

Temporal Event Stereo via Joint Learning with Stereoscopic Flow

Hoonhee Cho^{*}, Jae-Young Kang^{*}, and Kuk-Jin Yoon

Korea Advanced Institute of Science and Technology
{gnsgnsgml, mickeykang, kjoyoon}@kaist.ac.kr

Abstract. Event cameras are dynamic vision sensors inspired by the biological retina, characterized by their high dynamic range, high temporal resolution, and low power consumption. These features make them capable of perceiving 3D environments even in extreme conditions. Event data is continuous across the time dimension, which allows a detailed description of each pixel’s movements. To fully utilize the temporally dense and continuous nature of event cameras, we propose a novel temporal event stereo, a framework that continuously uses information from previous time steps. This is accomplished through the simultaneous training of an event stereo matching network alongside stereoscopic flow, a new concept that captures all pixel movements from stereo cameras. Since obtaining ground truth for optical flow during training is challenging, we propose a method that uses only disparity maps to train the stereoscopic flow. The performance of event-based stereo matching is enhanced by temporally aggregating information using the flows. We have achieved state-of-the-art performance on the MVSEC and the DSEC datasets. The method is computationally efficient, as it stacks previous information in a cascading manner. The code is available at <https://github.com/mickeykang16/TemporalEventStereo>.

Keywords: Stereo matching · Event camera · 3D from other sensors

1 Introduction

Depth estimation [40,66] is considered a fundamental problem within the field of computer vision, which can be adapted to a wide range of practical applications, including robotics and autonomous driving. Stereo matching [8,11,29,36,42,80,87–89,92,98], a prominent technique in depth estimation, identifies correspondences between rectified stereo images to compute depth. Recent stereo matching algorithms are tailored to conventional frame-based cameras. Yet, these sensors have limitations including high power consumption, limited dynamic range, and low data rate, constraining their use in mobile devices and edge computing platforms. On the other hand, drawing inspiration from the biological retina, the recently emerged dynamical vision sensor, also known as an event camera [6,43], addresses these issues by transmitting events representing instantaneous changes

^{*} Equal contribution.

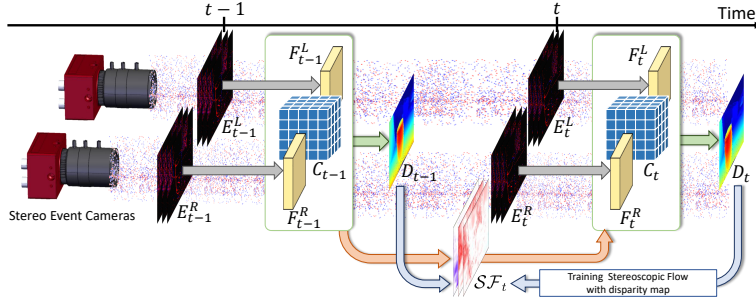


Fig. 1: Overview of the temporal event stereo. We can accurately and efficiently estimate dense disparity by propagating previously computed information to the present through stereoscopic flow (\mathcal{SF}).

in pixel intensity with millisecond-level temporal resolution. Event cameras can avoid motion blur caused by rapid object movement, and their ability to operate with high dynamic range and low power consumption could make them a game changer in autonomous driving and mobility systems.

By the nature of event cameras, transmitted events are sparse in space while containing dense temporal information. To accommodate the sparsity of the spatial domain, existing methods actively utilize the time information of events in the task of stereo matching. For example, Se-CFF [53] proposed a method that stacks events from a small to a large number and concentrates them into one, through an attention-based network. Another example, DTC [94] adopts a recurrent structure [22] to encode events sequentially, enriching current features. Although they have improved performance from temporal aggregation approaches, they lack a specialized approach to stereo matching. Specifically, temporal information transfer is done independently in left and right event, and limited to low-level representations (*e.g.*, feature map). Therefore, the limited data transfer results in sub-optimal performance. We made a breakthrough by proposing temporal event stereo, which uses a flow-based temporal aggregation.

Unlike the typical optical flow [4, 26, 32–34, 46, 69, 72, 81, 102] that calculates pixel displacement of a single camera, we define “stereoscopic flow” considering the specificity of stereo matching. Each corresponding point from two stereo cameras lies on a single horizontal epipolar line. Even if there is vertical pixel displacement over consecutive time steps, the matching points are placed on an epipolar line. Stereoscopic flow involves a triplet flow, defined in rectified stereo image pairs, consisting of two independent horizontal pixel displacements for each camera and a single vertical pixel displacement. As shown in Fig. 1, we reuse previous features and cost volumes by temporally propagating them through stereoscopic flow. The previous model [94], which uses recurrent structure, requires independent networks to aggregate different types of representations. However, our approach enables the aggregation of not only low-level features but also stereo-dependent information such as a cost volume with mini-

mal effort. Instead of requiring independent modules for each component, we can efficiently retrieve previous information through a warping operation using the once-calculated stereoscopic flow. Furthermore, without using additional and independent two optical flow networks, we can calculate temporal pixel movement with just a single stereoscopic flow network composed of very few convolutions as we jointly learn using high-level representation features learned from the stereo network. To ensure the stable training of stereoscopic flow, we introduce a novel and effective temporal disparity consistency loss. This loss facilitates the training of stereoscopic flow when two stereo cameras and a disparity map are available, offering a strategy to train stereoscopic flow in the absence of optical flow ground truth. Finally, we propose feature warping and cost volume warping to merge previous information with current inputs effectively, achieving state-of-the-art (SOTA) performance on the MVSEC and DSEC datasets.

Our contributions are summarized as follows: 1) We propose feature warping and cost volume warping to fuse previous information through stereoscopic flow effectively. 2) We propose a framework that jointly trains stereoscopic flow, as defined by a pair of stereo cameras, and stereo matching. 3) We introduce a temporal disparity consistency loss that enables the learning of stereoscopic flow even without optical flow ground truth. 4) Our approach significantly improves performance while being computationally efficient on MVSEC and DSEC benchmarks.

2 Related Works

Event-based Stereo Matching. Unlike traditional cameras, event cameras [23] capture the scene as a stream of sparse events, representing per-pixel intensity changes exceeding a set threshold. Their asynchronous nature makes them less susceptible to motion blur [38, 44, 56, 70, 71, 84, 95, 97] from rapid scene changes or movements. Event cameras, with their enhanced dynamic range [30, 61, 64], capture scenes that traditional cameras may miss, such as night scenes [3, 18, 78, 83, 96], making them ideal for driving scenarios.

Early event-based stereo matching utilized the hand-crafted method [7, 13, 39, 62, 63, 65, 99, 104, 105] to determine corresponding events between asynchronous streams. Recent works [1, 12, 53, 94] adopt deep learning-based approaches by processing events into image-like forms [17–19, 41, 46, 48–51, 54, 79, 102, 103] to accommodate the property of the sparsity of event data and to benefit from the knowledge of image-based deep learning stereo methods. To overcome the shortcomings of an event density imbalance across spatial dimensions, Se-CFF [53] proposed the multi-density stack event representation. Furthermore, DTC [94] aggregated sequential event information at the feature level using the discrete time convolutional module. As they have focused on low-level feature implementation, there is still potential for further advancement. It is possible to expand existing methods [53, 94] so that they can deal with higher-level representations (*e.g.* cost volume) by adding aggregation modules. However, the more information crosses from the past, the more computational cost should be added to the

model, which forces a trade-off between information and complexity. To make a breakthrough, we propose a method of flow-based temporal aggregation, a cost-effective method to gather multiple types of information. Our flow-based aggregation shares a one-time estimated flow in both feature warping and cost volume warping, adding minimal computation.

Learning with Event-based Optical Flow. Event cameras efficiently capture per-pixel changes despite their sparse data, proving highly useful for optical flow estimation [2, 5, 20, 27, 31, 45, 47, 55, 59, 76, 77, 101]. Their capability extends to capturing motion effortlessly, aiding in tasks like video frame interpolation [16, 37, 73, 75], image reconstruction [15, 21, 57], and motion deblurring [14, 35, 85]. These works focus on a main task while setting optical flow from events as an auxiliary task for joint learning to assist. The optical flow is used for temporal propagation and alignment, boosting the performance of the main task.

In our work, stereoscopic flow is estimated from stereo events to determine the pixel movement of stereo cameras. Stereoscopic flow is used to temporally transfer the features and cost volume to the next time step. Previous works [57, 85] designed independent flow estimation networks. However, inspired by the fact that both stereo and flow networks match features between frames, our flow network shares features from the stereo network, allowing for cost-effective design. To the best of our knowledge, our method is the first to propose the joint learning of event-based optical flow for stereo matching.

Event-based Optical Flow without Ground-truth. Events are continuous points containing temporal information, and by analyzing their motion over time, the displacement of each pixel can be effectively calculated. Some studies [24, 58, 68, 90, 103] have demonstrated distinction in optical flow through unsupervised learning by effectively utilizing the properties of events. Loss functions that solely rely on events face convergence issues depending on the initialization, as mentioned in [53, 67, 68]. This problem occurs when trained in an end-to-end manner with a stereo matching network. Also, as it is an edge-centric loss, we propose a new temporal disparity consistency loss.

3 Methodology

3.1 Preliminary

Event streams of left and right cameras between $t-1$ and t are represented as \mathcal{E}_t^L and \mathcal{E}_t^R , respectively. An event stream is a sequence of four-dimensional vectors (i, j, p, τ) . (i, j) represents pixel coordinate while $p \in \{-1, 1\}$ and $\tau \in (t-1, t]$ represents polarity and time respectively.

Event Representation. Following [102], event streams $\mathcal{E}_t = \{(i, j, p, \tau) | t-1 \leq \tau < t\}$ are stacked to generate an event voxel grid $E_t(b, i, j) \in \mathbb{R}^{B \times H \times W}$ with discretized time dimension B , scaled to $b \in [0, B-1]$.

Stereoscopic Flow. Given a pair of rectified stereo event voxels (E_t^L, E_t^R) , matching points in stereo voxels are on an epipolar line, which is a single horizontal line for rectified stereo. Since the epipolar constraint remains for all stereo

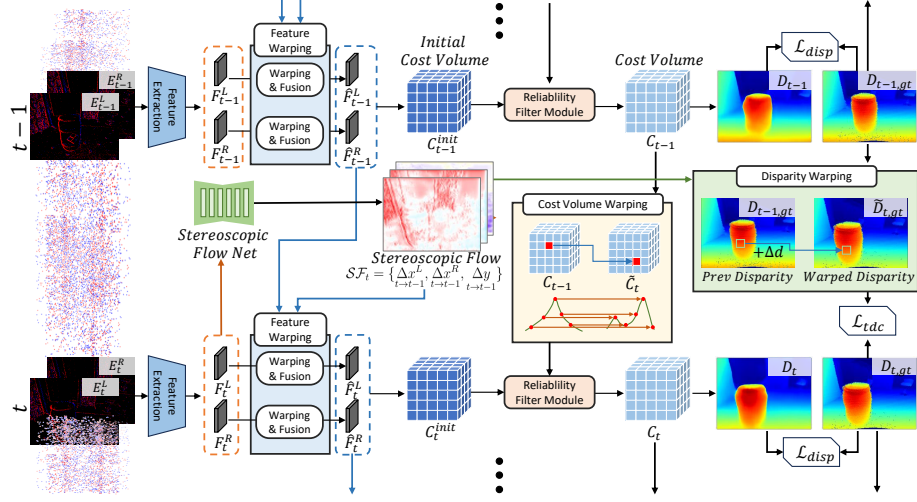


Fig. 2: Architecture overview of the proposed temporal event stereo. In each time step, information from the preceding moment is warped via stereoscopic flow and fused with the current information, boosting the intermediate representations such as feature map and cost volume. Stereoscopic flow training does not require the ground truth optical flow but receives a supervision signal from ground truth disparity.

pairs, it can be assumed that each matching pixels in left and right voxels experience the same vertical shift, while their horizontal shift may differ due to depth changes. Therefore, the flow of stereo voxel pairs can be expressed as a set of two horizontal flows for each voxel and a single shared vertical flow. We use the notation of $S\mathcal{F}_t = \{\Delta x_{t \rightarrow t-1}^L, \Delta x_{t \rightarrow t-1}^R, \Delta y\}_{t \rightarrow t-1} \in \mathbb{R}^{3 \times H \times W}$, for our stereoscopic flow.

Please note that all flows are defined as backward flows, from time frame t to $t-1$, to utilize backward warping. The terms stereoscopic flow, stereo flow, and $S\mathcal{F}$ will be used interchangeably throughout the paper.

3.2 Overall Architecture Design

As shown in Fig. 2, we design the joint framework for the temporal event-based stereo and flow network, that can be trained in an end-to-end manner. Following the well-working art [9, 29, 86, 88, 91], the main streamline of the stereo network is composed of four commonly used parts: 1) feature extraction, 2) cost volume generation, 3) cost aggregation, and 4) cost refinement. Conventional stereo modules are omitted in Fig. 2 for brevity and ease of understanding.

Feature Extraction. At each time step, stereo event voxel pair (E_t^L, E_t^R) is fed into the weight-shared feature encoder. These initial features (F_t^L, F_t^R) are concatenated and used to predict stereoscopic flow. With stereoscopic flow, the previous event features are warped and fused with the current features.

Stereoscopic Flow Network. Stereoscopic flow, the key feature of our design, plays an important role when aggregating various representations from the previous time step. Unlike the RNN base temporal aggregation methods from [94] that require extra RNN blocks for different representations, flow can be used to warp various types (*e.g.*, feature, cost volume, reliability map, and disparity) with minimal extra computational cost. The stereoscopic flow network shares encoded event features from the stereo network, allowing for using high representation features in feature matching. We can estimate stereo flow with a single network consisting of only 9 layers of 3×3 2D convolutions, enjoying the temporally dense property of event data. F_t^L and F_t^R are concatenated, and fed into the stereoscopic flow network to obtain the stereoscopic flow.

The stereoscopic flow is trained with a proposed temporal disparity consistency loss (see Eq.(7)) of ground-truth (GT) disparity. Temporal disparity consistency loss is a special flow loss that we introduce to train the stereoscopic flow and disparity collaboratively without using auxiliary GT optical flow. The error between previous GT disparity, warped from previous time $t - 1$ to current time t , and current GT disparity defines the disparity consistency loss.

Cost Volume Construction and Aggregation. Following [36], a pair of temporally fused features is channel-wise concatenated across all disparity levels, resulting in an initial cost volume $C_t^{init} \in \mathbb{R}^{C \times D \times H \times W}$. It passes two cost aggregation modules which are lightweight versions of the 3D hourglass [29]. Initial disparities, D_t^0 and D_t^1 are predicted after each module. Note that D_t^0 and D_t^1 are generated for supervision only during the training process for robust cost volume generation and are not computed during inference.

Cost Refinement via Temporal Aggregation. The cost volume is a key representation module that holds task-specific 3D information crucial for stereo matching. Therefore, we propose a way to benefit from the past cost volume. The 4D past cost is warped to the current time frame with our novel stereoscopic flow. Then, the entropy-based reliability fusion module efficiently fuses current and past costs. An hourglass-like refinement block finally refines the fused cost, making a precise disparity prediction.

Disparity Regression. A channel aggregation block squeezes the channel of the fused cost volume, resulting in a probability volume. Finally, the predicted disparity is computed by the soft argmin function [36].

3.3 Benefit from the Past

Warping and fusion of representations from the previous time steps enrich the intermediate representations (*e.g.* feature, cost volume), improving stereo matching performance cost-effectively. In our design, we propose two modules for temporal information warping: feature warping and cost volume warping. Intermediate representations of $t - 1$ are warped to current time t with stereoscopic flow.

Feature Warping. Events are triggered by changes in intensity, which are mostly generated by ego-motion. Often, due to little motion or static scenes, the current scene may receive too few events. This can be compensated by temporally accumulating information from the past. At each time step, previous features are




Fig. 3: Cost Volume Warping. Cost volume warping with 3-dimensional flow $\{\Delta d, \Delta x^L, \Delta y\}$ (left). Relation between disparity flow and stereoscopic flow (right).

warped in spatial (H, W) dimensions, keeping the channel value of each pixel. We can obtain the warped feature \hat{F}_t^L by warping past feature using the backward flow f_t^L and warping function \mathcal{W}_s , as follows: $\hat{F}_t^L = \mathcal{W}_s(F_{t-1}^L, f_t^L)$. The spatial warping function \mathcal{W}_s , which is a typical backward warping, with backward flow components $f_t^L = \{\Delta x^L, \Delta y\}_{t \rightarrow t-1}$ from \mathcal{SF}_t can be defined as:

$$\mathcal{W}_s(F_{t-1}^L, f_t^L)(c, h, w) = F_{t-1}^L(c, h + \Delta y_{t \rightarrow t-1}(h, w), w + \Delta x^L_{t \rightarrow t-1}(h, w)), \quad (1)$$

where $F_{t-1}^L, \hat{F}_t^L \in \mathbb{R}^{C \times H \times W}$. The same function \mathcal{W}_s is applied to the previous right feature F_{t-1}^R with flow $f_t^R = \{\Delta x^R, \Delta y^R\}_{t \rightarrow t-1}$. $\Delta y^R \in \mathbb{R}^{H \times W}$ is a vertical flow of the right frame, estimated along with \mathcal{SF}_t for feature warping. Warped features, \hat{F}_t^L and \hat{F}_t^R , are concatenated in channel dimension with the current event features F_t^L and F_t^R , respectively. Then, 3×3 2D convolution is applied to squeeze the channel to its original size.

Cost Volume Warping. Cost volume is a core structure that encloses the information on the matching similarity of a stereo pair. Hence, cost volume summarizes 3D information about the scene and undergoes extensive computations to achieve an optimal aggregation. Many existing works [53, 86, 94] discard the cost volume after computing it for a static time, even if most of the 3D depiction of a scene remains in future frames. We continuously reuse and temporally aggregate cost volumes to achieve a cost-effective design.

The previous cost volume $C_{t-1} \in \mathbb{R}^{C \times D \times H \times W}$ cannot be directly reused without warping, as it is misaligned to the current initial cost volume C_t^{init} . Two cost volumes in different time steps are misaligned in two ways. First, as the relative pose between an object and cameras differs, there should be spatial location misalignments on two cost volumes. This spatial movement can be easily estimated as $\{\Delta x^L, \Delta y\}_{t \rightarrow t-1}$ since a cost volume is defined on the left frame coordinate similar to the feature warping process. Second, as the object distance changes, the D-dimension of cost volume suffers a shift. Assume a 3D coordinate $(\delta_{t-1}, i_{t-1}, j_{t-1})$ of a previous cost volume C_{t-1} where $\delta_{t-1} \in [1, D], i_{t-1} \in [1, H], j_{t-1} \in [1, W]$. A coordinate (δ_t, i_t, j_t) refers to the matching score

of stereo pairs for a disparity candidate δ_{t-1} . Therefore, the disparity change of an object would redistribute the matching scores across all disparity candidates on D -dimension.

To compensate misalignments on cost volumes, we propose a 3D flow $f_t^C = \{ \Delta d_{t \rightarrow t-1}, \Delta x_{t \rightarrow t-1}^L, \Delta y_{t \rightarrow t-1} \}$, which is a backward flow defined on $D \times H \times W$ dimensions of the current cost volume. $\Delta d_{t \rightarrow t-1} \in \mathbb{R}^{D \times H \times W}$ and $\Delta x_{t \rightarrow t-1}^L, \Delta y_{t \rightarrow t-1} \in \mathbb{R}^{H \times W}$ because all δ in a spatial coordinate (i, j) shares a spatial flow. The concept diagram of 3D flow f_t^C can be found on the left side of Fig. 3. For understanding, arrows are aligned to the forward flow only within the figure.

In ordinary cases, it is impossible to estimate the disparity flow Δd for all disparity candidates $\delta \in [1, D]$ since there is no clue of where an arbitrary disparity candidate δ_{t-1} would be in time t . However, as a stereoscopic flow defines the independent flow for all coordinates in left and right event voxel in our model, we can define a relation between disparity flow and stereoscopic flow. For arbitrary point (i_{t-1}, j_{t-1}) on a left frame and a matching point with arbitrary disparity δ_{t-1} on a right frame, which is located in $(i_{t-1} - \delta_{t-1}, j_{t-1})$, we can track the movement of both points across time with stereoscopic flow. Knowing the locations of two points in $t-1$ and t indicates that we know the disparity changes for arbitrary coordinates and disparities across time. The right side of Fig. 3 illustrates that the change of disparity can be expressed as the difference of Δx^L and Δx^R , which can be written as Eq.(2). For clarity, all flows are defined as backward flows to use backward warping.

$$\Delta d_{t \rightarrow t-1}(\delta_t, i_t, j_t) = \Delta x_{t \rightarrow t-1}^L(i_t, j_t) - \Delta x_{t \rightarrow t-1}^R(i_t - \delta_t, j_t) \quad (2)$$

At every depth estimation time step, the previous time cost volume C_{t-1} is aligned to the current cost volume in three dimensions D, H , and W , using the 3D backward flow f_t^C . 3-dimensional backward warping is applied to the previous cost volume, resulting in a temporally aligned cost \tilde{C}_t , as depicted in Eq.(3). This novel approach allows us to fully utilize past stereo knowledge. As we recycle the stereoscopic flow, the only additional computational cost is the warping function itself.

$$\begin{aligned} \mathcal{W}_c(C_{t-1}, f_t^C)(c, d, h, w) \\ = C_{t-1}(c, d + \Delta d_{t \rightarrow t-1}(d, h, w), h + \Delta y_{t \rightarrow t-1}(h, w), w + \Delta x_{t \rightarrow t-1}^L(h, w)) \end{aligned} \quad (3)$$

Temporally aligned cost volume \tilde{C}_t is selectively fused with the current cost volume C_t based on their reliability. To balance computational cost and performance, we use the reliability filter module which is shown to be useful in [10, 60, 82, 93]. A disparity probability volume $P_t \in \mathbb{R}^{D \times H \times W}$ is calculated from the initial cost volume C_t^{init} by channel aggregation block, which is composed of 3D convolutions. Then, the entropy map is calculated as the Eq.(4).

$$E_t(i, j) = \sum_{d=0}^{D-1} -p_{i,j} \cdot \log p_{i,j} \quad (4)$$

The entropy map of the previous cost, $E_{t-1} \in \mathbb{R}^{H \times W}$, is calculated in the same way and warped to the current time frame with the warping function \mathcal{W}_s and flow f_t^L . The previous and the current entropy map is concatenated and processed with a few 2D convolution layers to calculate the pixel-wise weight of two cost volumes. Weighted sum of C_t and \tilde{C}_t finalizes the cost fusion process.

3.4 Temporal Disparity Consistency Loss

We propose the temporal disparity consistency (TDC) loss to provide a supervision signal for the stereoscopic flow network without extra flow ground truth. At a training phase, a previous GT disparity is warped to a current frame, compared with a current ground truth. When warping a disparity map, there should be a change in disparity value in a disparity map due to the relative pose of the camera and object as well as spatial coordinate shift. Therefore, the change in disparity value, which we define as residual disparity, is to be added to the warped disparity. As shown in the right side of Fig. 3, the residual disparity between two time steps can be calculated as the difference in pixel movement along the horizontal line between matching points in stereo cameras, which can be derived from the stereoscopic flow. Specifically, we apply the backward function \mathcal{W}_s to $\Delta x_{t \rightarrow t-1}^R$, which is defined in the right frame coordinate, to align with the left frame coordinate where disparity and cost volume are defined. Then, the residual disparity, $\Delta D_{t,gt}$, between $t-1$ and t can be calculated using current disparity ground-truth, $D_{t,gt}$, as:

$$\Delta D_{t,gt} = \mathcal{W}_s(\Delta x_{t \rightarrow t-1}^R, \{-D_{t,gt}, 0\}) - \Delta x_{t \rightarrow t-1}^L \quad (5)$$

The previous GT disparity $D_{t-1,gt} \in \mathbb{R}^{H \times W}$ can be aligned to the current time step using left camera flow components $\Delta x_{t \rightarrow t-1}^L, \Delta y_{t \rightarrow t-1}$ from \mathcal{SF}_t . As shown in Eq.(6), the first term is in charge of spatial pixel movement, while the residual disparity term deals with the disparity value change of each pixel.

$$\tilde{D}_{t,gt} = \mathcal{W}_s(D_{t-1,gt}, \{\Delta x_{t \rightarrow t-1}^L, \Delta y_{t \rightarrow t-1}\}) + \Delta D_{t,gt} \quad (6)$$

We define the error between current GT disparity and warped previous GT disparity as the TDC loss using smooth L1 loss [28] as:

$$\mathcal{L}_{tdc} = \text{Smooth}_{L_1}(D_{t,gt}, \tilde{D}_{t,gt}) \quad (7)$$

3.5 Objective Functions

Flow Loss. We utilize the proposed \mathcal{L}_{tdc} (Eq.(7)) to optimize stereoscopic flow from events along with stereo matching loss, even in the absence of ground truth for optical flow. For better flow estimation on edges, contrast maximization loss [57, 102], \mathcal{L}_c , is also used as follows:

$$\mathcal{L}_{flow} = \lambda_t \cdot \mathcal{L}_{tdc} + \lambda_c \cdot \mathcal{L}_c(e_t) \quad (8)$$

Table 1: Results for event-based stereo on MVSEC and DSEC datasets. * indicates that images were also used in the training. - indicates that official results are not provided in the original paper. The FPS indicates the GPU synchronized inference time on the MVSEC, using an NVIDIA TITAN Xp. The FPS measures are an average of 1000 iterations after 100 warm-up loops. “Sp” is an abbreviation for “split”.

Method	MVSEC								DSEC				Cost	
	Mean depth error [cm] ↓		Median depth error [cm] ↓		Mean disp. error [pix] ↓		1PA [%] ↑		1PE	2PE	MAE	RMSE	FPS	GFlops
	Sp 1 Sp 3		Sp 1 Sp 3		Sp 1 Sp 3		Sp 1 Sp 3		↓	↓	↓	↓	↑	↓
	Sp 1	Sp 3	Sp 1	Sp 3	Sp 1	Sp 3	Sp 1	Sp 3						
DDES [74]	16.7	27.8	6.8	14.7	0.59	0.94	89.8	74.8	10.915	2.905	0.576	1.381	34.3	75.8
DTC-PDS [94]	15.3	18.6	6.4	8.7	0.56	0.65	91.5	88.7	9.517	2.356	0.527	1.264	33.2	63.9
CTC-PDS [94]	14.9	20.6	6.4	10.6	0.53	0.73	91.6	88.2	-	-	-	-	-	-
EITNet* [1]	14.2	19.4	5.9	10.4	0.55	0.75	92.1	89.6	-	-	-	-	-	-
Se-CFF [53]	-	-	-	-	-	-	-	-	9.583	2.620	0.519	1.231	22.2	139.4
EIS-E [52]	13.3	25.7	-	-	-	-	80.6	68.3	9.958	2.645	0.529	1.222	-	-
DTC-SPADE [94]	13.5	17.1	5.2	7.9	0.46	0.60	93.0	89.7	9.277	2.416	0.527	1.291	11.6	237.4
Ours	13.0	15.0	5.0	5.8	0.46	0.49	92.9	92.6	8.662	2.259	0.493	1.172	71.6	57.4

Stereo Matching Loss. There are two pseudo predictions that are only inferred in training time using the current event, and a temporally refined final prediction. Smooth L1 loss is applied for three predictions.

$$\mathcal{L}_{stereo} = \sum_{i \in \{0,1\}} \lambda_i \cdot \text{Smooth}_{L_1}(D_t^i, D_{t,gt}) + \lambda_f \cdot \text{Smooth}_{L_1}(D_t, D_{t,gt}) \quad (9)$$

In summary, the total objective function \mathcal{L} is a summation of flow loss (\mathcal{L}_{flow}) and stereo loss (\mathcal{L}_{stereo}) for joint learning with stereoscopic flow as:

$$\mathcal{L} = \mathcal{L}_{flow} + \mathcal{L}_{stereo} \quad (10)$$

4 Experiments

4.1 Datasets

Our model is verified in MVSEC dataset [100] following the previous works. The dataset consists of event streams from calibrated stereo event cameras, which have a resolution of 346×260 , and ground truth depth. For a fair comparison, we adhered to the protocol [1, 74, 94] by conducting training and evaluation on sequences 1 and 3 of the Indoor Flying scene. For the metric, mean depth error, median depth error, mean disparity error, and one-pixel-accuracy (1PA) are used. Computational complexity is measured in terms of frame rate and FLOPs.

Furthermore, to validate the generalizability of the model, experiments are conducted on DSEC dataset [25]. It consists of 53 different driving scenes with stereo event cameras of 640×480 resolution. One-pixel-error (1PE), two-pixel-error, mean absolute error (MAE), and root mean square error (RMSE) are measured following the previous works and the benchmark website.




Fig. 4: Qualitative results on the Indoor Flying dataset of MVSEC. From top to bottom, the rows display #245 from sequence 1, #591 from sequence 1, #200 from sequence 3, and #1585 from sequence 3, respectively. The author reproduced qualitative results with publicly available code [74, 94].

4.2 Implementation Details

The models used in MVSEC and DSEC have identical structures, but we set the maximum disparity to 48 and 192, respectively. Additionally, to accommodate the disparity range, we increased the feature encoder channels. More details on implementation are given in the supplementary material.

4.3 Quantitative Results

Table 1 shows the dense disparity estimation results on both MVSEC and DSEC datasets. Focusing on precise depth estimation performance only using asynchronous and sparse events, the table compares event-only stereo methods tested either on MVSEC or DSEC. The results of other methods on MVSEC are taken from their original paper, and the results of DSEC are derived from the DSEC benchmark. The DTC-SPADE [94] achieved remarkable performance outperforming EITNet [1] which uses images in training, by introducing RNN structure for feature embedding. However, as a price of temporal encoding, computational cost increased significantly. Our main model outperforms previous SOTA except 1PA, while minimizing computation cost. A lightweight stereoscopic flow network and warping process enable our model to propagate stereoscopic knowledge across time with minimal burden. Since enough information is passed from the previous time step, a small feature encoder is sufficient to make a precise prediction. For the results on DSEC dataset, our results are uploaded and evaluated




Fig. 5: Qualitative results on the DSEC dataset. We use a pre-trained model provided by the author for Se-CFF [53] and train DTC-SPADE [94] using public code.

on the DSEC benchmark website¹. Our method achieves the best performance across all metrics and outperforms other methods by a significant margin.

4.4 Qualitative Results

Qualitative results of the proposed model and the previous event stereos are shown in Fig. 4 and Fig. 5. DDES [74] and Se-CFF [53], single pair event-based stereo matchings, often fail to reconstruct scenes due to a lack of sufficient event information. In contrast, DTC-SPADE [94] shows better qualitative results and stable depth reconstruction across the entire scene thanks to temporally aggregated features, yet it still struggles with edges and boundaries. Conversely, our method can compensate for the lack of feature information and 3D cues from past cost volumes, demonstrating remarkably sharp qualitative results for objects.

4.5 Ablation Study and Analysis

Ablation Study on Temporal Modules. A performance gain of each module and the effect of temporal aggregation are shown in Table 2. The baseline is a single-frame implementation of our model, which does not use past knowledge. Adding a feature warping module improves performance to a sub-optimal level by enriching the representation of each scene. Also, adding the cost volume warping module to the baseline improved all metrics by a large margin, by providing dense and abstract knowledge related to stereo matching to the next time step. Finally, the best performance is achieved when both scene-specific feature information and stereo-centric costs are utilized.

Reliability Filter. There can be many different methods to fuse two separate cost volumes from different time sources. As a cost volume represents a matching score for each disparity candidate, naive approaches like addition or channel-wise concatenation do not exploit the benefit of cost volume warping. Therefore, to achieve optimal performance, we applied cost volume fusion using

¹ <https://dsec.ifi.uzh.ch/uzh/disparity-benchmark/>

Table 2: Ablation studies of the proposed modules.

Ablation Settings	Mean depth error [cm] ↓	Mean disp error [pix] ↓	1PA [%] ↑
Single	15.2	0.53	91.4
+ Feature Warping	14.7	0.50	91.9
+ Cost Volume Warping	<u>13.6</u>	<u>0.49</u>	<u>92.3</u>
Full Model	13.0	0.46	92.9

Table 3: The ablation study of reliability filter.

Settings	Mean Depth ↓	Mean Disparity ↓	1PA ↑
Add	14.0	0.48	92.6
Concat	<u>13.5</u>	<u>0.47</u>	<u>92.7</u>
Cost Filter	<u>13.5</u>	<u>0.47</u>	92.6
Ours	13.0	0.46	92.9

Table 4: The analysis of the stereoscopic flow.

Settings	Epipolar Constraint	Mean Depth ↓	Mean Disp ↓	1PA ↑	FPS ↑
EV-Flow [101]	Soft	14.1	0.51	91.8	38.7
EV-Flow [101]	Hard (\mathcal{SF})	13.7	0.49	<u>92.2</u>	51.0
Ours	Soft	<u>13.3</u>	<u>0.47</u>	92.9	<u>69.7</u>
Ours	Hard (\mathcal{SF})	13.0	0.46	92.9	71.6

Table 5: The ablation study of flow loss (Eq. 8).

Ablation Settings	Mean Depth ↓	Mean Disp ↓	1PA ↑
\mathcal{L}_c	13.6	<u>0.47</u>	<u>92.4</u>
\mathcal{L}_{tdc}	<u>13.1</u>	<u>0.47</u>	92.9
$\mathcal{L}_c + \mathcal{L}_{tdc}$	13.0	0.46	92.9

various approaches, as shown in Table 3. The addition method averages two cost volumes, while concatenation uses a convolution layer to reduce the channel. The cost filter method employs a pixel-wise weight estimation CNN block that takes concatenated cost volumes as input. Among various fusion methods, the entropy filter achieves the best performance, so we adopt this approach.

Stereoscopic Flow. We conducted ablations in two ways, regarding the effect and efficiency of the proposed stereoscopic flow network. First, to validate the effect of stereo feature sharing, we fed event voxels directly to a large flow network, EV-Flow [103], capable of high-quality flow estimation. Second, we made a twin optical flow network that separately estimates left and right optical flows. Smooth L1 loss is applied to constrain vertical flows, and it is denoted as soft epipolar constraint in Table 4. Even if the added soft constraints trained two independent flow networks to estimate the same vertical flow, adding a redundant degree of freedom makes training inefficient. Our design, which uses a single stereoscopic flow network sharing stereo feature, outperforms other settings.

Temporal Disparity Consistency Loss. Table 5 shows the ablation study of the proposed temporal disparity consistency loss. As explored in previous studies [24], event-based optical flow can be learned with just contrast maximization \mathcal{L}_c , but joint training with other tasks can be unstable [68] and potentially conflict with those tasks. In contrast, our \mathcal{L}_{tdc} is not a fully self-supervised learning approach and is optimized by cues from disparity, allowing for stable training. \mathcal{L}_{tdc} alone can achieve satisfactory performance, but the best performance is attained when trained together with contrast.

4.6 Performance Gain from Temporal Consistency

To further validate the effectiveness of temporal aggregation, we provide Fig. 6, which shows the sequential qualitative results to deliver intuition of the impact of




Fig. 6: Qualitative results for temporal consistency on split 1 of MVSEC dataset.

knowing the past. Due to the small motion, event data lacks information on the overall scene, *e.g.*, cylinder object. The single frame estimation methods, DDES and our single baseline, fail to estimate disparity in some time steps that lack event streams. DTC-SPADE estimates rough disparity for challenging scenes but lacks edge details due to the limited temporal aggregation. Our model estimates consistent disparities using past stereo information, which is the behavior to improve stereo performance reflecting the real world.

5 Conclusion

We propose a new paradigm of a temporal event stereo network that aggregates past features and cost volume to enhance stereo representations using a stereoscopic flow, which is a stereo-targeted novel flow representation. Previous features and cost volume are aligned to the current time step, enriching the stereo representations. Especially, our model warps the entire cost volume, a stereo-specific representation, across the time frame achieving significant performance improvement with minimal computational cost. We also propose a TDC loss and use encoder-sharing architecture to train stereoscopic flow and stereo jointly without using flow GT. From this effort, we achieve state-of-the-art performance in both MVSEC and DSEC datasets with minimal computation.

Limitations and Future Work. Our framework demonstrates training stereo flow without ground truth, but self-supervised stereo matching using events alone remains unexplored. Our approach could inspire future research in this area.

Acknowledgements. This work was supported by the Technology Innovation Program (1415187329,20024355, Development of autonomous driving connectivity technology based on sensor-infrastructure cooperation) funded by the Ministry of Trade, Industry & Energy(MOTIE, Korea) and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (NRF2022R1A2B5B03002636).

References

1. Ahmed, S.H., Jang, H.W., Uddin, S.N., Jung, Y.J.: Deep event stereo leveraged by event-to-image translation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 882–890 (2021)
2. Almatrafi, M., Baldwin, R., Aizawa, K., Hirakawa, K.: Distance surface for event-based optical flow. *IEEE transactions on pattern analysis and machine intelligence* **42**(7), 1547–1556 (2020)
3. Alonso, I., Murillo, A.C.: Ev-segnet: Semantic segmentation for event-based cameras. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
4. 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)
5. Benosman, R., Clercq, C., Lagorce, X., Ieng, S.H., Bartolozzi, C.: Event-based visual flow. *IEEE transactions on neural networks and learning systems* **25**(2), 407–417 (2013)
6. Brandli, C., Berner, R., Yang, M., Liu, S.C., Delbruck, T.: A 240×180 130 db 3 μ s latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits* **49**, 2333–2341 (2014)
7. Camunas-Mesa, L.A., Serrano-Gotarredona, T., Ieng, S.H., Benosman, R.B., Linares-Barranco, B.: On the use of orientation filters for 3d reconstruction in event-driven stereo vision. *Frontiers in neuroscience* **8**, 48 (2014)
8. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 5410–5418 (2018)
9. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5410–5418 (2018)
10. Cheng, S., Xu, Z., Zhu, S., Li, Z., Li, L.E., Ramamoorthi, R., Su, H.: Deep stereo using adaptive thin volume representation with uncertainty awareness. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2524–2534 (2020)
11. Cheng, X., Zhong, Y., Harandi, M.T., Dai, Y., Chang, X., Drummond, T., Li, H., Ge, Z.: Hierarchical neural architecture search for deep stereo matching. *ArXiv abs/2010.13501* (2020)
12. Cho, H., Cho, J., Yoon, K.J.: Learning adaptive dense event stereo from the image domain. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17797–17807 (2023)
13. Cho, H., Jeong, J., Yoon, K.J.: Eomvs: Event-based omnidirectional multi-view stereo. *IEEE Robotics and Automation Letters* **6**(4), 6709–6716 (2021). <https://doi.org/10.1109/LRA.2021.3096161>

14. Cho, H., Jeong, Y., Kim, T., Yoon, K.J.: Non-coaxial event-guided motion deblurring with spatial alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12492–12503 (2023)
15. Cho, H., Kim, H., Chae, Y., Yoon, K.J.: Label-free event-based object recognition via joint learning with image reconstruction from events. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19866–19877 (2023)
16. Cho, H., Kim, T., Jeong, Y., Yoon, K.J.: Tta-evf: Test-time adaptation for event-based video frame interpolation via reliable pixel and sample estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25701–25711 (2024)
17. Cho, H., Yoon, K.J.: Event-image fusion stereo using cross-modality feature propagation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 454–462 (2022)
18. Cho, H., Yoon, K.J.: Selection and cross similarity for event-image deep stereo. In: European Conference on Computer Vision. pp. 470–486. Springer (2022)
19. Choi, J., Yoon, K.J., et al.: Learning to super resolve intensity images from events. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2768–2776 (2020)
20. Ding, Z., Zhao, R., Zhang, J., Gao, T., Xiong, R., Yu, Z., Huang, T.: Spatio-temporal recurrent networks for event-based optical flow estimation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 36, pp. 525–533 (2022)
21. Fox, G., Pan, X., Tewari, A., Elgharib, M., Theobalt, C.: Unsupervised event-based video reconstruction. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4179–4188 (2024)
22. Funahashi, K.i., Nakamura, Y.: Approximation of dynamical systems by continuous time recurrent neural networks. *Neural networks* **6**(6), 801–806 (1993)
23. Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A.J., Conradt, J., Daniilidis, K., et al.: Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence* **44**(1), 154–180 (2020)
24. Gallego, G., Rebecq, H., Scaramuzza, D.: A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3867–3876 (2018)
25. Gehrig, M., Aarents, W., Gehrig, D., Scaramuzza, D.: Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters* **6**(3), 4947–4954 (2021)
26. Gehrig, M., Millhäusler, M., Gehrig, D., Scaramuzza, D.: E-raft: Dense optical flow from event cameras. In: 2021 International Conference on 3D Vision (3DV). pp. 197–206. IEEE (2021)
27. Gehrig, M., Muglikar, M., Scaramuzza, D.: Dense continuous-time optical flow from event cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024)
28. Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
29. Guo, X., Yang, K., Yang, W., Wang, X., Li, H.: Group-wise correlation stereo network. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 3268–3277 (2019)

30. Han, J., Zhou, C., Duan, P., Tang, Y., Xu, C., Xu, C., Huang, T., Shi, B.: Neuromorphic camera guided high dynamic range imaging. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1730–1739 (2020)
31. Hu, L., Zhao, R., Ding, Z., Ma, L., Shi, B., Xiong, R., Huang, T.: Optical flow estimation for spiking camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17844–17853 (2022)
32. Hur, J., Roth, S.: Iterative residual refinement for joint optical flow and occlusion estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5754–5763 (2019)
33. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2462–2470 (2017)
34. Jiang, S., Campbell, D., Lu, Y., Li, H., Hartley, R.: Learning to estimate hidden motions with global motion aggregation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9772–9781 (2021)
35. Jiang, Z., Zhang, Y., Zou, D., Ren, J., Lv, J., Liu, Y.: Learning event-based motion deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3320–3329 (2020)
36. Kendall, A., Martirosyan, H., Dasgupta, S., Henry, P.: End-to-end learning of geometry and context for deep stereo regression. 2017 IEEE International Conference on Computer Vision (ICCV) pp. 66–75 (2017)
37. Kim, T., Chae, Y., Jang, H.K., Yoon, K.J.: Event-based video frame interpolation with cross-modal asymmetric bidirectional motion fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18032–18042 (June 2023)
38. Kim, T., Cho, H., Yoon, K.J.: Frequency-aware event-based video deblurring for real-world motion blur. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24966–24976 (2024)
39. Kogler, J., Humenberger, M., Sulzbachner, C.: Event-based stereo matching approaches for frameless address event stereo data. In: International Symposium on Visual Computing. pp. 674–685. Springer (2011)
40. Laga, H., Jospin, L.V., Boussaïd, F., Bennamoun: A survey on deep learning techniques for stereo-based depth estimation. *IEEE transactions on pattern analysis and machine intelligence* **PP** (2020)
41. Lagorce, X., Orchard, G., Galluppi, F., Shi, B.E., Benosman, R.B.: Hots: a hierarchy of event-based time-surfaces for pattern recognition. *IEEE transactions on pattern analysis and machine intelligence* **39**(7), 1346–1359 (2016)
42. Li, J., Wang, P., Xiong, P., Cai, T., Yan, Z., Yang, L., Liu, J., Fan, H., Liu, S.: Practical stereo matching via cascaded recurrent network with adaptive correlation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16263–16272 (2022)
43. Lichtsteiner, P., Posch, C., Delbrück, T.: A 128×128 120 db 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits* **43**, 566–576 (2008)
44. Lin, S., Zhang, J., Pan, J., Jiang, Z., Zou, D., Wang, Y., Chen, J., Ren, J.S.: Learning event-driven video deblurring and interpolation. In: ECCV (8). pp. 695–710 (2020)
45. Liu, H., Chen, G., Qu, S., Zhang, Y., Li, Z., Knoll, A., Jiang, C.: Tma: Temporal motion aggregation for event-based optical flow. In: ICCV (2023)

46. Liu, M., Delbruck, T.: Adaptive time-slice block-matching optical flow algorithm for dynamic vision sensors. *BMVC* (2018)
47. Luo, X., Luo, K., Luo, A., Wang, Z., Tan, P., Liu, S.: Learning optical flow from event camera with rendered dataset. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 9847–9857 (October 2023)
48. Manderscheid, J., Sironi, A., Bourdis, N., Migliore, D., Lepetit, V.: Speed invariant time surface for learning to detect corner points with event-based cameras. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10245–10254 (2019)
49. Maqueda, A.I., Loquercio, A., Gallego, G., García, N., Scaramuzza, D.: 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*. pp. 5419–5427 (2018)
50. Mostafavi, M., Nam, Y., Choi, J., Yoon, K.J.: E2sri: Learning to super-resolve intensity images from events. *IEEE transactions on pattern analysis and machine intelligence* **44**(10), 6890–6909 (2021)
51. Mostafavi, M., Wang, L., Yoon, K.J.: Learning to reconstruct hdr images from events, with applications to depth and flow prediction. *International Journal of Computer Vision* **129**, 900–920 (2021)
52. Mostafavi, M., Yoon, K.J., Choi, J.: Event-intensity stereo: Estimating depth by the best of both worlds. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4258–4267 (2021)
53. Nam, Y., Mostafavi, M., Yoon, K.J., Choi, J.: Stereo depth from events cameras: Concentrate and focus on the future. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022)
54. Nguyen, A., Do, T.T., Caldwell, D.G., Tsagarakis, N.G.: Real-time 6dof pose relocalization for event cameras with stacked spatial lstm networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. pp. 0–0 (2019)
55. Pan, L., Liu, M., Hartley, R.: Single image optical flow estimation with an event camera. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1669–1678. IEEE (2020)
56. Pan, L., Scheerlinck, C., Yu, X., Hartley, R., Liu, M., Dai, Y.: Bringing a blurry frame alive at high frame-rate with an event camera. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6820–6829 (2019)
57. Paredes-Vallés, F., de Croon, G.C.: Back to event basics: Self-supervised learning of image reconstruction for event cameras via photometric constancy. *arXiv preprint arXiv:2009.08283* (2020)
58. Paredes-Vallés, F., Scheper, K.Y., De Croon, G.C.: Unsupervised learning of a hierarchical spiking neural network for optical flow estimation: From events to global motion perception. *IEEE transactions on pattern analysis and machine intelligence* **42**(8), 2051–2064 (2019)
59. Paredes-Vallés, F., Scheper, K.Y., De Wagter, C., de Croon, G.C.: Taming contrast maximization for learning sequential, low-latency, event-based optical flow. *arXiv preprint arXiv:2303.05214* (2023)
60. Park, M.G., Yoon, K.J.: Learning and selecting confidence measures for robust stereo matching. *IEEE transactions on pattern analysis and machine intelligence* **41**(6), 1397–1411 (2018)

61. Perot, E., De Tournemire, P., Nitti, D., Masci, J., Sironi, A.: Learning to detect objects with a 1 megapixel event camera. *Advances in Neural Information Processing Systems* **33**, 16639–16652 (2020)
62. Piatkowska, E., Belbachir, A., Gelautz, M.: Asynchronous stereo vision for event-driven dynamic stereo sensor using an adaptive cooperative approach. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. pp. 45–50 (2013)
63. Piatkowska, E., Kogler, J., Belbachir, N., Gelautz, M.: Improved cooperative stereo matching for dynamic vision sensors with ground truth evaluation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. pp. 53–60 (2017)
64. 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* (2019)
65. Rogister, P., Benosman, R., Ieng, S.H., Lichtsteiner, P., Delbruck, T.: Asynchronous event-based binocular stereo matching. *IEEE Transactions on Neural Networks and Learning Systems* **23**(2), 347–353 (2011)
66. Scharstein, D., Szeliski, R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* **47**, 7–42 (2004)
67. Shiba, S., Aoki, Y., Gallego, G.: Event collapse in contrast maximization frameworks. *Sensors (Basel, Switzerland)* **22** (2022), <https://api.semanticscholar.org/CorpusID:250408092>
68. Shiba, S., Aoki, Y., Gallego, G.: Secrets of event-based optical flow. In: *European Conference on Computer Vision*. pp. 628–645. Springer (2022)
69. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 8934–8943 (2018)
70. Sun, L., Sakaridis, C., Liang, J., Jiang, Q., Yang, K., Sun, P., Ye, Y., Wang, K., Van Gool, L.: Event-based fusion for motion deblurring with cross-modal attention. In: *European Conference on Computer Vision (ECCV)* (2022)
71. Sun, L., Sakaridis, C., Liang, J., Sun, P., Cao, J., Zhang, K., Jiang, Q., Wang, K., Van Gool, L.: Event-based frame interpolation with ad-hoc deblurring. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18043–18052 (2023)
72. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II* 16. pp. 402–419. Springer (2020)
73. Tulyakov, S., Boicchio, A., Gehrig, D., Georgoulis, S., Li, Y., Scaramuzza, D.: Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17755–17764 (2022)
74. Tulyakov, S., Fleuret, F., Kiefel, M., Gehler, P., Hirsch, M.: Learning an event sequence embedding for dense event-based deep stereo. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1527–1537 (2019)
75. Tulyakov, S., Gehrig, D., Georgoulis, S., Erbach, J., Gehrig, M., Li, Y., Scaramuzza, D.: Time lens: Event-based video frame interpolation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16155–16164 (2021)
76. Wan, Z., Dai, Y., Mao, Y.: Learning dense and continuous optical flow from an event camera. *IEEE Transactions on Image Processing* **31**, 7237–7251 (2022)

77. Wan, Z., Mao, Y., Zhang, J., Dai, Y.: Rpeflow: Multimodal fusion of rgb-pointcloud-event for joint optical flow and scene flow estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10030–10040 (October 2023)
78. Wang, L., Chae, Y., Yoon, K.J.: Dual transfer learning for event-based end-task prediction via pluggable event to image translation. In: ICCV (2021)
79. Wang, L., Ho, Y.S., Yoon, K.J., et al.: Event-based high dynamic range image and very high frame rate video generation using conditional generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10081–10090 (2019)
80. Weinzaepfel, P., Lucas, T., Leroy, V., Cabon, Y., Arora, V., Brégier, R., Csurka, G., Antsfeld, L., Chidlovskii, B., Revaud, J.: Croco v2: Improved cross-view completion pre-training for stereo matching and optical flow. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17969–17980 (2023)
81. Weinzaepfel, P., Revaud, J., Harchaoui, Z., Schmid, C.: Deepflow: Large displacement optical flow with deep matching. In: Proceedings of the IEEE international conference on computer vision. pp. 1385–1392 (2013)
82. Won, C., Ryu, J., Lim, J.: End-to-end learning for omnidirectional stereo matching with uncertainty prior. *IEEE transactions on pattern analysis and machine intelligence* **43**(11), 3850–3862 (2020)
83. Xia, R., Zhao, C., Zheng, M., Wu, Z., Sun, Q., Tang, Y.: Cmda: Cross-modality domain adaptation for nighttime semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21572–21581 (2023)
84. Xu, F., Yu, L., Wang, B., Yang, W., Xia, G.S., Jia, X., Qiao, Z., Liu, J.: Motion deblurring with real events. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2583–2592 (October 2021)
85. Xu, F., Yu, L., Wang, B., Yang, W., Xia, G.S., Jia, X., Qiao, Z., Liu, J.: Motion deblurring with real events. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2583–2592 (2021)
86. Xu, G., Cheng, J., Guo, P., Yang, X.: Attention concatenation volume for accurate and efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12981–12990 (2022)
87. Xu, G., Wang, X., Ding, X., Yang, X.: Iterative geometry encoding volume for stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21919–21928 (2023)
88. Xu, H., Zhang, J.: Aanet: Adaptive aggregation network for efficient stereo matching. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 1956–1965 (2020)
89. Xu, H., Zhang, J.: Aanet: Adaptive aggregation network for efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1959–1968 (2020)
90. Ye, C., Mitrokhin, A., Fermüller, C., Yorke, J.A., Aloimonos, Y.: Unsupervised learning of dense optical flow, depth and egomotion from sparse event data. *arXiv preprint arXiv:1809.08625* (2018)
91. Zhang, F., Prisacariu, V., Yang, R., Torr, P.H.: Ga-net: Guided aggregation net for end-to-end stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 185–194 (2019)

92. Zhang, F., Prisacariu, V.A., Yang, R., Torr, P.H.S.: Ga-net: Guided aggregation net for end-to-end stereo matching. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 185–194 (2019)
93. Zhang, J., Li, S., Luo, Z., Fang, T., Yao, Y.: Vis-mvsnet: Visibility-aware multi-view stereo network. *International Journal of Computer Vision* **131**(1), 199–214 (2023)
94. Zhang, K., Che, K., Zhang, J., Cheng, J., Zhang, Z., Guo, Q., Leng, L.: Discrete time convolution for fast event-based stereo. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8676–8686 (2022)
95. Zhang, L., Zhang, H., Zhu, C., Guo, S., Chen, J., Wang, L.: Fine-grained video deblurring with event camera. In: *ICMM*. pp. 352–364. Springer (2021)
96. Zhang, S., Zhang, Y., Jiang, Z., Zou, D., Ren, J., Zhou, B.: Learning to see in the dark with events. In: *European Conference on Computer Vision*. pp. 666–682. Springer (2020)
97. Zhang, X., Yu, L.: Unifying motion deblurring and frame interpolation with events. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17765–17774 (2022)
98. Zhang, Y., Poggi, M., Mattoccia, S.: Temporalstereo: Efficient spatial-temporal stereo matching network. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 9528–9535. IEEE (2023)
99. Zhu, A.Z., Chen, Y., Daniilidis, K.: Realtime time synchronized event-based stereo. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 433–447 (2018)
100. 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>
101. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Ev-flownet: Self-supervised optical flow estimation for event-based cameras. *arXiv preprint arXiv:1802.06898* (2018)
102. Zhu, A.Z., Yuan, L., Chaney, K., Daniilidis, K.: Unsupervised event-based learning of optical flow, depth, and egomotion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 989–997 (2019)
103. 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)
104. Zou, D., Guo, P., Wang, Q., Wang, X., Shao, G., Shi, F., Li, J., Park, P.K.: Context-aware event-driven stereo matching. In: *2016 IEEE International Conference on Image Processing (ICIP)*. pp. 1076–1080. IEEE (2016)
105. Zou, D., Shi, F., Liu, W., Li, J., Wang, Q., Park, P.K., Shi, C.W., Roh, Y.J., Ryu, H.E.: Robust dense depth map estimation from sparse dvs stereos. In: *British Mach. Vis. Conf.(BMVC)*. vol. 1 (2017)