

LEAR: Learning Edge-Aware Representations for Event-to-LiDAR Localization

Kuangyi Chen Jun Zhang Yuxi Hu Yi Zhou Friedrich Fraundorfer

Abstract—Event cameras offer high-temporal-resolution sensing that remains reliable under high-speed motion and challenging lighting, making them promising for localization from LiDAR point clouds in GPS-denied and visually degraded environments. However, aligning sparse, asynchronous events with dense LiDAR maps is fundamentally ill-posed, as direct correspondence estimation suffers from modality gaps. We propose LEAR, a dual-task learning framework that jointly estimates edge structures and dense event–depth flow fields to bridge the sensing-modality divide. Instead of treating edges as a post-hoc aid, LEAR couples them with flow estimation through a cross-modal fusion mechanism that injects modality-invariant geometric cues into the motion representation, and an iterative refinement strategy that enforces mutual consistency between the two tasks over multiple update steps. This synergy produces edge-aware, depth-aligned flow fields that enable more robust and accurate pose recovery via Perspective-n-Point (PnP) solvers. On several popular and challenging datasets, LEAR achieves superior performance over the best prior method. The source code, trained models, and demo videos are made publicly available online¹.

I. INTRODUCTION

Visual localization aims to estimate the six degrees of freedom (6-DoF) pose of a camera within a reference map. It is a promising and widely adopted approach, as cameras are lightweight and cost-effective, making them particularly well-suited for deployment on compact robotic platforms such as Unmanned Aerial Vehicles (UAVs). However, standard cameras are susceptible to motion blur and lighting variations. Event cameras, a new class of bio-inspired sensors, offer compelling advantages—including high dynamic range, low latency, and asynchronous operation—which make them robust to such challenges. Consequently, event cameras have attracted growing interest as an alternative to standard cameras for visual localization.

A core challenge in event-based localization lies in establishing correspondences between events and reference maps, of which the signal or data representations are significantly different, with events usually being sparse but texture-aware, while point clouds typically being dense and geometrically informative. Traditional approaches often rely on task-specific reference maps, such as geo-tagged event

Kuangyi Chen, Jun Zhang, Yuxi Hu, and Friedrich Fraundorfer are with the Institute of Visual Computing, Graz University of Technology, 8010 Graz, Austria. E-mail: {kuangyi.chen, jun.zhang, yuxi.hu, friedrich.fraundorfer}@tugraz.at

Yi Zhou is with the Neuromorphic Automation and Intelligence Lab (NAIL), School of Artificial Intelligence and Robotics, Hunan University, Changsha 410012, China. E-mail: eeyzhou@hnu.edu.cn

¹<https://github.com/EasonChen99/LEAR>

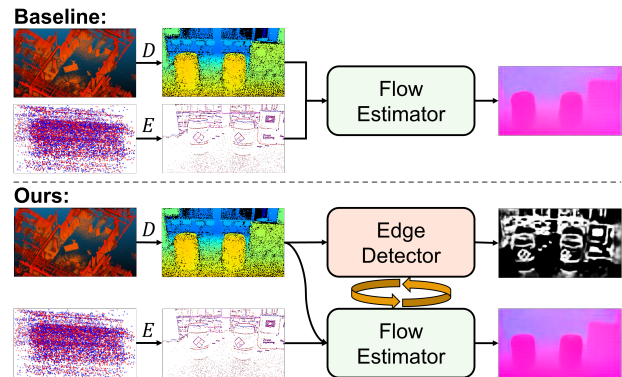


Fig. 1. *Our method (bottom)* integrates a flow estimator with an edge detector within a mutually reinforcing cycle, enabling more accurate 2D–3D matching (i.e., event–depth flow) compared to the flow-only baseline *EVLoc* [19] (top).

frames [1]–[3], handcrafted feature databases [4], [5], or learned descriptors [6]–[8]. While being effective for place recognition, these approaches typically yield coarse localization and often require large effort to build and maintain. To achieve more accurate localization, other methods construct 3D line maps [9] or semi-dense reconstructions [10]–[12], aligning events with projected map edges. However, their performance suffers under noise, occlusion, or weak texture. Photometric map-based methods [13], [14] attempt to match predicted intensity changes to observed events, but they lack view-dependent appearance modeling, struggling in dynamic or complex environments. Direct pose regression methods [15]–[18] offer efficient alternatives by implicitly encoding scene geometry within network parameters. However, they exhibit limited generalization to novel environments.

Unlike these task-specific representations, LiDAR point clouds offer a compelling alternative as reference maps: they are geometrically precise, robust to lighting, and often readily available via third-party providers. However, directly aligning events with LiDAR remains challenging due to their inherently different nature. *EVLoc* [19], to our knowledge, makes the first attempt to bridge this gap by formulating the problem as dense flow estimation between events and projected LiDAR points in 2D space. The flow estimates are then used to determine 2D-3D correspondences for camera pose recovery. While promising results are demonstrated, the approach functions largely as a black box, offering limited interpretability into the underlying mechanism.

Our key insight is that edge geometry forms a modality-invariant bridge between events and LiDAR: events are

triggered at intensity gradients, while depth discontinuities in LiDAR capture structural boundaries. Incorporating this shared structure into the representation can reduce the burden on cross-modal feature encoding and make correspondence estimation more robust. In our case, masking LiDAR point clouds using event-derived edge maps (as shown in Fig. 4, *Ground Truth Edge Map*) leads to significantly improved performance compared to using raw point clouds alone, as demonstrated by the *oracle* setting in Table I. However, generating such aligned edge masks requires knowledge of the relative pose—i.e., the associated correspondences we seek to estimate. This highlights a mutual dependency: accurate edge extraction depends on known correspondences, while good-quality correspondences require well-aligned, edge-aware encoded features.

Motivated by this insight, we introduce LEAR, a dual-task framework for event-based localization using LiDAR maps, as shown in Fig. 1. Given an initial pose guess—obtained through coarse global localization methods such as visual place recognition—LEAR refines it to the precise 6-DoF camera pose. LEAR couples a dense event–depth flow estimator with an edge detector in a mutually reinforcing architecture. A Cross-task Feature Fusion (CFF) module embeds modality-invariant edge cues into the motion estimation pipeline at the feature level, improving correspondence reliability across modalities. An Iterative Feature Refinement (IFR) module then recurrently updates both edge and flow features, enforcing geometric consistency and progressively improving prediction quality. This synergy yields edge-aware, depth-aligned flow fields that produce more stable and accurate pose recovery via a RANSAC-based PnP solver. Overall, our main contributions are summarized as follows:

- We propose a novel dual-task learning framework for event-based visual localization from LiDAR maps, which integrates a flow estimator and an edge detector, enabling edge-aware representations that enhance cross-modal consistency for event-LiDAR data alignment.
- We introduce the CFF and IFR modules to facilitate the learning of both task branches in a mutually reinforcing manner, improving localization and edge prediction.
- We conduct extensive experiments on public datasets of varying scenarios, and demonstrate the effectiveness of our proposed edge feature enhancement modules and state-of-the-art performance in localization accuracy.

II. RELATED WORK

Event-based visual localization has been approached from multiple perspectives. In this section, we review key developments along the mainstream directions.

Retrieval methods avoid cross-modal matching by converting event streams into queryable representations, such as images [1], [3], [4] or edge maps [20], enabling similarity-based matching against a reference database. Others construct handcrafted [5] or learned descriptors [6], [8], or encode implicit representations in network weights [7]. However, these approaches typically require costly database construction and

maintenance, and achieve only coarse localization, which can serve as an initialization step for our method.

Another line of work focuses on training neural networks to regress camera poses directly from events [15]–[18], encoding environmental information implicitly. For example, Nguyen et al. [15] estimate poses from short event bursts using a CNN and Stacked Spatial LSTM. Although these methods eliminate explicit maps, they suffer from limited scalability and poor generalization to unseen environments.

Pose refinement approaches assume a known 3D scene representation and an initial estimate of the event camera’s state. Based on the type of scene representation, these approaches can be broadly categorized into geometric and photometric methods. Geometric methods typically construct a 3D map enriched with edge features and optimize the camera pose by aligning events with these features. For instance, Yuan et al. [9] build a scene map consisting solely of vertical lines and achieve 3-DoF localization by matching 2D event lines with their 3D counterparts. More recent approaches [10]–[12] construct semi-dense maps using depth sensors or Structure from Motion (SfM) techniques, and estimate 6-DoF poses through polarity-aware geometric alignment between events and the semi-dense maps. However, since these methods rely heavily on explicit edge features, their performance degrades in scenes with noisy or weak texture. Photometric methods infer intensity changes from scene maps that contain color information [13], [14], [21]. For example, Bryner et al. [14] propose a method that estimates poses by aligning observed and predicted intensity change images, using a generative event model and a photometric 3D map. However, traditional photometric maps typically lack view-dependent appearance modeling, limiting their expressiveness under varying viewpoints. To address this, Liu et al. [21] introduce a Gaussian Splatting-based map, which is able to encode more complex ambient lighting effects. Nevertheless, rendering artifacts or inconsistencies in Gaussian Splatting can directly impact localization accuracy.

Compared to the aforementioned approaches, LiDAR maps offer dense, accurate, and lighting-invariant geometry, making them attractive for event-based localization. Additionally, LiDAR maps are often readily available from third-party providers, eliminating the need for task-specific scene construction. Despite this potential, few methods have addressed the challenge of cross-modal matching between events and LiDAR point clouds. [22], [23] are among the first attempts to handle this for event–LiDAR extrinsic calibration. However, they typically rely on task-specific markers or controlled environments, limiting their scalability. E2PNet [24] introduces a learning-based, image-like event feature representation and applies it to existing RGB-based methods [25], [26] for event-to-point cloud localization. However, it reports only coarse localization performance. The most closely related work with our method is EVLoc [19], which tackles alignment by directly learning dense flow fields between events and LiDAR maps. However, it doesn’t explicitly exploit modality-invariant features that can strengthen cross-modal correspondence. In contrast, we

propose a dual-task framework that explicitly targets this representational bottleneck by jointly learning an edge detector and a flow estimator in a mutually reinforcing manner, enabling consistency in feature space and significantly improving event–LiDAR localization robustness and accuracy.

III. PRELIMINARIES

In this section, we describe the common modules inherited from the baseline EVLoc [19], including data preprocessing and the flow estimation network.

1) **Data Preprocessing:** Given an initial camera pose \mathbf{T}_{init} , we transform the raw LiDAR points P_w into the camera coordinate system and project them onto image plane using a pinhole camera model to obtain the depth map I_D . Meanwhile, a structurally clear event image within a given time window is generated. The event image generation in [19] involves two steps: deblurring and denoising. In our implementation, we retain the deblurring step but omit denoising due to its high computational cost. Instead, we apply a Spatio-Temporal Contrast Filter and a Trail Filter from the Prophesee SDK [27] to preprocess the raw event stream. These filters effectively reduce both noise and event count. After filtering, we apply the event image generation algorithm from [19] to generate the final event image I_{EV} .

2) **Flow Estimation Network:** Same as [19], we formulate the cross-modal event–depth matching problem as a dense event–depth flow prediction task using a modified RAFT [28] network, trained with a masked L2 loss:

$$\mathcal{L}_{\text{flow}} = \frac{\sum_{u,v} \mathbf{1}[\mathbf{f}_{\text{gt}}(u,v) \neq 0] \|\mathbf{f}(u,v) - \mathbf{f}_{\text{gt}}(u,v)\|_2}{\sum_{u,v} \mathbf{1}[\mathbf{f}_{\text{gt}}(u,v) \neq 0]} + \epsilon, \quad (1)$$

where \mathbf{f} is the predicted event–depth flow. \mathbf{f}_{gt} denotes the ground-truth event–depth flow, obtained by projecting 3D points into the image plane using the ground truth pose \mathbf{T}_{gt} and the initial pose \mathbf{T}_{init} , and then computing the difference between the corresponding 2D projections. $\mathbf{1}[\mathbf{f}_{\text{gt}}(u,v) \neq 0]$ returns 1 if the ground-truth flow is valid (non-zero), and 0 otherwise. ϵ is a small constant to ensure numerical stability.

Remarks: EVLoc [19] demonstrates the feasibility of using a flow estimation network to predict dense correspondences between events and LiDAR maps. However, the method provides limited analysis of cross-modal consistency between events and LiDAR point clouds. The primary challenge in this task stems from the modality and appearance gap between depth maps and event images, making cross-modal matching substantially more difficult than intra-modal alignment. This naturally suggests that reducing this gap by making depth maps more “event-like” could simplify the matching process. Motivated by this insight, we redesign the framework to explicitly enhance cross-modal consistency and thereby further improve localization accuracy. The details of our method are presented in the following section.

IV. METHODOLOGY

In this section, we introduce our novel framework integrating the flow estimator alongside the edge detector (Sec. IV-A) in a mutually reinforcing manner, which is achieved via

two novel cross-task feature fusion modules (Sec. IV-B and IV-C). An overview of our framework is illustrated in Fig. 2.

A. Edge Detection Network

It’s known that events are typically triggered in gradient-rich regions—namely, along edges in event images. This insight motivates the idea of extracting edges from depth maps to facilitate cross-modal alignment. However, depth maps encode only geometric edge structures and lack the texture edge information present in event images, making such alignment not always reliable, particularly in complex or noisy environments. To address this, we propose to learn an encoder-decoder edge detector that extracts edge features from depth, and implicitly enhances feature-level consistency between modalities, as detailed in the Sec. IV-B and IV-C.

Due to the lack of learning-based edge detectors specifically designed for depth maps, we adopt the widely-used HED network [29], originally developed for RGB images. HED is a convolutional network that extracts five feature maps at resolutions $(1, 1/2, 1/4, 1/8, 1/16)$ from intermediate layers. Each feature map is passed through a 1×1 convolution and upsampled to match the input resolution, producing five side outputs. These outputs are then fused via a 1×1 convolution to generate the final edge map.

We use the input event image to generate ground truth edge map. Specifically, since the depth map generated from ground truth pose is aligned with the event image, we first apply dilation to the depth map to ensure coverage of sparse event pixels. Then, for each event pixel, we retrieve its corresponding depth value d . Using the pinhole camera model, we reconstruct the corresponding 3D points P_{edge} in the camera coordinate system. Next, we reproject the 3D edge points using the projection function π , based on the relative transformation between the ground truth pose and the initial pose:

$$I'_{EV} \triangleq \pi(P_{\text{edge}}, \mathbf{T}_{\text{init}}^{-1} \mathbf{T}_{\text{gt}}) \quad (2)$$

Finally, we generate the ground truth edge map I_{ED} by applying binarization to the reprojected event image I'_{EV} .

The network is trained using a class-balanced cross-entropy loss function [29], defined as:

$$\mathcal{L}_{\text{edge}} = - \sum_{u,v} \left[w I_{ED}(u,v) \log(p(u,v)) + (1-w)(1 - I_{ED}(u,v)) \log(1 - p(u,v)) \right] \quad (3)$$

Here, $p(u,v)$ denotes the predicted edge probability at pixel (u,v) , and w the class-balancing weight calculated based on the proportion of non-edge pixels in the ground truth edge map I_{ED} . This loss addresses the imbalance between edge and non-edge pixels by assigning a higher weight to the minority class (edge pixels).

B. Cross-task Feature Fusion

As outlined above, the generated depth map and event image are provided to the flow estimator and edge detector, which initially produce the event features F_{EV} ,

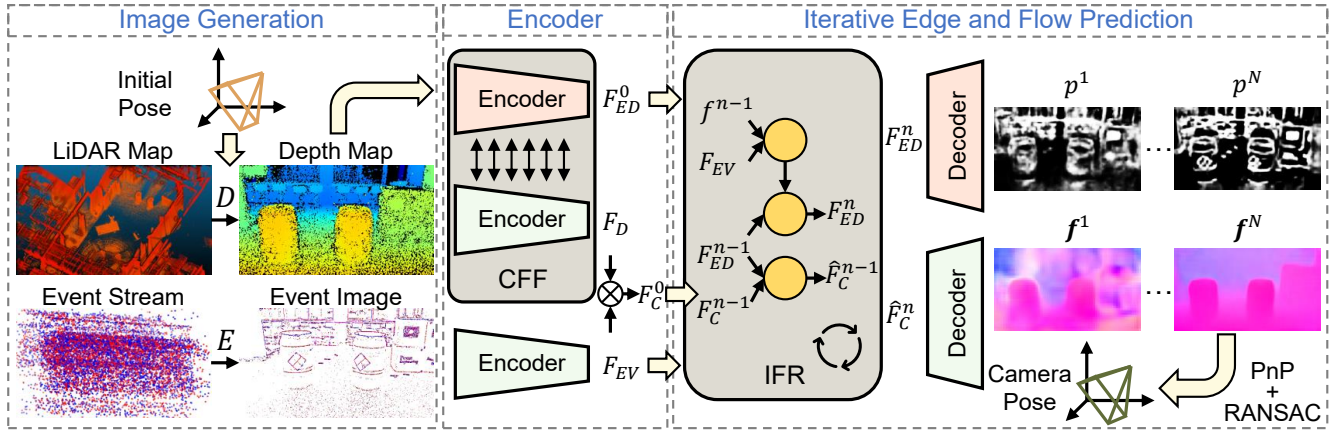


Fig. 2. The three stages of the proposed method: 1) *Image Generation*: a depth map is generated using a pinhole camera model D based on a given initial pose guess, and an event image is constructed from the raw event stream using an event image generator E ; 2) *Encoder*: the depth map and event image are fed into the flow estimation branch (**green module**), and the depth map also serves as input to the edge detection branch (**red module**). Encoded depth features F_D and edge features F_{ED} are fused across scales via the CFF module to produce edge-aware representations; 3) *Iterative Edge and Flow Prediction*: these, along with event features F_{EV} and correlation volumes F_C (constructed by indexing the element-wise products of F_D and F_{EV} using the current flow; see Fig. 5), are passed to the IFR module, where flow f and edge representations F_{ED} are refined iteratively and jointly, ultimately producing flow estimates $f^{1 \rightarrow N}$ and edge estimates $p^{1 \rightarrow N}$ from the corresponding decoders. The final flow estimate f^N is used to recover the camera pose via a PnP solver within a RANSAC loop. The yellow circles in the IFR module denote specific operations. See Sec. IV-C and Fig. 5 for details.




Fig. 3. Architecture of the CFF module. Both the edge detection and flow estimation branches employ five-layer encoders with different spatial resolutions at each layer. Feature fusion is performed at the middle three layers, as illustrated in the diagram.

depth features F_D , and edge features F_{ED} . The depth encoders in both branches are integrated through our proposed CFF module, as illustrated in Fig. 3. Multi-scale feature maps from the encoders of the edge detector and flow estimator are progressively fused and propagated. The flow estimator encoder extracts five feature maps at resolutions $(1/2, 1/2, 1/4, 1/8, 1/8)$ from the input depth map, while the edge detector encoder produces maps at $(1, 1/2, 1/4, 1/8, 1/16)$. Due to the resolution mismatch, fusion is performed only at the middle three layers. Features at layer i are concatenated and passed through two separate 1×1 convolutions to restore the original channel dimensions for each task branch, formulated as follows:

$$\hat{F}_s^i = \text{Conv}_{1 \times 1}^s(\text{Concat}(F_{ED}^i, F_D^i)), \quad (4)$$

where $s \in \{\text{ED}, \text{D}\}$ indicates the task branch (edge detection or depth), \hat{F}_{ED}^i and \hat{F}_D^i are enhanced features. Although both encoders operate on the same input depth map, they focus on different aspects: the edge detection encoder emphasizes fine-grained, low-level edge cues, while the flow estimation

encoder primarily captures motion-related features, focusing on both local correspondences and global motion context. Through cross-task feature fusion, the depth features gain enhanced sensitivity to local texture details, while the edge features become more geometrically structured, enhancing cross-modal consistency. As illustrated in Fig. 4, the depth feature map obtained after applying the CFF module contains richer texture details that align with structures present in the ground truth edge map.

C. Iterative Feature Refinement

The obtained features are fed into the IFR module, as shown in Fig. 5. This module consists of two components: (1) an edge update branch (**top**) that refines edge features by integrating information from aligned event features, and (2) a flow update branch (**bottom**) that refines flow estimates by incorporating edge features into the correlation volumes.

We first construct a correlation pyramid by computing the dot product between all pairs of depth features F_D and event features F_{EV} . The initial flow f^0 is set to zero. At iteration n ($n \in [1, N]$), the current flow estimate f^{n-1} is used to index into the correlation pyramid and retrieve a corresponding correlation volume F_C^n . In the flow update branch, the correlation volume is first concatenated with the current edge feature F_{ED}^n along the channel dimension, then passed through a 1×1 convolution to restore the original channel size. A GRU cell updates the hidden state using \hat{F}_C^n , the current flow, and context features F_{ctx} (with F_{ctx} and h^0 from the EVLoc-style context encoder [19]), and the resulting state is mapped by a flow head H to predict an incremental flow update. Then, in the edge update branch, the updated flow is used to warp the event feature F_{EV} , spatially aligning it with the edge feature F_{ED} . The current edge feature F_{ED}^n , extracted from the fourth encoder layer to match the resolution of the event feature, is split into a context feature c_{ED} and a hidden state h_{ED}^n . A GRU cell

takes as input the context feature, the hidden state, and the aligned event features, and outputs an updated hidden state. The updated hidden state is concatenated with the original context feature to form the refined edge feature F_{ED}^{n+1} .

$$\hat{F}_C^n = \text{Conv}_{1 \times 1}(\text{Concat}(F_C^n, F_{ED}^n)), \quad (5)$$

$$\mathbf{f}^n = \mathbf{f}^{n-1} + H(\text{GRU}(h^{n-1}, \hat{F}_C^n, \mathbf{f}^{n-1}, F_{\text{ctx}})), \quad (6)$$

$$F_{ED}^{n+1} = \text{Concat}(c_{ED}, \text{GRU}(h_{ED}^n, W(F_{EV}, \mathbf{f}^n), c_{ED})). \quad (7)$$

By integrating edge features in this way, the correlation volume is enriched with local structural details, complementing the global motion context and enabling the model to predict finer-grained flow. Notably, the two branches are mutually reinforcing: improved flow leads to better-aligned event features, while more accurate edge features enhance the fused correlation representation, further guiding the flow refinement process. See Fig. 6 the bottom rows.

D. Other Details

Loss function: As described above, the edge detector and flow estimator are trained jointly in a mutually reinforcing manner. The final loss is defined as a weighted combination of the two terms:

$$\mathcal{L} = \alpha \times \mathcal{L}_{\text{flow}} + \beta \times \mathcal{L}_{\text{edge}}, \quad (8)$$

where α is 1 and β is 100 in our experiments.

Pose estimation: We treat pixels with non-zero depth values as valid and extract flow predictions (i.e., 2D-3D correspondences) only at these locations. We use the publicly available PoseLib library [30], with a maximal reprojection error of 12.0 and Huber loss for optimization, while keeping all other parameters at their default values.

V. EXPERIMENTS

A. Datasets and Implementation Details

1) *Datasets:* M3ED dataset [31] is a typical dataset that provides high-definition LiDAR maps and ground truth trajectories, so we primarily evaluate our method on it. For

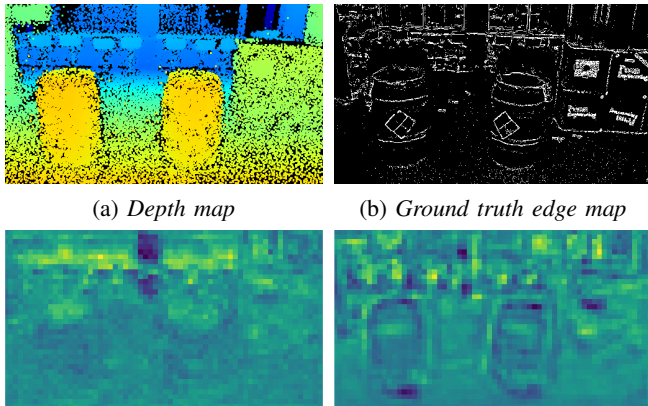


Fig. 4. Visualization of the encoded depth feature map. Each pixel value in the visualized feature map indicates the mean value of the corresponding feature vector at that location. The depth feature map after applying the CFF module contains richer texture details that align with structures.

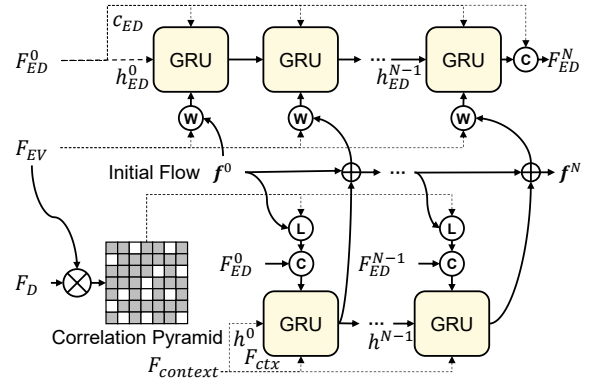


Fig. 5. Architecture of the IFR module. Edge features ($F_{ED}^0 \rightarrow F_{ED}^N$) and flow estimates ($\mathbf{f}^0 \rightarrow \mathbf{f}^N$) are progressively refined through mutual reinforcement. W, L, and C denote warping, indexing and concatenating, respectively.

TABLE I

ABLATION STUDY RESULTS. BOLD NUMBERS INDICATE THE BEST PERFORMANCE EXCEPT THE UPPER BOUND *Oracle*.

| | CFF | IFR | EPE | Mean Error | | Median Error | |
|---------------|-----|-----|-------------|---------------|-------------|---------------|-------------|
| | | | | Transl.[cm] ↓ | Rot.[°] ↓ | Transl.[cm] ↓ | Rot.[°] ↓ |
| Oracle | | | 1.66 | 4.14 | 0.39 | 3.31 | 0.22 |
| Baseline [19] | | | 6.10 | 9.01 | 1.10 | 8.11 | 0.97 |
| Direct Mask | | | 6.11 | 8.99 | 1.02 | 7.80 | 0.85 |
| (a) | ✓ | | 5.75 | 8.06 | 0.94 | 7.20 | 0.83 |
| (b) | ✓ | ✓ | 5.52 | 8.07 | 0.93 | 6.85 | 0.74 |

additional evaluation, we incorporate the DSEC dataset [32], which introduces more complex driving scenarios not present in M3ED. However, DSEC only offers disparity images. To facilitate evaluation, we reconstruct a 3D point cloud for each sample by back-projecting the disparity-derived depth into 3D space using the pinhole camera model. Additionally, since our method relies on high-quality, pre-built LiDAR maps to produce dense and accurate depth estimates, we exclude sequences heavily affected by dynamic objects, as their motion introduces significant noisy artifacts into the reconstructed maps.

2) *Data Processing:* To simulate uncertainty in coarse localization, we generate an initial pose estimate by perturbing the ground truth (identity matrix for DSEC) through random uniform sampling, with translation offsets within ± 50 cm and rotation offsets within $\pm 5^\circ$. To ensure consistency and efficiency, all input images are cropped and resized to fixed resolutions specific to each dataset. For the M3ED dataset, images are resized and cropped to 512×288 ; for DSEC, images are cropped to 480×360 .

3) *Training Details:* For M3ED, a separate model is trained per scene, as the dataset covers heterogeneous environments (indoor flight, outdoor parking, forest, etc.); for example, in the *falcon_indoor_flight* scene, the first two sequences are used for training and the last for evaluation. For scenes with only one sequence (e.g., *spot_indoor_building_loop*), we adopt a temporal split, using the first 70% of data for training and the remaining 30% for evaluation, ensuring that training and testing do




Fig. 6. Visualization of the predicted edges and flows. **Left column:** edge predictions; **Right column:** flow estimates.

not overlap. In contrast, a single model is trained across all DSEC scenes—since they all depict driving environments with similar sensor configurations and motion patterns—with the final sequence of each scene used for evaluation. In each experiment, the model is trained for 100 epochs with a batch size of 2, weight decay of 1×10^{-4} , an initial learning rate of 4×10^{-5} , and a *OneCycleLR* learning rate scheduler, using the Adam optimizer [33]. The IFR module is configured to run for 12 iterations during training and 24 iterations during testing. All experiments are conducted on a single NVIDIA RTX 4090 GPU.

4) *Evaluation Metrics:* For evaluation, localization accuracy is assessed using the mean and median errors of the predicted poses, while the average End-Point Error (EPE) is used to evaluate the accuracy of the estimated flow.

B. Ablation Study

To evaluate the effectiveness of the proposed modules, we conduct a series of ablation experiments, as demonstrated in Table I and Fig. 6. The *baseline* is EVLoc [19], which we retrain using the same image resolution and experimental settings as our method. In the *oracle* setting, we use the ground truth edge map to mask the depth input before both training and testing, producing a masked depth map

that closely resembles the texture of the input event image. This minimizes the appearance gap between modalities. As expected, the model significantly outperforms the baseline, demonstrating the theoretical upper bound of flow-based matching when perfect edge information is available. Next, we explore a progressive pipeline *direct mask* where the predicted edge maps are directly used to mask the depth input. The edge detector is trained independently, and an example of the predicted edge map is shown in the 2nd row of Fig. 6, this configuration slightly outperforms the baseline, validating that even imperfect edge predictions can improve depth–event matching.

Building upon this, we incrementally incorporate the proposed modules. By fusing edge features into the flow estimator through the CFF module, the *model (a)* achieves further improvements beyond simple masking. This suggests that implicit feature-level integration is more effective than direct masking. Additionally, as shown in the 3rd row of Fig. 6, the predicted edges exhibit finer details compared to those produced by a standalone edge detector, indicating that the CFF module also enhances the edge detection branch by integrating complementary information. Adding the IFR module yields further gains in localization accuracy (*model (b)*), confirming its contribution to refining flow estimates. Compared to a standalone edge detector HED (Fig. 6, 2nd row), our method produces sharper and more detailed edge predictions, demonstrating the IFR module’s ability to effectively fuse local texture cues from the event features into the edge representations. Furthermore, as iterations progress, both the predicted edges and the refined flow estimates become increasingly complete and accurate, as illustrated in Fig. 6 the last three rows, highlighting the mutually reinforcing nature of the iterative refinement process.

C. Results Analysis

In this section, we compare our method against the state-of-the-art event–LiDAR localization approach EVLoc [19] under the same experimental setup, as well as the RGB-based image-to-LiDAR method I2D-Loc [34]. Results are shown in Table II. M3ED does not provide ground truth poses for standard cameras. Although approximate poses can be obtained via interpolation, the resulting depth maps exhibit significant misalignment with the camera frames. In contrast, DSEC provides ground truth disparity maps for each RGB image, allowing us to apply the same ground truth generation method described above for the experiments on standard cameras.

As shown in Table II, our proposed method consistently outperforms EVLoc across all sequences under different lighting conditions, demonstrating the overall superiority of the proposed approach. On M3ED, in indoor environments (*falcon_indoor_flight* and *spot_indoor_building_loop*), LEAR reduces translation error by roughly 15–20% and rotation error by about 20–25% compared to EVLoc. In more challenging outdoor scenarios, both methods perform worse overall compared to indoors, suggesting that limited or unstructured scene texture makes cross-modal matching more

TABLE II

COMPARISON OF LOCALIZATION PERFORMANCE AMONG I2D-LOC, EVLOC, AND THE PROPOSED METHOD ON M3ED AND DSEC. MEDIAN ERRORS IN TRANSLATION AND ROTATION ARE USED AS EVALUATION METRICS. BOLD NUMBERS INDICATE THE BEST PERFORMANCE.

| Dataset | Test Scene | Time | I2D-Loc [34] | | EVLoc [19] | | Ours | |
|-----------------------|--|-----------|--------------|----------|--------------|-------------|--------------|-------------|
| | | | Transl.[cm]↓ | Rot.[°]↓ | Transl.[cm]↓ | Rot.[°]↓ | Transl.[cm]↓ | Rot.[°]↓ |
| M3ED [31] | <i>falcon_indoor_flight</i> | - | - | - | 8.11 | 0.97 | 6.85 | 0.74 |
| | <i>spot_indoor_building_loop</i> | - | - | - | 12.88 | 2.07 | 10.94 | 1.63 |
| | <i>falcon_outdoor_day_penno_parking</i> | day | - | - | 25.11 | 1.32 | 22.40 | 1.18 |
| | <i>falcon_forest_into_forest</i> | day | - | - | 20.46 | 1.87 | 19.35 | 1.84 |
| | <i>spot_outdoor_day_penno_short_loop</i> | day | - | - | 17.89 | 0.88 | 14.37 | 0.74 |
| | <i>spot_forest_road</i> | day | - | - | 13.55 | 0.60 | 12.96 | 0.57 |
| | <i>car_urban_day_penno</i> | day | - | - | 14.15 | 0.60 | 11.88 | 0.56 |
| | <i>car_forest_into_ponds</i> | day | - | - | 19.76 | 0.86 | 16.23 | 0.75 |
| | <i>spot_outdoor_night_penno_short_loop</i> | night | - | - | 34.85 | 1.71 | 24.97 | 1.04 |
| | <i>car_urban_night_penno</i> | night | - | - | 20.84 | 0.72 | 19.15 | 0.87 |
| DSEC [32] | <i>interlaken_00</i> | dawn/dusk | 10.45 | 0.42 | 7.11 | 0.30 | 5.56 | 0.25 |
| | <i>thun_00</i> | dawn/dusk | 11.15 | 0.55 | 8.26 | 0.42 | 8.08 | 0.41 |
| | <i>zurich_city_00</i> | dawn/dusk | 11.30 | 0.52 | 7.94 | 0.37 | 6.78 | 0.34 |
| | <i>zurich_city_01</i> | dawn/dusk | 9.88 | 0.36 | 7.63 | 0.30 | 6.15 | 0.28 |
| | <i>zurich_city_02</i> | dawn/dusk | 8.68 | 0.33 | 7.93 | 0.31 | 6.68 | 0.27 |
| | <i>zurich_city_03</i> | night | 14.01 | 0.74 | 8.79 | 0.61 | 8.14 | 0.51 |
| | <i>zurich_city_04</i> | dawn/dusk | 11.82 | 0.51 | 7.49 | 0.30 | 6.16 | 0.26 |
| | <i>zurich_city_05</i> | day | 9.27 | 0.37 | 6.42 | 0.28 | 5.36 | 0.24 |
| | <i>zurich_city_06</i> | day | 14.26 | 0.62 | 8.12 | 0.39 | 6.98 | 0.37 |
| | <i>zurich_city_07</i> | day | 11.33 | 0.43 | 7.27 | 0.34 | 6.25 | 0.32 |
| | <i>zurich_city_08</i> | dawn/dusk | 8.72 | 0.36 | 6.62 | 0.27 | 5.81 | 0.26 |
| | <i>zurich_city_09</i> | night | 8.05 | 0.32 | 6.38 | 0.29 | 5.92 | 0.28 |
| <i>zurich_city_10</i> | night | 11.35 | 0.49 | 8.37 | 0.35 | 7.35 | 0.34 | |
| <i>zurich_city_11</i> | dawn/dusk | 9.56 | 0.41 | 6.14 | 0.26 | 5.21 | 0.23 | |

difficult. Nevertheless, our method does not degrade relative to the baseline and still achieves modest gains, indicating that the model continues to learn useful features from complex environments. Nighttime sequences are particularly challenging: even after denoising and deblurring, the generated event frames contain little useful information of the captured scene but substantial noise, and the high resolution of the event camera further amplifies the noise effects. Nevertheless, LEAR achieves performance comparable to EVLoc, despite the lack of useful structural information.




Fig. 7. Comparison of ground-truth poses (green), initial poses (blue), and poses predicted by our method (red) on M3ED. Our predictions align closely with the ground truth, demonstrating high accuracy.

DSEC contains only driving scenarios but under diverse lighting conditions. LEAR consistently outperforms EVLoc and achieves up to 36% lower translation and 29% lower rotation errors, and also surpassing I2D-Loc with reductions of up to 61% in translation and 44% in rotation errors, respectively. The significant improvement over I2D-Loc highlights the benefits of event cameras for cross-modal data association. We believe this could be attributed to the much larger domain gap between daytime and nighttime RGB images, which causes the models to converge to a local

minimum. In contrast, event images remain consistent across both daytime and nighttime conditions.

Figure 7 visualizes the localization results of one scene in M3ED. Starting from a coarse initial pose (blue), LEAR refines it to a precise prediction (red) that closely aligns with the ground truth (green).

TABLE III
RUNTIME AND MEMORY CONSUMPTION COMPARISON.

| Methods | Stage | Runtime [ms] | Memory [MB] | Transl.[cm]/Rot.[°] |
|------------|----------------------------|--------------|-------------|---------------------|
| EVLoc [19] | Feature encoding | 4.33 | 1202 | 8.11/0.97 |
| | Flow estimation (24 iters) | 51.80 | | |
| Ours | Feature encoding | 7.71 | 1232 | 6.85/0.74 |
| | Flow estimation (24 iters) | 100.43 | | |
| | Feature encoding | 7.71 | 1232 | 7.01/0.75 |
| | Flow estimation (12 iters) | 51.30 | | |

Table III summarizes the runtime and memory consumption of the baseline and our method on M3ED. Under the same number of iterations, our method requires negligible extra cost on memory, while more runtime (primarily in flow estimation) than EVLoc. Nevertheless, this overhead remains practical and enables significant gains in localization accuracy. Furthermore, our method at 12 iterations still outperforms EVLoc at 24 iterations, while they require a comparable runtime in this case. These results demonstrate the effectiveness of the enhanced feature representation and iterative refinement introduced in our framework.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we propose a dual-branch framework for event-based visual localization in LiDAR maps. Our approach integrates an optical flow estimation network with an edge detection network, and introduces a cross-task feature fusion module along with an iterative feature refinement

module. These components enable implicit feature-level fusion between the two branches, enhancing cross-modal feature consistency in finding accurate correspondences for localization. Extensive experiments on multiple public datasets demonstrate the effectiveness of the proposed method.

Limitations: Despite its advantages, our method relies on local flow estimation in 2D space, which requires sufficient overlap between the input depth map and the event image. Moreover, successful matching depends heavily on the presence of sufficient depth variation. In scenarios with degenerate depth information—such as flat ground or smooth walls—our method may fail due to the lack of distinguishable geometric features. In the future, we plan to investigate direct event-to-point cloud registration without projection, which may offer greater robustness in cases with limited initial pose accuracy or poor geometric variation.

VII. ACKNOWLEDGEMENT

This research project has been supported by the Austrian Science Fund (FWF) under project agreement (I 6747-N) EVELOC.

REFERENCES

- [1] T. Fischer and M. Milford, "Event-based visual place recognition with ensembles of temporal windows," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6924–6931, 2020.
- [2] A. J. Lee and A. Kim, "Eventvlad: Visual place recognition with reconstructed edges from event cameras," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 2247–2252.
- [3] G. B. Nair, M. Milford, and T. Fischer, "Enhancing visual place recognition via fast and slow adaptive biasing in event cameras," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2024, pp. 3356–3363.
- [4] M. Milford, H. Kim, S. Leutenegger, and A. Davison, "Towards visual slam with event-based cameras," in *Robotics: Science and System Workshops*, 2015.
- [5] T. Fischer and M. Milford, "How many events do you need? event-based visual place recognition using sparse but varying pixels," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12275–12282, 2022.
- [6] D. Kong, Z. Fang, K. Hou, H. Li, J. Jiang, S. Coleman, and D. Kerr, "Event-vpr: End-to-end weakly supervised deep network architecture for visual place recognition using event-based vision sensor," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–18, 2022.
- [7] S. Hussaini, M. Milford, and T. Fischer, "Applications of spiking neural networks in visual place recognition," *IEEE Transactions on Robotics*, 2024.
- [8] O. Claxton, C. Malone, H. Carson, J. J. Ford, G. Bolton, I. Shames, and M. Milford, "Improving visual place recognition based robot navigation by verifying localization estimates," *IEEE Robotics and Automation Letters*, 2024.
- [9] W. Yuan and S. Ramalingam, "Fast localization and tracking using event sensors," in *IEEE International Conference on Robotics and Automation*. IEEE, 2016, pp. 4564–4571.
- [10] R. Yuan, T. Liu, Z. Dai, Y.-F. Zuo, and L. Kneip, "Evit: Event-based visual-inertial tracking in semi-dense maps using windowed nonlinear optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2024, pp. 10656–10663.
- [11] Y.-F. Zuo, W. Xu, X. Wang, Y. Wang, and L. Kneip, "Cross-modal semidense 6-dof tracking of an event camera in challenging conditions," *IEEE Transactions on Robotics*, vol. 40, pp. 1600–1616, 2024.
- [12] N. Xu, L. Wang, Z. Yao, and T. Okatani, "Mets: Motion-encoded time-surface for event-based high-speed pose tracking," *International Journal of Computer Vision*, pp. 1–19, 2025.
- [13] G. Gallego, J. E. Lund, E. Mueggler, H. Rebecq, T. Delbruck, and D. Scaramuzza, "Event-based, 6-dof camera tracking from photometric depth maps," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 10, pp. 2402–2412, 2017.
- [14] S. Bryner, G. Gallego, H. Rebecq, and D. Scaramuzza, "Event-based, direct camera tracking from a photometric 3d map using nonlinear optimization," in *IEEE International Conference on Robotics and Automation*. IEEE, 2019, pp. 325–331.
- [15] A. Nguyen, T.-T. Do, D. G. Caldwell, and N. G. Tsagarakis, "Real-time 6dof pose relocalization for event cameras with stacked spatial lstm networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 0–0.
- [16] Y. Jin, L. Yu, G. Li, and S. Fei, "A 6-dofs event-based camera relocalization system by cnn-lstm and image denoising," *Expert Systems with Applications*, vol. 170, p. 114535, 2021.
- [17] H. Lin, M. Li, Q. Xia, Y. Fei, B. Yin, and X. Yang, "6-dof pose relocalization for event cameras with entropy frame and attention networks," in *Proceedings of the 18th ACM SIGGRAPH International Conference on Virtual-Reality Continuum and Its Applications in Industry*, 2022, pp. 1–8.
- [18] H. Ren, J. Zhu, Y. Zhou, H. Fu, Y. Huang, and B. Cheng, "A simple and effective point-based network for event camera 6-dofs pose relocalization," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 18 112–18 121.
- [19] K. Chen, J. Zhang, and F. Fraundorfer, "Evloc: Event-based visual localization in lidar maps via event-depth registration," in *IEEE International Conference on Robotics and Automation*. IEEE, 2025, pp. 5751–5757.
- [20] A. J. Lee and A. Kim, "Eventvlad: Visual place recognition with reconstructed edges from event cameras," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021, pp. 2247–2252.
- [21] T. Liu, R. Yuan, Y. Ju, X. Xu, J. Yang, X. Meng, X. Lagorce, and L. Kneip, "Gs-evt: Cross-modal event camera tracking based on gaussian splatting," in *IEEE International Conference on Robotics and Automation*. IEEE, 2025, pp. 4587–4593.
- [22] K. Ta, D. Bruggemann, T. Brödermann, C. Sakaridis, and L. Van Gool, "L2e: Lasers to events for 6-dof extrinsic calibration of lidars and event cameras," in *IEEE International Conference on Robotics and Automation*. IEEE, 2023, pp. 11 425–11 431.
- [23] W. Xing, S. Lin, L. Yang, and J. Pan, "Target-free extrinsic calibration of event-lidar dyad using edge correspondences," *IEEE Robotics and Automation Letters*, vol. 8, no. 7, pp. 4020–4027, 2023.
- [24] X. Lin, C. Qiu, S. Shen, Y. Zang, W. Liu, X. Bian, M. Müller, C. Wang, et al., "E2pnet: event to point cloud registration with spatio-temporal representation learning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 18 076–18 089, 2023.
- [25] J. Li and G. H. Lee, "Deepi2p: Image-to-point cloud registration via deep classification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 960–15 969.
- [26] Q.-H. Pham, M. A. Uy, B.-S. Hua, D. T. Nguyen, G. Roig, and S.-K. Yeung, "Lcd: Learned cross-domain descriptors for 2d-3d matching," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, 2020, pp. 11 856–11 864.
- [27] Prophesee, "Prophesee metavision sdk," available at <https://docs.prophesee.ai/stable/index.html>.
- [28] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *European Conference on Computer Vision*. Springer, 2020, pp. 402–419.
- [29] S. Xie and Z. Tu, "Holistically-nested edge detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2015, pp. 1395–1403.
- [30] V. Larsson and contributors, "PoseLib - minimal solvers for camera pose estimation," 2020. [Online]. Available: <https://github.com/vlarsson/PoseLib>
- [31] K. Chaney, F. Cladera, Z. Wang, A. Bisulco, M. A. Hsieh, C. Korpela, V. Kumar, C. J. Taylor, and K. Daniilidis, "M3ed: Multi-robot, multi-sensor, multi-environment event dataset," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, June 2023, pp. 4015–4022.
- [32] M. Gehrig, W. Aarents, D. Gehrig, 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.
- [33] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [34] K. Chen, H. Yu, W. Yang, L. Yu, S. Scherer, and G.-S. Xia, "I2d-loc: Camera localization via image to lidar depth flow," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 194, pp. 209–221, 2022.