

Comparing Representations in Tracking for Event Camera-based SLAM

Jianhao Jiao[†] Huaiyang Huang[†] Liang Li[‡] Zhijian He[†] Yilong Zhu[†] Ming Liu[†]

[†] The Hong Kong University of Science and Technology*

[‡] The University of Hong Kong

Abstract

This paper investigates two typical image-type representations for event camera-based tracking: time surface (TS) and event map (EM). Based on the original TS-based tracker, we make use of these two representations' complementary strengths to develop an enhanced version. The proposed tracker consists of a general strategy to evaluate the optimization problem's degeneracy online and then switch proper representations. Both TS and EM are motion- and scene-dependent, and thus it is important to figure out their limitations in tracking. We develop six tracker variations and conduct a thorough comparison of them on sequences covering various scenarios and motion complexities. We release our implementations and detailed results to benefit the research community on event cameras: https://github.com/gogojjh/ESVO_extension.

1. Introduction

Event cameras are novel and bio-inspired sensors. Different from conventional frame cameras that capture images at a fixed rate, they asynchronously capture the per-pixel *intensity changes* and output a stream of *events*. Each event is encoded with information, including the triggered time, pixel localization, and the sign of the intensity change. As summarized in [5], event cameras have high temporal resolution (μs -level), high dynamic range (140dB v.s. 60dB of standard cameras), and low power consumption. These characteristics enable event cameras to have great potential for several computer vision and robotic tasks e.g., high-speed motion estimation [2, 6] as well as feature tracking [9, 32] and high dynamic range perception [17, 24], which are difficult to frame cameras. However, research on standard vision problems with event cameras is still preliminary.

This is because event cameras work in a fundamentally different way from frame cameras, which measure *intensity changes* asynchronously rather than the *absolute intensity* (i.e., grayscale data) at a constant rate. Thus, novel algorithms to process events must be investigated.

In this paper, we restrict the scope of literature to the SLAM problem using only event cameras, which does not include methods that use frame cameras [27] or structured light [14] to provide additional color/depth information. Such a SLAM problem has been addressed step by step in scenarios with an increasing complexity. Depending on the complexity of this problem, recent works can be categorized along with three axes: 1) problem dimensionality: from individually handling the localization and 3D reconstruction subproblem [7] to solving the complete tracking-and-mapping problem [22]; 2) type of motion: estimating from constrained motions such as pure rotation or planar motion [8, 12, 18] to arbitrary 6-DoF motions [30]; 3) type of scenes: from artificial patterns [16] to natural scenes [23].

Two event camera-based methods, *EVO* [22] and *ESVO* [30], stand out as solving the SLAM problem in the most general setting (6-DoF motion and natural 3D scenes). Different from the earlier work [11] that needs to recover absolute intensity on a dedicated hardware (GPU), both *EVO* and *ESVO* directly run on a CPU in real-time.

EVO was proposed to solve the monocular event-based state estimation, which consists of a plane sweep-based mapper and a tracker that warps a semi-dense 3D map onto a binary event map (EM). The tracker and mapper work in a parallel manner. *ESVO* also follows this tracking-and-mapping philosophy to design the stereo event-based Visual Odometry (VO) system. It exploits the novel time surface (TS), which encodes spatio-temporal constraints, to estimate a semi-dense depth map for each stereo observation pair. It utilizes the “negative” TS, which is aligned by the 3D map, to optimize pose parameters in tracking. In summary, trackers in both *EVO* and *ESVO* resemble the similar frame-based paradigm [4]: they use non-linear optimization to solve the 3D-2D alignment problem on image-type representations. Although the EM and TS have been separately applied in above event-based VO systems, the current liter-

*This work was supported by Collaborative Research Fund by Research Grants Council Hong Kong, under Project No. C4063-18G, Department of Science and Technology of Guangdong Province Fund, under Project No. GDST20EG54 and Zhongshan Municipal Science and Technology Bureau Fund, under project ZSST21EG06, awarded to Prof. Ming Liu. Email: jjiao@connect.ust.hk

ature has not offered comparative results of them.

The motivation of this paper is to address this deficiency by performing a comprehensive evaluation of the two representations. Furthermore, we enhance the original TS-based tracker by integrating it with the EM counterpart. This tracker has a degeneracy evaluation module to determine which representations are used. All implementations will be publicly released, which are developed from the open-source ESVO¹. We conduct experiments on simulated and real-world sequences, covering various scenarios from artificial patterns to natural scenes under motion in different complexities. Besides the tracking accuracy, we also analyze the limitations of the EM and TS by presenting several failure cases. We hope that our results and conclusions may reference researchers on event-based SLAM and indicate possible directions to improve the current state-of-the-art (SOTA) methods.

2. Overview of Camera Tracking

2.1. Formulation

The tracking module estimates 6-DoF poses of the event camera by taking the image-type event representations (also called event frames) as input. It assumes that a semi-dense 3D map of the environment is given. We denote the reference camera frame as $()^r$ and the current camera frame as $()^c$. We use M to indicate the template image in $()^r$ which is projected by the 3D map \mathcal{M} and \bar{I} to denote the current image. As explained in the following sections, \bar{I} is defined as the “negative” EM or TS. The goal of the tracking is to find the transformation $T(\theta)$ that optimally aligns \mathcal{M} onto \bar{I} . This “alignemnt” is done through the warping function:

$$W(\mathbf{x}, d; \theta) \doteq \pi(T(\theta) \cdot \pi^{-1}(\mathbf{x}, d)), \quad \mathbf{x} \in M, \quad (1)$$


where θ are the pose parameters; $T(\theta)$ turns θ into a transformation matrix from $()^r$ to $()^c$; $\mathbf{x} = [u, v]$ is the localization of a valid pixel on the template image; $\pi(\cdot)$ projects a 3D point onto the image plane, while $\pi^{-1}(\cdot)$ back-projects a pixel into 3D space with the known depth d . We minimize the objective function to find the optimal θ^* as

$$\arg \min_{\theta} \sum_{\mathbf{x} \in M} \rho(\bar{I}(W(\mathbf{x}, d; \theta))^2), \quad (2)$$

where $\bar{I}(\cdot)$ is the pixel value and $\rho(\cdot)$ is the robust loss. Following ESVO, we reformulate Problem (2) using the forward compositional Lucas-Kanade method, which iteratively refines the incremental pose parameters $\Delta\theta$ by minimizing the objective as

$$\arg \min_{\Delta\theta} \sum_{\mathbf{x} \in M} \rho(\bar{I}(W(W(\mathbf{x}, d; \Delta\theta); \theta))^2), \quad (3)$$

¹<https://github.com/HKUST-Aerial-Robotics/ESVO>



(a) An example of the time surface which is triggered at every 10ms (b) An example of the event map which aggregates 4000 events

Figure 1. An example of the TS and EM.

where the warping function is updated at each iteration

$$W(\mathbf{x}, d; \theta) \leftarrow W(W(\mathbf{x}, d; \Delta\theta); \theta). \quad (4)$$

The compositional approach is more efficient than the original Lucas-Kanade algorithm. This is because the Jacobian $\frac{\partial W}{\partial \theta}$ of the warping function is evaluated at a fixed point: $(\mathbf{x}, 0)$ which can be pre-computed. Another reformulation is called the inverse compositional Lucas-Kanade algorithm, which is used in EVO. It switches the roles of the template and current image. In other words, we need to evaluate the gradient ∇M of M as the part of the Jacobian, not of \bar{I} in the forward approach. But this paper considers the forward approach since it exploits the implicit “slope” feature of the TS. Moreover, the definition of θ varies according to different parameterization ways, including the Lie algebra [1] or Euler angles/Quaternions [26]/Cayley parameters [3] for rotation with translation vectors.

2.2. Two Event Representations

2.2.1 Event Map

Besides EVO, the binary EM has also been applied in visual-inertial odometry [21] where FAST corners [25] are directly extracted and tracked. Constructing the EM is simple: a group of events within a temporal neighborhood is aggregated onto the image. The pixel is set to 255 where an event is fired; otherwise, it is set to zero. Therefore, EMs are output in an asynchronous manner. But due to the data-driven nature of event cameras, the number of events in the group should be tuned for a specific dataset. Otherwise if the number is small, the frequency of producing EMs increases when the camera moves rapidly or scenes are high-contrast. This will raise the computational burden to the tracker.

2.2.2 Time Surface

The TS is also an image where each pixel stores a value. Using an exponential decay kernel, TSs can emphasize recent events over past events. Given an arbitrary timestamp

t , the value at a pixel \mathbf{x} is defined as

$$I(\mathbf{x}, t) \doteq \exp\left(-\frac{t - t_{last}(\mathbf{x})}{\delta}\right), \quad (5)$$

where $t_{last}(\mathbf{x}) \leq t$ is the timestamp of the last event at \mathbf{x} and δ is the constant decay rate parameter (i.e., $30ms$). This equation converts events into an image whose intensity is a function of the motion history at that location. Larger values correspond to a more recent motion. Online adjusting the value of t can signal the creation of TSs synchronously or asynchronously. In experiments, we output TSs in a synchronous way. For convenient visualization and processing, each TS is rescaled from $[0, 1]$ to the range $[0, 255]$.

2.2.3 Summary

We summarize the pros and cons of these representations from recent literature [5]. Generally, EMs implicitly represent the edge map since events are mostly triggered by edge patterns. They are the universal data structure compatible with conventional computer vision. Also, generating an EM is very fast (i.e., $< 0.5ms$). Nevertheless, EMs are highly sensitive to motion blur if the number of aggregating events is not set well. In contrast, TSs are more informative than EMs even though they require around $5 - 10ms$ for the synthesis. In experiments, we will further identify several limitations of both EMs and TSs in challenging scenarios.

The tracker operates the “negative” event frames i.e., $\bar{I}(\mathbf{x}) = 255 - I(\mathbf{x})$. This is based on a crucial observation: the objective (3) becomes minimum if the template image is “perfectly” aligned on dark areas in the current frame. An example can be seen in Fig. 4(c). Specifically, the negative TSs can be interpreted as an anisotropic distance field as presented in edge-based VO [31]. To enlarge the convergence basin’s width in optimization, a Gaussian blur (kernel size = 5 pixels) is applied to negative event frames.

2.3. Enhanced Tracker with Degeneracy Check

The synchronous pose output of the TS-based tracker is beneficial to sensor fusion like the visual-inertial odometry or other robotic applications (e.g., drone racing). However, unlike the EM that aggregates a fixed number of events, the TS becomes unreliable if few events are triggered (see Section 3). This issue frequently occurs where event cameras are working in textureless scenarios or relatively static to scenes. The “jump” in motion estimates is inevitable due to unconstrained degrees of freedom.

In this paper, we enhance the original TS-based tracker in robustness by taking EMs as the backup representation. The proposed tracker is based on a test that online evaluates the degeneracy of the optimization problem (3). Inspired by [28], we use the *factor* λ as the degeneracy quantification metric. λ is defined as the minimum eigenvalue of the

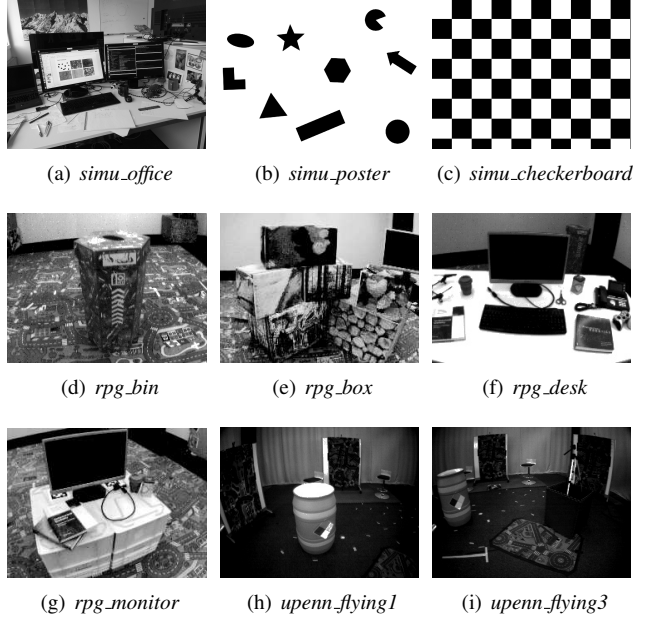


Figure 2. Scene images of *simu*, *rpg*, and *upenn* sequences. The camera resolutions are 346×260 , 240×180 , and 346×260 , respectively.

Hessian matrix which is computed by linearizing (3) as a normal equation: $\mathbf{H}\Delta\theta = \mathbf{g}$. We take both TSs and ESs as input, but first utilize the TS-based tracker. Before optimization, we compute λ and check if it is smaller than a threshold λ_{th} . If $\lambda < \lambda_{th}$, problem (3) may degenerate, and we use the EM-based tracker to optimize the pose parameters. The EM is constructed by aggregates the recent 4000 events. In experiments, we empirically set $\lambda_{th} = 100$.

3. Experiment

3.1. Implementation Details

We use the Eigen library [10] to solve the nonlinear optimization problem. The tracker requires a known 3D map, which is computed by the mapping module of ESVO in real-time. We use two public datasets, and the simulator [20] to generate sequences for evaluation. All data were collected with stereo event cameras, but the motion and scenes present various difficulties. The simulator provides synthetic data with a planar structure and an “ideal”, noise-free camera model. We collect sequences (*simu_X*) using stereo event cameras before a wall with various backgrounds. We category the collected simulated sequences into two classes according to the motion complexity: 1) slow ($\approx 0.3m/s$) and planar motion; 2) fast ($\approx 1.0m/s$) and 6-DoF motion. Real-world data in [30] (*rpg_X*) were collected in an office, while that in [33] (*upenn_X*) were collected by mounting cameras on a flying drone in a capacious indoor area. Note that the “X” indicates the type of scenes and motions. The




Figure 3. Failure cases of the EM-based tracker which aggregates 3000 events. Sequences from left to right are: *simu_office_6DoF*, *rpg_bin*, and *rpg_box*.

Sequence	TS	EM ₂₀₀₀	EM ₃₀₀₀	EM ₄₀₀₀	EM ₅₀₀₀	TSEM ₄₀₀₀
<i>simu_office_planar</i>	4.7	4.0	3.9	3.7	4.1	4.9
<i>simu_poster_planar</i>	4.7	3.7	4.3	4.6	5.0	5.2
<i>simu_checkerboard_planar</i>	4.2	2.9	2.2	2.3	2.4	3.5
<i>simu_office_6DoF</i>	9.1	25.3	21.0	16.6	15.8	9.4
<i>simu_poster_6DoF</i>	18.2	15.4	16.3	16.8	17.4	17.8
<i>simu_checkerboard_6DoF</i>	23.0	17.0	14.0	15.1	13.4	22.4
<i>rpg_bin_6DoF</i>	3.4	22.4	16.6	8.0	14.1	3.5
<i>rpg_box_6DoF</i>	6.5	5.3	17.1	13.7	9.8	22.1
<i>rpg_desk_6DoF</i>	3.4	2.9	3.3	3.2	2.9	3.6
<i>rpg_monitor_6DoF</i>	7.2	5.3	5.2	7.4	7.3	7.3
<i>upenn_indoor_flying1_6DoF</i>	18.5	22.0	16.7	16.0	22.1	14.9
<i>upenn_indoor_flying3_6DoF</i>	20.9	10.8	11.9	14.0	15.0	10.6

Table 1. Mean Absolute Trajectory Error in translation [cm] under 10-trials. The first two lowest errors are marked as bold.

ground-truth camera poses are provided. Fig. 2 shows the scene images. In experiments, the algorithm is executed on a desktop with an i7 CPU@4.20 GHz.

3.2. Tracking Evaluation

We implement six tracker variations according to the used representations. They are denoted by “TS”, “EM₂₀₀₀”, “EM₃₀₀₀”, “EM₄₀₀₀”, “EM₅₀₀₀”, and “TSEM₄₀₀₀” respectively. The subscript of “EM” indicates the number of events to generate an EM. “TSEM₄₀₀₀” is the enhanced tracker with the degeneracy check. To evaluate their performance, we compute the pose accuracy using the absolute trajectory error (ATE) in translation [29] w.r.t. the ground truth. We conduct 10-trial SLAM tests on each sequence. The mean ATE is reported in Table 1.


All trackers achieve similar accuracy on the simplest sequences (*simu_X_planar*), while they show worse accuracy on more complex sequences. We cannot directly determine which representations are better for tracking according to the accuracy in Table 1 since all trackers may have large er-

TS	EM ₂₀₀₀	EM ₃₀₀₀	EM ₄₀₀₀	EM ₅₀₀₀	TSEM ₄₀₀₀
8 ± 5	8 ± 6	7 ± 4	7 ± 4	7 ± 4	7 ± 3


Table 2. Average computation time [ms] of trackers.

rors on specific sequences. We should analyze the results by considering the characteristics of scenes and motion. The TS-based tracker outperform others on *simu_office_6DoF*, *rpg_bin*, and *rpg_box*, where high-contrast textures such as lines and rectangles are presented (see Fig. 2). In contrast, we observe that EM-based trackers reach another local minima which is far from the groundtruth during the optimization process and thus result in low accuracy. Fig. 3 illustrates three failure cases.

On the other hand, EM-based trackers consistently have better performance than their TS-based counterpart on *upenn_indoor_flying3*. We note that few events are generated when the drone stops and hovers at some time. This causes the generation of unreliable TSs which are triggered



(a) The value of the degeneracy factor λ on different sequences



(b) The 3D map is warped on the negative Time surface (TS). (c) The 3D map is warped on the negative Event map (EM₄₀₀₀).

Figure 4. We plot the values of λ on three sequences in (a). At the time of 8.2s of *upenn_indoor_flying1* (b) where the TS-based tracker degenerates, (c) the 3D map is well aligned on “dark areas” of the negative EM with sufficient constraints.

at a constant rate. For the tracker, not enough constraints (i.e., “dark areas” on TSs) make the optimization degenerate or fail. An example is shown in Fig. 4, with the plot displaying the value of the *degeneracy factor* λ on sequences: *simu_office_6DoF*, *rpg_box*, and *upenn_indoor_flying3*. Regarding *upenn_indoor_flying3*, at the time of 8.2s, the TS-based tracker degenerates due to the unreliable TS (see Fig. 4(b)). The corresponding *factor* λ is smaller than the threshold. Therefore, we can switch the stable EM and then redo the tracking, improving the tracker’s results.

The trajectories produced by the TS-based tracker on all sequences are compared in Fig. 5, Fig. 6, and Fig. 7. Due to limited space, only the ground truth trajectories and trajectories of the tracker with the highest accuracy are shown as the reference. Through these figures, readers can have a better understanding of the above explanations. For example, we observe that the poses of the TS-based tracker “jump” on *upenn_indoor_flying3* from Fig. 7. The unreliable TS causes degrades this performance.

3.3. Computation Time

The average computational time of tracker variations is reported in Table 2, where all of them take around 7-8ms to solve the pose estimation problem. Depending on the sensor resolution, the creation of the TS takes about 5-10ms [30]. The degeneracy evaluation takes about 0.5-1.5ms.

3.4. Discussion

All the tested trackers have their own limitations. First, although the enhanced TSEM-based tracker can mitigate the undesirable effect caused by the unreliable TS, it suffers shortcomings of both the TS and EM if λ_{th} is not well tuned. This explains why the TSEM has bad accuracy on *simu_poster_6DoF*, *simu_checkerboard_6DoF*, and *rpg_box*. Hence, the utilization of the TSEM is subtle: tuning a good parameter requires to know more about dataset, while the degeneracy check is used to evaluate the conditions in unfamiliar environments.

Second, we find that the estimated trajectories of the current tracker are not smooth. As an example, the TS-based tracker’s trajectory on *rpg_monitor* is visualized in Fig. 8. This is due to the fact that the tracker is subjected to large noise in optimization. Introducing the motion prior [4], considering historical measurements (sliding window-based framework [19], full batch optimization [16], or regularization term [13] could be possible solution to solve this issue. Moreover, as done in [21, 27, 34], the aid of inertial measurements have significantly extended the working scope of monocular event camera. Exploring stereo event-based VIO is also a promising direction for this problem.

Finally, the image-type representations quantize event timestamps. The tracking results can be further optimized by introducing spline-based optimization framework [9, 15].

4. Conclusion

In this work, we have conducted extensive comparisons of two image-type representations for event camera-based tracking. We also introduce the degeneracy check and propose an enhanced tracker to make use of their complementary strengths. Furthermore, we implement six tracker variations. In evaluating these trackers, our goal is to benchmark their performance in different-complexity scenarios. The results presented in Section 3 suggest that each representation has its own advantages and limitations for tracking. Such event-based representations are scene- and motion-dependent: both of them may degenerate in non-ideal situations. We hope that these results and conclusions shown in this paper may provide a reference for researchers on event-based SLAM and indicate possible directions to improve the state-of-the-art (SOTA) methods.

References

- [1] Timothy D Barfoot. *State estimation for robotics*. Cambridge University Press, 2017. 2
- [2] Samuel Bryner, Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. Event-based, direct camera tracking from a photometric 3d map using nonlinear optimization. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 325–331. IEEE, 2019. 1

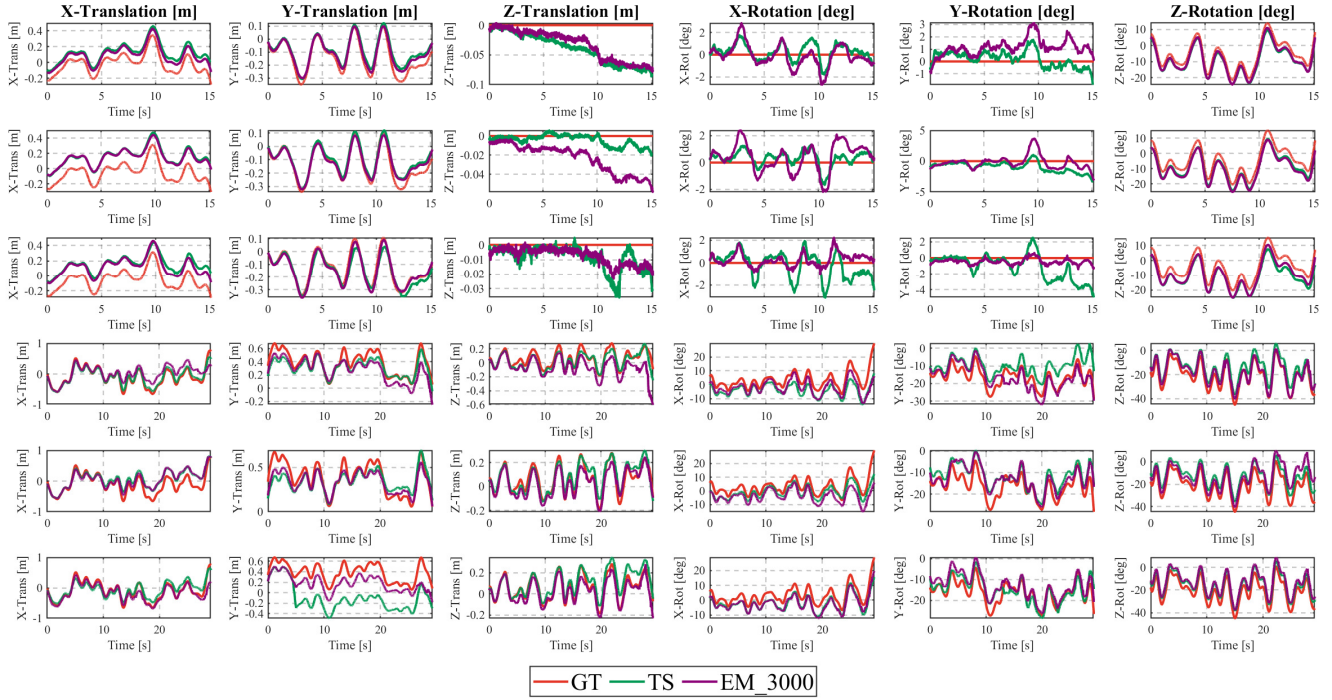


Figure 5. 6-DoF poses comparison of the TS-based (“TS”) and EM-based tracker with the highest accuracy (“EM₃₀₀₀”) against the ground truth (“GT”) on **simulated sequences**. Each row corresponds to a different sequence: *simu_office_planar*, *simu_poster_planar*, *simu_checkerboard_planar*, *simu_office_6DoF*, *simu_poster_6DoF*, *simu_checkerboard_6DoF*. These sequences are captured with stereo event cameras under random motions.

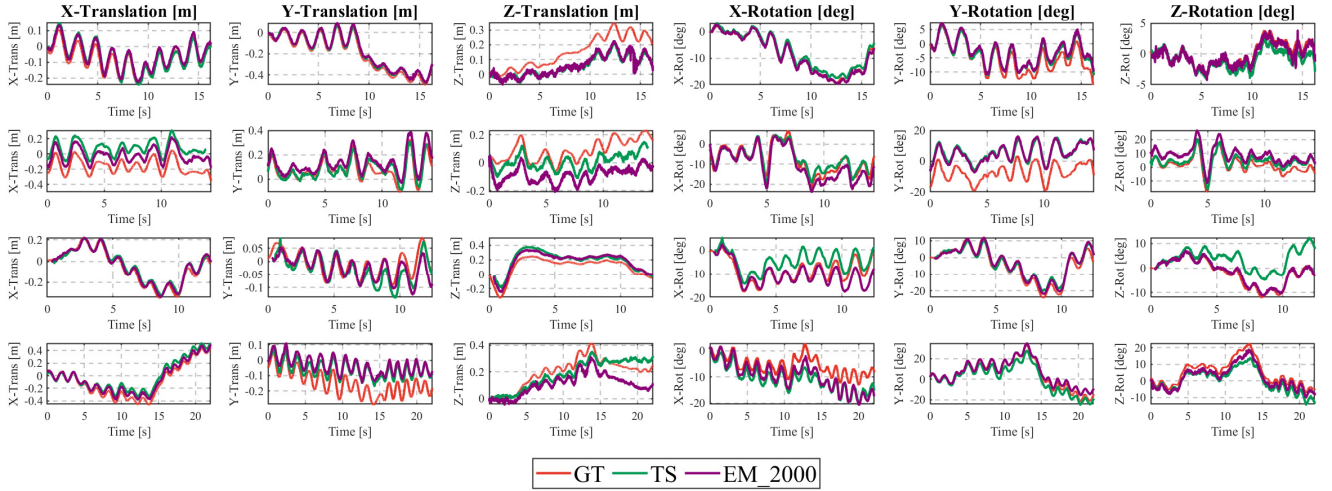


Figure 6. 6-DoF poses comparison of the TS-based (“TS”) and EM-based tracker with the highest accuracy (“EM₂₀₀₀”) against the ground truth (“GT”) on **rpg sequences**. Each row corresponds to a different sequence: *rpg_bin*, *rpg_box*, *rpg_desk*, *rpg_monitor*. These sequences are captured with a handheld stereo rig moving under a locally loopy behavior.

- [3] Arthur Cayley. About the algebraic structure of the orthogonal group and the other classical groups in a field of characteristic zero or a prime characteristic. *Reine Angewandte Mathematik*, 32(1846):6, 1846. 2
- [4] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual odometry for monocular and multicamera systems. *IEEE*

- Transactions on Robotics*, 33(2):249–265, 2016. 1, 5
- [5] G. Gallego, T. Delbruck, G. M. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2020. 1, 3
- [6] Guillermo Gallego, Jon EA Lund, Elias Mueggler, Henri Re-




Figure 7. 6-DoF poses comparison of the TS-based (“TS”) and TSEM-based tracker (“TSEM”) against the ground truth (“GT”) on **upenn sequences**. Each row corresponds to a different sequence: *upenn_indoor_flying1*, *upenn_indoor_flying3*. These sequences are captured with a stereo rig mounted on a drone.




Figure 8. TS’s trajectory on *rpg_monitor*.

becq, Tobi Delbruck, and Davide Scaramuzza. Event-based, 6-dof camera tracking from photometric depth maps. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2402–2412, 2017. 1

- [7] Guillermo Gallego, Jon EA Lund, Elias Mueggler, Henri Rebecq, Tobi Delbruck, and Davide Scaramuzza. Event-based, 6-dof camera tracking from photometric depth maps. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2402–2412, 2017. 1
- [8] Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. A unifying contrast maximization framework for event cameras, with applications to motion, depth, and optical flow estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3867–3876, 2018. 1
- [9] Cedric Le Gentil, Florian Tschopp, Ignacio Alzugaray, Teresa Vidal-Calleja, Roland Siegwart, and Juan Nieto. Idol: A framework for imu-dvs odometry using lines. *arXiv preprint arXiv:2008.05749*, 2020. 1, 5
- [10] Gaël Guennebaud, Benoit Jacob, et al. Eigen. URL: <http://eigen.tuxfamily.org>, 2010. 3
- [11] Hanme Kim, Stefan Leutenegger, and Andrew J Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *European Conference on Computer Vision*, pages 349–364. Springer, 2016. 1
- [12] Daqi Liu, Alvaro Parra, and Tat-Jun Chin. Globally optimal contrast maximisation for event-based motion estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6349–6358, 2020. 1
- [13] Stephan Liwicki, Christopher Zach, Ondrej Miksik, and Philip HS Torr. Coarse-to-fine planar regularization for dense monocular depth estimation. In *European Conference on Computer Vision*, pages 458–474. Springer, 2016. 5
- [14] Julien NP Martel, Jonathan Müller, Jörg Conradt, and Yulia Sandamirskaya. An active approach to solving the stereo matching problem using event-based sensors. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2018. 1
- [15] Elias Mueggler, Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. Continuous-time visual-inertial odometry for event cameras. *IEEE Transactions on Robotics*, 34(6):1425–1440, 2018. 5
- [16] Elias Mueggler, Guillermo Gallego, and Davide Scaramuzza. Continuous-time trajectory estimation for event-based vision sensors. Technical report, 2015. 1, 5
- [17] Liyuan Pan, Cedric Scheerlinck, Xin Yu, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Bringing a blurry frame alive at high frame-rate with an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6820–6829, 2019. 1
- [18] Xin Peng, Yifu Wang, Ling Gao, and Laurent Kneip. Globally-optimal event camera motion estimation. In *European Conference on Computer Vision*, pages 51–67. Springer, 2020. 1
- [19] Tong Qin, Peiliang Li, and Shaojie Shen. Vins-mono: A robust and versatile monocular visual-inertial state estimator. *IEEE Transactions on Robotics*, 34(4):1004–1020, 2018. 5
- [20] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. Esim: an open event camera simulator. In *Conference on Robot Learning*, pages 969–982. PMLR, 2018. 3
- [21] Henri Rebecq, Timo Horstschaefer, and Davide Scaramuzza. Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization. 2017. 2, 5
- [22] Henri Rebecq, Timo Horstschaefer, 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. 1
- [23] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. Events-to-video: Bringing modern computer vision to event cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3857–3866, 2019. 1
- [24] Henri Rebecq, René Ranftl, Vladlen Koltun, and Davide Scaramuzza. High speed and high dynamic range video with an event camera. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 1

- [25] Edward Rosten and Tom Drummond. Machine learning for high-speed corner detection. In European conference on computer vision, pages 430–443. Springer, 2006. 2
- [26] Joan Sola. Quaternion kinematics for the error-state kalman filter. arXiv preprint arXiv:1711.02508, 2017. 2
- [27] Antoni Rosinol Vidal, Henri Rebecq, Timo Horstschaefer, 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. 1, 5
- [28] Ji Zhang, Michael Kaess, and Sanjiv Singh. On degeneracy of optimization-based state estimation problems. In 2016 IEEE International Conference on Robotics and Automation (ICRA), pages 809–816. IEEE, 2016. 3
- [29] Zichao Zhang and Davide Scaramuzza. A tutorial on quantitative trajectory evaluation for visual (-inertial) odometry. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 7244–7251. IEEE, 2018. 4
- [30] Yi Zhou, Guillermo Gallego, and Shaojie Shen. Event-based stereo visual odometry. IEEE Transactions on Robotics, 2021. 1, 3, 5
- [31] Yi Zhou, Hongdong Li, and Laurent Kneip. Canny-vo: Visual odometry with rgb-d cameras based on geometric 3-d-2-d edge alignment. IEEE Transactions on Robotics, 35(1):184–199, 2018. 3
- [32] 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. 1
- [33] 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. 3
- [34] Alex Zihao Zhu, Nikolay Atanasov, and Kostas Daniilidis. Event-based visual inertial odometry. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 5391–5399, 2017. 5