

Temporal Residual Guided Diffusion Framework for Event-Driven Video Reconstruction

Lin Zhu¹, Yunlong Zheng¹, Yijun Zhang², Xiao Wang³, Lizhi Wang¹, and Hua Huang^{4,*}

¹ Beijing Institute of Technology, Beijing, China

² China Mobile (Suzhou) Software Technology Co., Ltd., Jiangsu, China

³ Anhui University, Anhui, China

⁴ Beijing Normal University, Beijing, China

Abstract. Event-based video reconstruction has garnered increasing attention due to its advantages, such as high dynamic range and rapid motion capture capabilities. However, current methods often prioritize the extraction of temporal information from continuous event flow, leading to an overemphasis on low-frequency texture features in the scene, resulting in over-smoothing and blurry artifacts. Addressing this challenge necessitates the integration of conditional information, encompassing temporal features, low-frequency texture, and high-frequency events, to guide the Denoising Diffusion Probabilistic Model (DDPM) in producing accurate and natural outputs. To tackle this issue, we introduce a novel approach, the Temporal Residual Guided Diffusion Framework, which effectively leverages both temporal and frequency-based event priors. Our framework incorporates three key conditioning modules: a pre-trained low-frequency intensity estimation module, a temporal recurrent encoder module, and an attention-based high-frequency prior enhancement module. In order to capture temporal scene variations from the events at the current moment, we employ a temporal-domain residual image as the target for the diffusion model. Through the combination of these three conditioning paths and the temporal residual framework, our framework excels in reconstructing high-quality videos from event flow, mitigating issues such as artifacts and over-smoothing commonly observed in previous approaches. Extensive experiments conducted on multiple benchmark datasets validate the superior performance of our framework compared to prior event-based reconstruction methods.

Keywords: Event camera · diffusion model · video reconstruction

1 Introduction

Event-based video reconstruction has emerged as a prominent area of research, fueled by the unique advantages offered by event cameras, such as high dynamic range and rapid motion capture capabilities [20, 44]. However, existing methods [2, 28, 29, 33, 41, 49, 50] tend to prioritize temporal information extraction from continuous event flow, leading to an overemphasis on low-frequency texture features and resulting in issues

* Corresponding author.

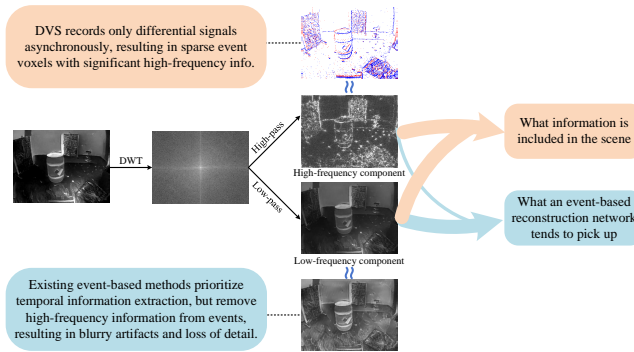


Fig. 1: Existing methods often emphasize low-frequency texture, causing over-smoothing and loss of high-frequency details in image reconstruction. This motivated us to explore a framework that strategically incorporates both temporal and high-frequency event priors.

like over-smoothing and artifacts in reconstructed scenes. It is crucial to recognize the unique nature of event cameras, which operate differently from conventional cameras by measuring high-frequency changes in intensity asynchronously at the time they occur. However, despite their potential, the temporal nature of real event cameras complicates the reconstruction problem [11, 43], necessitating innovative solutions. Fig. 1 illustrates a comparison between the frequency domain components of the scene and the reconstruction results of event-based reconstruction methods.

Recently, Diffusion Probabilistic Models (DPMs) [3, 13, 17, 26, 30, 36, 45] have shown remarkable progress in image generation, opening avenues for event-based reconstruction tasks by using conditional DPMs to generate realistic images with fine details. The integration of an initial predictor and a diffusion model ensures precise and adaptive conditioning, resulting in perceptual quality improvements of restored images. However, challenges persist in effectively integrating degraded images and other conditional information into DPMs for improved generative capacity. Incorporating various streams of conditional information into clusters of DPMs is essential for facilitating spatiotemporal adaptive conditioning throughout the reconstruction process. To address this challenge, we propose to integrate conditional information, encompassing temporal, low-frequency texture, and high-frequency events, to guide the Denoising Diffusion Probabilistic Model (DDPM) in generating accurate and natural outputs.

In this paper, we introduce a novel approach named the temporal residual guided diffusion framework. This framework strategically leverages both temporal and frequency-based event priors and incorporates three key conditioning modules: a pre-trained low-frequency intensity estimation module, a temporal recurrent encoder module, and an attention-based high-frequency prior enhancement module. To capture temporal scene variations effectively, we introduce the use of a temporal-domain residual image as the target for the diffusion model. By combining these three conditioning paths with the temporal residual framework, our proposed framework excels in reconstructing high-quality videos from event flow, addressing issues such as artifacts and over-smoothing commonly observed in previous approaches. In summary, this paper introduces innova-

tive frameworks in the realms of event-based video reconstruction and diffusion-based image restoration. Our proposed temporal residual guided diffusion video framework and unified conditional framework for image restoration significantly advance the state-of-the-art, addressing key challenges and showcasing superior performance through extensive experiments. The contributions of our paper can be summarized as follows:

1) We introduce a novel temporal residual guided diffusion framework. This framework combines temporal and frequency-based event priors to effectively capture temporal scene variations in video reconstruction.

2) The framework strategically incorporates three conditioning modules, namely a low-frequency intensity estimation module, a temporal recurrent encoder module, and an attention-based high-frequency prior enhancement module. Through the amalgamation of these conditioning paths with the temporal residual framework, our proposed approach excels in reconstructing high-quality videos from event flow.

3) Extensive experiments demonstrate that our framework is effective in overcoming issues like artifacts and over-smoothing, establishing itself as a noteworthy advancement in both event-based video reconstruction and diffusion-based image restoration.

2 Related Works

Event-based Video Reconstruction. In virtue of its high dynamic range, high temporal resolution, and low power consumption, event sensors excel in many visual application scenarios. Video reconstruction is a fundamental and popular topic in the event-based vision literature. Early event-based video reconstruction approaches relied on the representational similarity between events and gradients [5, 18, 25] or optical flow [1] to reconstruct scene structure directly. Nevertheless, these early methods fell short in achieving sufficiently realistic final reconstructed intensity images, primarily due to a lack of exploration into the prior information embedded in long-term data.

Since the widespread adoption of deep learning methods, many data-driven approaches have demonstrated significant potential. E2VID [28, 29] utilizes LSTM to accumulate temporal features of events and learns on large datasets with optical flow constraints, resulting in significant improvements in reconstruction quality. FireNet [33] employs a lighter network architecture to achieve faster reconstruction speeds. Sparse-E2VID [9] only emphasizes high-frequency by learning gradients with lightweight networks and then utilizes these gradients to generate images following traditional reconstruction approaches. Sparse-E2VID [9] emphasizes high-frequency by learning gradients with lightweight networks and then utilizes these gradients to generate images following traditional reconstruction approaches. SPADE-E2VID [2] uses previously reconstructed images to conditionally modulate the activations on a layer, significantly improving the effectiveness of sparse event reconstruction. ETNet [41] presents a hybrid CNN-Transformer [7] structure to reconstruct video, achieving the best reconstruction results to date. However, these methods do not effectively balance the relationship between long-term and short-term event features, resulting in reconstructed images often containing motion artifacts and blur.

Conditional Diffusion Model in Image/video Restoration. The process from intensity images to events can be modeled by a degradation model dominated by a differential

operator. Therefore, event-based video reconstruction task can be viewed as a class of video restoration tasks, and the prior relevant work has greatly benefited our research. In the wake of DDPM [13] demonstrating the impressive capability of diffusion models to generate images from random noise, applying diffusion model to image restoration tasks has become a popular focus in the literature. SR3 [32] and Platte [31] use degraded images as conditions, directly concatenating them with noisy images as the overall input to the network, demonstrating the compatibility and efficiency of diffusion model for image restoration. ShadowDiffusion [12], IDM [10], and DeS3 [14] utilize preprocessing features from existing models as conditions for the diffusion model. Resdiff [34] and UCDIR [47] employ additional preprocessing models to roughly restore degraded images before utilizing conditional diffusion models to generate residuals, significantly reducing the processing complexity of the task. IR-SDE [22] and Refusion [23] alter the diffusion process of DDPM itself, allowing sampling to commence from noisy degraded images rather than pure random noise, which significantly reduces both training and sampling costs. Inspired by multimodal tasks, Refusion [23] and DiffIR [42] encode images into a latent space to perform diffusion processes, accelerating the process of single-shot training and sampling. Event-Diffusion [19] attempts to apply the diffusion model to an event-based image reconstruction task with the input of reconstructed images and original events.

The various methods mentioned above involve directly retraining the diffusion model to handle image restoration tasks, which can be computationally expensive. However, there is a class of methods that leverage a pre-trained diffusion model and can handle image restoration tasks at an extremely low cost. SNIPS [16], DDRM [15], and DDNM [39] employ well-known degradation models to perform decomposition, making full use of the denoising prior of the diffusion model. Inspired by the use of classifier gradients in DMBG [6] for image synthesis, DPS [4] and GDP [8] calculate the posterior distribution based on Bayesian theory to guide the mean of each step in the sampling process.

In this paper, we introduce a novel temporal residual guided diffusion video framework, which integrates conditional information, encompassing temporal low-frequency texture and high-frequency events, to guide DDPM in event-based video reconstruction.

3 Methodology

3.1 Problem Statement

Event-based reconstruction tasks aim to reconstruct images from a series of events \mathcal{E} :

$$\mathcal{E} = \{e^i\}_{i=0}^{N-1} = \{(x^i, y^i, t^i, p^i)\}_{i=0}^{N-1}, \quad (1)$$

where x^i, y^i represent the pixel positions of the event e^i , and t^i represents the timestamp. Additionally, p^i represents the polarity of e^i , indicating the relationship between the scene brightness and time variation, which can be calculated as follows:

$$p^i = \begin{cases} +1, & L(x, y, t^i) - L(x, y, t^i - \delta t) \geq \varphi_+ \\ -1, & L(x, y, t^i) - L(x, y, t^i - \delta t) \leq \varphi_- \end{cases}, \quad (2)$$




Fig. 2: Comparison of different strategies. (a) Directly predict intensity images from the accumulation features of past events, such as E2VID [28, 29], ETNet [41]. (b) Jointly reconstruction from event feature accumulations and prediction from the previous frame, e.g., SPADE-E2VID [2]. (c) Our temporal residual guided diffusion framework. While most methods adopt the initial two strategies, the inherent temporal feature extracting in these approaches results in the forfeiture of high-frequency information from the events. Our approach effectively tackles this contradiction by generating high-frequency temporal residuals through a conditional diffusion model.

where L represents the scene intensity in the logarithmic domain, δt is the time interval from the previous event occurrence, and φ_+ , φ_- respectively denote the intensity thresholds for positive and negative polarities.

In a manner similar to the application of voxel-based organization in [51], we also transform events into two-dimensional images along the temporal channel with B bins:

$$V^t(b) = \sum_{i=1}^N p_i \max(0, 1 - |b - \frac{t^i - t^0}{t^{N-1} - t^0}(B - 1)|). \quad (3)$$

Considering sparse events contain more high-frequency information, we segment events between two frames based on event density, ultimately ensuring that the event density for each processing step does not exceed 0.25. Then, we can further transform the problem we need to address: Given a set of voxel grids $\{V^t\}_{t=0}^{T-1}$, estimate the scene video corresponding to each moment $\{\tilde{I}^t\}_{t=0}^{T-1}$.

3.2 Frequency-based Event Priors Analysis

Fig. 2 illustrates the current mainstream strategies for event-based reconstruction. The reconstructed image becomes more realistic with a more balanced representation of low-frequency and high-frequency information. Most methods [28, 33, 41] adopt the approach corresponding to (a) in Fig. 2, directly accumulating event information in the time domain to estimate scene intensity. This is extremely challenging, and they can achieve different approximations of realistic quality based on different architectures and algorithm complexities. SPADE-E2VID [2] adopts the approach corresponding to (b) in Fig. 2, building upon the predictions from the previous frame to simplify the task and achieve more realistic images with less computational cost. Our method, as illustrated by (c) in Fig. 2, is based on the initial intensity estimation from the previous frame. It incorporates certain low-frequency details while generating high-frequency temporal residual images through high-frequency random noise and events at the current time. Therefore, our algorithm is capable of achieving more realistic reconstruction results.

We perform frequency domain analysis on the reconstructed results of different methods, as shown in Fig. 3. Due to the limitations of the approach, both ETNet [41] and

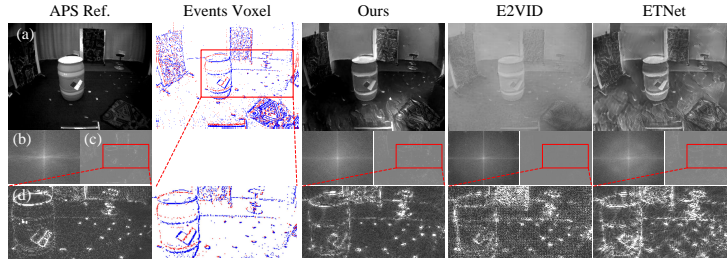


Fig. 3: Frequency domain analysis of reconstructed results. (a) Original image; (b) Fourier spectrum chart of intensity image; (c) High-frequency components of intensity image; (d) Local magnification diagram (scaled for representation). Despite the events being very similar to the high-frequency map of the scene, E2VID and ETNet cannot reconstruct precise high-frequency details.

E2VID [28] produce results with a noticeable blur at the edges, along with a reduction in high-frequency components and increased susceptibility to noise interference. Instead, the high-frequency components of our results closely resemble the events voxel, demonstrating that our method effectively utilizes the high-frequency features of events at the current time.

3.3 Temporal Residual Diffusion Framework

Fig. 4 shows the training phase of the proposed temporal residual diffusion framework. Event data represents the change in intensity over time and constitutes a differential signal. In contrast, intensity images depict the normalized brightness of the scene and constitute an integral signal. Due to their different modalities, there is a gap in data distribution, so using event data as a condition for directly generating intensity images is theoretically suboptimal, which has been confirmed in our experiments as well. To simplify the generation task and harness the temporal variations in the scene represented by the event voxel grid at the current moment, we use temporal-domain residual image x_0^t as the target images for the diffusion model:

$$x_0^t = I^t - \tilde{I}^{t-1}, \quad \tilde{I}^{t-1} = \mathcal{E}2\mathcal{V}(V^{t-1}), \quad (4)$$

where I^t represents the true intensity image of the scene at time t , \tilde{I}^{t-1} represents the estimated intensity image of the scene at time $t - 1$, and $\mathcal{E}2\mathcal{V}(\cdot)$ denotes the initial intensity predictor. From another perspective, the event voxel could be theoretically approximated as a degradation emanating from the temporal-domain residual image. Consequently, our framework embarks on a task analogous to image restoration, thereby harnessing the full generative capability of DDPM.

In terms of generating target residual image x_0^t under certain conditions using a diffusion model, the forward Markovian diffusion process q adds Gaussian noise to the data at intervals, similar to DDPM [13]:

$$q(x_\tau^t | x_{\tau-1}^t) = \mathcal{N}(x_\tau^t; \sqrt{1 - \beta_\tau} x_{\tau-1}^t, \beta_\tau \mathbf{I}), \quad (5)$$

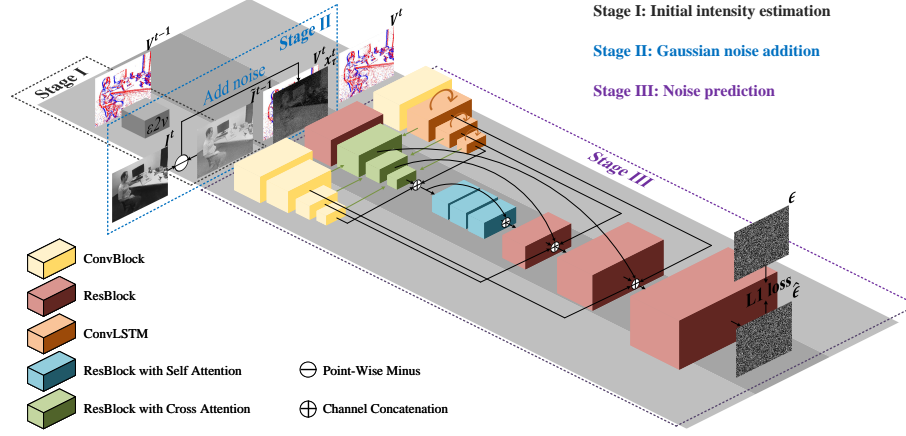


Fig. 4: Overview of temporal residual diffusion framework. At Stage I, a pre-trained intensity predictor generates initial low-frequency estimation; At Stage II, the residual image is computed in the time domain and noise is added; At Stage III, a triple-path conditional model is used to predict noise. Please refer to Fig. 5 for specific details on the ResBlock with Cross Attention.

where $\beta_\tau \in (0, 1)$ for all $\tau = 1, \dots, \mathcal{T}$. β_τ are pre-chosen hyperparameters which determine the variance of the noise added at each iteration. \mathcal{T} is the number of steps in the iteration. \mathbf{I} represents the identity matrix. Incorporating intermediate steps, we can obtain the distribution of x_τ^t given x_0^t , simplifying the noise addition:

$$q(x_\tau^t | x_0^t) = \mathcal{N}(x_\tau^t; \sqrt{\bar{\alpha}_\tau} x_0^t, (1 - \bar{\alpha}_\tau) \mathbf{I}), \quad (6)$$

where $\alpha_\tau := 1 - \beta_\tau$ and $\bar{\alpha}_\tau := \prod_{i=1}^\tau \alpha_i$. In order to extract high-frequency information from short-term events, we use the V^t , \tilde{I}^{t-1} , and the intermediate states of the ConvLSTM s^{t-1} as conditioning inputs to train the diffusion model. The pseudo-code for training is shown in Alg. 1.

The diffusion model generates content through a step-by-step denoising process executed in a reverse Markov chain. According to Bayesian theory, conditional probabilities can be derived as follows:

$$p_\theta(x_{\tau-1}^t | x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}) = \mathcal{N}(x_{\tau-1}^t; \mu_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau), \sigma_\tau^2 \mathbf{I}), \quad (7)$$

$$\mu_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau) = \frac{1}{\sqrt{\alpha_\tau}} \left(x_\tau^t - \frac{\beta_\tau}{\sqrt{1 - \bar{\alpha}_\tau}} \epsilon_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau) \right), \quad (8)$$

where $\sigma_\tau^2 = \frac{1 - \bar{\alpha}_{\tau-1}}{1 - \bar{\alpha}_\tau} \beta_\tau$, θ denotes model parameters which are optimized by maximizing the variational lower bound (VLB) during the training phase. In accordance with this conditional probability distribution, we can progressively generate the predicted intensity image at the sampling stage, and subsequently obtain the video as demonstrated in Alg. 2 and Fig. 5.

Algorithm 1: Training on a scene

Input: A set of events voxel grids $\{V^t\}_{t=0}^{T-1}$; A set of intensity images corresponding to events $\{I^t\}_{t=0}^{T-1}$; A initial intensity predictor $\mathcal{E}2\mathcal{V}$

```

1  $s^0 = None$ 
2 for  $t = 1, \dots, T - 1$  do
3    $\tilde{I}^{t-1} = \mathcal{E}2\mathcal{V}(V^{t-1})$ ;
4    $x_0^t = I^t - \tilde{I}^{t-1}$ ;
5    $\tau \sim Uniform(\{1, \dots, T\})$ ;
6    $\epsilon \sim \mathcal{N}(0, I)$ ;
7    $x_\tau^t = \sqrt{\alpha_\tau}x_0^t + \sqrt{1 - \alpha_\tau}\epsilon$ ;
8   Take gradient descent step on  $\nabla_\theta \|\epsilon - \epsilon_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau)\|_1$ ;
9    $s^t = s_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau)$ ;
10 end

```

Algorithm 2: Sampling on a scene

Input: A set of events voxel grids $\{V^t\}_{t=0}^{T-1}$; A initial intensity predictor $\mathcal{E}2\mathcal{V}$; A trained denoising model with weight θ

Result: A predicted video $\{\tilde{I}^0, \hat{I}^t\}_{t=1}^{T-1}$

```

1  $s^0 = None$ 
2 for  $t = 1, \dots, T - 1$  do
3    $\tilde{I}^{t-1} = \mathcal{E}2\mathcal{V}(V^{t-1})$ ;
4    $x_\tau^t \sim \mathcal{N}(0, I)$ ;
5   for  $\tau = T - 1, \dots, 1$  do
6     if  $\tau > 1$  then
7        $z \sim \mathcal{N}(0, I)$ ;
8     else
9        $z = 0$ ;
10    end
11     $x_{\tau-1}^t = \mu_\theta + \sqrt{\sigma_\theta}z$ 
12  end
13   $\hat{I}^t = \tilde{I}^{t-1} + x_0^t$ ;
14   $s^t = s_\theta(x_\tau^t, V^t, \tilde{I}^{t-1}, s^{t-1}, \tau)$ ;
15 end

```

3.4 Triple-path Conditional Model Architecture

Low-Frequency Intensity Estimation. As mentioned in Sec. 3.3, we attempt to address gap of data distribution by modifying the diffusion objective to include the difference image in the time domain. However, temporal residual guided diffusion without any other conditions relies too much on the intensity image at the previous time step, which is not quite accurate. We feed the generated results from the previous moment of $\mathcal{E}2\mathcal{V}$ as the initial intensity estimation into the network, and eliminate the error of the pre-trained $\mathcal{E}2\mathcal{V}$ through the accumulation of events in time domain. Specifically, we use a simple multi-scale convolutional layer to extract features of initial intensity estimation, as shown in the following equation:

$$\mathcal{F}(\tilde{I}^{t-1})_{l+1} = \text{ConvBlock}_l(\mathcal{F}(\tilde{I}^{t-1})_l), \quad (9)$$

where $\mathcal{F}(\cdot)_l$ represents the features at layer l . Each ConvBlock consists of two convolutional layers and a downsampling layer (except for the last one).

Recurrent Encoder for Temporal Motion Information. To accurately estimate the low-frequency information of the scene and predict the scene brightness more realistically, it is necessary to fully accumulate event data in the time domain. We use multi-scale ConvLSTM [35] to extract long-term and short-term features of events voxels:

$$\mathcal{F}(V^t)_{l+1} = \text{ConvLSTM}_l(\mathcal{F}(V^t)_l, s^{t-1}), \quad (10)$$

where s^{t-1} represents the hidden features of the events voxel from the previous time step.

Attention-based High-Frequency Prior Enhancement. We concatenate events voxel V^t with the noisy intensity differential image x_τ^t , both serving as representations of scene brightness variations physically, and input them into the same encoder to generate high-frequency details. The encoder feature extraction process for this path can be expressed by the following formula:

$$\mathcal{F}(x_\tau^t, V^t)_{l+1} = \text{CAtt}_l(\mathcal{F}(x_\tau^t, V^t)_l, \mathcal{F}(V^t)_l, \mathcal{F}(\tilde{I}^{t-1})_l). \quad (11)$$

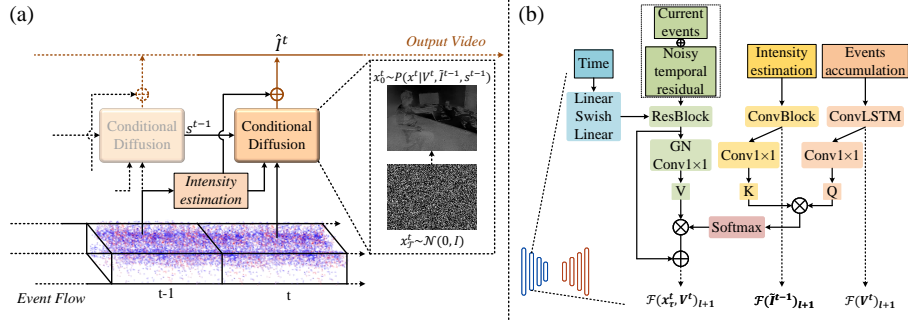


Fig. 5: (a). Overview of sampling a video. The conditional diffusion model utilizes current events, intensity estimation, and features from accumulated events in the previous moment as guidance. This process generates high-frequency temporal residuals, contributing to the intensity image for each frame when added to the initial intensity estimation. (b). Overview of ResBlock with Cross Attention. Focus on events accumulation and intensity estimation features on the noisy temporal residuals, where GN denotes group normalization.

Except for the highest resolution scale, where features are extracted using ResBlock, all other scales utilize ResBlock with cross attention as illustrated in Fig. 5. In order to fully exploit the low-frequency features of temporally accumulated information and initial intensity estimation, we employ a cross-attention mechanism between three encoders. Specifically, performing linear mapping on three different features separately using 1×1 convolutions without bias, resulting in Q, K, V :

$$Q = \text{Conv}_{1 \times 1}(\mathcal{F}(V^t)_l), K = \text{Conv}_{1 \times 1}(\mathcal{F}(\tilde{I}^{t-1})_l), V = \text{Conv}_{1 \times 1}(\mathcal{F}(x_\tau^t, V^t)_l). \quad (12)$$

Then, we calculate attention [7, 38] across encoders as shown in the following equation and add it to the original feature $\mathcal{F}(x_\tau^t, V^t)_l$:

$$\text{att} = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (13)$$

where d_k denotes the number of channels of Q .

We utilize residual blocks with self-attention to aggregate low-level features from three encoders. Skip connections and upsampling are then employed for decoding the features to obtain the final reconstruction. For specific details, please refer to the supplementary materials.

4 Experiments

4.1 Experimental Setup

During training, we use the same dataset as [29], which generates events on the COCO [21] dataset using the event simulator ESIM [27]. An alternative simulated dataset [37] could showcase scenes with a wider dynamic range and a threshold distribution more

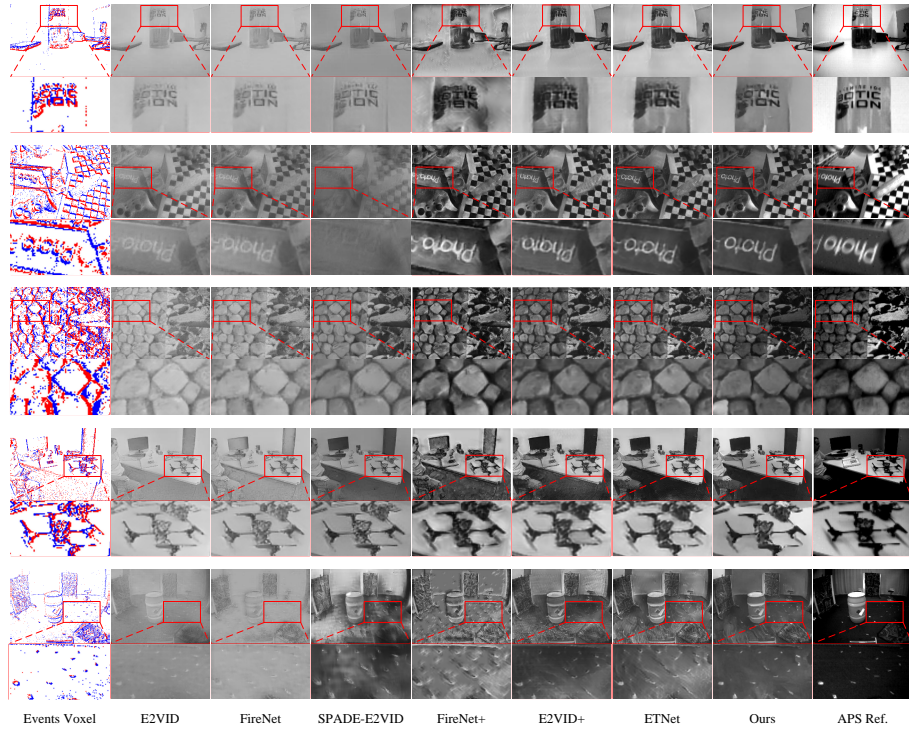


Fig. 6: Qualitative comparison on HQF (Row 1&2), IJRR (Row 3&4), and MVSEC (Row 5). Our results possess brightness distributions closest to the APS reference frame and the sharpest texture details.

akin to existing datasets like IJRR and MVSEC. This similarity might bolster the performance of reconstruction models such as E2VID+ [37] and FireNet+ [37]. Nevertheless, it may not be suitable for training our temporal residual framework. This constraint is linked to the inclusion of masked images in the dataset, aimed at simulating foreground and background motion. This methodology may introduce occlusion events, reminiscent of the occlusion challenges encountered in optical flow tasks.

In our experiments, to confirm the generalization performance of our method on real-world events, we use multiple real datasets to test, including IJRR [24], HQF [37] and MVSEC [48]. The iteration number \mathcal{T} takes the value of 2,000 during the training and testing phases. We use ETNet [41] as the initial intensity predictor. Then, we conduct an ablation study to find out the effectiveness of different strategies. Please refer to the supplementary materials to find more details about the experiments.

4.2 Comparison with the State-of-the-Art Methods

We compare the proposed method against various event-based reconstruction methods, including E2VID [29], E2VID+ [37], FireNet [33], FireNet+ [37], SPADE-E2VID [2]

Table 1: Quantitative evaluation on multiple real datasets. Best in bold.

Methods	IJRR [24]			HQF [37]			MVSEC [48]		
	MSE↓	SSIM↑	LPIPS↓	MSE↓	SSIM↑	LPIPS↓	MSE↓	SSIM↑	LPIPS↓
FireNet [33]	0.1226	0.3448	0.2340	0.0735	0.4035	0.2495	0.1613	0.2355	0.3189
E2VID [28]	0.1291	0.3490	0.2140	0.0644	0.4157	0.2288	0.1643	0.2496	0.2900
SPADE-E2VID [2]	0.0773	0.3696	0.2142	0.0648	0.3564	0.2683	0.0994	0.2759	0.2936
FireNet+ [37]	0.0954	0.3011	0.2547	0.0512	0.3382	0.2900	0.1271	0.2088	0.3843
E2VID+ [37]	0.0646	0.3688	0.1929	0.0363	0.4327	0.2160	0.0817	0.2762	0.2823
ETNet [41]	0.0694	0.3673	0.1863	0.0405	0.4004	0.2269	0.0889	0.2581	0.3016
Ours	0.0597	0.3856	0.1694	0.0441	0.4496	0.2039	0.0715	0.2912	0.2789

and ETNet [41]. The reconstructed images generated by each method are compared with the APS reference frames, and the mean square error (MSE), structural similarity (SSIM [40]) and perceptual loss (LPIPS [46]) indicators are used to quantitatively evaluate. We perform tests on all sequences of each dataset, Table 1 shows the comparison of average metrics results for each dataset, and Fig. 6 shows a visual comparison on some of the sequences. Of note, it is essential to recognize that the APS reference frames within the existing real dataset are captured by an integral camera operating within the same optical path. Due to limitations in dynamic range and differences in sampling mechanisms, there exists a considerable gap of the brightness distribution of the scene as described between APS frame and the intensity image based on event reconstruction. To ensure experimental fairness, all reconstructed images and APS reference frames undergo histogram equalization before evaluation like SPADE-E2VID [2].

Our approach has achieved state-of-the-art performance across almost all metrics. In particular, attributed to a thorough exploration of the high-frequency priors inherent in events, our results with superior structural details, exhibiting the highest SSIM metrics across all datasets, achieving improvements of 4.3%, 3.9% and 5.4% over the second-best on IJRR, HQF and MVSEC, respectively. In terms of visual comparison, our approach emphasizes detailed reconstruction, as evidenced by the clear textual outlines in the first and second rows of Fig. 6, distinctly separated from the background. The other methods either suffer from low contrast or lack completeness in reconstructing details, often containing artifacts and blur.

4.3 Ablation Study

To validate the rationality of the various strategies proposed, we conduct multiple sets of ablation experiments. The iterations for different models are kept identical and tested on the MVSEC dataset for experimental consistency. Given the sparse event flow, it presents greater challenges for reconstruction.

Effect of Temporal Residual Diffusion. In line with Sec. 3.3, due to the disparate modalities between event flow and intensity images, employing event data as a condition to guide the diffusion model for directly generating reconstructed images is theoretically unfeasible. In keeping all other conditions constant, we alter the optimization target of the diffusion model to intensity image, and the quantitative results of the testing are shown in the first row of Table 2. Please refer to supplementary materials for further visual comparison. Due to variations in the data distribution of the events and

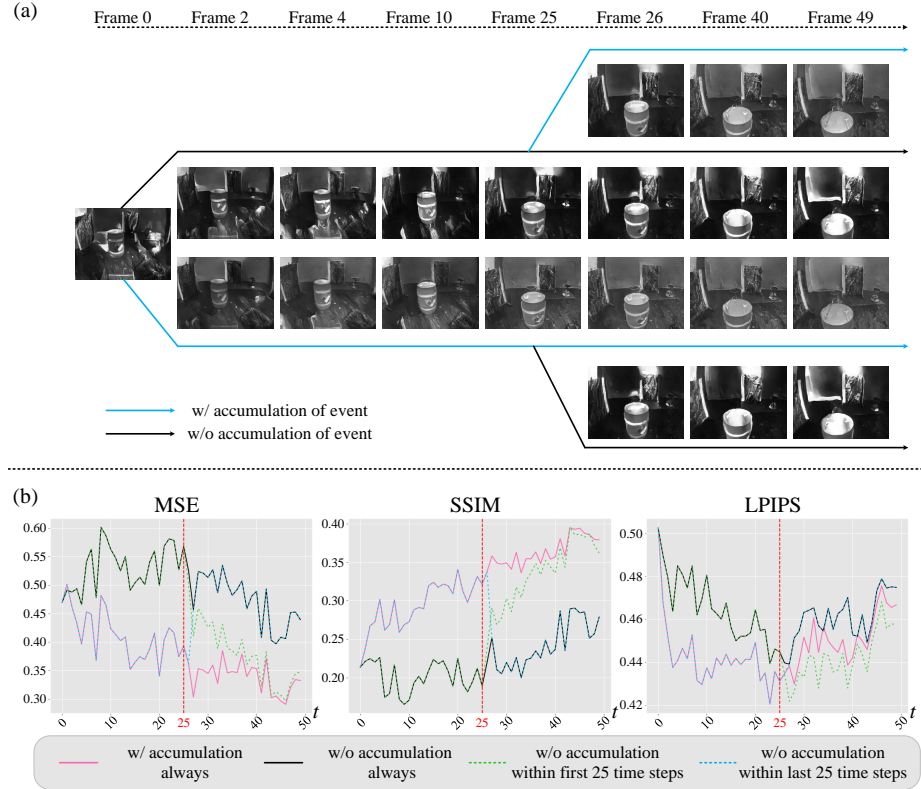


Fig. 7: (a) Visual comparison between w/o and w/ event temporal accumulation. (b) Quantitative results of (a). We test different event accumulation ways across multiple metrics over time (Best viewed with zoom in). ‘w/o accumulation within first 25 time steps’ represents the process in the first row of (a). ‘w/o accumulation within last 25 time steps’ represents the process in the fourth row of (a).

light intensity, directly generating intensity image results in a lack of sufficient texture details, especially in areas with dimly lit scenes.

Effect of Cross Encoder Attention Mechanism. In Sec. 3.4, we mentioned the use of a cross attention mechanism to focus on features across different encoders. We propose this strategy aiming to utilize temporal event features for correcting the errors from intensity estimation. The method that does not employ cross-attention mechanisms during the feature extraction stage has been retrained, and results are presented in the third row of Table 2. The visual comparison indicates that the absence of cross-attention mechanisms results in lower contrast. The brightness described differs significantly from the real scene, albeit with good reconstruction of edge details.

Effect of Recurrent Encoder for Event Accumulation. In line with Sec. 3.4, we employ ConvLSTM to aggregate temporal information for more accurate intensity esti-

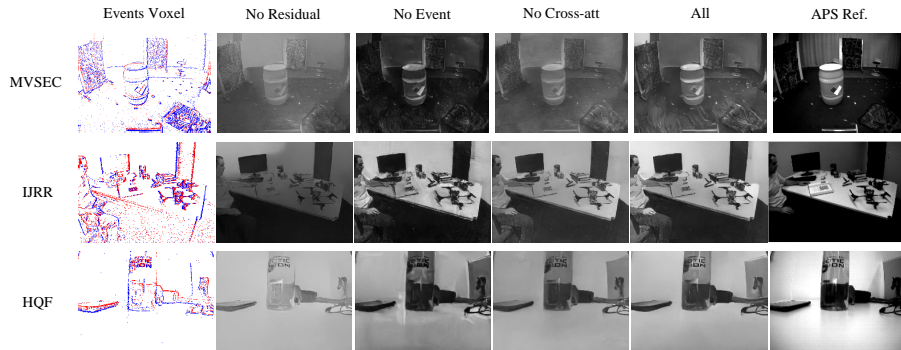


Fig. 8: The visual comparison of ablation experiments. ‘No residual’ in the second col signifies that the diffusion model generates target intensity image directly rather than temporal residual. ‘No event’ in the third col denotes the usage of initial intensity estimation as a only condition to guide the diffusion model. ‘No cross-att’ in the fourth col indicates the absence of cross-encoder attention mechanisms during the encoding phase. ‘All’ in the fifth col represents the results of our final model.




Fig. 9: Comparison with different inference steps on multiple datasets.

mation. Fig. 7 presents a comparison of whether temporal event features are utilized at different time steps. Furthermore, the lack of accumulated events indicates that the hidden state and cell state of the ConvLSTM input remains empty. The two middle rows depict the results generated respectively without using temporal cumulative event features at all and with complete usage always. In the second row, the overall brightness between each frame remains nearly constant, yet there are significant local brightness differences. This indicates the unreliability of relying solely on input intensity estimation, as there exists a considerable gap between it and the actual intensity. The third row demonstrates the accumulation of temporal events voxel. Here, the global brightness gradually changes, and areas initially overexposed or underexposed tend towards moderate brightness, validating the theory behind our design of the temporal recurrent encoder for correcting intensity errors. The first row, beginning the accumulation of events at a midpoint, gradually enhances contrast in the reconstructed image, and the corresponding metrics tend to approximate the results of continuous event accumula-

Table 2: Quantitative comparison of different strategies. Best in bold.

Strategies				MVSEC		
Residual	Recurrent	Cross-att	Event	MSE↓	SSIM↑	LPIPS↓
	✓	✓	✓	0.0902	0.2565	0.2977
✓		✓	✓	0.1021	0.2472	0.3023
✓	✓		✓	0.0812	0.2653	0.2958
✓		✓		0.0960	0.2309	0.3189
✓	✓	✓	✓	0.0707	0.2768	0.2892

tion. In the meantime, in the fourth row, abruptly ceasing to use the preceding temporal features at the midpoint causes a sudden deterioration in the reconstructed results, with the corresponding metrics plummeting to the level of non-cumulative event results across the entire duration. The quantitative results shown in subplot (b) of Fig. 7 also confirm the above analysis.

Effect of Event-based Conditional Diffusion. Our proposed method aims to exploit the high-frequency information inherent in sparse events to enhance texture details, as demonstrated in our comparative experiments where our results yield sharper edge regions. We retrain a diffusion model guided solely by initial intensity estimates, and its results are shown in the fourth row of Table 2. Visual assessments indicate a notably poor reconstruction quality, struggling not only to approximate the real scene’s brightness distribution but also exhibiting blurriness, artifacts, and a lack of sufficient texture details at local regions.

Effect of Different Inference Steps. Fig. 9 presents the metric results with varying inference steps across different datasets. It can be observed that with an increase in noise step size, the metrics gradually improve, especially the SSIM and LPIPS indicators show a notable enhancement. Due to the MSE metric’s greater emphasis on brightness differences between the image and its reference, while the noise step focuses on generating details, convergence is notably rapid for the MSE metric. Considering all measurement results, convergence is approached when the iterations exceed 1,000.

5 Conclusion

The advantage of event-based video reconstruction lies in high dynamic range and rapid motion capture. However, prevailing methods overly prioritize temporal information, resulting in over-smoothing and blurry artifacts. Our solution, the temporal residual guided diffusion Framework, adeptly integrates temporal features, low-frequency texture, and high-frequency event features. Three key conditioning modules enhance the Denoising Diffusion Probabilistic Model, ensuring accurate reconstructions. By leveraging the temporal-domain residual features, our model captures both temporal and high frequency event information. Our framework excels in mitigating over-smoothing and artifacts, evident in extensive benchmark experiments, surpassing prior methods. This marks a significant stride toward high-quality event-based video reconstruction, addressing persistent challenges in the field.

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. Bardow, P., Davison, A.J., Leutenegger, S.: Simultaneous optical flow and intensity estimation from an event camera. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 884–892 (2016)
2. Cadena, P.R.G., Qian, Y., Wang, C., Yang, M.: Spade-e2vid: Spatially-adaptive denormalization for event-based video reconstruction. *IEEE Transactions on Image Processing* **30**, 2488–2500 (2021)
3. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938* (2021)
4. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687* (2022)
5. Cook, M., Gugelmann, L., Jug, F., Krautz, C., Steger, A.: Interacting maps for fast visual interpretation. In: The 2011 International Joint Conference on Neural Networks. pp. 770–776. IEEE (2011)
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
8. Fei, B., Lyu, Z., Pan, L., Zhang, J., Yang, W., Luo, T., Zhang, B., Dai, B.: Generative diffusion prior for unified image restoration and enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9935–9946 (2023)
9. Gantier Cadena, P.R., Qian, Y., Wang, C., Yang, M.: Sparse-e2vid: A sparse convolutional model for event-based video reconstruction trained with real event noise. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 4150–4158 (2023). <https://doi.org/10.1109/CVPRW59228.2023.00437>
10. Gao, S., Liu, X., Zeng, B., Xu, S., Li, Y., Luo, X., Liu, J., Zhen, X., Zhang, B.: Implicit diffusion models for continuous super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10021–10030 (2023)
11. Gu, D., Li, J., Zhu, L., Zhang, Y., Ren, J.S.: Reliable event generation with invertible conditional normalizing flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023)
12. Guo, L., Wang, C., Yang, W., Huang, S., Wang, Y., Pfister, H., Wen, B.: Shadowdiffusion: When degradation prior meets diffusion model for shadow removal. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14049–14058 (2023)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
14. Jin, Y., Yang, W., Ye, W., Yuan, Y., Tan, R.T.: Shadowdiffusion: Diffusion-based shadow removal using classifier-driven attention and structure preservation. *arXiv preprint arXiv:2211.08089* (2022)

15. Kawar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. *Advances in Neural Information Processing Systems* **35**, 23593–23606 (2022)
16. Kawar, B., Vaksman, G., Elad, M.: Snips: Solving noisy inverse problems stochastically. *Advances in Neural Information Processing Systems* **34**, 21757–21769 (2021)
17. Kim, G., Kwon, T., Ye, J.C.: Diffusionclip: Text-guided diffusion models for robust image manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2426–2435 (2022)
18. Kim, H., Handa, A., Benosman, R., Ieng, S.H., Davison, A.J.: Simultaneous mosaicing and tracking with an event camera. *J. Solid State Circ* **43**, 566–576 (2008)
19. Liang, Q., Zheng, X., Huang, K., Zhang, Y., Chen, J., Tian, Y.: Event-diffusion: Event-based image reconstruction and restoration with diffusion models. In: *Proceedings of the 31st ACM International Conference on Multimedia*. pp. 3837–3846 (2023)
20. Lichtsteiner, P., Posch, C., Delbruck, T.: A 128×128 120 db 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits* **43**(2), 566–576 (2008). <https://doi.org/10.1109/JSSC.2007.914337>
21. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. pp. 740–755. Springer (2014)
22. Luo, Z., Gustafsson, F.K., Zhao, Z., Sjölund, J., Schön, T.B.: Image restoration with mean-reverting stochastic differential equations. *arXiv preprint arXiv:2301.11699* (2023)
23. Luo, Z., Gustafsson, F.K., Zhao, Z., Sjölund, J., Schön, T.B.: Refusion: Enabling large-size realistic image restoration with latent-space diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1680–1691 (2023)
24. Mueggler, E., Rebecq, H., Gallego, G., Delbruck, T., Scaramuzza, D.: The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and slam. *The International Journal of Robotics Research* **36**(2), 142–149 (2017)
25. Munda, G., Reinbacher, C., Pock, T.: Real-time intensity-image reconstruction for event cameras using manifold regularisation. *International Journal of Computer Vision* **126**, 1381–1393 (2018)
26. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: *International Conference on Machine Learning*. pp. 8821–8831. PMLR (2021)
27. Rebecq, H., Gehrig, D., Scaramuzza, D.: Esim: an open event camera simulator. In: *Conference on robot learning*. pp. 969–982. PMLR (2018)
28. Rebecq, H., Ranftl, R., Koltun, V., Scaramuzza, D.: Events-to-video: Bringing modern computer vision to event cameras. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3857–3866 (2019)
29. Rebecq, H., Ranftl, R., Koltun, V., Scaramuzza, D.: High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence* **43**(6), 1964–1980 (2019)
30. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
31. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–10 (2022)
32. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(4), 4713–4726 (2022)

33. Scheerlinck, C., Rebecq, H., Gehrig, D., Barnes, N., Mahony, R., Scaramuzza, D.: Fast image reconstruction with an event camera. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 156–163 (2020)
34. Shang, S., Shan, Z., Liu, G., Zhang, J.: Resdiff: Combining cnn and diffusion model for image super-resolution. arXiv preprint arXiv:2303.08714 (2023)
35. Shi, X., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems* **28** (2015)
36. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 32. Curran Associates, Inc. (2019), https://proceedings.neurips.cc/paper_files/paper/2019/file/3001ef257407d5a371a96dcd947c7d93-Paper.pdf
37. Stoffregen, T., Scheerlinck, C., Scaramuzza, D., Drummond, T., Barnes, N., Kleeman, L., Mahony, R.: Reducing the sim-to-real gap for event cameras. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16*. pp. 534–549. Springer (2020)
38. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
39. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. arXiv preprint arXiv:2212.00490 (2022)
40. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
41. Weng, W., Zhang, Y., Xiong, Z.: Event-based video reconstruction using transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2563–2572 (2021)
42. Xia, B., Zhang, Y., Wang, S., Wang, Y., Wu, X., Tian, Y., Yang, W., Van Gool, L.: Diffir: Efficient diffusion model for image restoration. arXiv preprint arXiv:2303.09472 (2023)
43. Xiang, X., Zhu, L., Li, J., Tian, Y., Huang, T.: Temporal up-sampling for asynchronous events. In: *2022 IEEE International Conference on Multimedia and Expo (ICME)*. pp. 01–06. IEEE (2022)
44. Yang, M., Liu, S.C., Delbruck, T.: A dynamic vision sensor with 1% temporal contrast sensitivity and in-pixel asynchronous delta modulator for event encoding. *IEEE Journal of Solid-State Circuits* **50**(9), 2149–2160 (2015). <https://doi.org/10.1109/JSSC.2015.2425886>
45. Zeng, Z., Yang, F., Liu, H., Satoh, S.: Improving deep metric learning via self-distillation and online batch diffusion process. *Visual Intelligence* **2**(1), 1–13 (2024)
46. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
47. Zhang, Y., Shi, X., Li, D., Wang, X., Wang, J., Li, H.: A unified conditional framework for diffusion-based image restoration. arXiv preprint arXiv:2305.20049 (2023)
48. Zhu, A.Z., Thakur, D., Özaslan, T., Pfrommer, B., Kumar, V., Daniilidis, K.: The multivehicle stereo event camera dataset: An event camera dataset for 3d perception. *IEEE Robotics and Automation Letters* **3**(3), 2032–2039 (2018). <https://doi.org/10.1109/LRA.2018.2800793>
49. Zhu, L., Li, J., Wang, X., Huang, T., Tian, Y.: Neuspikes-net: High speed video reconstruction via bio-inspired neuromorphic cameras. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2400–2409 (2021)

50. Zhu, L., Wang, X., Chang, Y., Li, J., Huang, T., Tian, Y.: Event-based video reconstruction via potential-assisted spiking neural network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3594–3604 (2022)
51. Zihao Zhu, A., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based optical flow using motion compensation. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. pp. 0–0 (2018)