

ETAP: Event-based Tracking of Any Point

Friedhelm Hamann^{1,4,6}, Daniel Gehrig², Filbert Febryanto^{1,4,6},
 Kostas Daniilidis^{2,5}, and Guillermo Gallego^{1,3,4,6}.

¹ Technische Universität Berlin, ² University of Pennsylvania, ³ Einstein Center for Digital Future,
⁴ Robotics Institute Germany, ⁵ Archimedes, Athena RC, ⁶ Science of Intelligence Excellence Cluster.

Abstract

Tracking any point (TAP) recently shifted the motion estimation paradigm from focusing on individual salient points with local templates to tracking arbitrary points with global image contexts. However, while research has mostly focused on driving the accuracy of models in nominal settings, addressing scenarios with difficult lighting conditions and high-speed motions remains out of reach due to the limitations of the sensor. This work addresses this challenge with the first event camera-based TAP method. It leverages the high temporal resolution and high dynamic range of event cameras for robust high-speed tracking, and the global contexts in TAP methods to handle asynchronous and sparse event measurements. We further extend the TAP framework to handle event feature variations induced by motion — thereby addressing an open challenge in purely event-based tracking — with a novel feature alignment-loss which ensures the learning of motion-robust features. Our method is trained with data from a new data generation pipeline and systematically ablated across all design decisions. Our method shows strong cross-dataset generalization and performs 136% better on the average Jaccard metric than the baselines. Moreover, on an established feature tracking benchmark, it achieves a 20% improvement over the previous best event-only method and even surpasses the previous best events-and-frames method by 4.1%. Our code is available at <https://github.com/tub-rip/ETAP>.

1. Introduction

Understanding scene motion from a video remains a fundamental challenge in computer vision, with renewed interest through its formulation as tracking any point (TAP) [13, 25, 54]. A new class of powerful methods has been quickly adopted for downstream tasks, e.g., in robotics [7, 59]. However, existing methods focus on tracking in nominal settings due to the fundamental limitations of the sensor.

Event cameras represent a novel class of visual sensors offering high temporal resolution, high dynamic range

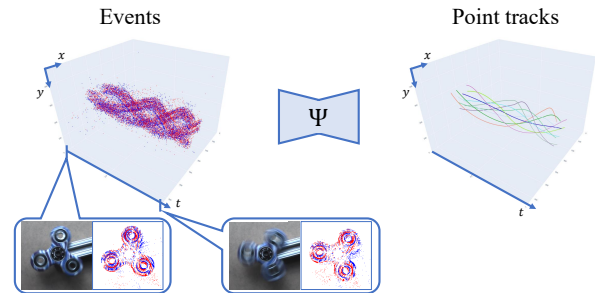


Figure 1. *Event-only Tracking of Any Point.* Our method uses only events to track semi-dense long-range point trajectories, working in conditions where frame-based methods fail.

(HDR), and low power consumption, characteristics that make them valuable stand-alone sensors for various robotic perception tasks. These innovative sensors address several limitations of conventional cameras, particularly handling motion blur and high-speed movements. To fully leverage these advantages we focus on developing an event-only method for tracking arbitrary 2D points within the data stream without additional sensor input.

Estimating point trajectories for arbitrary scene motion presents significant challenges, with recent solutions emerging primarily through deep neural networks trained on synthetic data. Progress in dense optical flow [15, 58] and subsequently point tracking have been driven by supervised learning on rendered datasets, which provide ground truth (GT) scene motion. Event simulators have enabled a conceptually similar workflow for event-based vision algorithms. Using the event generation model, events - representing pixel-wise intensity changes - can be synthesized from high-frame-rate videos. The input video can either be real (temporally upsampled through video interpolation) or synthetic. For motion estimation tasks, this approach has been applied to dense optical flow estimation [22] and sparse feature tracking [44]. However, the synthetic datasets used for training are simplistic warps of 2D objects lacking realism and limiting performance.

The feasibility of training motion-estimation networks

using synthetic data has been established for conventional cameras. Networks demonstrate good generalization, partly by exploiting the correlation of appearance features between timesteps. With event simulators, similar approaches could be extended to event cameras, though several unique questions must be addressed. A key challenge in event-based tracking using synthetic data stems from the immaturity of simulation tools. While it is possible to combine frame-based physically based rendering (PBR) pipelines with video-to-event conversion tools, this process involves numerous parameters that require tuning to achieve optimal results. A second challenge with event camera data is its inherent motion dependence. Consider a simple scenario, illustrated in Fig. 2, showing two recordings of an identical scene (e.g. shapes on a wall) with perpendicular camera motions. In the first recording, the camera moves horizontally, in the second, vertically. With a conventional camera, both scenarios would capture nearly identical images (aside from a slight offset). However, an event camera produces markedly different signals in each case due to its motion-dependent nature. This poses a unique challenge for algorithms that rely on feature correlation, as the feature extractor must be invariant not only to appearance changes and geometric transformations but *also* to scene motion.

We present, to the best of our knowledge, the first model for event-based tracking of any point (ETAP). The model tracks several points in parallel, iteratively updating position and appearance features for each point through spatial and temporal attention blocks. It processes event stacks (grid representations compatible with convolutional feature encoders), constructed at specified tracking timesteps.

The network is trained on a newly developed synthetic dataset. Our data generation pipeline combines Kubric [23] and Vid2e [20]. Through a systematic evaluation of each design decision (including threshold selection, scene dynamics, and render frame rate) we demonstrate that our EventKubric dataset improves performance by 8% (measured by feature age on the Event-aided Direct Sparse Odometry dataset [26] (EDS)) compared to the same model trained on the strong baseline approach using pre-rendered Kubric Multi-Object Video (MOVi) - F dataset.

We also introduce a novel contrastive loss that promotes motion-robust feature extraction in our network. For each training sample, we generate a variant with inverted time and random rotation while preserving appearance. This transformation maintains the scene structure but inverts motion direction. We extract spatial feature maps from both representations, interpolate features at tracked points, and reward high cosine similarity between corresponding feature vectors. Our feature alignment-loss encourages the generation of motion-invariant correlation features.

We evaluate two tasks, TAP (Task 1) and additionally on feature tracking (Task 2) for comparison with previous




Figure 2. *The motion dependence problem.* Many tracking methods rely on the correspondence of features. While the appearance of frames (left) is independent of the scene movement, the event camera data depends on the motion direction. Image courtesy of [1].

event-based methods. TAP is evaluated on EventKubric, the Extreme Event Decompression Dataset (E2D2) [60] (for which we provide new ground truth), and on custom sequences recorded with a beamsplitter system for fair comparison between event- and frame-based algorithms. Feature tracking is evaluated on an established benchmark (comprising of EDS and the “event camera dataset” (EC)), where ETAP achieves significant improvements over previous event-only methods (20%) and surpassing the best method combining frames and events by 4.1%.

Our contributions are summarized as follows:

1. The first event-only tracking-any-point (TAP) method, with SOTA results on two tasks (TAP and feature tracking) and extensive evaluation on six datasets (EVIMO2, EDS, EC, E2D2, EventKubric, Aviary).
2. A new synthetic event dataset (EventKubric) that enables robust tracking performance, with a thorough empirical evaluation of key design decisions.
3. For evaluation, we release new ground truth for EVIMO2 and E2D2 sequences, as well as a challenging aviary sequence for qualitative evaluation.
4. A novel contrastive feature-alignment loss that promotes motion-robust feature extraction from event data.

The experiments show strong cross-dataset generalization to different camera types and resolutions, with outstanding tracking capabilities in a variety of conditions.

2. Related Work

Motion Estimation and Point Tracking. Visual motion estimation remains a central theme of general scene understanding, which has, throughout the years, developed into a diverse field of study. Early paradigms [4, 40] focused on estimating the long-term motion of distinctive patches throughout a set of images using an autoregressive template tracking framework. They successively estimated warps from template to target patches in each new image, by minimizing the change in appearance and then updated these templates at each step. In the real world, however, image patches often *do change in appearance* [42] or distort in complex ways that require the development of complex warping models [34]. Moreover, image patches with

Dataset	Source	Events	#Samples	Resolution	fps [Hz]	sample duration [s]	IMO	Annotations			
								optical flow	depth	point tracking	segmentations
TAP-Vid Kubric, MOVI-F	3D PBR	none	≈ 10000	512×512	12	2	✓	✓	✓	✓	✓
BlinkFlow [35]	3D PBR	synthetic	3587	640×480	10	1	✓	✓	✓	✗	✓
MultiFlow [22]	2D warp	synthetic	12100	512×384	100	0.5	✓	✓*	✗	✗	✗
EventKubric (Ours)	3D PBR	synthetic	10173	512×512	48	2	✓	✓	✓	✓	✓

Table 1. *Dataset comparison.* Overview of publicly available synthetic motion estimation event datasets.

few gradients often provide insufficient constraints to accurately estimate motion, due to the aperture problem. Increasing the context via variational approaches that optimize a global objective [27] can address this but at the cost of over-smoothing the result.

Since the advent of deep learning, this context is now captured via deep architectures with large receptive fields [28, 58] and regularized by implicit priors learned from data. This enabled the tracking of large semantic object bounding boxes [16, 48], action bounding boxes [8], or deformable semantic masks [31, 49, 61], but these are often constrained to specific object classes and not generally applicable to, for example, object parts or single points.

Recently, tracking *single points* has gained traction, due to its flexibility in addressing arbitrary structures. For a given point set in a frame, it predicts their corresponding positions in other frames jointly with explicit visibility. After the early model-based approach Particle Video [54] it was re-introduced by “Particle Video Revisited” (PIPs) [25]. While leveraging many of the early insights for long-term feature tracking such as feature correlation, appearance change modeling, and autoregressive tracking, it did so with modern tools like learning-based feature correlation and a lookup originally designed for optical flow [58].

A key driver of this field has been the curation of large-scale synthetic data: The usage of simulated data is scalable, supports dense motion annotations, provides controllable data complexity, and poses fewer problems regarding privacy and licensing. FlyingChairs [15] and FlyingThings [43] are early datasets widely used for training of optical flow methods, while Kubric provides a flexible dataset generator [23] for large-scale training of point tracking. Follow-up work [13, 14] introduced TAP-Vid, a set of synthetic and real-world datasets which form a common benchmark today, along with methods TAP-Net [13] and TAPIR [14] which innovated on the original design of PIPs. Since then PointOdyssey [62] appeared, which enhances the realism of the synthetic sequences, and provides additional annotations beyond point tracks. These benchmarks sparked the development of methods like LocoTrack [10] and CoTracker [29], which are the state-of-the-art in point tracking. The work in [29], for instance, uses a single model to track several points in parallel leveraging spatial attention between points to model interrelation between them.

Event Camera-based Motion Estimation. Despite these developments, image-based tracking still suffers from fundamental limitations of frame-based sensors, namely a limited framerate, motion blur, and saturation artifacts in challenging lighting conditions, which cause visual aliasing and algorithm degradation. Event cameras [36, 50] are relatively new vision sensors, which can address these issues with their higher dynamic range, limited motion blur, and ability to capture sparse and asynchronous *changes* in the visual data, also called *events*, in continuous time [17, 55].

Similar to image-based tracking, early methods for tracking with events focused on tracking blobs [37] or simple patterns [33, 47]. They use iterative closest point (ICP) [32] or expectation Maximization (EM) [63] to align small spatio-temporal event volumes or perform multi-hypothesis tracking [2] that predict the feature motion. A main challenge in event camera-based tracking is the dependence of feature appearance on camera and object motion, which limits the use of purely appearance-based trackers. To address this, appearance refinement [2, 57], auxiliary sensors providing motion-invariant appearance [18, 21, 32] and data-driven approaches have been explored [9, 38, 41, 44]. Despite their promise, these methods have their limitations: Refinement and learning-based point trackers still use simple synthetic datasets based on moving 2D planes [22, 38, 44] which show only a weak transfer to the real world, and thus necessitate self-supervised finetuning [24, 44]. On the other hand, methods using frames and events, such as [38] specifically combine events and frames in a data-driven approach for point tracking but inherit some of the shortcomings of frames during high-speed motion and in challenging lighting conditions.

In this work, we perform purely event-based tracking and are free of these limitations. Moreover, we provide a large-scale, realistic point-tracking dataset for events, summarized in Tab. 1. It enables the learning of powerful priors, together with our novel contrastive feature alignment loss to explicitly enforce motion-independent features across time.

3. Tracking Any Point With an Event Camera

Problem Formulation. Let us formalize the TAP task with event cameras. These sensors measure so-called *events*, i.e. per-pixel brightness changes $e_k \doteq (\mathbf{x}_k, \tau_k, p_k)$ where

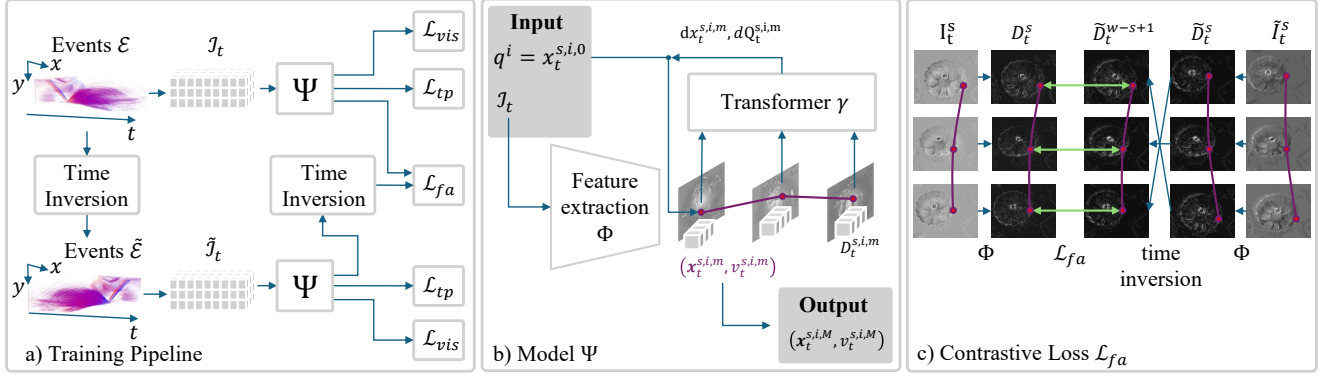


Figure 3. (a) During training each sample has a time inverted duplicate, model Ψ outputs tracks, visibility flags, and descriptors used to calculate the 3 loss terms. (b) The model Ψ takes query points and event stacks as input and extracts spatial feature maps at each timestep. tokens built per point and timestep are iteratively updated. (c) A visualization of the FA-loss. Descriptors of the inverted second sample are time-reversed to receive matching pairs for the FA-loss

$\mathbf{x}_k = (x_k, y_k)^\top$ is the pixel the event is produced, $\tau_k \in \mathbb{R}$ is its timestamp with μs resolution and $p_k \in \{-1, 1\}$ is its polarity (the sign of the brightness change). Each event e_k is triggered when the logarithmic brightness at pixel \mathbf{x}_k exceeds a threshold C , called contrast sensitivity. Over a time interval $\mathcal{T} = (\tau_s, \tau_e)$ the event camera thus outputs asynchronous events $E = \{e_k\}$ at different pixels.

Next, let $P(\tau) = \{(\mathbf{x}^i(\tau), Q^i(\tau), v^i(\tau))\}_{i=1}^{N_p}$ be a set of N_p points moving over time $\tau \in \mathcal{T}$, where $\mathbf{x}^i(\tau) \in \mathbb{R}^2$ is the pixel position of point i , $Q^i(\tau) \in \mathbb{R}^d$ is its descriptor and $v^i(\tau) \in \{0, 1\}$ is its visibility. The descriptors Q^i are 1D feature vectors used to estimate visual similarity with per-timestamp feature maps. Separate descriptors of the same point at different times, allow modeling appearance changes. Note that points may be initialized asynchronously, and thus the cardinality of $P(\tau)$ may not remain constant over time.

Following the formalism in [22] we focus on the point positions at discrete time instances τ_t with $t = 0, 1, \dots, T$ and define their position at these instances by $P_t \doteq P(\tau_t)$. Similarly, we select windows of events

$$E_t = \{e_k \mid \tau_k \in (\tau_t - \Delta\tau_t, \tau_t)\} \subset E \quad (1)$$

that are temporally aligned with τ_t , where $\Delta\tau_t$ is the time span of E_t , which contains a constant number of events N_e . We consider a sliding window of such point observations $\mathcal{P}_t \equiv P_{t-w+1:t}$, with window size $w = 8$, as well as the sequence $\mathcal{E}_t \equiv E_{t-w+1:t}$ defined similarly. We formalize TAP as finding the function Ψ that estimates the tracks \mathcal{P}_t from events \mathcal{E}_t and past tracks \mathcal{P}_{t-T_s}

$$\mathcal{P}_t = \Psi(\mathcal{P}_{t-T_s}, \mathcal{E}_t). \quad (2)$$

where $T_s = 4$ is the stride.

Feature Representations. In practice, each event window E_t is replaced by event representations $I_t = \mathcal{F}(E_t) \in$

$\mathbb{R}^{H \times W \times B}$ [46] where H and W are the sensor’s height and width, and $B = 10$ is the number of time bins. [12, 19, 64].

We extract multi-scale d -dimensional features $D_t^\lambda \in \mathbb{R}^{\frac{H}{k2^{\lambda-1}} \times \frac{W}{k2^{\lambda-1}} \times d}$ from tensors I_t using an encoder $\phi(I_t)$, with subsequent average pooling. $\lambda = 1, \dots, S$ (with $S = 4$) is the scale, and $k = 4$ an overall reduction in the resolution.

Initialization. We manually provide query points $q^i = \mathbf{x}_{t_i}^i$ at time indices t_i , and before the subsequently explained transformer-based refinement, broadcast the points to all timesteps $\mathbf{x}_t^i = \mathbf{x}_t$ of the sliding window, where a point is initialized. Similarly, the descriptors Q_t^i are initialized via the broadcast $Q_t^i = Q_{t_i}^i$, with $Q_{t_i}^i = \text{BilinearInterp}(D_{t_i}^\lambda, \mathbf{x}_{t_i}^i)$, where $\text{BilinearInterp}(\cdot, \cdot)$ samples the feature map $D_{t_i}^\lambda$ at continuous coordinates $\mathbf{x}_{t_i}^i$ using bilinear interpolation.

Tracker. We implement the tracker (2) Ψ following [29]. For simplicity, we omit the global timestep t and regard only point positions within one sliding window $P_s^i \doteq P_{t-w+s}^i$ with relative window index $s = 1, \dots, w$, and $\mathcal{P} \doteq P_{1:w}^i$. We denote $\mathcal{D} \doteq D_{t-w+s:t}^\lambda$ as the according feature maps defined in the same interval as the point tracks \mathcal{P} .

Specifically, the tracker iteratively refines pixel positions \mathbf{x}_s^i and descriptors \tilde{Q}_s^i of the i^{th} point via

$$(\text{d}\tilde{\mathbf{x}}_s^{i,m}, \text{d}\tilde{Q}_s^{i,m}) = \gamma(\mathcal{P}^m, \mathcal{D}) \quad (3)$$

$$\tilde{\mathbf{x}}_s^{i,m+1} = \tilde{\mathbf{x}}_s^{i,m} + \text{d}\tilde{\mathbf{x}}_s^{i,m} \quad (4)$$

$$\tilde{Q}_s^{i,m+1} = \tilde{Q}_s^{i,m} + \text{d}\tilde{Q}_s^{i,m} \quad (5)$$

Note that $\tilde{Q}_s^{i,m}$ and $\tilde{\mathbf{x}}_s^{i,m}$ denote the descriptor and position at iteration m , and \mathcal{P}^m are the tracks with updated descriptor and point position. The update step is iterated M times, to obtain the final position estimates $\tilde{\mathbf{x}}_s^{i,M}$. After the last iteration, visibilities are computed with a simple linear layer via $v_s^i = \sigma(\Theta \tilde{Q}_s^{i,M})$.

We implement γ as a transformer that operates on tokens $\mathcal{O}_s^{i,m}$ indexed by relative time s and point index i via alter-

nating intra-point attention (across index i), and temporal attention (across index s). At each iteration m we compute these tokens as the following concatenation:

$$\mathcal{O}_s^{i,m} = (\eta(\Delta \mathbf{x}_s^{i,m}), Q_s^{i,m}, C_s^{i,m}, v_s^i) + \eta'(\mathbf{x}_0^{i,m}) + \eta'(t)$$

with $\Delta \mathbf{x}_s^{i,m} = \mathbf{x}_s^{i,m} - \mathbf{x}_1^{i,m}$, positional encodings η, η' , and spatial correlation features $C_s^{i,m}$, which are discussed next.

Correlation Features. As part of the input tokens to the transformer, we provide information on the similarity of descriptors $\tilde{Q}_s^{i,m}$ to points of their surroundings. The correlation features within a patch $B_\Delta = \{\delta \in \mathbb{Z}^2 \mid \|\delta\|_\infty \leq \Delta\}$ are calculated via the inner products

$$C_s^{i,m} = \oplus_{\lambda=1}^S \oplus_{\delta \in B_\Delta} \left\langle \tilde{Q}_s^{i,m}, D(\tilde{\mathbf{x}}_s^{i,m}/k\lambda + \delta) \right\rangle \quad (6)$$

Here we are concatenating correlations within the patch B_Δ of size $|B_\Delta| = (2\Delta + 1)^2 = 49$ across four scales, resulting in a feature dimension of 196. At each iteration m of the transformer refinement, we use the updated point locations $\tilde{\mathbf{x}}_s^{i,m}$ to compute these features.

Motion Robust Event Features. Event camera data is inherently motion-dependent, unlike conventional cameras where the same scene produces the same signal regardless of motion (Fig. 2). Based on the linearized event generation model (LEGM), we can show that under time inversion, the events E_t and events of the inverted scene \tilde{E}_t are not the same (see Supplementary for mathematical derivation), and consequently, the corresponding descriptors D_t^s and \tilde{D}_t^{w-s+1} are different (note $w - s + 1$ is the inverted index). We explicitly enforce motion consistency by maximizing the similarity of descriptors $d_t^{s,i} \doteq D_t^s(\mathbf{x}_t^{s,i})$ and $\tilde{d}_t^{s,i} \doteq \tilde{D}_t^{w-s+1}(\mathbf{x}_t^{w-s+1,i})$ sampled at track positions $\mathbf{x}_k^{s,i}$ and $\mathbf{x}_k^{w-s+1,i}$, under different motions, and incorporate this into the loss function:

$$\mathcal{L}_{\text{fa}} = \sum_t \frac{1}{|\mathcal{P}_t|} \sum_{i,s} \left(1 - \langle \mathbf{u}(d_t^{s,i}), \mathbf{u}(\tilde{d}_t^{s,i}) \rangle \right)^2 \quad (7)$$

where $\mathbf{u}(\mathbf{a}) = \mathbf{a}/\|\mathbf{a}\|$ unitizes a vector, and $|\mathcal{P}_t|$ counts the number of points within the time window. To further enhance the diversity of motion, we randomly rotate the events \tilde{E} by an angle $\theta \in \{0, 90^\circ, 180^\circ, 270^\circ\}$.

We supplement this loss with the track prediction error \mathcal{L}_{tp} , which penalizes the absolute difference between predicted and GT tracks at each refinement step m weighted by 0.8^{M-m} . We also include the visibility loss \mathcal{L}_{vis} , which is the cross-entropy on predicted visibility flags [29]. Our total loss is calculated as $\mathcal{L} = 0.1\mathcal{L}_{\text{tp}} + \mathcal{L}_{\text{vis}} + 0.1\mathcal{L}_{\text{fa}}$.

4. Dataset Generation

Our model training relies on a new synthetic dataset created through a three-step process: First, we render short video




Figure 4. *Task 1 - TAP on EVIMO2.* Visualization of track predictions from first to last timestamp.




Figure 5. *Task 1 - TAP on EventKubric.* Semi-dense tracks are predicted for 2s-samples.

Method	Input	Metrics		
		AJ \uparrow	$\delta_{avg}^x \uparrow$	OA \uparrow
EVIMO2				
E2Vid [52] + CoTracker [29]	Events	0.531	0.663	0.861
ETAP w/o FA-loss (Ours)	Events	<u>0.655</u>	<u>0.787</u>	<u>0.884</u>
ETAP (Ours)	Events	0.661	0.789	0.895
EventKubric (synthetic)				
E2Vid [52] + CoTracker [29]	Events	0.236	0.331	0.815
ETAP w/o FA-loss (Ours)	Events	0.556	0.677	0.894
ETAP (Ours)	Events	<u>0.546</u>	<u>0.675</u>	<u>0.890</u>

Table 2. *Task 1 - TAP evaluation on EventKubric and EVIMO.*

clips using Kubric [23], then adaptively upsample them using FILM [53], and finally convert the resulting high-frame-rate video to events using ESIM [51]. For each sample, we generated 2048 point tracks derived from Kubric’s ground truth data. The complete dataset comprises 10,173 samples, split randomly into 80:15:5 ratios for training, validation, and testing. Representative samples from the training set are provided in the Supplementary material.

Physics-based rendering. Using Kubric [23] we render 2-s videos at 48 FPS resulting in 96 frames with 512×512 -px resolution. We opt for a higher FPS than the available Kubric datasets, and disable motion blur to reduce the error introduced through upsampling and event generation. Scenes contain approximately 20 3D rigid objects under gravity simulated with the BULLET [11] physics engine and ray-tracing with Blender [5]. We generate 60% of samples with linear camera movement and 40% with pan-

Method	Input	Metrics		
		AJ \uparrow	$\delta_{avg}^x \uparrow$	OA \uparrow
CoTracker [29]	F	0.007	0.117	0.1
E2Vid [52] + CoTracker [29]	E	0.183	0.262	0.938
ETAP w/o FA-loss (Ours)	E	<u>0.268</u>	<u>0.406</u>	<u>0.826</u>
ETAP (Ours)	E	0.308	0.466	0.769

Table 3. Task 1 - TAP on fidget spinner (E2D2 dataset).

ning movements as used in [14], mimicking natural camera movements found in many real datasets.

Synthetic Event Generation. Due to the computational cost of ray-traced rendering and Kubric’s fixed frame rate constraint, we adopt the VID2E [20] workflow, employing adaptive neural frame interpolation, such that the maximum optical flow between consecutive upsampled frames is at most one pixel, following [51]. After upsampling, we generate events using random contrast sensitivities $C \sim \mathcal{U}(0.16, 0.34)$ as in [30].

Point Track Generation. Since Kubric doesn’t directly provide point tracks, we compute them from the depth, segmentation, and surface normal GT. Following [13], we randomly sample 2048 GT tracks, ensuring adequate object coverage (compared to a high background pixels portion).

5. Experiments

After showing implementation details in Sec. 5.1 explains, we validate our method on two tasks: TAP (Sec. 5.2) and feature tracking (Sec. 5.3). The main technical difference of TAP is the explicit prediction of visibility flags. Additional evaluation of feature tracking allows for a thorough comparison with state-of-the-art event-based methods on established benchmarks. Lastly, we present ablations in Sec. 5.4.

5.1. Implementation Details

Our model was trained on 4 NVIDIA A100 80GB GPUs with a batch size of 2, resolution of 512×512 , and the AdamW optimizer [39] with a learning rate of $5 \cdot 10^{-4}$ and weight decay 10^{-4} . Each sample comprises 256 trajectories with length 24. For the first 10^5 steps, we optimize only the track prediction and visibility loss and then add \mathcal{L}_{fa} for $1.2 \cdot 10^5$ steps. For training, on EventKubric we use $N_e = 4 \cdot 10^5$ events (for more info see Suppl.Mat.).

The event stacks undergo std-mean normalization, computed across batch and time dimensions but independently for each channel to accommodate the varying event counts. We apply Gaussian noise augmentation with $\sigma = 0.1$ (event counts) for the first channel, scaling according to the event count for the remaining channels.

5.2. Task 1: TAP

Results on EVIMO2. We evaluate point tracking on real event data using EVIMO2 [6], creating new ground truth




Figure 6. Task 1 - TAP on E2D2. Shown are four timesteps of each sequence. At the beginning (t_0) the model is queried with the marked points. Input modality: F - frames, E - Events

tracks from its motion capture data. Our approach mirrors the EVIMO2 Continuous Flow Dataset [24] methodology, differing only in our generation of long-term tracks with occlusion flags. Tests use Samsung event camera data (640×480 px) featuring independently moving rigid objects in both dynamic scenarios and static conditions where event-based methods typically struggle due to their motion dependence. Figure 4 and Tab. 2 show the strong performance of our method predicting long tracks and occlusions.

Results on EventKubric. We evaluate performance on our EventKubric test split (501×2 s samples), comparing against CoTracker [29] operating on E2VID [52] images. We select 24 evenly distributed tracking timestamps within each sample and assess performance using standard TAP-metrics [13]: Average-pixel-within-threshold δ_{avg}^x measures the fractions of visible points within a threshold at several levels (1, 2, 4, 8, 16px in 512×512 resolution), occlusion accuracy (OA) is the fraction of correct visibility corrections v_t^i and average Jaccard (AJ), combines both, where a point is correctly predicted when it is within the threshold and has correct visibility prediction. Quantitative results on Tab. 2 demonstrate that our model learns stable track and occlusion prediction, surpassing the E2Vid+CoTracker baseline by 136%. Fig. 5 shows an example prediction.

Results on E2D2. To compare the limitations of frame and event-based algorithms, we design an experiment using the recent Extreme Event Decompression Dataset [60] (E2D2), which provides synchronized frames and events from a beamsplitter system. We utilize this setup for fair




Figure 7. *Task 1 - TAP qualitative result.* Tracking in a very demanding scenario: a small, low-textured, fast object with high deformation and an HDR background.

cross-modality comparison, specifically picking a scene of a rotating fidget spinner (Fig. 6, where angular velocity increases rapidly over 0.5 seconds, progressively challenging tracking performance). The scene, recorded under low light conditions, limits the frame-based camera to 10 Hz due to exposure constraints. Similarly, the event camera data exhibits significant noise and shadow-induced artifacts. We generate 330 Hz GT tracks for query points on the fidget spinner using angular velocity estimates (see Suppl.Mat.). We compare against two methods, one is CoTracker [29] run on the RGB frames and a second time on E2VID images reconstructed at the GT timestamps. The results show that the frame-based tracking method fails due to two factors: aliasing from the wheel’s rotation combined with low frame rate, and severe motion blur from extended exposure times. In contrast, the events capture information at sufficient temporal resolution. Quantitative comparison in Tab. 3 shows the comparison of the two event-based methods shows that our ETAP produces significantly more accurate tracks, surpassing the event-based baseline by 68% AJ.

Qualitative Analysis. We perform qualitative analysis on recordings from E2D2 and additional recordings that were done in a similar manner using a beamsplitter system. Fig. 7 shows one example. The sequence of a bird, with RGB frames at 66 Hz is very challenging with a small, fast-moving, highly deforming target with little structure and an HDR background. The comparison with CoTracker [29] a state-of-the-art frame-based model, shows that our model performs better in tracking the target and capturing details.




Figure 8. *Task 2 - Feature tracking on EDS.* Notably, our tracker captures points even after they leave the FOV and reenter.

Method	Input	EDS		EC	
		Feature Age \uparrow	Expected FA \uparrow	Feature Age \uparrow	Expected FA \uparrow
ICP [32]	E	0.060	0.040	0.256	0.245
EKLT [21]	E+F	0.325	0.205	0.811	0.775
DDFT [44]	E+F	0.576	0.472	0.825	0.818
FE-TAP [38]	E+F	0.676	0.589	0.844	0.838
EM-ICP [63]	E	0.161	0.120	0.337	0.334
HASTE [3]	E	0.096	0.063	0.442	0.427
DDFT E2VID [44]	E	0.589	0.495	0.794	0.786
ETAP w/o FA-loss (Ours)	E	<u>0.698</u>	0.599	<u>0.885</u>	<u>0.879</u>
ETAP (Ours)	E	0.704	<u>0.598</u>	0.888	0.883

Table 4. *Task 2 - feature tracking on EDS & EC.* Input: E+F (Events & Frames), E (Events only). More results in Suppl.mat.

It also captures more details like the wing flaps, showing the potential of event-based TAP and the improved performance against an event-based E2VID+CoTracker baseline.

5.3. Task 2: Feature Tracking

Results on EDS & EC. We evaluate our method on the EC [45] EDS [26] datasets, following standard feature tracking protocols. These datasets provide synchronized events and frames at resolutions of 240×180 and 640×480 px respectively. Performance is measured using *feature age* (FA) and *expected feature age*, which quantify the duration until a track deviates beyond a threshold distance from the GT. For detailed descriptions of the evaluation protocol and metrics, we refer to [44]. We evaluate our tracker against two categories of methods: those using only events and those using events and frames for enhanced information. Event and frame-based methods comprise ICP [32], “Event-based Kanade-Lucas-Tomasi” (EKLT) [21], which employs template patches extracted from grayscale frames with subsequent event-based tracking, “Data-driven feature tracking for Event Cameras” (DDFT) [44], a recent data-driven approach using similar principles, and “Frame-Event Fusion TAP” (FE-TAP) [38], which implements correlation-based point tracking. Event-only methods comparable to our ETAP comprise EM-ICP [63], HASTE [3], and DDFT E2VID [44], an adaptation of the combined method using E2VID- instead of grayscale-frames. Table 4 summarizes tracking results on the two datasets. Our method outperforms all other event-based methods by a large margin (20% on EDS). Remarkably, it also performs 4.1% better than the best method using frames and events combined.

Name	Resolution	Contrast thresholds	Augment	Dataset	Base fps	Varying dynamics	DNN input	EDS		EC	
								Feature Age \uparrow	Expected FA \uparrow	Feature Age \uparrow	Expected FA \uparrow
Baseline	(256,256)	0.2	–	MOVI-F	12	\times	event stack	0.598	0.500	0.780	0.775
High resolution	(512,512)	0.2	–	MOVI-F	12	\times	event stack	0.659	0.561	<u>0.808</u>	<u>0.802</u>
Random Thresholds as in [30]	(256,256)	$\sim \mathcal{U}(0.16, 0.34)$	–	MOVI-F	12	\times	event stack	0.627	0.531	0.836	0.830
Random Thresholds as in [56]	(256,256)	$\sim \mathcal{U}(0.20, 1.50)$	–	MOVI-F	12	\times	event stack	0.609	0.519	0.801	0.795
Frame Rate Influence	(256,256)	0.2	–	MOVI-F + EventKubric _{static}	12 - 48	\times	event stack	0.618	0.514	0.781	0.777
Varying Dynamics Influence	(256,256)	0.2	–	MOVI-F + EventKubric _{dynamic}	12 - 48	\checkmark	event stack	0.617	0.528	0.781	0.776
Noise Augmentation	(256,256)	0.2	Gauss. noise	MOVI-F	12	\times	event stack	0.631	0.530	0.822	0.816
Representation Influence	(256,256)	0.2	–	MOVI-F	12	\times	voxel grid	0.592	0.505	0.805	0.799
MultiFlow	(512,384)	–	–	MultiFlow [22]	1000	N/A	event stack	0.221	0.178	0.323	0.316
EventKubric (Ours)	(256,256)	$\sim \mathcal{U}(0.16, 0.34)$	–	EventKubric (Ours)	48	\checkmark	event stack	<u>0.646</u>	<u>0.550</u>	0.777	0.772

Table 5. *Sensitivity study.* Analysis of different parameter configurations and their impact on feature tracking performance. Higher values are better (\uparrow). Note that most models were trained on a smaller resolution than the final model (256×256 px).

Figure 8 shows that our method tracks points precisely and even further recovers tracks well when they leave the frame and reenter, which is often not captured by the ground truth and not reflected in the metrics.

5.4. Sensitivity and Ablations

Contrastive loss. Our final model was trained for 10^5 steps without the contrastive loss and further refined for $1.2 \cdot 10^5$ steps including it. We continue to train a comparison model without contrastive loss from the same checkpoint and provide results on all datasets (Tabs. 2 to 4) marked as ETAP w/o FA-loss (Ours). The loss gives a slight boost across all real event datasets (except the synthetic EventKubric) and helps the network to learn motion-robust features.

Dataset and Training Sensitivity. We provide the results of sensitivity studies conducted during dataset construction. In preparation of our self-rendered dataset, we created baseline datasets using MOVI-F, a freely available pre-rendered Kubric dataset. We applied the same event and point track generation procedures as in Sec. 4 with different versions for different contrast sensitivities. We run Vid2e on 512×512 px resolution but downsample at training time for most experiments. The datasets served as a development benchmark to validate design choices. All models were trained for $1.7 \cdot 10^5$ training steps on four Nvidia A6000 GPUs with an effective batch size of 8 (except high-res. on A100). Tab. 5 shows the conducted experiments. We used metrics on EDS for design decisions as it is more consistent than EC, where metrics often alternate between epochs.

The results show the biggest improvement for higher resolution and for choosing random thresholds for event generation of $\sim \mathcal{U}(0.16, 0.34)$ as reported in [30]. The column *base fps* indicates the rendered framerate (before FILM up-sampling). The influence of frame rate and varying dynamics were measured before rendering a whole dataset and therefore tested with only 3500 samples, respectively, and paired with samples of MOVI-F to match the number of baseline samples for comparability. We found both measures (increasing the base frame rate from 12 to 48 and us-

ing panning motion) effective, increasing the performance by $\approx 2\%$ each. The results confirm the effectiveness of the Gaussian noise augmentation and show a slight performance advantage for event stacks over voxel grids. Lastly, we trained our method on MultiFlow. It does not provide meta information to derive visibility flags, which we set to 1 for all tracks. In our tests, we only were able to achieve inferior results compared to training on EventKubric.

5.5. Limitations

Since high-resolution event cameras only provide monochrome information, they cannot yet leverage color information to establish appearance correspondences between points. Furthermore, we observe that our method relies on query times during scene motion. Track features Q_t^i initialized in the absence of motion, and therefore events, do not capture the scene appearance well. This is an inherent problem of event data and could be addressed by reinitializing track features as soon as motion is detected.

6. Conclusion

We introduce the first event-only method for tracking any point in a data stream. The method shows strong performance on five datasets, across different camera types and resolutions, and outperforms all compared methods on a common feature tracking benchmark by a large margin. Its capability is driven by the rigorous design of a new synthetic dataset and a contrastive loss providing robustness of correlation features. Results also show scenarios where our event-only method has advantages over frame-based ones.

Acknowledgements

We thank Dr. Fermüller and NeuroPAC for fostering collaborations within the event-based community (NSF OISE 2020624). Funded by the DFG (German Research Foundation) – EXC 2002/1 “Science of Intelligence” – project no. 390523135. We furthermore gratefully acknowledge the support of the following grants: NSF FRR 2220868, NSF IIS-RI 2212433, NSF TRIPODS 1934960, ONR N00014-22-1-2677, NSF NCS-FO 2124355, SNF 225354.

References

- [1] Ignacio Alzugaray. *Event-driven Feature Detection and Tracking for Visual SLAM*. PhD thesis, ETH Zurich, 2022. [2](#)
- [2] Ignacio Alzugaray and Margarita Chli. ACE: An efficient asynchronous corner tracker for event cameras. In *Int. Conf. 3D Vision (3DV)*, pages 653–661, 2018. [3](#)
- [3] Ignacio Alzugaray and Margarita Chli. Haste: multi-hypothesis asynchronous speeded-up tracking of events. In *British Mach. Vis. Conf. (BMVC)*, page 744, 2020. [7](#), [11](#)
- [4] Simon Baker, Ralph Gross, Ishikawa Takahiro, and Iain Matthews. Lucas-kanade 20 years on: A unifying framework: Part 2. *Technical Report CMU-RI-TR-03-01*, 2003. [2](#)
- [5] D Blender Online Community. Blender—a 3d modelling and rendering package. *Blender Foundation*, 2018. [5](#)
- [6] Levi Burner, Anton Mitrokhin, Cornelia Fermüller, and Yiannis Aloimonos. EVIMO2: An event camera dataset for motion segmentation, optical flow, structure from motion, and visual inertial odometry in indoor scenes with monocular or stereo algorithms. *arXiv e-prints*, 2022. [6](#)
- [7] Weirong Chen, Le Chen, Rui Wang, and Marc Pollefeys. LEAP-VO: Long-term effective any point tracking for visual odometry. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 19844–19853, 2024. [1](#)
- [8] Feng Cheng and Gedas Bertasius. TallFormer: Temporal action localization with a long-memory transformer. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 503–521, 2022. [3](#)
- [9] Philippe Chiberre, Etienne Perot, Amos Sironi, and Vincent Lepetit. Detecting stable keypoints from events through image gradient prediction. In *IEEE Conf. Comput. Vis. Pattern Recog. Workshops (CVPRW)*, 2021. [3](#)
- [10] Seokju Cho, Jiahui Huang, Jisu Nam, Honggyu An, Seungyong Kim, and Joon-Young Lee. Local all-pair correspondence for point tracking. *Eur. Conf. Comput. Vis. (ECCV)*, 2024. [3](#)
- [11] Erwin Coumans. Bullet physics simulation. In *ACM SIGGRAPH 2015 Courses*, 2015. [5](#)
- [12] Yongjian Deng, Hao Chen, Hai Liu, and Youfu Li. A voxel graph cnn for object classification with event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1172–1181, 2022. [4](#)
- [13] Carl Doersch, Ankush Gupta, Larisa Markeeva, Adria Recasens, Lucas Smaira, Yusuf Aytar, Joao Carreira, Andrew Zisserman, and Yi Yang. TAP-Vid: A benchmark for tracking any point in a video. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, pages 13610–13626, 2022. [1](#), [3](#), [6](#)
- [14] Carl Doersch, Yi Yang, Mel Vecerik, Dilara Gokay, Ankush Gupta, Yusuf Aytar, Joao Carreira, and Andrew Zisserman. Tapir: Tracking any point with per-frame initialization and temporal refinement. In *Int. Conf. Comput. Vis. (ICCV)*, pages 10061–10072, 2023. [3](#), [6](#)
- [15] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazırbaş, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *Int. Conf. Comput. Vis. (ICCV)*, pages 2758–2766, 2015. [1](#), [3](#)
- [16] Tobias Fischer, Jiangmiao Pang, Thomas E Huang, Linlu Qiu, Haofeng Chen, Trevor Darrell, and Fisher Yu. Qdtrack: Quasi-dense similarity learning for appearance-only multiple object tracking. *arXiv preprint arXiv:2210.06984*, 2022. [3](#)
- [17] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew Davison, Jörg Conradt, Kostas Daniilidis, and Davide Scaramuzza. Event-based vision: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(1):154–180, 2022. [3](#), [9](#)
- [18] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Asynchronous, photometric feature tracking using events and frames. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 766–781, 2018. [3](#)
- [19] Daniel Gehrig, Antonio Loquercio, Konstantinos G. Derpanis, and Davide Scaramuzza. End-to-end learning of representations for asynchronous event-based data. In *Int. Conf. Comput. Vis. (ICCV)*, pages 5632–5642, 2019. [4](#)
- [20] Daniel Gehrig, Mathias Gehrig, Javier Hidalgo-Carrió, and Davide Scaramuzza. Video to Events: Recycling video datasets for event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3583–3592, 2020. [2](#), [6](#)
- [21] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. EKLt: Asynchronous photometric feature tracking using events and frames. *Int. J. Comput. Vis.*, 128: 601–618, 2020. [3](#), [7](#), [11](#)
- [22] Mathias Gehrig, Manasi Muglikar, and Davide Scaramuzza. Dense continuous-time optical flow from event cameras. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(7):4736–4746, 2024. [1](#), [3](#), [4](#), [8](#)
- [23] Klaus Greff, Francois Belletti, Lucas Beyer, Carl Doersch, Yilun Du, Daniel Duckworth, David J Fleet, Dan Gnanapragasam, Florian Golemo, Charles Herrmann, et al. Kubric: A scalable dataset generator. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 3749–3761, 2022. [2](#), [3](#), [5](#)
- [24] Friedhelm Hamann, Ziyun Wang, Ioannis Asmanis, Kenneth Chaney, Guillermo Gallego, and Kostas Daniilidis. Motion-prior contrast maximization for dense continuous-time motion estimation. In *Eur. Conf. Comput. Vis. (ECCV)*, 2024. [3](#), [6](#)
- [25] Adam W Harley, Zhaoyuan Fang, and Katerina Fragkiadaki. Particle video revisited: Tracking through occlusions using point trajectories. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 59–75, 2022. [1](#), [3](#)
- [26] Javier Hidalgo-Carrió, Guillermo Gallego, and Davide Scaramuzza. Event-aided direct sparse odometry. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 5781–5790, 2022. [2](#), [7](#)
- [27] Berthold K.P. Horn and Brian G. Schunck. Determining optical flow. *J. Artificial Intell.*, 17(1):185 – 203, 1981. [3](#)
- [28] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1647–1655, 2017. [3](#)

- [29] Nikita Karaev, Ignacio Rocco, Benjamin Graham, Natalia Neverova, Andrea Vedaldi, and Christian Rupprecht. Co-Tracker: It is better to track together. *Eur. Conf. Comput. Vis. (ECCV)*, 2024. 3, 4, 5, 6, 7
- [30] Simon Klenk, Marvin Motzet, Lukas Koestler, and Daniel Cremers. Deep event visual odometry. In *Int. Conf. 3D Vision (3DV)*, pages 739–749, 2024. 6, 8
- [31] Matej Kristan, Jiří Matas, Aleš Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Cehovin, Alan Lukežič, et al. The ninth visual object tracking vot2021 challenge results. In *Int. Conf. Comput. Vis. (ICCV)*, pages 2711–2738, 2021. 3
- [32] Beat Kueng, Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Low-latency visual odometry using event-based feature tracks. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, pages 16–23, 2016. 3, 7, 11
- [33] Xavier Lagorce, Cédric Meyer, Sio-Hoi Ieng, David Filliat, and Ryad Benosman. Asynchronous event-based multikernel algorithm for high-speed visual features tracking. *IEEE Trans. Neural Netw. Learn. Syst.*, 26(8):1710–1720, 2015. 3
- [34] Xi Li, Weiming Hu, Chunhua Shen, Zhongfei Zhang, Anthony Dick, and Anton Van Den Hengel. A survey of appearance models in visual object tracking. *ACM transactions on Intelligent Systems and Technology (TIST)*, 4(4):1–48, 2013. 2
- [35] Yijin Li, Zhaoyang Huang, Shuo Chen, Xiaoyu Shi, Hongsheng Li, Hujun Bao, Zhaopeng Cui, and Guofeng Zhang. Blinkflow: A dataset to push the limits of event-based optical flow estimation. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, pages 3881–3888, 2023. 3
- [36] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *IEEE J. Solid-State Circuits*, 43(2):566–576, 2008. 3
- [37] Martin Litzenberger, Christoph Posch, D. Bauer, Ahmed Nabil Belbachir, P. Schön, B. Kohn, and H. Garn. Embedded vision system for real-time object tracking using an asynchronous transient vision sensor. In *Digital Signal Processing Workshop*, pages 173–178, 2006. 3
- [38] Jiaxiong Liu, Bo Wang, Zhen Tan, Jinpu Zhang, Hui Shen, and Dewen Hu. Tracking any point with frame-event fusion network at high frame rate. *arXiv preprint arXiv:2409.11953*, 2024. 3, 7, 11
- [39] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [40] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Int. Joint Conf. Artificial Intell. (IJCAI)*, pages 674–679, 1981. 2
- [41] 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 *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 3
- [42] Iain Matthews, Takahiro Ishikawa, and Simon Baker. The template update problem. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2004. 2
- [43] Nikolaus Mayer, Eddy Ilg, Philip Hausser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 4040–4048, 2016. 3
- [44] Nico Messikommer, Carter Fang, Mathias Gehrig, and Davide Scaramuzza. Data-driven feature tracking for event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023. 1, 3, 7, 11, 14
- [45] 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. *Int. J. Robot. Research*, 36(2):142–149, 2017. 7
- [46] Yeongwoo Nam, Mohammad Mostafavi, Kuk-Jin Yoon, and Jonghyun Choi. Stereo depth from events cameras: Concentrate and focus on the future. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 6104–6113, 2022. 4, 9
- [47] Zhenjiang Ni, Sio-Hoi Ieng, Christoph Posch, Stéphane Régnier, and Ryad Benosman. Visual tracking using neuromorphic asynchronous event-based cameras. *Neural Computation*, 27(4):925–953, 2015. 3
- [48] Jiangmiao Pang, Linlu Qiu, Xia Li, Haofeng Chen, Qi Li, Trevor Darrell, and Fisher Yu. Quasi-dense similarity learning for multiple object tracking. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2021. 3
- [49] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 DAVIS challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 3
- [50] Christoph Posch, Daniel Matolin, and Rainer Wohlgenannt. An asynchronous time-based image sensor. In *IEEE Int. Symp. Circuits Syst. (ISCAS)*, pages 2130–2133, 2008. 3
- [51] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. ESIM: an open event camera simulator. In *Conf. on Robotics Learning (CoRL)*, pages 969–982. PMLR, 2018. 5, 6
- [52] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. Events-to-video: Bringing modern computer vision to event cameras. In *IEEE Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2019. 5, 6
- [53] Fitsum Reda, Janne Kontkanen, Eric Tabellion, Deqing Sun, Caroline Pantofaru, and Brian Curless. FILM: Frame interpolation for large motion. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 250–266, 2022. 5
- [54] Peter Sand and Seth Teller. Particle video: Long-range motion estimation using point trajectories. *Int. J. Comput. Vis.*, 80:72–91, 2008. 1, 3
- [55] Shintaro Shiba, Friedhelm Hamann, Yoshimitsu Aoki, and Guillermo Gallego. Event-based background oriented schlieren. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(4): 2011–2026, 2024. 3
- [56] Timo Stoffregen, Cedric Scheerlinck, Davide Scaramuzza, Tom Drummond, Nick Barnes, Lindsay Kleeman, and Robert Mahony. Reducing the sim-to-real gap for event cameras. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 534–549, 2020. 8
- [57] David Tedaldi, Guillermo Gallego, Elias Mueggler, and Davide Scaramuzza. Feature detection and tracking with the dynamic and active-pixel vision sensor (DAVIS). In *Int. Conf.*

- Event-Based Control, Comm. Signal Proc. (EBCCSP)*, 2016. [3](#)
- [58] Zachary Teed and Jia Deng. RAFT: Recurrent all pairs field transforms for optical flow. In *Eur. Conf. Comput. Vis. (ECCV)*, pages 402–419, 2020. [1](#), [3](#)
 - [59] Mel Vecerik, Carl Doersch, Yi Yang, Todor Davchev, Yusuf Aytar, Guangyao Zhou, Raia Hadsell, Lourdes Agapito, and Jon Scholz. Robotap: Tracking arbitrary points for few-shot visual imitation. In *IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 5397–5403, 2024. [1](#)
 - [60] Ziyun Wang, Friedhelm Hamann, Kenneth Chaney, Wen Jiang, Guillermo Gallego, and Kostas Daniilidis. Event-based continuous color video decompression from single frames. *arXiv preprint arXiv:2312.00113*, 2023. [2](#), [6](#)
 - [61] Qiangqiang Wu, Tianyu Yang, Wei Wu, and Antoni B Chan. Scalable video object segmentation with simplified frame-work. In *Int. Conf. Comput. Vis. (ICCV)*, pages 13879–13889, 2023. [3](#)
 - [62] Yang Zheng, Adam W Harley, Bokui Shen, Gordon Wetstein, and Leonidas J Guibas. Pointodyssey: A large-scale synthetic dataset for long-term point tracking. In *Int. Conf. Comput. Vis. (ICCV)*, pages 19855–19865, 2023. [3](#)
 - [63] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based feature tracking with probabilistic data association. In *IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 4465–4470, 2017. [3](#), [7](#), [11](#)
 - [64] Nikola Zubic, Daniel Gehrig, Mathias Gehrig, and Davide Scaramuzza. From Chaos Comes Order: Ordering Event Representations for Object Recognition and Detection . In *Int. Conf. Comput. Vis. (ICCV)*, pages 12800–12810, 2023. [4](#)