

A Linear N-Point Solver for Structure and Motion from Asynchronous Tracks

Hang Su^{1*} Yunlong Feng^{1*} Daniel Gehrig² Panfeng Jiang¹ Ling Gao³
Xavier Lagorce¹ Laurent Kneip^{1,4}

¹ ShanghaiTech University, ² University of Pennsylvania, ³ Amap, Alibaba Group,
⁴ Shanghai Engineering Research Center of Intelligent Vision and Imaging

Abstract

Structure and continuous motion estimation from point correspondences is a fundamental problem in computer vision that has been powered by well-known algorithms such as the familiar 5-point or 8-point algorithm. However, despite their acclaim, these algorithms are limited to processing point correspondences originating from a pair of views each one representing an instantaneous capture of the scene. Yet, in the case of rolling shutter cameras, or more recently, event cameras, this synchronization breaks down. In this work, we present a unified approach for structure and linear motion estimation from 2D point correspondences with arbitrary timestamps, from an arbitrary set of views. By formulating the problem in terms of first-order dynamics and leveraging a constant velocity motion model, we derive a novel, linear point incidence relation allowing for the efficient recovery of both linear velocity and 3D points with predictable degeneracies and solution multiplicities. Owing to its general formulation, it can handle correspondences from a wide range of sensing modalities such as global shutter, rolling shutter, and event cameras, and can even combine correspondences from different collocated sensors. We validate the effectiveness of our solver on both simulated and real-world data, where we show consistent improvement across all modalities when compared to recent approaches. We believe our work opens the door to efficient structure and motion estimation from asynchronous data. Code can be found at <https://github.com/suhang99/AsyncTrack-Motion-Solver>.

1. Introduction

Finding the continuous motion of a single monocular camera is one of the most fundamental problems in the area of geometric computer vision. In the calibrated case, the core of the solution to the visual odometry problem en-

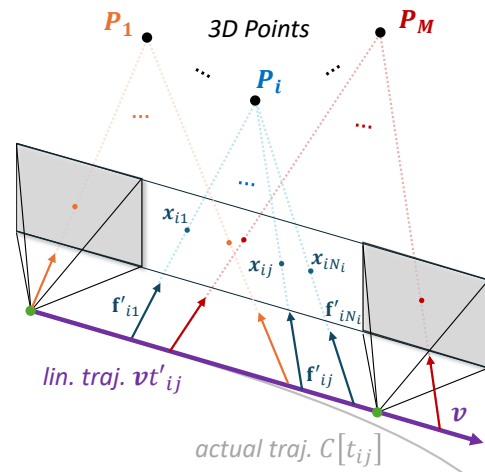


Figure 1. We develop a linear N-point solver for recovering 3D points P_i and the velocity \mathbf{v} of a camera undergoing quasi-linear motion, given a set of timestamped observations $(\mathbf{x}_{ij}, t_{ij})$. No assumptions are made about the temporal synchronization of observations \mathbf{x}_{ij} , yielding a general algorithm that can handle observations from global shutter, rolling-shutter, and fully asynchronous rolling sensors, such as event cameras. Observations \mathbf{x}_{ij} are converted into rotation compensated bearing vectors $\mathbf{f}'_{ij} \doteq \exp([\omega t'_{ij}]_{\times}) \mathbf{K}^{-1} \tilde{\mathbf{x}}_{ij}$, and used to construct a set of linear point incidence relations. Here $t'_{ij} = t_{ij} - t_s$ denotes relative time, t_s the reference time, \mathbf{K} the camera intrinsics, ω the angular rate, which can be given by an IMU or upstream estimation algorithm.

tails the identification of the *extrinsic parameters* relating a pair of views of a scene, meaning a Euclidean transformation consisting of a relative orientation and an up-to-scale translational displacement. Note however that in the video-based, continuous motion setting, there is only a marginal difference between finding the relative pose and finding local camera dynamics. Indeed, if the images come from a constant-rate video stream, we merely have to divide the relative displacement variables by the frame sampling period to obtain approximate local camera dynamics.¹

¹The matter may complicate if more jerky stop-and-go motion is involved, in which case key-frame selection may be required, and the op-

*indicates equal contribution

It might not seem obvious, but the primary reason why—in the classical visual odometry scenario—we stick to relative transformation parameters rather than first-order dynamics is the very nature of the sampling mechanism of traditional cameras: Images are sampled synchronously, and—at least in the common case—at a relatively low rate, which easily tends to temporally under-sample more agile camera motion. For this reason, motion estimation is traditionally framed as the recovery of delta transformations, instead of local first-order dynamics. However, with the introduction of temporally denser and notably asynchronously sampling sensors, the consideration of first-order dynamics and a constant velocity motion model becomes a practical necessity.

Important examples of such sensors are given by high-speed cameras, rolling shutter cameras, and event cameras. Rolling shutter cameras notably capture images one row at a time, in a quasi-asynchronous way, leading to timestamp differences at different rows of the image, and these differences need to be meaningfully addressed in a given motion estimation algorithm. The event camera on the other hand is a relatively new type of visual sensor [25] and has been a recent enabler of high-speed and low-power computer vision due to its unique working principle. It consists of independent pixels that return an asynchronous stream of per-pixel brightness change measurements, *i.e.* *events*. Each event indicates a discrete, signed change in perceived brightness timestamped by a clock with microsecond-level resolution.

In this work, we focus on motion estimation using geometric feature-based methods. While a plethora of such methods exist for global shutter and rolling shutter cameras, none manage to natively support features extracted at asynchronous timestamps from potentially row-wise or completely asynchronous sensors. We fill this gap by building upon a recently introduced geometric method for line feature-based motion and structure estimation from asynchronous measurements [10, 11], and extend it to operate on 3D points instead. Originally proposed for event cameras, the solver utilizes a small, sparse set of events generated anywhere along the reprojection of a straight line perceived under constant linear motion. Surprisingly, given sufficient measurements of a single line, it is possible to recover the full 3D location of the straight line as well as a projection of the 3D camera velocity.

The present work makes the following contributions:

- We propose a new solver for geometric motion and structure estimation from asynchronous, time-stamped reprojections of points measured under the assumption of constant linear motion with known rotation. We thus extend previous line-feature-based solvers by a novel, point-

timization problem typically transitions into a multi-view, absolute pose bundle adjustment scheme. However, the statement is approximately valid under the assumption of a simple, continuous trajectory.

feature based approach, and contribute to a growing theoretical body of geometric incidence relations that operate over dynamics and measurements sampled arbitrarily in space and time.

- The proposed method is a single-stage linear solver that operates over an arbitrary number of asynchronous point feature tracks and is highly efficient by employing the Schur complement trick. We furthermore outline the exact conditions under which the full linear velocity becomes observable. Surprisingly, under a linear motion model, three temporal observations of only a single point enable us to recover the full orientation of the displacement baseline as well as the corresponding 3D point.
- Through our experiments, we demonstrate general applicability of the theory to various types of cameras, including regular global shutter cameras, rolling shutter cameras, and event cameras.

2. Related Work

Geometric Solvers: The geometric solution of camera motion is undoubtedly one of the major success stories of computer vision. Based on epipolar geometry [15], a plethora of algorithms has been proposed to efficiently estimate scene properties and camera parameters from a set of 2D point correspondences (*i.e.* pixel observations) between two views [34, 43].

Relaxing the synchronicity assumption is already required when addressing correspondences captured from rolling shutter images. Important theory has been devised to estimate absolute pose and motion from known 2D-3D correspondences using either linear or polynomial solvers [1, 21, 22, 40]. In particular, Saurer *et al.* [40] use a similar point incidence relation as the one presented in this work. Unlike these works, our method uses only 2D point measurements and simultaneously solves motion and structure via a linear system. The work in [24] targets homography estimation from point correspondences derived from rolling shutter images, but requires the observed 3D points to be on a plane. Dai *et al.* [5] also relax the synchronicity assumption, however using a pair-wise epipolar co-planarity constraint instead of the N-point incidence relation proposed in this work. Further related approaches are given by n-linearities [15], visual-inertial bootstrapping approaches [20], or recent works on visual odometry on cars [16]. However—though able to process N-point feature tracks—these approaches are proposed for regularly sampling global shutter cameras, or simply multiple individual cameras. To the best of our knowledge, we propose the first theory that relies on a constant velocity motion model and permits for fully asynchronous measurements.

Event-based Motion Estimation: Various methods for

feature extraction and tracking—both data-driven [14, 30, 31], and traditional [12, 13, 23]—as well as visual odometry [18, 19, 27, 33, 36, 39, 45, 48] have already been proposed. A common strategy in many of these works, particularly in learning-based pipelines, is to aggregate events into synchronous, frame-like representations. This approach, however, sidesteps the advantage of event cameras. In parallel, some earlier research explicitly exploit the asynchronous nature of the data, developing methods for motion estimation that operate directly on the event stream. Despite these different approaches, the critical question of closed-form bootstrapping often remains unaddressed. An interesting alternative for event-based motion and structure estimation that processes raw events with their original timestamp is given by contrast maximization [8, 38, 44]. By employing a compact parametric image warping function, events are unwrapped into a reference view in which—owing to the sparsity of strong appearance gradients—the entropy of their distribution is minimized. Although the framework has been successfully applied to various problems such as motion estimation, optical flow, and depth estimation [14, 17, 26, 42], it is a computationally intensive approach that involves iterative batch optimization over many events and remains restricted to homographic warping scenarios (*e.g.* pure rotation, planar scene). Several works [4, 28, 35] also combine event cameras with IMU sensors to address high-speed maneuvers and challenging conditions, leveraging the complementary properties of both sensing modalities.

Recent approaches have explored efficient, sparse geometric solvers better suited to the asynchronous nature of event data. Peng *et al.* [37] and Xu *et al.* [46] utilize three-view geometry based on 3D lines for camera ego-motion estimation. Gao *et al.* [10, 11] improve this idea by developing an N-point linear solver for line and motion estimation, providing new insights into the manifold distribution of the events generated by the observation of a straight line under motion. Built upon this, Zhao *et al.* [47] propose a new solver for full-DoF motion estimation via rank minimization. Unlike these approaches, our work focuses on asynchronous point feature tracks and exploits their spatio-temporal characteristics for accurate motion estimation.

3. Methodology

We model the observations of M 3D points $\{\mathbf{P}_i\}_{i=1}^M$ by a calibrated camera undergoing an arbitrary 6 degrees of freedom (DoF) motion on the time interval $\mathcal{T} = [t_s - \Delta t, t_s + \Delta t]$, with reference time t_s , and half-width Δt . We denote the camera pose at time $t \in \mathcal{T}$ with $\mathbf{C}(t) \in SE(3)$. Throughout this duration we assume each point $\mathbf{P}_i \in \mathbb{R}^3$ to be observed N_i times by a point tracker, leading to spatio-temporal track observations $\mathcal{X}_i = \{(\mathbf{x}_{ij}, t_{ij})\}_{j=1}^{N_i}$ in the image plane. Each observation $(\mathbf{x}_{ij}, t_{ij})$ comprises the projec-

tion \mathbf{x}_{ij} of point \mathbf{P}_i at timestamp t_{ij} in the image plane.

Note that no assumption is made on the synchronicity of timestamps t_{ij} , leading to a very general formulation. This formulation can handle a wide range of tracking scenarios and sensing modalities:

1. Tracks derived from a sequence of global shutter images, which may or may not be temporally aligned due to loss of tracking or track initialization.
2. Row-wise synchronized tracks derived from a rolling shutter camera, where points at different row coordinates do not necessarily share the same timestamp t_{ij} .
3. Fully asynchronous feature tracks from an event camera. These tracks may be densely sampled in time and have little to no timestamp coherence.

We will show that our solver seamlessly handles all of these cases. In Sec. 3.1, we will present the relevant incidence relation, which forms a set of linear constraints between the point location, camera motion, and point observation. Then, in Sec. 3.2 we will present how to write a set of such constraints as a linear system, before solving it in Sec. 3.3. We discuss properties of the solver in Sec. 3.4. Sec. 3.5 concludes with implementation details.

3.1. Incidence Relationship

We will assume a configuration as illustrated in Fig. 1, where a camera with pose $\mathbf{C}(t)$ composed of orientation $\mathbf{R}(t)$ and position $\mathbf{p}(t)$ undergoes quasi-linear dynamics on the small interval $\mathcal{T} = [t_s - \Delta t, t_s + \Delta t]$.

Let $\tilde{\mathbf{x}}_{ij}$ be the 3D homogeneous coordinate of \mathbf{x} and $\mathbf{f}_{ij} = \mathbf{K}^{-1}\tilde{\mathbf{x}}_{ij}$ be the normalized coordinates (*i.e.* bearing vectors) of the feature tracks previously described. Our incidence relation leverages the fact that the 3D point \mathbf{P}_i observed at time t_{ij} should project onto the observed point \mathbf{f}_{ij} . We find the 3D point in camera coordinates \mathbf{P}'_{ij} at time t_{ij} as

$$\mathbf{P}'_{ij} = \mathbf{R}^\top(t_{ij}) (\mathbf{P}_i - \mathbf{p}(t_{ij})) . \quad (1)$$

The above constraint implies that \mathbf{f}_{ij} and \mathbf{P}'_{ij} are parallel, which can be formulated as a constraint on their cross-product

$$\mathbf{f}_{ij} \times (\mathbf{R}^\top(t_{ij}) (\mathbf{P}_i - \mathbf{p}(t_{ij}))) = \mathbf{0}_{3 \times 1} . \quad (2)$$

Using properties of cross products and introducing the rotated bearing $\mathbf{f}'_{ij} \doteq \mathbf{R}(t_{ij})\mathbf{f}_{ij}$ yields

$$\mathbf{f}'_{ij} \times (\mathbf{P}_i - \mathbf{p}(t_{ij})) = \mathbf{0}_{3 \times 1} , \quad (3)$$

which is our desired incidence relation, also treated in [40]. In the appendix we show that this incidence relation can be specialized to the epipolar constraint used in the familiar 5-point or 8-point algorithm [15] or the line incidence relation used in the recent line solver in [10, 11]. In the next section, we will describe how to use this relation to solve for the 3D points and camera motion.

3.2. Transition to Linear System

Expressing all quantities with respect to the reference frame $\mathbf{C}(t_s) = \mathbf{I}_{4 \times 4}$ at time t_s we expand the motion of the camera with a Taylor Series $\mathbf{R}(t_{ij}) \approx \exp([\boldsymbol{\omega} t'_{ij}]_{\times})$ and $\mathbf{p}(t_{ij}) \approx \mathbf{v} t'_{ij}$. We introduce angular rate $\boldsymbol{\omega}$, linear velocity \mathbf{v} and relative timestamp $t'_{ij} = t_{ij} - t_s$. The operation $[\cdot]_{\times}$ maps vectors to a skew-symmetric matrix. Note that the degree of expansion is arbitrary, and each chosen degree will yield a given system of equations that are linear in the 3D points and body rates. In what follows, however, we will focus on a linear expansion, and will present the arbitrary case in the appendix, and applications in the experiments.

As done in previous work, we focus on finding the linear velocity \mathbf{v} for a given $\boldsymbol{\omega}$, which we assume to be given either by an external IMU (as in [11]), or other rotation estimation algorithms [7, 47].

We make use of the fact that the cross product with \mathbf{f}'_{ij} can be rewritten as a product with $[\mathbf{f}'_{ij}]_{\times}$

$$[\mathbf{f}'_{ij}]_{\times} \mathbf{P}_i - t'_{ij} [\mathbf{f}'_{ij}]_{\times} \mathbf{v} = \mathbf{0}_{3 \times 1}. \quad (4)$$

We gather all such constraints that involve the point \mathbf{P}_i into a single system of equations.

$$\underbrace{\begin{bmatrix} [\mathbf{f}'_{i1}]_{\times} & -t'_{i1} [\mathbf{f}'_{i1}]_{\times} \\ \vdots & \vdots \\ [\mathbf{f}'_{iN_i}]_{\times} & -t'_{iN_i} [\mathbf{f}'_{iN_i}]_{\times} \end{bmatrix}}_{\doteq [\mathbf{F}_i \quad \mathbf{G}_i] \in \mathbb{R}^{3N_i \times 6}} \begin{bmatrix} \mathbf{P}_i \\ \mathbf{v} \end{bmatrix} = \mathbf{0}_{3N_i \times 1}. \quad (5)$$

In a last step, we stack all such constraints originating from different points \mathbf{P}_i into one large system yielding

$$\underbrace{\begin{bmatrix} \mathbf{F}_1 & & & \mathbf{G}_1 \\ & \mathbf{F}_2 & & \mathbf{G}_2 \\ & & \ddots & \vdots \\ & & & \mathbf{F}_M & \mathbf{G}_M \end{bmatrix}}_{\doteq \mathbf{A} \in \mathbb{R}^{3N \times (3M+3)}} \underbrace{\begin{bmatrix} \mathbf{P}_1 \\ \mathbf{P}_2 \\ \vdots \\ \mathbf{P}_M \\ \mathbf{v} \end{bmatrix}}_{\doteq \mathbf{x} \in \mathbb{R}^{3M+3}} = \mathbf{0}_{3N \times 1}, \quad (6)$$

where we call $N = \sum_i N_i$ the total number of observations. Finally, we notice that this system imposes a linear constraint on the unknown points \mathbf{P}_i and camera velocity \mathbf{v} , and thus it admits an efficient solver, discussed next.

3.3. Solver

The linear system above could be solved with standard tools, by employing a singular value decomposition (SVD) on the matrix \mathbf{A} and then recovering the last column of the orthogonal matrix \mathbf{V} corresponding to the smallest singular value of \mathbf{A} . However, as the number of observations increases, computing the singular value decomposition may

start to pose a computational burden. For this reason, we first limit our focus to only recovering \mathbf{v} , and then show how to find the \mathbf{P}_i . As we will see, the sparse structure of \mathbf{A} allows the derivation of an efficient solver. We start off by left multiplying the linear system by \mathbf{A}^T , and writing the resulting system as a block system of equations:

$$\underbrace{\mathbf{A}^T \mathbf{A}}_{\doteq \mathbf{M}} \mathbf{x} = \begin{bmatrix} \mathbf{M}_A & \mathbf{M}_B \\ \mathbf{M}_B^T & \mathbf{M}_D \end{bmatrix} \begin{bmatrix} \mathbf{P}_{1:M} \\ \mathbf{v} \end{bmatrix} = \mathbf{0}_{(3M+3) \times 1}, \quad (7)$$

where we have stacked \mathbf{P}_i into $\mathbf{P}_{1:M}$, and the dimensions of the subblocks are $\mathbf{M}_A \in \mathbb{R}^{3M \times 3M}$, $\mathbf{M}_B \in \mathbb{R}^{3M \times 3}$ and $\mathbf{M}_D \in \mathbb{R}^{3 \times 3}$. We write out the explicit form of these matrices in the appendix. We then employ the Shur-complement trick to write a system only in \mathbf{v} which has the form

$$\underbrace{(\mathbf{M}_D - \mathbf{M}_B^T \mathbf{M}_A^{-1} \mathbf{M}_B)}_{\doteq \mathbf{B} \in \mathbb{R}^{3 \times 3}} \mathbf{v} = \mathbf{0}. \quad (8)$$

This last equation can be efficiently solved by employing SVD on the matrix \mathbf{B} , finding the normalized velocity estimate $\hat{\mathbf{v}}$ as the principle direction corresponding with the smallest singular value of \mathbf{B} . Note that the velocity is normalized due to absence of scale in a monocular setup.

One may think that the inversion of \mathbf{M}_A in Eq. 8 is expensive since $\mathbf{M}_A \in \mathbb{R}^{3M \times 3M}$ leading naively to $O(M^3)$ complexity. However, the matrix \mathbf{M}_A is actually block diagonal with M blocks of size 3×3 , leading to an efficient inversion algorithm of complexity $O(M)$ instead. Moreover, all terms can be computed from a linear combination of terms $[\mathbf{f}'_{ij}]_{\times}^2 = \mathbf{f}'_{ij} \mathbf{f}'_{ij}{}^T - \|\mathbf{f}'_{ij}\|^2 \mathbf{I}_{3 \times 3}$, leading to significant sharing of computation. Finally, having found the estimate $\hat{\mathbf{v}}$ we can find the solution to \mathbf{P}_i as

$$\hat{\mathbf{P}}_i = -(\mathbf{F}_i^T \mathbf{F}_i)^{-1} \mathbf{F}_i^T \mathbf{G}_i \hat{\mathbf{v}}, \quad (9)$$

which can be done efficiently by reusing computation from Eq. 8. Now let us analyze the properties of our solver, how many solutions it generates, and when it may fail.

3.4. Solver Properties

Solution Multiplicity: We start off by discussing the solution multiplicity of the above solver. The SVD operation in Eq. 8 yields two possible unit vectors $\hat{\mathbf{v}}$ and $-\hat{\mathbf{v}}$ which entail the two possible solutions $\hat{\mathbf{P}}_{1:M}$ and $-\hat{\mathbf{P}}_{1:M}$. During deployment, we select the correct solution by recognizing that the recovered points $\hat{\mathbf{P}}_i$ must have positive depth. We thus test for the following condition

$$(\hat{\mathbf{P}}_i)_z = -((\mathbf{F}_i^T \mathbf{F}_i)^{-1} \mathbf{F}_i^T \mathbf{G}_i \hat{\mathbf{v}})_z > 0 \quad (10)$$

and invert the velocity if it is violated.

Degeneracy: Next, let us analyze when degenerate solutions may be encountered. We see that the solution of \mathbf{v}

in Eq. 8 depends on the inversion of \mathbf{M}_A which in turn depends on the inversion of block diagonal matrices of the form $\mathbf{F}_i^\top \mathbf{F}_i$. To successfully invert these matrices, we require that \mathbf{F}_i has a full rank. \mathbf{F}_i is composed of matrices $[\mathbf{f}'_{ij}]_\times$ which only have two independent rows. We may thus consider the reduced form of \mathbf{F}_i with size $2N_i \times 3$ where we have deleted every third row. Thus, to have full rank $N_i \geq 2$, *i.e.* we must enforce that every track has at least two different observations. In practice, we do this by simply discarding tracks with only one observation. Finally, enforcing $\text{rank}(\mathbf{B}) \geq 2$ ensures that the SVD step succeeds in Eq. 8. For a discussion on the rank of \mathbf{B} see the appendix. **Constraint Analysis:** The system in Eq. 6 has $3M + 2$ unknowns (number of variables minus one for unobservability of scale), and $3N$ constraints. However, each $[\mathbf{f}'_{ij}]_\times$ only adds two linearly independent constraints, so the number of constraints is actually $2N$. Thus recovering all unknowns needs

$$2N \geq 3M + 2 \implies N \geq \left\lceil \frac{3M}{2} \right\rceil + 1 \quad (11)$$

observations. As discussed above, we require at least two observations per 3D point for stable inversion, *i.e.*

$$N \geq 2M \quad (12)$$

We now consider four cases:

- $M = 1$: Here $N = N_1 \geq 3$ leads to an overconstrained system of at least 6 equations with 5 unknowns. Dropping one equation makes this case minimal.
- $M = 2$: Here $N \geq 4$. In particular $N_1 = N_2 = 2$ leads to a *minimal* solver with 8 equations in 8 unknowns. For larger N the system becomes overconstrained again.
- $M = 3$: Here $N \geq 6$ with $N_1 = N_2 = N_3 = 2$ yielding a *minimal* system of 12 equations in 12 unknowns.
- $M > 3$: Here $N \geq 2M$ leading to an overconstrained system of $2N$ equations in $3M + 2$ unknowns.

Interestingly, the first three cases all give rise to minimal 3 point, 4 point or 6 point algorithms. We summarize the complete algorithm in Alg. 1.

3.5. Implementation Details

We derive feature tracks from a variety of input modalities using off-the-shelf trackers, and then embed the proposed point solver into a RANSAC loop. This loop removes outliers from poor tracking, and produces a refined estimate based on the found inliers [6].

Feature Tracking: We use off-the-shelf trackers designed for (i) global shutter cameras, (ii) rolling shutter cameras, and (iii) event cameras. Each tracker provides observations \mathbf{x}_{ij} , which are converted into rotated bearing vectors \mathbf{f}'_{ij} , based on IMU angular rate readings $\boldsymbol{\omega}$. For global shutter cameras, the timestamp t_{ij} of the observation is simply the image timestamp. For rolling shutter cameras $t_{ij} = \frac{y_{ij}}{(H-1)T_s}$

Algorithm 1 N-Point Solver for Structure & Motion

Input: A set of track observations $(\mathbf{x}_{ij}, t_{ij})$, and angular rate $\boldsymbol{\omega}$ from an IMU. Reference time t_s .

Output: Estimates of points $\hat{\mathbf{P}}_i$ and linear velocity $\hat{\mathbf{v}}$.

- Compute rotated bearing vectors $\mathbf{f}'_{ij} = \exp([\boldsymbol{\omega}t'_{ij}]_\times) \mathbf{f}_{ij}$.
 - Compute \mathbf{F}_i and \mathbf{G}_i in Eq. 5. Ensure that $\text{rank}(\mathbf{F}_i) = 3$, otherwise terminate.
 - Compute \mathbf{B} in Eq. 8 and solve for $\hat{\mathbf{v}}$. Terminate if the rank of \mathbf{B} is smaller than 2.
 - Compute $\hat{\mathbf{P}}_i$ from Eq. 9.
 - Check the depth via inequality Eq. 10. If it is negative, invert the signs of $\hat{\mathbf{P}}_i$ and $\hat{\mathbf{v}}$.
-

is corrected by the row index y_{ij} scaled by the row scanning time T_s , and height in pixels H of the sensor. For event cameras, timestamps are assigned based on measured events, resulting in asynchronous tracks. Tracks shorter than 2 are pruned to avoid degeneracy in the solver.

RANSAC: In each iteration of RANSAC we perform the following three steps: First, to balance computational complexity with spatio-temporal distribution, we sample M feature tracks, and then $N_i = n \geq 2$ temporally distributed observations $(\mathbf{f}'_{ij}, t'_{ij})$. Then, we generate a velocity hypothesis $\hat{\mathbf{v}}$ based on the $N = \sum_i N_i$ observations following the solver in Sec. 3.3, and rejecting the solution yielding negative point depth. Next, inliers are identified via the consistency of $\hat{\mathbf{v}}$ with observations $(\mathbf{f}'_{ij}, t_{ij})$. For track i , we predict the 3D point $\hat{\mathbf{P}}_i$ from Eq. 9, map it into the frame at each time t_{ij} , resulting in $\hat{\mathbf{P}}'_{ij} = \hat{\mathbf{P}}_i - \hat{\mathbf{v}}t'_{ij}$ and then project it into the current frame, yielding bearing estimate $\hat{\mathbf{f}}'_{ij}$. We use the average angular residual $\bar{\theta}$ between the observed and estimated bearing vectors along the track as error metric

$$\bar{\theta}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \arccos \left(\frac{\mathbf{f}'_{ij}{}^\top \hat{\mathbf{f}}'_{ij}}{\|\mathbf{f}'_{ij}\| \|\hat{\mathbf{f}}'_{ij}\|} \right) \quad (13)$$

A track is classified as an inlier if $\bar{\theta}_i$ is lower than a certain threshold (*e.g.* 5°), and the hypothesis with the highest number of inliers is retained throughout iteration.

After termination, the inliers corresponding with the best hypothesis are used to estimate the hypothesis $\hat{\mathbf{v}}$ leading to a refined estimate. The inlier ratio serves as a confidence metric reflecting the solver's robustness to outlier tracks.

4. Experiments

We comprehensively evaluate the performance of the proposed point solver, in two stages: First, we validate our method in a simulated environment (Sec. 4.1), where we study its sensitivity to different noise source including timestamp jitter, pixel noise and noise on the angular rate readings. We also study its accuracy as a function of the number of tracks and number of observations per track.




Figure 2. **Left:** Pixel noise sensitivity; **Middle:** Timestamp jitter impact; **Right:** Rotation perturbation effects. Each plot compares three observation levels with the number of features and the number of observation per track: sparse (5-5), moderate (20-20), and dense (100-50). Shaded regions represent error bounds across 1000 trials.

In a second step, we then report results in real world settings (Sec. 4.2), where we study its application to tracks derived from global-shutter, rolling-shutter and event-based cameras. Throughout the experimental section, we will report the accuracy of the scale-less velocity similar to [11], which is defined as the angular error between the true velocity \mathbf{v}_{gt} and estimate $\hat{\mathbf{v}}$ by

$$\theta_{\text{err}} = \arccos \left(\frac{\mathbf{v}_{\text{gt}}^T \hat{\mathbf{v}}}{\|\mathbf{v}_{\text{true}}\| \|\hat{\mathbf{v}}\|} \right) \quad (14)$$

Errors in 3D point estimation are not reported, as they are typically subsumed in the velocity error. In the appendix, we also apply our method for normalized acceleration estimation, with a similar error metric as above.

4.1. Simulation Experiments

We set a virtual camera with a resolution of 640×480 and a focal length of 320 pixels. A velocity vector with fixed magnitude $\|\mathbf{v}\| = 1\text{m/s}$ and random direction is generated to simulate camera motion. Static 3D points are randomly distributed within a one-meter cubic volume positioned two meters in front of the camera. This ensures that no points cross the camera plane during the motion. Observations are generated over a sliding time window of 0.2 seconds, with timestamps uniformly sampled within this interval. Each scenario is repeated 1,000 times to ensure statistical significance. Our solver achieves a minimal-case runtime of $63\mu\text{s}$ on CPU (Intel Xeon Platinum 8352V@3.5GHz). In the following sections, we analyze three key factors: noise resilience, observation count and track length.

4.1.1. Analysis of Noise Resilience

We first evaluate the solver’s robustness under three noise sources: inaccurate point tracking, temporal misalignment, and orientation drift in camera pose. As illustrated in Fig. 2, experiments vary noise level across practical ranges: pixel noise (0 - 5 pixels), timestamp jitter (0 - 50ms), and rotational perturbation (angular velocity noise of 0 - 30 deg/s),

while testing three different settings of observation. The results demonstrate that the solver achieves sub- 5° error at moderate noise levels (1 pixels, 10 ms jitter, 5 deg/s), validating its feasibility in typical operational scenarios. Notably, timestamp jitter exhibits a nonlinear error escalation, with error rising sharply beyond 15 ms. In contrast, pixel noise induces near-linear error scaling, suggesting the tolerance to common feature-tracking inaccuracies, while rotation perturbation has linear impact on the performance. We also notice more observations can effectively mitigate errors due to pixel noise and timestamp jitter, yet yield limited improvement for rotation-induced errors. Nevertheless, external sensor such as IMU can address this limitation by providing rotation-compensated inputs.

4.1.2. Analysis of Spatial-temporal Observations

We also analyzed the effects of feature track count M and per-track observation number $N_i = n$ under two combined noise conditions: $1\text{px} + 1\text{ms} + 2\text{deg/s}$ and $2\text{px} + 2\text{ms} + 5\text{deg/s}$. In general, track count (M) correlates with spatial resolution—high-resolution cameras (*e.g.* regular frame-based sensors) can maintain hundreds of tracks spatially. Observation count ($N_i = n$) depends on temporal resolution: event cameras, with microsecond-level precision, can densely sample a track over short time windows, enabling high n values (*e.g.* 50 observations) even in constrained durations, which explains our choice of the upper bound. As shown in Fig. 3 (left), increasing tracks from 3 to 30 significantly reduces velocity errors under both noise levels, but has marginal gains beyond 30 tracks. In contrast, raising observation counts per track under fixed time windows (Fig. 3 right) shows limited efficacy. Thus, the interplay between M and n reflects sensor-specific spatiotemporal trade-offs: frame-based systems excel in spatial coverage (high M) with temporally sparse measurements, while event cameras leverage temporal uniformity (high n) despite lower track counts. This highlights how sensor archi-




Figure 3. Track count (left) vs. Observation count per track (right). Box colors indicate noise levels (green: low, violet: high).

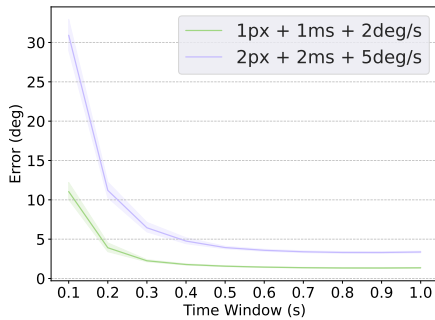


Figure 4. Velocity errors vs. time window length and noise.

texture shapes robustness under multi-source noise.

4.1.3. Analysis of Track Length

We further analyze the impact of the track length—determined by the temporal observation window—on velocity estimation. The time interval indicates the spatial displacement of feature tracks on the imaging plane: longer windows generally provide longer tracks, as features travel larger pixel distances under camera or scene motion. As illustrated in Fig. 4, we indirectly control track length by varying the size of the time window. Under combined noise conditions, larger time windows consistently improve the solver’s robustness. Longer tracks are able to average out high-frequency noise which enables the solver to recover stable velocity estimates. It also suggests that 3D points with smaller depth that induce larger apparent motion on the imaging plane can provide more reliable results. Notably, event cameras, which maintain temporally uniform observations, can benefit from larger time window.

4.2. Real Experiments

In a next step, we deploy our solver on real data collected with three different sensing modalities (as described in Sec. 3) using public datasets: The Event Camera Dataset [32] was collected using a DAVIS camera [2] which provides a pixel-aligned and time-synchronized stream of global shutter ($240 \times 180@24\text{Hz}$) and event camera data,

as well as IMU measurements. The VECtor dataset [9] provides event camera data and global shutter frames ($1224 \times 1024@30\text{Hz}$) from a stereo rig, as well as IMU data. Finally, the TUM dataset [41] provides rolling shutter images ($1280 \times 1024@20\text{Hz}$). We compare our method against eventail [11], an asynchronous line-based velocity solver. It uses clusters of events, each generated by a separate line to regress velocity components, and fuses these to a full normalized velocity via velocity averaging. On rolling shutter images we implement a similar baseline based on eventail, termed eventail + RS. eventail + RS first extracts Canny edges [3] from the rolling shutter images, and then treats the detected points as events with timestamps assigned according to their row index. We do not report acceleration estimation results on real data due to high noise sensitivity, and leave addressing this challenge for future work.

Experimental Setup: For global and rolling shutter images, we employ the standard Kanade-Lucas-Tomasi Tracker [29]. For global shutter images we assign timestamps at the midpoint of the exposure time. For rolling shutter images we compute timestamps based on the row index of the feature and scanning time. For event camera data we use the recent learning-based point tracker ETAP [14] to generate point tracks that preserve the high temporal resolution. To ensure the numerical stability of the solver, we filter out tracks shorter than 10 pixels, as these often amplify noise. All methods operate within identical time windows to enable fair comparison. In each sequence we use angular rates measured by an IMU to rotation compensate the bearing vectors. The RANSAC pipeline is configured with a maximum iteration of 200, with each iteration sampling $M = 4$ randomly selected tracks containing $N_i = 5$ observations. A point track is classified as an inlier if its angular error falls below 5° . To speed up convergence, we terminate RANSAC when the inlier ratio exceeds 0.9.

Application to Global Shutter Cameras We report the result of our method in Tab. 1. It can be seen that our method using global shutter images achieves a lower error than the eventail solver operating on events. This is due to two effects: First, the eventail solver relies on extracting events

Table 1. Mean / median velocity error in degrees on tracks from global shutter images. * and gray: results on subset with track inlier ratio > 0.9. eventail + E uses events from an event camera.

Seq.	eventail [11] + E	Ours	Ours*
<i>desk-normal</i>	22.7 / 23.4	15.1 / 8.5	10.2 / 7.3
<i>sofa-normal</i>	21.9 / 17.6	15.9 / 7.8	9.8 / 6.3
<i>mountain-normal</i>	25.2 / 21.4	17.1 / 7.5	10.9 / 6.1
<i>shapes_translation</i>	31.8 / 32.7	17.1 / 7.2	9.9 / 6.2
<i>boxes_translation</i>	34.8 / 34.1	16.5 / 11.6	13.3 / 10.7

Table 2. Mean / median velocity error in degrees on tracks from rolling shutter images. * and gray: results on subset with inlier ratio > 0.9. eventail + RS uses canny edges from rolling shutter images. “no correction”: no rolling shutter timestamp correction.

Seq.	with correction			no correction	
	eventail [11] + RS	Ours	Ours*	Ours	Ours*
<i>Seq 4</i>	43.8 / 40.8	27.5 / 20.1	22.6 / 17.4	28.1 / 22.9	22.8 / 15.7
<i>Seq 5</i>	45.5 / 44.8	24.7 / 17.0	19.3 / 13.8	27.0 / 18.4	19.2 / 14.6

generated by 3D lines in the scene, which is limited to highly geometric structures. By contrast our method can rely on tracks which can be extracted more easily. The second effect is the use of images vs. events. Events are known to suffer from changing appearance due motion changes, leading to drift in the feature tracks [12, 13]. We will see later that combining tracks from colocated GS cameras and events can improve results, even beyond the results using global shutter cameras alone. Next, we focus on the results marked with *, indicating evaluation on the subset where over 90% of tracks are termed inliers. On this subset, errors are further reduced, showing the importance of having geometrically consistent observations.

Application to Rolling Shutter Cameras We report the results of our method in Tab. 2. We see that our point solver yields a significant improvement with respect to the eventail solver. This is mainly due to the fact that eventail found few lines in the presented sequences. By contrast, our method relies on feature tracks which are more easily extracted. We also show a result without rolling shutter timestamp correction, denoted with “no correction”. It is visible that this reduces the accuracy of the method by a few degrees, showing the benefit of correct timestamp association, and the flexibility of our method to take non-synchronized feature tracks into account. Finally, as before we see that results on the subset marked with * are better, indicating the importance of geometrically consistent tracks for estimation.

Application to Event-based Cameras Finally, we apply our method to tracks derived from an event-based camera, and show results in Tab. 3. Our point solver running on events alone outperforms eventail by 10 – 30%, and this is again due to the use of point tracks instead of lines. Interestingly, frame-based tracks yield better perfor-

Table 3. Mean / median velocity error in degrees on tracks from events. * and gray: results on subset with track inlier ratio > 0.9. E stands for events from an event camera. E+GS refers to our method combining tracks from events and global shutter images. Note that for - no colocated GS and event sensor are available.

Seq.	eventail [11] + E	Ours + E	Ours* + E	Ours + E + GS	Ours* + E + GS
<i>desk-normal</i>	22.7 / 23.4	19.3 / 17.8	14.2 / 14.2	-	-
<i>sofa-normal</i>	21.9 / 17.6	19.0 / 18.5	16.3 / 14.9	-	-
<i>mountain-normal</i>	25.2 / 21.4	17.1 / 16.1	16.9 / 15.8	-	-
<i>shapes_translation</i>	31.8 / 32.7	16.8 / 10.1	13.0 / 9.1	14.4 / 7.5	7.0 / 6.7
<i>boxes_translation</i>	34.3 / 34.1	12.6 / 10.0	12.1 / 7.7	10.3 / 8.1	9.3 / 5.9

mance than event-based ones, particularly on the VECTOR sequences. This aligns with our simulation findings where higher feature density (from frame cameras’ 5× resolution advantage over event cameras) improves spatial sampling. We show the benefit of combining tracks from different sensing modalities, denoted with E+GS. In particular, the sequences *shapes_translation* and *boxes_translation* were recorded with a DAVIS camera [2] which features pixels that simultaneously record events and global shutter images. In this setting, we see that adding images significantly reduces errors. Moreover, comparing to Tab. 1 we also see that adding events improves over the global shutter result, highlighting the complementarity of the sensors. This result highlights the benefit of having an asynchronous point solver that can flexibly incorporate both global shutter and event-based observations.

5. Future Work and Conclusion

Future Work: While we believe that the proposed solver makes a significant stride toward handling asynchronous tracks in an efficient way, we acknowledge its dependence on available angular rates from an IMU. Initial steps have been made in incorporating angular rate estimation into existing solvers [10, 11], but further work is needed to make these solvers efficient. Finally, we only show linear acceleration estimation in simulation (see appendix), and found significant challenges with noise on real-world data. This indicates estimation stability issues for higher-order derivatives. Future work should aim to identify and reduce the effect of noise on higher-order derivative estimation.

Conclusion: We present a linear N-point solver for recovering structure and linear motion from asynchronous feature tracks. It generalizes solvers that rely on additional structure constraints such as points lying on a line, or time constraints, such as assuming synchronized timestamps. We showed experimentally that the motions recovered by our solver are more accurate than those produced by previous work, and also more robust in natural, line-deprived environments. We believe that our solver sets the stage for many new innovations to come by enabling the seamless integration of asynchronous feature tracks into geometric solvers.

Acknowledgments

This research has been supported by project 62250610225 by the Natural Science Foundation of China, as well as projects 22DZ1201900, and dfycbj-1 by the Natural Science Foundation of Shanghai.

References

- [1] Cenek Albl, Zuzana Kukelova, Viktor Larsson, and Tomas Pajdla. Rolling shutter camera absolute pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 42(6):1439–1452, 2020. 2
- [2] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. A 240×180 130db $3\mu\text{s}$ latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits*, 49(10):2333–2341, 2014. 7, 8
- [3] John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, PAMI-8(6):679–698, 1986. 7
- [4] William Chamorro, Joan Solà, and Juan Andrade-Cetto. Event-imu fusion strategies for faster-than-imu estimation throughput. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3976–3983, 2023. 3
- [5] Yuchao Dai, Hongdong Li, and Laurent Kneip. Rolling shutter camera relative pose: Generalized epipolar geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4132–4140, 2016. 2
- [6] 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. 5
- [7] Guillermo Gallego and Davide Scaramuzza. Accurate angular velocity estimation with an event camera. *IEEE Robotics and Automation Letters*, 2(2):632–639, 2017. 4
- [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 (CVPR)*, pages 3867–3876, 2018. 3
- [9] 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. 7
- [10] Ling Gao, Hang Su, Daniel Gehrig, Marco Cannici, Davide Scaramuzza, and Laurent Kneip. A 5-point minimal solver for event camera relative motion estimation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 8015–8025, 2023. 2, 3, 8
- [11] Ling Gao, Daniel Gehrig, Hang Su, Davide Scaramuzza, and Laurent Kneip. A linear n-point solver for line and motion estimation with event cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 4, 6, 7, 8
- [12] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Asynchronous, photometric feature tracking using events and frames. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 750–765, 2018. 3, 8
- [13] Daniel Gehrig, Henri Rebecq, Guillermo Gallego, and Davide Scaramuzza. Ekl: Asynchronous photometric feature tracking using events and frames. *International Journal of Computer Vision (IJCV)*, 128(3):601–618, 2020. 3, 8
- [14] Friedhelm Hamann, Daniel Gehrig, Filbert Febryanto, Kostas Daniilidis, and Guillermo Gallego. Event-based tracking of any point with motion-robust correlation features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18–37. Springer, 2025. 3, 7
- [15] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004. 2, 3
- [16] Kun Huang, Yifu Wang, and Laurent Kneip. Motion estimation of non-holonomic ground vehicles from a single feature correspondence measured over n views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12698–12707, 2019. 2
- [17] Haram Kim and H Jin Kim. Real-time rotational motion estimation with contrast maximization over globally aligned events. *IEEE Robotics and Automation Letters*, 6(3):6016–6023, 2021. 3
- [18] Hanme Kim, Stefan Leutenegger, and Andrew J. Davison. Real-time 3d reconstruction and 6-dof tracking with an event camera. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 349–364, 2016. 3
- [19] Simon Klenk, Marvin Motzet, Lukas Koestler, and Daniel Cremers. Deep event visual odometry. In *International Conference on 3D Vision (3DV)*, pages 739–749. IEEE, 2024. 3
- [20] Laurent Kneip, Agostino Martinelli, Stephan Weiss, Davide Scaramuzza, and Roland Siegwart. Closed-form solution for absolute scale velocity determination combining inertial measurements and a single feature correspondence. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 4546–4553. IEEE, 2011. 2
- [21] Zuzana Kukelova, Cenek Albl, Akihiro Sugimoto, and Tomas Pajdla. Linear solution to the minimal absolute pose rolling shutter problem. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 265–280. Cham, 2019. Springer International Publishing. 2
- [22] Zuzana Kukelova, Cenek Albl, Akihiro Sugimoto, Konrad Schindler, and Tomas Pajdla. Minimal rolling shutter absolute pose with unknown focal length and radial distortion. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 698–714. Springer, 2020. 2
- [23] Xavier Lagorce, Sio-Hoi Ieng, Xavier Clady, Michael Pfeiffer, and Ryad B Benosman. Spatiotemporal features for asynchronous event-based data. *Frontiers in neuroscience*, 9:46, 2015. 3
- [24] Yizhen Lao and Omar Ait-Aider. Rolling shutter homography and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 43(8):2780–2793, 2021. 2

- [25] Patrick Lichtensteiner, Christoph Posch, and Tobi Delbruck. A 128×128 120db $15\mu\text{s}$ latency asynchronous temporal contrast vision sensor. *IEEE Journal of Solid-State Circuits (JSSC)*, (2):566–576, 2008. 2
- [26] Daqi Liu, Alvaro Parra, and Tat-Jun Chin. Globally optimal contrast maximisation for event-based motion estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6348–6357, 2020. 3
- [27] Daqi Liu, Alvaro Parra, and Tat-Jun Chin. Spatiotemporal registration for event-based visual odometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4937–4946, 2021. 3
- [28] Xiuyuan Lu, Yi Zhou, Junkai Niu, Sheng Zhong, and Shaojie Shen. Event-based visual inertial velometer. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024. 3
- [29] Bruce D. Lucas and Takeo Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 674–679. William Kaufmann, 1981. 7
- [30] Nico Messikommer*, Carter Fang*, Mathias Gehrig, and Davide Scaramuzza. Data-driven feature tracking for event cameras. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [31] Nico Messikommer, Carter Fang, Mathias Gehrig, Giovanni Cioffi, and Davide Scaramuzza. Data-driven feature tracking for event cameras with and without frames. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2025. 3
- [32] 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. *International Journal of Robotics Research (IJRR)*, 36(2):142–149, 2017. 7
- [33] Elias Mueggler, Guillermo Gallego, Henri Rebecq, and Davide Scaramuzza. Continuous-time visual-inertial odometry for event cameras. *IEEE Transactions on Robotics (T-RO)*, 34(6):1425–1440, 2018. 3
- [34] D. Nister. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 26(6):756–770, 2004. 2
- [35] Junkai Niu, Sheng Zhong, Xiuyuan Lu, Shaojie Shen, Guillermo Gallego, and Yi Zhou. Esvo2: Direct visual-inertial odometry with stereo event cameras. *arXiv preprint arXiv:2410.09374*, 2024. 3
- [36] Roberto Pellerito, Marco Cannici, Daniel Gehrig, Joris Belhadj, Olivier Dubois-Matra, Massimo Casasco, and Davide Scaramuzza. Deep visual odometry with events and frames. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, 2024. 3
- [37] Xin Peng, Wanting Xu, Jiaqi Yang, and Laurent Kneip. Continuous event-line constraint for closed-form velocity initialization. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2021. 3
- [38] Xin Peng, Ling Gao, Yifu Wang, and Laurent Kneip. Globally-optimal contrast maximisation for event cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 44(7):3479–3495, 2022. 3
- [39] 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. 3
- [40] Olivier Saurer, Marc Pollefeys, and Gim Hee Lee. A minimal solution to the rolling shutter pose estimation problem. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 1328–1334, 2015. 2, 3
- [41] David Schubert, Nikolaus Demmel, Lukas Von Stumberg, Vladyslav Usenko, and Daniel Cremers. Rolling-shutter modelling for direct visual-inertial odometry. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, pages 2462–2469. IEEE, 2019. 7
- [42] Shintaro Shiba, Yannick Klose, Yoshimitsu Aoki, and Guillermo Gallego. Secrets of event-based optical flow, depth and ego-motion estimation by contrast maximization. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2024. 3
- [43] Henrik Stewénius, Christopher Engels, and David Nistér. Recent developments on direct relative orientation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 60(4):284–294, 2006. 2
- [44] Timo Stoffregen and Lindsay Kleeman. Event cameras, contrast maximization and reward functions: An analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12292–12300, 2019. 3
- [45] 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. 3
- [46] Wanting Xu, Xin Peng, and Laurent Kneip. Tight fusion of events and inertial measurements for direct velocity estimation. *IEEE Transactions on Robotics (T-RO)*, 40:240–256, 2024. 3
- [47] Ji Zhao, Banglei Guan, Zibin Liu, and Laurent Kneip. Full-dof egomotion estimation for event cameras using geometric solvers. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. 3, 4
- [48] 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 (CVPR)*, pages 5816–5824, 2017. 3