

Edged USLAM: Edge-Aware Event-Based SLAM with Learning-Based Depth Priors

Şebnem Sarıözkan ^{id}, Hürkan Şahin ^{id}, Olaya Álvarez-Tuñón ^{id}, and Erdal Kayacan ^{id}

Abstract— Conventional visual simultaneous localization and mapping (SLAM) algorithms often fail under rapid motion, low illumination, or abrupt lighting transitions due to motion blur and limited dynamic range. Event cameras mitigate these issues with high temporal resolution and high dynamic range (HDR), but their sparse, asynchronous outputs complicate feature extraction and integration with other sensors, e.g., inertial measurement units (IMUs) and standard cameras. We present Edged USLAM, a hybrid visual-inertial system that extends Ultimate SLAM (USLAM) with an edge-aware front-end and a lightweight depth module. The front-end enhances event frames for robust feature tracking and nonlinear motion compensation, while the depth module provides coarse, region-of-interest (ROI)-based scene depth to improve motion compensation and scale consistency. Evaluations across public benchmarks and real-world unmanned air vehicle (UAV) flights demonstrate that performance varies significantly by scenario. For instance, event-only methods like point-line event-based visual-inertial odometry (PL-EVIO) or learning-based pipelines such as deep event-based visual odometry (DEVO) excel in highly aggressive or extreme HDR conditions. In contrast, Edged USLAM provides superior stability and minimal drift in slow or structured trajectories, ensuring consistently accurate localization on real flights under challenging illumination. These findings highlight the complementary strengths of event-only, learning-based, and hybrid approaches, while positioning Edged USLAM as a robust solution for diverse aerial navigation tasks.

Index Terms— Visual-Inertial SLAM, Event-Based Vision, Deep Learning for Visual Perception, Perception in Aerial Systems.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) is a cornerstone of autonomous navigation, enabling robots to build a map of their surroundings while estimating their own motion within it. Despite its success in ground vehicles [1], augmented reality/virtual reality devices [2], and service robots [3], the demand for robust SLAM is particularly critical in aerial robotics, where drones operate in hazardous or inaccessible environments [4]. Examples include collapsed buildings after earthquakes [5], confined ship compartments during recycling and inspection tasks [6], industrial facilities with toxic atmospheres or underground tunnels and mines [7]. In these scenarios, global positioning system (GPS) signals are often unavailable, making onboard, real-time state estimation essential for safe and reliable operation.

*This work was supported by the Horizon Europe Grant Agreement No. 101136056.

Ş. Sarıözkan, H. Şahin, and E. Kayacan are with the Automatic Control Group (RAT), Paderborn University, 33098 Paderborn, Germany (e-mail: {sariozka, hursah}@mail.uni-paderborn.de and erdal.kayacan@uni-paderborn.de). O. Álvarez-Tuñón is with EIVA A/S, and IT University of Copenhagen (ITU), Denmark (e-mail: oltu@itu.dk).




Fig. 1. Overview of the proposed Edged USLAM framework. The UAV (right) equipped with a protective cage, event camera, and depth module navigates through a cluttered indoor environment (top). The pipeline integrates learning-based depth estimation (bottom left) with enhanced event frames for robust ORB-based tracking (bottom center-right), enabling accurate 3D trajectory reconstruction (top).

Conventional vision-based SLAM algorithms face significant limitations. Standard cameras, the primary sensing modality in most pipelines, struggle with fast motion [8] or sudden illumination changes due to motion blur, limited dynamic range, and high latency [9]. These shortcomings lead to unreliable localization in challenging environments where drones are most critical, such as dark, cluttered, or dynamic settings like collapsed structures or industrial sites. Such limitations highlight the need for alternative sensing approaches to ensure robust navigation in these demanding scenarios [10].

To address the mentioned limitations, event cameras provide a biologically inspired sensing alternative. Unlike conventional cameras, which capture frames at fixed intervals, event cameras asynchronously detect per-pixel brightness changes with microsecond latency and high dynamic range [11]. Consequently, event cameras are ideal for fast motion, HDR, and low-light conditions. However, their sparse and asynchronous output is fundamentally different from conventional RGB images, posing unique challenges for SLAM. In particular, the lack of dense grid-structured data leads to a representation mismatch, hindering the use of standard feature extraction and matching techniques [11], [12].

Additionally, the microsecond-level timestamps of events introduce asynchronous fusion difficulties when integrating with frame-based sensors or IMUs, requiring specialized synchronization and motion-compensation strategies as seen in [13], [14].

This work presents Edged USLAM¹ as illustrated in Fig. 1, a hybrid visual–inertial SLAM system that extends USLAM [13] by introducing a lightweight, edge-aware front-end tailored for event data. The primary motivation is to reduce reliance on conventional RGB frames by enhancing the quality of event imagery itself, enabling robust localization even when standard intensity images are unavailable. Our method combines motion-compensated event frames with efficient edge filtering and grid-based oriented FAST and rotated BRIEF (ORB) feature tracking, improving robustness in low-light, high-speed, and low-featured scenes. Unlike prior methods that degrade in low-gradient or low-event conditions [15], [16], the proposed pipeline enhances landmark consistency through stronger spatial saliency.

The proposed algorithm, Edged USLAM, is validated on public benchmarks and real-time UAV flights with motion capture ground truth. We introduce a new aerial dataset that features diverse motion and illumination conditions for a realistic evaluation. The experimental results presented show superior localization performance compared to ORB-SLAM3 and USLAM, achieving robustness in dark environments and during aggressive maneuvers.

The contributions of this paper are as follows:

- A novel edge-aware front-end for event-based visual–inertial SLAM is designed, improving robustness under low-light and fast-motion conditions.
- Motion-compensated event frames are integrated with lightweight edge filtering, which is enhanced by a depth estimation module and grid-based ORB tracking to enhance landmark consistency.
- Real-time performance is demonstrated on a drone platform equipped with onboard computation, showing resilience in cluttered indoor environments.
- A new aerial dataset with synchronized events, IMU, and ground truth trajectories is provided, enabling future benchmarking in GPS-denied scenarios.

The remainder of the paper is organized as follows. Section II reviews related work on event-based odometry and SLAM. Section III presents the Edged USLAM framework, detailing motion compensation, edge-aware feature extraction, and depth prior integration. Section IV describes the experimental setup and datasets, followed by quantitative and qualitative evaluations in Section V. Finally, Section VI concludes with a discussion of limitations and future research directions.

II. RELATED WORK

Visual SLAM has steadily evolved, beginning with filter-based monocular systems like MonoSLAM [17] and

advancing to modern multi-sensor frameworks such as ORB-SLAM3 [18] and VINS-Mono [19]. Existing approaches can be broadly categorized into three paradigms: feature-based, direct, and learning-based methods. Feature-based pipelines extract sparse landmarks and remain the most widely adopted due to their robustness, computational efficiency, and suitability for real-time robotics [20], [21]. Their main weakness lies in the reliance on frame-based cameras, which often degrade under motion blur and low-light conditions. Direct methods such as LSD-SLAM [22] and DSO [23] exploit pixel-level photometric consistency, enabling dense reconstructions without explicit feature extraction. However, they remain highly sensitive to changes in illumination and camera calibration errors. Learning-based methods, including VINet [24] and CNN-SLAM [25], leverage neural networks to infer depth or poses directly from raw images. While they offer strong adaptability and scene understanding, their performance is hindered by generalization issues, data requirements, and limited interpretability [26]. Given these limitations across all paradigms, event-based vision has emerged as a promising alternative, offering high temporal resolution, low latency, and exceptional dynamic range, making it particularly well-suited for fast motion and challenging illumination conditions.

Early efforts such as USLAM [13] marked a turning point by tightly coupling asynchronous events, standard frames, and IMU measurements within a single optimization pipeline. By generating motion-compensated event frames using depth from tracked landmarks, it enabled parallel feature tracking on both modalities and robust performance under aggressive motion. Building on this, PL-EVIO [27] incorporated line features from event time surfaces to better exploit structured environments like corridors, while ESVO [28] departed from feature-based tracking and proposed a direct stereo formulation. The latter eliminated explicit event matching and enabled semi-dense reconstructions even in extremely low-light settings. More recently, EVI-SAM [29] combined feature-based and direct residuals in a sliding-window framework, achieving accurate 6-DoF tracking and dense mapping by jointly exploiting event maps, time surfaces, and inertial cues.

A complementary direction is learning-based event SLAM. Methods such as DEIO [30] and DEVO [31] integrate neural depth estimation into the odometry pipeline, improving metric scale recovery and robustness in texture-poor or repetitive environments. This trend highlights how depth priors can act as stabilizing cues to compensate for the sparsity of event data. In this trajectory, our Edged USLAM is positioned as a lightweight yet robust system. While not fully learning-based, it introduces two key ideas: an edge-aware front-end to enhance motion-compensated event frames, and a compact depth prior module that reduces scale drift. Together, these design choices improve landmark consistency and enable reliable tracking in visually degraded, GPS-denied environments.

¹Project page at: <https://sebnem-byte.github.io/edged-uslam.github.io/>

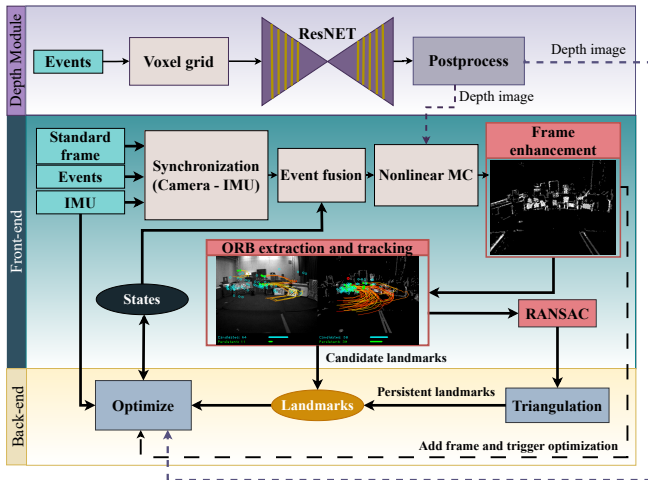


Fig. 2. Overview of the proposed Edged USLAM methodology. Events, standard frames, and IMU data are synchronized and fused with nonlinear motion compensation, followed by frame enhancement and ORB-based feature tracking. Landmarks are refined via RANSAC and triangulation, while depth priors from the ResNet module support consistent state estimation and optimization in the back-end.

III. METHODOLOGY

The proposed Edged USLAM architecture (Fig. 2) builds upon the original Ultimate SLAM framework [13], which fuses events, standard frames, and IMU data in a tightly coupled pipeline. We retain the synchronization strategy and augmented event packet generation, but enhance both the motion model and event preprocessing. Events are warped using nonlinear pose interpolation with $SE(3)$ transformations. Warped events are accumulated into sharper frames and enhanced via contrast and edge filters before ORB-based keypoints are extracted. In parallel, the depth module predicts dense event-based depth maps, from which scene-level priors are extracted. These priors support motion compensation in the front-end and are fused as soft constraints in the back-end optimization, improving scale consistency while maintaining real-time performance.

A. Nonlinear motion compensation

Event cameras output asynchronous brightness changes, which are grouped into overlapping event packets synchronized with IMU data. Let N denote the number of events in a packet and $i \in \{1, \dots, N\}$ be the event index. The timestamp of index i is denoted by t_i . To compensate for motion, the camera pose at each event timestamp t_i is obtained by interpolating between the start and end poses on $SE(3)$:

$$\mathbf{T}_i = \mathbf{T}_0 \cdot \exp(\alpha_i \cdot \log(\mathbf{T}_1 \mathbf{T}_0^{-1})), \quad \alpha_i = \frac{t_i - t_0}{t_1 - t_0}, \quad (1)$$

where $\mathbf{T}_0, \mathbf{T}_1 \in SE(3)$ denote the camera poses at the start and end of the event packet, and α_i is the normalized interpolation factor at time t_i .

Each event is represented as $e_i = (\mathbf{x}_i, t_i, p_i)$, where $\mathbf{x}_i = (x_i, y_i)$ is the pixel location, t_i the timestamp, and $p_i \in \{-1, +1\}$ the polarity. With an estimated depth d_i and camera intrinsic matrix \mathbf{K} , the event is lifted to 3D, warped

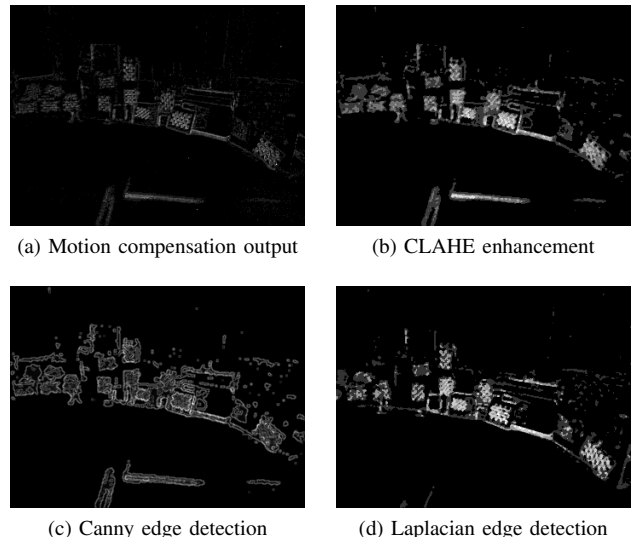


Fig. 3. Comparison of event frame enhancement methods. (a) Motion compensation sharpens asynchronous event frames into a more stable representation. (b) CLAHE improves local contrast under varying illumination. (c) Canny edge detection extracts fine-grained structural contours. (d) Laplacian filtering highlights high-frequency edges and textures.

with \mathbf{T}_i , and reprojected to obtain the motion-compensated pixel:

$$\hat{\mathbf{x}}_i = \pi(\mathbf{T}_i d_i \mathbf{K}^{-1} [x_i, y_i, 1]^\top) - \Delta \mathbf{x}_{\text{last}}, \quad (2)$$

where $\pi(\cdot)$ denotes the camera projection function and $\Delta \mathbf{x}_{\text{last}}$ is the previous alignment correction.

This depth-aware motion compensation suppresses temporal smearing and yields sharper event frames, enabling more reliable feature extraction and tracking. The nonlinear warping in $SE(3)$ is essential to accurately compensate rotational and translational motion, ensuring geometrically consistent reprojection.

B. Event image enhancement

Event cameras produce temporally precise but spatially sparse and noisy data. To enable reliable feature extraction, we introduce an edge-enhanced preprocessing pipeline after motion compensation. First, a Gaussian blur reduces isolated noise, followed by contrast limited adaptive histogram equalization (CLAHE) to improve local contrast. An optional sharpening step further refines structural details. Subsequently, edge detectors (Canny, Laplacian, or Sobel) extract contours, with morphological erosion used when finer localization is required. The resulting edge map is blended with the original motion-compensated frame, yielding sharper and more structured event images. This representation improves the repeatability and localization of ORB keypoints, particularly in low-contrast or low-motion conditions.

The complete pipeline is outlined in Algorithm 1 and Fig. 3 shows an illustrative example of the enhancement stages, and compares the outcomes of different edge extraction strategies.

Algorithm 1 Event Frame Enhancement

```
1: function ENHANCEEVENTFRAME( $I_{\text{event}}$ )
2:    $I_{\text{blur}} \leftarrow \text{GaussianBlur}(I_{\text{event}}, \sigma)$ 
3:    $I_{\text{CLAHE}} \leftarrow \text{CLAHE}(I_{\text{blur}})$ 
4:   if contrast_sharpening is enabled then
5:      $I_{\text{enh}} \leftarrow I_{\text{CLAHE}} + \lambda \cdot (I_{\text{CLAHE}} - I_{\text{blur}})$ 
6:   else
7:      $I_{\text{enh}} \leftarrow I_{\text{CLAHE}}$ 
8:   if method == “Canny” then
9:      $I_{\text{edge}} \leftarrow \text{Canny}(I_{\text{enh}})$ 
10:  else if method == “Laplacian” then
11:     $I_{\text{edge}} \leftarrow \text{Laplacian}(I_{\text{enh}})$ 
12:  else if method == “Sobel” then
13:     $I_{\text{edge}} \leftarrow \text{Sobel}(I_{\text{enh}})$ 
14:  if edge_thinning is enabled then
15:     $I_{\text{edge}} \leftarrow \text{Erode}(I_{\text{edge}}, B)$ 
16:  return  $\alpha \cdot I_{\text{event}} + \beta \cdot I_{\text{edge}}$ 
```

C. Feature extraction and tracking

Unlike USLAM, which relies on FAST features, and other event-based methods that often employ conventional keypoint detectors, our system adopts a grid-based ORB strategy to enforce spatial uniformity. By dividing the event frame into an $R \times C$ grid, where R and C denote the number of grid rows and columns, and retaining the strongest keypoint per cell, the method mitigates feature clustering in textured regions and ensures more balanced geometric coverage, leading to improved robustness in low-texture and high-speed scenarios. Let $\mathcal{K} = \{k_1, \dots, k_N\}$ be the set of all detected ORB keypoints in the image $I \in \mathbb{R}^{H \times W}$, where H and W denote the image height and width, and each keypoint $k_i = (\mathbf{x}_i, r_i)$ consists of a position \mathbf{x}_i and a response r_i . We denote by $\mathcal{K}_{i,j}$ the subset of keypoints within grid cell $G_{i,j}$, and by $k_{i,j}^*$ the selected keypoint with the maximum response in that cell. The final selection \mathcal{K}^* is then obtained as the union of all $k_{i,j}^*$ across the grid.

$$k_{i,j}^* = \arg \max_{k \in \mathcal{K}_{i,j}} \text{response}(k), \quad \mathcal{K}^* = \bigcup_{i=1}^R \bigcup_{j=1}^C \{k_{i,j}^*\}. \quad (3)$$

where $k_{i,j} \subset \mathcal{K}$ denotes keypoints within grid cell $G_{i,j}$. Feature tracking is performed with a pyramidal KLT optical flow. Each landmark is projected into the current frame using the relative pose, and the tracker refines its position by minimizing the photometric error over a local patch.

D. Depth module integration in Edged-USLAM

To enable real-time metric scale recovery, we integrate the learning-based event-to-depth model (e2depth) [32] into the Edged USLAM framework. The network is used in a lightweight configuration, taking event-voxel grids as input and providing a scene-depth estimate. For deployment, we implement a simplified ROS node that omits non-critical modules, makes the recurrent state optional, and applies fast voxel cropping and normalization, reducing latency for

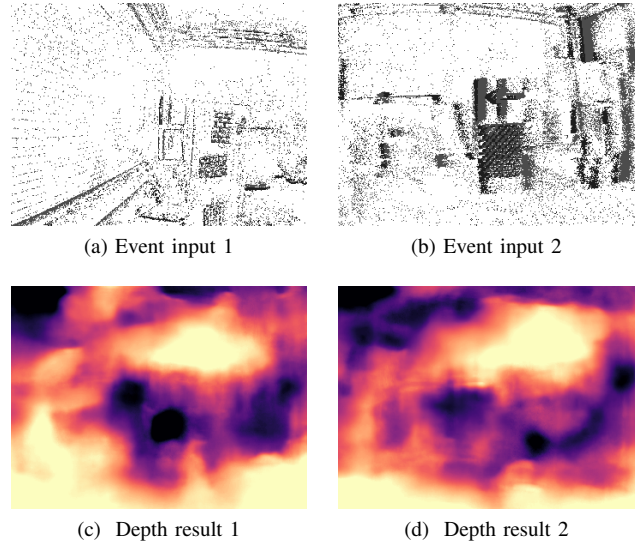


Fig. 4. Real-time demonstration of event-based depth estimation. Top row: sample event inputs (a–b). Bottom row: corresponding coarse depth estimation results (c–d) computed by a lightweight model for rapid inference, intended not for accurate reconstruction but for rapid inference of approximate object distances.

online inference. Specifically, we compute the mean depth \bar{d}_t within a central ROI and stabilize it with exponential smoothing:

$$d_t = \alpha \bar{d}_t + (1 - \alpha) d_{t-1}, \quad d_t \in (d_{\min}, d_{\max}), \quad (4)$$

where α is a smoothing factor and (d_{\min}, d_{\max}) are validity bounds. The resulting d_t is published as a scene-depth prior and fused with landmark-based triangulation inside the front-end.

In the back-end, the ROI-based scene depth d_t is converted to inverse depth ρ_t^{ROI} and used as a soft prior in the factor graph:

$$E_{\text{depth}} = \|\bar{\rho}_t - s \rho_t^{\text{ROI}}\|_{\Sigma_t}^2, \quad (5)$$

where $\bar{\rho}_t$ is the optimized average inverse depth in the current keyframe, s is a global scale parameter, and Σ_t encodes confidence.

Figure 4 illustrates the intermediate depth maps. Although coarse, these lightweight predictions are geometry-aware and enable real-time motion compensation and scale consistency.

IV. EXPERIMENTS

The proposed SLAM system was evaluated in real-world conditions using a quadrotor platform equipped with a PX4-based Pixhawk 6X flight controller. The onboard perception stack runs on an NVIDIA Jetson Orin NX mission computer, which processes data from a DAVIS 346 event camera for event-based visual-inertial odometry. Localization performance was assessed by comparing the SLAM-estimated poses against ground truth provided by a motion capture system, operating at high frequency with full 6-DoF accuracy.

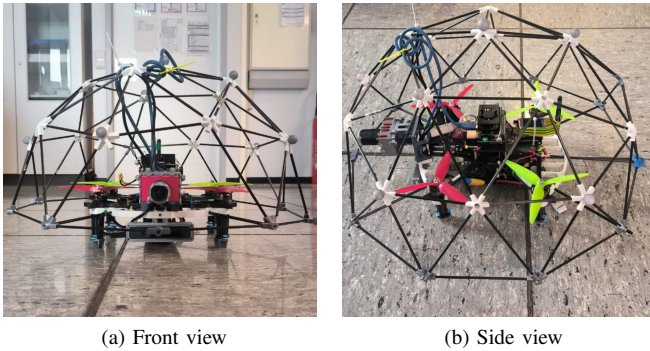


Fig. 5. TESTUDO an UAV platform equipped with a collision-tolerant protective cage and multi-sensor payload (event camera, tracking sensor, and IMU) for real-time SLAM evaluation in confined and hazardous environments

As shown in Fig. 5, the TESTUDO is designed as a collision-tolerant quadrotor for operation in confined and hazardous environments. The lightweight geodesic cage protects both the drone and its payload during impacts, allowing safe navigation in cluttered spaces.

A. Evaluation on benchmark sequences

For a detailed evaluation of our system, we considered sequences from multiple publicly available event-based datasets. Specifically, representative sequences from the Event-Camera Dataset and Simulator [33] were used to cover challenging motion and lighting conditions, including aggressive 6-DoF motion, dynamic objects, and pure translational motion. In addition, sequences from the HKU Event-based VO/VIO/SLAM dataset [34] were included to validate performance in indoor and fast-motion UAV scenarios. We further tested on the recently introduced VECtor dataset [35], which provides large-scale and diverse recordings under real-world conditions. Together, these datasets offer broad coverage of motion dynamics, illumination changes, and scene complexity, enabling a thorough evaluation of robustness and generalization.

The position RMSE results on the Event-Camera Dataset and Simulator are reported in Table I. Edged USLAM achieves the lowest average error (0.14 m) and provides the best accuracy in four out of seven sequences. Compared to USLAM [13], it consistently reduces drift, with the largest improvements observed in dynamic sequences. The improvement indicates that adding edge-aware event processing and

TABLE I
POSITION RMSE (M) ON THE EVENT-CAMERA DATASET AND SIMULATOR DATASET WITH EVENT-BASED METHODS.

Sequence	Edged(ours)	EVI-SAM [29]	USLAM [13]	PL-EVIO [27]	EKLT-VIO [36]
boxes_6dof	0.14	0.16	0.30	0.21	0.84
boxes_trans	0.16	0.11	0.27	0.06	0.48
hdr_boxes	0.22	0.13	0.37	0.10	0.46
dynamics_trans	0.10	0.30	0.18	0.24	0.40
poster_trans	0.07	0.34	0.12	0.54	0.35
hdr_poster	0.18	0.15	0.31	0.12	0.65
dynamic_6DOF	0.12	0.27	0.19	0.48	0.79
Average	0.14	0.17	0.23	0.23	0.56

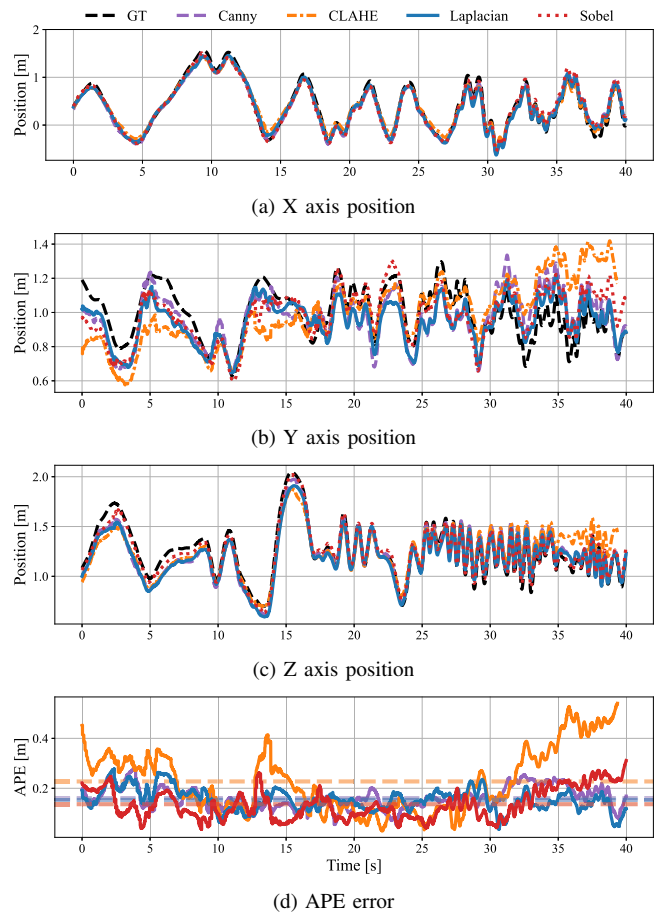


Fig. 6. Comparison of four event frame enhancement methods (Canny, CLAHE, Laplacian, Sobel) on the boxes_6dof sequence. (a–c) Estimated X, Y, and Z axis positions compared with ground truth (black dashed). (d) Absolute position error (APE) over time. Sobel achieves the lowest mean error, followed by Canny and Laplacian, while CLAHE yields the highest errors. Mean APE values: Sobel is 0.143 m, Canny is 0.167 m, Laplacian is 0.175 m, and CLAHE is 0.227 m.

depth cues provides stronger geometric consistency than USLAM’s original fusion of frames, events, and inertial data.

Compared to PL-EVIO [27], which is purely event-driven, our method performs better in dynamic and cluttered environments, where the absence of appearance and depth information causes larger drift for PL-EVIO. EVI-SAM [29] benefits from semantic priors and yields stable results in static scenes, but remains sensitive to fast motion and moving objects, where Edged USLAM maintains more reliable tracking. Finally, EKLT-VIO [36] shows the weakest performance overall due to the limited geometric constraints of event-based KLT tracking, while our method reduces drift substantially across all sequences.

To assess the impact of the edge-enhancement module, we compared four filtering strategies—Canny, CLAHE, Laplacian, and Sobel—on the boxes_6dof sequence. As shown in Figure 6, CLAHE, while primarily a local contrast enhancement technique rather than an edge detector, increases the number of detectable features by amplifying intensity differences. At the same time, it also amplifies high-frequency noise, which reduces stability under motion. Laplacian filter-

TABLE II
POSITION RMSE (M) ON HKU AND VECTOR DATASETS WITH CLASSICAL AND LEARNING-BASED METHODS.

Dataset	Sequence	Hybrid	Standard (Non-learning)				Learning-based	
		Edged(ours)	EVI-SAM [29]	USLAM [13]	PL-EVIO [27]	EVO [12]	DPVO [37]	DEVO [31]
HKU	agg_tran	0.13	0.17	0.59	0.07	<i>failed</i>	0.12	0.06
	agg_rot	0.25	0.24	3.14	0.23	<i>failed</i>	0.28	0.05
	agg_walk	0.19	0.42	2.00	0.42	<i>failed</i>	<i>failed</i>	0.9
	hdr_circle	0.49	0.13	1.32	0.14	<i>failed</i>	1.19	0.39
	hdr_slow	0.13	0.11	2.80	0.13	<i>failed</i>	0.36	0.08
VECTor	corner_slow	0.39	2.50	4.83	2.10	4.83	<i>failed</i>	0.59
	robot_norm	0.14	0.67	1.81	0.68	<i>failed</i>	0.22	0.17
	robot_fast	0.56	0.22	1.65	0.14	<i>failed</i>	0.78	0.13
	sofa_fast	0.67	0.98	2.54	0.17	<i>failed</i>	0.60	0.38
	sofa_norm	0.59	0.19	5.74	0.19	<i>failed</i>	0.22	0.22
Average		0.35	0.56	2.64	0.42	0.47	0.47	0.29

TABLE III
ABSOLUTE POSE ERROR (APE, M) AND RANGE (MIN-MAX) ACROSS MOTION CATEGORIES AND ILLUMINATION CONDITIONS ON OUR DATASET.

Category	Sequence	Edged USLAM (ours)			ORB-SLAM3 [18]			USLAM [13]		
		APE	Max	Min	APE	Max	Min	APE	Max	Min
Motion	circle	0.26	0.64	0.02	0.21	2.1	0.01	0.48	1.1	0.1
	aggressive	0.19	0.62	0.015	0.57	1.75	0.05	0.35	0.75	0.12
	aggressive-lowlit	0.28	0.58	0.05	<i>failed</i>	–	–	0.32	1.0	0.02
	square	0.32	0.72	0.06	0.19	0.96	0.035	0.71	2.1	0.10
	line	0.21	0.66	0.03	0.21	0.61	0.036	0.47	1.5	0.09
	turned	0.38	0.86	0.10	<i>failed</i>	–	–	0.55	1.35	0.10
Illumination	low-lit	0.34	0.72	0.014	<i>failed</i>	–	–	0.85	1.6	0.15
	30% lit	0.54	1.10	0.12	<i>failed</i>	–	–	0.58	2.1	0.18
	60% lit	0.45	0.96	0.08	1.7	3.8	0.06	0.54	1.17	0.06
	dynamic HDR	0.76	1.56	0.11	<i>failed</i>	–	–	1.49	3.2	0.25
	constant HDR	0.52	1.45	0.09	1.45	3.6	0.21	0.71	1.69	0.2
Average		0.38			0.72			0.64		

ing similarly produces a noisier representation, but provides more coherent edge structures than CLAHE, though still fewer reliable features compared to gradient-based operators. In contrast, Sobel and Canny produce sharper and cleaner edge responses with lower noise sensitivity, resulting in more stable feature tracking and improved trajectory estimates.

Table II summarizes results on the HKU and VECTor benchmarks. Edged USLAM achieves an average RMSE of 0.35 m, substantially improving over USLAM [13] (2.64 m) and showing more consistent accuracy than purely event-based pipelines such as PL-EVIO [27]. On the HKU sequences, which emphasize aggressive motion, the proposed method maintains stable performance across all trajectories, whereas USLAM degrades significantly under fast translations and rotations. While event-only or learning-based methods occasionally reach lower errors in specific cases (e.g., PL-EVIO on *agg_tran*, DEVO on *agg_rot*), they lack robustness across the full set of sequences. Furthermore, event-based methods like DEVO exhibit higher drift than non-event-based methods such as DPVO under insufficient motion, due to the inherent lack of event data.

On the VECTor dataset, which combines both slow and fast motion, Edged USLAM consistently outperforms USLAM and EVI-SAM, particularly in structured low-motion sequences such as *sofa_norm* and *sofa_fast*, where conventional SLAM pipelines suffer from sparse event information. At the same time, our method remains competitive under fast motion (e.g., *robot_fast*) thanks to edge-enhanced event frames and depth priors, which support more reliable landmark initialization. Compared to learning-based baselines (DPVO, DEVO), Edged USLAM offers slightly higher RMSE on average, but has the key advantage of requiring no pretraining or domain adaptation, making it directly applicable to unseen environments in real time.

B. Real-time flight evaluation

All real-time evaluations were conducted on a quadrotor platform in an indoor test arena equipped with a motion capture system for ground-truth trajectory acquisition. The dataset collected during these flights covers diverse motion profiles and illumination conditions, including takeoff, landing, and free maneuvers in environments with a re-

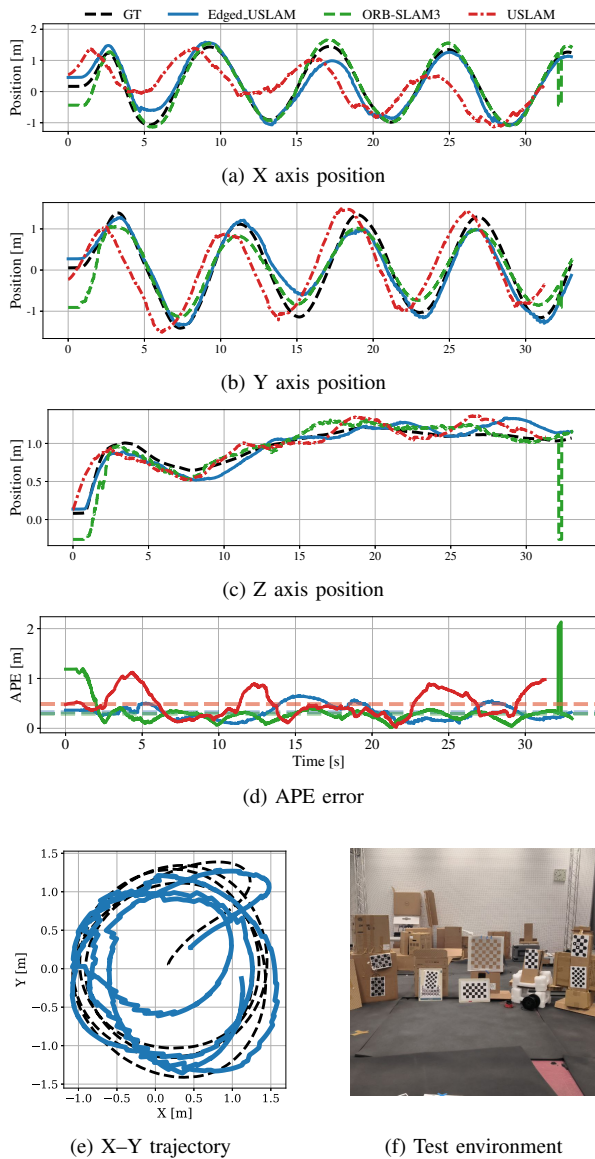


Fig. 7. Real-time circular motion tracking performance of different SLAM methods in well-lit conditions. Top: estimated axis positions (X, Y, Z) compared with ground truth (black dashed) and absolute position error (APE) over time. Bottom: XY trajectory overlay and test environment. Mean APE values: Edged USLAM is 0.26 m, ORB-SLAM3 is 0.21 m, and USLAM is 0.486 m.

alistic number of visual features (neither highly textured nor texture-sparse). The motion profiles include circular flights, aggressive manual trajectories with roll and lateral shifts, low-light aggressive flights, square trajectories through obstacles, straight-line paths, and aggressive yaw rotations. For illumination tests, adjustable room lighting was systematically varied: low light, 30% lit, and 60% lit. Additionally, a dynamic HDR condition was introduced by randomly switching lateral light sources on and off, while a constant HDR setting was created by opening a side window to let sunlight enter from one direction.

Table III summarizes the APE statistics across these conditions, which are part of our dataset. ORB-SLAM3 achieves strong performance in repetitive scenes (e.g., circle), benefit-

ing from loop closure and back-end optimization. However, its monocular setup introduces significant scale drift, and it fails in low-light and HDR scenarios where standard frames lose texture. It also struggles with aggressive yaw rotations, often losing tracking due to delayed relocalization. USLAM performs better than ORB-SLAM3 under low illumination and dynamic lighting, thanks to its reliance on events. Nonetheless, it suffers from higher drift and instability during aggressive maneuvers, with APE values frequently exceeding 0.5–1.0 m. In contrast, the proposed algorithm consistently improves over USLAM across both motion and illumination categories. The edge-enhanced event preprocessing provides more reliable feature generation, resulting in reduced APE (e.g., 0.19 m under aggressive motion, compared to 0.35 m with USLAM). Also, Edged USLAM maintains robustness under conditions where ORB-SLAM3 fails (low-light, dynamic HDR). Notably, the error under the low-lit condition (0.34 m) is lower than in a moderately lit scenario, such as the 30% lit case (0.54 m); this is because under this extreme low-light condition, the system effectively disregards the unreliable RGB stream, whereas under moderate lighting, the integration of noisy visual features introduces residual errors.

Figure 7 shows the XYZ position estimates and APE analysis for three methods. ORB-SLAM3 achieved the lowest APE (0.21 m), mainly due to its loop closure mechanism. However, it suffered from intermittent tracking failures under fast motion and lighting changes, visible as trajectory discontinuities and APE spikes. USLAM showed consistent lag and drift along the X-axis (depth), with a higher APE of 0.486 m. This was likely caused by delayed pose propagation and limited feature observability. In contrast, Edged USLAM achieved more stable tracking, particularly along the Y and Z axes, with reduced temporal delay. Its mean APE was 0.26 m. While some depth inaccuracies remained, the edge-enhanced front-end allowed better data association and smoother trajectories than USLAM.

V. CONCLUSION AND FUTURE WORK

The main contribution of this work is Edged USLAM, a hybrid visual-inertial SLAM framework that augments Ultimate SLAM with edge-aware event frame processing and lightweight depth priors. Through evaluations on the Event-Camera, HKU, and VECTOR benchmarks as well as real UAV flights, the experiments demonstrate that different SLAM methods excel under different motion and illumination conditions. Event-only pipelines such as PL-EVIO and learning-based methods like DEVO can achieve lower errors in aggressive maneuvers or extreme HDR scenes, highlighting their robustness to fast dynamics and illumination changes. In contrast, Edged USLAM consistently provides more stable tracking and reduced drift in structured and long-duration trajectories, where it outperforms classical baselines including ORB-SLAM3 and USLAM, while maintaining reliable operation in real UAV deployments. The evaluations reveal that while IMU-RGB pipelines remain strongest in slow-motion settings, our edge-aware event processing provides

advantages in HDR environments, and the depth-aware event frames improve both accuracy and landmark initialization during aggressive maneuvers and in dynamic environments.

In addition to releasing a drone-based dataset with diverse motion and illumination conditions, future work will focus on improving depth estimation beyond the current lightweight ROI-based design, potentially leveraging event-IMU fusion to reduce dependence on RGB frames. Further enhancements to the back-end optimization and the integration of loop closure are also planned to strengthen global consistency and mapping quality.

REFERENCES

- [1] M. Montemerlo, J. Becker, S. Bhat, H. Dahlkamp, D. Dolgov, S. Ettinger, D. Haehnel, T. Hilden, G. Hoffmann, B. Huhnke *et al.*, “Junior: The stanford entry in the urban challenge,” in *Journal of Field Robotics*, vol. 25, no. 9. Wiley Online Library, 2008, pp. 569–597.
- [2] G. Klein and D. Murray, “Parallel tracking and mapping for small AR workspaces,” in *Proc. IEEE/ACM Int. Symp. Mixed Augmented Reality (ISMAR)*, Nara, Japan, 2007, pp. 225–234.
- [3] C. Stachniss and W. Burgard, “Mobile robot mapping and localization in non-static environments,” in *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, vol. 3, July 2005, pp. 1324–1329.
- [4] H. X. Pham, H. I. Ugurlu, J. Le Fevre, D. Bardakci, and E. Kayacan, “Chapter 15 - Deep learning for vision-based navigation in autonomous drone racing,” in *Deep Learning for Robot Perception and Cognition*, A. Iosifidis and A. Tefas, Eds. Academic Press, 2022, pp. 371–406.
- [5] A. Kleiner and C. Dornhege, “Real-time localization and elevation mapping within urban search and rescue scenarios,” *Journal of Field Robotics*, vol. 24, no. 8-9, pp. 723–745, 2007.
- [6] A. Dey, O. A. Ejohwomu, and P. W. Chan, “Sustainability challenges and enablers in resource recovery industries: A systematic review of the ship-recycling studies and future directions,” *Journal of Cleaner Production*, vol. 329, p. 129787, 2021.
- [7] T. Dang, F. Mascariich, S. Khattak, H. Nguyen, H. Nguyen, S. Hirsh, R. Reinhart, C. Papachristos, and K. Alexis, “Autonomous search for underground mine rescue using aerial robots,” in *2020 IEEE Aerospace Conference*, 2020, pp. 1–8.
- [8] K. F. Andersen, H. X. Pham, H. I. Ugurlu, and E. Kayacan, “Event-based navigation for autonomous drone racing with sparse gated recurrent network,” in *2022 European Control Conference (ECC)*, 2022, pp. 1342–1348.
- [9] E. Mueggler, B. Huber, and D. Scaramuzza, “Event-based, 6-dof pose tracking for high-speed maneuvers,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 2761–2768.
- [10] O. Álvarez Tuñón, Y. Brodskiy, and E. Kayacan, “Monocular visual simultaneous localization and mapping: (r)evolution from geometry to deep learning-based pipelines,” *IEEE Access*, vol. 8, pp. 75 850–75 870, 2020.
- [11] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, “Event-based vision: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 154–180, 2022.
- [12] 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 Robotics and Automation Letters*, vol. 2, no. 2, pp. 593–600, 2017.
- [13] A. R. 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 Robotics and Automation Letters*, vol. 3, no. 2, pp. 994–1001, 2018.
- [14] D. Gehrig, M. Gehrig, J. Hidalgo-Carrió, and D. Scaramuzza, “DSEC: A stereo event camera dataset for driving scenarios,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4947–4954, 2021.
- [15] Y. He, J. Zhao, Y. Guo, W. He, and K. Yuan, “PL-VIO: Tightly-coupled monocular visual-inertial odometry using point and line features,” *Sensors*, vol. 18, no. 4, 2018. [Online]. Available: <https://www.mdpi.com/1424-8220/18/4/1159>
- [16] 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.
- [17] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, “MonoSLAM: Real-time single camera SLAM,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 6, pp. 1052–1067, 2007.
- [18] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [19] T. Qin, P. Li, and S. Shen, “VINS-Mono: A robust and versatile monocular visual-inertial state estimator,” *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [20] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [21] A. M. Barros, M. Michel, Y. Moline, G. Corre, and F. Carrel, “Comprehensive survey of visual slam algorithms,” *Robotics*, vol. 23, no. 1, p. 123, 2022.
- [22] J. Engel, T. Schöps, and D. Cremers, “LSD-SLAM: Large-scale direct monocular SLAM,” in *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer, 2014, pp. 834–849.
- [23] J. Engel, V. Koltun, and D. Cremers, “Direct sparse odometry,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 3, pp. 611–625, 2018.
- [24] R. Clark, S. Wang *et al.*, “VINet: Visual-inertial odometry as a sequence-to-sequence learning problem,” in *AAAI*, 2017.
- [25] K. Tateno, F. Tombari, I. Laina, and N. Navab, “CNN-SLAM: Real-time dense monocular SLAM with learned depth prediction,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017, pp. 6565–6574.
- [26] O. Álvarez Tuñón, Y. Brodskiy, and E. Kayacan, “Loss it right: Euclidean and riemannian metrics in learning-based visual odometry,” 2024. [Online]. Available: <https://arxiv.org/abs/2401.05396>
- [27] W. Guan, P. Chen, Y. Xie, and P. Lu, “PL-EVIO: Robust monocular event-based visual inertial odometry with point and line features,” *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 4, pp. 6277–6293, 2024.
- [28] Y. Zhou, G. Gallego, and S. Shen, “Event-based stereo visual odometry,” *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1433–1450, 2021.
- [29] W. Guan, P. Chen, H. Zhao, Y. Wang, and P. Lu, “EVI-SAM: Robust, real-time, tightly-coupled event-visual-inertial state estimation and 3D dense mapping,” *Advanced Intelligent Systems*, vol. 6, no. 12, p. 2400243, 2024.
- [30] W. Guan, F. Lin, P. Chen, and P. Lu, “DEIO: Deep event inertial odometry,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2025.
- [31] 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.
- [32] J. Hidalgo-Carrió, D. Gehrig, and D. Scaramuzza, “Learning monocular dense depth from events,” in *Proceedings of the IEEE International Conference on 3D Vision (3DV)*. IEEE, 2020, pp. 534–542.
- [33] E. Mueggler, H. Rebecq, G. Gallego, T. Delbruck, and D. Scaramuzza, “The event-camera dataset and simulator: Event-based data for pose estimation, visual odometry, and SLAM,” *International Journal of Robotics Research*, vol. 36, no. 2, pp. 142–149, 2017.
- [34] W. Guan, P. Chen, and P. Lu, “HKU Event-based VO/VIO/SLAM Dataset,” https://github.com/arclab-hku/Event_based_VO-VIO-SLAM, 2025, mIT license; accessed: Sep. 10, 2025.
- [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 Robotics and Automation Letters*, 2022, preprint accepted June 2022.
- [36] F. Mahlknecht, D. Gehrig, J. Nash, F. M. Rockenbauer, B. Morrell, J. Delaune, and D. Scaramuzza, “Exploring event camera-based odometry for planetary robots,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8651–8658, 2022.
- [37] Z. Teed, L. Lipson, and J. Deng, “Deep patch visual odometry,” in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS ’23. Curran Associates Inc., 2023.