

Deep Visual Odometry for Stereo Event Cameras

Sheng Zhong*, Junkai Niu*, Yi Zhou†

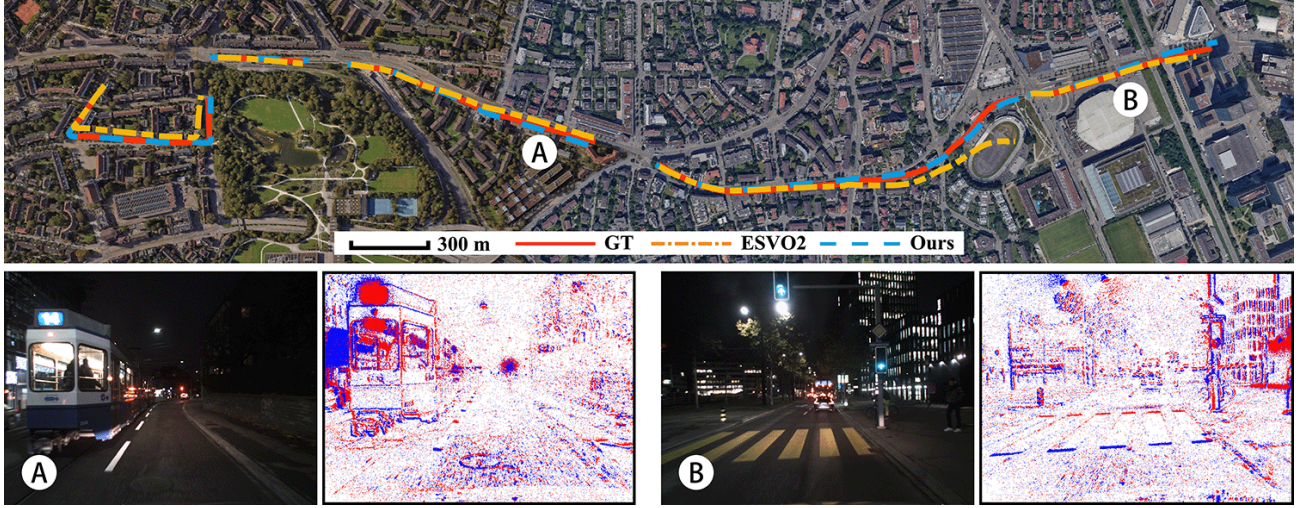


Fig. 1: The proposed system operates in real time on the *dsec_zurich_city09* [1] sequence, which encompasses large-scale, nighttime HDR scenario at VGA resolution. Top: The ground truth trajectory is provided by a LiDAR-IMU-based method [2] (red) and aligned with Google Maps. Due to the discontinuity of event data, the trajectory is segmented into five parts, with our method (blue) achieving excellent performance on each segment, outperforming the results of ESVO2 [3] (yellow). Bottom: The scene images and corresponding accumulated event maps (both for visualization purpose only).

Abstract—Event-based cameras are bio-inspired sensors with pixels that independently and asynchronously respond to brightness changes at microsecond resolution, offering the potential to handle state estimation tasks involving motion blur and high dynamic range (HDR) illumination conditions. However, the versatility of event-based visual odometry (VO) relying on handcrafted data association (either direct or indirect methods) is still unreliable, especially in field robot applications under low-light HDR conditions, where the dynamic range can be enormous and the signal-to-noise ratio is spatially-and-temporally varying. Leveraging deep neural networks offers new possibilities for overcoming these challenges. In this paper, we propose a learning-based stereo event visual odometry. Building upon Deep Event Visual Odometry (DEVO), our system (called Stereo-DEVO) introduces a novel and efficient static-stereo association strategy for sparse depth estimation with almost no additional computational burden. By integrating it into a tightly coupled bundle adjustment (BA) optimization scheme, and benefiting from the recurrent network’s ability to perform accurate optical flow estimation through voxel-based event representations to establish reliable patch associations, our system achieves high-precision pose estimation in metric scale. In contrast to the offline performance of DEVO, our system can process event data of Video Graphics Array (VGA) resolution in real time. Extensive evaluations on multiple public real-world datasets and self-collected data justify our system’s versatility, demonstrating superior performance compared to state-of-the-art event-based

VO methods. More importantly, our system achieves stable pose estimation even in large-scale nighttime HDR scenarios.

Index Terms—SLAM, Localization, Deep Learning for Visual Perception.

MULTIMEDIA MATERIAL

Video: <https://youtu.be/7UykRsmk3Zc>

Code: <https://github.com/NAIL-HNU/SDEVO.git>

I. INTRODUCTION

NEUROMORPHIC event-based cameras are bio-inspired visual sensors that asynchronously detect and transmit pixel-wise intensity changes, known as events. Consequently, they eliminate grayscale images and fixed capture rates, thereby reducing power consumption and bandwidth requirements. Endowed with microsecond temporal resolution and up to 160 dB dynamic range [4], event cameras are capable of tackling challenging scenarios that are beyond the reach of standard cameras, such as aggressive motion and/or HDR illumination tracking [5]–[7] and Simultaneous Localization and Mapping (SLAM) [8]–[11].

Despite the promising potential of event cameras, their unique output renders them incompatible with existing modules in visual VO/SLAM pipelines designed for standard cameras. How to establish accurate and stable data associations in event data or event representations is a primary challenge for event-based VO/SLAM systems. A widely adopted approach involves registering reconstructed local 3D maps to event representation images, commonly referred to as the direct method. In this context, Niu *et al.* [3] recently propose a stereo event-based VIO system called ESVO2 that achieves

Manuscript received: May 19, 2025; Accepted: August 29, 2025.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by the National Key Research and Development Project of China under Grant 2023YFB4706600.

All authors are with the Neuromorphic Automation and Intelligence Lab (NAIL) at School of Artificial Intelligence and Robotics, Hunan University, Changsha, China.

* equal contribution; † corresponding author (eeyzhou@hnu.edu.cn).

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

real-time performance at VGA resolution and demonstrates excellent performance across multiple event datasets, particularly in large-scale scenarios. However, ESVO2 still faces several limitations. Firstly, the inherent limitation of the direct method makes it difficult to handle aggressive motion scenarios. Secondly, it fails to operate effectively in complex nighttime HDR scenes, as the flicker of artificial light sources and the significant variation in light intensity with distance can severely disrupt the photometric consistency assumption and the spatio-temporal consistency constraints built upon it. As shown in Fig. 1, the excessive non-edge events exceed the capacity of ESVO2.

The introduction of deep networks may address these limitations and fully unlock the potential of event cameras in these challenging scenarios. Klenk *et al.* [12] recently propose a learning-based VO approach using a monocular event camera, called DEVO, demonstrating strong generalization and accuracy across extensive datasets from simulation to real environments. However, the scale ambiguity of the monocular system and its fully offline operation mode limits its real-world applications.

To overcome the limitations of DEVO [12], we extend it to a stereo VO system that enables online and real-time ego motion estimation in metric scale. Our system demonstrates superior performance in large-scale and HDR night scenes (see Fig. 1), where existing methods [3, 13]–[15] fail to operate or perform poorly. The contribution of the paper is summarized as follows.

- A learning-based visual odometry system using a stereo event camera, called “Stereo-DEVO”, which realizes high-precision ego motion estimation in metric scale. This is achieved by incorporating a novel and highly efficient static stereo association method, and integrating it into a tightly-coupled BA optimization (Sec. III).
- An online system that can operate in real time on event data of VGA resolution. This is achieved by refactoring the DEVO framework [12] on ROS, specifically by improving the processing of raw event data and establishing a reliable inter-process communication mechanism. (Sec. IV-F).
- A comprehensive evaluation on five publicly available datasets and a comparison against six existing methods (including five event-based methods and an image-aided method), demonstrating state-of-the-art accuracy in ego motion estimation and remarkable cross-scene generalization of the proposed method (Sec. IV). Our code will be released for further research.

Outline: The rest of the paper is organized as follows. First, a literature review is provided in Sec. II. Second, we provide a discussion of our method, particularly focusing on the entries listed in the contribution (Sec. III). Finally, experimental results are provided in Sec. IV, followed by the conclusion (Sec. V).

II. RELATED WORK

Like its standard-vision counterparts, event-based VO/SLAM can be broadly divided into model-based

methods and learning-based methods. The former estimates the camera trajectory by processing events and explicitly constructing photometric and geometric constraints, while the latter utilizes deep networks to learn implicit information from event streams or event representations for trajectory estimation.

Model-based Event VO/SLAM: From the perspective of event data processing, existing model-based methods can be mainly categorized into feature-based methods and direct methods. Leveraging the characteristics of event data, researchers modify the original Harris [16] and FAST [17] methods designed for standard cameras to develop new feature detection strategies, such as event corners [18, 19]. The geometric constraints provided by these feature can be directly utilized by mature geometric tools such as Perspective- n -Point (PnP) methods [20]. Despite the success of some feature-based methods (*e.g.*, [21]), event features lack robustness because the motion-dependent nature of event data often results in feature tracking failures during sudden changes in the event camera’s motion. Another mainstream strategy for feature detection and tracking is inspired by the motion compensation method [22], a unified pipeline for event-based model fitting. This strategy is widely employed in event-based visual-inertial odometry (VIO) pipelines [14, 23], where inertial measurement unit (IMU) measurements are fused to improve system performance through keyframe-based nonlinear optimization [24].

Unlike feature-based methods, direct methods refer to those that directly process either events or raw pixel information via some event representation. Based on the constant-brightness assumption in logarithmic scale, Kim *et al.* [25] propose the first direct method, consisting of three interleaved probabilistic filters to jointly estimate the camera pose, 3D map of the scene, and image intensity. To justify that intensity information recovery is not needed, Rebecq *et al.* [26] propose a geometric approach called EVO, which estimates the camera pose through 3D-2D registration, aligning the local 3D map reconstructed by EMVS [8] to the 2D event locations on the image plane. Recently, based on MC-EMVS [27], the successive work of EMVS, Ghosh *et al.* [15] employ a similar alignment strategy to develop a more accurate and efficient parallel tracking and mapping system. Unlike the aforementioned methods, Zhou *et al.* [13] propose the first stereo event-based VO pipeline, which exploits spatio-temporal consistency of the events across the image planes for mapping and localization. Building upon this pipeline, [3, 28] incorporate IMU measurements and improve event representations to enhance system performance. Among these, ESVO2 [3] achieves state-of-the-art performance in accuracy and efficiency across multiple event datasets by an efficient edge pixel sampling strategy and a tightly-coupled visual-inertial optimization. However, ESVO2 still suffers from limitations in robustness and generalization, particularly when dealing with aggressive motion and nighttime HDR scenarios. These challenges stem from both the inherent limitation of direct methods and the excessive non-edge events in nighttime scenes that compromise spatio-temporal consistency. The introduction of deep networks offers new possibilities for overcoming these challenges.

Learning-based Event VO/SLAM: Zhu *et al.* [29] and Ye

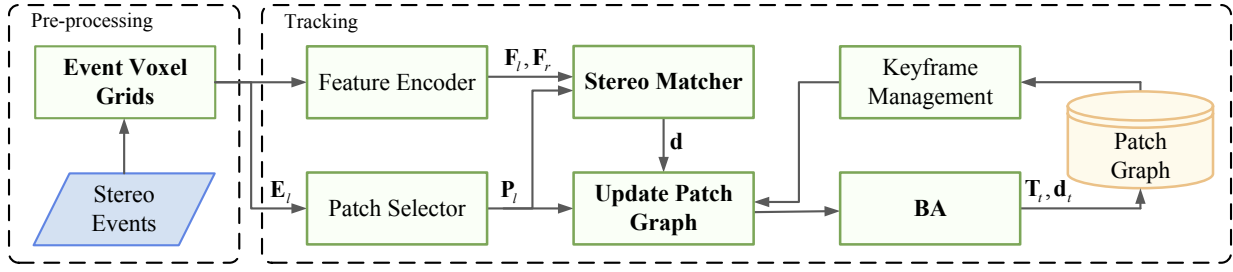


Fig. 2: Flowchart of the proposed Stereo-DEVO system. Our system builds on DEVO [12], with newly added and customized functions highlighted in bold. The two main modules (*i.e.*, pre-processing and tracking) run independently online.

et al. [30] employ a convolutional neural network (CNN) to predict optical flow, semi-dense depth, and ego-motion using unsupervised learning on the MVSEC [31] dataset. Despite achieving favorable results on test data, these methods fail to show generalization beyond the training scenarios. To address this limitation, Klenk *et al.* [12] extend DPVO [32] to the event modality by employing a novel patch selection mechanism for sparse event data, called DEVO. By training on a very large dataset, it demonstrates generalization from simulation to several real-world benchmarks and achieves competitive performance compared to both classical and deep-learning baselines. However, DEVO only takes a monocular event camera as input, resulting in trajectory with scale ambiguity. Building modules upon DEVO, Guan *et al.* [33] integrates an event-IMU BA to recover scale information and achieves success on certain small-scale datasets, but its lack of robustness in scale estimation leads to poor performance on larger-scale datasets such as DSEC [1].

Similar to [33], we also aim to eliminate the scale ambiguity of DEVO. However, our system extends it by introducing stereo event data to construct static stereo associations, achieving stable and accurate depth estimation, and integrating it into a tightly coupled BA optimization module to obtain metric-scale pose estimation. Our method is a learning-based stereo event VO system, demonstrating generalization and robustness across a variety of scenarios and achieving superior performance.

III. METHODOLOGY

A. System Overview

The goal is to deliver an online stereo event-based VO system that can track the 6-DoF motion of the camera in metric scale. To this end, we build on top of DEVO [12] and extend it to an stereo pipeline with additional non-trivial efforts, as shown in Fig. 2. The pre-processing module receives as input the stereo event stream and performs tri-linear interpolation counting according to the pixel coordinates and timestamps of the events ($e_k = (x_k, y_k, t_k, p_k)$) within a fixed time interval, thereby generating a pair of voxel grids $\mathbf{E}_l, \mathbf{E}_r \in \mathbb{R}^{H \times W \times 5}$. Here H, W denote the sensor's spatial resolution, and 5 represents that the time interval is divided into five equal segments.

In the tracking module, we reuse the network architecture of DEVO, and in particular, the following three functions: *patch selector*, *feature encoder*, and *iterative update operator*, which

are incorporated in the *update patch graph*. The *patch selector* predicts a score map based on the left voxel grid and uses a pooled multinomial sampling strategy to uniformly extract a fixed number of high-scoring patches \mathbf{P}_l . Each patch in \mathbf{P}_l can be denoted as

$$\mathbf{P}_i = (\mathbf{u}, \mathbf{v}, \mathbf{1}, \mathbf{d})^T \quad \mathbf{u}, \mathbf{v}, \mathbf{d} \in \mathbb{R}^{1 \times p^2}, \quad (1)$$

where \mathbf{u} and \mathbf{v} represent the pixel coordinates, while \mathbf{d} and p ($p = 3$ in our implementation) are the inverse depth and the patch size, respectively. The *feature encoder* establishes a two-level matching feature pyramid for the left voxel, and extracts only matching features from the right voxel. To simplify the notation, we refer to the feature map and the voxel collectively as a frame, denoted by \mathbf{F}_l and \mathbf{F}_r , respectively. Based on \mathbf{P}_l and \mathbf{F}_l , the *iterative update operator* establishes and updates the temporal associations of patches on the left frame within the patch graph by iteratively estimating the optical flow. Meanwhile, our efficient static stereo association method accurately recovers the depth of the patches and integrate them into the patch graph (Sec. III-B). Finally, a tightly-coupled BA module refines the frame poses and patch depths within the sliding window, which is maintained at a fixed size by keyframe management (Sec. III-C).

B. Depth Estimation and Update Patch Graph

Accurate static stereo associations are fundamental for resolving scale ambiguity through metric depth estimation in a stereo stereo VO system. Our system follows DEVO's utilization of event sparsity, thereby eliminating the need for dense or semi-dense depth maps as performed in [29, 30], which not only require the introduction of additional networks but also significantly increase the computational burden. To achieve accurate depth estimation while maintaining system efficiency, we propose a novel and efficient static stereo association method that exploits geometric constraints and the system's existing matching feature pyramids.

Given a patch \mathbf{P}_i , let its center coordinates on the first-level matching feature be (x, y) . Since the voxel grids have been stereo rectified, the corresponding matching point on the right feature map must lie between $(x - D, y)$ and (x, y) , where D is the maximum disparity. Let $\mathbf{g}_l(x, y)$ and $\mathbf{g}_r(x, y)$ represent the $p \times p$ patch centered at (x, y) on the first-level matching feature of \mathbf{F}_l and \mathbf{F}_r , respectively, then the depth estimation problem can be expressed as:

$$\delta_{best} = \arg \max_{1 \leq \delta \leq D} \langle \mathbf{g}_l(x, y), \mathbf{g}_r(x - \delta, y) \rangle, \quad (2)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 3: *Patch graph within the sliding window.* The patches selected from each frame are connected with the nearby left frames and current right frame.

where δ_{best} denotes the optimal disparity, corresponding to the static stereo association; and $\langle \cdot \rangle$ represents the dot product, used to measure the similarity between two patches. Then the inverse depth can be simply calculated from δ_{best} and the camera parameters via the operation of triangulation. This static stereo association strategy effectively utilizes the information contained in the feature pyramid as a similarity criterion, enabling accurate and robust depth estimation. More importantly, it introduces almost no additional computational burden to the system (Sec. IV-F), which greatly ensures the system’s real-time performance. Notably, our method performs searches only on first-level feature map, as the GPU’s high parallelism allows for faster execution compared to coarse-to-fine searches on two-level feature pyramids without compromising accuracy.

After depth estimation, the newly selected patches with established static stereo associations are incorporated into a fixed-length sliding window along with their corresponding left and right frames, as depicted in Fig. 3. The associations between consecutive left frames are constructed and updated by the *iterative update operator*. It is noteworthy that the *iterative update operator* is excluded from patch tracking on the right frame due to its reliance on optical flow estimation. Unlike consecutive frames, there is no optical flow constraint between the left and right frames at a synchronized timestamp. Similarly, methods that rely on similarity metrics such as zero-mean normalized cross-correlation (ZNCC) for block matching on voxel grids are not considered, as these methods are prone to mismatches. Experiments demonstrate that our static stereo association strategy outperforms these methods in terms of accuracy and robustness (Sec. IV-E).

C. Bundle Adjustment

Building upon the sliding window in Fig. 3, we formulate a tightly-coupled BA optimization problem that accurately and robustly estimates and refines both the camera’s metric-scale pose and the patches’ depths within the window by simultaneously integrating static and temporal stereo associations. The state variables in the optimization window are defined as follows:

$$\begin{aligned} \chi &= [\mathbf{x}_0, \mathbf{x}_1 \dots \mathbf{x}_n] \\ \mathbf{x}_k &= [\mathbf{R}_{c_k}^w, \mathbf{p}_{c_k}^w, \lambda_0^k, \lambda_1^k, \dots, \lambda_{m-1}^k], k \in [0, n]. \end{aligned} \quad (3)$$

\mathbf{x}_k represents the state of the k^{th} keyframe in the optimization window, including the camera’s position and orientation in the

TABLE I: Event camera parameters of the datasets used in the experiments.

| Dataset | Cameras | Resolution [px] | BL [cm] | FOV [°] |
|---------------------|------------------|-----------------|---------|---------|
| <i>rpg</i> [34] | DAVIS240C | 240×180 | 14.7 | 62.9 |
| <i>MVSEC</i> [31] | DAVIS346 | 346×260 | 10.0 | 74.8 |
| <i>DSEC</i> [1] | Prophesee Gen3.1 | 640×480 | 59.9 | 60.1 |
| <i>VECTOR</i> [35] | Prophesee Gen3 | 640×480 | 17.0 | 67.0 |
| <i>TUM-VIE</i> [36] | Prophesee Gen4 | 1280×720 | 11.8 | 65.0 |

world frame, as well as the inverse depth of all patches in that keyframe. n and m denote the number of keyframes and the number of patches per frame, respectively ($n = 10$ and $m = 96$ in our implementation). We define the optimization problem as minimizing the Mahalanobis norm of the reprojection residuals of patches across all temporal and static stereo associations:

$$\min_{\chi} \left\{ \sum_{(i,j) \in \mathcal{C}} \|r_C(\hat{z}_j^i, \chi)\|_{\Sigma_{ij}}^2 + \sum_{(l,r) \in \mathcal{S}} \|r_S(\hat{z}_r^l, \chi)\|_{\Sigma_{lr}}^2 \right\}. \quad (4)$$

$r_C(\hat{z}_j^i, \chi)$ and $r_S(\hat{z}_r^l, \chi)$ are reprojection residuals for temporal stereo and static stereo, respectively. The corresponding sets of associations within the window are denoted by \mathcal{C} and \mathcal{S} . We solve the optimization problem using two Gauss-Newton iterations to find the optimal camera poses and the inverse depth of corresponding patches. To improve computational efficiency, the Schur complement is employed to decompose the optimization problem and accelerate the Gauss-Newton iterations.

Each time new frames are added to the window and BA optimization is completed, we compute the optical flow magnitude between the $n-5$ and $n-3$ left frames using the patch graph. If this is below than a certain threshold, the $n-4$ frame is removed; otherwise, the oldest frame is discarded. These two frames are selected because the patch associations remain relatively stable after multiple updates, reducing the likelihood of misjudgments.

IV. EXPERIMENTS

In this section, we provide a comprehensive evaluation of our system. It is worth noting that all experiments are conducted without modifying or fine-tuning the network models. First, we introduce the publicly available stereo event-based dataset used in the evaluation (Sec. IV-A). Next, we compare the trajectory results of our system with five state-of-the-art event-based VO/VIO algorithms and one image-aided event-based VIO method, demonstrating superior trajectory accuracy compared to existing event-based methods (Sec. IV-B). Third, the robustness of our system in both large-scale scenarios (Sec. IV-C) and aggressive motion scenarios (Sec. IV-D) is evaluated using both self-collected data and open-source sequences. Finally, we investigate the accuracy of the novel static stereo associations (Sec. IV-E) and provide a computational-complexity analysis of each function in our system (Sec. IV-F).

A. Public Datasets Used

To comprehensively evaluate the performance of our system, we use five publicly available stereo event datasets with

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II: Relative pose error and Absolute trajectory error (RMS) on datasets of multiple resolutions [RPE:cm/s, ATE:cm]. Best and second best results are **highlighted** and underlined, respectively.

| Dataset | Sequence | ESVO [13] | | ES-PTAM [15] | | ESIO [14] | | ESVO2 [3] | | DEIO [33] | | Ours | |
|---------|------------------------|-----------|---------|--------------|--------------|-----------|---------------|---------------|----------------|---------------|---------------|---------------|---------------|
| | | RPE | ATE | RPE | ATE | RPE | ATE | RPE | ATE | RPE | ATE | RPE | ATE |
| rpg | <i>box</i> | 7.20 | 5.80 | 4.22 | <u>4.06</u> | 20.39 | 11.38 | <u>4.18</u> | 4.31 | – | – | 2.87 | 1.92 |
| | <i>monitor</i> | 3.20 | 3.30 | 2.59 | 2.34 | 13.17 | 7.87 | <u>1.69</u> | <u>2.31</u> | – | – | 1.65 | 0.91 |
| | <i>bin</i> | 3.10 | 2.80 | <u>1.82</u> | 2.57 | 12.10 | 7.08 | 2.53 | <u>2.27</u> | – | – | 1.54 | 1.16 |
| | <i>desk</i> | 4.50 | 3.20 | 3.99 | 2.84 | 6.18 | 3.16 | <u>3.53</u> | <u>1.57</u> | – | – | 3.33 | 1.46 |
| | <i>reader</i> | 5.60 | 6.60 | – | – | failed | failed | 2.12 | 2.68 | – | – | <u>3.59</u> | <u>3.78</u> |
| MVSEC | <i>indoor1_edited</i> | 7.38 | 16.59 | 6.89 | 15.02 | 228.84 | 820.36 | <u>5.05</u> | 7.63 | 66.62 | <u>7.52</u> | 3.73 | 3.76 |
| | <i>indoor2_edited</i> | 7.39 | 14.94 | – | – | 134.48 | 417.85 | <u>6.12</u> | 10.05 | 53.92 | 6.96 | 5.73 | 8.00 |
| | <i>indoor3_edited</i> | 5.97 | 10.03 | – | – | failed | failed | 4.75 | <u>7.35</u> | 37.55 | 29.75 | <u>5.18</u> | 5.57 |
| | <i>indoor4_edited</i> | failed | failed | – | – | 181.03 | 173.51 | 10.36 | <u>5.59</u> | 122.48 | 12.71 | <u>14.15</u> | 4.57 |
| DSEC | <i>city04_a</i> | 69.14 | 370.32 | 26.92 | 131.62 | 187.03 | 940.80 | 16.98 | 56.17 | 641.27 | 184.51 | <u>23.70</u> | 96.86 |
| | <i>city04_b</i> | 32.84 | 115.56 | 26.48 | <u>29.02</u> | 188.23 | 434.87 | <u>21.70</u> | 73.83 | 726.44 | 116.44 | 13.56 | 17.65 |
| | <i>city04_c</i> | 82.03 | 932.84 | 86.21 | 1184.37 | 145.64 | 1153.69 | <u>39.12</u> | <u>508.71</u> | 1501.47 | 9128.73 | 34.39 | 481.62 |
| | <i>city04_d</i> | 190.67 | 2676.11 | 68.02 | 1053.87 | 540.98 | 6822.53 | 18.00 | <u>546.58</u> | 1939.15 | 10970.03 | <u>31.18</u> | 406.20 |
| | <i>city09_a</i> | 271.62 | 2433.91 | failed | failed | failed | failed | <u>80.86</u> | <u>1183.04</u> | – | – | 15.73 | 196.97 |
| | <i>city09_b</i> | 138.33 | 742.19 | 68.92 | 221.35 | 561.41 | 2887.84 | 32.10 | 107.48 | – | – | <u>55.00</u> | <u>267.80</u> |
| | <i>city09_c</i> | 417.00 | 2962.40 | 1998.86 | 3673.80 | failed | failed | <u>361.56</u> | <u>1648.30</u> | – | – | 273.98 | 931.15 |
| | <i>city09_d</i> | 280.43 | 5398.48 | failed | failed | 105.62 | 1766.16 | <u>91.51</u> | 1920.36 | – | – | 39.38 | 625.81 |
| | <i>city09_e</i> | 339.24 | 3345.86 | 744.30 | 1480.79 | 801.43 | 4201.00 | <u>125.22</u> | <u>1075.09</u> | – | – | 38.56 | 246.26 |
| | <i>city11_a</i> | 56.82 | 366.22 | 25.14 | 117.86 | 54.53 | 107.36 | <u>12.94</u> | <u>48.77</u> | – | – | 10.32 | 47.56 |
| | <i>city11_b</i> | 188.97 | 3241.69 | failed | failed | 53.55 | <u>300.14</u> | 14.83 | 441.79 | – | – | <u>15.58</u> | 282.06 |
| VECTor | <i>robot_normal</i> | 16.79 | 7.32 | – | – | 3.38 | 5.17 | <u>3.16</u> | <u>4.81</u> | – | – | 1.84 | 3.16 |
| | <i>corner_slow</i> | 6.45 | 13.70 | – | – | 7.01 | 2.67 | <u>1.91</u> | <u>2.15</u> | 13.25 | 10.70 | 1.49 | 1.06 |
| | <i>hdr_normal</i> | 5.55 | 18.40 | – | – | 15.65 | 27.85 | <u>3.79</u> | <u>13.53</u> | – | – | 1.55 | 6.71 |
| | <i>sofa_normal</i> | failed | failed | – | – | 54.32 | 43.94 | <u>13.59</u> | <u>40.28</u> | – | – | 4.44 | 7.18 |
| | <i>desk_normal</i> | 5.12 | 20.81 | – | – | failed | failed | <u>3.50</u> | <u>16.47</u> | 86.55 | 254.76 | 2.00 | 6.48 |
| | <i>corridors_dolly</i> | failed | failed | – | – | failed | failed | failed | failed | <u>164.53</u> | <u>516.08</u> | 16.87 | 196.53 |
| | <i>units_dolly</i> | failed | failed | – | – | failed | failed | failed | failed | <u>138.73</u> | <u>826.73</u> | 32.43 | 444.51 |
| TUM-VIE | <i>1d_trans</i> | 8.41 | 12.54 | 0.71 | <u>1.05</u> | failed | failed | 1.06 | 3.33 | 24.49 | 1.85 | <u>0.89</u> | 1.05 |
| | <i>3d_trans</i> | 7.47 | 17.19 | 4.94 | 8.53 | failed | failed | <u>2.32</u> | 7.26 | 26.43 | <u>1.28</u> | 1.73 | 1.26 |
| | <i>6d_trans</i> | 14.28 | 13.46 | 12.02 | 10.25 | failed | failed | 2.49 | <u>3.21</u> | 26.31 | 1.45 | 2.55 | <u>1.69</u> |
| | <i>desk</i> | 5.87 | 12.92 | 2.53 | 2.50 | failed | failed | <u>3.41</u> | 6.16 | 30.75 | 1.49 | 2.24 | <u>2.44</u> |
| | <i>desk2</i> | 6.23 | 4.42 | 4.04 | 7.20 | failed | failed | <u>2.41</u> | 4.02 | 33.65 | 1.45 | 2.35 | <u>1.79</u> |
| Others | <i>hnu_campus</i> | – | 1321.80 | – | – | – | 1420.48 | – | <u>181.77</u> | – | – | – | 151.81 |
| | <i>drone_fast</i> | failed | failed | – | – | 192.13 | 154.04 | failed | failed | – | – | 19.39 | 13.29 |

* “–” represents the lack of experimental results of the algorithm (according to the respective publication). RPE evaluation on *hnu_campus* is unavailable due to missing pose ground truth.

varying resolutions and scene geometries, including *rpg* [34], *MVSEC* [31], *DSEC* [1], *VECTor* [35] and *TUM-VIE* [36]. The parameters of the event cameras used in each dataset are shown in Tab. I.

RPG, *MVSEC*, *VECTor*, and *TUM-VIE* are all collected in indoor scenes, with four different resolutions ranging from 240×180 to 1280×720 . Except for *MVSEC*, which is captured using a stereo event camera mounted on a drone, the other three datasets are collected with a handheld one. These datasets primarily involve reciprocal motion within the same scene, whereas only *VECTor* provides some sequences featuring long-distance indoor motion. *DSEC* consists of a large-scale collection of outdoor driving scenes, featuring a variety of complex scenarios (e.g., moving object). Since the *DSEC* dataset does not provide ground truth trajectories, we use the output of the LiDAR algorithm (LIO-SAM [2]) as the ground truth for subsequent evaluation.

B. Full System Evaluation

To evaluate the camera pose tracking performance of our system, we conduct extensive comparisons with five event-based VO/VIO systems across five stereo event datasets, as well as with an image-aided stereo event VIO algorithm. Among them, ESVO [13] is the first event-based stereo VO

TABLE III: Comparison of trajectory results on the *DSEC* and *VECTor* datasets [RPE:cm/s, ATE:cm].

| Dataset | Sequence | ESVIO [14] | | Ours | |
|--------------------|------------------------|-------------|--------------|---------------|---------------|
| | | RPE | ATE | RPE | ATE |
| DSEC | <i>city04_a</i> | 130.73 | 508.17 | 23.70 | 96.86 |
| | <i>city04_b</i> | 260.02 | 129.70 | 13.56 | 17.65 |
| | <i>city04_c</i> | 279.17 | 2938.09 | 34.39 | 481.62 |
| | <i>city04_d</i> | 196.22 | 2217.39 | 31.18 | 406.20 |
| | <i>city09_a</i> | 34.99 | 328.51 | 15.73 | 196.97 |
| | <i>city09_b</i> | 856.53 | 3481.49 | 55.00 | 267.80 |
| | <i>city09_c</i> | 632.14 | 8043.46 | 273.98 | 931.15 |
| | <i>city09_d</i> | 128.96 | 1677.13 | 39.39 | 564.33 |
| | <i>city09_e</i> | 54.11 | 438.19 | 36.49 | 246.26 |
| VECTor | <i>robot_normal</i> | 2.20 | 4.95 | 1.84 | 3.16 |
| | <i>corner_slow</i> | 1.67 | 1.43 | 1.49 | 1.06 |
| | <i>hdr_normal</i> | 0.90 | 1.96 | 1.55 | 6.71 |
| | <i>sofa_normal</i> | 4.55 | 5.06 | 4.44 | 7.18 |
| | <i>desk_normal</i> | 2.24 | 6.00 | 2.00 | 6.48 |
| | <i>corridors_dolly</i> | 17.32 | 92.04 | 16.87 | 196.53 |
| <i>units_dolly</i> | 101.99 | 872.07 | 32.43 | 444.51 | |

algorithm using a direct method. Building upon it, ESVO2 [3] also utilizes a direct method and incorporates IMU measurements to enhance the system. ES-PTAM [15] is the latest event-only direct stereo VO system, differing from ESVO in that it employs Multi-Camera Event-based Multi-View Stereo (MC-EMVS) [27] for mapping. DEIO [33], similar to our system, is a successive work of DEVO [12], incorporating an IMU to estimate metric-scale poses. We do not compare

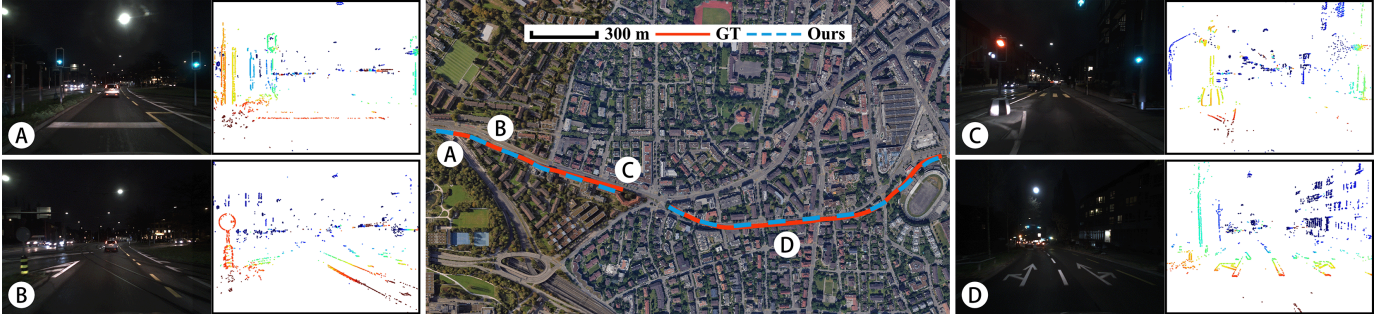


Fig. 4: Estimated camera trajectories and MC-EMVS reconstructed depth maps on partial sequences of *dsec_zurich_city09*. Inverse depth maps are color-coded, from red (close) to blue (far) over a white background. The scene images are for visualization purposes only.




Fig. 5: *Large-scale loop scenario*. This scenario is adopted from the *hnu_campus* sequence of ESVO2 [3], where the trajectory ground truth is provided by real-time kinematic (RTK) positioning and aligned with Google Maps.

with DEVO because its poses suffer from scale ambiguity. ESVO [14] is a feature-based VIO, leverages both standard cameras and event cameras. They also provide an event-only version (denoted by ESIO), and we compare our system with both of them. All trajectories used for evaluation are obtained from either the raw data or by running the open-sourced systems. We use “-” to indicate the absence of raw data, while “failed” denotes that the open-sourced systems fail to run successfully.

To conduct qualitative trajectory evaluation, we employ two standard evaluation metrics: relative pose error (RPE) and absolute trajectory error (ATE) [37]. The results of all methods are obtained using the EVO tool [38] and can be found in Tab. II and Tab. III. Our system achieves either the best or the second-best accuracy across all sequences.

Benefiting from the high-precision depth results provided by the multi-view stereo mapping method, ES-PTAM [15] performs well in small-scale, short-distance motion scenarios. In contrast, ESVO2 [3], which incorporates a more efficient edge pixel sampling strategy and integrates a tightly-coupled IMU backend, achieves significantly high accuracy in large-scale scenarios like *DSEC*. However, due to the inherent limitations of the direct method, ESVO2 struggles to handle scenes with significant optical flow (e.g. *corridors_dolly*). ESIO [14] is the only feature-based event VIO presented in the




Fig. 6: *Aggressive motion scenario*. Left column: grayscale images (for visualization only) at two different positions. Right column: trajectory results of our system.

table, employing ARC* [18] for feature detection and tracking on motion-compensated event images. Notably, ESIO fails to function properly on the *TUM-VIE* dataset, which we attribute to ARC*’s failure to handle excessively blurred edges in high-resolution event images, causing the system to malfunction. Unlike these methods, our system employs a deep network for patch selection and update on event voxel grids, fully leveraging the optical flow information contained in the events, which enables our system to achieve favorable results on multiple datasets that exhibit diverse geometric characteristics and motion patterns.

Building upon DEVO [12], DEIO [33] incorporates IMU measurements and integrates it into a tightly-coupled event-IMU BA unit to recover metric-scale poses, achieving outstanding performance on the *TUM-VIE* dataset. However, it exhibits significant scale errors on other datasets, presumably resulting from the improper weight between IMU and visual constraints. In comparison, our system eliminates scale ambiguity using static stereo associations, achieving greater robustness. Additionally, the initial poses in DEIO’s raw data differ from those employed by other algorithms, leading to significant anomalies in its RPE results.

To further demonstrate the performance of our system, we also compare it with ESVO [14], an image-aided event-based stereo VIO system, on the *DSEC* and *VECToR* datasets. The results presented in Tab. III confirm that our system’s accuracy

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE IV: Results of ablation study [ATE:cm].

| Sequence | Update | ZNCC | Ours |
|-----------------------------|--------|--------|---------------|
| <i>rpg_bin</i> | 1.69 | 3.54 | 1.16 |
| <i>mvsec_indoor1_edited</i> | 7.30 | 6.57 | 3.76 |
| <i>dsec_city04_d</i> | 528.22 | 544.18 | 406.20 |
| <i>vector_robot_normal</i> | 48.67 | 4.06 | 3.16 |
| <i>vector_units_dolly</i> | 574.43 | 491.40 | 444.51 |
| <i>tum_6d_trans</i> | 385.08 | 2.40 | 1.69 |

is comparable to that of image-based methods.

C. Evaluation in Large-Scale Scenarios

We conduct experiments in two large-scale scenarios to evaluate the scalability of our system, as well as its robustness in nighttime HDR environments. The first scenario uses a sequence called *hnu_campus* from ESVO2 [3], where a stereo-event camera mounted on an electric bicycle captures a closed-loop trajectory returning to the origin in a daytime campus environment, with moving objects introducing dynamic interference. As demonstrated in Fig. 5, our system’s estimated trajectory closely aligns with the ground truth. The second scenario is the nighttime long-distance driving sequence *zurich_city09*, provided by DSEC [1]. In such a challenging setting, the dim lighting conditions and noise events caused by artificial light sources make it difficult for image-based methods and other event-based approaches to function properly. However, our system leverages the rich optical flow information from voxel-based event representations and precise patch association mechanisms to achieve state-of-the-art performance. Specifically, we attain a mean absolute trajectory error of 4.52 m over a 2,510 m trajectory (Fig. 1), with detailed quantitative comparisons against baseline methods provided in Tab. II. To further validate trajectory accuracy, we implement 3D scene reconstruction using MC-EMVS [27] on partial sequences of *zurich_city09*. As illustrated in Fig. 4, leveraging the high-precision trajectory results provided by our method, MC-EMVS achieves faithful reconstruction of near-range scenes under challenging nighttime HDR conditions, while far-range scenes remain unrecoverable due to sparse event data.

D. Evaluation in Aggressive Motion Scenarios

To further assess the performance of our system in aggressive motion scenarios, Fig. 6 presents the trajectory estimation results in a high-speed indoor sequence called *drone_fast*, captured by a pair of DAVIS-346 event cameras mounted on a drone. The drone reaches a maximum linear velocity of 3.9 m/s and performs elliptical motions in the gravity direction. Under such aggressive motion, standard cameras exhibit significant motion blur, whereas our system achieves stable pose estimation. The quantitative results can be found in Tab. II, where our system yields an absolute trajectory error of only 0.13 m over a 55.4 m trajectory, outperforming other comparison methods.

TABLE V: Computational performance.[Time: ms]

| Node | Function | Desktop | Laptop | NVIDIA Jetson Orin NX |
|----------------|--------------------|---------|--------|-----------------------|
| Pre-processing | Event voxels | 13.52 | 40.41 | 121.40 |
| Tracking | Feature encoder | 18.39 | 34.13 | 170.51 |
| | Patch selector | 2.92 | 22.28 | 117.60 |
| | Depth estimation | 0.46 | 4.45 | 5.17 |
| | Patch graph update | 12.10 | 92.49 | 433.64 |
| | Bundle adjustment | 1.53 | 3.03 | 17.40 |
| | Others | 27.33 | 42.84 | 201.51 |
| | Subtotal | 62.73 | 199.22 | 945.83 |

E. Ablation Study

We conduct ablation experiments on sequences from multiple datasets to demonstrate that our static stereo association strategy is more accurate and robust than other methods. Specifically, as shown in Tab. IV, “Ours” refers to using the proposed static stereo association strategy for depth estimation, while “Update” refers to using the update operator, originally used for temporal stereo association only, to iteratively perform patch tracking twice on the right frame for depth estimation, while “ZNCC” refers to using ZNCC for block matching along the epipolar lines in the intermediate layer of voxel grids. “Update” achieves satisfactory results only in scenarios with small disparities caused by low resolution or large scene depth. However, when disparities become large, the update operator fails to establish accurate static stereo associations. Additionally, “ZNCC” is prone to mismatches, resulting in generally lower trajectory accuracy compared to our method.

F. Computational Efficiency

Unlike DEVO [12], which is a fully offline system, Stereo-DEVO is an online system implemented on ROS. To assess the computational performance of our system, we evaluate our algorithm on three different platforms: a desktop (Intel Core i7-12700K CPU, RTX 3090 GPU), a laptop (Intel Core i5-12500H CPU, RTX 2050 GPU), and an embedded platform (NVIDIA Jetson Orin NX), with compute precision set to mixed-precision, using the DSEC dataset. The average run time for each function is presented in Tab. V.

Our system consists of two processes: The pre-processing node, implemented in C++, converts the events into voxel grids, with the entire process executed on the CPU. The computational bottleneck lies in the tracking node, which is primarily implemented in Python and involves *feature encoder*, *patch selector*, *depth estimation*, *patch graph update*, and *bundle adjustment* on the voxel grids to estimate the trajectory. All these operations are accelerated using CUDA on the GPU. The “Others” entry includes operations such as data transfer and keyframe management, with most of the time spent transferring voxel grids from the CPU to the GPU. The subtotal runtime of the tracking node confirms that our system can stably operate in real time at 10 Hz on stereo event datasets with VGA resolution on the desktop platform, and at approximately 5 Hz on the laptop platform. However, the embedded platform’s performance is severely constrained by its 25W power budget, resulting in significantly higher latency.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

V. CONCLUSION

We present a learning-based stereo visual odometry system, which has won 1st place in the M3ED SLAM Challenge at the 5th CVPR Workshop on event-based vision. The proposed system is built on top of DEVO [12], a monocular event-only VO system that leverages a learned patch selector and a pooled multinomial sampling for tracking sparse event patches. Our system effectively resolves DEVO's scale ambiguity by integrating stereo event data, establishing robust static stereo correspondences with minimal computational burden. These correspondences are then incorporated into a tightly-coupled bundle adjustment optimization framework, thereby enabling precise metric-scale pose estimation. Furthermore, we enhance DEVO's architecture to transform it from a purely offline system into an online and real-time framework operating on VGA resolution. This advancement represents a significant step forward in enabling practical applications of learning-based event visual odometry. Extensive evaluations on real-world datasets demonstrate our method's superior generalization capability and robustness compared to existing systems. Our system particularly excels in large-scale nighttime HDR environments, effectively showcasing the inherent advantages of event cameras in handling such demanding conditions.

ACKNOWLEDGMENT

We thank Dr. Xiuyuan Lu for the help in the experiment. We also thank Dr. Yi Yu for proofreading.

REFERENCES

- [1] M. Gehrig, W. Aarents, D. Gehrig, and D. Scaramuzza, "DSEC: A stereo event camera dataset for driving scenarios," *IEEE Robot. Autom. Lett.*, vol. 6, no. 3, pp. 4947–4954, 2021.
- [2] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "LIO-SAM: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2020, pp. 5135–5142.
- [3] J. Niu, S. Zhong, X. Lu, S. Shen, G. Gallego, and Y. Zhou, "Esvo2: Direct visual-inertial odometry with stereo event cameras," *IEEE Transactions on Robotics*, 2025.
- [4] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor," *IEEE J. Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, 2008.
- [5] G. Gallego, J. E. A. Lund, E. Mueggler, H. Rebecq, T. Delbruck, and D. Scaramuzza, "Event-based, 6-DOF camera tracking from photometric depth maps," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2402–2412, Oct. 2018.
- [6] E. Mueggler, G. Gallego, H. Rebecq, and D. Scaramuzza, "Continuous-time visual-inertial odometry for event cameras," *IEEE Trans. Robot.*, vol. 34, no. 6, pp. 1425–1440, Dec. 2018.
- [7] Z. Wang, T. Molloy, P. van Goor, and R. Mahony, "Asynchronous blob tracker for event cameras," *IEEE Trans. Robot.*, 2024.
- [8] H. Rebecq, G. Gallego, E. Mueggler, and D. Scaramuzza, "EMVS: Event-based multi-view stereo—3D reconstruction with an event camera in real-time," *Int. J. Comput. Vis.*, pp. 1–21, Nov. 2017.
- [9] A. Rosinol Vidal, H. Rebecq, T. Horstschaefer, and D. Scaramuzza, "Ultimate SLAM? combining events, images, and IMU for robust visual SLAM in HDR and high speed scenarios," *IEEE Robot. Autom. Lett.*, vol. 3, no. 2, pp. 994–1001, Apr. 2018.
- [10] D. G. Javier Hidalgo-Carrio and D. Scaramuzza, "Learning monocular dense depth from events," in *3D Vision (3DV)*, Nov. 2020, pp. 534–542.
- [11] S. Guo and G. Gallego, "CMax-SLAM: Event-based rotational-motion bundle adjustment and SLAM system using contrast maximization," *IEEE Trans. Robot.*, vol. 40, pp. 2442–2461, 2024.
- [12] S. Klenk, M. Motzet, L. Koestler, and D. Cremers, "Deep event visual odometry," in *International Conference on 3D Vision, 3DV 2024, Davos, Switzerland, March 18-21, 2024*. IEEE, 2024, pp. 739–749.
- [13] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual odometry," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1433–1450, 2021.
- [14] P. Chen, W. Guan, and P. Lu, "ESVIO: Event-based stereo visual inertial odometry," *IEEE Robot. Autom. Lett.*, vol. 8, no. 6, pp. 3661–3668, 2023.
- [15] S. Ghosh, V. Cavinato, and G. Gallego, "ES-PTAM: Event-based stereo parallel tracking and mapping," in *Eur. Conf. Comput. Vis. Workshops (ECCVW)*, 2024.
- [16] C. Harris and M. Stephens, "A combined corner and edge detector," in *Proc. Fourth Alvey Vision Conf.*, vol. 15, 1988, pp. 147–151.
- [17] E. Rosten and T. Drummond, "Machine learning for high-speed corner detection," in *Eur. Conf. Comput. Vis. (ECCV)*, 2006, pp. 430–443.
- [18] I. Alzugaray and M. Chli, "Asynchronous corner detection and tracking for event cameras in real time," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3177–3184, Oct. 2018.
- [19] R. Li, D. Shi, Y. Zhang, K. Li, and R. Li, "FA-Harris: A fast and asynchronous corner detector for event cameras," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2019.
- [20] S. Li, C. Xu, and M. Xie, "A robust o (n) solution to the perspective-n-point problem," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 7, pp. 1444–1450, 2012.
- [21] A. Hadviger, I. Cvišić, I. Marković, S. Vražić, and I. Petrović, "Feature-based event stereo visual odometry," in *Eur. Conf. Mobile Robots (ECMR)*, 2021, pp. 1–6.
- [22] G. Gallego, H. Rebecq, and D. Scaramuzza, "A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2018, pp. 3867–3876.
- [23] A. Z. Zhu, N. Atanasov, and K. Daniilidis, "Event-based visual inertial odometry," in *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017, pp. 5816–5824.
- [24] S. Leutenegger, P. Furgale, V. Rabaud, M. Chli, K. Konolige, and R. Siegwart, "Keyframe-based visual-inertial SLAM using nonlinear optimization," in *Robotics: Science and Systems (RSS)*, 2013.
- [25] H. Kim, S. Leutenegger, and A. J. Davison, "Real-time 3D reconstruction and 6-DoF tracking with an event camera," in *Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 349–364.
- [26] H. Rebecq, T. Horstschäfer, G. Gallego, and D. Scaramuzza, "EVO: A geometric approach to event-based 6-DOF parallel tracking and mapping in real-time," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 593–600, 2017.
- [27] S. Ghosh and G. Gallego, "Multi-event-camera depth estimation and outlier rejection by refocused events fusion," *Adv. Intell. Syst.*, vol. 4, no. 12, p. 2200221, 2022.
- [28] J. Niu, S. Zhong, and Y. Zhou, "IMU-aided event-based stereo visual odometry," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2024, pp. 11 977–11 983.
- [29] A. Z. Zhu, L. Yuan, K. Chaney, and K. Daniilidis, "Unsupervised event-based learning of optical flow, depth, and egomotion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 989–997.
- [30] C. Ye, A. Mitrokhin, C. Fermüller, J. A. Yorke, and Y. Aloimonos, "Unsupervised learning of dense optical flow, depth and egomotion from sparse event data," *arXiv preprint arXiv:1809.08625*, 2018.
- [31] A. Z. Zhu, D. Thakur, T. Ozaslan, B. Pfrommer, V. Kumar, and K. Daniilidis, "The multivehicle stereo event camera dataset: An event camera dataset for 3D perception," *IEEE Robot. Autom. Lett.*, vol. 3, no. 3, pp. 2032–2039, July 2018.
- [32] Z. Teed, L. Lipson, and J. Deng, "Deep patch visual odometry," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [33] W. Guan, F. Lin, P. Chen, and P. Lu, "Deio: Deep event inertial odometry," *arXiv preprint arXiv:2411.03928*, 2024.
- [34] Y. Zhou, G. Gallego, H. Rebecq, L. Kneip, H. Li, and D. Scaramuzza, "Semi-dense 3D reconstruction with a stereo event camera," in *Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 242–258.
- [35] L. Gao, Y. Liang, J. Yang, S. Wu, C. Wang, J. Chen, and L. Kneip, "VECTOR: A versatile event-centric benchmark for multi-sensor slam," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 8217–8224, 2022.
- [36] S. Klenk, J. Chui, N. Demmel, and D. Cremers, "Tum-vie: The tum stereo visual-inertial event dataset," in *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [37] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Oct. 2012.
- [38] M. Grupp, "evo: Python package for the evaluation of odometry and slam." <https://github.com/MichaelGrupp/evo>, 2017.