
Ours	$\mathcal{L}_i^*$	13.77/16.63	0.120/0.793	0.670/0.093	13.76/16.11	0.096/0.769	0.694/0.110	13.76/16.37	0.108/0.781	0.682/0.101
	$\mathcal{L}_i$	18.55/19.76	0.237/0.878	0.527/0.051	16.63/17.79	0.167/0.797	0.593/0.097	17.59/18.77	0.202/0.837	0.560/0.074
	$\mathcal{L}_s + \mathcal{L}_i$	18.50/19.59	0.237/0.876	0.524/0.051	17.61/18.41	0.179/0.816	<b>0.590/0.089</b>	18.05/19.00	0.208/0.846	<b>0.557/0.070</b>
	$\mathcal{L}_s$	<b>18.72/19.38</b>	<b>0.251/0.880</b>	<b>0.517/0.050</b>	<b>18.16/18.80</b>	<b>0.187/0.820</b>	0.601/0.093	<b>18.44/19.09</b>	<b>0.219/0.850</b>	0.559/0.072

Table 3. Quantitative evaluation of different light intensities on synthetic dataset.

Method	Light intensity (16)			Light intensity (32)			Light intensity (64)		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
TFP(32)+NeRF	15.27/15.57	0.154/0.785	0.588/0.065	17.81/18.55	0.191/0.827	0.562/0.077	20.63/21.86	0.239/0.872	0.528/0.067
TFP(256)+NeRF	14.44/16.41	0.127/0.794	0.652/0.088	15.82/18.68	0.170/0.835	0.616/0.077	17.28/21.27	0.189/0.852	0.596/0.081
TFI+NeRF	13.77/16.63	0.120/0.793	0.670/0.093	15.07/18.53	0.156/0.825	0.634/0.083	16.16/20.56	0.187/0.852	0.602/0.077
Spk2img+NeRF	13.42/14.05	0.066/0.726	0.724/0.129	14.59/15.39	0.109/0.768	0.651/0.093	16.37/17.73	0.169/0.837	0.595/0.073
Ours	<b>18.72/19.38</b>	<b>0.251/0.880</b>	<b>0.517/0.050</b>	<b>22.05/23.66</b>	<b>0.300/0.926</b>	<b>0.477/0.050</b>	<b>23.89/24.46</b>	<b>0.411/0.929</b>	<b>0.428/0.049</b>




Figure 5. Quantitative results on real-world spike data.

we render 16 high-resolution images captured by the camera and feed them into the spike generator to replicate the dynamic recording process of the spike camera. Each scene comprises 100 sets of images and their corresponding event data. Further details about the synthetic data can be found in our supplementary material.

**Real-world Spike Data.** We employ the spike camera [15] to record real-world spike data. This camera is capable of capturing spike streams with a spatial resolution of  $250 \times 400$  and a temporal resolution of 20,000 Hz. Holding the camera by hand, we capture data from five real-world scenarios, each featuring texture details under different illumination conditions. Each dataset comprises approximately 35 images from various viewpoints along with their corresponding spike data. To facilitate a comprehensive comparison of different cameras, we also use the DAVIS346 color event camera [30] to capture event and RGB data. This camera is equipped to capture spatial-temporally aligned event data and RGB frames, with a resolution of  $346 \times 260$  and an exposure time set to 33ms for the RGB frames.

**Baselines.** Given the absence of existing spike-based NeRF methods, we compare our approach with three spike-based image reconstruction methods: TFI [46, 50], TFP [46, 50], and Spk2imgNet [45]. In a two-stage process, we reconstruct images using TFI and TFP (with window sizes of 32 and 256), which are commonly employed spike-based reconstruction methods, and Spk2imgNet, a learning-based method leveraging neural networks. While there are other learning-based methods for spike-based reconstruction [3], those methods aim to simultaneously estimate flow and image. The approach by [51] requires long spike sequences for recurrently optimizing the reconstruction process. Consequently, Spk2imgNet is well-suited for the NeRF setting, demonstrating stable performance. To ensure a fair comparison, all methods adopt a common NeRF backbone.

**Training Details.** For the simulated data, we use the poses provided by the Blender and the generated spike stream as the input of the NeRF framework. For real-world data, we first reconstruct the spike stream into images, and then use COLMAP to estimate the poses. For a fair comparison, all comparison methods use the same poses as the input of the network. At each optimization iteration, we randomly sample a batch of camera rays from the set of all pixels in the dataset. We use the Adam optimizer with a learning rate that begins at  $5 \times 10^{-4}$  and decays exponentially to  $5 \times 10^{-5}$ . The optimization for a single scene typically takes about 20 hours to converge on a single NVIDIA 3090 GPU.

## 4.2. Quantitative Experiment

We utilize three standard metrics: PSNR, SSIM [34], and LPIPS [42]. These metrics are employed to quantify the similarity between the synthesized novel views and the provided target novel views. The quantitative results on the




Figure 6. Quantitative results on different light intensities.

Table 4. Quantitative evaluation on real-world data.

Camera	Method	Brisque ↓	NIQE ↓
APS	Image+NeRF	52.07	9.06
APS	Image+NeRF*	53.05	10.06
DVS	Event+ENeRF	83.98	12.71
Spike camera	TFP+NeRF	42.64	6.16
Spike camera	TFI+NeRF	47.92	7.49
Spike camera	Ours	<b>35.94</b>	<b>6.09</b>

synthetic dataset are presented in Table 1. Since only the object will produce firing spikes in the synthetic dataset, we provide two measurement approaches: the left one utilizes the entire image for measurement, while the right one focuses solely on measurements within the object region. The results show that our model performs better than other two-stage methods, proving the effectiveness of the proposed framework. Table 4 presents the results on real-world data. As the real data lacks ground truth, we assess the results using no-reference quality assessment methods, specifically Brique [24] and NIQE [25]. The results indicate that our method outperforms other models. Additional experimental results can be found in our supplementary materials.

### 4.3. Qualitative Experiment

We perform qualitative experiments on both synthetic and real-world data. As shown in Fig. 6, the rendered views from our method show markedly improved sharpness. TFP (W=32) struggles to generate novel views for the “mic” and “materials” scenes due to their low light intensity. The short window used by TFP fails to gather sufficient information. Similarly, Spk2img also exhibits suboptimal performance. This learning-based method is designed for ideal illumination conditions, and its performance sharply degrades when confronted with noisy spikes. In Fig. 5, three real-world scenes with varying illumination conditions are depicted. “toys” exhibits high light intensity, “grid” features medium

light, and “box” has low illumination. The results illustrate that our model performs effectively across different scenes and lighting conditions in real spike data.

### 4.4. Ablation Study

We perform ablation studies on the proposed long-term spike rendering loss, and the results are presented in Table 2. In the table,  $\mathcal{L}_s$  represents the proposed spike rendering loss,  $\mathcal{L}_i$  involves using the light intensity of rays plus threshold correction, and the reconstructed images from spikes as supervision, while  $\mathcal{L}_i^*$  directly employs the estimated light intensity and reconstructed images as supervision. The experiments are conducted in both static and dynamic settings, affirming the effectiveness of our framework. Furthermore, we assess the performance under varying light intensities on the synthetic dataset. The results, presented in Table 3 and Fig. 6, demonstrate the robustness of our framework across different illumination conditions.

## 5. Conclusion

This paper introduces SpikeNeRF as a pioneering approach to derive a volumetric scene representation from spike camera data using NeRF. With a focus on purely spike-based supervision, SpikeNeRF preserves texture and motion details in high temporal resolution, addressing the challenges associated with real-world spike sequences. Our evaluations on a newly curated dataset of synthetic and real spike sequences demonstrate the efficacy of SpikeNeRF for novel view synthesis. We hope that our work will shed light on the research of high quality 3D representation learning with novel spike stream techniques.

**Acknowledgement** This work is partially supported by National Natural Science Foundation of China under Grant No.62302041, 62322204, 62131003, China National Postdoctoral Program under contract No.BX20230469, and Beijing Institute of Technology Research Fund Program for Young Scholars.



## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 2
- [2] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A  $240 \times 180$  130 db 3  $\mu$ s latency global shutter spatiotemporal vision sensor. *IEEE J. Solid-State Circuits*, 49(10):2333–2341, 2014. 1
- [3] Shiyang Chen, Chaoteng Duan, Zhaofei Yu, Ruiqin Xiong, and Tiejun Huang. Self-supervised mutual learning for dynamic scene reconstruction of spiking camera. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 2859–2866. International Joint Conferences on Artificial Intelligence Organization, 2022. Main Track. 3, 7
- [4] Tobi Delbrück, Bernabe Linares-Barranco, Eugenio Culurciello, and Christoph Posch. Activity-driven, event-based vision sensors. In *IEEE Int. Symp. Circuits Syst. (ISCAS)*, pages 2426–2429, 2010. 1
- [5] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. 1, 2
- [6] Siwei Dong, Zhichao Bi, Yonghong Tian, and Tiejun Huang. Spike coding for dynamic vision sensor in intelligent driving. *IEEE Internet of Things Journal*, 6(1):60–71, 2018. 5
- [7] S. Dong, L. Zhu, D. Xu, Y. Tian, and T. Huang. An efficient coding method for spike camera using inter-spike intervals. In *2019 Data Compression Conference (DCC)*, pages 568–568, 2019. 3
- [8] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021. 1, 2
- [9] Mathias Gehrig, Mario Millhausler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *2021 International Conference on 3D Vision (3DV)*, pages 197–206. IEEE, 2021. 1
- [10] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014. 5
- [11] Daxin Gu, Jia Li, Lin Zhu, Yu Zhang, and Jimmy S Ren. Reliable event generation with invertible conditional normalizing flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 1
- [12] Jin Han, Chu Zhou, Peiqi Duan, Yehui Tang, Chang Xu, Chao Xu, Tiejun Huang, and Boxin Shi. Neuromorphic camera guided high dynamic range imaging. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2020. 1, 3
- [13] Javier Hidalgo-Carrio, Daniel Gehrig, and Davide Scaramuzza. Learning monocular dense depth from events. In *2020 International Conference on 3D Vision (3DV)*, pages 534–542. IEEE, 2020. 1
- [14] Liwen Hu, Rui Zhao, Ziluo Ding, Lei Ma, Boxin Shi, Ruiqin Xiong, and Tiejun Huang. Optical flow estimation for spiking camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17844–17853, 2022. 1
- [15] Tiejun Huang, Yajing Zheng, Zhaofei Yu, Rui Chen, Yuan Li, Ruiqin Xiong, Lei Ma, Junwei Zhao, Siwei Dong, Lin Zhu, et al. 1000 $\times$  faster camera and machine vision with ordinary devices. *Engineering*, 25:110–119, 2023. 1, 7
- [16] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18398–18408, 2022. 2
- [17] Inwoo Hwang, Junho Kim, and Young Min Kim. Ev-nerf: Event based neural radiance field. pages 837–847, 2023. 1, 3
- [18] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021. 1, 2
- [19] Weng Fei Low and Gim Hee Lee. Robust e-nerf: Nerf from sparse & noisy events under non-uniform motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18335–18346, 2023. 3
- [20] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12861–12870, 2022. 1, 2
- [21] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 2
- [22] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 3
- [23] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16190–16199, 2022. 2
- [24] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12):4695–4708, 2012. 8

- [25] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. 8
- [26] Emre O Neftci, Hesham Mostafa, and Friedemann Zenke. 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, 2019. 6
- [27] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 1, 2
- [28] Yunshan Qi, Lin Zhu, Yu Zhang, and Jia Li. E2nerf: Event enhanced neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13254–13264, 2023. 3
- [29] Viktor Rudnev, Mohamed Elgharib, Christian Theobalt, and Vladislav Golyanik. Eventnerf: Neural radiance fields from a single colour event camera. *arXiv preprint arXiv:2206.11896*, 2022. 1, 2, 3
- [30] Gemma Taverni, Diederik Paul Moeys, Chenghan Li, Celso Cavaco, Vasyl Motsnyi, David San Segundo Bello, and Tobi Delbruck. Front and back illuminated dynamic and active pixel vision sensors comparison. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 65(5):677–681, 2018. 7
- [31] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza. Time lens: Event-based video frame interpolation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16155–16164, 2021. 1
- [32] Xiao Wang, Jianing Li, Lin Zhu, Zhipeng Zhang, Zhe Chen, Xin Li, Yaowei Wang, Yonghong Tian, and Feng Wu. Visevent: Reliable object tracking via collaboration of frame and event flows. *IEEE Transactions on Cybernetics*, 2023. 1
- [33] Yixuan Wang, Jianing Li, Lin Zhu, Xijie Xiang, Tiejun Huang, and Yonghong Tian. Learning stereo depth estimation with bio-inspired spike cameras. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2022. 1
- [34] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 7
- [35] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. 6
- [36] Jibin Wu, Chenglin Xu, Xiao Han, Daquan Zhou, Malu Zhang, Haizhou Li, and Kay Chen Tan. Progressive tandem learning for pattern recognition with deep spiking neural networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7824–7840, 2021. 5
- [37] Xijie Xiang, Lin Zhu, Jianing Li, Yixuan Wang, Tiejun Huang, and Yonghong Tian. Learning super-resolution reconstruction for high temporal resolution spike stream. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(1):16–29, 2021. 1, 3
- [38] Lan Xu, Weipeng Xu, Vladislav Golyanik, Marc Habermann, Lu Fang, and Christian Theobalt. Eventcap: Monocular 3d capture of high-speed human motions using an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4968–4978, 2020. 1
- [39] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 2
- [40] Malu Zhang, Hong Qu, Ammar Belatreche, Yi Chen, and Zhang Yi. A highly effective and robust membrane potential-driven supervised learning method for spiking neurons. *IEEE transactions on neural networks and learning systems*, 30(1):123–137, 2018. 6
- [41] 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 postsynaptic potential function for backpropagation in deep spiking neural networks. *IEEE transactions on neural networks and learning systems*, 33(5):1947–1958, 2021. 6
- [42] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 7
- [43] J. Zhao, R. Xiong, and T. Huang. High-speed motion scene reconstruction for spike camera via motion aligned filtering. In *IEEE Int. Symp. Circuits Syst. (ISCAS)*, pages 1–5, 2020. 3
- [44] Jing Zhao, Jiyu Xie, Ruiqin Xiong, Jian Zhang, Zhaofei Yu, and Tiejun Huang. Super resolve dynamic scene from continuous spike streams. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2533–2542, 2021. 3
- [45] Jing Zhao, Ruiqin Xiong, Hangfan Liu, Jian Zhang, and Tiejun Huang. Spk2imgnet: Learning to reconstruct dynamic scene from continuous spike stream. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11996–12005, 2021. 1, 2, 3, 7
- [46] Lin Zhu, Siwei Dong, Tiejun Huang, and Yonghong Tian. A retina-inspired sampling method for visual texture reconstruction. In *IEEE Int. Conf. Multimedia Expo. (ICME)*, pages 1432–1437, 2019. 1, 2, 3, 7
- [47] Lin Zhu, Siwei Dong, Tiejun Huang, and Yonghong Tian. Hybrid coding of spatiotemporal spike data for a bio-inspired camera. *IEEE Trans. on Circuits and Systems for Video Technology*, 2020. 1, 5
- [48] Lin Zhu, Siwei Dong, Jianing Li, Tiejun Huang, and Yonghong Tian. Retina-like visual image reconstruction via spiking neural model. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1438–1446, 2020. 3
- [49] Lin Zhu, Jianing Li, Xiao Wang, Tiejun Huang, and Yonghong Tian. Neuspikes-net: High speed video reconstruction

- tion via bio-inspired neuromorphic cameras. In *IEEE Int. Conf. Comput. Vis. (ICCV)*, pages 2400–2409, 2021. [1](#), [3](#), [4](#)
- [50] Lin Zhu, Siwei Dong, Jianing Li, Tiejun Huang, and Yonghong Tian. Ultra-high temporal resolution visual reconstruction from a fovea-like spike camera via spiking neuron model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1233–1249, 2022. [7](#)
- [51] Lin Zhu, Yunlong Zheng, Mengyue Geng, Lizhi Wang, and Hua Huang. Recurrent spike-based image restoration under general illumination. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8251–8260, 2023. [1](#), [3](#), [4](#), [5](#), [6](#), [7](#)