

Event-Based Stereo Visual Odometry With Native Temporal Resolution via Continuous-Time Gaussian Process Regression

Jianeng Wang¹ and Jonathan D. Gammell¹

Abstract—Event-based cameras asynchronously capture individual visual changes in a scene. This makes them more robust than traditional frame-based cameras to highly dynamic motions and poor illumination. It also means that every measurement in a scene can occur at a unique time. Handling these different measurement times is a major challenge of using event-based cameras. It is often addressed in visual odometry (VO) pipelines by approximating temporally close measurements as occurring at one common time. This grouping simplifies the estimation problem but, absent additional sensors, sacrifices the inherent temporal resolution of event-based cameras. This paper instead presents a complete stereo VO pipeline that estimates directly with individual event-measurement times without requiring any grouping or approximation in the estimation state. It uses continuous-time trajectory estimation to maintain the temporal fidelity and asynchronous nature of event-based cameras through Gaussian process regression with a physically motivated prior. Its performance is evaluated on the MVSEC dataset, where it achieves $7.9 \cdot 10^{-3}$ and $5.9 \cdot 10^{-3}$ RMS relative error on two independent sequences, outperforming the existing publicly available event-based stereo VO pipeline by two and four times, respectively.

Index Terms—Event-based Visual Odometry, Vision-Based Navigation, Localization, SLAM.

I. INTRODUCTION

VISUAL Odometry (VO) is a technique to estimate ego-motion in robotics [1], [2], [3], [4], [5]. VO systems using traditional frame-based cameras often struggle in scenarios with high speed motion and poor illumination. In these scenarios, the motion blur and poor image contrast of frame-based cameras result in bad estimation performance.

Event-based cameras perform better than traditional cameras in these challenging scenarios. They detect pixelwise intensity change and report the time at which the change occurs asynchronously. This gives them high temporal resolution and high dynamic range avoiding the limitations of frame-based cameras and providing the potential for more accurate VO systems [6]. Any event-based VO system must address the asynchronous event times. Many pipelines do this by grouping similar feature times to a common time [7], [8], [9]. This allows for the

Manuscript received 1 June 2023; accepted 17 August 2023. Date of publication 4 September 2023; date of current version 12 September 2023. This letter was recommended for publication by Associate Editor L. Zhang and Editor J. Civera upon evaluation of the reviewers' comments. (Corresponding author: Jianeng Wang.)

The authors are with the Estimation, Search, and Planning (ESP) Group, Oxford Robotics Institute (ORI), University of Oxford, OX1 4AR Oxford, U.K. (e-mail: jianeng@robots.ox.ac.uk; gammell@robots.ox.ac.uk).

Digital Object Identifier 10.1109/LRA.2023.3311374




Fig. 1. Illustration of the continuous-time trajectory estimation pipeline. Event clusters are defined by an asynchronous event data stream in discrete windows based on number of events and their times (e.g., $t_1 \leq t \leq t_2$). Features (red) are detected from the resulting clusters and matched with features occurring in other clusters (blue) to create tracklets. Each event feature in the tracklet, $y_{j,k}$, is a measurement of landmark, p_j , and defines a trajectory state in the estimation problem, $x_k = \{\mathbf{T}_{k,1}, \boldsymbol{\omega}_k\}$, at the measurement time, t_k . The camera motion is estimated as a continuous-time trajectory function, $x(t)$, defined by the discrete states and a physically founded motion prior.

direct application of frame-based VO pipelines but sacrifices the temporal resolution of event cameras.

This paper instead presents an event-based VO system that uses the unique asynchronous timestamps directly in the estimation problem without grouping or approximation. It estimates the camera motion as a continuous-time trajectory represented by states at unique feature times and a white-noise-on-acceleration (WNOA) motion prior. The trajectory is estimated using non-parametric Gaussian process regression. This results in a continuous, physically founded trajectory that exploits the temporal resolution of asynchronous event cameras and can estimate complex, real-world motions. (Fig. 1). This paper presents a complete event-based stereo VO pipeline using continuous-time Gaussian process regression. It is compatible with any feature detector and tracker, including frame-based methods for traditional cameras, without reducing the temporal resolution of event-based cameras. It also uses Motion-Compensated RANSAC (MC-RANSAC) [10] to consider the unique measurement times during outlier rejection and independently provide better tracklets and initial conditions for estimation. The resulting continuous-time trajectory provides estimates of camera poses at any and all timestamps in the estimation window.

It is evaluated on the publicly available Multi Vehicle Stereo Event Camera (MVSEC) dataset [11], where it obtains a more accurate and smoother trajectory estimate than the state-of-the-art Event-based Stereo Visual Odometry (ESVO) [9]. It achieves

$7.9 \cdot 10^{-3}$ and $5.9 \cdot 10^{-3}$ root-mean-squared (RMS) relative error in $SE(3)$ and 5.78% and 4.93% final global translational error as a percent of path length on two independent MVSEC sequences. This outperforms the publicly available ESVO on these sequences, especially in terms of trajectory smoothness and RMS relative error where it is two- and four-times better, respectively. The rest of the paper is organized as follows. Section II summarizes the existing literature on event-based VO. Section III presents the complete pipeline of Gaussian-process continuous-time VO. Section IV evaluates the system and ESVO on the MVSEC dataset and discusses the results. Section V presents the summary of the work.

II. RELATED WORK

Event-based motion estimation techniques can be described by their handling of asynchronous event times as either grouping times into discrete frames (Section II-A) or considering times individually (Section II-B).

A. Grouped-Time Approaches

Grouping event data together creates data frames like traditional cameras. Traditional VO pipelines can then be used to estimate camera poses at the discrete times assigned to the event frames. These frame times reduce the temporal resolution of the data by replacing all the individual events in a frame with a single time.

Initial event-based VO research focuses on simplified scenarios, e.g., 2D planar motion [12] or rotation-only motion [13]. Research extends estimation to $SE(3)$ motion by incorporating complementary sensors into the estimation pipeline. Kueng et al. [14] fuse both events and traditional-camera frames to detect features on the image frames and track them using events. Weikersdorfer et al. [15] use both event and RGB-D cameras to provide depth information for each event and create a 3D map for localization.

Event-based Visual Inertial Odometry (VIO) systems specifically include inertial measurement units (IMUs) and are a popular area of research. Zhu et al. [16] and Rebecq et al. [17] accumulate events in a spatial-temporal window to reconstruct frames for feature tracking. The resulting tracklets are used to minimize reprojection and initial error for pose estimation. Ultimate-SLAM [7] extends [17] to use the IMU to generate motion-compensated event frames and reformulates the cost function for camera egomotion estimation. Chen et al. [18] extract features using an asynchronous feature detector (Arc* [19]) from a stereo pair of event cameras and adopt an estimation pipeline similar to [7]. IMU Dynamic Vision Sensor Odometry using Lines (IDOL) [20] uses an alternative VIO paradigm. It uses Gaussian process regression to preintegrate IMU measurements and associate each event with timely accurate IMU data. This maintains the temporal resolution of the event data but still estimates the trajectory at the discrete states of the event frames.

Kim et al. [21] estimate $SE(3)$ motion using only event-based cameras. They interleave three filters to estimate camera motion, intensity-based frames and scene depth. EVO [8] improves computational performance by interleaving geometric semidense mapping [22] and image-to-model alignment for pose estimation. ESVO [9] uses parallel tracking and mapping to estimate the egomotion trajectory and a semidense 3D scene reconstruction. Hadviger et al. [23] present a feature-based stereo VO pipeline using events, which group events into frames

and then adopts a similar estimation framework to a traditional frame-based stereo VO pipeline [3].

Discrete-time event-based VO groups multiple event times into a single time. This is helpful for feature detection and tracking but approximates the time of individual event features and reduces the temporal resolution of the measured data. This approximation is inconsistent with the asynchronous nature of event cameras and introduces potential measurement errors. This paper instead presents a VO pipeline that estimates the camera trajectory from individual event times and maintains the temporal resolution of event cameras. It uses only an event-camera stereo pair and can be implemented with either frame-based or event-based feature detection and tracking.

B. Individual-Time Approaches

Including all the individual, possibly unique event times in the estimation maintains the temporal resolution of event cameras but defines an underconstrained problem. Similar estimation problems are solved for rolling-shutter cameras and scanning lidars using continuous-time estimation [24], [25]. These techniques estimate the camera trajectory as a continuous function where pose can be queried at any time in the estimation window. A comparison of discrete and continuous-time trajectories can be found in [26].

Mueggler et al. [27] use a continuous-time pose estimation framework that uses IMU measurements and represents the trajectory as cumulative cubic B-splines. This avoids grouping event feature times and maintains their temporal resolution, but requires preprocessing to obtain a scene map. Wang et al. [28] use volumetric contrast maximization for continuous-time estimation. The trajectory is initialized with an Ackermann motion model [29], and globally optimized with a B-spline-based continuous-time estimation framework. This allows the estimator to maintain the native temporal resolution of the event data but limits it to planar motion.

Liu et al. [30] also estimate continuous-time trajectories with Gaussian process regression and a WNOA motion prior, but in a monocular VO system. Their estimation runs asynchronously and considers individual event times but, unlike traditional VO pipelines, couples tracklet outlier rejection with the motion estimation which may introduce different sources of error. The evaluation of their algorithm on real-world event datasets is also limited to five-second sequences.

This paper presents a complete continuous-time event-based stereo VO pipeline that maintains individual event times in the trajectory estimation. In contrast to these existing works, it maintains temporal resolution with either frame-based or event-based feature detection and tracking and with a RANSAC formulation [10] that separates outlier rejection from estimation. It uses a WNOA motion prior to estimate a full $SE(3)$ trajectory directly from event tracklets and their unique timestamps. This approach takes full advantage of the asynchronous nature of the event cameras and allows pose to be queried at any time in the estimation window.

III. METHODOLOGY

This paper presents an event-based stereo VO system that uses the native temporal resolution of event-based cameras for estimation (Fig. 2). Features are detected by clustering the event streams of each camera while maintaining the unique event times and tracked in a traditional frame-based manner (Section III-A).




Fig. 2. Overview of the presented VO pipeline. The system takes an asynchronous stereo event stream and clusters the events into event frames and SAEs. Features are detected and tracked in the event frames and each feature is assigned an event time from the SAE. The resulting asynchronous event feature tracklets are filtered with a motion-compensated RANSAC that accounts for their asynchronous times. This gives a consistent inlier tracklet set and motion prior for the continuous-time estimator.

This allows for the use of any frame-based feature detection and tracking method or could be directly replaced by event-based approaches. The resulting asynchronous tracklets are then filtered for outliers with a motion-compensated RANSAC (Section III-B). This removes outliers more accurately than other methods by accounting for the different tracklet times. The camera trajectory is then estimated from all the unique tracklet states using Gaussian process regression with a WNOA prior (Section III-C). This results in a continuous-time VO system that estimates the camera pose at each unique tracklet time and can be queried for the pose at any other time in the estimation window.

A. Event Feature Extraction and Matching

Geometric features are extracted from clustered events (Section III-A1) using traditional frame-based feature detection techniques, and then matched between the stereo pair and through time to construct feature tracklets (Section III-A2). The temporal resolution of event cameras is maintained by assigning (possibly unique) times to each feature from the corresponding event in a Surface of Active Events (SAE) [31]. This allows the pipeline to use any frame-based feature detector and tracker and still create asynchronous tracklets for estimation. It could also be directly replaced with event-based methods, e.g., Arc* [19] or HASTE [32].

1) *Event Clustering*: The stereo event stream is rectified and clustered to construct new SAEs and new binary event frames (Fig. 2). The left and right event streams are synchronously clustered by registering events within a user-specified time window or until either camera registers more than a user-specified number of events in that time. These thresholds define a minimum effective frame rate when the scene changes slowly (e.g., small motion) and a faster frame rate when it changes rapidly (e.g., high-speed motion), while keeping the left and right frames synchronized for feature matching. The SAE records the most recent event timestamp of each pixel location and is used to maintain the asynchronous nature of the event camera. The binary event frame is a grey image where white pixels denote events, regardless of polarity.

2) *Feature Detection and Matching*: Features are independently detected in the left and right binary event frames for each stereo pair of clusters (Fig. 2). These features are then matched using a quad-matching scheme. Features in the current left frame are first matched to the current right frame. The matched features are then successively matched from the current right frame to the previous right frame, from the previous right frame to previous left frame and finally from the previous left frame back to the current left frame. Features successfully matched to all these pairs are kept as tracklets. This feature detection and tracking can use any traditional frame-based approach.

The timestamp of each tracklet state is assigned from the nearest event in the associated SAE. This allows traditional frame-based feature detection methods to detect asynchronous tracklets that maintain the original temporal resolution of the event data. Event-based feature detection and tracking algorithms generate these asynchronous tracklets directly and could be used instead.

Tracklets are extended beyond consecutive frames by matching the new tracklets to previously detected tracklets in earlier frames. This is done by independently checking each new tracklet against a user-specified number of previous frames and recording any additional matches.

The tracklets are filtered using user-specified thresholds to discard those that have

- 1) a large time difference between stereo features,
- 2) a small disparity between stereo features,
- 3) a short length, and
- 4) a short time.

These filtering schemes remove incorrect matches and improve tracklet quality. This feature extraction and matching process is performed independently at different image resolutions to detect and track features of different size. The resulting tracklets are then processed together to remove outliers.

B. Outlier Rejection

Traditional VO pipeline uses Random Sample Consensus (RANSAC) [33] to remove tracklet outliers before estimation. Traditional RANSAC assumes tracklet states occur at common times and uses a discrete transform motion model. This assumption is incorrect for asynchronous event tracklets and this paper instead uses MC-RANSAC [10], which makes no assumptions about common state times and uses a constant-velocity model in $SE(3)$. The fast version of MC-RANSAC (Section III-B1) is used to find an initial inlier set by repeatedly selecting tracklets, calculating the constant-velocity model, and segmenting the tracklets into inliers and outliers based on a user-specified threshold. This process is repeated a user-specified number of times and then the largest inlier set is refined using the full iterative MC-RANSAC (Section III-B2) to find the final inlier and outlier segmentation (Fig. 2). Both versions of MC-RANSAC are compared to traditional RANSAC in [10].

1) *Fast MC-RANSAC*: The set of tracklets between two event clusters is a number, M_{clst} , of stereo measurements of different landmarks at possibly unique times. MC-RANSAC segments these tracklets into inliers and outliers by finding the most tracklets that can be explained by a single velocity. The constant $SE(3)$ velocity of the sensor, $\varpi \in \mathbb{R}^{6 \times 1}$, is

$$\varpi = \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\omega} \end{bmatrix},$$

where $\mathbf{v}, \boldsymbol{\omega} \in \mathbb{R}^{3 \times 1}$ are linear and angular velocity components, respectively. This constant velocity over a time, Δt , gives the relative $SE(3)$ transformation,

$$\mathbf{T} = \exp(\Delta t \boldsymbol{\omega}^\wedge),$$

where $\exp(\cdot)$ is the matrix exponential and $(\cdot)^\wedge$ is the lifting operator that converts $\mathbb{R}^{6 \times 1}$ to $\mathbb{R}^{4 \times 4}$ [34].

A velocity can be calculated from tracklets faster but less accurately by minimizing the error in Euclidean space,

$$J_{\text{fast}}(\boldsymbol{\omega}) = \frac{1}{2} \sum_{j=1}^{M_{\text{rand}}} \mathbf{e}_{\text{fast},j,k}^T \mathbf{e}_{\text{fast},j,k}, \quad (1)$$

where $M_{\text{rand}} \leq M_{\text{clst}}$ is the the number of randomly selected tracklets and $\mathbf{e}_{\text{fast},j,k} \in \mathbb{R}^{3 \times 1}$ is the motion-model error,

$$\mathbf{e}_{\text{fast},j,k} = \mathbf{P} \left(\mathbf{p}_k^{j,k} - \mathbf{z} \left(\mathbf{T}_{k,k'} \mathbf{p}_{k'}^{j,k'} \right) \right), \quad (2)$$

where \mathbf{P} is the mapping from homogeneous to 3D coordinates, $\mathbf{p}_k^{j,k} \in \mathbb{R}^{4 \times 1}$ is the position of the j^{th} landmark relative to the camera pose at time t_k , $\mathbf{T}_{k,k'}$ is the transformation from timestamp $t_{k'}$ to t_k , $t_{k'} < t_k$ is the earlier time of the tracklet segment, and $\mathbf{z}(\cdot)$ is the constant-velocity motion model,

$$\mathbf{z} \left(\mathbf{T}_{k,k'} \mathbf{p}_{k'}^{j,k'} \right) = \mathbf{T}_{k,k'} \mathbf{p}_{k'}^{j,k'}$$

Assuming a short duration tracklet, the corresponding small transformation is approximated as

$$\mathbf{T} \approx \mathbf{1} + \boldsymbol{\xi}^\wedge = \mathbf{1} + \Delta t \boldsymbol{\omega}^\wedge, \quad (3)$$

where $\mathbf{1} \in \mathbb{R}^{4 \times 4}$ is the identity matrix, $\boldsymbol{\xi} \in \mathbb{R}^{6 \times 1}$ is the pose in vector form and Δt is the duration of the tracklet. Substituting (3) into (2) approximates the error term as,

$$\begin{aligned} \mathbf{e}_{\text{fast},j,k} &\approx \mathbf{P} \left(\mathbf{p}_k^{j,k} - (\mathbf{1} + \Delta t_{k,k'} \boldsymbol{\omega}^\wedge) \mathbf{p}_{k'}^{j,k'} \right) \\ &= \mathbf{d}_{j,k} - \Delta t_{k,k'} \mathbf{D}_{j,k} \boldsymbol{\omega}, \end{aligned} \quad (4)$$

where

$$\begin{aligned} \Delta t_{k,k'} &= t_k - t_{k'}, \\ \mathbf{d}_{j,k} &= \mathbf{P} \left(\mathbf{p}_k^{j,k} - \mathbf{p}_{k'}^{j,k'} \right), \\ \mathbf{D}_{j,k} &= \mathbf{P} \left(\mathbf{p}_{k'}^{j,k'} \right)^\odot, \end{aligned}$$

and $(\cdot)^\odot$ is the $\mathbb{R}^{4 \times 1}$ to $\mathbb{R}^{4 \times 6}$ operator [34]. Substituting (4) into the Euclidean space cost function in (1), differentiating with respect to $\boldsymbol{\omega}$, and setting the result equal to zero gives the velocity best describing a set of M_{rand} tracklets,

$$\boldsymbol{\omega} = \left(\sum_{j=1}^{M_{\text{rand}}} \Delta t_{j,k}^2 \mathbf{D}_{j,k}^T \mathbf{D}_{j,k} \right)^{-1} \left(\sum_{j=1}^{M_{\text{rand}}} \Delta t_{j,k} \mathbf{D}_{j,k}^T \mathbf{d}_{j,k} \right).$$

This velocity can then be used to segment all M_{clst} tracklets into inliers and outliers by comparing their reprojection error relative to tracklet length,

$$e_{\text{rel}} = \frac{\mathbf{y}_{j,k} - \mathbf{s} \left(\mathbf{z} \left(\exp(\Delta t_{k,k'} \boldsymbol{\omega}^\wedge), \mathbf{p}_{k'}^{j,k'} \right) \right)}{\|\mathbf{y}_{j,k} - \mathbf{y}_{j,k'}\|}, \quad (5)$$

to a user-specified threshold, where $\mathbf{y}_{j,k}$ is a measurement of the j^{th} landmark at time t_k and $\mathbf{s}(\cdot)$ is the nonlinear camera projection model from a 3D landmark to a 2D image point.

The process of randomly selecting a small number of tracklets, quickly calculating a velocity from them, and then using this velocity to classify all the tracklets as inliers or outliers is repeated a user-specified number of times. The largest inlier set found is used as the initial segmentation for the full iterative MC-RANSAC.

2) *Iterative MC-RANSAC*: A more accurate iterative MC-RANSAC approach minimizes the reprojection error of each tracklet in image space with the cost function,

$$J_{\text{iter}}(\boldsymbol{\omega}) = \frac{1}{2} \sum_{j=1}^{M_{\text{fast}}} \mathbf{e}_{\text{iter},j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{e}_{\text{iter},j,k}, \quad (6)$$

where $\mathbf{R}_{j,k}$ is the covariance matrix for the measurements of the j^{th} tracklet and M_{fast} is the number of inliers found by the fast MC-RANSAC. The reprojection error is

$$\mathbf{e}_{\text{iter},j,k} = \mathbf{y}_{j,k} - \mathbf{s} \left(\mathbf{z} \left(\exp(\Delta t_{k,k'} \boldsymbol{\omega}^\wedge), \mathbf{p}_{k'}^{j,k'} \right) \right). \quad (7)$$

The error function is linearized by representing the velocity as a nominal value, $\bar{\boldsymbol{\omega}}$, and a small perturbation, $\delta \boldsymbol{\omega}$,

$$\boldsymbol{\omega} = \bar{\boldsymbol{\omega}} + \delta \boldsymbol{\omega}. \quad (8)$$

Substituting (7) and (8) into (6), differentiating with respect to $\delta \boldsymbol{\omega}$, and setting the result equal to zero gives the perturbation that minimizes the linearization,

$$\delta \boldsymbol{\omega}^* = \left(\sum_{j=1}^{M_{\text{fast}}} \mathbf{H}_{j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{H}_{j,k} \right)^{-1} \left(\sum_{j=1}^{M_{\text{fast}}} \mathbf{H}_{j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{e}_{\text{iter},j,k} \right),$$

where $\mathbf{H}_{j,k}$ is the Jacobian of error function in (7),

$$\mathbf{H}_{j,k} = \left. \frac{\partial \mathbf{s}}{\partial \mathbf{z}} \right|_{\bar{\mathbf{z}}} \Delta t_{k,k'} \mathbf{T}_{k,k'} \left(\mathbf{p}_{k'}^{j,k'} \right)^\odot \mathcal{T}_{k,k'}^{-1} \mathcal{J}_{k,k'},$$

where the partial derivative of the sensor model is evaluated at the nominal value, $\mathcal{T}_{k,k'}$ is the adjoint of $SE(3)$ and $\mathcal{J}_{k,k'}$ is the left Jacobian of $SE(3)$ [34].

This process is iterated until convergence to find the velocity best describing the initial inlier set found by fast MC-RANSAC,

$$\bar{\boldsymbol{\omega}} \leftarrow \bar{\boldsymbol{\omega}} + \delta \boldsymbol{\omega}^*.$$

This velocity is then used to segment all M_{clst} tracklets into inliers and outliers by comparing their reprojection error using (5) to a user-specified threshold. This is more accurate outlier rejection than using fast MC-RANSAC alone.

C. Continuous-Time Trajectory Optimization

Traditional discrete-time estimation requires at least three measurements at every estimation state. This is often achieved in event-based VO by grouping event features to common times, which reduces the temporal resolution of event cameras. Continuous-time trajectory estimation can instead operate directly on asynchronous tracklets, which may result in estimation states with less than three measurements, by incorporating a motion prior or basis function. This allows continuous-time estimation techniques to maintain the temporal resolution of the event cameras.

The inlier tracklets from MC-RANSAC are used to define the trajectory optimization problem (Fig. 2). Each unique tracklet timestamp defines a state in the estimation problem. The continuous-time trajectory is estimated from these states using Gaussian process regression with a WNOA prior [24] (Section III-C1). This iterative process uses the velocities found during MC-RANSAC as an initial condition.

The WNOA motion prior is physically founded and accounts for real-world kinematics, unlike other continuous-time parametrizations that enforce mathematical smoothness independent of its physical plausibility. It is compared quantitatively to other estimation techniques in [24], [35]. The resulting trajectory can be queried for the camera pose at any timestamp in the estimation window (Section III-C2).

1) *WNOA Estimator*: The estimated states are defined as

$$\mathbf{x} = \left\{ \mathbf{T}_{k,1}, \boldsymbol{\omega}_k, \mathbf{p}_1^{j,1} \right\}_{j=1, \dots, M, k=1, \dots, K},$$

where $\mathbf{T}_{k,1} \in SE(3)$ is the pose at the time t_k relative to the initial pose, $\boldsymbol{\omega}_k$ is the corresponding 6DOF body-centric velocity and $\mathbf{p}_1^{j,1}$ is the position of the j^{th} landmark relative to the initial pose. The camera trajectory is represented by discrete estimated states, $\mathbf{T}_{k,1}$ and $\boldsymbol{\omega}_k$, which can be denoted with a slight abuse of notation as $\mathbf{x}_k = \{\mathbf{T}_{k,1}, \boldsymbol{\omega}_k\}$. The acceleration, $\ddot{\boldsymbol{\omega}}$, is assumed to be a zero-mean, white-noise Gaussian process,

$$\ddot{\boldsymbol{\omega}} \sim \mathcal{GP}(\mathbf{0}, \mathbf{Q}_c \delta(t - t')), \quad (9)$$

where $\mathbf{Q}_c \in \mathbb{R}^{6 \times 6}$ is a diagonal power spectral density matrix, and $\delta(\cdot)$ is the Dirac delta function.

The WNOA assumption defines a locally constant velocity motion model. A local trajectory state, $\boldsymbol{\gamma}_k$, can be defined as a continuous-time function with respect to the global trajectory state,

$$\boldsymbol{\gamma}_k(\tau) = \begin{bmatrix} \boldsymbol{\xi}_k(\tau) \\ \dot{\boldsymbol{\xi}}_k(\tau) \end{bmatrix} = \begin{bmatrix} \ln(\mathbf{T}_k(\tau) \mathbf{T}_k^{-1})^\vee \\ \mathcal{J}(\ln(\mathbf{T}_k(\tau) \mathbf{T}_k^{-1})^\vee)^{-1} \boldsymbol{\omega}_k(\tau) \end{bmatrix},$$

where $\mathbf{T}_k(\tau)$ is the pose at time $t_k \leq \tau \leq t_{k+1}$, $\mathcal{J}(\cdot)^{-1}$ is the inverse left Jacobian function, $\ln(\cdot)$ is the inverse exponential map, and $(\cdot)^\vee$ is the inverse lifting operator [34].

The estimator minimizes a joint cost function,

$$J_{\text{joint}}(\mathbf{x}) = J_{\text{prior}}(\mathbf{x}) + J_{\text{meas}}(\mathbf{x}), \quad (10)$$

where $J_{\text{prior}}(\mathbf{x})$ is the motion prior cost function and $J_{\text{meas}}(\mathbf{x})$ is the measurement cost term. The prior cost function penalizes trajectory states that deviate from WNOA assumption. The prior cost function is

$$J_{\text{prior}}(\mathbf{x}) = \frac{1}{2} \sum_{k=1}^{K-1} \mathbf{e}_{\text{prior},k+1,k}^T \mathbf{Q}_k^{-1}(t_{k+1}) \mathbf{e}_{\text{prior},k+1,k},$$

where the prior error is

$$\mathbf{e}_{\text{prior},k+1,k} = \begin{bmatrix} \ln(\mathbf{T}_{k+1,1} \mathbf{T}_{k,1}^{-1})^\vee - (t_{k+1} - t_k) \boldsymbol{\omega}_k \\ \mathcal{J}(\ln(\mathbf{T}_{k+1,1} \mathbf{T}_{k,1}^{-1})^\vee)^{-1} \boldsymbol{\omega}_{k+1} - \boldsymbol{\omega}_k \end{bmatrix}, \quad (11)$$

and the prior covariance matrix at t_τ from t_k is,

$$\mathbf{Q}_k(\tau) = \begin{bmatrix} \frac{1}{3} \Delta t_{\tau,k}^3 \mathbf{Q}_c & \frac{1}{2} \Delta t_{\tau,k}^2 \mathbf{Q}_c \\ \frac{1}{2} \Delta t_{\tau,k}^2 \mathbf{Q}_c & \Delta t_{\tau,k} \mathbf{Q}_c \end{bmatrix}, \quad (12)$$

where $\Delta t_{\tau,k} = \tau - t_k$ and $\mathbf{Q}_c \in \mathbb{R}^{6 \times 6}$ is the power spectral density matrix defined in (9).

The measurement cost function that minimizes the feature tracklet reprojection error is

$$J_{\text{meas}}(\mathbf{x}) = \frac{1}{2} \sum_{j,k} \mathbf{e}_{\text{meas},j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{e}_{\text{meas},j,k},$$

where $\mathbf{R}_{j,k} \in \mathbb{R}^{3 \times 3}$ is the measurement covariance matrix of the j^{th} landmark viewed from the k^{th} state and the tracklet reprojection error is

$$\mathbf{e}_{\text{meas},j,k} = \mathbf{y}_{j,k} - \mathbf{s} \left(\mathbf{z} \left(\mathbf{T}_{k,1}, \mathbf{p}_1^{j,1} \right) \right). \quad (13)$$

The optimal camera trajectory is obtained by optimizing the joint cost function,

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \{ J_{\text{meas}}(\mathbf{x}) + J_{\text{prior}}(\mathbf{x}) \},$$

using the Gauss-Newton method. The states are approximated as operating points, $\mathbf{x}_{\text{op}} = \{\mathbf{T}_{\text{op}}, \boldsymbol{\omega}_{\text{op}}, \mathbf{p}_{\text{op}}\}$, and perturbations, $\delta \mathbf{x} = \{\delta \boldsymbol{\xi}, \delta \boldsymbol{\omega}, \delta \boldsymbol{\zeta}\}$, linearizing (10) as

$$J_{\text{joint}}(\mathbf{x}) = J_{\text{joint}}(\mathbf{x}_{\text{op}}) - \mathbf{b}^T \delta \mathbf{x} + \frac{1}{2} \delta \mathbf{x}^T \mathbf{A} \delta \mathbf{x}, \quad (14)$$

where

$$\begin{aligned} \mathbf{A} &= \sum_{j,k} \mathbf{P}_{j,k}^T \mathbf{G}_{j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{G}_{j,k} \mathbf{P}_{j,k} \\ &+ \sum_k \mathbf{P}_k^T \mathbf{E}_k^T \mathbf{Q}_k^{-1}(t_{k+1}) \mathbf{E}_k \mathbf{P}_k, \\ \mathbf{b} &= \sum_{j,k} \mathbf{P}_{j,k}^T \mathbf{G}_{j,k}^T \mathbf{R}_{j,k}^{-1} \mathbf{e}_{\text{meas},j,k} \\ &+ \sum_k \mathbf{P}_k^T \mathbf{E}_k^T \mathbf{Q}_k^{-1}(t_{k+1}) \mathbf{e}_{\text{prior},k+1,k}, \end{aligned}$$

where $\mathbf{P}_{j,k}$ and \mathbf{P}_k are matrices to pick the specific components of the total perturbation, $\mathbf{G}_{j,k}$ is the Jacobian of (13),

$$\mathbf{G}_{j,k} = \frac{\partial \mathbf{s}}{\partial \mathbf{z}} \Big|_{\mathbf{z}} \left[\left(\mathbf{T}_{\text{op},k,1} \mathbf{p}_{\text{op},1}^{j,1} \right)^\odot \quad \mathbf{0} \quad \mathbf{T}_{\text{op},k,1} \begin{bmatrix} \mathbf{1} \\ \mathbf{0}^T \end{bmatrix} \right],$$

and \mathbf{E}_k is the Jacobian of the prior error function in (11),

$$\mathbf{E}_k = \begin{bmatrix} \mathbf{E}_{11} & \Delta t_{k+1,k} \mathbf{1} & \mathbf{E}_{13} & \mathbf{0} \\ \mathbf{E}_{21} & \mathbf{1} & \mathbf{E}_{23} & \mathbf{E}_{24} \end{bmatrix},$$

where

$$\begin{aligned} \mathbf{E}_{11} &= \mathcal{J}_{k+1,k}^{-1} \boldsymbol{\tau}_{k+1,k}, & \mathbf{E}_{13} &= \mathbf{E}_{24} = -\mathcal{J}_{k+1,k}^{-1}, \\ \mathbf{E}_{21} &= \frac{1}{2} \boldsymbol{\omega}_{k+1}^\wedge \mathcal{J}_{k+1,k}^{-1} \boldsymbol{\tau}_{k+1,k}, & \mathbf{E}_{23} &= -\frac{1}{2} \boldsymbol{\omega}_{k+1}^\wedge \mathcal{J}_{k+1,k}^{-1}, \end{aligned}$$

and $(\cdot)^\wedge$ is the $\mathbb{R}^{6 \times 1}$ to $\mathbb{R}^{6 \times 6}$ operator [34].

Taking the derivative of (14) with respect to $\delta \mathbf{x}^*$, setting the result equal to zero and solving the resulting linear system, $\mathbf{A} \delta \mathbf{x}^* = \mathbf{b}$, gives the perturbation that minimizes the linearization. The estimation states are updated,

$$\begin{aligned} \mathbf{T}_{\text{op},k,1} &\leftarrow \exp(\delta \boldsymbol{\xi}^*) \mathbf{T}_{\text{op},k,1}, \\ \boldsymbol{\omega}_{\text{op},k} &\leftarrow \boldsymbol{\omega}_{\text{op},k} + \delta \boldsymbol{\omega}^*, \\ \mathbf{p}_{\text{op},j} &\leftarrow \mathbf{p}_{\text{op},j} + \delta \boldsymbol{\zeta}^*. \end{aligned}$$

and the process is iterated until convergence. The final estimate is the landmark position and the continuous-time trajectory represented as discrete state poses, discrete local velocities, and the WNOA prior.

2) *Querying the Continuous-Time Trajectory*: The continuous-time trajectory can be queried for the camera pose at any time during the estimation window, $t_1 \leq \tau \leq t_K$. The pose is interpolated between the states at the two closest times, $t_m \leq \tau \leq t_n$, using the WNOA prior. The local trajectory state, $\gamma_m(\tau)$, is interpolated [24] as

$$\gamma_m(\tau) = \Lambda(\tau)\gamma_m(t_m) + \Omega(\tau)\gamma_n(t_n),$$

where

$$\begin{aligned} \Lambda(\tau) &= \Phi(\tau, t_m) - \Omega(\tau)\Phi(t_n, t_m) \\ \Omega(\tau) &= \mathbf{Q}_m(\tau)\Phi(t_n, \tau)^T \mathbf{Q}_m(t_n)^{-1} \\ \Phi(\tau, t_m) &= \begin{bmatrix} \mathbf{1} & (\tau - t_m)\mathbf{1} \\ \mathbf{0} & \mathbf{1} \end{bmatrix}, \end{aligned}$$

and $\mathbf{Q}_m(\cdot)$ is defined in (12).

This interpolation can also be used to define the estimation states at a subset of the measurement times [24], [35].

IV. EXPERIMENTS

The presented pipeline is evaluated on the MVSEC dataset [11] and compared against the publicly available ESVO [9], a discrete event-based stereo VO pipeline. MVSEC consists of complex, nonconstant-velocity motion with stereo event camera data and 100 Hz ground truth poses for indoor scenes. The *indoor1* and *indoor3* sequences are used since ESVO provides tuning for these sequences. The performance of both algorithms is evaluated using global and relative error (Section IV-A) and the results are discussed in Section IV-B.

A sliding-window version of the system is implemented in MATLAB using LIBVISO2 [36] for feature detection and tracking. The sliding window width is five and LIBVISO2 is run on both full- and half-resolution images. Tracklets are filtered out if they

- 1) have more than 20 ms time difference between stereo features,
- 2) have less than 2px disparity,
- 3) have less than 2px length, and
- 4) last less than 40 ms in time.

Outlier rejection was done with 10000 iterations of fast MC-RANSAC followed with one call of iterative MC-RANSAC, both using an inlier threshold of 5%. The estimator uses the covariances $\mathbf{Q}_c^{-1} = 50\text{diag}(1, 1, 1, 10, 10, 10)$ and $\mathbf{R}_{j,k}^{-1} = 0.1\text{diag}(5, 5, 1)$ and terminates when the cost change between two iterations is less than 1%.

A. Evaluation Metrics

Global error quantifies the estimator accuracy relative to the initial pose. Relative error quantifies the amount of error in each estimate and is often used to calculate aggregate values over the trajectory. Both can be calculated from ground truth using a single general equation [37],

$$\text{err}(t_m, t_n) = \ln \left(\mathbf{T}_{m^{\text{GT}}, m} \mathbf{T}_{n, m} \mathbf{T}_{m^{\text{GT}}, m}^{-1} \mathbf{T}_{n^{\text{GT}}, m^{\text{GT}}}^{-1} \right)^{\vee},$$

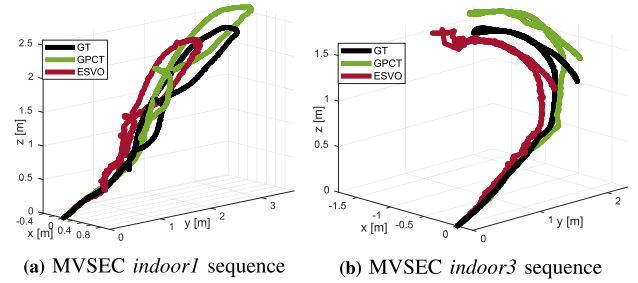


Fig. 3. Trajectory plots of the presented Gaussian process continuous-time approach (GPCT), ESVO and ground truth (GT) results in 3D space. GPCT performs better in challenging scenarios like rotation and back-and-forth motions. It also has a smoother trajectory due to the WNOA motion prior.

where $\mathbf{T}_{n, m}$ is the estimated transform to the n^{th} frame from the m^{th} frame, $\mathbf{T}_{n^{\text{GT}}, m^{\text{GT}}}$ is the ground truth transform to the n^{th} frame from the m^{th} frame, and $\mathbf{T}_{m^{\text{GT}}, m}$ is the transform to the ground truth m^{th} frame from the estimate of the m^{th} frame. The global error at a time, t_k , is then defined relative to the initial pose,

$$\text{GE}(t_k) = \text{err}(t_k, t_1),$$

and the relative error is defined relative to the previous pose,

$$\text{RE}(t_k) = \text{err}(t_k, t_{k-1}).$$

The error of the presented continuous-time system is calculated using the timestamps of ESVO. For global error, the continuous-time trajectory is aligned with ESVO's initial pose and then queried at its own state times. For relative error, the continuous-time trajectory is only queried at the discrete state times of ESVO so that both estimators have the same durations between estimates and their relative errors can be compared directly. The high frequency ground truth trajectory is linearly interpolated to the timestamps of the estimated states.

The prototype implementation of the continuous-time system does not run in real time. The system's computational performance should be improved by implementing it in a more efficient language, such as C++, and using keytimes to reduce the number of the estimation states [24], [35].

B. Results

The estimated trajectories are evaluated qualitatively relative to the ground truth in 3D (Fig. 3). They are evaluated quantitatively by calculating the global and relative errors with respect to ground truth (Fig. 4, Tables I and II). The errors are evaluated statistically using root-mean-squared (RMS), standard deviation (St. Dev.) and maximum error (Max). The maximum global error and final global error are also presented as a percentage of the integration of the norm of the relevant ground-truth component (e.g., path length). The qualitative results demonstrate the benefits of the WNOA prior, with the presented system having a smoother estimated trajectory than that of ESVO. The presented system also quantitatively has smaller RMS relative error and better or equivalent maximum global error than ESVO.

1) *MVSEC indoor1*: The presented system has a similar performance in global error to ESVO (Fig. 4(a)). The relative error is smaller than ESVO (Fig. 4(c)) with an RMS value that is two-times better (Table II). This illustrates the better local

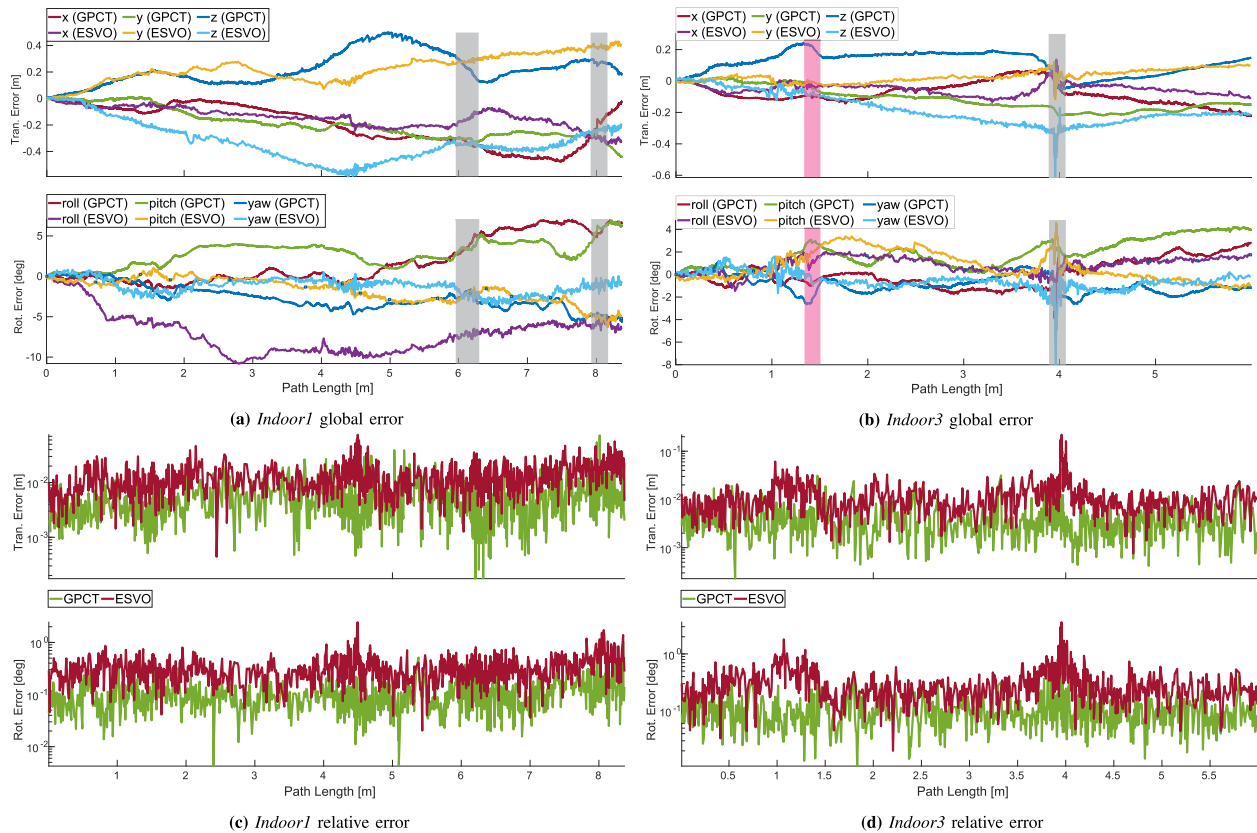


Fig. 4. Global and relative error of the presented Gaussian process continuous-time approach (GPCT) and ESVO as a function of path length on MVSEC *indoor1* and *indoor3*. The translational and rotational global error are plotted separately and the relative error is plotted as a single $SE(3)$ quantity. Note that the relative error y-axis is a logarithmic scale which suppresses spikes. The grey shaded areas represent complex camera motions that result in poor feature quality. The pink areas indicate featureless regions where the presented technique relies solely on the motion prior.

TABLE I
GLOBAL ERROR OF THE PRESENTED GAUSSIAN PROCESS CONTINUOUS-TIME APPROACH (GPCT) AND ESVO

		<i>Indoor1</i>					<i>Indoor3</i>				
		Max	Max %	Final %	RMS	St. Dev.	Max	Max %	Final %	RMS	St. Dev.
GPCT	tran.	0.635	7.58%	5.73%	0.372	0.192	0.306	5.11%	5.08%	0.218	0.059
	rota.	0.195	7.98%	7.6%	0.087	0.05	0.088	4.23%	4.13%	0.051	0.023
	$SE(3)$	0.639	7.18%	5.78%	0.382	0.196	0.319	4.95%	4.93%	0.224	0.061
ESVO	tran.	0.642	7.66%	6.61%	0.477	0.152	0.683	11.4%	4.33%	0.232	0.104
	rota.	0.189	8.87%	6.8%	0.139	0.039	0.165	8.66%	2.39%	0.038	0.02
	$SE(3)$	0.663	7.55%	6.51%	0.497	0.155	0.697	10.9%	4.12%	0.235	0.103

TABLE II
RELATIVE ERROR OF THE PRESENTED GAUSSIAN PROCESS CONTINUOUS-TIME APPROACH (GPCT) AND ESVO

		<i>Indoor1</i>			<i>Indoor3</i>		
		Max	RMS	St. Dev.	Max	RMS	St. Dev.
GPCT	tran.	$73 \cdot 10^{-3}$	$7.5 \cdot 10^{-3}$	$5.3 \cdot 10^{-3}$	$32 \cdot 10^{-3}$	$5.5 \cdot 10^{-3}$	$3.6 \cdot 10^{-3}$
	rota.	$18 \cdot 10^{-3}$	$2.4 \cdot 10^{-3}$	$1.4 \cdot 10^{-3}$	$13 \cdot 10^{-3}$	$2.3 \cdot 10^{-3}$	$1.3 \cdot 10^{-3}$
	$SE(3)$	$74 \cdot 10^{-3}$	$7.9 \cdot 10^{-3}$	$5.3 \cdot 10^{-3}$	$32 \cdot 10^{-3}$	$5.9 \cdot 10^{-3}$	$3.6 \cdot 10^{-3}$
ESVO	tran.	$75 \cdot 10^{-3}$	$17 \cdot 10^{-3}$	$9.2 \cdot 10^{-3}$	$219 \cdot 10^{-3}$	$24 \cdot 10^{-3}$	$19 \cdot 10^{-3}$
	rota.	$43 \cdot 10^{-3}$	$7.2 \cdot 10^{-3}$	$3.9 \cdot 10^{-3}$	$63 \cdot 10^{-3}$	$8.3 \cdot 10^{-3}$	$5.6 \cdot 10^{-3}$
	$SE(3)$	$86 \cdot 10^{-3}$	$18 \cdot 10^{-3}$	$9.2 \cdot 10^{-3}$	$220 \cdot 10^{-3}$	$25 \cdot 10^{-3}$	$19 \cdot 10^{-3}$

consistency of the trajectory estimate and explains the smooth trajectory plot in Fig. 3(a).

The grey shaded areas in Fig. 4(a) denote complex motions where significant error occurs. When the camera undergoes

large motions (e.g., large rotation) the observed scene changes drastically. This reduces the quality and quantity of tracklets found by the clustered feature detection and tracking and as a result the quality of the trajectory estimation. This can be improved with better feature tracking, likely specifically designed for event-based cameras.

2) *MVSEC indoor3*: The presented system estimates a smoother trajectory than ESVO in the *indoor3* segment. The estimator qualitatively describes the ground truth motion and is locally consistent. It has almost a four-times better RMS relative error than ESVO (Table II)

Challenging camera motions are marked in Fig. 4(b) as pink and grey shaded areas. The grey areas correspond to a back-and-forth camera motion and results in poor performance for both techniques. ESVO loses tracking and has to reinitialize, resulting in large global and relative error spikes (Fig. 4(a) and (d)). The

presented system performs better than ESVO in this area, but feature quality also decreases which increases both global and relative error.

The pink areas correspond to a large rotational motion. Feature tracking fails for the presented system during this motion but the WNOA prior carries the estimate through without causing significant error. This demonstrates the robustness of the presented system to feature tracking failure and the potential to improve performance with ongoing research on event-based feature tracking.

V. CONCLUSION

This paper presents a complete event-based continuous-time VO pipeline that maintains the temporal resolution of event cameras throughout the estimation. This pipeline can use either traditional frame-based or new event-based feature detectors and trackers to generate asynchronous tracklets. These tracklets are filtered for outliers using a motion-compensated RANSAC that accounts for the unique tracklet times. The pipeline estimates a continuous-time trajectory using nonparametric Gaussian process regression with a physically founded WNOA motion prior that can be queried for the camera pose at any time within the estimation window. The system's performance is evaluated on the publicly available MVSEC dataset where it achieves better performance than the publicly available ESVO pipeline, especially in terms of RMS relative error.

REFERENCES

- [1] H. P. Moravec, *Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover*. Stanford, CA, USA: Stanford Univ., 1980.
- [2] D. Nistér, O. Naroditsky, and J. Bergen, "Visual odometry," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2004.
- [3] Y. Cheng, M. Maimone, and L. Matthies, "Visual Odometry on the Mars Exploration Rovers," in *Proc. IEEE Int. Conf. Syst., Man Cybern.*, 2005, pp. 903–910.
- [4] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 3, pp. 611–625, Mar. 2018.
- [5] C. Forster, M. Pizzoli, and D. Scaramuzza, "SVO: Fast Semi-Direct Monocular Visual Odometry," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 15–22.
- [6] G. Gallego et al., "Event-based vision: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 1, pp. 154–180, Jan. 2022.
- [7] A. R. Vidal, H. Rebecq, T. Horstschäfer, and D. Scaramuzza, "Ultimate SLAM? combining events, images, and IMU for robust visual SLAM in HDR and high-speed scenarios," *IEEE Robot. Automat. Lett.*, vol. 3, no. 2, pp. 994–1001, Apr. 2018.
- [8] H. Rebecq, T. Horstschäfer, G. Gallego, and D. Scaramuzza, "EVO: A geometric approach to event-based 6-DoF parallel tracking and mapping in real time," *IEEE Robot. Automat. Lett.*, vol. 2, no. 2, pp. 593–600, Apr. 2017.
- [9] Y. Zhou, G. Gallego, and S. Shen, "Event-based stereo visual odometry," *T-RO*, vol. 37, no. 5, pp. 1433–1450, 2021.
- [10] S. Anderson and T. D. Barfoot, "RANSAC for motion-distorted 3D visual sensors," in *Proc. Int. Conf. Intell. Robots Syst.*, 2013, pp. 2093–2099.
- [11] A. Z. Zhu, D. Thakur, T. Özaslan, B. Pfommer, V. Kumar, and K. Daniilidis, "The multivehicle stereo event camera dataset: An event camera dataset for 3D perception," *IEEE Robot. Automat. Lett.*, vol. 3, no. 3, pp. 2032–2039, Jul. 2018.
- [12] D. Weikersdorfer, R. Hoffmann, and J. Conradt, "Simultaneous localization and mapping for event-based vision systems," in *Proc. Int. Conf. Comput. Vis. Syst.*, 2013, pp. 133–142.
- [13] H. Kim, A. Handa, R. Benosman, S.-H. Ieng, and A. Davison, "Simultaneous mosaicing and tracking with an event camera," in *Proc. Brit. Mach. Conf.*, 2014, doi: 10.5244/c.28.26.
- [14] B. Kueng, E. Mueggler, G. Gallego, and D. Scaramuzza, "Low-latency visual odometry using event-based feature tracks," in *Proc. IEEE Int. Conf. Intell. Robots Syst.*, 2016, pp. 16–23.
- [15] D. Weikersdorfer, D. B. Adrian, D. Cremers, and J. Conradt, "Event-based 3D SLAM with a depth-augmented dynamic vision sensor," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 359–364.
- [16] A. Zihao Zhu, N. Atanasov, and K. Daniilidis, "Event-based visual inertial odometry," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5816–5824.
- [17] H. Rebecq, T. Horstschäfer, and D. Scaramuzza, "Real-time visual-inertial odometry for event cameras using keyframe-based nonlinear optimization," in *Proc. Brit. Mach. Vis. Conf.*, 2017, doi: 10.5244/c.31.16.
- [18] P. Chen, W. Guan, and P. Lu, "ESVIO: Event-based stereo visual inertial odometry," *IEEE Robot. Automat. Lett.*, vol. 8, no. 6, pp. 3661–3668, Jun. 2022.
- [19] I. Alzugaray and M. Chli, "Asynchronous corner detection and tracking for event cameras in real time," *Robot. Automat. Lett.*, vol. 3, no. 4, pp. 3177–3184, Oct. 2018.
- [20] C. Le Gentil, F. Tschopp, I. Alzugaray, T. Vidal-Calleja, R. Siegwart, and J. Nieto, "IDOL: A framework for IMU-DVS odometry using lines," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 5863–5870.
- [21] H. Kim, S. Leutenegger, and A. J. Davison, "Real-time 3D reconstruction and 6-DoF tracking with an event camera," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 349–364.
- [22] H. Rebecq, G. Gallego, E. Mueggler, and D. Scaramuzza, "EMVS: Event-based multi-view stereo-3D reconstruction with an event camera in real-time," in *Proc. Int. J. Comput. Vis.*, 2018, pp. 1394–1414.
- [23] A. Hadviger, I. Cvišić, I. Marković, S. Vražić, and I. Petrović, "Feature-based event stereo visual odometry," in *Proc. Eur. Conf. Mobile Robots*, 2021, pp. 1–6.
- [24] S. Anderson and T. D. Barfoot, "Full STEAM ahead: Exactly sparse Gaussian process regression for batch continuous-time trajectory estimation on SE(3)," in *Proc. Int. Conf. Intell. Robots Syst.*, 2015, pp. 157–164.
- [25] P. Furgale, C. H. Tong, T. D. Barfoot, and G. Sibley, "Continuous-time batch trajectory estimation using temporal basis functions," *IJRR*, vol. 34, no. 14, pp. 2088–2095, 2015.
- [26] G. Cioffi, T. Cieslewski, and D. Scaramuzza, "Continuous-time vs. discrete-time vision-based SLAM: A comparative study," *RA-L*, vol. 7, no. 2, pp. 2399–2406, Apr. 2022.
- [27] E. Mueggler, G. Gallego, H. Rebecq, and D. Scaramuzza, "Continuous-time visual-inertial odometry for event cameras," *IEEE Trans. Robot.*, vol. 34, no. 6, pp. 1425–1440, Dec. 2018.
- [28] Y. Wang et al., "Visual odometry with an event camera using continuous ray warping and volumetric contrast maximization," *Sensors*, vol. 22, no. 15, 2022, Art. no. 5687.
- [29] K. Huang, Y. Wang, and L. Kneip, "Motion estimation of non-holonomic ground vehicles from a single feature correspondence measured over n views," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 12706–12715.
- [30] D. Liu, A. Parra, Y. Latif, B. Chen, T.-J. Chin, and I. Reid, "Asynchronous optimisation for event-based visual odometry," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 9432–9438.
- [31] R. Benosman, C. Clercq, X. Lagorce, S.-H. Ieng, and C. Bartolozzi, "Event-based visual flow," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 25, no. 2, pp. 407–417, Feb. 2013.
- [32] I. Alzugaray and M. Chli, "HASTE: Multi-hypothesis asynchronous speeded-up tracking of events," in *Proc. Brit. Mach. Vis. Conf.*, 2020, Art. no. 744.
- [33] B. Kitt, A. Geiger, and H. Lategahn, "Visual odometry based on stereo image sequences with RANSAC-based outlier rejection scheme," in *Proc. IEEE Intell. Veh. Symp.*, 2010, pp. 486–492.
- [34] T. D. Barfoot, *State Estimation for Robotics*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [35] C. H. Tong, P. Furgale, and T. D. Barfoot, "Gaussian process gauss-newton for non-parametric simultaneous localization and mapping," *IJRR*, vol. 32, no. 5, pp. 507–525, 2013.
- [36] A. Geiger, J. Ziegler, and C. Stiller, "Stereoscan: Dense 3D reconstruction in real-time," in *Proc. IEEE Intell. Veh. Symp.*, 2011, pp. 963–968.
- [37] K. M. Judd and J. D. Gammell, "Multimotion visual odometry (MVO)," 2021, *arXiv:2110.15169*.