

SuperEvent: Cross-Modal Learning of Event-based Keypoint Detection for SLAM

Yannick Burkhardt^{1,2,3}Simon Schaefer^{1,2,3}Stefan Leutenegger^{1,2,3}

Abstract

Event-based keypoint detection and matching holds significant potential, enabling the integration of event sensors into highly optimized Visual SLAM systems developed for frame cameras over decades of research. Unfortunately, existing approaches struggle with the motion-dependent appearance of keypoints and the complex noise prevalent in event streams, resulting in severely limited feature matching capabilities and poor performance on downstream tasks. To mitigate this problem, we propose SuperEvent, a data-driven approach to predict stable keypoints with expressive descriptors. Due to the absence of event datasets with ground truth keypoint labels, we leverage existing frame-based keypoint detectors on readily available event-aligned and synchronized gray-scale frames for self-supervision: we generate temporally sparse keypoint pseudo-labels considering that events are a product of both scene appearance and camera motion. Combined with our novel, information-rich event representation, we enable SuperEvent to effectively learn robust keypoint detection and description in event streams. Finally, we demonstrate the usefulness of SuperEvent by its integration into a modern sparse keypoint and descriptor-based SLAM framework originally developed for traditional cameras, surpassing the state-of-the-art in event-based SLAM by a wide margin. Source code is available at [ethz-mrl.github.io/SuperEvent](https://github.com/ethz-mrl/SuperEvent).

1. Introduction

Event cameras offer exciting advantages over their frame-based counterparts, such as an increased temporal resolution, little motion blur, and a high dynamic range. These properties promise to enhance robustness for robot estimation and perception tasks under fast motion and uncontrollable lighting conditions. However, processing the sparse and asynchronous output of event cameras requires fundamentally different algorithms than traditional frames. The decades-long research edge in frame-based vision resulted




Figure 1. Detections on the sequence *rec1499023756* of the DDD20 [31] dataset (not used for training). Top: Pseudo-labels from SuperPoint [15] and SuperGlue [59] on the gray-scale frames. Bottom: Matched keypoints from SuperEvent in the event stream at the corresponding time stamps.

in most hardware and software being highly optimized for frame processing. A common approach for tasks like Structure-from-Motion, Visual Odometry (VO), Simultaneous Localization and Mapping (SLAM), or Place Recognition is to detect and match keypoints across frames to gain a geometric scene understanding [7, 40]. Consequently, a robust keypoint detector and descriptor working on the event stream would unlock the potential of integrating event cameras into extensive and highly-developed algorithms developed initially for frame-based vision.

While several authors already proposed event-based approaches to detect keypoints [9, 10, 19, 32, 33, 47], their matching abilities still show limitations. Most employ nearest-neighbor matching in pixel space, which restricts their usability to visually simple scenes. In an attempt to overcome this issue, more recent approaches learn descriptors in a data-driven fashion. To generate ground truth labels – inspired by frame-based methods [15] – they warp their event representations with a random homography and penalize non-consistent model predictions. This approach forces the models to learn descriptors to match visually similar keypoints. However, it fundamentally ignores the true

¹Technical University of Munich,
{[yannick.burkhardt](mailto:yannick.burkhardt@tum.de), [simon.k.schaefer](mailto:simon.k.schaefer@tum.de)}@tum.de

²ETH Zürich, lestefan@ethz.ch

³Munich Center for Machine Learning (MCML)

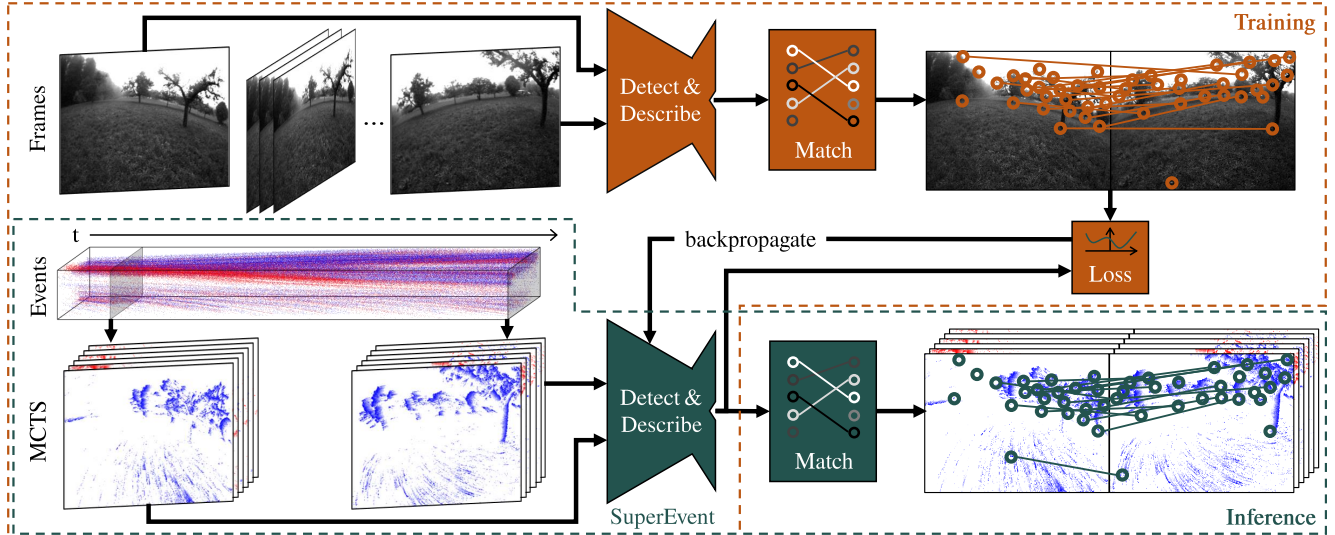


Figure 2. Data processing pipeline of SuperEvent: For training, pseudo-labels are generated using a frame-based detector and matcher on the gray-scale frame pairs. We generate spatially synchronized MCTS at the same timestamps and feed them to SuperEvent. The network predictions are compared to the pseudo-labels and the network weights are optimized using backpropagation. During inference, the network predictions can be used to detect keypoints and match their descriptors on the event stream only.

nature of the event generation process that conflates camera motion with scene appearance and geometry. In essence, matches are generated only from one original set of events per keypoint, thereby producing only limited and potentially also unrealistic training data. We hypothesize that this simplification ultimately results in unstable matching.

In contrast, we propose training our model to predict descriptors for the keypoint correspondences at *different* points in time – therefore training with *different* and actually *recorded* rather than warped events. Since generating keypoint and match labels manually is intractable, we design a scalable method to exploit the great capabilities of frame-based vision models: generating pseudo-labels in aligned frames and then using them to train our model’s predictions. An example is shown in Figure 1.

To provide the model with rich input data, we propose the novel event representation Multi-Channel Time Surfaces (MCTS) which generalizes the commonly used time surface representation [38] to an n -channel tensor of time surfaces for both polarities and different time window sizes. Since an optimal window size depends inversely on the motion magnitude, this approach further increases the model’s robustness to both slow and fast motion. Additionally, we revisit network architecture choices of frame- and event-based vision to optimize our model’s performance. Figure 2 shows the complete data processing pipeline.

In absence of evaluation procedures to test event-based keypoint and descriptor stability, we develop a benchmark to assess the matched keypoint’s quality for pose estimation, similar to benchmarks commonly used in frame-based vision [22, 59, 62, 64]. On two datasets recorded with dif-

ferent sensors, settings, resolutions, and scenes, we outperform other data-driven and handcrafted approaches in terms of pose estimation by a large margin, demonstrating the robustness of our approach. We carry out ablation studies to confirm the effectiveness of our design choices.

Finally, we integrate our model into a state-of-the-art (SOTA) stereo visual inertial SLAM (VI-SLAM) framework [40], replacing the frame-based keypoint detector and descriptor extractor with SuperEvent. Our results surpass the current SOTA on event-based Visual Odometry and SLAM. This demonstrates the great potential of an event-sensor integration into existing, well-developed algorithms for frame-based vision.

We name our method *SuperEvent* in honor of the great influence of the works SuperPoint [15] and SuperGlue [59] on the research field of keypoint detection as well as on our model directly: ultimately, their frame-based detections and matches are used as pseudo-labels to teach keypoint detection and description to SuperEvent.

Contributions: (1) We propose a scalable method to generate training data for interest point detection and matching, leveraging robust existing frame-based keypoint detectors. The resulting data incorporates the real temporal dependence of event generation using real event camera recordings. This enables data-driven models to effectively learn robust keypoint detection and description. (2) We develop the model for event-based vision through key design choices: Our novel MCTS event representation increases the performance of data-driven models by reducing their dependence on motion speed of the scene. A transformer backbone shows to further improve the model’s feature ex-

traction. In our conducted experiments, our approach outperforms SOTA detection and matching by a large margin. We publicly release our code to support future event-based vision research. (3) Showcasing the practicality of our method, we integrate the trained model into an existing stereo VI-SLAM approach developed for frame-based vision. The resulting event inertial SLAM system outperforms SOTA approaches on the commonly used TUM-VIE [35] and VECtor [18] datasets by a large margin.

2. Related Work

We begin with a short summary of keypoint detectors for frame-based vision since our work employs some of these findings for model design and training data generation. We then group the event-based works into handcrafted and data-driven keypoint detection, as well as tracking methods. Since our focus in this work lies on keypoint detection and matching based on event data exclusively, we do not consider approaches that rely on combination with other sensors, e.g., frame cameras. For an extensive overview of event-based vision, we refer the interested reader to [17].

Frame-based keypoint detection: there exists a vast variety of approaches to detecting and describing keypoints in frames. Traditional, handcrafted methods [5, 6, 28, 41, 45, 58, 60] are still widely used, since they are well studied and usually work robustly with little overhead. More recent, data-driven methods outperform the traditional keypoint detectors with impressive results. SuperPoint [15] is the first method trained with homographic adaptation, a self-supervised training technique of warping input images and their detected labels by randomized homographies. SuperPoint’s architecture consists of a shared VGG [61]-backbone with a detector and descriptor head, predicting descriptors in a discrete grid which is interpolated to full resolution. The authors later publish the extension SuperGlue [59], a graph neural network that improves the descriptor matching. Following works propose different architectures trained with homographic adaptation [22, 62, 64].

Event-based keypoint detection: handcrafted approaches can be broadly divided into filters and frame reconstruction. Filtering methods work directly on the event stream, attempting to find events originating from corners [1, 11]. While their asynchronous nature offers little processing overhead, downstream tasks are required to process asynchronous data. Existing well-developed downstream solutions for frame-based vision therefore require major modifications.

Frame reconstruction methods collect events to construct a frame representation, such as binary event frames or time surfaces [38]. Adopting methods from frame-based vision, such as the Harris corner detector [28] in [23, 65] or the FAST detector [57] in [49], these approaches detect corners in these event frames. However, they suffer from the in-

herent differences between reconstructed event frames and traditional frames, resulting in limited robustness of the algorithms developed for frame-based vision.

Additionally, these handcrafted methods usually employ simple nearest-neighbor matching in pixel space to track keypoints over time. This approach is only reliable for simple scenes since cluttered geometry, high variance in depth, and fast motions cause keypoints to become occluded, tightly clustered, or to jump. Also, the developed heuristics struggle with the ubiquitous and complex noise inherent to event cameras. Usually, they require manual parameter tuning for different camera models, datasets, or sometimes even scenes.

Event-based keypoint detection: data-driven approaches promise to cope better with the event data and its noise. A series of works [9, 10, 47] train data-driven corner detectors using labels by detecting keypoints on corresponding gray-scale frames that are synthetic or from an HVGA ATIS sensor. While these approaches learn data-driven detection, they only track keypoints by nearest-neighbor matching in pixel space, yielding the aforementioned drawback of limited robustness in complex scenes. Additionally, the employed corner detector for ground truth generation as well as the machine-learning models have limited performance due to their simplicity.

For these reasons, EventPoint [32] and SD2Event [19] employ more complex neural networks that additionally predict descriptors. While EventPoint uses the same architecture as SuperPoint [15] and finetunes its pre-trained weights, SD2Event employs agent-based attention to learn detection and description. Both approaches are trained with homographic adaptation and varying time window sizes. While this results in rotation-, distortion-, and scale-invariant models, the main challenge of keypoint description remains largely unsolved: Because most events result from scene or camera motion, corresponding keypoints vary in visual appearance. Since robust keypoint matching requires similar descriptor pairs, this motion dependence must be considered. However, homographic adaptation creates static frame pairs with little to no variance in scene motion within the samples. This contradicts the dynamic and motion-dependent nature of events, resulting in unstable keypoint matching between different timestamps.

Additionally, the authors of EventPoint propose the event representation Tencode, extending time surfaces with two additional channels to encode the events’ polarities. However, to match SuperPoint’s architecture with single-channel input, the Tencode tensor is then converted to one channel, disregarding the previously introduced separation of polarities.

Our method enforces robust descriptor matching by employing training data containing pseudo-labels for keypoints and descriptors under varying motion. We avoid

synthetic event data since current simulators struggle to model realistic noise. Furthermore, our MCTS representation strictly separates the polarities and provides a range of time window sizes, providing additional information and omitting the time window parameter choice.

Event-based keypoint tracking: while early works fall back on batching events and processing events in a discretized way [37, 52, 68], a promising idea is to exploit their spatial continuous visual flow [2, 13, 30]. Among these approaches is HASTE [3] which calculates the most likely motion hypothesis from a fixed set of translations and rotations for every incoming event. Since HASTE relies on external keypoint initialization, RATE [33] employs a Shi-Tomasi [60] corner detector on the binary event frame, initializing several HASTE instances to track multiple keypoints in parallel.

A major drawback of the described tracking approaches is their computational overhead when processing each event individually. Also, most methods only track one keypoint, resulting in an almost linear growth of computational cost with increasing number of tracker instances. E.g., RATE tracks only 20 features in real-time for an event stream of 180×40 pixels resolution – which is hardly sufficient for downstream applications such as SLAM. Additionally, these handcrafted approaches struggle with environmental and hardware changes, as well as the event cameras’ noise.

Our method SuperEvent works reliably in real-time with adjustable frequency and without re-training or excessive parameter tuning for different camera models. It’s dense keypoint heatmap predictions barely result in overhead for more detections. In contrast to the tracking approaches, SuperEvent can be directly integrated into downstream algorithms relying on synchronous output and descriptors, e.g., for SLAM with place recognition, as we demonstrate.

Event-based Odometry: existing Event Odometry (EO) approaches are developed specifically for event processing. While some approaches combine events with frames [8, 25, 26, 46, 66], event-only approaches can be classified as monocular EO [36, 55], monocular EO with IMU [24, 27], stereo EO [21, 67], and stereo EO with IMU [44, 53, 54]. Due to the short history of event cameras, these systems require extensive research and development efforts to work reliably in practice.

In contrast, SuperEvent can be integrated into existing frame-based VO and SLAM systems. We thereby achieve event-based VI-SLAM leveraging highly accurate frame-based systems, emerging from years of research [7, 40].

3. Method

In this section, we describe how we train SuperEvent, introduce the advantages and construction of the MCTS event representation, and detail the network design choices.

3.1. Training Data Generation

To train SuperEvent, we exclusively employ real event camera recordings. Thus, our model learns to cope with the complex event stream properties, e.g., artifacts and noise, which are difficult to simulate realistically.

However, manually annotating keypoints in the asynchronous event data stream is very time-consuming and does not scale. Therefore, we propose an approach to generate ground truth data automatically. We take advantage of the pixel-synchronized events and frames of iniVation DAVIS¹ event cameras. Employing a highly robust frame-based keypoint detector and matcher such as SuperPoint [15] + SuperGlue [59] on the gray-scale frames, we can obtain temporally sparse pseudo-labels. To generate a diverse set of training data with reliable pseudo-labels, we process every frame i of a sequence in the following way:

- Firstly, we compute the distances in pixel-space of all key points to their respective matches in frame $i + 1$. We only consider frames i for training where the distances’ median surpasses a threshold c_m . A small median distance indicates that the scene is nearly static, resulting in insufficient event data for a useful prediction.
- To expose the model to diverse camera pose changes and motion directions, we generate pseudo-labels by matching the descriptors with a subsequent frame $i + j$. We increase j recursively, resulting in multiple pseudo-labels with increasing time interval between the frames:

$$j \leftarrow j + j_s, j_s \sim \mathcal{U}(1, j_{\max}), \quad (1)$$

with the increment j_s drawn from the uniform distribution $\mathcal{U}(1, j_{\max})$. j is initialized with 0 before its first update. The maximum step distance j_{\max} is manually adjusted depending on the properties and texture of the sequence, e.g., for sequences with fast motion, j_{\max} is decreased to yield sufficient matches.

- If the number of detected matches falls below the threshold c_m , these pseudo-labels are discarded, and the next frame $i + 1$ becomes the new reference frame. This procedure ensures sufficient visual overlap and enough texture in the two frames.

Frames with under- or overexposure, or strong motion blur harm the prediction quality of the frame-based model. While these frames are usually filtered out since there are few matches (if c_d is chosen adequately), we also manually remove scenes that mostly contain too dark, too bright, or blurred frames.

3.2. Multi-channel Time Surface

To allow for efficient batch processing of event data, the asynchronous stream is converted to a tensor with fixed dimensions. In the literature, various variants have been proposed to encode the event stream. An ablation study in [48]

¹<https://inivation.com/>

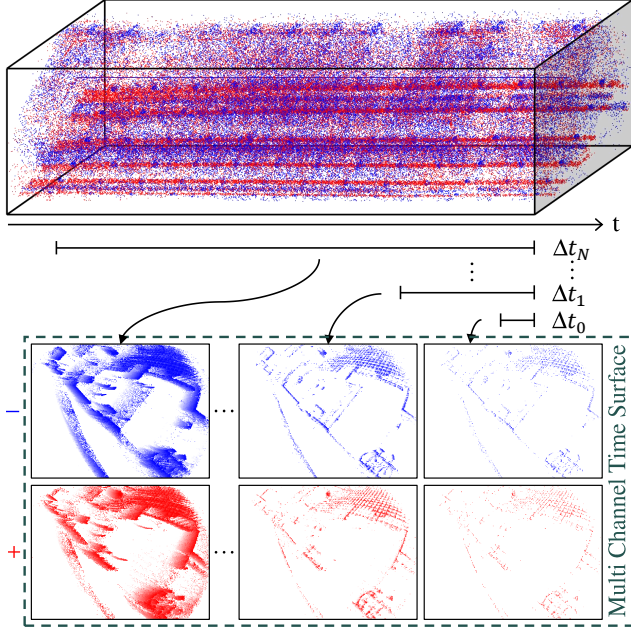


Figure 3. Schematic depiction of the MCTS generation process.

compares some of the most commonly used event representations: binary event frames, time surfaces [38], and voxel grids [70]. For their data-driven feature tracking approach, time surfaces yield the best model performance. The time surface variant *Tencode* with two more channels encoding the event polarity is proposed in [32]. It is used as model input after conversion to a single-channel frame.

Time surfaces and its variant *Tencode* have proven to be an efficient input representation for keypoint detection and tracking algorithms [19, 32, 48], since they assign the highest values to the pixel locations of the most recent events that are usually caused by visual edges. This facilitates finding keypoints along these edges. However, these encodings have shortcomings:

- Time surfaces are one-channel frames that ignore the event polarity and thereby discard potentially useful information. The *Tencode* representation is an attempt to overcome this issue. However, since it is converted back to a single channel tensor, there is still no clear separation between positive and negative events.
- Time surfaces and the *Tencode* representation depend on the time window length Δt . Increasing this parameter not only results in more events being considered but also decreases the contrast between the most recent events and past events. For high optical flow, this negatively impacts the detection of edges. Conversely, a smaller time window increases this contrast but includes fewer events, which can fail to sufficiently capture edges in areas with little optical flow. Choosing a suitable Δt is challenging, as it depends on the camera and scene motion which might be unknown *a priori*. The bottom part of Figure 3

shows the effect of varying Δt : Only the time surfaces in the middle clearly display the net on the top right.

Our proposed MCTS representation mitigates these problems: It combines time surfaces for each polarity and for various time windows of temporal size Δt_n , $n \in \{1, \dots, N\}$. We distribute the time window sizes logarithmically to trade off data preprocessing and coverage of a wide variety of visual motion speeds. Events are encoded by their timestamp t , their pixel coordinates x and y , and their polarity p . Mathematically, the MCTS tensor at time τ considering all i events $\mathbf{e}_i = (t_i, x_i, y_i, p_i)$ between time $\tau - \Delta t_n$ and τ is defined as

$$\mathbf{MCTS} = (\mathbf{TS}_{-1, \Delta t_1}, \dots, \mathbf{TS}_{-1, \Delta t_N}, \mathbf{TS}_{+1, \Delta t_1}, \dots, \mathbf{TS}_{+1, \Delta t_N}), \quad (2)$$

with each time surface \mathbf{TS} , initialized with zeros, defined by polarity $p \in \{-1, +1\}$ and Δt_n

$$\mathbf{TS}_{p, \Delta t_n}(x_i, y_i) = \max_{i, p_i=p} \left(1 - \frac{\tau - t_i}{\Delta t_n} \right). \quad (3)$$

The MCTS generation process is visualized in Figure 3.

3.3. Network Architecture & Loss

As commonly employed in frame-based keypoint detection models [15, 22, 64], we combine a shared backbone with a detector and a descriptor head. Starting with the basic SuperPoint [15] architecture with adjusted loss functions for our temporal matching with pseudo-labels, we train various models with different backbones and hyperparameters and compare their performance. As an alternative approach, we replace the grid-based descriptor head with full resolution prediction and the respective loss functions as in [22, 64]. Please refer to the full results of this ablation study in the **Supplementary Material**.

Our resulting architecture consists of a 3-layer MaxViT [63] backbone with a Feature Pyramid Network (FPN) [43] which is also employed in a highly efficient, event-based detector [20]. Combined with grid-based VGG [61] detector and descriptor heads, SuperEvents fully convolutional components can process various input resolutions without retraining. Figure 4 shows the final network architecture.

The loss function evaluates the network prediction for two corresponding training tensors $\{0, 1\}$. It combines the individual detector losses $L_{s,0}$ and $L_{s,1}$ with the joint descriptor loss L_d weighted with the constant c_λ

$$L = L_{s,0} + L_{s,1} + c_\lambda L_d. \quad (4)$$

Following [15], the detection is implemented as a classification problem of the pixels in image patches, i.e., the pixel indices represent the “classes”. The detector loss L_{CE}

is thus defined as cross-entropy between a predicted patch of keypoint scores $\mathbf{y}_{s_{h,w}}$ (dimension 8×8 plus no-keypoint-dustbin, thus 65) per grid cell (indices h, w) and the pseudo-labels $\hat{\mathbf{y}}_{s_{h,w}}$ in one-hot encoding

$$L_s = \frac{1}{H_c W_c} \sum_{h=1}^{H_c} \sum_{w=1}^{W_c} L_{CE}(\mathbf{y}_{s_{h,w}}, \hat{\mathbf{y}}_{s_{h,w}}), \quad (5)$$

for $H_c \times W_c$ grid cells. In the case of multiple keypoint labels in a grid cell, one is randomly selected.

Unlike other models trained with homographic adaptation, our sparse pseudo-labels do not allow for dense descriptor training for every pixel or grid cell. As in [15], we only consider a single descriptor (256 channels) per grid cell, $\mathbf{y}_{d_{h_0, w_0}}$ and $\mathbf{y}_{d_{h_1, w_1}}$ (normalized to length one across the channels) with (h_0, w_0) and (h_1, w_1) denoting the grid indices of tensor 0 and 1. Successful keypoint matching requires similar descriptors for correspondences and distinct descriptors for unrelated keypoints. We measure descriptor similarity with the dot product $d_{h_1, w_1}^{h_0, w_0} = \mathbf{y}_{d_{h_0, w_0}} \cdot \mathbf{y}_{d_{h_1, w_1}}$. Our pseudo label $\hat{y}_{h_1, w_1}^{h_0, w_0} = 1$ indicates that two descriptors at the respective grid cell locations correspond, otherwise $\hat{y}_{h_1, w_1}^{h_0, w_0} = 0$. Since the majority of the descriptors do not correspond, the partial loss of descriptor matches is weighted with the constant c_d . Mathematically, we formulate the descriptor loss as

$$L_d = \sum_{h_0=1}^{H_c} \sum_{w_0=1}^{W_c} \sum_{h_1=1}^{H_c} \sum_{w_1=1}^{W_c} \begin{cases} c_d \max(0, c_p - d_{h_1, w_1}^{h_0, w_0}), & \hat{y}_{h_1, w_1}^{h_0, w_0} = 1, \\ \max(0, d_{h_1, w_1}^{h_0, w_0} - c_n), & \hat{y}_{h_1, w_1}^{h_0, w_0} = 0. \end{cases} \quad (6)$$

The constants c_p and c_n lead to saturation for maximal desired descriptor (dis-)similarity. All descriptors of grid cells without a keypoint label index ($\mathbf{y}_{d_{h_0, w_0}} = 0$ or $\mathbf{y}_{d_{h_1, w_1}} = 0$) do not contribute to the loss.

During inference, we employ a simplified, yet more effective Non-Minimum-Suppression (NMS) [51] approach in the context of keypoint detection. Instead of the common approach to model keypoints as fixed-size boxes and then apply classical NMS as in [15, 32], we find local maxima in the predicted keypoint score heatmap. In contrast to classical NMS, our method only considers local maxima and not all keypoint candidates sufficiently far from other points with higher scores. Additionally, it employs parallelizable operations instead of consecutive updates of the keypoints sorted by score. Every predicted score s at the pixel coordinates (u, v) is considered for a potential keypoint if

$$s[u, v] > s[u + \Delta u, v + \Delta v] \quad \forall \Delta u, \Delta v \in [-c_b, \dots, -1, 1, \dots, c_b]. \quad (7)$$

with c_b determining the size of the local neighborhood. The keypoint candidates are further filtered by a minimal required score c_s , or the n_s highest scores are selected.

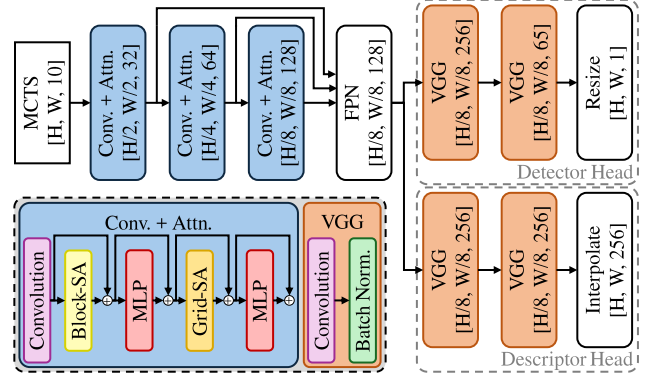


Figure 4. SuperEvent network architecture and tensor dimensions: a shared transformer backbone is combined with a detector and a descriptor head. The components of the Convolution-Attention blocks (*Conv.* + *Attn.*) and the VGG blocks are displayed on the bottom left. Activations are omitted for simplicity.

Table 1. Inference times of our compiled PyTorch model on GPU for common event-stream resolutions (cropped to compatible shape) using float16 precision (same accuracy as float32 in our experiments). NMS requires an additional 0.4 ms on average.

160×240	240×320	480×640	720×1280
1.3 ms	2.6 ms	7.8 ms	23.2 ms

4. Experiments

To validate the effectiveness of the presented approach, we conduct a series of experiments and ablation studies. Firstly, we evaluate the detection and description of keypoints by estimating the camera pose changes based on the matched points. We outperform other keypoint detectors and trackers [10, 32, 33] by a large margin.

In our second experiment, we integrate SuperEvent into the SLAM framework OKVIS2 [40] originally developed for frame-based cameras. We achieve convincing results on the widely used event-SLAM benchmarks [18, 35] without extensive tuning, showcasing SuperEvent’s versatility.

For ablation studies, more qualitative examples, a detailed description of the pose estimation benchmark, and a comparison to frame-based matching for high-speed/HDR scenarios, please refer to our **Supplementary Material**.

4.1. Implementation Details

We generated training data as introduced in Section 3.1. To expose the model to various scenes and camera intrinsics, we combine multiple datasets [14, 31, 39, 56, 69], obtaining more than 100k training sample pairs with pseudo-labels. We implement SuperEvent in PyTorch [4] and train it with the Adam [34] optimizer (learning rate $1e^{-4}$, betas 0.9, 0.999) for 10 epochs.

Table 2. Pose estimation on the Event Camera dataset.

Method	ECD: Pose Estimation AUC in %		
	@5°	@10°	@20°
LLAK [10]	0.7	1.4	2.1
RATE [33]	<u>3.3</u>	<u>8.4</u>	<u>18.0</u>
EventPoint [32]	1.6	3.0	5.4
SuperEvent (ours)	22.7	35.8	46.7

We employ the network architecture as in Section 3.3 with the loss function hyperparameters $c_\lambda = 10$, $c_d = 0.5$, $c_p = 1.0$, and $c_n = 0.2$. We encode events as a 10-channel MCTS with 5 time window sizes Δt logarithmically distributed from 0.001 s to 0.1 s. The detection threshold is set to $c_s = 0.01$ and the NMS neighborhood $c_b = 2$ pixels.

We evaluate on a desktop PC with an Intel Core i5-13600 processor, an NVIDIA GeForce GTX 4070 GPU, and 32 GB RAM. Inference times are shown in Table 1.

4.2. Pose Estimation

A common benchmark for evaluating frame-based keypoint detectors is pose estimation [22, 59, 62, 64]. These works use datasets such as ScanNet [12] and MegaDepth [42] containing various images of the same scene with the associated ground truth camera poses. To create a similar event-based benchmark, we employ sequences with ground truth poses from real-world dynamic datasets. We search for uniformly distributed pose changes between 1° and 45°. The upper bound is chosen to increase the probability of visual overlap. However, since we do not know the ground truth geometry of the scene, it cannot be guaranteed that every sample contains sufficient information to recover the ground truth pose. Next, we process the event stream with various approaches to detect and match keypoints. After outlier removal with Random Sample Consensus (RANSAC) [16], we estimate the camera pose and calculate the error with respect to the ground truth. Finally, we report the area-under-curve (AUC) metric for different thresholds.

We compare the pose estimation capability of SuperEvent to the 3 approaches LLAK [10], RATE [33], and EventPoint [32]. While we use the open-source implementations of LLAK and RATE, we re-implemented and trained EventPoint by carefully following the steps outlined in the paper. We were unable to evaluate the approach SD2Event [19] due to unpublished code and missing implementation details in the respective work.

The experiment is conducted on two diverse datasets: Firstly, the Event Camera Dataset (ECD) [50] is widely used to evaluate event-based approaches, as it was one of the first publicly available datasets. It is challenging due to the low resolution (180×240), the noisy output of the initiation DAVIS240C event camera, and fast motion changes. Secondly, we use the more recent dataset Event-aided Di-

Table 3. Pose estimation on Event-aided Direct Sparse Odometry.

Method	EDS: Pose Estimation AUC in %		
	@5°	@10°	@20°
LLAK [10]	0.5	0.7	1.0
RATE [33]	<u>2.1</u>	<u>5.1</u>	<u>10.3</u>
EventPoint [32]	1.6	2.8	5.2
SuperEvent (ours)	15.2	26.4	40.1

rect Sparse Odometry (EDS) [29]. In contrast to our training data, this dataset is recorded with a Prophesee Gen 3.1² event camera with a higher resolution of 480×640 pixels and an over-proportionally increased event rate. Fast, inconsistent camera motions and a high dynamic range make pose estimation challenging.

Table 2 shows that SuperEvent outperforms all baselines by a large margin. On the ECD dataset, its advantage is the greatest for the high-precision estimations below 5° error, confirming the effectiveness of our method. In contrast, we found that LLAK detects many keypoints but cannot track the vast majority of them. EventPoint’s matching capability is severely limited, possibly due to their static assumptions during training inherently contradicting the dynamic nature of event cameras. Finally, RATE achieves decent results for estimations with up to 20° error, but since it cannot recover from tracking loss, it depends on consistent and precise tracking of most keypoints. As seen in the results, this requirement cannot be stably fulfilled.

Table 3 indicates a slight decrease in SuperEvent’s pose estimation performance on the EDS dataset compared to ECD. Since all approaches perform worse, we attribute this effect to the dataset being more challenging. Interestingly, the performance of SuperEvent for 20° precision is still in the same region as for the ECD dataset. Considering the different properties of the Prophesee camera compared to the DAVIS camera employed in the training sequences, these results confirm the generalization capability of SuperEvent.

4.3. Stereo Event-Visual Inertial SLAM

We demonstrate the usefulness of SuperEvent with a downstream experiment. To date, frame-based visual odometry and SLAM systems largely outperform their event-based counterparts, attributed to their decades-long lead in development. By integrating SuperEvent into the frame-based VI-SLAM system OKVIS2 [40], we enable this framework to process the keypoints detected in the event stream. We compare the performance of the resulting event-only system to other recent event-based Odometry and SLAM approaches on the TUM-VIE dataset [35] in Table 4. Moreover, note that the next best results of the data-driven Mono E(I)O systems DEVO and DEIO are achieved only

²<https://www.prophesee.ai/>

Table 4. Results on the small scale *mocap* sequences of TUM-VIE dataset, ATE in cm. The first three gray results require scale alignment with the ground truth; all other approaches estimate the absolute scale. Baseline numbers taken from [27, 53].

Method	Modality	1d-trans	3d-trans	6dof	desk	desk2	Average
EVO [55] w/ scale alignment	Mono E	7.50	12.50	85.50	54.10	75.20	46.96
DEVO [36] w/ scale alignment	Mono E	0.50	1.10	1.60	1.70	1.00	1.18
DEIO [27] w/ scale alignment	Mono E + IMU	0.42	1.11	1.37	1.36	0.73	1.00
DEIO [27]	Mono E + IMU	1.08	<u>1.12</u>	<u>1.39</u>	<u>1.41</u>	<u>1.19</u>	<u>1.24</u>
ESVO [67]	Stereo E	12.54	17.19	13.46	12.92	4.42	12.11
ES-PTAM [21]	Stereo E	<u>1.05</u>	8.53	10.25	2.50	7.20	5.91
ICRA'24 [54]	Stereo E + IMU	3.85	18.90	failed	8.99	9.47	–
ESVO2 [53]	Stereo E + IMU	3.33	7.26	3.21	6.16	4.02	4.78
OKVIS2 [40] + SuperEvent (ours)	Stereo E + IMU	0.44	0.89	0.43	0.58	0.41	0.55
(ours without loop closure)		(0.43)	(0.89)	(0.43)	(0.70)	(0.40)	(0.57)

Table 5. Results on VECtor [18] large scale sequences. Our results are obtained with online calibration. Baseline numbers are taken from [27] and [36]. DEVO [36] cannot recover the absolute scale. Sequences marked ‘–’ are not evaluated in [27]. ATE in cm.

Method	corr.- dolly	corr.- walk	units- dolly	units- scooter	school- dolly	school- scooter
Ours	33.13	133.16	122.61	59.05	69.96	39.07
DEIO [27]	<u>492.65</u>	<u>325.00</u>	<u>826.38</u>	<u>304.14</u>	–	–
ESVO [67]	failed	failed	failed	failed	<u>1371.0</u>	<u>983.0</u>

after aligning the scale of the estimated trajectories with ground truth. However, real-world applications, e.g., VR and robotics, typically require accurate metric scale. Results on the VECtor dataset [18] are listed in Table 5.

OKVIS2’s loop closure capability relies on long-term descriptor matching. While its effect on the small-scale sequences is limited, we evaluate OKVIS2 + SuperEvent on the *loop-floor* sequences of the TUM-VIE dataset, since these are explicitly designed to test a system’s loop-closure capabilities. Thanks to SuperEvent, the OKVIS2 loop-closure feature works reliably with events, as shown in Table 6. Figure 5 qualitatively compares the estimated trajectories of the sequence *loop-floor0*. Without loop closure, there is significant drift – albeit not out of the ordinary compared to frame-based VI-SLAM. However, with loop closure enabled, the starting room is recognized and the drift is corrected. This demonstrates the robustness of SuperEvent, achieving successful descriptor matching even with minutes of time between observations.

5. Conclusion

We presented SuperEvent, a novel keypoint detector and descriptor for event streams. During its training, we consider the time-dependent keypoint appearance and employ pseudo-labels detected on frames. Combined with the flexible MCTS representation and a transformer back-

Table 6. Effect of loop closure of OKVIS2 [40] + SuperEvent on TUM-VIE *loop-floor* sequences, ATE in cm.

loop-floor	0	1	2	3
Estimated length	349 m	316 m	279 m	303 m
Loop closure	4.96	4.64	8.92	4.74
W/o loop closure	132.11	161.92	116.00	129.17




Figure 5. OKVIS2 + SuperEvent’s estimated trajectories of TUM-VIE sequence *loop-floor0* with and without loop closure.

bone, it achieves SOTA performance on keypoint-based pose estimation. SuperEvent’s integration into frame-based downstream applications enables event vision in existing keypoint-based systems.

As future work, SuperEvent can be combined with a frame-based detector to realize systems that simultaneously process frames and events, exploiting their complementary advantages. Also, with a similar approach to SuperEvent, an event-based line detector could improve the performance of downstream applications in man-made environments.

Finally, our training data is generated to enable consistent feature description independent of their motion-dependent appearance. However, since we employ real sequences, consecutive samples are biased to having similar motion directions. This is reflected in the performance of SuperEvent, whose descriptor matchability suffers under strong motion changes. While this is an inherent property of event cameras, a different approach to generating training data could mitigate it.

References

- [1] Ignacio Alzugaray and Margarita Chli. Asynchronous corner detection and tracking for event cameras in real time. *IEEE Robotics and Automation Letters*, 3(4):3177–3184, 2018. 3
- [2] Ignacio Alzugaray and Margarita Chli. Ace: An efficient asynchronous corner tracker for event cameras. In *2018 International Conference on 3D Vision (3DV)*, pages 653–661. IEEE, 2018. 4
- [3] Ignacio Alzugaray and Margarita Chli. HASTE: Multi-hypothesis asynchronous speeded-up tracking of events. In *31st British Machine Vision Virtual Conference (BMVC 2020)*, page 744. ETH Zurich, Institute of Robotics and Intelligent Systems, 2020. 4
- [4] Jason Ansel, Edward Yang, Horace He, Natalia Gimelshein, Animesh Jain, Michael Voznesensky, Bin Bao, Peter Bell, David Berard, Evgeni Burovski, Geeta Chauhan, Anjali Chourdia, Will Constable, Alban Desmaison, Zachary DeVito, Elias Ellison, Will Feng, Jiong Gong, Michael Gschwind, Brian Hirsh, Sherlock Huang, Kshiteej Kalambarakar, Laurent Kirsch, Michael Lazos, Mario Lezcano, Yanbo Liang, Jason Liang, Yinghai Lu, CK Luk, Bert Maher, Yunjie Pan, Christian Puhrsch, Matthias Reso, Mark Saroufim, Marcos Yukio Siraichi, Helen Suk, Michael Suo, Phil Tillet, Eikan Wang, Xiaodong Wang, William Wen, Shunting Zhang, Xu Zhao, Keren Zhou, Richard Zou, Ajit Mathews, Gregory Chanan, Peng Wu, and Soumith Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)*. ACM, 2024. 6
- [5] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3):346–359, 2008. 3
- [6] Michael Calonder, Vincent Lepetit, Christoph Strecha, and Pascal Fua. Brief: Binary robust independent elementary features. In *Computer Vision–ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part IV 11*, pages 778–792. Springer, 2010. 3
- [7] Carlos Campos, Richard Elvira, Juan J. Gómez, José M. M. Montiel, and Juan D. Tardós. ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021. 1, 4
- [8] Peiyu Chen, Weipeng Guan, and Peng Lu. ESVIO: Event-based stereo visual inertial odometry. *IEEE Robotics and Automation Letters*, 8(6):3661–3668, 2023. 4
- [9] Philippe Chiberre, Etienne Perot, Amos Sironi, and Vincent Lepetit. Detecting stable keypoints from events through image gradient prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1387–1394, 2021. 1, 3
- [10] Philippe Chiberre, Etienne Perot, Amos Sironi, and Vincent Lepetit. Long-lived accurate keypoints in event streams. *arXiv preprint arXiv:2209.10385*, 2022. 1, 3, 6, 7
- [11] Xavier Clady, Sio-Hoi Ieng, and Ryad Benosman. Asynchronous event-based corner detection and matching. *Neural Networks*, 66:91–106, 2015. 3
- [12] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 7
- [13] Laurent Dardet, Ryad Benosman, and Sio-Hoi Ieng. An event-by-event feature detection and tracking invariant to motion direction and velocity. *Authorea Preprints*, 2023. 4
- [14] Jeffrey Delmerico, Titus Cieslewski, Henri Rebecq, Matthias Faessler, and Davide Scaramuzza. Are we ready for autonomous drone racing? the uzh-fpv drone racing dataset. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6713–6719. IEEE, 2019. 6
- [15] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperPoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 1, 2, 3, 4, 5, 6
- [16] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 7
- [17] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conradt, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180, 2020. 3
- [18] Ling Gao, Yuxuan Liang, Jiaqi Yang, Shaoxun Wu, Chenyu Wang, Jiaben Chen, and Laurent Kneip. Vector: A versatile event-centric benchmark for multi-sensor slam. *IEEE Robotics and Automation Letters*, 7(3):8217–8224, 2022. 3, 6, 8
- [19] Yuan Gao, Yuqing Zhu, Xinjun Li, Yimin Du, and Tianzhu Zhang. SD2Event: Self-supervised learning of dynamic detectors and contextual descriptors for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3055–3064, 2024. 1, 3, 5, 7
- [20] Mathias Gehrig and Davide Scaramuzza. Recurrent vision transformers for object detection with event cameras. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13884–13893, 2023. 5
- [21] Suman Ghosh, Valentina Cavinato, and Guillermo Gallego. ES-PTAM: Event-based stereo parallel tracking and mapping. In *European Conference on Computer Vision (ECCV) Workshops*, 2024. 4, 8
- [22] Pierre Gleize, Weiyao Wang, and Matt Feiszli. SiLK: Simple learned keypoints. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 22442–22451, 2023. 2, 3, 5, 7
- [23] Arren Glover, Aiko Dinale, Leandro De Souza Rosa, Simeon Bamford, and Chiara Bartolozzi. IuvHarris: A practical corner detector for event-cameras. *IEEE Transactions on*

- Pattern Analysis and Machine Intelligence*, 44(12):10087–10098, 2021. 3
- [24] Weipeng Guan and Peng Lu. Monocular event visual inertial odometry based on event-corner using sliding windows graph-based optimization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2438–2445, 2022. 4
- [25] Weipeng Guan, Peiyu Chen, Yuhan Xie, and Peng Lu. PL-EVIO: Robust monocular event-based visual inertial odometry with point and line features. *IEEE Transactions on Automation Science and Engineering*, 2023. 4
- [26] Weipeng Guan, Peiyu Chen, Huibin Zhao, Yu Wang, and Peng Lu. EVI-SAM: Robust, real-time, tightly-coupled event–visual–inertial state estimation and 3d dense mapping. *Advanced Intelligent Systems*, 6(12):2400243, 2024. 4
- [27] Weipeng Guan, Fuling Lin, Peiyu Chen, and Peng Lu. DEIO: Deep event inertial odometry. *arXiv preprint arXiv:2411.03928*, 2024. 4, 8
- [28] Christopher G. Harris and M. J. Stephens. A combined corner and edge detector. In *Alvey Vision Conference*, 1988. 3
- [29] Javier Hidalgo-Carrió, Guillermo Gallego, and Davide Scaramuzza. Event-aided direct sparse odometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5781–5790, 2022. 7
- [30] Sumin Hu, Yeeun Kim, Hyungtae Lim, Alex Junho Lee, and Hyun Myung. eCDT: Event clustering for simultaneous feature detection and tracking. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3808–3815. IEEE, 2022. 4
- [31] Yuhuang Hu, Jonathan Binas, Daniel Neil, Shih-Chii Liu, and Tobi Delbruck. DDD20 end-to-end event camera driving dataset: Fusing frames and events with deep learning for improved steering prediction. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE, 2020. 1, 6
- [32] Ze Huang, Li Sun, Cheng Zhao, Song Li, and Songzhi Su. EventPoint: Self-supervised interest point detection and description for event-based camera. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5396–5405, 2023. 1, 3, 5, 6, 7
- [33] Mikihiro Ikura, Cedric Le Gentil, Marcus G. Müller, Florian Schuler, Atsushi Yamashita, and Wolfgang Stürzl. RATE: Real-time asynchronous feature tracking with event cameras. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11662–11669, 2024. 1, 4, 6, 7
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 6
- [35] Simon Klenk, Jason Chui, Nikolaus Demmel, and Daniel Cremers. TUM-VIE: The TUM stereo visual-inertial event dataset. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8601–8608. IEEE, 2021. 3, 6, 7
- [36] Simon Klenk, Marvin Motzet, Lukas Koestler, and Daniel Cremers. Deep event visual odometry. In *2024 International Conference on 3D Vision (3DV)*, pages 739–749. IEEE, 2024. 4, 8
- [37] Beat Kueng, Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Low-latency visual odometry using event-based feature tracks. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 16–23. IEEE, 2016. 4
- [38] Xavier Lagorce, Garrick Orchard, Francesco Galluppi, Bertram E. Shi, and Ryad B. Benosman. HOTS: A hierarchy of event-based time-surfaces for pattern recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(7):1346–1359, 2017. 2, 3, 5
- [39] Alex Junho Lee, Younggun Cho, Young-sik Shin, Ayoun Kim, and Hyun Myung. ViViD++: Vision for visibility dataset. *IEEE Robotics and Automation Letters*, 7(3):6282–6289, 2022. 6
- [40] Stefan Leutenegger. OKVIS2: Realtime scalable visual-inertial SLAM with loop closure, 2022. 1, 2, 4, 6, 7, 8
- [41] Stefan Leutenegger, Margarita Chli, and Roland Y. Siegwart. BRISK: Binary robust invariant scalable keypoints. In *2011 International Conference on Computer Vision*, pages 2548–2555, 2011. 3
- [42] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2041–2050, 2018. 7
- [43] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 5
- [44] Zhe Liu, Dianxi Shi, Ruihao Li, and Shaowu Yang. ESVIO: event-based stereo visual-inertial odometry. *Sensors*, 23(4):1998, 2023. 4
- [45] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004. 3
- [46] Florian Mahlknecht, Daniel Gehrig, Jeremy Nash, Friedrich M Rockenbauer, Benjamin Morrell, Jeff DeLaune, and Davide Scaramuzza. Exploring event camera-based odometry for planetary robots. *IEEE Robotics and Automation Letters*, 7(4):8651–8658, 2022. 4
- [47] Jacques Manderscheid, Amos Sironi, Nicolas Bourdis, Davide Migliore, and Vincent Lepetit. Speed invariant time surface for learning to detect corner points with event-based cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10245–10254, 2019. 1, 3
- [48] Nico Messikommer, Carter Fang, Mathias Gehrig, and Davide Scaramuzza. Data-driven feature tracking for event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5642–5651, 2023. 4, 5
- [49] Elias Mueggler, Chiara Bartolozzi, and Davide Scaramuzza. Fast event-based corner detection. *Proceedings of the British Machine Vision Conference*, 2017. 3
- [50] Elias Mueggler, Henri Rebecq, Guillermo Gallego, Tobi Delbruck, and Davide Scaramuzza. The event-camera dataset and simulator: Event-based data for pose estimation, visual

- odometry, and slam. *The International Journal of Robotics Research*, 36(2):142–149, 2017. 7
- [51] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *18th international conference on pattern recognition (ICPR'06)*, pages 850–855. IEEE, 2006. 6
- [52] Zhenjiang Ni, Aude Bolepion, Joël Agnus, Ryad Benosman, and Stéphane Régnier. Asynchronous event-based visual shape tracking for stable haptic feedback in microrobotics. *IEEE Transactions on Robotics*, 28(5):1081–1089, 2012. 4
- [53] Junkai Niu, Sheng Zhong, Xiuyuan Lu, Shaojie Shen, Guillermo Gallego, and Yi Zhou. ESVO2: Direct visual-inertial odometry with stereo event cameras. *arXiv e-prints*, pages arXiv:2410, 2024. 4, 8
- [54] Junkai Niu, Sheng Zhong, and Yi Zhou. Imu-aided event-based stereo visual odometry. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11977–11983, 2024. 4, 8
- [55] Henri Rebecq, Timo Horstschäfer, Guillermo Gallego, and Davide Scaramuzza. EVO: A geometric approach to event-based 6-dof parallel tracking and mapping in real time. *IEEE Robotics and Automation Letters*, 2(2):593–600, 2016. 4, 8
- [56] Juan Pablo Rodríguez-Gómez, Raul Tapia, Julio L Paneque, Pedro Grau, Augusto Gómez Eguíluz, Jose Ramiro Martínez-de Dios, and Anibal Ollero. The GRIFFIN perception dataset: Bridging the gap between flapping-wing flight and robotic perception. *IEEE Robotics and Automation Letters*, 6(2):1066–1073, 2021. 6
- [57] Edward Rosten and Tom Drummond. Machine learning for high-speed corner detection. In *Computer Vision—ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006. Proceedings, Part I 9*, pages 430–443. Springer, 2006. 3
- [58] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. ORB: An efficient alternative to SIFT or SURF. In *2011 International Conference on Computer Vision*, pages 2564–2571, 2011. 3
- [59] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. 1, 2, 3, 4, 7
- [60] Jianbo Shi and Tomasi. Good features to track. In *1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 593–600, 1994. 3, 4
- [61] Karen Simonyan. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3, 5
- [62] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8922–8931, 2021. 2, 3, 7
- [63] Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. MaxViT: Multi-axis vision transformer. In *European conference on computer vision*, pages 459–479. Springer, 2022. 5
- [64] Michał Tyszkiewicz, Pascal Fua, and Eduard Trulls. DISK: Learning local features with policy gradient. *Advances in Neural Information Processing Systems*, 33:14254–14265, 2020. 2, 3, 5, 7
- [65] Valentina Vasco, Arren Glover, and Chiara Bartolozzi. Fast event-based harris corner detection exploiting the advantages of event-driven cameras. In *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4144–4149. IEEE, 2016. 3
- [66] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschäfer, and Davide Scaramuzza. Ultimate SLAM? Combining events, images, and IMU for robust visual slam in hdr and high-speed scenarios. *IEEE Robotics and Automation Letters*, 3(2):994–1001, 2018. 4
- [67] Yi Zhou, Guillermo Gallego, and Shaojie Shen. Event-based stereo visual odometry. *IEEE Transactions on Robotics*, 37(5):1433–1450, 2021. 4, 8
- [68] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based feature tracking with probabilistic data association. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4465–4470. IEEE, 2017. 4
- [69] Alex Zihao Zhu, Dinesh Thakur, Tolga Özaslan, Bernd Pfrommer, Vijay Kumar, and Kostas Daniilidis. The multi-vehicle stereo event camera dataset: An event camera dataset for 3d perception. *IEEE Robotics and Automation Letters*, 3(3):2032–2039, 2018. 6
- [70] Alex Zihao Zhu, Liangzhe Yuan, Kenneth Chaney, and Kostas Daniilidis. Unsupervised event-based optical flow using motion compensation. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018. 5