


# EMoTive: Event-guided Trajectory Modeling for 3D Motion Estimation

Zengyu Wan, Wei Zhai , Yang Cao, Zhengjun Zha  
University of Science and Technology of China

wanzengyu@mail.ustc.edu.cn, {wzhai056@, forrest@, zhazj@}ustc.edu.cn

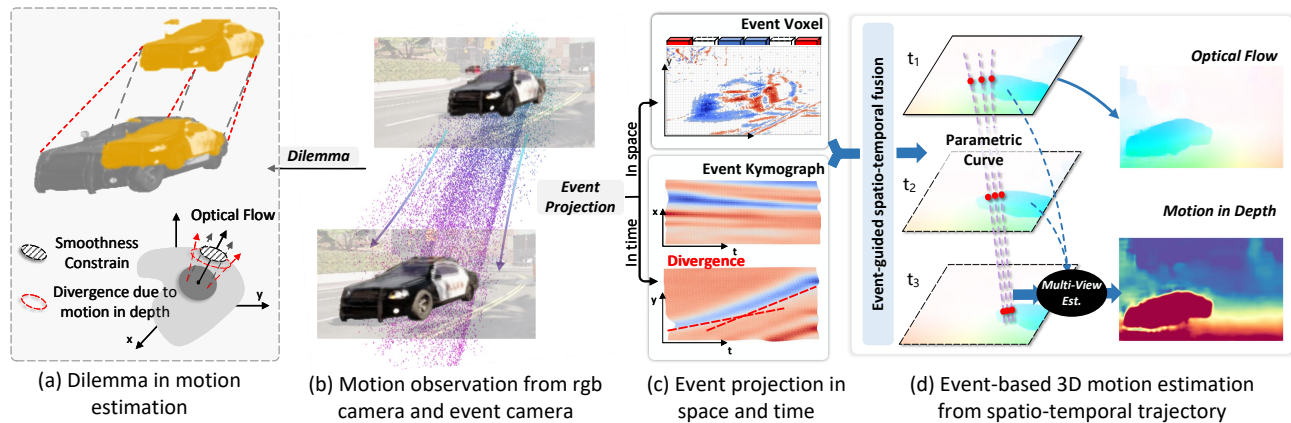


Figure 1. (a) The dilemma faced by motion estimation algorithms between the local smoothness assumption and the real divergence due to motion in depth; (b) (c) The spatio-temporal projection of events (Event Kymograph) provides an clear observation of temporal evolution, enabling heterogeneous spatio-temporal motion analysis; (d) 3D motion can be inferred by modelling spatio-temporal trajectories via event-guided non-uniform parametric curves.

## Abstract


Visual 3D motion estimation aims to infer the motion of 2D pixels in 3D space based on visual cues. The key challenge arises from depth variation induced spatio-temporal motion inconsistencies, disrupting the assumptions of local spatial or temporal motion smoothness in previous motion estimation frameworks. In contrast, event cameras offer new possibilities for 3D motion estimation through continuous adaptive pixel-level responses to scene changes. This paper presents **EMoTive**, a novel event-based framework that models spatio-temporal trajectories via event-guided non-uniform parametric curves, effectively characterizing locally heterogeneous spatio-temporal motion. Specifically, we first introduce **Event Kymograph** - an event projection method that leverages a continuous temporal projection kernel and decouples spatial observations to encode fine-grained temporal evolution explicitly. For motion representation, we introduce a density-aware adaptation mechanism to fuse spatial and temporal features under event guidance,

coupled with a non-uniform rational curve parameterization framework to adaptively model heterogeneous trajectories. The final 3D motion estimation is achieved through multi-temporal sampling of parametric trajectories, yielding optical flow and depth motion fields. To facilitate evaluation, we introduce **CarlaEvent3D**, a multi-dynamic synthetic dataset for comprehensive validation. Extensive experiments on both this dataset and a real-world benchmark demonstrate the effectiveness of the proposed method.

## 1. Introduction

Visual 3D motion estimation infers the movement of 2D pixels in 3D space, playing a crucial role in spatial intelligence for autonomous systems and enabling diverse applications in dynamic scene understanding [2, 5, 21, 37]. The key challenge of 3D motion estimation lies in maintaining motion representation consistency throughout the spatial deformation process induced by depth variation.

Existing approaches employ frame-based cameras or LiDAR for 3D motion perception; their reliance on fixed temporal synchronization limits performance under complex

: Corresponding author.

motion patterns and depth variations. For example, most optical flow methods [9, 32, 36, 39, 52] estimate planar motion through inter-frame correspondences of pixels/point clouds, but their insufficient temporal observation necessitates restrictive local spatial motion consistency priors or long temporal consistency to regularize the solution space. Recent efforts [23, 26, 28] attempt to decouple planar and depth motion estimation in dual spaces (feature space and scale space), yet the inherent contradiction persists as both spaces originate from the same pixel domain. The root cause stems from heterogeneous spatio-temporal motion: object projections in image space undergo local deformations due to motion in depth, invalidating local temporal invariance assumptions.

Event cameras - bio-inspired vision sensors that asynchronously report pixel-level brightness changes through event streams [12, 50] - offer microsecond temporal resolution, high dynamic range, and minimal motion blur, making them promising for accurate 3D motion estimation. Specifically, event streams can form Event Kymograph through x-t and y-t decoupled projections to capture temporal evolution in microseconds, as shown in Fig. 1. This fine observation enables non-stationary modeling for 3D motion estimation while avoiding the motion dilemma.

Based on the spatio-temporal clues from event cameras, we propose EMoTive, an event-guided trajectory parameterization framework, to achieve 3D motion estimation. This framework models motion trajectories through non-uniform parametric curves whose deformation characteristics encode 3D motion parameters. Specifically, it first constructs spatial representations (Event Voxel) through x-y plane projections and temporal representations (Event Kymograph) through x-t and y-t decoupled plane projections. Then, dual cost volumes are constructed on both spatial and temporal representations to guide iterative trajectory refinement. Meanwhile, a density-aware adaptation mechanism is proposed to update the curve parameters by non-uniform sampling of the local event density, which is correlated with the motion strength. The resulting trajectories parameterize pixel motion over time, from which we derive optical flow at arbitrary timestamps and compute depth motion fields through temporal analysis. To enhance depth estimation robustness, we further propose a multi-view fusion algorithm that aggregates observations across adjacent timestamps. To better simulate real-world motion scenarios, we construct the CarlaEvent3D dataset using the CARLA autonomous driving simulator, which includes diverse environments (night, rain, fog, *etc.*) and realistic motions for training and evaluation.

The contributions can be summarized as follows:

1. For event-based 3D motion estimation, the **EMoTive** model is introduced via event-guided non-uniform parametric curve modeling to characterize heterogeneous

spatio-temporal motion patterns.

2. To encode fine-grained temporal evolution from events, the **Event Kymograph** is proposed by leveraging a continuous temporal projection kernel from decoupled spatial observations.
3. During the progress of trajectory formation, a **Density-aware adaptation** mechanism is designed to guide non-uniform spatio-temporal fusion and curve parameter updating depending on the motion sensitivity of events.
4. For comprehensive evaluation, the **CarlaEvent3D** dataset is introduced, which simulates real-world motion under multiple complex scenarios. Extensive experiments on both synthetic and real-world benchmarks reveal the competitive performance of the EMoTive with superior efficiency.

## 2. Related Work

### 2.1. 3D Motion Estimation

Scene flow was first introduced by Vedula *et al.* [45] to estimate the 3D motion of objects in dense scenes. Various methods have been proposed depending on the available data modalities, such as monocular/stereo images, LiDAR, or RGB-D data. Stereo-based methods [6, 31, 49] employ a two-stage pipeline: (1) disparity estimation for 3D reconstruction followed by (2) temporal correspondence matching for scene flow computation. LiDAR-based techniques [17, 22, 30, 33] leverage point cloud registration across frames using geometric matching algorithms. However, these approaches exhibit modality-specific limitations [47, 53]: Stereo systems require precise extrinsic calibration, while LiDAR suffers from sparse sampling artifacts due to non-uniform beam distribution and surface reflectivity dependencies.

Recent advances in monocular 3D motion estimation have demonstrated promising results through hybrid architectures. Early variational methods [4, 10] jointly optimized motion and depth under restrictive priors [16, 46], which limited their generalization capability. A paradigm shift occurred with Yang and Ramanan’s normalized scene flow formulation [51], enabling depth-agnostic 3D motion estimation through optical expansion analysis. Subsequent works by Han *et al.* [26–28] introduced cross-scale feature fusion mechanisms, unifying optical flow estimation and depth motion prediction in an end-to-end framework. Nevertheless, these image-based methods remain susceptible to spatial deformation artifacts.

### 2.2. Event-Based Motion Estimation

Event cameras offer unique advantages for motion estimation through low latency ( $\mu\text{s}$ -level) and high dynamic range ( $> 120\text{dB}$ ). Early solutions on event-based motion estimation [7, 19, 35] employed physical constraints like in-

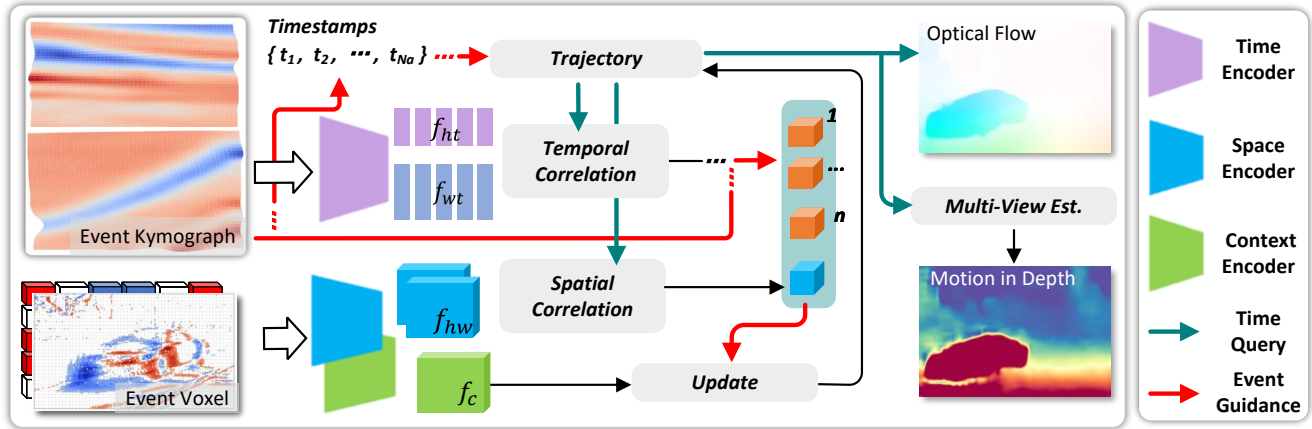


Figure 2. Overall pipeline of EMoTive with event-guided trajectory formation scheme. The Event Voxel and Event Kymograph will be first utilized to construct spatio-temporal dual correlations for motion representation. A density-aware adaptation mechanism is then implemented to fuse spatial and temporal features under event guidance to model trajectories adaptively. The final 3D motion estimation is achieved through multi-temporal sampling of parametric trajectories, yielding optical flow and depth motion field

tensity constancy and local motion smoothness. Plane fitting techniques [1, 3, 34] estimate normal flow under local smooth motion assumptions, while contrast maximization frameworks [11, 18, 20, 38, 42, 48] optimize event motion through spatio-temporal alignment analysis. However, these methods face fundamental limitations in handling complex motion patterns and avoiding local optima. Another line of research pursues learning-based approaches, with architectures like EV-FlowNet [54] and E-RAFT [14] demonstrating robust 2D optical flow estimation. Recent multi-modal approaches [47, 53] integrate events with RGB and radar data, exploiting event-driven temporal cues to enhance 3D motion modeling in high-dynamic-range scenarios. Notably, Li *et al.* [24] developed BlinkVision, a synthetic dataset enabling event-based scene flow learning through novel mesh processing pipelines. Despite progress, the challenge of locally heterogeneous spatio-temporal motion modeling persists.

### 3. Method

To address the challenge of 3D motion estimation, this paper proposes **EMoTive**, which models heterogeneous spatio-temporal motion process through non-uniform parametric curve representations of spatio-temporal trajectories  $\mathcal{T}$ . Specifically, we introduce **Event Kymograph**, a temporal representation from decoupled space, to leverage temporal correlations from event streams explicitly. The proposed spatio-temporal trajectory  $\mathcal{T}$ , guided by event streams, is formally defined as:  $\mathcal{T}(t, x, y) \mapsto \mathbb{R}^2$ . This continuous representation describes pixel-level motion evolution throughout temporal sequences. Crucially, the trajectory formulation enables the acquisition of pixel motion in image space at arbitrary timestamps, yielding optical flow  $\mathcal{O}$ .

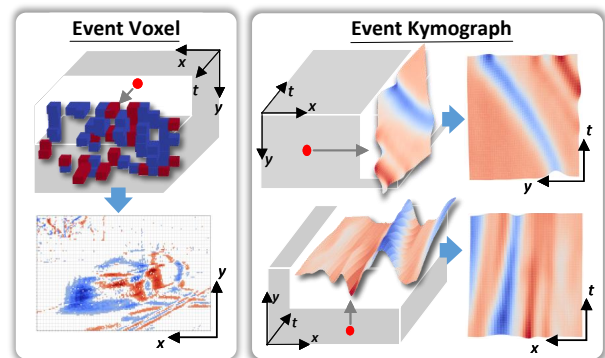


Figure 3. Projection axes of Event Voxel and Event Kymograph. The Event Kymograph captures fine temporal cues in decoupled spatial axes.

Furthermore, by integrating temporal derivatives, we can establish the motion in depth  $\mathcal{M}$ , providing comprehensive 3D motion estimation for dynamic scenes.

#### 3.1. Event Representation

Asynchronous event streams pose significant challenges for direct analysis and computation. A prevalent solution involves spatio-temporal event aggregation through planar projections, noted as Event Voxel [8]. This approach aggregates events through temporal quantized, combined with spatio-temporal triangular kernels, formally expressed as:

$$V(x, y, b) = \sum_i p_i k(x - x_i) k(y - y_i) k(b - t_i^*), \quad (1)$$

where the event streams  $(x_i, y_i, t_i, p_i)_{i \in [1, N]}$  is temporally quantized with bin duration  $T_B$ :  $t_i^* = (t_i - t_1)/T_B$ , and  $k(a)$  denotes the triangular sampling kernel:  $k(a) = \max(0, 1 - |a|)$ . This formulation yields a voxelized tensor  $V : \mathbb{R}^{B \times H \times W}$  through spatio-temporal aggregation.

However, conventional event voxel encoding inevitably discards fine cues due to coarse temporal quantization. To overcome these limitations, we propose Event Kymograph projection – a novel spatio-temporal mapping that preserves temporal fidelity through the  $x-t, y-t$  spacial axis decoupled projection, formulated as:

$$\begin{pmatrix} K_x \\ K_y \end{pmatrix} = \begin{pmatrix} \sum_i p_i k(x - x_i) g(t - t_i | \sigma) \\ \sum_i p_i k(y - y_i) g(t - t_i | \sigma) \end{pmatrix}, \quad (2)$$

where  $\sigma$  controls the temporal smoothing scale of the Gaussian kernel  $g$ :  $g(a|\sigma) = \exp(-(a/\sigma)^2)$ . This projection strategically decouples spatial dimensions while maintaining fine temporal precision (e.g.,  $10\mu s$ ) through continuous Gaussian temporal encoding, thereby enabling high-resolution observation of transient processes.

Our framework collaboratively combines both projection paradigms as spatio-temporal event representation: 1) Event Voxel captures spatially dominant features, and 2) Event Kymograph preserves temporally sensitive patterns.

### 3.2. Spatio-temporal Dual Cost Volumes

**Spatial Correlation.** Building upon established architectures in optical flow estimation [14, 43], EMoTive employs two same 2D convolution networks as space encoder to extract latent spatial representation  $f_{hw} : \mathbb{R}^{D_s \times H_D \times W_D}$  and context encoder to obtain context representation  $f_c$  from the Event Voxel. And given latent features  $f_{hw}^p, f_{hw}^t$  from temporally adjacent voxel projections, we generate an  $N$ -level cost pyramid via multi-resolution inner-product:

$$C_{hw}^m = \frac{1}{D_s} \sum_{c=1}^{D_s} f_{hw}^p(c, i, j) \cdot \mathcal{A}_{2d}^{2m}(f_{hw}^t(c, k, l)), \quad (3)$$

where  $\mathcal{A}_{2d}^{2m}$  denotes 2D average pooling with kernel size  $2^m$  for coarse-to-fine matching,  $m \in \{1, \dots, N\}$  indexes the pyramid level, and we adopt a 2-level cost pyramid.

**Temporal Correlation.** For temporal feature extraction, we apply a 1D convolution encoder operating on the Event Kymograph  $(K_x, K_y)$ . The projection will firstly be evenly divided into subblocks based on the  $N_a$  temporal anchors to form the sequence  $K_{x|y, s} = \{K_{x|y, s}(1), \dots, K_{x|y, s}(N_a)\}$ . The encoder operates on each subblock and outputs two complementary temporal feature tensors:  $f_{ht} : \mathbb{R}^{N_a \times D_t \times H_D}$  and  $f_{wt} : \mathbb{R}^{N_a \times D_t \times W_D}$ , preserving axis-specific motion pattern. To further capture motion dynamics, we formulate a spatial-decoupled temporal cost pyramid through cross-subblock correlation along

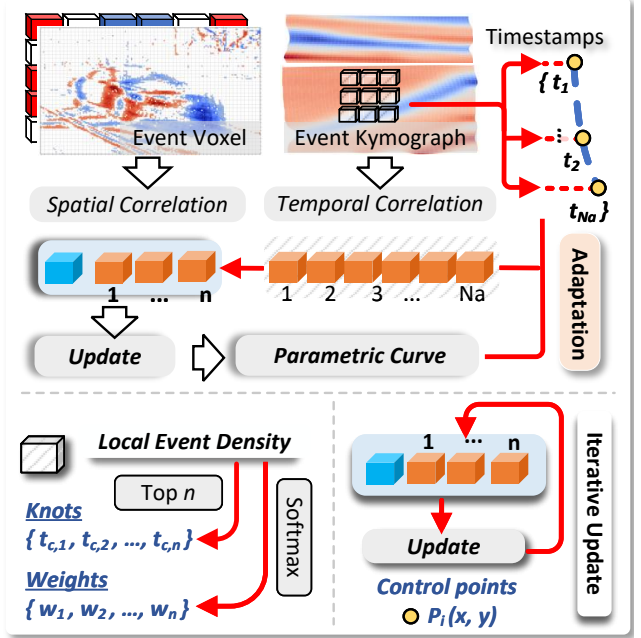


Figure 4. Event-guided trajectory formation process. Local event density is introduced to adaptively adjust the knots and weights while fusing spatio-temporal features for control points update.

orthogonal spatio-temporal projections:

$$\begin{pmatrix} C_{wt}^m \\ C_{ht}^m \end{pmatrix} = \begin{pmatrix} \frac{1}{D_t} \sum_c f_{wt}(1, c, i) \cdot \mathcal{A}_{1d}^{2m}(f_{wt}(n, c, j)) \\ \frac{1}{D_t} \sum_c f_{ht}(1, c, k) \cdot \mathcal{A}_{1d}^{2m}(f_{ht}(n, c, l)) \end{pmatrix}, \quad (4)$$

where,  $\mathcal{A}_{1d}^{2m}$  implements 1D temporal pooling, and  $n \in \{2, \dots, N_a\}$  indexes temporal subblocks.

For subsequent efficient correlation sampling, the temporal correlation on the joint space axis is then synthesized through tensor product fusion:

$$C_t^m(n, i, k, j, l) \doteq C_{ht}^m(n, i, j) \otimes C_{wt}^m(n, k, l), \quad (5)$$

where  $\otimes$  denotes outer product along spacial dimensions.

### 3.3. Spatio-Temporal Trajectory

To model heterogeneous motion patterns from event stream, we propose a non-uniform rational B-spline (NURBS) based spatio-temporal trajectory representation. This parametric formulation extends prior uniform B-spline approaches [15, 44] by incorporating event-driven non-uniformity through adaptive knot vectors and weights, addressing limitations in the motion expressiveness of globally controlled curves. The trajectory map  $\mathcal{T} : \mathbb{R}^3 \mapsto \mathbb{R}^2$  is defined as:

$$\mathcal{T}(t, x, y) = \frac{\sum_i^n N_{i,p}(t) w_i \mathbf{P}_i(x, y)}{\sum_i^n N_{i,p}(t) w_i}, \quad (6)$$

where  $N_{i,p}(t)$  is the  $p$ -th degree B-spline basis function with non-uniform knot vector  $\mathbf{T} = \{t_1, \dots, t_m\}$ ,  $w_i \in \mathbb{R}^+$  represents the event-adaptive weight and  $\mathbf{P}_i \in \mathbb{R}^2$  is the  $i$ -th control point from the set  $\mathcal{P} = \{\mathbf{P}_1, \dots, \mathbf{P}_n\}$ .

### 3.3.1. Density-aware Adaptation

The expressiveness of NURBS curves is primarily governed by knot vector distribution and control point weighting. Higher knot density, coupled with elevated weights, enhances the curve’s capacity to capture rapid motion patterns. Building on the motion sensitivity of event [12], we formulate an event-density guided adaptation mechanism. The adaptation pipeline comprises three phases:

(1) Spatio-temporal density estimation: Compute the distribution tensor  $E_s \in \mathbb{R}^{T \times H \times W}$  from Event Kymograph:  $E_s(i_k, x, y) = K_{s,x}(i_k, x) \odot K_{s,y}(i_k, y)$ , where  $i_k \in [1, N_a]$  indexes temporal blocks. A 3D average pooling operator  $\mathcal{A}_{3d}$  generates the density  $D_s = \mathcal{A}_{3d}(E_s)$ .

(2) Key parameter selection: Extract top- $n$  temporal indices  $\{i_{t,j}\}_{j=1}^n$  from  $D_s$  along the temporal axis, where  $n$  matches the control point count. Then convert these to normalized timestamps:  $t_{l,j} = \frac{i_{t,j}}{N_a}, \forall j \in 1, \dots, n$ .

(3) Curve parameter update: For clamped NURBS constraints, compute adjustable knots  $t_{c,i}$  via sliding-window averaging and weight values from density normalization:

$$t_{c,i} = \frac{1}{p} \sum_{k=0}^{p-1} t_{l,i+k}, \quad i \in 1, \dots, n-p+1, \quad (7)$$

$$\{w_j\}_{j=1}^n = \sigma \left( D_s(i_{t,j})_{j=1}^n \right), \quad (8)$$

where  $\sigma(\cdot)$  denotes softmax normalization.

### 3.3.2. Control Points

The trajectory control points are dynamically updated through an event-guided spatiotemporal feature fusion process, as illustrated in Fig. 4. Our update mechanism operates in three phases:

**Spatio-temporal Query:** For each pixel  $\mathbf{p} = (x, y)$  at reference time  $t = 0$ , we sample its warped position  $\mathbf{p}_t = \mathbf{p} + \mathcal{T}(t, x, y)$  along the NURBS trajectory. At each queried timestamp  $\{t_{l,j}\}_{j=1}^n$  from density adaptation, we extract neighborhood correlation within radius  $r$  from the cost pyramid. This constructs 4D temporal/spatial cost volume  $C \in \mathbb{R}^{H \times W \times n \times (2r+1)^2}$ , capturing motion patterns.

**Feature Fusion:** We combine three information streams: (1) Temporal cost volumes at adaptively sampled  $\{t_{l,j}\}$ ; (2) Spatial cost volume from end-time displacement  $\mathcal{T}(1, x, y)$ ; (3) Context representation  $f_c$  from spatial feature encoding. These are concatenated into a spatio-temporal feature  $\mathcal{F} \in \mathbb{R}^{H \times W \times d}$  through convolutions.

**Iterative Refinement:** Initializing  $\mathcal{P}^{(0)} = \mathbf{0}$ , we iteratively update control points via feature  $\mathcal{F}$ :  $\mathcal{P}^{(j+1)} = \mathcal{P}^{(j)} + \text{GRU}(\mathcal{F}, \mathcal{P}^{(j)})$ , GRU is gated recurrent unit.

## 3.4. Scene Flow Estimation

Building upon the depth-normalized scene flow  $\hat{S}_f$  [51], we decompose it into two complementary components: optical flow  $\mathcal{O}$  and motion in depth  $\mathcal{M}$ :

$$\hat{S}_f = K \cdot \frac{S_f}{Z} = ((\mathcal{M} - 1) \cdot \mathbf{u} + \mathcal{M} \cdot \mathcal{O}), \quad (9)$$

where  $K$  represents the camera intrinsic matrix,  $Z$  denotes the relative scene depth,  $\mathbf{u}$  indicates the homogeneous image coordinates, and  $S_f$  corresponds to the scene flow. The optical flow component at time  $t$  is directly computable through trajectory queries:  $\mathcal{O}(x, y) = \mathcal{T}(t, x, y)$ .

The motion in depth component  $\mathcal{M}$  can be derived from the temporal gradient of the trajectory  $\mathcal{T}$ . Assuming a non-rotating rigid body under perspective projection with constant velocity in world coordinates, we establish the depth motion relationship:

$$\frac{Z_1}{Z_0} = \mathcal{M} = \frac{v_0 \Delta t + \Delta x}{v_1 \Delta t + \Delta x}, \quad \Delta t = t_1 - t_0, \quad (10)$$

where  $\Delta x$  and  $(v_0, v_1)$  represent the displacement and velocities of the object along the x-axis between times  $t_0$  and  $t_1$ , respectively (complete derivation provided in supplementary material).

By estimating the trajectory gradient  $\mathcal{T}'(t)$  [41] to obtain instantaneous velocity  $v$ , we compute  $\mathcal{M}$  via Eq. (10). Enforcing physical motion consistency across temporal observations yields the multi-view relationship:  $\mathcal{M}_2 = \frac{t_2}{t_1}(\mathcal{M}_1 - 1) + 1$ . This leads to our temporal stability constraint through viewpoint aggregation:

$$\mathcal{M}_k = \frac{1}{k} \sum_i \frac{t_k}{t_i} (\mathcal{M}_i - 1) + 1. \quad (11)$$

## 3.5. Supervision Framework

The proposed method employs a multi-task learning objective combining optical flow supervision, motion in depth supervision and temporal gradient regularization. The composite loss function is formulated as:

$$L = L_{\text{flow}} + L_{\text{depth}} + \lambda L_t, \quad \lambda = 10^{-7}. \quad (12)$$

**Optical Flow Supervision.** We implement an exponentially weighted L1 loss across iterative refinements:

$$L_{\text{flow}} = \sum_{k=1}^N \gamma^{N-k} \left( |\mathcal{O}_x^{(k)} - \mathcal{O}_x^{gt}|_1 + |\mathcal{O}_y^{(k)} - \mathcal{O}_y^{gt}|_1 \right), \quad (13)$$

**Motion in Depth Supervision.** Following similar principles, the motion in depth loss can be obtained as:

$$L_{\text{depth}} = \sum_{k=1}^N \gamma^{N-k} |\mathcal{M}^{(k)} - \mathcal{M}^{gt}|_1. \quad (14)$$




Figure 5. Characteristics of the data distribution in the CarlaEvent3D dataset




Figure 6. Event/image from CarlaEvent3D under multiple scenes.

**Temporal Gradient Regularization.** To prevent high-order trajectory distortions (*e.g.*, spiral artifacts in linear motion), we impose a first-order smoothness constraint on temporal derivatives of trajectory:

$$L_t = \sum_{i=1}^{T-1} |\mathcal{T}'(t_{i+1}) - \mathcal{T}'(t_i)|_1. \quad (15)$$

This regularization term enforces smooth temporal evolution of trajectory, effectively suppressing non-physical transformation while preserving motion consistency.

## 4. Experiment

### 4.1. Experimental Setup

**Dataset.** Current event-based 3D motion datasets predominantly derive ground truth through LiDAR in point cloud coordinates [47, 53], resulting in sparse and inconsistent signals. These datasets exhibit several critical limitations, including a lack of dynamic object annotations and limited environmental diversity (daytime only). The DSEC benchmark [13] provides improved event-aligned optical flow in corrected camera coordinates, supporting day-night scenarios. However, its filtering of moving objects restricts applicability to motion estimation. In this paper, we develop CarlaEvent3D dataset which provides various driving scenes and ground truth 3D motion annotations with Carla simulator. For event generation, we utilize this simulator to generate high-frame-rate video, then simulate events using the DVS Voltmeter [25] algorithm. The dataset contains 22,125 synchronized event-label tuples (resolution:

320×960) across 6 environmental conditions (Fig. 5), partitioned into: 1.Training: 45 sequences (13,275 samples); 2.Validation: 15 sequences (4,425 samples); 3.Testing: 15 sequences (4,425 samples). The collected objects guarantee a wide range of motion distribution, with velocity over 100 pixels / 100 ms (highway speeds). Furthermore, the distribution of events in various environments shows significant differences, from 0.5 - 5.0 Mev/s (foggy *v.s.* night).

**Experimental Setup.** We employ the AdamW optimizer with gradient clipping and the OneCycle learning rate policy over 60,000 iterations to ensure stable convergence. For the Event Kymograph projection kernel, we set the smoothing scale  $\sigma = 10$  to balance event density and temporal resolution. And the temporal subblock count  $N_a$  is set to 6 for efficient feature aggregation. The hyperparameters of the trajectory will be defined through the ablation study.

The scene flow validation is partitioned into optical flow and motion in depth in CarlaEvent3D. For optical flow estimation, we adopt Average End-Point Error (EPE) and Fl-outlier ratio (Fl) [14] metrics, while motion in depth accuracy is quantified through logarithmic error (log-mid) [51]. Furthermore, we validate methods on the DSEC benchmark for scene flow estimation and employ 3D End-Point Error (EPE<sub>3D</sub>) and Accuracy-at-5cm (ACC<sub>0.05</sub> metrics[47].

### 4.2. Validation on the CarlaEvent3D Dataset

In this section, we conducted comprehensive motion estimation benchmarking on the CarlaEvent3D dataset, with comparisons detailed in Tab. 1. To ensure fair comparison under event-driven paradigms, we re-implement four image-based methods – Expansion [51], TPCV [51], Scale-Flow [27], and Scale++ [26] – using Event Voxel as input with consistent training protocols.

The results reveal that the methods depending on spatial correlation exhibit significant performance degradation (5.274 - 1.963 EPE increase) due to event data’s inherent low spatial redundancy, On the other hand, the ETTCM [40] model, which estimates motion per event, is limited by the asynchronous nature of the algorithm and struggles to achieve ideal performance in complex motion scenarios. Our EMoTive framework addresses these

Model	flow@D		mid@D	flow@S		mid@S	params.(M)	time(s)
	EPE $\downarrow$	f1 $\downarrow$	log-mid $\downarrow$	EPE $\downarrow$	f1 $\downarrow$	log-mid $\downarrow$		
ERaft [14]	2.781	24.604	-	3.067	27.326	-	5.04	0.049
Expansion [51]	7.821	57.653	171.237	8.266	61.819	173.781	12.13	0.300
TPCV [28]	4.510	42.230	285.253	5.298	47.335	335.034	10.33	0.085
ScaleFlow [23]	4.518	42.885	268.050	5.288	47.956	321.067	10.70	0.090
Scale++ [26]	5.242	40.081	260.165	6.044	46.391	298.588	42.96	0.119
ETTCM <sup>1</sup> [40]	-	-	-	27.596	96.538	518.365	-	10 <sup>-6</sup> /ev
<b>EMoTive(UniForm)</b>	2.669	24.607	122.023	3.060	27.367	147.991	5.67	0.040
<b>EMoTive</b>	<b>2.547</b>	<b>22.866</b>	<b>113.593</b>	<b>2.850</b>	<b>24.659</b>	<b>138.402</b>	5.61	0.040

Table 1. The validation results on the CarlaEvent3D dataset. @D and @S represent dense labels and event-masked sparse labels.

Model	flow		mid
	EPE $\downarrow$	f1 $\downarrow$	log-mid $\downarrow$
ERaft [14]	0.451	0.773	-
Expansion [51]	1.590	6.808	179.044
TPCV [28]	0.782	2.353	379.688
ScaleFlow [23]	1.047	4.733	367.939
Scale++ [26]	1.106	6.464	345.220
ETTCM [26]	16.917	84.318	433.3
<b>EMoTive(UniForm)</b>	0.497	1.083	156.164
<b>EMoTive</b>	<b>0.439</b>	<b>0.742</b>	<b>152.722</b>

Table 2. Validation of 3D motion estimation on the DSEC dataset

limitations through fine temporal observation and event-guided spatio-temporal trajectory modeling, outperforms spatial-dependent methods by 43.5% in optical flow EPE, and 33.6% in motion in depth log-mid. Furthermore, the event-guided non-uniform trajectory design specifically contributes 4.6% (optical flow) and 6.9% (depth motion) accuracy gains compared to uniform B-spline parameterization.

In terms of parameter count and inference time, EMoTive with only 5.61M learnable parameters (47.6% fewer than ScaleFlow’s 10.7M baseline), achieves 40 ms inference speed over 100 ms event data on an NVIDIA A6000 GPU, representing 52.9% speed improvement.

### 4.3. Validation on the DSEC Dataset

To validate our method’s practical efficacy, we conduct 3D motion and scene flow estimation on the DSEC dataset. Following the dataset split of [47], we train ERaft [14] from scratch due to its unavailable official implementation. All methods are evaluated under identical conditions.

**3D Motion Estimation**. As shown in Tab. 2, EMoTive shows superior performance with 0.439px EPE (43.9% improvement over TPCV), 0.742% f1 ratio (1.611% reduction), and 152.722 log-mid error (14.7% better than Expansion). The spatial sparsity challenges methods like TPCV, whose performance degrades by 33.1% in motion in depth due to its spatial reliance.

**Scene Flow Analysis**. The scene flow evaluation in

Model	Input	EPE <sub>3D</sub> $\downarrow$	ACC <sub>0.05</sub> $\uparrow$
CamLiFlow [29]	RGB+PC	0.120	53.5%
RPEFlow [47]	RGB+PC+Ev	<b>0.104</b>	<b>60.5%</b>
Expansion [51]	Ev	0.489	4.3%
Scale++ [26]	Ev	0.603	1.4%
<b>EMoTive</b>	Ev	0.205	25.5%

Table 3. Quantitative evaluation of scene flow on DSEC datasets. "PC" refers to the radar point cloud, and "Ev" refers to the event.




Figure 7. Scene flow estimation. The green point cloud represents the initial position, the red is the position at the target time, and the red arrow denotes the real scene flow. The blue point cloud and arrows are the estimated position and corresponding scene flow.




Figure 8. Impact of the trajectory parameters.

Tab. 3 reveals that EMoTive surpasses monocular event-based methods by 58%, achieves  $EPE_{3D} = 0.205\text{cm}$  and  $ACC_{0.05} = 25.5\%$ . However, the multimodal fusion methods (CamliFlow/RPEFlow) maintain 41.5-49.3% advantage due to privileged depth from Lidar [47]. Furthermore, the scene flow projections in Fig. 7 demonstrate EMoTive’s capability to catch motion and structure in 3D.

### 4.4. Ablation Study

To systematically validate EMoTive’s architectural design, we conduct several ablation experiments on Car-

	w/o $L_t$	w/ $L_t$		flow	mid
$L_t$				EPE	log-mid
✗				3.148	139.409
✓				3.118	138.585

Table 4. Impact of trajectory regularization on EMOtive.

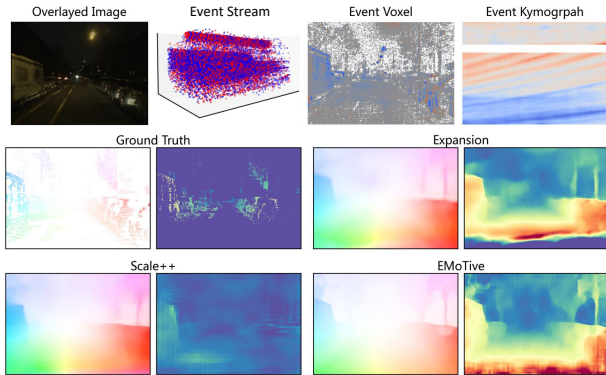


Figure 9. Visual comparison of motion estimation results (DSEC). The left is optical flow and the right is motion in depth estimation.

laEvent3D’s challenging scenarios (*e.g.*, rainy days, night scenes, *etc.*). Four critical components are analyzed: temporal resolution, trajectory degrees, control point number, and constraint loss ( $L_t$ ). All variants are trained from scratch with identical hyper-parameters.

**Temporal Resolution Analysis.** The temporal resolution in the Event Kymograph affects motion characterization. As shown in Fig. 8, we observe that the performance keeps improving with finer resolution as it better captures motion. However, when beyond 120 bins, noise amplification in temporal gradients degrades accuracy. Therefore, the optimal time granularity is 120 bins / 100ms.

**Trajectory degree.** The curve degree  $p$  of event-guided trajectories governs motion representation fidelity. As shown in Fig. 8, we observe that 3rd-order is optimal, where higher orders introduce overfitting in optical flow while lower orders underfitting in motion in depth.

**Control Point.** The control point count  $N_c$  balances representational capacity and over-smoothing risks. As quantified in Fig. 8, the optimal  $N_c = 5$  achieves 3.118px EPE (0.024px decrease from  $N_c = 4$ ) and 138.585 log-mid error (1.6% improvement over  $N_c = 6$ ).

**Trajectory Regularization.** The temporal gradient consistency constraint suppresses high-order distortions in trajectory learning. Without this constraint, as visualized in Tab. 4, trajectories exhibit undesirable curling due to unbounded higher-order derivatives. Quantitative results in Tab. 4 reveal subtle but critical improvements, 0.03 error decrease in EPE and 0.824 decrease in log-mid.




Figure 10. Visualization of trajectory estimated from EMOtive. Optical flow and motion in depth are derived from the trajectory.

## 4.5. Result Visualization

To intuitively demonstrate the motion estimation performance, we visualized the results on the DSEC dataset in a nighttime scene, with comparative results in Fig. 9. Due to dark noise, there is severe degradation in the event projections, which leads to suboptimal motion estimation from the model only relying on spatial features. However, EMOtive maintains sharp motion boundaries, yielding better motion results. This robustness stems from event-guided dual spatio-temporal features fusion and parametric representation of trajectories, introducing motion knowledge on the spatio-temporal process.

Furthermore, we visualized the trajectories estimated by EMOtive in Fig. 10. The trajectories not only provide motion in depth but also allow querying of the motion at any moment. Typical optical flow results assume that planar motion remains first-order invariant, only capturing linear motion. In contrast, the trajectory curves from EMOtive express high-order motion, better reflecting the heterogeneous spatio-temporal changes.

## 5. Conclusion

This paper addresses the fundamental challenge of local spatio-temporal inconsistent motion in 3D motion estimation through a novel trajectory-centric framework. The proposed EMOtive framework establishes temporal continuity in motion representation by developing parametric curve trajectories guided by events. The technical foundation is built upon two pillars: 1) A fine-grained temporal projection, termed Event Kymograph, that leverages a continuous temporal projection kernel while decoupling spatial observations, and 2) A non-uniform rational curve parameterization framework that adaptively models trajectories under event guidance. The introduced temporal gradient regularization further ensures the physical plausibility of the learned trajectories without sacrificing representational capacity. Comprehensive evaluation on both synthetic and real-world benchmarks demonstrates superior performance in challenging scenarios. Future research will focus on combining this paradigm with other sensors and exploring its integration with predictive control systems.

## 6. Acknowledgments

This work is supported by the National Natural Science Foundation of China (NSFC) under Grants 62225207, 62436008 and 62306295.

## References

- [1] Myo Tun Aung, Rodney Teo, and Garrick Orchard. Event-based plane-fitting optical flow for dynamic vision sensors in fpga. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2018. 3
- [2] Abhishek Badki, Orazio Gallo, Jan Kautz, and Pradeep Sen. Binary ttc: A temporal geofence for autonomous navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12946–12955, 2021. 1
- [3] Ryad Benosman, Charles Clercq, Xavier Lagorce, Sio-Hoi Ieng, and Chiara Bartolozzi. Event-based visual flow. *IEEE transactions on neural networks and learning systems*, 25(2):407–417, 2013. 3
- [4] Fabian Brickwedde, Steffen Abraham, and Rudolf Mester. Mono-sf: Multi-view geometry meets single-view depth for monocular scene flow estimation of dynamic traffic scenes. In *ICCV*, pages 2780–2790, 2019. 2
- [5] Jeffrey Byrne and Camillo J Taylor. Expansion segmentation for visual collision detection and estimation. In *2009 IEEE International Conference on Robotics and Automation*, pages 875–882. IEEE, 2009. 1
- [6] Jan Čech, Jordi Sanchez-Riera, and Radu Horaud. Scene flow estimation by growing correspondence seeds. In *CVPR*, pages 3129–3136. IEEE, 2011. 2
- [7] Matthew Cook, Luca Gugelmann, Florian Jug, Christoph Krautz, and Angelika Steger. Interacting maps for fast visual interpretation. In *The 2011 International Joint Conference on Neural Networks*, pages 770–776. IEEE, 2011. 2
- [8] Yongjian Deng, Hao Chen, Hai Liu, and Youfu Li. A voxel graph cnn for object classification with event cameras. In *CVPR*, pages 1172–1181, 2022. 3
- [9] Qiaole Dong and Yanwei Fu. Memflow: Optical flow estimation and prediction with memory. In *CVPR*, pages 19068–19078, 2024. 2
- [10] David Ferstl, Christian Reinbacher, Gernot Riegler, Matthias Rüther, and Horst Bischof. atgv-sf: Dense variational scene flow through projective warping and higher order regularization. In *2014 2nd International Conference on 3D Vision*, pages 285–292. IEEE, 2014. 2
- [11] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *CVPR*, pages 3867–3876, 2018. 3
- [12] 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 TPAMI*, 44(1):154–180, 2020. 2, 5
- [13] Mathias Gehrig, Willem Aarents, Daniel Gehrig, and Davide Scaramuzza. Dsec: A stereo event camera dataset for driving scenarios. *IEEE Robotics and Automation Letters*, 6(3):4947–4954, 2021. 6
- [14] Mathias Gehrig, Mario Millhäusler, Daniel Gehrig, and Davide Scaramuzza. E-raft: Dense optical flow from event cameras. In *2021 International Conference on 3D Vision (3DV)*, pages 197–206. IEEE, 2021. 3, 4, 6, 7
- [15] Mathias Gehrig, Manasi Muglikar, and Davide Scaramuzza. Dense continuous-time optical flow from event cameras. *IEEE TPAMI*, pages 1–12, 2024. 4
- [16] Paulo FU Gotardo and Aleix M Martinez. Non-rigid structure from motion with complementary rank-3 spaces. In *CVPR*, pages 3065–3072. IEEE, 2011. 2
- [17] Xiuye Gu, Yijie Wang, Chongruo Wu, Yong Jae Lee, and Panqu Wang. Hplflownet: Hierarchical permutohedral lattice flownet for scene flow estimation on large-scale point clouds. In *CVPR*, pages 3254–3263, 2019. 2
- [18] Han Han, Wei Zhai, Yang Cao, Bin Li, and Zheng-jun Zha. Event-based tracking any point with motion-augmented temporal consistency. *arXiv preprint arXiv:2412.01300*, 2024. 3
- [19] Berthold KP Horn and Brian G Schunck. Determining optical flow. *Artificial intelligence*, 17(1-3):185–203, 1981. 2
- [20] Xueyan Huang, Yueyi Zhang, and Zhiwei Xiong. Progressive spatio-temporal alignment for efficient event-based motion estimation. In *CVPR*, pages 1537–1546, 2023. 3
- [21] Hanme Kim, Stefan Leutenegger, and Andrew J Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *European conference on computer vision*, pages 349–364. Springer, 2016. 1
- [22] Yair Kittenplon, Yonina C Eldar, and Dan Raviv. Flow-step3d: Model unrolling for self-supervised scene flow estimation. In *CVPR*, pages 4114–4123, 2021. 2
- [23] Yuyang Leng, Renyuan Liu, Hongpeng Guo, Songqing Chen, and Shuochoao Yao. Scaleflow: Efficient deep vision pipeline with closed-loop scale-adaptive inference. In *ACM MM*, pages 1698–1706, 2023. 2, 7
- [24] Yijin Li, Yichen Shen, Zhaoyang Huang, Shuo Chen, Weikang Bian, Xiaoyu Shi, Fu-Yun Wang, Keqiang Sun, Hujun Bao, Zhaopeng Cui, et al. Blinkvision: A benchmark for optical flow, scene flow and point tracking estimation using rgb frames and events. In *ECCV*, pages 19–36. Springer, 2024. 3
- [25] Songnan Lin, Ye Ma, Zhenhua Guo, and Bihan Wen. Dvs-voltmeter: Stochastic process-based event simulator for dynamic vision sensors. In *ECCV*, pages 578–593. Springer, 2022. 6
- [26] Han Ling and Quansen Sun. Scaleflow++: Robust and accurate estimation of 3d motion from video. *arXiv preprint arXiv:2407.09797*, 2024. 2, 6, 7
- [27] Han Ling, Quansen Sun, Zhenwen Ren, Yazhou Liu, Hongyuan Wang, and Zichen Wang. Scale-flow: Estimating 3d motion from video. In *ACM MM*, pages 6530–6538, 2022. 6
- [28] Han Ling, Yinghui Sun, Quansen Sun, and Zhenwen Ren. Learning optical expansion from scale matching. In *CVPR*, pages 5445–5454, 2023. 2, 7
- [29] Haisong Liu, Tao Lu, Yihui Xu, Jia Liu, Wenjie Li, and Lijun Chen. Camliflow: bidirectional camera-lidar fusion for

- joint optical flow and scene flow estimation. In *CVPR*, pages 5791–5801, 2022. 7
- [30] Jiuming Liu, Guangming Wang, Weicai Ye, Chaokang Jiang, Jinru Han, Zhe Liu, Guofeng Zhang, Dalong Du, and Hesheng Wang. Diffflow3d: Toward robust uncertainty-aware scene flow estimation with iterative diffusion-based refinement. In *CVPR*, pages 15109–15119, 2024. 2
- [31] Liang Liu, Guangyao Zhai, Wenlong Ye, and Yong Liu. Unsupervised learning of scene flow estimation fusing with local rigidity. In *IJCAI*, pages 876–882, 2019. 2
- [32] Pengpeng Liu, Michael R Lyu, Irwin King, and Jia Xu. Learning by distillation: a self-supervised learning framework for optical flow estimation. *IEEE TPAMI*, 44(9):5026–5041, 2021. 2
- [33] Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. Meteor-net: Deep learning on dynamic 3d point cloud sequences. In *ICCV*, pages 9246–9255, 2019. 2
- [34] Weng Fei Low, Zhi Gao, Cheng Xiang, and Bharath Ramesh. Sofea: A non-iterative and robust optical flow estimation algorithm for dynamic vision sensors. In *CVPR*, pages 82–83, 2020. 3
- [35] Bruce D Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *IJCAI*, pages 674–679, 1981. 2
- [36] Kunming Luo, Chuan Wang, Shuaicheng Liu, Haoqiang Fan, Jue Wang, and Jian Sun. Upflow: Upsampling pyramid for unsupervised optical flow learning. In *CVPR*, pages 1045–1054, 2021. 2
- [37] Thiago Marinho, Massinissa Amrouche, Venanzio Cichella, Dušan Stipanović, and Naira Hovakimyan. Guaranteed collision avoidance based on line-of-sight angle and time-to-collision. In *2018 Annual American Control Conference (ACC)*, pages 4305–4310. IEEE, 2018. 1
- [38] Jun Nagata, Yusuke Sekikawa, and Yoshimitsu Aoki. Optical flow estimation by matching time surface with event-based cameras. *Sensors*, 21(4):1150, 2021. 3
- [39] Ngoc Long Nguyen, Jérémy Anger, Axel Davy, Pablo Arias, and Gabriele Facciolo. Self-supervised super-resolution for multi-exposure push-frame satellites. In *CVPR*, pages 1858–1868, 2022. 2
- [40] Urbano Miguel Nunes, Laurent Udo Perrinet, and Sio-Hoi Ieng. Time-to-contact map by joint estimation of up-to-scale inverse depth and global motion using a single event camera. In *ICCV*, pages 23653–23663, 2023. 6, 7
- [41] Les Piegl and Wayne Tiller. The nurbs book. 5
- [42] Shintaro Shiba, Yoshimitsu Aoki, and Guillermo Gallego. Secrets of event-based optical flow. In *ECCV*, pages 628–645. Springer, 2022. 3
- [43] Zachary Teed and Jia Deng. 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*, pages 402–419. Springer, 2020. 4
- [44] Stepan Tulyakov, Alfredo Bochicchio, Daniel Gehrig, Stamatios Georgoulis, Yuanyou Li, and Davide Scaramuzza. Time lens++: Event-based frame interpolation with parametric non-linear flow and multi-scale fusion. In *CVPR*, pages 17755–17764, 2022. 4
- [45] Sundar Vedula, Simon Baker, Peter Rander, Robert Collins, and Takeo Kanade. Three-dimensional scene flow. In *ICCV*, pages 722–729. IEEE, 1999. 2
- [46] René Vidal, Yi Ma, Stefano Soatto, and Shankar Sastry. Two-view multibody structure from motion. *IJCV*, 68(1):7–25, 2006. 2
- [47] Zhexiong Wan, Yuxin Mao, Jing Zhang, and Yuchao Dai. Rpeflow: Multimodal fusion of rgb-pointcloud-event for joint optical flow and scene flow estimation. In *ICCV*, pages 10030–10040, 2023. 2, 3, 6, 7
- [48] Zengyu Wan, Yang Wang, Zhai Wei, Ganchao Tan, Yang Cao, and Zheng-Jun Zha. Event-based optical flow via transforming into motion-dependent view. *IEEE Transactions on Image Processing*, 2024. 3
- [49] Andreas Wedel, Clemens Rabe, Tobi Vaudrey, Thomas Brox, Uwe Franke, and Daniel Cremers. Efficient dense scene flow from sparse or dense stereo data. In *ECCV*, pages 739–751. Springer, 2008. 2
- [50] Yuliang Wu, Ganchao Tan, Jinze Chen, Wei Zhai, Yang Cao, and Zheng-Jun Zha. Event-based asynchronous hdr imaging by temporal incident light modulation. *Optics Express*, 32(11):18527–18538, 2024. 2
- [51] Gengshan Yang and Deva Ramanan. Upgrading optical flow to 3d scene flow through optical expansion. In *CVPR*, pages 1334–1343, 2020. 2, 5, 6, 7
- [52] Shuai Yuan, Lei Luo, Zhuo Hui, Can Pu, Xiaoyu Xiang, Rakesh Ranjan, and Denis Demandolx. Unsamflow: Unsupervised optical flow guided by segment anything model. In *CVPR*, pages 19027–19037, 2024. 2
- [53] Hanyu Zhou, Yi Chang, and Zhiwei Shi. Bring event into rgb and lidar: Hierarchical visual-motion fusion for scene flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26477–26486, 2024. 2, 3, 6
- [54] Alex Zihao Zhu and Liangzhe Yuan. Ev-flownet: Self-supervised optical flow estimation for event-based cameras. In *Robotics: Science and Systems*, 2018. 3