

Efficient Event Stream Super-Resolution with Recursive Multi-Branch Fusion

Quanmin Liang^{1,2*}, Zhilin Huang^{2,3*}, Xiawu Zheng², Feidiao Yang²,
Jun Peng², Kai Huang^{1†} and Yonghong Tian^{2,4†}

¹School of Computer Science and Engineering, Sun Yat-Sen University

²Peng Cheng Laboratory

³Shenzhen International Graduate School, Tsinghua University

⁴Peking University

liangqm5@mail2.sysu.edu.cn, {zerinhwang03, yhtian}@pku.edu.cn, zhengxiawu@xmu.edu.cn,
{yangf, pengj01}@pcl.ac.cn, huangk36@mail.sysu.edu.cn

Abstract

Current Event Stream Super-Resolution (ESR) methods overlook the redundant and complementary information present in positive and negative events within the event stream, employing a direct mixing approach for super-resolution, which may lead to detail loss and inefficiency. To address these issues, we propose an efficient Recursive Multi-Branch Information Fusion Network (RMFNET) that separates positive and negative events for complementary information extraction, followed by mutual supplementation and refinement. Particularly, we introduce Feature Fusion Modules (FFM) and Feature Exchange Modules (FEM). FFM is designed for the fusion of contextual information within neighboring event streams, leveraging the coupling relationship between positive and negative events to alleviate the misleading of noises in the respective branches. FEM efficiently promotes the fusion and exchange of information between positive and negative branches, enabling superior local information enhancement and global information complementation. Experimental results demonstrate that our approach achieves over **17%** and **31%** improvement on synthetic and real datasets, accompanied by a **2.3** \times acceleration. Furthermore, we evaluate our method on two downstream event-driven applications, *i.e.*, object recognition and video reconstruction, achieving remarkable results that outperform existing methods. Our code and Supplementary Material are available at <https://github.com/Lqm26/RMFNet>.

1 Introduction

Event cameras are biologically inspired asynchronous sensors [Brandli *et al.*, 2014]. Unlike traditional cameras, event cameras register only the changes in brightness for each pixel over time. These are known as “events”, which are categorized as positive or negative, depending on whether the brightness increases or decreases, respectively. This characteristic significantly reduces the amount of recorded informa-

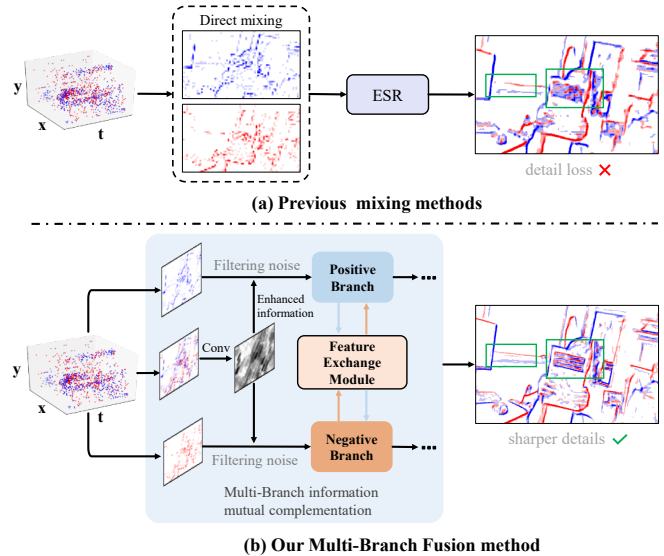


Figure 1: Compared to previous ESR methods that directly mix positive and negative events, our multi-branch approach effectively extracts and integrates features from positive and negative events, achieving a more complete and clearer details (see the **green box**).

tion, resulting in advantages such as high temporal resolution, low power consumption, and a high dynamic range (HDR) [Gallego *et al.*, 2020]. However, as the application scenarios become more complex, the spatial resolution of existing event cameras is insufficient [Li *et al.*, 2021]. Increasing spatial resolution at the hardware level presents challenges in implementing asynchronous circuits [Gallego *et al.*, 2020], making it difficult to maintain the low power consumption and high temporal resolution advantages of event cameras [Weng *et al.*, 2022; Gehrig and Scaramuzza, 2022]. Therefore, some researchers propose to address this issue at the software level, *e.g.* by leveraging advanced algorithms, which is referred to as Event Stream Super-Resolution (ESR).

Current research on ESR can be mainly divided into two directions. One approach aims to directly generate high-resolution event data from low-resolution event streams by spiking neural networks [Li *et al.*, 2019a; Li *et al.*, 2021] or

frame-assisted methods [Wang *et al.*, 2020b]. However, these methods often require significant memory [Li *et al.*, 2019a; Li *et al.*, 2021] and high-quality images as assistance [Wang *et al.*, 2020b], which complicates the training process and hinders achieving large-factor super-resolution. Hence, researchers have proposed stacking event streams into either event frames [Rebecq *et al.*, 2017] or event count images [Maqueda *et al.*, 2018; Zhu *et al.*, 2018] and subsequently applying learning-based methods for ESR [Duan *et al.*, 2021; Weng *et al.*, 2022]. Within event streams, there exist spatiotemporally inconsistent positive and negative events [Gehrig *et al.*, 2019]. These events do not perfectly align on a 2D plane but contain complementary information. Merging them into an event frame results in partial cancellation between positive and negative events, forming a new representation. As positive and negative events typically do not occur independently, the event frame helps filter out some naturally occurring noise in the event stream. Consequently, positive events, negative events, and event frames each contain different information about the event stream. However, previous methods did not effectively distinguish and fully utilize this information. They simply mix them and input them into the ESR model, leading to the loss of fine details in the super-resolved (SR) event stream (Figure 1(a)).

To address these issues, we propose an efficient Recursive Multi-Branch Information Fusion Network (RMFNET). As illustrated in Figure 1(b), this network processes positive events, negative events, and event frames in a multi-branch fashion. Positive and negative events contain the majority of information in the event stream, while the event frame provides the guidance to filter noises. Therefore, we design a Feature Fusion Module (FFM) to highlight the valuable information in positive and negative streams according to the event frame at the initial stage. Specifically, this module calculates attention weight maps from features of different branches, facilitating the fusion of contextual information and aiding the positive and negative branches in noise removal. Subsequently, the positive and negative branches conduct feature extraction for positive and negative events, respectively, employing a Feature Exchange Module (FEM) for the adaptive fusion and exchange. Through capturing complementary information and long-range dependencies between positive and negative events, FEM improves the integration and exchange of information across different branches.

The main contributions of our work are as follows:

- We introduce an efficient Recursive Multi-Branch Information Fusion Network capable of effectively merging positive events, negative events, and event frames, thereby obtaining high-quality SR event images.
- We design Feature Fusion Modules and Feature Exchange Modules, which enhance the positive and negative event streams while effectively fusing and complementing information across different branches.
- We explored the impact of existing data augmentation methods on ESR tasks and proposed an effective data augmentation strategy to enhance the model’s robustness and performance.
- Our method achieves over 17% and 31% improvement

on synthetic and real datasets, accompanied by a $2.3\times$ acceleration. In downstream event-based recognition and reconstruction tasks, our method effectively enhances performance, further validating the effectiveness of our approach.

2 Related Work

Event Stream Super-resolution. Due to the unique spatiotemporal characteristics of event streams, event stream super-resolution (ESR) tasks are often more challenging. Initially, Li *et al.* [Li *et al.*, 2019a] introduced the Event Count Map (ECM) as a method to describe event spatial distribution. They established a spatiotemporal filter to generate a time-rate function and employed a non-homogeneous Poisson distribution to model events on each pixel. However, this approach encounters inaccuracies in estimating spatiotemporal distributions when performing high-factor super-resolution. To address this issue, Wang *et al.* [Wang *et al.*, 2020b] proposed a novel optimization framework called GEF, which utilizes motion correlation probabilities to filter event noise. The optimization maximizes the structural correspondence between low-resolution events and high-resolution image signals, facilitating event stream super-resolution in conjunction with image frames. Despite performing well in certain scenarios, the GEF method exhibits performance degradation when image frame quality deteriorates. Building upon this, Li *et al.* [Li *et al.*, 2021] proposed a spatio-temporal constraint learning method based on the spiking neural network (SNN) characteristics to simultaneously learn temporal and spatial features in event streams. On the other hand, Duan *et al.* [Duan *et al.*, 2021] transformed event streams into a 2D event frames format and designed a 3D U-Net-based network for ESR. While both methods demonstrated excellent performance in small-scale super-resolution tasks, they faced challenges of excessive memory requirements and training difficulties in large-factor super-resolution. To effectively address the challenges of large-factor super-resolution, Weng *et al.* [Weng *et al.*, 2022] introduced an event-based super-resolution method based on Recurrent Neural Networks. They initially transformed event streams into coarse-grained high-resolution event streams using coordinate relocation, followed by super-resolution through recurrent networks. This approach not only handles high-factor super-resolution effectively but also mitigates training challenges posed by excessive memory requirements.

However, these methods did not account for the spatiotemporal inconsistencies and complementarities between positive and negative events in the event stream. Directly mixing them may lead to the loss of details. Therefore, we propose RMFNET, employing a multi-branch approach to mutually fuse and complement positive and negative events, which effectively enhances the performance of ESR.

3 Method

In this section, we first introduced the data representation methods for event cameras in Section 3.1. Subsequently, we presented the proposed Recursive Multi-Branch Information




Figure 2: Architecture of our proposed Recursive Multi-Branch Information Fusion Network (RMFNET). Initially, the event frame is fused into positive and negative branches along with the previous output O_{t-1} and state h_{t-1} using the Feature Fusion Module (**bottom left**). Subsequently, each branch independently extracts features through Residual Blocks, and a Feature Exchange Module (**bottom middle**) facilitates the exchange of information between the branches. Finally, the features from the positive and negative branches are concatenated, and high-resolution event count images O_t are obtained through Pixel Shuffle operation.

Fusion Network in Section 3.2. Finally, we described the data augmentation methods for ESR in Section 3.3.

3.1 Event Data Representation

A set of event streams can be represented as $\mathcal{E} = \{e_k\}_{k=1}^N$, where N is the number of events, each event $e_k \in \mathcal{E}$ can be denoted by a tuple (x_i, y_i, t_i, p_i) , representing spatial coordinates, timestamp and polarity respectively. Subsequently, we partition $\{e_k\}_{k=1}^N$ into positive events $\{e_k\}_{k=1}^{N_p}$ and negative events $\{e_k\}_{k=1}^{N_n}$ based on their polarity $p_i = \pm 1$. Specifically, we stack $\{e_k\}_{k=1}^{N_p}$ and $\{e_k\}_{k=1}^{N_n}$ into event count images [Maqueda *et al.*, 2018; Zhu *et al.*, 2018] according to the following equations:

$$h(x, y) = \sum_{e_k \in \mathcal{E}} \delta(x - x_i, y - y_i) \quad (1)$$

where δ represents the Kronecker delta. Thus, we can build up two event count images from $\{e_k\}_{k=1}^{N_p}$ and $\{e_k\}_{k=1}^{N_n}$: positive $\mathbf{p}_t \in \mathbb{R}^{H \times W}$, and negative $\mathbf{n}_t \in \mathbb{R}^{H \times W}$. And the event frame [Rebecq *et al.*, 2017] is obtained by stacking all events (including positive and negative events) using equation (1), resulting in $\mathbf{f}_t \in \mathbb{R}^{H \times W}$.

3.2 Multi-Branch Fusion Networks

The framework of our proposed RMFNET is depicted in Figure 2. The main inputs of this network include positive events, negative events, and event frames. Additionally, following a recursive approach [Schuster and Paliwal, 1997], we introduce the previous output O_{t-1} and state h_{t-1} into the input, aiding in better capturing features from adjacent event

streams and achieving contextual fusion of event stream information. Given that positive and negative events contain the majority of information in the event stream, we process them separately through dedicated positive and negative branches. The event frame, serving as a coupled representation of positive and negative events, assists in filtering out noise (as positive and negative events do not occur independently). Thus, in the initial stage of RMFNET, we fuse \mathbf{f}_t with O_{t-1} and h_{t-1} as event-enhanced information $F_{Enh} \in \mathbb{R}^{C \times H \times W}$, which is passed to the positive and negative branches through the Feature Fusion Module (FFM). The positive and negative branches utilize Residual Block [He *et al.*, 2016] as the backbone for feature extraction from the fused positive and negative events, respectively. Subsequently, a Feature Exchange Module (FEM) is employed to facilitate the fusion and exchange of information between the positive and negative branches. Finally, the features from positive and negative events are concatenated, outputting the hidden state h_t , and the SR event count images O_t are obtained through pixel shuffle [Shi *et al.*, 2016].

Feature Fusion Module

As depicted in the bottom left corner of Figure 2, the FFM is tasked with transmitting event-enhanced information F_{Enh} to the positive and negative branches without compromising their distinctive features. Denoting the features extracted by convolutional layers for positive and negative events as $F_t^p \in \mathbb{R}^{C \times H \times W}$ and $F_t^n \in \mathbb{R}^{C \times H \times W}$, respectively. Initially, we concatenate F_t^p with F_{Enh} , followed by a preliminary fusion through a Basic Block, resulting in F_t^{fuse} . The event-enhanced information encompasses details from adja-

cent event streams and coupling information between positive and negative events, effectively guiding the positive branch in detail recovery. To integrate these features seamlessly, we utilize the fused feature F_t^{fuse} to compute two attention weights. The first is local attention weight:

$$\mathbf{A}_t^{loc} = BN \left(C_{1 \times 1} \left(\mathbf{R} \left(BN \left(C_{1 \times 1} \left(F_t^{fuse} \right) \right) \right) \right) \right) \quad (2)$$

where BN denotes batch normalization, $C_{1 \times 1}$ represents a 1×1 convolution operation, and \mathbf{R} represents the ReLU activation function.

The second is global attention weight $\mathbf{A}_t^{glo} \in \mathbb{R}^{C \times 1 \times 1}$, which is computed channel-wise. Specifically, we incorporate global average pooling to process F_t^{fuse} along the spatial dimensions:

$$\mathbf{A}_t^{glo} = \mathbf{f}^{att} \left(GAP \left(F_t^{fuse} \right) \right) \quad (3)$$

where \mathbf{f}^{att} represents the function given in equation (2), and GAP stands for global average pooling.

Finally, we combine the global and local attention, apply it to F_t^{fuse} , and add it to the previous features of positive events F_t^p , thereby integrating the event-enhanced information into the positive branch:

$$F_t^{out} = F_t^p + F_t^{fuse} \otimes \left(\sigma \left(\mathbf{A}_t^{glo} \oplus \mathbf{A}_t^{loc} \right) \right) \quad (4)$$

where \otimes represents element-wise product, σ denotes the sigmoid activation function, and \oplus signifies broadcasting addition. Our network is entirely symmetric with respect to positive and negative branches, so the negative branch follows the same process.

Feature Exchange Module

Considering the complementary and redundant information present in positive and negative events, directly integrating feature information from these branches may have adverse effects. To address this, we introduce a Feature Exchange Module (depicted below Figure 2), which utilizes attention mechanisms to automatically select and enhance crucial features, facilitating efficient information exchange between the two branches.

Firstly, to reduce the redundancy in individual branch features and emphasize important features, we apply spatial attention separately to both branches:

$$\tilde{\mathbf{F}}_t = Conv_{basic} \left(F_t^{in} \right) \quad (5)$$

$$\tilde{\mathbf{F}}_t^P = Conv \left(\tilde{\mathbf{F}}_t \right) \otimes \tilde{\mathbf{F}}_t + Conv \left(\tilde{\mathbf{F}}_t \right) \quad (6)$$

where F_t^{in} represents the input features from the positive and negative branches, $Conv_{basic}$ denotes the Basic Block, and $Conv$ represents the convolutional operation. The $Conv(\tilde{\mathbf{F}}_t)$ in Eq.(6) respectively serves as the weight and bias, adjusting the weights of branch features. $\tilde{\mathbf{F}}_t^P$ is the output of the positive branch, and $\tilde{\mathbf{F}}_t^N$ is obtained similarly from the negative branch.

Subsequently, inspired by self-attention mechanisms [Vaswani *et al.*, 2017; Wang *et al.*, 2018], we design two symmetrical Attention Blocks to capture complementary information from the positive and negative branches. Taking

the positive branch as an example, we use $C_{1 \times 1}$ to obtain $\mathbf{V} \in \mathbb{R}^{C \times (HW)}$ for the positive branch, and $\mathbf{Q} \in \mathbb{R}^{C_1 \times (HW)}$ and $\mathbf{K} \in \mathbb{R}^{C_1 \times (HW)}$ for the negative branch. Here, C represents the number of channels, and C_1 is set to $1/8$ of C for enhanced computational efficiency. Therefore, the output of the positive branch, fused with features from the negative branch, can be represented as:

$$\mathbf{F}_{fuse}^P = \mathbf{V} \otimes \left(\sigma \left(\mathbf{Q}^T \otimes \mathbf{K} \right) \right) \quad (7)$$

where \mathbf{Q}^T represents the transpose of \mathbf{Q} . Through the two symmetrical attention modules, we can achieve a complementary fusion of features from the positive and negative branches, effectively enhancing the performance of ESR.

Training Objectives

We partition the event stream into multiple sequences of length T for training our method, following the approach of Weng *et al.* [Weng *et al.*, 2022]. We set $T = 9$ and use Mean Squared Error (MSE) to calculate the loss:

$$\mathcal{L} = \sum_{t=1}^T MSE \left(O_t^{SR}, ECI_t^{HR} \right) \quad (8)$$

where O_t^{SR} represents the event count images of the final SR event stream, ECI_t^{HR} represents the ground truth event count images, and MSE is the mean square error function.

3.3 Data Augmentation for ESR

Previous research in the field of image or video super-resolution has shown that methods involving operations or augmentations in the pixel space [Zhang *et al.*, 2017; Yun *et al.*, 2019] can effectively enhance task performance [Yoo *et al.*, 2020], as they preserve the spatial relationships within the images. In the realm of ESR, there is currently a lack of systematic investigation into Event Stream Super-Resolution Data Augmentation (ESRDA). To address this gap, we adapt and refine data augmentation methods from some event stream studies [Gu *et al.*, 2021; Barchid *et al.*, 2023] and RGB image domains, exploring the impact of data augmentation on the ESR task. We experiment with the following methods:

- **Polarity flipping.**
- **RandomFlip** [Simonyan and Zisserman, 2015].
- **Drop by time** [Gu *et al.*, 2021].
- **Random drop** [Gu *et al.*, 2021].
- **Drop by area** [Gu *et al.*, 2021].
- **Random drop or add noise.**
- **Static Translation.**
- **RandomResizedCrop** [He *et al.*, 2016].

Regarding the details and parameters for data augmentation operations, please refer to the **Supplementary Material**.

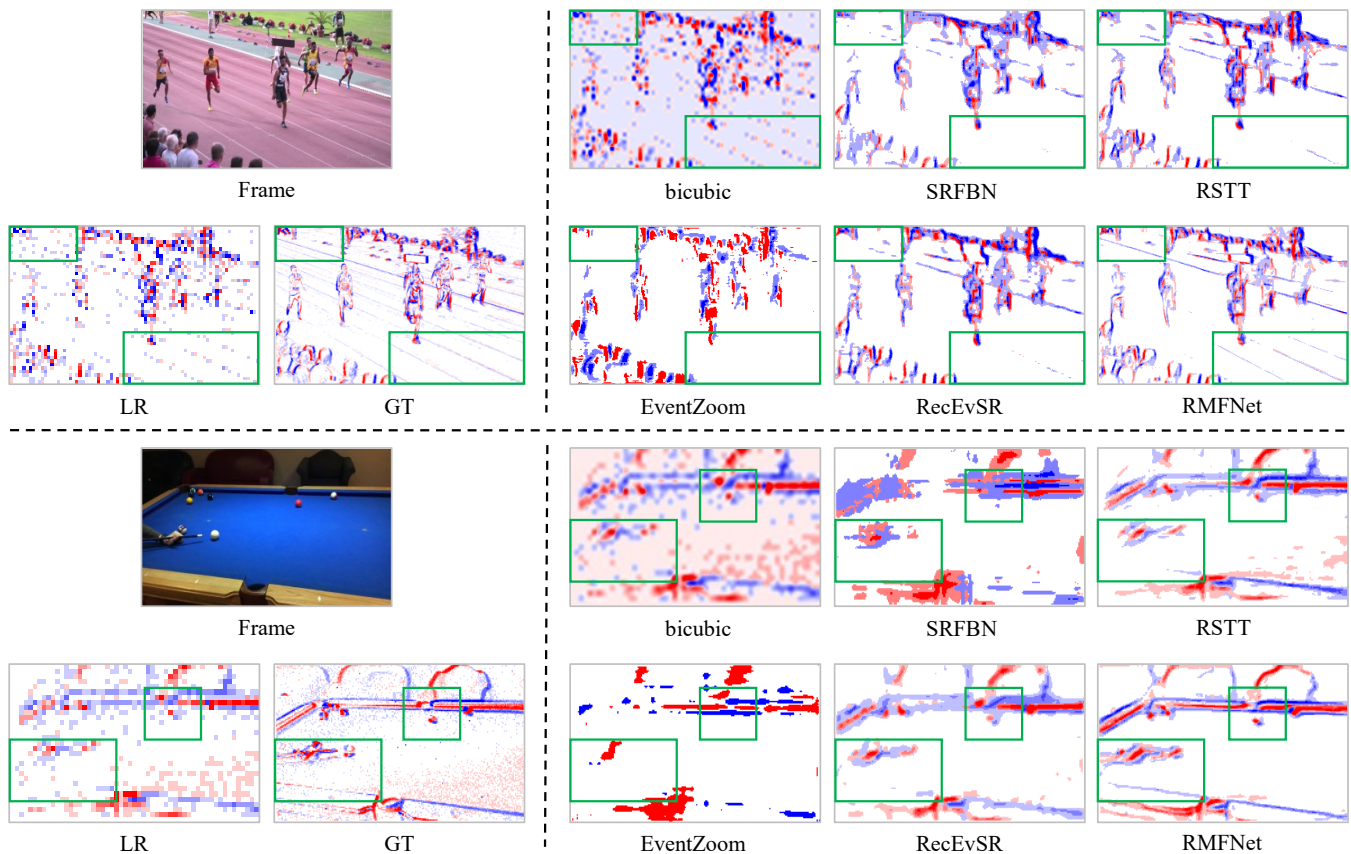


Figure 3: Qualitative analysis comparison on synthetic and real datasets. The upper and lower figures represent $4\times$ SR results on the NFS-syn and EventNFS datasets, respectively. It is evident that our RMFNET excels in recovering finer details of the event streams on both datasets (see the green box), resulting in sharper edges. Positive events are in blue, negative events in red. **Zoom in for the best view.**

4 Experiments

4.1 Datasets and Training Settings

Obtaining event data is challenging, and the availability of event datasets containing LR-HR pairs at multiple scales is limited. To address this scarcity, similar to many event-based tasks [Weng *et al.*, 2022; Rebecq *et al.*, 2019; Wang *et al.*, 2020a], we employed synthetic simulation datasets to enrich our training data. EventNFS [Duan *et al.*, 2021] is the first dataset to include LR-HR pairs captured through a designed display-camera system, capturing rapidly displayed images on a monitor. However, due to device resolution limitations, the minimum resolution is 55×31 , and there are only $4\times$ data pairs at the maximum. Moreover, the data at the smallest resolution suffers severe degradation due to its low resolution. To overcome these issues, we utilized an event simulator [Lin *et al.*, 2022] to transform the NFS dataset [Kiani Galoogahi *et al.*, 2017] and RGB-DAVIS dataset [Wang *et al.*, 2020b] into event data, resulting in NFS-syn and RGB-syn datasets. We selected these datasets because of their high temporal resolution, which can better simulate real-world event streams. For further details, please refer to the **Supplementary Material**.

For a fair comparison, we maintained training settings consistent with [Weng *et al.*, 2022]. We used *MSE* as the evaluation metric for our models. All experiments were conducted

on a Tesla V100 GPU.

4.2 Comparison with State-of-the-Art Models

In this work, we primarily compared our proposed RMFNET with two previous learning-based approaches, EventZoom [Duan *et al.*, 2021] and RecEvSR [Weng *et al.*, 2022]. Other ESR methods [Li *et al.*, 2021; Wang *et al.*, 2020b; Li *et al.*, 2019a] relying on real frames as assistance or prone to failure in complex scenes, pose challenges for fair comparisons. EventZoom, being the first learning-based event stream super-resolution method, faces challenges in training for large-scale SR due to its 3D-Unet architecture, making it difficult and computationally expensive. To address this, following previous practices [Weng *et al.*, 2022], we ran EventZoom- $2\times$ multiple times to obtain results for larger SR factors. Additionally, we include classic image super-resolution methods such as bicubic and SRFBN [Li *et al.*, 2019b], as well as a transformer-based video super-resolution method, RSTT [Geng *et al.*, 2022], for comparison. We randomly split the real dataset EventNFS for training and testing. To evaluate the model’s generalization, we select a subset of NFS-syn data for $2(4, 8)\times$ SR training and then validate on both the NFS-syn and RGB-syn datasets.

Qualitative Analysis Results. As depicted in Figure 3, we present the $4\times$ SR results of various methods on both syn-

Methods	NFS-syn			RGB-syn		EventNFS-real		Param (M)			Inference time (ms)		
	2×	4×	8×	2×	4×	2×	4×	2×	4×	8×	2×	4×	8×
bicubic	0.616	0.531	0.545	0.1197	0.1429	0.760	0.899	-	-	-	-	-	-
SRFBN	0.411	0.394	0.394	0.1051	0.1010	0.415	0.545	<u>2.1</u>	3.6	7.9	37.3	54.8	65.4
RSTT	<u>0.389</u>	<u>0.366</u>	0.365	<u>0.0954</u>	<u>0.0909</u>	<u>0.310</u>	<u>0.399</u>	3.8	4.1	4.3	61.4	61.1	73.0
EventZoom	0.806	1.049	1.239	0.4462	1.2232	0.778	1.248	11.5	11.5	11.5	17.4	70.1	396.6
RecEvSR	0.430	0.368	<u>0.332</u>	0.5013	0.3360	0.376	0.449	1.8	1.8	1.8	<u>13.2</u>	<u>18.9</u>	<u>19.2</u>
Ours	0.316	0.300	0.305	0.0899	0.0865	0.250	0.316	3.0	<u>3.1</u>	<u>3.6</u>	7.0	7.5	7.8

Table 1: Quantitative analysis comparison on real and synthetic datasets. Mean Squared Error (MSE) is used as the evaluation metric. Model Parameters (Param) and Inference time are calculated on the NFS-syn dataset. **Bold** and underline indicate the best and second-best results.

thetic and real data (for more results, please refer to the **Supplementary Material**). It can be observed that traditional image super-resolution methods such as bicubic and SRFBN [Li *et al.*, 2019b] struggle to achieve satisfactory visual results in ESR tasks, exhibiting blurry edges and significant detail loss. This may be attributed to the gap between event stream and RGB images. EventZoom [Duan *et al.*, 2021], on the other hand, exhibits numerous detail losses, likely due to error accumulation from multiple runs of EventZoom-2×. In comparison, RSTT [Geng *et al.*, 2022] and RecEvSR [Weng *et al.*, 2022] produce event images of higher quality, yet they still fall short in detail restoration and supplementation. In contrast, our proposed RMFNET can better extract detailed information from positive and negative event streams and complement each other, resulting in more comprehensive details and clearer edge information.

Quantitative Analysis Results. As shown in Table 1, compared to the previously SOTA ESR method RecEvSR, RMFNET achieves an average MSE improvement of 17.7% and 31.6% on NFS-syn and EventNFS, respectively. On RGB-syn, RecEvSR exhibits fragile generalization, while our RMFNET maintains good generalization with an average MSE improvement of 78%. Additionally, the average inference speed is improved by 2.3×. Compared to the video super-resolution method RSTT, RMFNET achieves an average MSE improvement of 17.8% and 20% on NFS-syn and EventNFS, respectively. On RGB-syn, while RSTT maintains good generalization, our RMFNET still outperforms it with a 5.3% improvement. Furthermore, our inference speed is improved by 8.7×. These results demonstrate the efficiency and robustness of our proposed RMFNET.

4.3 Analysis of ESRDA Methods

As shown in Table 2, we compared the impact of different DA methods on our RMFNET for the 4× ESR task. To better highlight the influence of data augmentation methods on the generalization of our model, we only conducted training on NFS-syn and performed testing on NFS-syn, RGB-syn, and EventNFS. It can be observed that Polarity flipping, RandomFlip, and Drop by time all contribute to performance gains in the ESR task. However, Static translation leads to training instability, while RandomResizedCrop and Drop by area result in a decline in the performance and generalization of our RMFNET. This suggests that altering the relative spatial rela-

Method	NFS-syn	RGB-syn	EventNFS
RMFNET (w/o DA)	0.304	0.0874	0.795
Polarity flipping	0.304	0.0865 ↓	0.793 ↓
RandomFlip	0.302 ↓	0.0868 ↓	0.793 ↓
Drop by time	0.304	0.0873 ↓	0.790 ↓
Random drop	0.306 ↑	0.0880 ↑	0.785 ↓
Drop by area	0.306 ↑	0.0881 ↑	0.783 ↓
Random drop or add noise	0.303 ↓	0.0871 ↓	0.786 ↓
Static Translation	-	-	-
RandomResizedCrop	0.330 ↑	0.0921 ↑	0.799 ↑
Selected DA's (random)	0.300 ↓	0.0865 ↓	0.771 ↓

Table 2: Comparison of different data augmentation methods in ESR task. Training is conducted on the NFS-syn dataset, and 4× SR testing is performed on NFS-syn, RGB-syn, and EventNFS datasets.

tionships between events may adversely affect the ESR task. This could be attributed to the sparse and unidimensional nature of event streams, lacking important features such as color and intensity. Therefore, disrupting the relative spatial relationships among event streams significantly impacts the overall structure, introducing additional noise and consequently leading to a decline in model performance.

Random drop only discards a portion of events in the LR event stream, introducing potential biases in model fitting. To address this, we propose Random drop or add noise, where events are not only dropped with a certain probability but noise is also added simultaneously, mitigating this issue and enhancing model robustness. Lastly, inspired by RandAugment [Cubuk *et al.*, 2020], we combine Polarity flipping, RandomFlip, Drop by time, and Random drop or add noise into a data augmentation ensemble, from which one augmentation is randomly selected (Selected DA). Experimental results demonstrate that our DA strategy effectively enhances the performance and generalization of our model in the ESR task. For more related experiments, please refer to the **Supplementary Material**.

4.4 Ablation Study

To validate the effectiveness of different components in our proposed RMFNET, we conducted experiments with four different variants and compared the 4× SR results on the

Model	Multi-Branch	FFM	FEM	NFS-syn	EventNFS
model#A	✗	✗	✗	0.329	0.347
model#B	✓	✗	✗	0.317	0.331
model#C	✓	✗	✓	0.309	0.322
model#D	✓	✓	✗	0.313	0.326
model#E	✓	✓	✓	0.300	0.316

Table 3: Ablation results for different components of our RMFNET.

NFS-syn and EventNFS datasets.

As shown in Table 3, we compared RMFNET with several variants with different settings: 1) **model#A**: using a single-branch model, concatenating event images and state h_t at the initial stage, and then inputting them into the model. 2) **model#B**: discarding FFM and FEM modules, using lateral connections [Feichtenhofer *et al.*, 2019; Christoph and Pinz, 2016] between branches as an alternative. 3) **model#C**: discarding the FFM module, using lateral connections between branches as an alternative. 4) **model#D**: discarding the FEM module, using lateral connections between positive and negative branches as an alternative.

According to the results in Table 3, the multi-branch model significantly outperforms the single-branch model, as it effectively decouples different parts of the event stream, allowing for fine-grained learning of each part’s features. Additionally, the FFM and FEM designed in our model efficiently fuse and exchange features from different branches, promoting information complementarity between positive and negative event streams, outperforming methods that directly mix features from different branches. For more details about model hyperparameter ablation experiments, please refer to the **Supplementary Material**.

4.5 Event-based Applications

Video Reconstruction. Video reconstruction is a crucial task within event-based applications [Rebecq *et al.*, 2019; Stoffregen *et al.*, 2020; Weng *et al.*, 2021; Liang *et al.*, 2023; Yang *et al.*, 2023]. Firstly, we conducted $2(4, 8)\times$ SR on NFS-syn using bicubic, SRFBN [Li *et al.*, 2019b], RSTT [Geng *et al.*, 2022], EventZoom [Duan *et al.*, 2021], RecEvSR [Weng *et al.*, 2022], and our RMFNET. Subsequently, we adopt E2VID [Rebecq *et al.*, 2019] as the benchmark algorithm for event-based video reconstruction and utilize the structural similarity (SSIM) [Wang *et al.*, 2004] and the perceptual similarity (LPIPS) [Zhang *et al.*, 2018] as evaluation metrics for reconstruction quality. Table 4 presents the quantitative results for event-based video reconstruction, indicating that our method surpasses others in both SSIM and LPIPS metrics, and exhibits more visually satisfying details (see **Supplementary Material**). This further underscores our method’s capability to better restore details in the LR event stream.

Object Recognition. We also perform a comparison of all models and methods in the event-based object recognition task. In this context, following the methodology of Weng *et al.* [Weng *et al.*, 2021], we employ the NCars dataset [Sironi *et al.*, 2018] for experimentation and leverage the classifier proposed by Gehrig *et al.* [Gehrig *et al.*, 2019]

Methods	Video Reconstruction					
	2 \times		4 \times		8 \times	
	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow	SSIM \uparrow	LPIPS \downarrow
bicubic	0.568	0.395	0.609	0.522	0.598	0.545
SRFBN	0.608	0.389	0.618	0.455	0.612	0.489
RSTT	<u>0.627</u>	<u>0.359</u>	<u>0.639</u>	<u>0.424</u>	0.622	0.472
EventZoom	0.542	0.429	0.575	0.488	0.574	0.542
RecEvSR	0.611	0.371	0.637	0.426	<u>0.630</u>	<u>0.466</u>
RMFNET	0.648	0.339	0.667	0.409	0.653	0.450

Methods	Object Recognition					
	ACC \uparrow	AUC \uparrow	ACC \uparrow	AUC \uparrow	ACC \uparrow	AUC \uparrow
bicubic	56.67	57.43	56.01	56.89	49.95	50.77
SRFBN	61.12	61.94	60.89	61.03	50.02	50.86
RSTT	<u>63.51</u>	<u>63.96</u>	<u>63.02</u>	<u>64.29</u>	52.97	54.07
EventZoom	54.68	56.03	49.56	50.45	47.96	48.74
RecEvSR	62.91	63.47	62.37	63.07	<u>53.57</u>	<u>54.48</u>
RMFNET	68.75	69.56	69.52	69.80	58.16	59.05

GT	85.16	84.99	93.44	93.52	94.96	94.81
----	-------	-------	-------	-------	-------	-------

Table 4: Quantitative comparison for event-based video reconstruction and object recognition. Video reconstruction is conducted on the NFS-syn dataset, while object recognition is performed on the NCars dataset [Sironi *et al.*, 2018]. AUC and ACC represent accuracy and area under the curve, respectively. GT denotes the result obtained by directly using downsampled event streams for recognition. **Bold** and underline indicate the best and second-best results.

for object recognition. Specifically, we first performed $8\times$ downsampling on the NCars dataset through coordinate relocation. Subsequently, we employ different models to conduct $2(4, 8)\times$ super-resolution on the event stream and employ the object recognition method for identification. Table 4 illustrates the results of object recognition comparison. We evaluate using accuracy (ACC) and area under the curve (AUC). GT signifies the utilization of results directly obtained from downsampled raw event streams. It can be observed that our method outperforms other approaches consistently across $2(4, 8)\times$ super-resolution scales. In comparison with previous methods, our approach achieves an average improvement of over 9% in terms of ACC and AUC. These results demonstrate the superior detail restoration capability of our method.

5 Conclusion

In this paper, we introduced an efficient Recursive Multi-Branch Information Fusion Network (RMFNET) for ESR tasks. RMFNET leverages a carefully designed multi-branch network architecture, taking decoupled positive and negative events as well as coupled event frames as input to achieve super resolution of event streams. Additionally, we introduced attention-based Feature Fusion Module and Feature Exchange Module, which effectively integrate contextual information from neighboring event streams and facilitate the exchange of complementary information between positive and negative events. Furthermore, we explored the impact of data augmentation methods on ESR tasks and proposed an effective data augmentation strategy to enhance model robustness and performance. Results on both real and synthetic datasets demonstrated that our approach outperforms previous ESR methods across various metrics.

Acknowledgments

This work was supported in part by the Guangxi Key R & D Program (No. GuikeAB24010324), in part by the National Natural Science Foundation of China (No. 62088102, No. 62425101), in part by the Key-Area Research and Development Program of Guangdong Province (No. 2021B0101400002), and the Major Key Project of PCL (No. PCL2021A13).

Contribution Statement

Quanmin Liang and Zhilin Huang made equal contributions. Kai Huang and Yonghong Tian are Corresponding Author. All the authors participated in designing research, analyzing data, and writing the paper.

References

- [Barchid *et al.*, 2023] Sami Barchid, José Mennesson, and Chaabane Djéraba. Exploring joint embedding architectures and data augmentations for self-supervised representation learning in event-based vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3902–3911, 2023.
- [Brandli *et al.*, 2014] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240×180 130 db $3 \mu\text{s}$ latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, 2014.
- [Christoph and Pinz, 2016] R Christoph and Feichtenhofer Axel Pinz. Spatiotemporal residual networks for video action recognition. *Advances in neural information processing systems*, 2, 2016.
- [Cubuk *et al.*, 2020] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020.
- [Duan *et al.*, 2021] Peiqi Duan, Zihao W Wang, Xinyu Zhou, Yi Ma, and Boxin Shi. Eventzoom: Learning to denoise and super resolve neuromorphic events. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12824–12833, 2021.
- [Feichtenhofer *et al.*, 2019] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019.
- [Gallego *et al.*, 2020] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2020.
- [Gehrig and Scaramuzza, 2022] Daniel Gehrig and Davide Scaramuzza. Are high-resolution event cameras really needed? *arXiv preprint arXiv:2203.14672*, 2022.
- [Gehrig *et al.*, 2019] Daniel Gehrig, Antonio Loquercio, Konstantinos G Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5633–5643, 2019.
- [Geng *et al.*, 2022] Zhicheng Geng, Luming Liang, Tianyu Ding, and Ilya Zharkov. Rstt: Real-time spatial temporal transformer for space-time video super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17441–17451, 2022.
- [Gu *et al.*, 2021] Fuqiang Gu, Weicong Sng, Xuke Hu, and Fangwen Yu. Eventdrop: Data augmentation for event-based learning. *arXiv preprint arXiv:2106.05836*, 2021.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Kiani Galoogahi *et al.*, 2017] Hamed Kiani Galoogahi, Ashton Fagg, Chen Huang, Deva Ramanan, and Simon Lucey. Need for speed: A benchmark for higher frame rate object tracking. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1125–1134, 2017.
- [Li *et al.*, 2019a] Hongmin Li, Guoqi Li, and Luping Shi. Super-resolution of spatiotemporal event-stream image. *Neurocomputing*, 335:206–214, 2019.
- [Li *et al.*, 2019b] Zhen Li, Jinglei Yang, Zheng Liu, Xiaomin Yang, Gwanggil Jeon, and Wei Wu. Feedback network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3867–3876, 2019.
- [Li *et al.*, 2021] Siqu Li, Yutong Feng, Yipeng Li, Yu Jiang, Changqing Zou, and Yue Gao. Event stream super-resolution via spatiotemporal constraint learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4480–4489, 2021.
- [Liang *et al.*, 2023] Quanmin Liang, Xiawu Zheng, Kai Huang, Yan Zhang, Jie Chen, and Yonghong Tian. Event-diffusion: Event-based image reconstruction and restoration with diffusion models. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3837–3846, 2023.
- [Lin *et al.*, 2022] Songnan Lin, Ye Ma, Zhenhua Guo, and Bihan Wen. Dvs-voltmeter: Stochastic process-based event simulator for dynamic vision sensors. In *European Conference on Computer Vision*, pages 578–593. Springer, 2022.
- [Maqueda *et al.*, 2018] Ana I Maqueda, Antonio Loquercio, Guillermo Gallego, Narciso García, and Davide Scaramuzza. Event-based vision meets deep learning on steering prediction for self-driving cars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5419–5427, 2018.
- [Rebecq *et al.*, 2017] Henri Rebecq, Timo Horstschafer, and Davide Scaramuzza. Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization. 2017.

- [Rebecq *et al.*, 2019] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence*, 43(6):1964–1980, 2019.
- [Schuster and Paliwal, 1997] Mike Schuster and Kuldip K Paliwal. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681, 1997.
- [Shi *et al.*, 2016] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [Sironi *et al.*, 2018] Amos Sironi, Manuele Brambilla, Nicolas Bourdis, Xavier Lagorce, and Ryad Benosman. Hats: Histograms of averaged time surfaces for robust event-based object classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1731–1740, 2018.
- [Stoffregen *et al.*, 2020] Timo Stoffregen, Cedric Scheerlinck, Davide Scaramuzza, Tom Drummond, Nick Barnes, Lindsay Kleeman, and Robert Mahony. 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*, pages 534–549. Springer, 2020.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [Wang *et al.*, 2004] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [Wang *et al.*, 2020a] Lin Wang, Tae-Kyun Kim, and Kuk-Jin Yoon. Eventsr: From asynchronous events to image reconstruction, restoration, and super-resolution via end-to-end adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8315–8325, 2020.
- [Wang *et al.*, 2020b] Zihao W Wang, Peiqi Duan, Oliver Cossairt, Aggelos Katsaggelos, Tiejun Huang, and Boxin Shi. Joint filtering of intensity images and neuromorphic events for high-resolution noise-robust imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1609–1619, 2020.
- [Weng *et al.*, 2021] Wenming Weng, Yueyi Zhang, and Zhiwei Xiong. Event-based video reconstruction using transformer. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 2543–2552. IEEE, 2021.
- [Weng *et al.*, 2022] Wenming Weng, Yueyi Zhang, and Zhiwei Xiong. Boosting event stream super-resolution with a recurrent neural network. In *European Conference on Computer Vision*, pages 470–488. Springer, 2022.
- [Yang *et al.*, 2023] Yixin Yang, Jin Han, Jinxiu Liang, Imari Sato, and Boxin Shi. Learning event guided high dynamic range video reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13924–13934, 2023.
- [Yoo *et al.*, 2020] Jaejun Yoo, Namhyuk Ahn, and Kyung-Ah Sohn. Rethinking data augmentation for image super-resolution: A comprehensive analysis and a new strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8375–8384, 2020.
- [Yun *et al.*, 2019] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.
- [Zhang *et al.*, 2017] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- [Zhang *et al.*, 2018] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 586–595. Computer Vision Foundation / IEEE Computer Society, 2018.
- [Zhu *et al.*, 2018] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Ev-flownet: Self-supervised optical flow estimation for event-based cameras. *arXiv preprint arXiv:1802.06898*, 2018.