

Occlusion-Robust MVO: Multimotion Estimation Through Occlusion Via Motion Closure

Kevin M. Judd¹ and Jonathan D. Gammell¹

Abstract—Visual motion estimation is an integral and well-studied challenge in autonomous navigation. Recent work has focused on addressing multimotion estimation, which is especially challenging in highly dynamic environments. Such environments not only comprise multiple, complex motions but also tend to exhibit significant occlusion.

Previous work in object tracking focuses on maintaining the integrity of object tracks but usually relies on specific appearance-based descriptors or constrained motion models. These approaches are very effective in specific applications but do not generalize to the full multimotion estimation problem.

This paper presents a pipeline for estimating multiple motions, including the camera egomotion, in the presence of occlusions. This approach uses an expressive motion prior to estimate the $SE(3)$ trajectory of every motion in the scene, even during temporary occlusions, and identify the reappearance of motions through *motion closure*. The performance of this occlusion-robust multimotion visual odometry (MVO) pipeline is evaluated on real-world data and the Oxford Multimotion Dataset.

I. INTRODUCTION

The ability to safely navigate through a dynamic environment is a crucial task in autonomous robotics. Visual odometry (VO) is widely used to estimate the egomotion of a camera by isolating the static parts of a scene [1], but it is more challenging to segment multiple motions in a complex, dynamic scene. Recent work has focused on extending VO approaches to this *multimotion estimation problem* [2]. These highly dynamic scenes not only pose difficult motion estimation challenges but also tend to include significant amounts of occlusion.


An occlusion can be defined as any lack of direct observations of a scene, regardless of the cause. Direct occlusions occur when an object is obscured by another or leaves the sensor’s field of view. Indirect occlusions result from sensor limitations or algorithmic failure, such as when motion blur corrupts feature matching or object detection. Consistently estimating multiple, continuous motions in the presence of both direct and indirect occlusions is necessary for autonomous navigation in complex, dynamic environments.

Multiple object tracking (MOT) approaches address occlusions using appearance-based object representations to recognize objects when they become unoccluded [3]. These approaches employ specialized detectors and simple motion models that limit their ability to track general objects and estimate the full $SE(3)$ pose of each object.

Our multimotion visual odometry (MVO) pipeline [2] simultaneously estimates the full $SE(3)$ trajectory of every motion in a scene, including the camera egomotion, without a



(a)



(b)

Fig. 1. Motion trajectories produced by the occlusion-robust MVO system for a real-world street scene before (a) and after (b) an occlusion. The camera pans to follow a cyclist (cyan) moving from left to right across the scene as they are temporarily occluded by a truck (purple) moving from right to left. The egomotion of the camera is estimated simultaneously with the trajectories of the cyclist and truck. The object trajectories are estimated directly when they are visible (solid lines), and a white-noise-on-acceleration prior is used to extrapolate the trajectory of the cyclist while it is occluded by the truck (dashed line). When the cyclist becomes unoccluded, it is detected through *motion closure* and its extrapolated trajectory is rectified through interpolation. The result is a consistent $SE(3)$ trajectory of every motion in the scene, even during temporary occlusions.

priori assumptions about object appearance. It addresses the multimotion estimation problem by applying multilabeling techniques to the traditional VO pipeline using only a rigid-motion assumption. By emphasizing motion over appearance, we contend that it is generally more important to understand *how* things are moving in the environment than to understand *what* they are.

This paper presents the occlusion-robust MVO pipeline, which extends the MVO framework to estimate multiple $SE(3)$ motions through occlusion (Fig. 1). It exploits a

¹K. M. Judd and J. D. Gammell are with the Estimation, Search, and Planning (ESP) research group at the Oxford Robotics Institute (ORI), University of Oxford, United Kingdom. {kjudd, gammell}@robots.ox.ac.uk




Fig. 2. Trajectory estimates from the occlusion-robust MVO system before (a), during (b), and after (c) an occlusion in the `occlusion_2_unconstrained` segment of the OMD [4]. The trajectory of the swinging block (4, red) is directly estimated when it is visible in (a) and (c) and is extrapolated using the constant-velocity motion prior (dashed line) when it is occluded by the block tower (1, blue) in (b). When the block becomes unoccluded in (c), it is rediscovered through *motion closure* and the estimates are interpolated to match the directly estimated trajectory.

continuous $SE(3)$ motion prior for egomotion estimation and extends it to geocentric, third-party motion estimation. This prior allows occlusion-robust MVO to both extrapolate temporarily occluded motions and recognize when they reappear (Fig. 2). This process of *motion closure* improves trajectory consistency and accuracy without relying on specific appearance-based models.

The accuracy of this approach is evaluated both quantitatively and qualitatively. The metric estimation error is measured against ground-truth data from the Oxford Multimotion Dataset (OMD) [4]. The applicability of both the rigid-motion and constant-velocity assumptions is demonstrated on a real-world scenario with a truck occluding a cyclist (i.e., a rigid motion occluding a piecewise-rigid motion).

II. BACKGROUND

Motion estimation and object tracking are integral to a wide range of computer vision applications. Tracking and estimating the motion of an individual object through a scene has been widely explored but complex, dynamic scenes make this task significantly more difficult due to frequent occlusions. Successfully addressing the multimotion estimation problem requires the ability to both track and estimate multiple motions in the presence of direct and indirect occlusions.

A. Multimotion Estimation

Many multimotion estimation approaches only solve a subset of the rigid multimotion estimation problem by applying application-specific constraints and assumptions. This limits their applicability for real-world multimotion estimation problems.

Costeira and Kanade [5] use matrix decomposition to determine the motion and shape of each dynamic object. This factorization usually requires points to be tracked for the entire estimation window, which is difficult due to direct and indirect occlusions. Some techniques allow for missing data points [6] but are not designed for many short feature tracks, as is commonly encountered in practice.

Torr [7] uses a recursive RANSAC framework to find and remove dominant motion models from the remaining feature

points. This framework is efficient at finding the dominant models in a scene, but the likelihood of sampling consistent models decreases as models are removed and the signal-to-noise ratio of the remaining points decreases. Sabzevari and Scaramuzza [8] improve this likelihood by applying geometric and kinematic constraints specific to driving scenarios, but these constraints do not generalize well to other applications.

Ozden et al. [9] consider many practical challenges in multimotion estimation, such as incomplete feature tracks. They propose a model selection framework that explicitly models the merging and splitting of motions, but it relies on separate egomotion estimation and does not address direct occlusions.

Qiu et al. [10] and Eckenhoff et al. [11] attempt to address the multimotion estimation problem using monocular visual-inertial odometry techniques. Monocular cameras cannot directly estimate depth, so an IMU must be used to disambiguate the scale of the scene. These limitations lead to several degenerate cases in multimotion estimation and these techniques do not address occlusion.

The original MVO pipeline [2] addresses the multimotion estimation problem by applying multilabeling techniques to the traditional VO pipeline. MVO simultaneously estimates the full $SE(3)$ trajectory of the camera and every third-party motion in a scene without *a priori* assumptions about object appearance. It relies on direct observations and can estimate through some partial and indirect occlusions, but it is unable to handle significant observation dropouts.

B. Multiple Object Tracking

Most visual MOT techniques follow the *tracking-by-detection* paradigm. They use a variety of specific, appearance-based object models to detect targets in each frame and focus on accurately associating those detections across multiple frames [12].

Target detectors often use bounding-box representations rather than the full target pose, so objects are usually tracked in image or Cartesian space using simple motion models [3]. Other approaches define highly specialized models that do not generalize well to other domains [13–16]. These constraints limit their ability to track general objects and estimate the full $SE(3)$ pose of each object.

Tracking and data association are often performed using recursive filters or global energy minimizations. Kalman [12, 17] and particle [18] filters use simple motion models to recursively predict the location of the target and update the current state based on current observations. The problem of assigning tracks to new detections or other tracks can be solved using the Hungarian algorithm [19] or other greedy alternatives [13, 18]. Tracking-by-detection techniques are limited by the quality of the detectors they use, and recursive methods often fail due to occlusions and lighting changes.

Energy-based techniques incorporate object appearance, motion, and interaction models in a cost functional. The functional is defined over a graph where vertices represent detections and edges represent transitions between frames, and it is minimized using flow-based techniques [20] or optimization frameworks [21]. These specialized approaches must define representative cost functionals and do not generalize well to other applications.

C. Tracking Through Occlusion

Accurate data association is more difficult in highly dynamic environments with significant occlusion. Direct occlusions can be predicted by modeling object overlaps [16] or using scene understanding [22], which can help to avoid misassociated detections. Yang et al. [23] propose a learning-based conditional random field model that considers the interdependence of observed motions, especially in the presence of occlusion. These prediction methods can be used for direct occlusions, but indirect occlusions are more difficult to predict.

Partial occlusions are challenging for appearance-based techniques because they change the observed shape of the occluded object. Feature-based techniques track targets through partial occlusions when a sufficient number of feature points can be tracked [24], but grouping features into distinct objects is difficult if their bulk motion is similar. Other techniques define specific, part-based appearance models to infer the pose of the entire object from the portions that are visible [13, 14, 25].

Full occlusions are often overcome by using motion priors to extrapolate trajectories in the absence of direct observations. Zhang et al. [20] explicitly incorporate occlusion hypotheses into a flow minimization. This *hypothesize-and-test* paradigm works well in the presence of short or partial occlusions but fails under long occlusions as there is no information available to prune hypotheses [12]. Ryoo et al. [26] avoid this impractical growth with their *observe-and-explain* paradigm, which delays hypothesizing occluded motions until an unoccluded detection is observed near the source of occlusion. Mitzel et al. [16] extrapolate unobserved target trajectories for a set number of frames to allow for reassociation when the target becomes unoccluded. These techniques are limited by the object representations used by the target detectors, so they tend to be constrained to simple motion models that do not adequately represent the motion of an object. They also focus on rediscovering an object based on appearance and position, tending to ignore its current motion.

This paper extends our previous work in multimotion estimation [2] to estimate and track motions through occlusions. It introduces a continuous motion prior that is used to extrapolate motions through both direct and indirect occlusions and allows for *motion closure* to reacquire motions and improve their occluded estimates. The full $SE(3)$ trajectory of every motion in the scene is estimated through occlusions, and the approach is evaluated on real-world data and an occlusion dataset containing ground-truth trajectories for all motions in the scene.

III. MOTION PRIORS

The way a motion is modeled directly affects the performance of multimotion estimation. Overly-simplified models can reduce estimation accuracy, but complex models are often inefficient. A good compromise is rigid motion, which reduces the motion trajectory space to $SE(3)$ while maintaining fidelity with most real-world motions. Many approaches sacrifice this fidelity for simplicity by further constraining motions to $SE(2)$ [8, 16], \mathbb{R}^3 [20, 21], or image space [26]. Others sacrifice generality by defining high-dimensional models for applications such as human tracking [13–15].

Models can also be defined discretely or continuously. Discrete models represent a trajectory as a sparse set of states, which is well-suited for synchronized sensors such as globally shuttered cameras. Continuous models smoothly represent a trajectory at all times using an assumption, or *prior*, about how objects move through the world [27]. These models incorporate this smooth prior directly into the representation, which is preferable for scanning or high-rate sensors. This prior can also be exploited to intelligently estimate occluded trajectories.

Motions can also be defined in arbitrary frames, and a simple motion in one frame may become complex when expressed in another. Two bodies, each with some known motion relative to a static reference frame, do not exhibit the same type of motion relative to *each other*. A model that is expressed egocentrically may be appropriate to estimate the egomotion of a camera relative to its static environment but not relative to other dynamic objects. It is therefore often necessary to express models in an inertial or quasi-inertial (e.g., geocentric) frame.

A. White-Noise-on-Acceleration Motion Prior

This paper extends the $SE(3)$ white-noise-on-acceleration (i.e., locally-constant-velocity) motion prior described by Anderson et al. [27] to all motions in the scene. This prior effectively penalizes the trajectory’s deviation from a constant body-centric velocity. It is physically founded because objects tend to move smoothly (i.e., with differentiable velocity) and drastic changes in direction, especially those that coincide with occlusions, are relatively rare in many autonomous navigation scenarios.

The continuous-time trajectory of the ℓ -th motion in a scene, $\mathcal{S}_\ell(t) := \{\mathbf{T}_\ell(t), \boldsymbol{\omega}_\ell(t)\}$, is defined as both the $SE(3)$ poses, $\mathbf{T}_\ell(t)$, and the local, body-centric velocities,



Fig. 3. Occlusion-aware MVO extends the original MVO pipeline to accurately estimate trajectories through occlusions. The multimotion-fitting section of the pipeline both segments a set of tracklets according to their motion and estimates the trajectories explaining that motion. Once the segmentation converges, the camera egomotion is identified and used to estimate the geocentric trajectories of all other objects in the scene. In *motion closure*, the white-noise-on-acceleration prior is used to extrapolate occluded trajectories and determine if newly discovered motions can be explained by the reappearance of an occluded object.

$\varpi_\ell(t)$. The trajectory state is assumed to vary smoothly over time in the Lie algebra, $\mathfrak{se}(3)$. This prior takes the form

$$\begin{aligned} \dot{\mathbf{T}}_\ell(t) &= \varpi_\ell(t)^\wedge \mathbf{T}_\ell(t) \\ \dot{\varpi}_\ell(t) &= \mathbf{w}'(t), \quad \mathbf{w}'(t) \sim \mathcal{GP}(\mathbf{0}, \mathbf{Q}'_c \delta(t - t')), \end{aligned} \quad (1)$$

where \mathbf{w}' is a zero-mean, white-noise Gaussian process with power spectral density matrix, $\mathbf{Q}'_c \in \mathbb{R}^{6 \times 6}$, and ϖ^\wedge is the $\mathfrak{se}(3)$ representation of $\varpi \in \mathbb{R}^6$ as defined in [28].

This continuous-time trajectory can be estimated at a collection of discrete time steps, t_1, \dots, t_K , such that,

$$\mathcal{S}_{\ell_k} := \{\mathbf{T}_{\ell_k \ell_1}, \varpi_{\ell_k}\} \equiv \mathcal{S}_\ell(t_k), \quad t_1 \leq t_k \leq t_K,$$

where $\mathbf{T}_{\ell_k \ell_1} := \mathbf{T}_\ell(t_k)$ and $\varpi_{\ell_k} := \varpi_\ell(t_k)$. These time steps correspond to observation times when measurements of the scene are collected.

The system in (1) is nonlinear and finding a numerical solution is costly. By assuming the motion between measurement times is small, the system can be recast as a set of local, linear time-invariant stochastic differential equations. Under the constant-velocity assumption with no exogenous input, these equations take the form,

$$\gamma_{\ell_k}(\tau) = \Phi(\tau, t_k) \gamma_{\ell_k}(t_k) = \begin{bmatrix} \mathbf{1} & (\tau - t_k) \mathbf{1} \\ \mathbf{0} & \mathbf{1} \end{bmatrix} \gamma_{\ell_k}(t_k), \quad (2)$$

where $\Phi(\tau, t_k)$ is the state transition function from t_k to τ and $\mathbf{1}$ is the identity matrix. The local state, γ_{ℓ_k} , is defined as

$$\gamma_{\ell_k}(t) := \begin{bmatrix} \ln(\mathbf{T}_\ell(t) \mathbf{T}_{\ell_k \ell_1}^{-1})^\vee \\ \mathcal{J}(\ln(\mathbf{T}_\ell(t) \mathbf{T}_{\ell_k \ell_1}^{-1})^\vee)^{-1} \varpi_\ell(t) \end{bmatrix},$$

where $\mathcal{J}(\cdot)$ is the left Jacobian of $SE(3)$, and $t_k \leq t \leq t_{k+1}$. Applying the prior *locally* at each time, k , casts the *global* nonlinear system as a sequence of linear, time-invariant systems.

IV. METHODOLOGY

The original MVO pipeline [2] extends VO to *multimotion* segmentation and estimation. As with traditional stereo VO pipelines, a set of tracklets $\mathcal{P} := \{p\}$ is generated by matching salient image points across rectified stereo image pairs and temporally across consecutive stereo frames. The motion segmentation and estimation are then cast as a multilabeling problem where each label, $\ell \in \mathcal{L}$, represents a motion hypothesis, \mathcal{S}_ℓ , calculated from a subset of tracklets, $\mathcal{P}_\ell \subseteq \mathcal{P}$. The labeling is found using CORAL [29], a convex optimization approach to the multilabeling problem. All motion hypotheses are initially treated as potentially belonging to the static portions of the scene (i.e., represent the camera's egomotion). Geocentric trajectories are found in a final step where a label is selected to represent the motion of the camera.

The occlusion-robust MVO pipeline (Fig. 3) extends MVO to handle both direct and indirect occlusions by employing a continuous $SE(3)$ motion prior. The prior is used both to estimate directly observed trajectories and to extrapolate occluded motions. All motion hypotheses are treated as potentially describing the egomotion until the segmentation converges (Sec. IV-A). The egomotion label is identified before performing a full-batch estimation of each trajectory in a geocentric frame (Sec. IV-A.2). The motion prior is used to extrapolate previously estimated trajectories that are not found in the current frame due to occlusion or estimation failure (Sec. IV-B). These extrapolated trajectories are then used in *motion closure* to determine if any new trajectory can be explained by the reappearance of a previously observed motion. Trajectories found to belong to the same motion are used to correct occluded estimates through interpolation (Sec. IV-C).

A. Continuous-Time Geocentric Estimation

The original MVO pipeline estimates motion trajectories without assuming which motion represents the camera [2], but this is not appropriate for the white-noise-on-acceleration prior. This locally-constant-velocity assumption is not generally valid for two noninertial frames because two bodies moving with constant velocity relative to an inertial frame do not generally have zero acceleration relative to *each other* (Sec. III). The camera egomotion is estimated from the static background but estimates of third-party motions must account for camera motion, which may result in nonconstant relative velocity. Addressing this requires identifying and estimating the camera egomotion first before using it to estimate the other trajectories in a geometric frame.

The egomotion label, C , is chosen using prior information or heuristics. It can be initialized as the largest label, as in VO, after which it can be propagated forward in time by choosing the label that maximizes the overlap in support with the previous egomotion label. Motion-based similarity metrics can also be used to identify or validate the choice at each timestep and maintain a consistent egomotion trajectory. 1) *Egomotion Estimation*: The egomotion of the camera is estimated using the approach described in [27]. The system state, \mathbf{x} , comprises the estimated pose transforms and body-centric velocities, $\{\mathcal{S}_{C_k}\}_{k=1, \dots, K}$, and the labeled landmark points, $\{\mathbf{p}_{C_1}^{j_1 C_1}\}_{j=1, \dots, |\mathcal{P}_C|}$. The estimated state, \mathbf{x} , is found by minimizing an objective function, $J(\mathbf{x}) = J_y(\mathbf{x}) + J_p(\mathbf{x})$, consisting of the measurement and prior terms.

The measurement term, $J_y(\mathbf{x})$, constrains the estimated state using the observations,

$$J_y(\mathbf{x}) := \frac{1}{2} \sum_{jk} \mathbf{e}_{y,jk}(\mathbf{x})^T \mathbf{R}_{jk}^{-1} \mathbf{e}_{y,jk}(\mathbf{x}),$$

where the error is the residual of the measurement model, $\mathbf{g}(\cdot)$, compared to the observations, \mathbf{y}_{jk} .

$$\mathbf{e}_{y,jk}(\mathbf{x}) := \mathbf{y}_{jk} - \mathbf{g}(\mathbf{x}_{jk}) = \mathbf{y}_{jk} - \mathbf{s}(\mathbf{z}(\mathbf{x}_{jk})). \quad (3)$$

and $\mathbf{x}_{jk} := \{\mathcal{S}_{C_k}, \mathbf{p}_{C_1}^{j_1 C_1}\}$ is the estimation state. The measurement model applies the perspective camera model, $\mathbf{s}(\cdot)$, to landmark points transformed by the transform model, $\mathbf{z}(\mathbf{x}_{jk}) := \mathbf{T}_{C_k C_1} \mathbf{p}_{C_1}^{j_1 C_1}$, and \mathbf{R}_{jk} is the covariance of the zero-mean additive Gaussian noise in the measurements.

The prior term, $J_p(\mathbf{x})$, constrains the current trajectory estimate by the previous velocity,

$$J_p(\mathbf{x}) := \frac{1}{2} \sum_k \mathbf{e}_{p,k}(\mathbf{x})^T \mathbf{Q}_k(t_{k+1})^{-1} \mathbf{e}_{p,k}(\mathbf{x}),$$

where the error penalizes deviation from the constant-velocity prior defined in (2),

$$\mathbf{e}_{p,k}(\mathbf{x}) = \gamma_k(t_{k+1}) - \Phi(t_{k+1}, t_k) \gamma_k(t_k), \quad (4)$$

and the covariance block matrix is defined in terms of the power spectral density, \mathbf{Q}_c , of the local white noise process,

$$\mathbf{Q}_k(t) := \begin{bmatrix} \frac{1}{3}(t-t_k)^3 \mathbf{Q}_c & \frac{1}{2}(t-t_k)^2 \mathbf{Q}_c \\ \frac{1}{2}(t-t_k)^2 \mathbf{Q}_c & (t-t_k) \mathbf{Q}_c \end{bmatrix}.$$

The total cost, $J(\mathbf{x})$, is minimized by linearizing the error about an operating point, \mathbf{x}_{op} . The operating point is perturbed according to the transform perturbations, $\{\epsilon_k \in \mathbb{R}^6\}$, velocity perturbations $\{\psi_k \in \mathbb{R}^6\}$, and landmark perturbations, $\{\zeta_j \in \mathbb{R}^3\}$, which are stacked to form the full state perturbation, $\delta \mathbf{x}$.

Linearizing the cost function requires linearizing (3) and (4). Using the Jacobians of the measurement error function, \mathbf{G}_{jk} , and the prior error function, \mathbf{E}_k , the linearized cost is given by

$$J(\mathbf{x}) \approx J(\mathbf{x}_{\text{op}}) - \mathbf{b}^T \delta \mathbf{x} + \frac{1}{2} \delta \mathbf{x}^T \mathbf{A} \delta \mathbf{x}, \quad (5)$$

where,

$$\begin{aligned} \mathbf{b} &= \sum_{jk} \mathbf{P}_{jk}^T \mathbf{G}_{jk}^T \mathbf{R}_{jk}^{-1} \mathbf{e}_{y,jk}(\mathbf{x}_{\text{op}}) + \sum_k \mathbf{P}_k^T \mathbf{E}_k^T \mathbf{Q}_k^{-1} \mathbf{e}_{p,k}(\mathbf{x}_{\text{op}}), \\ \mathbf{A} &= \sum_{jk} \mathbf{P}_{jk}^T \mathbf{G}_{jk}^T \mathbf{R}_{jk}^{-1} \mathbf{G}_{jk} \mathbf{P}_{jk} + \sum_k \mathbf{P}_k^T \mathbf{E}_k^T \mathbf{Q}_k^{-1} \mathbf{E}_k \mathbf{P}_k. \end{aligned}$$

The indicator matrices \mathbf{P}_{jk} and \mathbf{P}_k are defined such that $\delta \mathbf{x}_{jk} = \mathbf{P}_{jk} \delta \mathbf{x}$ and $\delta \mathbf{x}_k = \mathbf{P}_k \delta \mathbf{x}$, respectively.

The Jacobian of the measurement function is given by

$$\mathbf{G}_{jk} := \left. \frac{\partial \mathbf{g}}{\partial \mathbf{x}} \right|_{\mathbf{x}_{\text{op},jk}} = \left. \frac{\partial \mathbf{s}}{\partial \mathbf{z}} \right|_{\mathbf{z}(\mathbf{x}_{\text{op},jk})} \left. \frac{\partial \mathbf{z}}{\partial \mathbf{x}} \right|_{\mathbf{x}_{\text{op},jk}}, \quad (6)$$

where

$$\left. \frac{\partial \mathbf{z}}{\partial \mathbf{x}} \right|_{\mathbf{x}_{\text{op},jk}} = \left[\left(\mathbf{T}_{\text{op},C_k C_1} \mathbf{p}_{\text{op},C_1}^{j_1 C_1} \right)^\odot \quad \mathbf{0} \quad \mathbf{T}_{\text{op},C_k C_1} \mathbf{D} \right],$$

$\mathbf{D} = [\mathbf{1} \quad \mathbf{0}]^T$, and $(\cdot)^\odot$ is defined in [28].

The Jacobian of the prior error function is

$$\mathbf{E}_k = \begin{bmatrix} \mathcal{J}_{k+1,k}^{-1} \bar{\mathcal{T}}_{k+1,k} & \Delta t_k \mathbf{1} & -\mathcal{J}_{k+1,k}^{-1} & \mathbf{0} \\ \frac{1}{2} \bar{\omega}_{k+1}^\wedge \mathcal{J}_{k+1,k}^{-1} \bar{\mathcal{T}}_{k+1,k} & \mathbf{1} & -\frac{1}{2} \bar{\omega}_{k+1}^\wedge \mathcal{J}_{k+1,k}^{-1} & -\mathcal{J}_{k+1,k}^{-1} \end{bmatrix},$$

where $\mathcal{J}_{k+1,k}^{-1} := \mathcal{J} \left(\ln(\mathbf{T}_{C_{k+1} C_k})^\vee \right)^{-1}$, $\Delta t_k := t_{k+1} - t_k$, and $(\cdot)^\wedge$ and $\bar{\mathcal{T}}_{k+1,k} \in \mathbb{R}^{6 \times 6}$, the adjoint of $\mathbf{T}_{C_{k+1} C_k}$, are defined in [28].

The optimal perturbation, $\delta \mathbf{x}^*$, to minimize the linearized cost, $J(\mathbf{x})$, is the solution to $\mathbf{A} \delta \mathbf{x}^* = \mathbf{b}$. Each element of the operating point is then updated using

$$\begin{aligned} \mathbf{T}_{\text{op},C_k C_1} &\leftarrow \exp(\epsilon_k^*) \mathbf{T}_{\text{op},C_k C_1}, \\ \bar{\omega}_{C_k} &\leftarrow \bar{\omega}_{C_k} + \psi_k^*, \\ \mathbf{p}_{\text{op},C_1}^{j_1 C_1} &\leftarrow \mathbf{p}_{\text{op},C_1}^{j_1 C_1} + \mathbf{D} \zeta_j^*, \end{aligned}$$

and the cost is relinearized about the updated operating point. The process iterates until the state converges and $\mathbf{x} \leftarrow \mathbf{x}_{\text{op}}$.

2) *Third-Party Estimation*: The geocentric motions not belonging to the camera, $\{\mathcal{S}_\ell\}_{\ell \in \mathcal{L} \setminus C}$, are calculated using the estimated egomotion, \mathcal{S}_C . As in Section IV-A, each label's state, \mathbf{x} , comprises its pose and velocity estimates, $\{\mathcal{S}_{\ell_k}\}_{k=1,\dots,K}$, and its associated landmark points, $\{\mathbf{p}_{C_1}^{j_1 C_1}\}_{j=1,\dots,|\mathcal{P}_\ell|}$.

The transform model, \mathbf{z} , used by [27] is invalid for third-party motions and must be adjusted to transform egocentrically observed points through a geocentrically estimated state,

$$\mathbf{z}'(\mathbf{x}_{jk}) := \mathbf{T}_{C_k C_1} \mathbf{T}_{\ell_1 C_1}^{-1} \mathbf{T}_{\ell_k \ell_1}^{-1} \mathbf{F}_{\ell_k \ell_1} \mathbf{T}_{\ell_1 C_1} \mathbf{p}_{C_1}^{j_1 C_1},$$

where $\mathbf{F}_{\ell_k \ell_1}$ is the object deformation matrix (identity for rigid bodies), and $\mathbf{T}_{C_k C_1}$ is the camera egomotion as estimated in Section IV-A. The transform from the camera to the object centroid is given by

$$\mathbf{T}_{\ell_1 C_1} = \begin{bmatrix} \mathbf{C}_{\ell_1 C_1} & \mathbf{r}_{\ell_1}^{C_1 \ell_1} \\ \mathbf{0}^T & \mathbf{1} \end{bmatrix} = \begin{bmatrix} \mathbf{C}_{\ell_1 C_1} & -\mathbf{C}_{\ell_1 C_1} \mathbf{r}_{C_1}^{\ell_1 C_1} \\ \mathbf{0}^T & \mathbf{1} \end{bmatrix}. \quad (7)$$

The rotation, $\mathbf{C}_{\ell_1 C_1}$, is arbitrary and initially assumed to be identity for new motions. The translation, $\mathbf{r}_{C_1}^{\ell_1 C_1}$, is calculated as the centroid of the labeled points, $p^j \in \mathcal{P}_\ell$, observed in the first frame. In a sliding-window pipeline, this transform can be determined from the previously estimated trajectory estimates.

The motion model part of the measurement Jacobian is given by the block-row vector,

$$\left. \frac{\partial \mathbf{z}'}{\partial \mathbf{x}} \right|_{\mathbf{x}_{\text{op},jk}} = \left[-\mathbf{T}_{C_k C_1} \mathbf{T}_{\ell_1 C_1}^{-1} \mathbf{T}_{\text{op},\ell_k \ell_1}^{-1} \left(\mathbf{F}_{\ell_k \ell_1} \mathbf{T}_{\ell_1 C_1} \mathbf{p}_{\text{op},C_1}^{j_1 C_1} \right)^\odot \quad \mathbf{0} \quad \mathbf{T}_{C_k C_1} \mathbf{T}_{\ell_1 C_1}^{-1} \mathbf{T}_{\text{op},\ell_k \ell_1}^{-1} \mathbf{F}_{\ell_k \ell_1} \mathbf{T}_{\ell_1 C_1} \mathbf{D} \right]$$

This Jacobian is used to estimate \mathbf{G}_{jk} in (6), and (5) is used to estimate the continuous-time geocentric trajectory, \mathcal{S}_ℓ , of every third-party motion in the scene.

B. Trajectory Extrapolation

Motion priors are used to extrapolate motions in the presence of occlusions. The local state, γ_k , at time t_k can be used in (2) to estimate the extrapolated state, $\hat{\gamma}_k$, at time τ [27]. The extrapolated state is then transformed to the global state, consisting of the extrapolated transform, $\mathbf{T}_{\hat{\ell}}$, and velocity, $\bar{\omega}_{\hat{\ell}}$, via

$$\begin{aligned} \mathbf{T}_{\hat{\ell}}(\tau) &= \exp([\mathbf{1} \quad \mathbf{0}] \hat{\gamma}_k(\tau)) \mathbf{T}_\ell(t_k), \\ \bar{\omega}_{\hat{\ell}}(\tau) &= [\mathbf{0} \quad \mathbf{1}] \hat{\gamma}_k(\tau). \end{aligned}$$

Estimates can be extrapolated forward or backward in time. Their accuracy depends on the fidelity of the motion prior to



Fig. 4. Motion segmentation produced by occlusion-robust MVO. The egomotion of the camera is estimated from the static points in the scene shown in black. The motions of the swinging block (4, red) and the block tower (1, blue) are segmented and estimated simultaneously with the egomotion.

the true object motions, i.e., if their velocity is near-constant. As the length of the extrapolation grows, the estimates will drift from the true motion of the object, especially if it exhibits significant changes in velocity. If the extrapolation error is large, the trajectory will be treated as a newly discovered motion when it becomes unoccluded.

C. Motion Closure and Interpolation

The extrapolated estimates of previously seen motions can be used to identify unoccluded motions through *motion closure*. If a newly discovered motion is sufficiently similar to an occluded motion, then it is taken to be the reappearance of that previous motion.

Each newly discovered trajectory, $\mathcal{S}_{\ell'_k} := \{\mathbf{T}_{\ell'_k \ell'_1}, \boldsymbol{\omega}_{\ell'_k}\}$, is compared to an occluded motion's extrapolated trajectory at time t_k using both position and velocity [30]. If this motion-based similarity is less than a user-defined threshold, the motions are closed; otherwise, the new trajectory is considered a new motion. The corrected trajectory, $\mathcal{S}_{\ell_k} := \{\mathbf{T}_{\ell_k \ell_1}, \boldsymbol{\omega}_{\ell_k}\}$, is then estimated from the extrapolated trajectory, $\mathcal{S}_{\ell_k} := \{\mathbf{T}_{\hat{\ell}_k \hat{\ell}_1}, \boldsymbol{\omega}_{\ell_k}\}$, and the closure transform, $\mathbf{T}_{\ell_k \hat{\ell}_k}$,

$$\mathbf{T}_{\ell_k \ell_1} = \mathbf{T}_{\ell_k \hat{\ell}_k} \mathbf{T}_{\hat{\ell}_k \hat{\ell}_1} \mathbf{T}_{\hat{\ell}_1 \ell_1}, \quad \boldsymbol{\omega}_{\ell_k} = \boldsymbol{\omega}_{\ell'_k},$$

where $\mathbf{T}_{\hat{\ell}_1 \ell_1}$ is identity because the corrected and extrapolated trajectories are equivalent before the occlusion. The corrected velocity is taken directly from the newly discovered velocity.

The closure transform,

$$\mathbf{T}_{\ell_k \hat{\ell}_k} = \begin{bmatrix} \mathbf{C}_{\ell_k \hat{\ell}_k} & \mathbf{r}_{\ell_k \hat{\ell}_k} \\ \mathbf{0}^T & 1 \end{bmatrix},$$

shifts the extrapolated trajectory position, $\mathbf{r}_{\hat{\ell}_k C_k}^{\ell_k}$, to the observed centroid of the newly discovered trajectory, $\mathbf{r}_{C_k}^{\ell'_k}$.

The closure rotation, $\mathbf{C}_{\ell_k \hat{\ell}_k}$, and the closure translation, $\mathbf{r}_{\ell_k \hat{\ell}_k}^{\ell_k}$, are calculated independently. The closure rotation is taken as identity because it is difficult to determine the




Fig. 5. The translational and rotational errors in the estimated motion of the camera (a), the block tower (b), and the swinging block (c) for a 300-frame section of `occlusion2 unconstrained` segment of the OMD. The occlusion-robust MVO errors are displayed in red-green-blue, and the original MVO errors are displayed in magenta-yellow-cyan. Each object is compared to ground-truth trajectory data and errors are reported in an arbitrary geocentric frame with the z-axis up and arbitrary x- and y-axes. Grey regions show when the swinging block was occluded by the tower, or when the tower was stationary and effectively part of the background. Dashed lines represent the extrapolation error and the solid lines represent the error in the direct or interpolated estimates. As expected, the original MVO pipeline failed to track the blocks past their first occlusions. This also led to significant errors in the egomotion estimate that have been cropped in (a).

change in rotation of the object after the occlusion without using appearance-based metrics. This carries forward the extrapolation of the original motion's rotation.

The closure translation is given as

$$\begin{aligned} \mathbf{r}_{\ell_k}^{\hat{\ell}_k} &= \mathbf{C}_{\ell_k C_k} \mathbf{r}_{C_k}^{\hat{\ell}_k} = \mathbf{C}_{\ell_k C_k} \mathbf{r}_{C_k}^{\hat{\ell}_k \ell'_k} \\ &= \mathbf{C}_{\ell_k \hat{\ell}_k} \mathbf{C}_{\hat{\ell}_k C_k} \left(\mathbf{r}_{C_k}^{\hat{\ell}_k C_k} - \mathbf{r}_{C_k}^{\ell'_k C_k} \right). \end{aligned}$$

The extrapolated camera-object rotation, $\mathbf{C}_{\hat{\ell}_k C_k}$, is taken from

$$\mathbf{T}_{C_k \hat{\ell}_k} = \begin{bmatrix} \mathbf{C}_{\hat{\ell}_k C_k}^T & \mathbf{r}_{C_k}^{\hat{\ell}_k C_k} \\ \mathbf{0}^T & 1 \end{bmatrix} = \mathbf{T}_{C_k C_1} \mathbf{T}_{\ell_1 C_1}^{-1} \mathbf{T}_{\hat{\ell}_1 \ell_1}^{-1} \mathbf{T}_{\hat{\ell}_k \hat{\ell}_1}^{-1}.$$

The corrected pose and velocity, \mathcal{S}_{ℓ_k} , are then used to interpolate from the beginning of the occlusion at time

t_{j+1} and correct the extrapolated estimates. Given the state transition function defined in (2), the interpolation between t_j and t_k is defined as

$$\hat{\gamma}_j(\tau) = \Lambda(\tau) \gamma_j(t_j) + \Omega(\tau) \gamma_j(t_k),$$

where

$$\begin{aligned} \Lambda(\tau) &= \Phi(\tau, t_j) - \Omega(\tau) \Phi(t_k, t_j), \\ \Omega(\tau) &= \mathbf{Q}_j(\tau) \Phi(t_k, \tau)^T \mathbf{Q}_j(t_k)^{-1}, \end{aligned}$$

and $t_j < \tau < t_k$ [27]. This interpolation can explain the occluded motion of the object better than extrapolation because it includes direct estimates on both sides of the occlusion.

V. EXPERIMENTS AND RESULTS

The performance of occlusion-robust MVO is evaluated both quantitatively and qualitatively. The OMD [4] was used to metrically evaluate the estimation accuracy of the pipeline using ground-truth trajectory data (Figs. 4 and 5). Although the OMD consists of controlled indoor experiments, its ground-truth trajectory data for *every* motion in the scene makes it uniquely appropriate for quantitatively evaluating the accuracy of MVO. The pipeline is also qualitatively evaluated on complex, piece-wise rigid motions in a real-world scenario for which there is no ground truth (Fig. 1).

The results were produced from sequences of stereo camera data using a 16-frame sliding window. Tracklets are generated using LIBVISO2 [31] and the minimization was performed analytically with Ceres [32]. The quantitative results were run on a 300-frame segment from the `occlusion_2.unconstrained` segment of the OMD. The motion closure similarity was calculated using the Euclidean distance between the extrapolated and observed transforms in the $\mathfrak{se}(3)$ tangent space [30] and the threshold was set to 1. The transforms between the ground-truth Vicor frames and our estimated frames are arbitrary, so the first 15 poses are used to calibrate these transforms [33]. All errors are reported for geocentric trajectory estimates.

The occlusion-robust MVO egomotion estimates (Fig. 5a) exhibited a maximum total drift of 0.26 m (over a 7.48 m path). This error (3.48%) is reasonable compared to the drift in other similar, camera-only VO systems [34]. The centroid estimates used in motion closure occasionally shifted significantly and caused the interpolated error of the block tower and the swinging block to be worse than the extrapolated error. The block tower exhibited a maximum total drift of 0.66 m (over a 3.89 m path), and the swinging block exhibited a maximum total drift of 1.58 m (over a 14.12 m path).

As expected, the original MVO pipeline failed to track the motions of blocks through the occlusions. The egomotion estimation error was also higher (1.10 m) because the block tower may be mistakenly labeled as the static background when it begins to move again. In contrast, the occlusion-robust MVO pipeline both tracks the blocks through the occlusions and estimates the egomotion more accurately because motion closure helps disambiguate between the static background and motions starting from rest.

The qualitative results were run on data collected from a real-world street scenario for which ground-truth trajectory data is prohibitively difficult to collect (Fig. 1). The camera pans to follow a cyclist moving from left to right across the scene as they are temporarily occluded by a truck moving from right to left. Both objects are consistently segmented and estimated, along with the camera egomotion, despite the cyclist exhibiting some nonrigid motion. The constant-velocity assumption proved reasonable in this scenario as the cyclist's motion was accurately extrapolated through the occlusion and motion closure was successful.

VI. DISCUSSION

The occlusion-robust pipeline consistently segments the motions of the camera, blocks, cyclist, and truck while also estimating the trajectories through occlusions; however, it is still a sparse, feature-based technique. The observable shape (and centroid) of an object changes as it moves, affecting the geocentric estimate of its trajectory. As the object becomes more occluded, the quality of the trajectory estimation will degrade. It is unlikely that an object will become occluded or unoccluded instantaneously, but this can be mitigated by predicting occlusions as in [16].

The reliance of the pipeline on motion means that objects that temporarily have the same motion will be given the same label, such as when a dynamic object becomes stationary. This is often desirable, as it implicitly handles trajectory merging, but many applications might require a form of *motion permanence*. This could be introduced by using appearance-based techniques or by explicitly modeling trajectory merging.

The accuracy of the object centroid, and therefore the trajectory, depends on the distribution of observed features. The newly calculated centroid of a motion after an occlusion does not always match the original centroid, especially if it is still partially occluded. This discrepancy can cause jumps in the trajectory and is a major source of error in the estimates of the block tower (Fig. 5b) because it is often partially outside the view of the camera. The original MVO pipeline partially mitigated this through a rolling-average calculation of the centroid, but this is difficult in geocentric estimation as the centroid is required for the calculation of third-party trajectories. More robust centroid calculation remains an area of ongoing work, including considering part-to-whole extrapolation techniques [13, 14, 25].

The qualitative results (Fig. 1) illustrate that occlusion-robust MVO is readily applicable to real-world scenarios. The pipeline is able to segment and estimate the motions of both the cyclist and the truck simultaneously with the egomotion of the camera. When the cyclist is occluded, their trajectory is extrapolated using the constant-velocity prior until they are unoccluded and motion closure is successful.

Piecewise-rigid motions consist of a number of small rigid motions. If the motions are largely consistent, MVO groups them into a bulk motion; otherwise, they are rejected as noise. Though the cyclist in Fig. 1 exhibits some piece-wise rigid or nonrigid motion, the overall motion is largely rigid and the robustness of MVO means the trajectory is reliably estimated.

The applicability of the white-noise-on-acceleration prior is limited in scenes where objects change direction or speed. Few objects move with truly constant velocity, but these results illustrate that it is robust to small changes in velocity and allows velocities to evolve over time. Future work will consider a white-noise-on-jerk prior [35] to better model smoothly varying accelerations, and more expressive priors can be used to address complex dynamic motions, such as when the swinging block changed direction at the end of the segment in Fig. 5c.

VII. CONCLUSION

This paper extends the MVO pipeline to complex, dynamic environments with significant occlusions. Occlusion-robust MVO uses an $SE(3)$ motion prior to metrically estimate both camera egomotion and third-party motions even in the presence of temporary occlusions. It uses motion closure to recognize the reappearance of previously observed motions and improve estimates of their unobserved trajectories. Its performance is evaluated on a challenging segment from the Oxford Multimotion Dataset (OMD) [4] exhibiting significant occlusion and complex $SE(3)$ motions, as well as a real-world scenario with piecewise-rigid motion. We are currently exploring extensions to the pipeline as discussed in Section VI, as well as its applicability to other sensors, such as RGB-D and event cameras.

REFERENCES

- [1] H. P. Moravec, "Obstacle avoidance and navigation in the real world by a seeing robot rover," Ph.D. dissertation, Stanford, CA, USA, 1980.
- [2] K. M. Judd, J. D. Gammell, and P. Newman, "Multimotion visual odometry (MVO): Simultaneous estimation of camera and third-party motions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3949–3956.
- [3] A. Yilmaz, O. Javed, and M. Shah, "Object tracking: A survey," *ACM Computing Surveys*, vol. 38, no. 4, p. 13, 2006.
- [4] K. M. Judd and J. D. Gammell, "The Oxford multimotion dataset: Multiple $SE(3)$ motions with ground truth," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 800–807, 2019.
- [5] J. P. Costeira and T. Kanade, "A multibody factorization method for independently moving objects," *International Journal of Computer Vision*, vol. 29, no. 3, pp. 159–179, 1998.
- [6] R. Vidal and R. Hartley, "Motion segmentation with missing data using power factorization and GPCA," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2004, pp. 310–316.
- [7] P. H. Torr, "Geometric motion segmentation and model selection," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 356, no. 1740, pp. 1321–1340, 1998.
- [8] R. Sabzevari and D. Scaramuzza, "Multi-body motion estimation from monocular vehicle-mounted cameras," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 638–651, 2016.
- [9] K. E. Ozden, K. Schindler, and L. V. Gool, "Multibody structure-from-motion in practice," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, no. 6, pp. 1134–1141, 2010.
- [10] K. Qiu, T. Qin, W. Gao, and S. Shen, "Tracking 3-d motion of dynamic objects using monocular visual-inertial sensing," *IEEE Transactions on Robotics (T-RO)*, vol. 35, no. 4, pp. 799–816, 8 2019.
- [11] K. Eickenhoff, Y. Yang, P. Geneva, and G. Huang, "Tightly-coupled visual-inertial localization and 3-d rigid-body target tracking," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1541–1548, 4 2019.
- [12] D. Reid, "An algorithm for tracking multiple targets," *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [13] B. Wu and R. Nevatia, "Detection and tracking of multiple, partially occluded humans by Bayesian combination of edgelet based part detectors," *International Journal of Computer Vision*, vol. 75, no. 2, pp. 247–266, 2007.
- [14] G. Shu, A. Dehghan, O. Oreifej, E. Hand, and M. Shah, "Part-based multiple-person tracking with partial occlusion handling," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2012, pp. 1815–1821.
- [15] B. Yang and R. Nevatia, "Online learned discriminative part-based appearance models for multi-human tracking," in *European Conference on Computer Vision*. Springer, 2012, pp. 484–498.
- [16] D. Mitzel, E. Horbert, A. Ess, and B. Leibe, "Multi-person tracking with sparse detection and continuous segmentation," in *European Conference on Computer Vision*. Springer, 2010, pp. 397–410.
- [17] Z. Khan, T. Balch, and F. Dellaert, "An MCMC-based particle filter for tracking multiple interacting targets," in *European Conference on Computer Vision*. Springer, 2004, pp. 279–290.
- [18] M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. V. Gool, "Robust tracking-by-detection using a detector confidence particle filter," in *IEEE International Conference on Computer Vision*, 2009, pp. 1515–1522.
- [19] J. Xing, H. Ai, and S. Lao, "Multi-object tracking through occlusions by local tracklets filtering and global tracklets association with detection responses," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009, pp. 1200–1207.
- [20] L. Zhang, Y. Li, and R. Nevatia, "Global data association for multi-object tracking using network flows," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.
- [21] M. Byeon, H. Yoo, K. Kim, S. Oh, and J. Y. Choi, "Unified optimization framework for localization and tracking of multiple targets with multiple cameras," *Computer Vision and Image Understanding*, vol. 166, pp. 51 – 65, 2018.
- [22] R. Kaucic, A. G. Amitha Perera, G. Brooksby, J. Kaufhold, and A. Hoogs, "A unified framework for tracking through occlusions and across sensor gaps," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2005, pp. 990–997 vol. 1.
- [23] B. Yang, C. Huang, and R. Nevatia, "Learning affinities and dependencies for multi-target tracking using a CRF model," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2011, pp. 1233–1240.
- [24] D. Sugimura, K. M. Kitani, T. Okabe, Y. Sato, and A. Sugimoto, "Using individuality to track individuals: Clustering individual trajectories in crowds using local appearance and frequency trait," in *IEEE International Conference on Computer Vision*. 2009, pp. 1467–1474.
- [25] W. Hu, X. Li, W. Luo, X. Zhang, S. Maybank, and Z. Zhang, "Single and multiple object tracking using log-Euclidean Riemannian subspace and block-division appearance model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 12, pp. 2420–2440, 2012.
- [26] M. S. Ryoo and J. K. Aggarwal, "Observe-and-explain: A new approach for multiple hypotheses tracking of humans and objects," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 2008, pp. 1–8.
- [27] S. Anderson and T. D. Barfoot, "Full STEAM ahead: Exactly sparse Gaussian process regression for batch continuous-time trajectory estimation on $SE(3)$," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2015, pp. 157–164.
- [28] T. D. Barfoot, *State Estimation for Robotics*. Cambridge University Press, 2017.
- [29] P. Amayo, P. Piniés, L. M. Paz, and P. Newman, "Geometric multi-model fitting with a convex relaxation algorithm," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [30] K. M. Judd, "Unifying Motion Segmentation, Estimation, and Tracking for Complex Dynamic Scenes," Ph.D. dissertation, Oxford, UK, 2019.
- [31] A. Geiger, J. Ziegler, and C. Stiller, "StereoScan: dense 3D reconstruction in real-time," in *IEEE Intelligent Vehicles Symposium*, 2011.
- [32] S. Agarwal, K. Mierle, and Others, "Ceres Solver."
- [33] Z. Zhang and D. Scaramuzza, "A tutorial on quantitative trajectory evaluation for visual(-inertial) odometry," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. 2018, pp. 7244–7251.
- [34] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3354–3361.
- [35] T. Y. Tang, D. J. Yoon, and T. D. Barfoot, "A white-noise-on-jerk motion prior for continuous-time trajectory estimation on $SE(3)$," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 594–601, 2019.