Joint Pose and Shape Estimation of 3D Extended Objects Using Recursive Tangent-Space Bayesian Filtering on SO(3)

Jiachen Zhou, Harald Kruggel-Emden, and Uwe D. Hanebeck

***Abstract*— In this paper, we propose a recursive estimation framework for simultaneous pose and shape reconstruction of three-dimensional extended objects from point cloud data. While shape and other kinematic parameters, e.g., position, naturally reside in a Euclidean vector space, rigid-body rotation evolves on the nonlinear manifold, the special orthogonal group $\mathrm{SO}(3)$. Fusing both through a joint density leads to significant theoretical and practical challenges. To address this, we reformulate the rotation estimation task in the Euclidean tangent space of $\mathrm{SO}(3)$. Rather than directly filtering on the manifold, we track a rotation perturbation in tangent space, represented by a three-dimensional rotation vector, relative to a time-varying reference orientation. This enables the use of novel nonlinear Bayesian filtering techniques. After updating the stochastic rotation vector, a reset operation is performed to integrate the estimated rotation perturbation into the reference orientation. This procedure adjusts the deviation statistics while rigorously following the group operations on rotation manifolds.**

***Index Terms*— Bayesian inference, extended object tracking, orientation estimation, Gaussian assumed density filter**

I. INTRODUCTION

Context: We consider the problem of three-dimensional extended object tracking (EOT), where the main objective is to simultaneously estimate the target's kinematic state and its spatial extent (or shape) recursively over time, using noisy position measurements obtained from the target's surface. While EOT is well established in robotics and autonomous navigation, it also plays a pivotal role in advancing particle measurement techniques. Precise reconstruction of particulate solids' shapes is essential, as their morphology directly determines their mechanical, thermal, and chemical behavior in various industrial processing stages. When integrated with advanced tracking algorithms capable of estimating each particle's time-varying three-dimensional position and orientation, this approach can significantly enhance measurement accuracy across a broad spectrum of particulate applications.

Recent advances in sensor resolution and multi-modal integration have dramatically increased the quantity and quality of tracking data, which benefits EOT, where targets are modeled by their spatial extent rather than as points. These rich measurements enable a more accurate joint estimation of both the object's pose and its shape via probabilistic nonlinear filtering. In such a framework, the system state typically includes the target's position, orientation, shape parameters, and optionally velocities or accelerations. Since most of these components are usually not directly observable, they must be inferred from available measurements. A central objective of EOT is therefore to effectively interpret and process these data by integrating reasonable measurement models with advanced nonlinear filtering techniques.

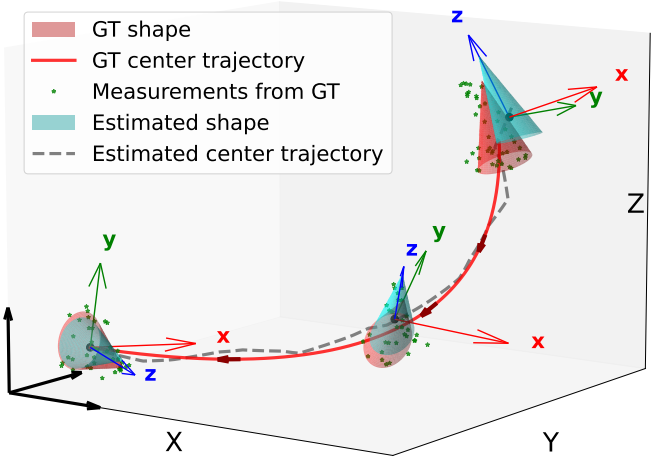

Fig. 1: Illustrative example of three-dimensional EOT, where the target's pose and shape are recursively estimated from noisy surface measurements (green dots). The true cone (red) is shown at three representative time steps with its center trajectory (solid red). The estimated shape appears in cyan, and the dashed gray line traces the inferred center trajectory. Local body-fixed axes at each estimated pose visualize the recovered orientations in the inertial coordinate system.

State-of-the-art: Relatively few studies have directly addressed the challenge of EOT in three-dimensional environments, despite the widespread availability of point cloud data from depth cameras and LiDAR sensors [1], [2], [3]. Targets of interest typically undergo full spatial rigid-body motion, exhibiting six degrees of freedom (DoF). This motion comprises both translation and rotation that must be explicitly modeled for joint pose–shape estimation [4], [5]. While translation evolves in Euclidean space and can be handled within standard filtering frameworks, the rotational component lies on the special orthogonal group $\mathrm{SO}(3)$ and is particularly challenging due to its non-Euclidean geometry. This geometry prevents the straightforward use of Euclidean-state nonlinear estimators [6].

An extended object's three-dimensional orientation can be represented in several ways, including rotation matrices, unit

Jiachen Zhou and Uwe D. Hanebeck are with the Intelligent Sensor-Actuator-Systems Laboratory (ISAS), Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology (KIT), Germany (e-mail: jiachen.zhou@kit.edu; uwe.hanebeck@kit.edu).

Harald Kruggel-Emden is with the Chair of Mechanical Process Engineering and Solids Processing, Institute of Chemical and Process Engineering, Technische Universität Berlin (TUB), Germany (e-mail: kruggel-emden@tu-berlin.de).

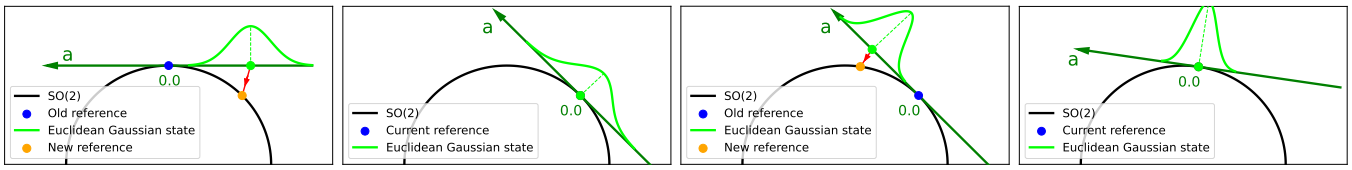

(a) Prediction step (b) Post-update orientation reset (c) Measurement update (d) Post-update orientation reset

Fig. 2: Example illustration of the proposed recursive tangent-space filtering for planar rotation perturbations on $\mathrm{SO}(2)$. The deterministic reference orientation in (a) appears as a blue point on the rotation manifold $\mathrm{SO}(2)$ (the unit circle $\mathbb{S}^1$, shown in black). After each prediction step, the predicted probability density function (PDF) (lime) of the state $\boldsymbol{a}$, used to parameterize the rotation deviation, is calculated in the Euclidean tangent space (dark green) at this reference. The mean of that distribution is then transferred back onto $\mathrm{SO}(2)$ via the exponential map (red), yielding the updated reference orientation (orange). This step equivalently resets the tangent-space mean to zero at the new reference in (b). Subsequent measurement updates in (c) follow the same pattern, recentering the tangent-space distribution around zero in (d). Crucially, we also explicitly parallel-transport the state's uncertainty between successive tangent spaces after each processing step, a procedure that is often overlooked or mishandled but essential to preserve uncertainty consistency when dealing with manifolds.

quaternions, or Euler angles. Each representation exhibits inherent trade-offs between minimality, singularity, and redundancy. Euler angles provide an intuitive and compact representation, but suffer from singularities, such as gimbal lock, which can lead to loss of one rotational DoF. In EOT, Euler angles remain useful when motion is kinematically constrained. For instance, [3] models a sailing boat's orientation using only a single yaw angle. In [2], a cuboid target under planar motion is represented by yaw and roll, enabling recursive estimation of both shape and orientation. Although reducing the number of angles mitigates gimbal-lock risks, such parameterizations cannot represent arbitrary full spatial rotations. From an estimation standpoint, probabilistic filtering must provide orientation with uncertainty. However, placing a Gaussian directly on Euclidean angle coordinates ignores both the non-Euclidean geometry of $\mathrm{SO}(3)$ and the periodic nature of Euler angles. This Euclidean–Gaussian assumption, made explicitly in prior EOT studies, e.g., [3], [2], is a further limitation.

While non-singular parameterizations avoid discontinuities, they introduce redundancy and require constraints. A full 3×3 rotation matrix uses nine parameters to represent a 3-DoF rotation, i.e., it introduces six redundant parameters that must be constrained by explicitly enforcing orthonormality. Such constraints can cause numerical instabilities and increase computational cost [7], [8]. In contrast, unit quaternions incur only a single redundant DoF and are free of singularities [9], [10]. Recursive filtering of unit quaternions must, however, respect their nonlinear manifold geometry, namely the unit 3-sphere $\mathbb{S}^3$ embedded in $\mathbb{R}^4$ [11], [12].

Extending EOT to fully three-dimensional settings is considerably more challenging. Orientation must be estimated on the curved hypersphere $\mathbb{S}^3$, while shape and other kinematic parameters are simultaneously inferred in other multidimensional Euclidean space. To avoid the challenge of defining a joint density across fundamentally different manifolds, a hierarchical orientation representation is employed in [4], [5], which is closely related to our approach. A deterministic unit quaternion provides a globally non-singular reference orientation, while the rotation deviation is parameterized by a three-dimensional Gibbs vector relative to this reference. Both the Gibbs vector and the object's shape parameters are then estimated jointly in a common Euclidean space, enabling standard Bayesian filters for vector-valued states. The reference orientation gets updated after each processing step. However, the non-Euclidean geometry of the $\mathrm{SO}(3)$ rotation manifold is overlooked in the aforementioned approaches. As a result, the uncertainty associated with this stochastic Gibbs vector fails to be properly parallel-transported alongside the reference orientation update.

Contribution: In this article, we develop a unified framework for tracking three-dimensional targets while simultaneously inferring their geometries from noisy point cloud data. Our novelty lies in the reformulation of the spatial orientation estimation problem by embedding successive orientation updates within the Euclidean tangent space of the $\mathrm{SO}(3)$ rotation manifold. Instead of filtering directly on this curved nonlinear manifold, our method tracks a compact, three-dimensional rotation vector, also known as the axis–angle vector, in the Euclidean tangent space about a deterministic time-varying reference orientation. The reference orientation is encoded as a globally non-singular 3×3 rotation matrix and is incrementally updated by composing it with the expected value of the rotation perturbation parameterized by this stochastic rotation vector.

In the context of three-dimensional EOT, this approach bridges the gap between manifold-valued spatial orientation and Euclidean shape and kinematic states. We achieve this by jointly estimating the rotation vector together with the target's shape and kinematic parameters in a single Euclidean state vector. Consequently, no intractable joint density over mixed manifolds is required. As a result, various nonlinear Bayesian filtering algorithms (usually defined in Euclidean space) can be directly employed.

After each update of the stochastic rotation vector, we apply a post-update reset that absorbs the estimated rotation perturbation into the reference orientation. Concretely, the mean of the tangent-space distribution is transferred back

onto the $\mathrm{SO}(3)$ manifold via the exponential map, yielding a new reference. Simultaneously, these perturbation statistics are re-centered to zero in the updated tangent space.

Moreover, propagating the full uncertainty of the tangent-space random vector across successive tangent spaces along a manifold is frequently overlooked or mishandled. In this paper, we exploit the intrinsic geometry of $\mathrm{SO}(3)$ to derive closed-form expressions for the parallel-transport operator that propagates rotation-vector covariances along manifold geodesics. By explicitly incorporating these manifold-aware operations into our recursive filter, we ensure that uncertainty remains correctly aligned with the local linearization at each reference orientation. For enhanced clarity and intuitive understanding, a detailed two-dimensional example on the $\mathrm{SO}(2)$ manifold is provided in Fig. 2 to illustrate the core concepts of our $\mathrm{SO}(3)$ filtering framework.

Furthermore, our full spatial orientation estimator can be seamlessly integrated as a higher-quality plug-in replacement into any three-dimensional EOT framework, regardless of whether the measurement model employs a parametric or non-parametric shape representation. We evaluate the proposed EOT framework in a simulated free-fall of an elliptic cone undergoing unconstrained rigid-body motion. The object is initialized with a nonzero angular velocity to enable unrestricted three-dimensional rotation. This setup provides a realistic benchmark for particle measurement research, facilitating rigorous evaluation of full pose-shape reconstruction. The results demonstrate superior performance over state-of-the-art methods, with our approach accurately recovering both the target's spatial extent and its time-varying orientation in the inertial frame.

II. Problem Formulation

We aim to estimate the parameters of an unknown three-dimensional extended object, in particular its shape, position, and orientation, based on noisy position measurements sequentially collected from its surface. As outlined above, all recursively estimated variables are embedded in a single Euclidean space. At discrete time step k, the Euclidean system state vector $\underline{\boldsymbol{x}}_k$, which encapsulates all necessary kinematic and extent parameters, is denoted by

$$\underline{\boldsymbol{x}}_k = \left[(\underline{\boldsymbol{x}}_k^{\text{kin.}})^\top,\ (\underline{\boldsymbol{x}}_k^{\text{ext.}})^\top \right]^\top . \tag{1}$$

The kinematic component $\underline{\boldsymbol{x}}_k^{\text{kin.}}$ comprises the position of the target's center of mass in the inertial frame ${}^I\underline{\boldsymbol{c}}_k$ (which also defines the body-fixed frame origin), velocity ${}^I\underline{\boldsymbol{\varphi}}_k$, a three-dimensional rotation vector $\underline{\boldsymbol{\delta}}_{B,\,k}$, and the angular velocity of the body with respect to the inertial frame ${}^I\underline{\boldsymbol{\omega}}_{B,\,k}$. The rotation vector $\underline{\boldsymbol{\delta}}_{B,k}$ is defined about the current deterministic reference orientation ${}^I\mathbf{R}^{\text{ref.}}_{B,\,k-1}$ and parameterizes the rotational deviation ${}^B\mathbf{R}^{\text{pert.}}_k\left(\underline{\boldsymbol{\delta}}_{B,\,k}\right)$. The extent component $\underline{\boldsymbol{x}}_k^{\text{ext.}}$ contains the object's shape parameters. Given the wealth of available shape models [13], [14], we adopt an existing parameterization and focus on the proposed filtering framework rather than introducing a novel representation.

During each measurement update, we assume the availability of a measurement set $\mathcal{Y}_k = \{{}^I\underline{\tilde{y}}_{k,i}\}_{i=1}^{N_k}$, consisting of N_k individual Cartesian position measurements expressed in the inertial frame. The total number of measurements N_k can vary over time. These measurements are highly informative to infer the hidden system state, as they not only provide evidence about the target's kinematic state, but also reflect characteristics of its spatial extent. Each individual measurement ${}^I\underline{\tilde{y}}_{k,i}$ is modeled as a noisy observation of an unknown three-dimensional point in the inertial frame ${}^I\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i}$, referred to as the measurement source, which is assumed to lie on the target boundary, i.e.,

$$\begin{aligned}
{}^I\underline{\tilde{y}}_{k,i} &= {}^I\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i} + \underline{\boldsymbol{v}}_{k,i}\,, \\
{}^I\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i} &= {}^I\underline{\boldsymbol{c}}_k + {}^I\mathbf{R}^{\text{ref.}}_{B,\,k}\ {}^B\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i} \\
&= {}^I\underline{\boldsymbol{c}}_k + {}^I\mathbf{R}^{\text{ref.}}_{B,\,k-1}\ {}^B\mathbf{R}^{\text{pert.}}_k\left(\underline{\boldsymbol{\delta}}_{B,\,k}\right)\ {}^B\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i}\,, \\
& \qquad i = 1,\,\ldots\,,\,N_k\,,
\end{aligned} \tag{2}$$

where ${}^B\underline{\boldsymbol{z}}^{\boldsymbol{x}}_{k,i}$ is the measurement source expressed in the local body-fixed frame, $\underline{\boldsymbol{v}}_{k,i}$ denotes additive zero-mean Gaussian white noise with a known probability distribution. Measurement noises are mutually independent and independent of the state.

We represent the probability density function (PDF) of the state $\underline{\boldsymbol{x}}_k$ at time step k conditioned on the k received measurement sets $\mathcal{Y}_1, \ldots, \mathcal{Y}_k$ as

$$f^e_{\boldsymbol{x}_k}(\underline{x}_k) = f_{\boldsymbol{x}_k}\left(\underline{x}_k \mid \mathcal{Y}_1, \ldots, \mathcal{Y}_k\right) = f_{\boldsymbol{x}_k}\left(\underline{x}_k \mid \mathcal{Y}_{1:k}\right)\ , \tag{3}$$

and the predicted state density, that is the state PDF at time step k conditioned only on $\mathcal{Y}_1, \ldots, \mathcal{Y}_{k-1}$ as

$$f^p_{\boldsymbol{x}_k}(\underline{x}_k) = f_{\boldsymbol{x}_k}\left(\underline{x}_k \mid \mathcal{Y}_{1:k-1}\right)\ . \tag{4}$$

One of our primary objectives is to correct the predicted state estimate $f^p_{\boldsymbol{x}_k}(\underline{x}_k)$ using the newly arrived measurements $\mathcal{Y}_k$. Typically, this refinement is performed by applying Bayes' rule and assuming that the current measurement set $\mathcal{Y}_k$ is conditionally independent of the already processed measurement sets $\mathcal{Y}_{1:k-1}$ given the predicted state estimate

$$f\left(\mathcal{Y}_k \mid \underline{x}_k,\ \mathcal{Y}_{1:k-1}\right) = f\left(\mathcal{Y}_k \mid \underline{x}_k\right)\ . \tag{5}$$

The corrected state density then follows as

$$f^e_{\boldsymbol{x}_k}(\underline{x}_k) \propto f(\mathcal{Y}_k \mid \underline{x}_k)\ f^p_{\boldsymbol{x}_k}(\underline{x}_k)\ . \tag{6}$$

By exploiting the mutual independence of the noise terms, we obtain a likelihood function that simultaneously processes all the measurements within a single filter step

$$f(\mathcal{Y}_k \mid \underline{x}_k) = \prod_{i=1}^{N_k} f_{\boldsymbol{y}_k}({}^I\underline{\tilde{y}}_{k,i} \mid \underline{x}_k)\ . \tag{7}$$

A key benefit of this factorization is that each measurement ${}^I\underline{\tilde{y}}_{k,i}$ can be processed independently. First, we express the generative measurement model (2) in probabilistic form as the conditional density $f_{\boldsymbol{y}_k}({}^I\underline{y}_k \mid\ \underline{x}_k)$. We then define the likelihood function for a specific observation ${}^I\underline{\tilde{y}}_{k,i}$ as $f^{\mathrm{L}}_{k,i}(\underline{x}_k) \overset{\text{def}}{=} f_{\boldsymbol{y}_k}({}^I\underline{\tilde{y}}_{k,i} \mid \underline{x}_k)$. Once this likelihood function $f^{\mathrm{L}}_{k,i}(\underline{x}_k)$ is explicitly obtained, it can be incorporated into a recursive Bayesian filter for extended object tracking.

III. Key Ideas and Groundwork

We present a brief overview and notation of the basic mathematical preliminaries that are necessary to understand the Riemannian geometry used in this work. This treatment is primarily based on [15], [16], and partially on [17].

A. Rotation Manifold $\mathrm{SO}(3)$ and Its Tangent Spaces

In three dimensions, any finite rigid-body rotation can be rigorously described as an element of the special orthogonal group $\mathrm{SO}(3)$, which is the curved manifold of all valid rotation matrices $\mathbf{R}$

$$\mathrm{SO}(3) := \left\{ \mathbf{R} \in \mathbb{R}^{3\times 3} \;\middle|\; \mathbf{R}^\top \mathbf{R} = \mathbf{I}_3,\, \det \mathbf{R} = 1 \right\} . \quad (8)$$

This Lie group forms the natural geometric framework for analyzing and composing three-dimensional rotations.

While rotation matrices and unit quaternions provide global, non-singular representations, it is often convenient in practice to employ minimal three-parameter coordinates, such as the rotation vector. By definition, it is derived from the tangent space $T_{\mathbf{I}_3}\mathrm{SO}(3)$ at the identity matrix $\mathbf{I}_3$ that represents zero rotation [18]. This canonical choice yields a compact Euclidean parameter for estimation. The resulting rotation parameterization is detailed in the next subsection.

To formalize this tangent-space perspective, we briefly recall the notion of tangent vectors on $\mathrm{SO}(3)$. Let $\mathbf{\Phi}(\eta) : \mathbb{R} \to \mathrm{SO}(3)$ be a smooth curve with $\mathbf{\Phi}(0) = \mathbf{R}$. Its tangent vector at $\mathbf{R}$ is defined as

$$\dot{\mathbf{\Phi}}(0) := \lim_{\eta \to 0} \frac{\mathbf{\Phi}(\eta) - \mathbf{\Phi}(0)}{\eta} \;\in\; T_{\mathbf{R}}\mathrm{SO}(3)\ ,\ \eta \in \mathbb{R}\,, \quad (9)$$

characterizing the instantaneous rate of change of the curve at $\mathbf{R}$. The tangent space $T_{\mathbf{R}}\mathrm{SO}(3)$ consists of all such tangent vectors and provides a linear approximation of the manifold in a neighborhood of $\mathbf{R}$. This characterization forms the basis for the minimal rotation parameterizations used in this work.

B. Rotation Parameterization

Building on the tangent–space framework introduced above, we formalize the rotation vector parameterization and its mapping between the Euclidean space $\mathbb{R}^3$, the Lie algebra $\mathfrak{so}(3)$, and the rotation manifold $\mathrm{SO}(3)$. Given a rotation vector $\underline{\delta} = (\delta_1, \delta_2, \delta_3)^\top \in \mathbb{R}^3$, whose direction specifies the axis of rotation and whose magnitude $\|\underline{\delta}\|_2$ equals the rotation angle. We interpret $\underline{\delta}$ as an element of the tangent space $T_{\mathbf{I}_3}\mathrm{SO}(3) =$ Lie algebra $\mathfrak{so}(3)$ at $\mathbf{I}_3$. The correspondence between $\mathbb{R}^3$ and $\mathfrak{so}(3)$ is established via the operator $[\cdot]_\times : \mathbb{R}^3 \to \mathfrak{so}(3)$, $\underline{\delta} \mapsto [\underline{\delta}]_\times$. This operator gives rise to a skew-symmetric matrix

$$[\underline{\delta}]_\times \;=\; \begin{bmatrix} 0 & -\delta_3 & \delta_2 \\ \delta_3 & 0 & -\delta_1 \\ -\delta_2 & \delta_1 & 0 \end{bmatrix} \in \mathfrak{so}(3)\,. \quad (10)$$

Similarly, the operator $(\cdot)^\vee : \mathfrak{so}(3) \to \mathbb{R}^3$ recovers the rotation vector from its skew-symmetric matrix [5].

A rotation matrix $\mathbf{R}$ and its corresponding rotation vector $\underline{\delta}$ are related via the matrix-exponential series [15]

$$\mathbf{R} = \exp\left([\underline{\delta}]_\times\right) := \mathbf{I}_3 + [\underline{\delta}]_\times + \frac{1}{2!}\,[\underline{\delta}]_\times^2 + \frac{1}{3!}\,[\underline{\delta}]_\times^3 + \cdots .$$

For $\|\underline{\delta}\|_2 \neq 0$, setting $\mathbf{A} = [\underline{\delta}]_\times$, the cubic identity $\mathbf{A}^3 = -\|\underline{\delta}\|_2^2\,\mathbf{A}$ collapses all higher powers so that the odd- and even-order terms sum to $\sin\big(\|\underline{\delta}\|_2\big)$ and $\cos\big(\|\underline{\delta}\|_2\big)$, respectively, giving the Rodrigues' formula in closed-form

$$\mathbf{R} = \mathbf{I}_3 \;+\; \frac{\sin\big(\|\underline{\delta}\|_2\big)}{\|\underline{\delta}\|_2}\,[\underline{\delta}]_\times \;+\; \frac{1-\cos\big(\|\underline{\delta}\|_2\big)}{\|\underline{\delta}\|_2^2}\,[\underline{\delta}]_\times^2\,. \quad (11)$$

The logarithm map returns a unique skew-symmetric matrix in the Lie algebra with rotation angle in the interval $(0, \pi)$ from a rotation matrix $\mathbf{R} \neq \mathbf{I}_3$

$$\begin{aligned} &\log(\cdot) : \mathrm{SO}(3) \;\longrightarrow\; \left\{ [\underline{\delta}]_\times \in \mathfrak{so}(3) \mid 0 < \|\underline{\delta}\|_2 < \pi \right\}, \\ &[\underline{\delta}]_\times = \frac{f\,(\mathbf{R})}{2\,\sin\,(f\,(\mathbf{R}))}\,\left(\mathbf{R} - \mathbf{R}^\top\right)\,, \\ &f\,(\mathbf{R}) = \arccos\!\Big(\tfrac{\mathrm{tr}(\mathbf{R})-1}{2}\Big)\,. \end{aligned} \quad (12)$$

Because rotations differing by multiples of 2π coincide, one obtains $\log\big(\exp([\underline{\delta}]_\times)\big) = [\underline{\delta}]_\times$, if and only if $\|\underline{\delta}\|_2 < \pi$, while for $\|\underline{\delta}\|_2 \geq \pi$, the logarithm map projects the rotation vector into the principal interval, resulting in angle wrapping.

In this paper, we represent a full spatial orientation ${}^I\mathbf{R}_B$ of the body with respect to the inertial coordinate system

$${}^I\mathbf{R}_B \;:=\; {}^I\mathbf{R}_B^{\text{ref.}}\,{}^B\mathbf{R}^{\text{pert.}} = {}^I\mathbf{R}_B^{\text{ref.}}\,\exp\left([\underline{\delta}_B]_\times\right)\,, \quad (13)$$

where $\underline{\delta}_B \in \mathbb{R}^3$ is a perturbation rotation vector in the tangent space at reference ${}^I\mathbf{R}_B^{\text{ref.}}$. More precisely, $\underline{\delta}_B$ is mapped onto the Lie algebra at the identity $\mathbf{I}_3$ via the $[\cdot]_\times$ operator in (10), yielding $[\underline{\delta}_B]_\times \in \mathfrak{so}(3)$. The exponential map (11) projects this element to a perturbation rotation on $\mathrm{SO}(3)$, which is then applied by left multiplication with the reference ${}^I\mathbf{R}_B^{\text{ref.}}$. A globally valid orientation is obtained, which we adopt as the new reference. This operation is known as *left translation*. For an intuitive illustration, see the $\mathrm{SO}(2)$ examples in Fig. 2 (a) and (c). The same concepts generalize directly to the $\mathrm{SO}(3)$ setting, where the green point marks $\underline{\delta}_B$ on the tangent space at the old reference and the orange point denotes the updated reference through the exponential mapping shown by the red arrow.

C. Rotation Vector Kinematics

In three-dimensional rigid-body kinematics, the rotation vector $\underline{\delta}_B$ is related to the angular velocity $\underline{\omega}_B$ via the time derivative of the exponential map (11). In a discrete-time setting, we assume $\underline{\omega}_{B,\,k}$ is constant over $[t_k,\, t_{k+1}]$ of duration Δ_t. Under this assumption, the exact temporal evolution of the rotation vector is governed by the rotation vector composition rule introduced in [19] as

$$\underline{\delta}_{B,\,k+1} = \Big(\log\left[\exp([\underline{\delta}_{B,\,k}]_\times)\,\exp([\underline{\omega}_{B,\,k}\,\Delta_t]_\times)\right]\Big)^\vee\,, \quad (14)$$

where $\exp([\underline{\omega}_{B,\,k}\,\Delta_t]_\times)$ represents the incremental rotation resulting from the constant angular velocity. All quantities are expressed in the body-fixed coordinate system.

IV. Recursive Tangent-Space Bayesian Filtering

So far, we have established the foundations. In this section, we develop our recursive tangent-space Bayesian filtering framework for three-dimensional EOT. We first study uncertainty propagation for a tangent-space random vector across successive tangent spaces on a manifold.

A. Gaussian Assumed Density Filter

EOT is conventionally formulated as a nonlinear state estimation problem and addressed via recursive Bayesian filtering. Most standard off-the-shelf filtering approaches require a Euclidean state space. Our unified formulation directly meets this assumption, in which all the random variables, including the rotation vector, object shape, and kinematic parameters are jointly represented within a single Euclidean state vector. However, state PDFs are often multi-modal or non-Gaussian, making closed-form Bayesian updates intractable. Therefore, we employ simplifications in the form of Gaussian Assumed Density Filters (GADFs), which, at each time and measurement update, approximate the complex true state PDF with a single Gaussian.

B. Uncertainty Propagation via Parallel Transport

We model the entire hidden system state as a joint Gaussian in Euclidean space, in which the stochastic rotation vector appears as one of the Gaussian components and is interpreted in the tangent space at a deterministic reference orientation. Its mean is projected onto $\mathrm{SO}(3)$ via the exponential map after each processing step. Concurrently, its uncertainty, captured by the covariance matrix, must also be explicitly and correctly transported between the corresponding Euclidean tangent spaces. This dual mapping of mean and covariance is one of the fundamental points for the construction of our filters.

At each $\mathbf{R} \in \mathrm{SO}(3)$, the tangent space $T_{\mathbf{R}}\mathrm{SO}(3)$ is a three-dimensional real vector space isomorphic to $\mathbb{R}^3$. A tangent vector $\underline{v} \in T_{\mathbf{R}}\mathrm{SO}(3)$ is anchored at the base point $\mathbf{R}$. Since $T_{\mathbf{R_1}}\mathrm{SO}(3)$ and $T_{\mathbf{R_2}}\mathrm{SO}(3)$ are distinct vector spaces for $\mathbf{R_1} \neq \mathbf{R_2}$, vectors based at different points are not directly comparable or subtractable. They must first be transported to a common tangent space. *Parallel transport* provides a principled mechanism to move tangent vectors along a smooth curve $\alpha\colon [0,1] \to \mathrm{SO}(3)$. For each $t \in [0,1]$, the parallel–transport operator

$$\mathrm{PT}_{\alpha(0)\to\alpha(t)}(\cdot)\,:\, T_{\alpha(0)}\mathrm{SO}(3) \to T_{\alpha(t)}\mathrm{SO}(3)\,, \tag{15}$$

is a linear isomorphism that carries vectors from the source to the target tangent space. The key property of parallel transport for Lie groups is the preservation of inner products between the two Euclidean tangent spaces.

The subsequent question is how to define a well-posed curve on the manifold. We leverage the intrinsic geometry of the $\mathrm{SO}(3)$ manifold to derive its *geodesics* in closed form and then adopt these geodesics as the curves along which parallel transport is carried out. Geodesics are the local minimizer of the length of a curve between two points on a manifold $\mathcal{M}$ computed with a Riemannian metric. On $\mathrm{SO}(3)$, geodesics are described by rotations carried out at a constant angular velocity [20]. When orientations are parameterized by unit quaternions, geodesics are great circles on a unit 3-sphere $\mathbb{S}^3$. Fig. 3 shows an analogous great circle on a $\mathbb{S}^2$.

Definition 1: Let $\mathbf{R}_0, \mathbf{R}_1 \in \mathrm{SO}(3)$ and define the relative rotation $\mathbf{\Delta} := \mathbf{R}_0^{-1}\,\mathbf{R}_1 = \mathbf{R}_0^{\top}\,\mathbf{R}_1 \in \mathrm{SO}(3)$, where we have used $\mathbf{R}_0^{-1} = \mathbf{R}_0^{\top}$, since rotation matrices are orthogonal.

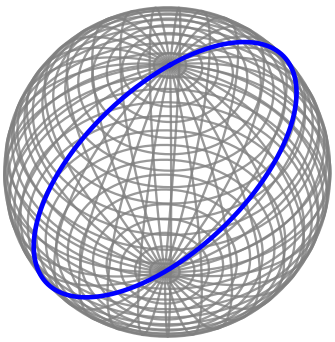

Fig. 3: A great circle (blue) on the unit 2-sphere $\mathbb{S}^2$, shown as a lower-dimensional analogy for geodesics on $\mathrm{SO}(3)$. Under unit-quaternion coordinates, geodesics on $\mathrm{SO}(3)$ are great circles on $\mathbb{S}^3$. The gray $\mathbb{S}^2$ is for intuition only.

Let $\mathbf{A} := \log(\mathbf{\Delta}) \in \mathfrak{so}(3)$ denote the matrix logarithm. Equivalently, there exists a unique rotation vector $\underline{\omega} \in \mathbb{R}^3$ with $\|\underline{\omega}\|_2 \in [0,\pi)$ such that $\mathbf{A} = [\underline{\omega}]_\times$. The curve on $\mathrm{SO}(3)$

$$\gamma(t) \;=\; \mathbf{R}_0\,\exp\!\big(t\,\mathbf{A}\big) \in \mathrm{SO}(3)\,,\; t \in [0,1] \subset \mathbb{R}\,, \tag{16}$$

is the unique *geodesic* connecting $\mathbf{R}_0$ and $\mathbf{R}_1$ with $\gamma(0) = \mathbf{R}_0$ and $\gamma(1) = \mathbf{R}_1$. When $\|\underline{\omega}\|_2 = 0$, i.e., $\mathbf{R}_0 = \mathbf{R}_1$, γ reduces to the constant curve. □

We equip $\mathrm{SO}(3)$ with the canonical Riemannian metric [16]. For $\mathbf{R} \in \mathrm{SO}(3)$ and $\underline{u}_1, \underline{u}_2 \in T_{\mathbf{R}}\mathrm{SO}(3)$, define

$$\langle \underline{u}_1,\, \underline{u}_2 \rangle_{\mathbf{R}} := \langle \mathbf{R}^{-1}\underline{u}_1,\, \mathbf{R}^{-1}\underline{u}_2 \rangle_{\mathbf{I}_3}\,. \tag{17}$$

The induced norm is $\|\underline{u}\|_{\mathbf{R}}^2 = \langle \underline{u},\, \underline{u} \rangle_{\mathbf{R}}$. Under the definition of the geodesic and using $\gamma'(\tau) = \gamma(\tau)\,\mathbf{A}$, $\gamma(\tau)^{-1}\,\gamma(\tau) = \mathbf{I}_3$, its energy functional reduces to a constant, highlighting an important structural property of geodesics.

$$\begin{aligned} E(\gamma) \;&=\; \frac{1}{2}\int_0^1 \underbrace{\big\|\gamma'(\tau)\big\|_{\gamma(\tau)}^2}_{=\,\langle \gamma'(\tau),\gamma'(\tau)\rangle_{\gamma(\tau)}^2}\;\mathrm{d}\tau \\ &=\; \frac{1}{2}\int_0^1 \underbrace{\big\langle \gamma(\tau)^{-1}\,\gamma'(\tau),\, \gamma(\tau)^{-1}\,\gamma'(\tau)\big\rangle_{\mathbf{I}_3}}_{=\,\big\|\gamma(\tau)^{-1}\,\gamma'(\tau)\big\|_{\mathbf{I}_3}^2}\,\mathrm{d}\tau \\ &=\; \frac{1}{2}\int_0^1 \big\|\gamma(\tau)^{-1}\,\gamma(\tau)\,\mathbf{A}\big\|_{\mathbf{I}_3}^2\,\mathrm{d}\tau = \text{const.}\,. \end{aligned} \tag{18}$$

Below we give the definition of parallel transport on $\mathrm{SO}(3)$.

Definition 2: Let $\mathbf{R}_0, \mathbf{R}_1 \in \mathrm{SO}(3)$ and set $\gamma(t) = \mathbf{R}_0\,\exp\!\big(t\,\mathbf{A}\big)$, $\mathbf{A} = \log\!\big(\mathbf{R}_0^{\top}\,\mathbf{R}_1\big)$, $t \in [0,1]$ be the geodesic joining $\mathbf{R}_0$ to $\mathbf{R}_1$. For any tangent vector $\underline{v} \in T_{\mathbf{R}_0}\mathrm{SO}(3)$, write $\underline{v} = \mathbf{R}_0\,\underline{n}$, $\underline{n} \in \mathfrak{so}(3)$, $\underline{n} = \mathbf{R}_0^{\top}\,\underline{v}$. Then the *left-trivialized parallel transport* of $\underline{v}$ along γ is defined by

$$\mathrm{PT}_{\mathbf{R}_0\to\gamma(t)}(\underline{v}) \;=\; \gamma(t)\,\underline{n} \;=\; \mathbf{R}_0\,\exp\!\big(t\,\mathbf{A}\big)\,\big(\mathbf{R}_0^{\top}\,\underline{v}\big). \tag{19}$$

This map is a linear isometry under the canonical bi-invariant Riemannian metric. □

At $t = 1$ with $\gamma(1) = \mathbf{R}_1$, i.e., after transporting from $t = 0$ to $t = 1$ along the full geodesic, we obtain

$$\begin{aligned}\mathrm{PT}_{\mathbf{R}_0 \to \mathbf{R}_1}(\underline{v}) &= \mathbf{R}_0 \, \exp\big(\mathbf{A}\big) \, \big(\mathbf{R}_0^\top \, \underline{v}\big) \\ &= \mathbf{R}_0 \, \exp\big(\log(\mathbf{R}_0^\top \, \mathbf{R}_1)\big) \, \big(\mathbf{R}_0^\top \, \underline{v}\big) \\ &= \mathbf{R}_1 \, \mathbf{R}_0^\top \, \underline{v} \, . \end{aligned} \tag{20}$$

Hence, transporting $\underline{v}$ along the full geodesic from $\mathbf{R}_0$ to $\mathbf{R}_1$ is equivalent to applying the single rotation $\mathbf{R}_1 \, \mathbf{R}_0^\top$ to $\underline{v}$.

Beyond vectors, parallel transport applies naturally to covariance matrices, yielding a rigorous mechanism for uncertainty propagation on manifold-valued trajectories. We adapt the proposition in [21] to the specific case of $\mathrm{SO}(3)$.

Proposition 1: Let $\mathbf{\Sigma}_0 \in \mathbb{R}^{3\times3}$ be a symmetric, and positive definite covariance matrix on the tangent space $T_{\mathbf{R}_0}\mathrm{SO}(3)$ at $\mathbf{R}_0$. $\mathbf{\Sigma}_0$ admits the eigendecomposition

$$\mathbf{\Sigma}_0 \; = \; \sum_{m=1}^{3} \lambda_m \, \underline{v}_m \, \underline{v}_m^\top, \tag{21}$$

where $\{\underline{v}_m\}_{m=1}^3$ are the eigenvectors of $\mathbf{\Sigma}_0$, forming an orthonormal basis of $T_{\mathbf{R}_0}\mathrm{SO}(3)$, and λ_m are the eigenvalues. Let $\alpha\colon [0,1] \to \mathrm{SO}(3)$ be a geodesic on $\mathrm{SO}(3)$ with $\alpha(0) = \mathbf{R}_0$ and denote the parallel transport of each eigenvector $\underline{v}_m$ along α for $t \in [0,1]$ as $\underline{v}'_m(t) = \mathrm{PT}_{\mathbf{R}_0 \to \alpha(t)}(\underline{v}_m)$. According to definition 2, the updated covariance matrix

$$\begin{aligned}\mathbf{\Sigma}_t \; &= \; \sum_{m=1}^{3} \lambda_m \, \underline{v}'_m(t) \, \underline{v}'_m(t)^\top \\ &= \; \sum_{m=1}^{3} \lambda_m \big[\mathbf{R}_0 \, \exp\big(t\mathbf{A}\big) \, \mathbf{R}_0^\top\big] \, \underline{v}_m \, \underline{v}_m^\top \big[\mathbf{R}_0 \, \exp\big(t\mathbf{A}\big) \, \mathbf{R}_0^\top\big]^\top \\ &= \big[\mathbf{R}_0 \, \exp\big(t\mathbf{A}\big) \, \mathbf{R}_0^\top\big] \, \mathbf{\Sigma}_0 \big[\mathbf{R}_0 \, \exp\big(t\mathbf{A}\big) \, \mathbf{R}_0^\top\big]^\top \end{aligned} \tag{22}$$

is precisely the parallel transport of $\mathbf{\Sigma}_0$ along curve α. □

Building on the results above, we have thus achieved our key objective for uncertainty propagation along $\mathrm{SO}(3)$. After mapping the mean rotation vector from $T_{\mathbf{R}_0}\mathrm{SO}(3)$ onto a new point $\mathbf{R}_1$ of $\mathrm{SO}(3)$ via the exponential map, we transport its covariance matrix $\mathbf{\Sigma}_0$ in closed form as

$$\mathbf{\Sigma}_1 = \mathbf{R}_1 \, \mathbf{R}_0^\top \, \mathbf{\Sigma}_0 \, \mathbf{R}_0 \, \mathbf{R}_1^\top \, . \tag{23}$$

This procedure ensures that the Gaussian uncertainty remains correctly aligned with the local Euclidean approximation at each updated reference orientation.

C. Post-Update Orientation Reset

After each prediction step or measurement update, our framework executes three core operations to maintain consistency between the manifold-valued mean and its Euclidean-space covariance. Given the predicted or posterior tangent-space mean vector ${}^{B}\hat{\underline{\delta}}_{k|1:k-1}$, ${}^{B}\hat{\underline{\delta}}_{k|1:k} \in \mathbb{R}^3$, we retract it onto $\mathrm{SO}(3)$ via the exponential map

$$\begin{aligned}{}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k-1} &= {}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k-1|1:k-1} \, \exp\left(\left[{}^{B}\hat{\underline{\delta}}_{k|1:k-1}\right]_\times\right) \text{ or} \\ {}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k} &= {}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k-1} \, \exp\left(\left[{}^{B}\hat{\underline{\delta}}_{k|1:k}\right]_\times\right) , \end{aligned} \tag{24}$$

respectively, thus yielding a new deterministic reference orientation on $\mathrm{SO}(3)$ after each processing step.

Moreover, we transport the covariance from the previous tangent space into the tangent space at the updated reference orientation via parallel transport. We construct an orthogonal update matrix after prediction step or measurement update

$$\begin{aligned}\mathbf{P}^{\text{Pred.}} &= {}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k-1} \left[{}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k-1|1:k-1}\right]^\top \text{ or} \\ \mathbf{P}^{\text{Filt.}} &= {}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k} \left[{}^{I}\mathbf{R}^{\text{ref.}}_{B,\,k|1:k-1}\right]^\top , \end{aligned} \tag{25}$$

and then update the covariance matrix of $\underline{\boldsymbol{\delta}}$ by

$$\begin{aligned}\mathbf{\Sigma}^{\boldsymbol{\delta}}_{k|1:k-1} &= \mathbf{P}^{\text{Pred.}} \, \mathbf{\Sigma}^{\boldsymbol{\delta}}_{k-1|1:k-1} \, \left[\mathbf{P}^{\text{Pred.}}\right]^\top \text{ or} \\ \mathbf{\Sigma}^{\boldsymbol{\delta}}_{k|1:k} &= \mathbf{P}^{\text{Filt.}} \, \mathbf{\Sigma}^{\boldsymbol{\delta}}_{k|1:k-1} \, \left[\mathbf{P}^{\text{Filt.}}\right]^\top . \end{aligned} \tag{26}$$

Finally, we reset the tangent-space mean to zero

$${}^{B}\hat{\underline{\delta}}_{k|1:k-1} \; = \; \underline{0} \;\; \text{or} \;\; {}^{B}\hat{\underline{\delta}}_{k|1:k} \; = \; \underline{0} \, . \tag{27}$$

so that subsequent filtering steps always operate around a zero-mean Gaussian in the newly updated tangent space.

These three steps, including exponential mapping, covariance transport, and mean value re-centering, are repeated after each prediction and measurement update, guaranteeing that both mean and covariance remain correctly aligned with the local linearization at every reference orientation.

D. Nonlinear Bayesian Filtering

To address the nonlinear state estimation problem, we adopt the advanced particle-flow-based Progressive Gaussian Filters (PGFs) [22], [23], which outperforms conventional nonlinear Kalman filters. Unlike methods that rely on linearizing the nonlinear measurement equation, PGF employs a progression mechanism that splits each measurement update into a sequence of sub-updates and gradually incorporates observation information into the state estimate. Within this framework, closed-form likelihoods for a specific target tracking scenario have been derived [24]. In this work, we choose PGF [23] for recursive Bayesian inference due to its proven accuracy and high computational efficiency.

V. Evaluation

We evaluate our proposed EOT framework on a simulated free-fall experiment involving an elliptic cone released from a known height. During its descent, the cone undergoes fully unconstrained rigid body motion. It is endowed with an unknown initial angular velocity about an axis through its center of mass and subsequently rotates freely. This scenario provides a realistic and representative benchmark for simultaneous estimation of the full pose with six DoF and the object's geometric parameters.

To keep the simulation tractable, we fix the cone's center of mass as the reference point for all force and torque analyses. We assume a uniform gravitational field acting exclusively through this point, so that gravity induces no net torque. We further neglect any other external torques and model the cone as a perfectly rigid body of homogeneous density, yielding a constant inertia tensor in the body frame.

Under these assumptions, translational and rotational dynamics decouple completely and may be treated independently. Our evaluation thus concentrates on the EOT framework's capacity to recover both the full pose trajectory and the cone's shape under idealized free-fall conditions.

We benchmark our framework against state-of-the-art orientation estimators. In [5], approaches based on the Extended Kalman Filter (EKF) and the Unscented Kalman Filter (UKF) are introduced. Like our approach, these filters encode rotational perturbations as rotation vectors within the state. However, while we compute the update matrix in (25) by parallel transport, they propagate uncertainty between tangent spaces using different mechanisms. We also evaluate a zero-order approximation technique [4], [25], in which the covariance matrix of the stochastic rotation vector is kept unchanged after each update. In other words, our covariance-update procedure defined in (26) is not employed in these methods. To ensure a fair comparison among all frameworks for orientation estimation, we employ the same shape model and the same nonlinear filtering technique.

The experiment evaluates multiple EOT frameworks in a recursive state estimation scenario. We generate 100 independent ground-truth state trajectories, each comprising 500 discrete time steps sampled at a constant interval of 0.01 seconds. The target is an elliptic cone with height $h = 10$ cm and base semi-axes $r_x = 4.5$ cm, $r_y = 2.5$ cm. At each time step, we observe surface point measurements that are corrupted by additive, zero-mean isotropic Gaussian noise with a standard deviation of 0.3 cm per coordinate. We quantify the error of orientation estimation between two rotation matrices, namely a true one $\mathbf{R}^{\mathrm{GT}}$ and an estimated one $\mathbf{R}^{\mathrm{Est.}}$, by the smallest relative rotation angle that aligns them. Concretely, let $\mathbf{R}^{\mathrm{GT}}, \mathbf{R}^{\mathrm{Est.}} \in \mathrm{SO}(3)$ be two valid rotation matrices. We define the relative rotation $\mathbf{R}^{\mathrm{Err.}} = \left(\mathbf{R}^{\mathrm{GT}}\right)^{\top} \mathbf{R}^{\mathrm{Est.}}$, compute $c = \frac{\mathrm{tr}\left(\mathbf{R}^{\mathrm{Err.}}\right)-1}{2}$, and then clamp c to the interval $[-1, 1]$. The angular error in degrees is then $\theta_{\mathrm{err}} = \frac{180}{\pi} \arccos(c)$. To expedite execution, all simulations were implemented in `Julia` and deployed on the HAICORE KIT cloud platform. Each Monte Carlo trial was assigned a dedicated CPU core on an Intel Xeon Platinum 8368 CPU with 2.40 GHz, thereby fully exploiting the available parallel-computing infrastructure and reducing overall runtime.

The mean rotation error in degrees between the ground-truth and estimated orientations is computed at each of the 500 time steps over 100 Monte Carlo runs. The resulting error curves, shown in Fig. 4, reveal distinct behaviors for the three methods under comparison. The zero-order scheme used in [25], [4] can be viewed as a straightforward extension of the classical Multiplicative Extended Kalman Filter (MEKF). In our experiments, it achieves only a marginal initial error reduction before rapidly diverging. By step 200, the mean error exceeds 14° and continues to grow rapidly thereafter, indicating an inability to correct the orientation error effectively. The reason for that behavior is the ignorance of the explicit uncertainty propagation. The full-order reset approach [5] exhibits little improvement. The error drops to around 9.5° by step 100, but then steadily climbs. This behavior may be attributed to incorrect propagation of uncertainty between successive tangent-space updates, leading to continuous filter degradation. In contrast, our method quickly reduces the mean orientation error and maintains it at a relative low and nearly constant level until the end of the simulation. Although exact reconstruction of the ground-truth orientation trajectory is not achievable due to measurement noise and model inaccuracies, the error consistently converges to a low magnitude, demonstrating the effectiveness of the proposed framework. Fig. 5 plots the Root Mean Square Error (RMSE) between the ground-truth and estimated shape parameters of the elliptic cone, i.e., its height and the two base semi-axes, over the same trials. We show only our method, as the two baselines diverge in orientation and thus yield uninformative geometry estimates. Our framework achieves low, convergent RMSEs, indicating accurate geometric reconstruction alongside reliable orientation estimation.

VI. Conclusion

In this work, we present a recursive Bayesian filter for joint pose–shape estimation that treats rotations on $\mathrm{SO}(3)$ via a stochastic tangent-space parameterization with explicit parallel transport of uncertainty. By estimating a stochastic rotation vector in the Euclidean tangent space about a deterministic time-varying reference, the method integrates with standard GADFs while consistently propagating uncertainty on $\mathrm{SO}(3)$. However, working in tangent spaces is only local and thus works best under small rotation vector uncertainty around the current reference. Large rotations between successive updates can cause the local, reference-centered linearization to break down. In future work, we will exploit the state's product structure, modeled as a Cartesian product of a nonlinear manifold and a Euclidean space.

VII. Acknowledgments

The IGF project 01IF22901N of the research association Forschungs-Gesellschaft Verfahrenstechnik e.V. (GVT) is supported in a program to promote Industrial Collective Research (IGF) by the Federal Ministry of Economic Affairs and Climate Action on the basis of a decision of the German Bundestag. This work is also supported by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT partition.

References

[1] Jiaying Lin, Giovanni Campa, Christian-Eike Framing, Jan-Jöran Gehrt, René Zweigel, and Dirk Abel. "Adaptive shape fitting for LiDAR object detection and tracking in maritime applications". In: *International Journal of Transport Development and Integration* 5.2 (2021), pp. 105–117.

[2] Tim Baur, Johannes Reuter, Antonio Zea, and Uwe D. Hanebeck. "Shape Estimation and Tracking using Spherical Double Fourier Series for Three-Dimensional Range Sensors". In: *Proceedings of the 2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2021)*. Karlsruhe, Germany, Sept. 2021. DOI: 10.1109/MFI52462.2021.9591169.

[3] Tim Baur, Johannes Reuter, Antonio Zea, and Uwe D. Hanebeck. "Extent Estimation of Sailing Boats Applying Elliptic Cones to 3D LiDAR Data". In: *Proceedings of the 25th International Conference on Information Fusion (Fusion 2022)*. Linköping, Sweden, July 2022. DOI: 10.23919/FUSION49751.2022.9841265.

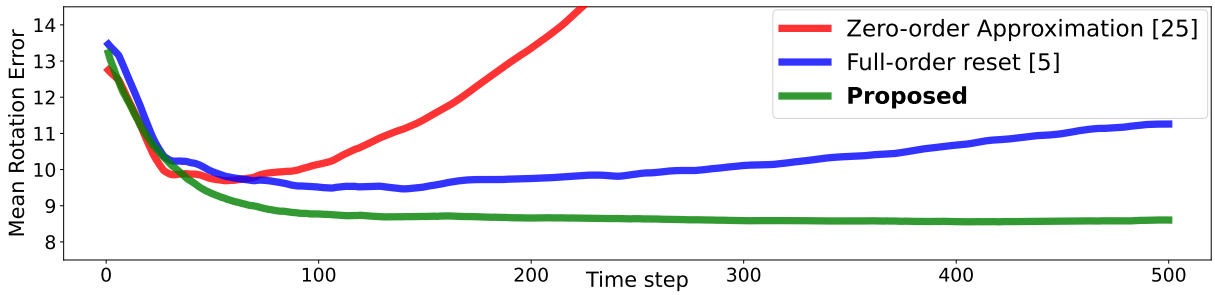

Fig. 4: Monte Carlo mean rotational error (in degrees) over 500 time steps under free-fall conditions, comparing the zero-order [25], full-order reset [5], and proposed filters.

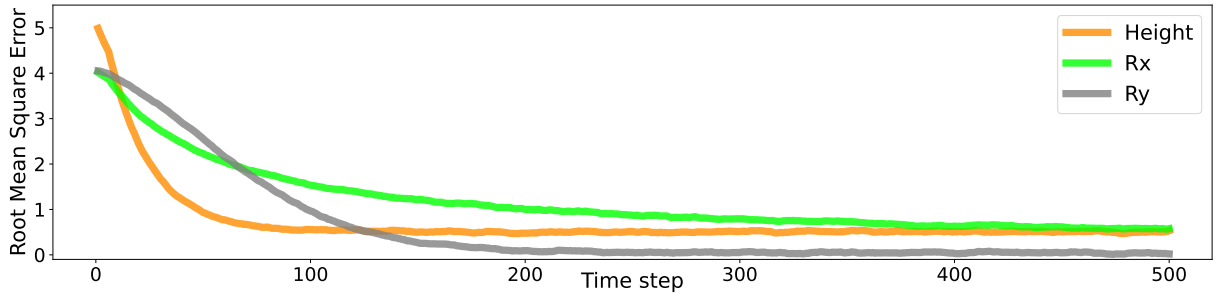

Fig. 5: Root-mean-square error (in centimeters) of estimated shape parameters over time for the proposed framework.

[4] Murat Kumru and Emre Özkan. "Three-dimensional extended object tracking and shape learning using Gaussian processes". In: *IEEE Transactions on Aerospace and Electronic Systems* 57.5 (2021), pp. 2795–2814.

[5] Rajan Gill, Mark W Mueller, and Raffaello D'Andrea. "Full-order solution to the attitude reset problem for Kalman filtering of attitudes". In: *Journal of Guidance, Control, and Dynamics* 43.7 (2020), pp. 1232–1246.

[6] Florian Pfaff, Kailai Li, and Uwe D. Hanebeck. "The State Space Subdivision Filter for SE(3)". In: *Proceedings of the 25th International Conference on Information Fusion (Fusion 2022)*. Linköping, Sweden, July 2022. DOI: 10.23919/FUSION49751.2022.9841384.

[7] Daniel Martinec and Tomas Pajdla. "Robust rotation and translation estimation in multiview reconstruction". In: *2007 IEEE conference on computer vision and pattern recognition*. IEEE. 2007, pp. 1–8.

[8] Robert Mahony, Tarek Hamel, and Jean-Michel Pflimlin. "Nonlinear complementary filters on the special orthogonal group". In: *IEEE Transactions on automatic control* 53.5 (2008), pp. 1203–1218.

[9] F Landis Markley. "Attitude estimation or quaternion estimation?" In: *The Journal of the Astronautical Sciences* 52 (2004), pp. 221–238.

[10] Urban Maeder and Manfred Morari. "Attitude estimation for vehicles with partial inertial measurement". In: *IEEE Transactions on Vehicular Technology* 60.4 (2011), pp. 1496–1504.

[11] Gerhard Kurz, Florian Pfaff, and Uwe D. Hanebeck. "Discretization of SO(3) Using Recursive Tesseract Subdivision". In: *Proceedings of the 2017 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2017)*. Daegu, Republic of Korea, Nov. 2017. DOI: 10.1109/MFI.2017.8170406.

[12] Kailai Li, Florian Pfaff, and Uwe D. Hanebeck. "Grid-Based Quaternion Filter for SO(3) Estimation". In: *Proceedings of the 2020 European Control Conference (ECC 2020)*. Virtual, May 2020. DOI: 10.23919/ECC51009.2020.9143723.

[13] Antonio Zea, Florian Faion, and Uwe D. Hanebeck. "Tracking Extended Objects using Extrusion Random Hypersurface Models". In: *Proceedings of the IEEE ISIF Workshop on Sensor Data Fusion: Trends, Solutions, Applications (SDF 2014)*. Bonn, Germany, Oct. 2014. DOI: 10.1109/SDF.2014.6954722.

[14] Florian Faion, Antonio Zea, Jannik Steinbring, Marcus Baum, and Uwe D. Hanebeck. "Recursive Bayesian Pose and Shape Estimation of 3D Objects Using Transformed Plane Curves". In: *Proceedings of the IEEE ISIF Workshop on Sensor Data Fusion: Trends, Solutions, Applications (SDF 2015)*. Bonn, Germany, Oct. 2015. DOI: 10.1109/SDF.2015.7347698.

[15] Jari Mäkinen. "Rotation manifold SO (3) and its tangential vectors". In: *Computational Mechanics* 42 (2008), pp. 907–919.

[16] Manfredo Perdigao Do Carmo and J Flaherty Francis. *Riemannian geometry*. Vol. 2. Springer, 1992.

[17] Xavier Pennec. "Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements". In: *Journal of Mathematical Imaging and Vision* 25 (2006), pp. 127–154.

[18] Ralph Abraham, Jerrold E Marsden, and Tudor Ratiu. *Manifolds, tensor analysis, and applications*. Vol. 75. Springer Science & Business Media, 2012.

[19] Malcolm D Shuster et al. "A survey of attitude representations". In: *Navigation* 8.9 (1993), pp. 439–517.

[20] Alyssa Novelia and Oliver M O'Reilly. "On geodesics of the rotation group SO (3)". In: *Regular and Chaotic Dynamics* 20 (2015), pp. 729–738.

[21] Søren Hauberg, François Lauze, and Kim Steenstrup Pedersen. "Unscented Kalman filtering on Riemannian manifolds". In: *Journal of mathematical imaging and vision* 46 (2013), pp. 103–120.

[22] Uwe D. Hanebeck. "PGF 42: Progressive Gaussian Filtering with a Twist". In: *Proceedings of the 16th International Conference on Information Fusion (Fusion 2013)*. Istanbul, Turkey, July 2013.

[23] Jannik Steinbring and Uwe D. Hanebeck. "Progressive Gaussian Filtering Using Explicit Likelihoods". In: *Proceedings of the 17th International Conference on Information Fusion (Fusion 2014)*. Salamanca, Spain, July 2014.

[24] Jannik Steinbring, Marcus Baum, Antonio Zea, Florian Faion, and Uwe D. Hanebeck. "A Closed-Form Likelihood for Particle Filters to Track Extended Objects with Star-Convex RHMs". In: *Proceedings of the 2015 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 2015)*. San Diego, California, USA, Sept. 2015. DOI: 10.1109/MFI.2015.7295740.

[25] F Landis Markley. "Attitude error representations for Kalman filtering". In: *Journal of guidance, control, and dynamics* 26.2 (2003), pp. 311–317.