Beyond Kalman Filtering: a Dual Quaternion Geometric Observer for Robust 6DoF Pose Estimation

Ilya Afanasyev
Innopolis University
Innopolis, Russia
i.afanasyev@innopolis.ru
Saint Petersburg Electrotechnical University "LETI"
St. Petersburg, Russia
imafanasev@etu.ru

***Abstract*—Autonomous robots and wearable devices operating in dynamic 3D environments require robust pose estimation that preserves the geometric structure of rigid-body motion. Traditional filtering approaches, such as Extended Kalman Filters (EKF), often decouple rotation and translation, leading to linearization artifacts and kinematic inconsistencies. This paper presents a comprehensive benchmark of a Dual Quaternion Geometric Observer (GeoDQ) against state-of-the-art manifold-aware filters (UKF-M) and standard error-state filters (ESKF). We introduce an optimized architecture featuring a Screw Linear Interpolation (SCLERP) fusion mechanism, augmented with a heuristic tilt correction, PID-based feedback loops, and Numba-based JIT compilation for real-time performance. Evaluated on 35 distinct trajectories from the RoNIN dataset under wearable smartphone configurations with unstable IMU orientations, the proposed GeoDQ (JIT) achieves a Root Mean Square Error (RMSE) of 0.0043 m. By interpolating directly on the $SE(3)$ manifold, GeoDQ outperforms the standard ESKF (0.1409 m) and UKF-M (0.2312 m) by orders of magnitude. Furthermore, robustness analysis demonstrates that the geometric observer exhibits smooth error growth under severe measurement decimation, maintaining sub-meter accuracy where filter-based methods rapidly diverge. Despite its mathematical rigor, the JIT-optimized GeoDQ executes in just 3.0 s per trajectory, making it significantly faster than standard Python-based filters and highly suitable for edge computing applications, wearable devices, and resource-constrained platforms.**

I. Introduction

Robust six-degree-of-freedom (6DoF) pose estimation is fundamental for autonomous robots, from wearable devices to embodied platforms, operating in complex, GPS-denied environments. Contemporary mobile platforms combine inertial measurement units (IMU) with visual odometry (VO), yet conventional fusion methods often decouple translation and rotation, yielding kinematic inconsistencies, singularities (e.g., gimbal lock), and long-term drift. Standard approaches like the Extended Kalman Filter (ESKF) ignore the intrinsic geometry of the Special Euclidean group $SE(3)$, leading to linearization artifacts and suboptimal performance [1], [2]. While Visual-Inertial Odometry (VIO) has advanced the state-of-the-art [3], [4], these methods remain computationally expensive for resource-constrained embedded robotic systems.

Recent directions exploit Lie group theory to improve robustness [5], [6]. In particular, dual quaternions offer an elegant alternative by representing pose as a unified entity on the $SE(3)$ manifold, enabling coordinate-free uncertainty propagation and avoiding the singularity issues inherent in Euler angles. While advanced filters like the Unscented Kalman Filter on Manifold (UKF-M) provide high precision through sigma-point transformations, their computational footprint remains prohibitive for embedded hardware.

To address these challenges, we propose a **Dual Quaternion Geometric Observer (GeoDQ)** that leverages Screw Linear Interpolation (SCLERP) [7] for geometric state fusion. Our approach treats pose as a single geometric entity, integrating high-rate IMU propagation (200 Hz) with visual odometry corrections. Following standard practice for the RoNIN dataset, the original 5 Hz VO measurements are temporally aligned with the IMU stream via interpolation, enabling synchronous fusion while preserving the asynchronous nature of real-world sensor inputs. By combining this geometric formulation with Numba-based JIT compilation, we achieve both superior accuracy and real-time performance suitable for resource-constrained platforms.

Evaluated on 35 distinct trajectories from the RoNIN (Robust Neural Inertial Navigation) dataset [8], the proposed GeoDQ achieves an RMSE of **0.0043 m**, outperforming both ESKF (0.1409 m) and UKF-M (0.2312 m) by orders of magnitude. Moreover, the JIT-optimized implementation processes an entire trajectory in just **3.05 s**, demonstrating a favorable trade-off between estimation quality and computational efficiency.

Our key contributions are fourfold:

1) **Geometric SE(3) Formulation:** We reformulate 6DoF sensor fusion as a pure geometric estimation problem using unit dual quaternions. This unified representation eliminates the artificial decoupling of rotation and translation, resolving classical singularities (e.g., gimbal

lock) while intrinsically preserving the rigid-body motion structure.

2) **Adaptive SCLERP Fusion Algorithm:** We introduce a novel update mechanism based on Screw Linear Interpolation (SCLERP) with robust Huber-weighted innovation gating. Augmented by heuristic tilt correction and PID-based bias feedback, our observer estimates 6DoF poses without requiring Jacobians, Euler angles, or complex covariance matrices, ensuring kinematically valid updates.
3) **Edge-Optimized Implementation:** We provide a highly efficient software architecture leveraging Numba JIT compilation, achieving a $\sim 7\times$ speedup over standard Python execution. Processing an entire trajectory in just **3.05 s**, the system is explicitly tailored for high-frequency, real-time operation on resource-constrained wearable and IoT edge platforms.
4) **Benchmark and Robustness Validation:** We conduct a comprehensive evaluation on 35 diverse trajectories from the RoNIN dataset. GeoDQ not only achieves state-of-the-art accuracy (RMSE of **0.0043 m**) and outperforms JIT-optimized ESKF and UKF-M baselines, but crucially demonstrates exceptional resilience to measurement sparsity. Unlike traditional filters that rapidly diverge during VO dropouts, the geometric observer maintains sub-meter accuracy with smooth, gradual error degradation.

By grounding fusion in the intrinsic geometry of $SE(3)$ and combining it with practical optimization techniques, this work offers a lightweight, drift-resistant pose estimation pipeline well suited for resource-constrained autonomous robots operating in dynamic, GPS-denied environments. The computational efficiency of GeoDQ also enables continuous deployment on low-power edge IoT devices and lightweight wearable systems where preserving battery life is critical.

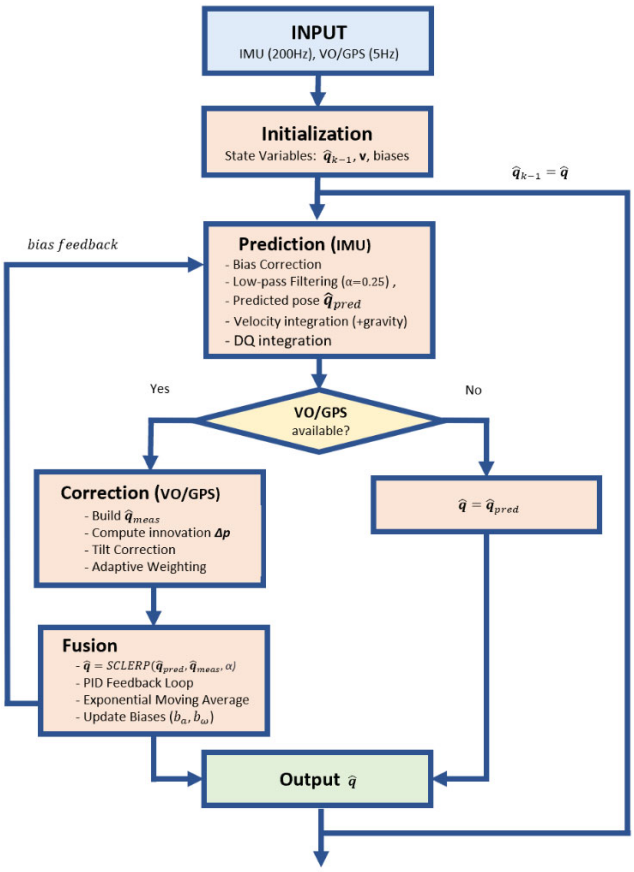

Fig. 1. GeoDQ Observer Architecture: IMU prediction corrected via SCLERP fusion with bias feedback loops.

II. Related Work

Accurate 6DoF pose estimation has long relied on sensor fusion frameworks combining inertial, magnetic, and global navigation data. The Extended Kalman Filter (EKF) and its variants dominate mobile robotic systems, typically representing orientation via quaternions and position as a separate 3D vector. As discussed in Section I, such decoupled state representations introduce artificial separation between rotation and translation, leading to linearization artifacts and kinematic inconsistencies. Complementary filters offer computationally lightweight alternatives but generally sacrifice long-term accuracy and effective drift suppression.

In visual-inertial navigation and SLAM, tightly coupled methods such as Visual-Inertial Odometry (VIO) have advanced the state of the art by integrating camera and IMU data [3]. However, these approaches remain computationally expensive and often treat translation and rotation in separate estimation spaces [4], limiting geometric fidelity. Smartphone-based pose tracking systems have demonstrated practical performance in AR applications [9], and augmented reality navigation methods combining image segmentation with multi-sensor fusion achieve high building recognition rates [10], yet these solutions still rely on decoupled estimation schemes.

A more principled direction emerges from exploiting Lie group theory for state estimation. The Invariant Extended Kalman Filter (IEKF) leverages group symmetries to reduce linearization errors [5], but its complexity and dependency on specialized frameworks limit widespread adoption. Extensions such as LG-EKF [6], matrix Lie group $\mathrm{SE}_K(3)$ [11], the Iterated IEKF [12], and pedestrian dead-reckoning via IEKF on smartphones [13] offer improved robustness, yet practical implementations on consumer devices remain scarce. Robust estimation techniques, such as RAESKF for MEMS-IMU attitude [14] and quaternion-based fusion under magnetic disturbances [15], further enhance filter resilience. Smartphone performance evaluations underscore the challenges of consumer-grade sensors in GNSS-denied settings, motivating lightweight geometric approaches [16].

Dual quaternions provide a compact algebraic representation for unified 6DoF pose encoding, enabling Screw Linear Interpolation (SCLERP) which generalizes spherical linear interpolation to full rigid-body motion while preserving the underlying screw axis. Despite mathematical elegance and demonstrated benefits in robotics manipulation and animation

[7], dual quaternions have seen limited use in mobile sensor fusion. Recent surveys and applications in manipulator dynamics highlight their potential [17], [18], but real-time smartphone implementations are lacking. Prior work on dual quaternion filtering includes spacecraft navigation via a dual quaternion EKF [19], robotic control algorithms based on screw theory [20], and theoretical analyses of geometric state estimation on SE(3) [1], [2]. However, these efforts focus primarily on theoretical formulations or high-computational platforms rather than resource-constrained embedded robotic systems.

In contrast to these efforts, our work delivers a practical dual quaternion SCLERP fusion pipeline optimized for smartphone inertial and visual odometry data. By treating pose as a single entity on SE(3) and leveraging coordinate-free interpolation with JIT optimization, we achieve drift-resistant, geometrically consistent trajectories in real time on resource-constrained devices. Our comprehensive benchmark against both conventional filters (ESKF) and manifold-aware methods (UKF-M) demonstrates superior accuracy-efficiency trade-offs. While contemporary learning-based VIO methods show high accuracy, they often exceed the computational budget of edge devices; thus, we focus on benchmarking against computationally comparable state-of-the-art manifold filters, addressing the gap between theoretical geometric approaches and practical deployment requirements. The computational efficiency of GeoDQ also paves the way for extended autonomous navigation on extreme edge-computing hardware.

III. Mathematical Background

A. Manifold-Aware Unscented Kalman Filter (UKF-M)

The UKF-M addresses nonlinear state evolution through deterministic sigma-point sampling on the $\mathrm{SO}(3) \times \mathbb{R}^3$ manifold [21]. Our implementation maintains a 15-dimensional error state: $\mathbf{x} = [\mathbf{p}^T, \mathbf{v}^T, \boldsymbol{\theta}^T, \mathbf{b}_g^T, \mathbf{b}_a^T]^T$, where $\mathbf{p}$ is position, $\mathbf{v}$ is velocity, $\boldsymbol{\theta}$ represents rotation error in the $\mathfrak{so}(3)$ Lie algebra, and $\mathbf{b}_g, \mathbf{b}_a$ are gyroscope and accelerometer biases.

Sigma-point generation uses the scaled unscented transform with Van der Merwe parameters ($\alpha = 0.1$, $\beta = 2.0$, $\kappa = 0$) [22]:

$$\mathcal{X}_0 = \mathbf{0}, \quad \mathcal{X}_i = \pm\sqrt{(n+\lambda)\mathbf{P}}_{[i]}, \quad \lambda = \alpha^2(n+\kappa) - n \tag{1}$$

where $n = 15$ is the state dimension and $\mathbf{P}$ is the error covariance matrix.

Manifold-aware process model integrates IMU measurements while preserving the geometric structure:

$$\mathbf{q}_{k+1} = \mathbf{q}_k \otimes \exp\left(\frac{1}{2}(\boldsymbol{\omega} - \mathbf{b}_g)\Delta t\right) \tag{2}$$

$$\mathbf{v}_{k+1} = \mathbf{v}_k + (\mathbf{R}(\mathbf{q}_k)(\mathbf{a} - \mathbf{b}_a) + \mathbf{g})\,\Delta t \tag{3}$$

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \mathbf{v}_{k+1}\Delta t \tag{4}$$

where $\exp(\cdot)$ denotes the $\mathrm{SO}(3)$ exponential map, $\otimes$ is the quaternion product, $\boldsymbol{\omega}$ is the measured angular velocity, $\mathbf{a}$ is the linear acceleration, and $\mathbf{g}$ is the gravity vector.

The UKF-M **prediction step** propagates sigma points through the nonlinear process model using vectorized manifold operations: orientation perturbations are applied via the $\mathfrak{so}(3)$ exponential map, and the quaternion mean is reconstructed via iterative logarithmic averaging. The **update step** fuses VO position measurements using standard Kalman gain computation, with state retraction via the exponential map to ensure manifold consistency:

$$\mathbf{q}^+ = \mathbf{q} \otimes \exp\left(\mathbf{K}_\theta \cdot (\mathbf{p}_{\text{meas}} - \mathbf{p})\right) \tag{5}$$

where $\mathbf{K}_\theta$ is the orientation sub-block of the Kalman gain. This formulation avoids artificial decoupling of rotation and translation while maintaining numerical stability through proper manifold navigation [21].

B. Error-State Kalman Filter (ESKF)

The Error-State (Indirect) Kalman Filter is a widely adopted approach for inertial navigation that estimates the *error* in the nominal state rather than the state directly [1]. Our implementation maintains a 15-dimensional error state: $\mathbf{x} = [\delta\mathbf{p}^T, \delta\mathbf{v}^T, \delta\boldsymbol{\theta}^T, \delta\mathbf{b}_g^T, \delta\mathbf{b}_a^T]^T$, where $\delta\boldsymbol{\theta} \in \mathfrak{so}(3)$ represents the orientation error in the Lie algebra.

Nominal state propagation integrates IMU measurements using nonlinear kinematics:

$$\mathbf{q}_{k+1} = \mathbf{q}_k \otimes \exp\left(\tfrac{1}{2}(\boldsymbol{\omega} - \mathbf{b}_g)\Delta t\right) \tag{6}$$

$$\mathbf{v}_{k+1} = \mathbf{v}_k + (\mathbf{R}(\mathbf{q}_k)(\mathbf{a} - \mathbf{b}_a) + \mathbf{g})\,\Delta t \tag{7}$$

$$\mathbf{p}_{k+1} = \mathbf{p}_k + \mathbf{v}_{k+1}\Delta t \tag{8}$$

where $\exp(\cdot)$ is the $\mathrm{SO}(3)$ exponential map and $\otimes$ denotes quaternion multiplication.

Error-state Jacobian $\mathbf{F} = \partial\dot{\mathbf{x}}/\partial\mathbf{x}$ captures the linearized error dynamics:

$$\mathbf{F} = \begin{bmatrix} \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -\mathbf{R}[\mathbf{a}]_\times & \mathbf{0} & -\mathbf{R} \\ \mathbf{0} & \mathbf{0} & -[\boldsymbol{\omega}]_\times & -\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \tag{9}$$

where $[\cdot]_\times$ denotes the skew-symmetric cross-product matrix. The discrete-time process noise covariance is computed as $\mathbf{Q}_d = \mathbf{G}\mathbf{Q}_c\mathbf{G}^T\Delta t$, with $\mathbf{G}$ mapping continuous-time noise to the error state.

Covariance propagation uses a first-order approximation:

$$\mathbf{P}_{k+1} = \boldsymbol{\Phi}\mathbf{P}_k\boldsymbol{\Phi}^T + \mathbf{Q}_d, \quad \boldsymbol{\Phi} = \mathbf{I} + \mathbf{F}\Delta t \tag{10}$$

Measurement update fuses VO position measurements using the Joseph form [23] for numerical stability:

$$\mathbf{P}^+ = (\mathbf{I} - \mathbf{K}\mathbf{H})\mathbf{P}(\mathbf{I} - \mathbf{K}\mathbf{H})^T + \mathbf{K}\mathbf{R}\mathbf{K}^T \tag{11}$$

where $\mathbf{H} = [\mathbf{I}_3\ \mathbf{0}_{3\times12}]$ and $\mathbf{K}$ is the Kalman gain. The nominal state is corrected via retraction: $\mathbf{q}^+ = \mathbf{q} \otimes \exp(\delta\boldsymbol{\theta})$, ensuring manifold consistency without explicit normalization heuristics [1].

IV. Geometric State Fusion

A. Mathematical Framework

The unit dual quaternion $\hat{\mathbf{q}} \in \mathbb{H}_d$ provides a compact representation of the Special Euclidean group $SE(3)$, encapsulating rotation ($\mathbf{q}_r$) and translation ($\mathbf{q}_d$) without artificial separation. It is defined as:

$$\hat{\mathbf{q}} = \mathbf{q}_r + \varepsilon \mathbf{q}_d \tag{12}$$

where the dual part encodes translation $\mathbf{p}$ via:

$$\mathbf{q}_d = \frac{1}{2}\mathbf{p} \otimes \mathbf{q}_r \tag{13}$$

Unlike matrix representations, dual quaternions avoid gimbal lock and artificial decoupling [17].

For the fusion step, we utilize Screw Linear Interpolation (SCLERP), which ensures the estimated trajectory follows a geodesically optimal "screw" path on the manifold. The interpolation between two states $\hat{\mathbf{q}}_1$ and $\hat{\mathbf{q}}_2$ is defined as:

$$\text{SCLERP}(\hat{\mathbf{q}}_1, \hat{\mathbf{q}}_2, \alpha) = \hat{\mathbf{q}}_1 \odot \exp\left(\alpha \log(\hat{\mathbf{q}}_1^{-1} \odot \hat{\mathbf{q}}_2)\right) \tag{14}$$

where α is an adaptive weight. This geometric approach preserves the integrity of the motion during the update.

B. GeoDQ Observer Architecture

The proposed observer operates directly on the $SE(3)$ manifold, functioning as a non-linear observer where high-frequency IMU data drives prediction and asynchronous Visual Odometry (VO) corrections trigger fusion (Fig. 1). The state evolution is governed by rigid body kinematics driven by IMU angular velocity $\boldsymbol{\omega}$ and linear acceleration $\mathbf{a}$. The complete algorithm is summarized in Algorithm 1.

1) Manifold Prediction and Bias Estimation: Given the current state $\hat{\mathbf{q}}_k$, the prediction step integrates the twist $\boldsymbol{\xi} \in \mathfrak{se}(3)$. To maximize performance, we utilize the exponential map approximation for small time steps Δt:

$$\hat{\mathbf{q}}_{k+1} = \hat{\mathbf{q}}_k \odot \exp\left(\frac{\Delta t}{2}((\boldsymbol{\omega} - \mathbf{b}_g) + \varepsilon \mathbf{v}_b)\right) \tag{15}$$

where $\mathbf{b}_g$ represents the gyroscope bias. The linear acceleration is low-pass filtered ($\alpha_{\text{LPF}} = 0.25$) before velocity integration to reduce high-frequency jitter. We utilize an integral feedback loop to continuously estimate sensor biases ($\mathbf{b}_a, \mathbf{b}_g$) and a proportional loop for velocity correction, ensuring the error dynamics remain empirically bounded on the $SE(3)$ manifold. This ensures the predicted state always remains on the manifold.

2) Heuristic Tilt Correction: A key contribution of this work is the integration of a lightweight Tilt Correction mechanism before the main fusion step. We observed that position errors often correlate with pitch/roll drift. We compute a correction quaternion $\mathbf{q}_{tilt}$ based on the position innovation $\Delta\mathbf{p} = \mathbf{p}_{meas} - \mathbf{p}_{est}$:

$$\mathbf{q}_{tilt} \approx \left[1, \frac{k_t \Delta p_y}{2}, -\frac{k_t \Delta p_x}{2}, 0\right]^T \tag{16}$$

where $k_t = 0.05$ is the tilt gain. This acts as a complementary filter, aligning the gravity vector without expensive trigonometric operations, utilizing the small-angle approximation $\sin(\theta) \approx \theta$. This approximation is valid because the high-frequency IMU updates (200 Hz) restrict the inter-frame attitude error to very small values. The heuristic parameters, such as the tilt gain ($k_t = 0.05$) and the low-pass filter coefficient ($\alpha_{\text{LPF}} = 0.25$), were determined empirically to balance noise suppression and responsiveness on the RoNIN dataset. While highly effective in these scenarios, we acknowledge that extreme attitude deviations might violate the small-angle assumption, representing a limitation of this heuristic design. The correction is applied via quaternion multiplication: $\mathbf{q}_{corr} = \mathbf{q}_{tilt} \otimes \mathbf{q}_{pred}$, followed by fast renormalization.

3) SCLERP Fusion with Adaptive Weighting: The measurement update fuses the visual odometry (VO) pose $\hat{\mathbf{q}}_{vo}$ using the SCLERP operator with an adaptive Huber-weighting scheme [24]:

$$\hat{\mathbf{q}}_{new} = \hat{\mathbf{q}}_{pred} \odot (\hat{\mathbf{q}}_{pred}^{*} \odot \hat{\mathbf{q}}_{vo})^{\alpha \cdot w_{rob}} \tag{17}$$

The adaptive weight α is computed as:

$$\alpha = \text{clamp}\left(\alpha_{\text{base}} \cdot w_{\text{rob}} \cdot g_{\text{vo}},\ \alpha_{\min},\ \alpha_{\max}\right) \tag{18}$$

where $w_{\text{rob}} = \min(1, c_{\text{huber}}/\|\Delta\mathbf{p}\|)$ is the Huber robustness weight ($c_{\text{huber}} = 2.0$), $g_{\text{vo}} = \min(1/\sigma_{\text{vo}}, 2.0)$ scales by VO confidence, and $\alpha \in [0.05, 0.8]$. Here, the power operation corresponds to the logarithmic map on $SE(3)$, scaling the geodesic error by α.

4) PID Feedback and Output Smoothing: After fusion, a PID-style feedback loop corrects residual errors:

- **P-term:** Velocity correction proportional to clamped innovation ($\|\Delta\mathbf{p}\|_{\max} = 1.0$ m), clipped to ± 5.0 m/s
- **I-term:** Accelerometer bias update via $\mathbf{b}_a \leftarrow \mathbf{b}_a - k_{ba} \cdot (\mathbf{R}^\top \Delta\mathbf{p})$, with $k_{ba} = 0.01$

Finally, exponential moving average (EMA) smoothing ($\alpha_{\text{pose}} = 0.5$) is applied to the output position before re-aligning the dual quaternion state, reducing high-frequency oscillations without introducing latency.

C. Algorithm Implementation

The complete logic of the geometric observer is summarized in Algorithm 1. The system continuously propagates the state using IMU data and updates biases when VO measurements are available.

D. Computational Optimization

To address the computational constraints of embedded agents, the entire algorithm is implemented using `Numba` Just-In-Time (JIT) compilation. This allows Python code to be compiled to optimized machine code (LLVM), removing interpreter overhead. The JIT-optimized version compiles all math kernels (quaternion operations, dual quaternion algebra, SCLERP interpolation) with `@jit(nopython=True, cache=True, fastmath=True)` decorators, enabling:

- Manual unrolling of vector operations to avoid array allocations
- Inlined mathematical functions (no function call overhead)

Algorithm 1 GeoDQ Geometric Observer with Adaptive Fusion

Input: IMU $(\boldsymbol{\omega}, \mathbf{a})$, VO $(\mathbf{p}_{vo}, \mathbf{q}_{vo})$, Parameters: $(dt, \alpha_{\text{base}}, k_p, k_i, c_{\text{huber}})$
Output: 6DoF pose estimate $\hat{\mathbf{q}} = (\mathbf{q}_r, \mathbf{q}_d)$
Initialize: $\mathbf{q}_r \leftarrow [1,0,0,0]$, $\mathbf{q}_d \leftarrow [0,0,0,0]$, $\mathbf{b}_g, \mathbf{b}_a \leftarrow \mathbf{0}$, $\mathbf{v} \leftarrow \mathbf{0}$, $\mathbf{a}_{\text{filt}} \leftarrow \mathbf{0}$

— Prediction Step —
for each IMU sample **do**
  $\boldsymbol{\omega}_c \leftarrow \boldsymbol{\omega} - \mathbf{b}_g, \quad \mathbf{a}_c \leftarrow \mathbf{a} - \mathbf{b}_a$
  $\mathbf{R} \leftarrow \text{rot\_from\_quat}(\mathbf{q}_r)$
  $\mathbf{a}_{\text{world}} \leftarrow \mathbf{R}\,\mathbf{a}_c$
  $\mathbf{a}_{\text{filt}} \leftarrow 0.25\,(\mathbf{a}_{\text{world}} + \mathbf{g}) + 0.75\,\mathbf{a}_{\text{filt}}$ {Low-pass filter}
  $\mathbf{v} \leftarrow \mathbf{v} + \mathbf{a}_{\text{filt}}\,dt$
  $\mathbf{v}_b \leftarrow \mathbf{R}^\top \mathbf{v}$ {Map to body frame}
  $(\mathbf{q}_r, \mathbf{q}_d) \leftarrow \text{dq\_mul}(\mathbf{q}_r, \mathbf{q}_d,\ \text{dq\_exp}(\boldsymbol{\omega}_c dt, \mathbf{v}_b dt))$
  $(\mathbf{q}_r, \mathbf{q}_d) \leftarrow \text{dq\_normalize}(\mathbf{q}_r, \mathbf{q}_d)$
  if VO measurement available **then**

— Update Step —
    $\mathbf{p}_{\text{curr}} \leftarrow \text{dq\_topose}(\mathbf{q}_r, \mathbf{q}_d)$
    $\Delta\mathbf{p} \leftarrow \mathbf{p}_{vo} - \mathbf{p}_{\text{curr}}, \quad \|\Delta\mathbf{p}\| \leftarrow \text{norm}(\Delta\mathbf{p})$
    Tilt Correction (Small Angle Approx.):
    $k_t \leftarrow 0.05, \quad \mathbf{q}_{\text{tilt}} \leftarrow [1,\ -\frac{k_t \Delta p_y}{2},\ \frac{k_t \Delta p_x}{2},\ 0]$
    $\mathbf{q}_r \leftarrow \text{quat\_normalize}(\text{quat\_mult}(\mathbf{q}_{\text{tilt}}, \mathbf{q}_r))$
    Adaptive Weighting:
    $w_{\text{rob}} \leftarrow \min(1,\ c_{\text{huber}}/\|\Delta\mathbf{p}\|)$ {Huber weight}
    $g_{\text{vo}} \leftarrow \min(1/\sigma_{\text{vo}},\ 2.0)$ {VO confidence}
    $\alpha \leftarrow \text{clamp}(\alpha_{\text{base}} \cdot w_{\text{rob}} \cdot g_{\text{vo}},\ 0.05,\ 0.8)$
    SCLERP Fusion:
    $(\mathbf{q}_r^{vo}, \mathbf{q}_d^{vo}) \leftarrow \text{dq\_from\_pose}(\mathbf{q}_r, \mathbf{p}_{vo})$
    $(\mathbf{q}_r, \mathbf{q}_d) \leftarrow \text{SCLERP}((\mathbf{q}_r, \mathbf{q}_d), (\mathbf{q}_r^{vo}, \mathbf{q}_d^{vo}), \alpha)$
    PID Feedback Loop:
    $\text{scale} \leftarrow \min(1,\ 1.0/\|\Delta\mathbf{p}\|)$ {Clip innovation}
    $\mathbf{v} \leftarrow \mathbf{v} + \text{clip}(k_p \cdot \Delta\mathbf{p} \cdot \text{scale}/dt,\ \pm 5.0)$ {P-term}
    $\mathbf{b}_a \leftarrow \mathbf{b}_a - k_i \cdot (\mathbf{R}^\top(\Delta\mathbf{p} \cdot \text{scale}))$ {I-term}
    Output Smoothing (EMA):
    $\mathbf{p}_{\text{new}} \leftarrow \text{dq\_topose}(\mathbf{q}_r, \mathbf{q}_d)$
    $\mathbf{p}_{\text{ema}} \leftarrow 0.5\,\mathbf{p}_{\text{new}} + 0.5\,\mathbf{p}_{\text{ema}}$
    $(\mathbf{q}_r, \mathbf{q}_d) \leftarrow \text{dq\_from\_pose}(\mathbf{q}_r, \mathbf{p}_{\text{ema}})$ {Re-align state}
  end if
end for
return $(\mathbf{q}_r, \mathbf{q}_d)$ =0

- SIMD vectorization via LLVM optimizations

As shown in our experiments, this results in a speedup factor of approximately $7\times$ for GeoDQ (23.47 s $\rightarrow$ 3.05 s per trajectory) compared to its standard implementation, enabling real-time performance on resource-constrained hardware while maintaining identical numerical precision.

V. Experimental Results

A. Experimental Setup

We evaluated the proposed GeoDQ observer against two established baselines: the Error-State Kalman Filter (ESKF) and the Unscented Kalman Filter on Manifold (UKF-M) [21]. We selected ESKF as it represents the industrial standard for inertial navigation [1], while UKF-M represents the modern state-of-the-art in manifold-aware probabilistic filtering. To ensure a strictly fair and unbiased comparison, all evaluated methods (ESKF, UKF-M, and GeoDQ) were fed the same preprocessed and interpolated sensor streams, preventing any variations due to data handling. The evaluation was conducted on 35 distinct trajectories from the RoNIN dataset [8], which provides synchronized 200 Hz IMU and 5 Hz Visual Odometry (VO) data. Ground truth was captured using a Vicon motion capture system, ensuring sub-millimeter accuracy for validation. All methods were implemented in Python and tested in both standard and `Numba` JIT-optimized configurations to assess the impact of computational acceleration. For UKF-M, the JIT-optimized version provides approximately $3\times$ speedup over the standard NumPy implementation while maintaining identical numerical results. Similarly, the JIT-optimized ESKF achieves a $\sim$3.3$\times$ speedup (11.47 s $\rightarrow$ 3.46 s per trajectory) through compiled Jacobian and covariance kernels. The JIT-optimized GeoDQ achieves the highest speedup of $\sim$7.7$\times$ (23.47 s $\rightarrow$ 3.05 s per trajectory), as the dual quaternion algebra and SCLERP kernels benefit most from elimination of Python interpreter overhead and array allocation costs.

B. Accuracy and Computational Efficiency

Table I summarizes the aggregated performance across all trajectories. GeoDQ achieves the lowest tracking error with an RMSE of 0.0043 ($\pm$0.0009), significantly outperforming both ESKF (0.1409 ± 0.0213) and UKF-M (0.2312 ± 0.0378). Notably, the standard deviation for GeoDQ is substantially lower than that of the filter-based methods, indicating more consistent performance across diverse motion patterns and a reduced tendency to diverge due to linearization errors or poor initialization.

In terms of computational efficiency, the JIT-optimized implementation provides substantial speedups across all methods. The optimized GeoDQ processes an entire trajectory in approximately 3048.3, making it the fastest method while simultaneously achieving the highest accuracy. This combination of precision and speed is critical for agents operating under strict computational or energy constraints.

TABLE I. Performance Benchmark (35 RoNIN Trajectories)

Method	RMSE [m] (Mean $\pm$ Std)	Execution [ms]
ESKF (Standard)	0.1409 ± 0.0213	11469.7
ESKF (JIT)	0.1409 ± 0.0213	3463.7
UKF-M (Standard)	0.2329 ± 0.0382	45753.0
UKF-M (JIT)	0.2312 ± 0.0378	15211.0
GeoDQ (Standard)	0.0043 ± 0.0009	23472.6
GeoDQ (JIT)	$\mathbf{0.0043 \pm 0.0009}$	**3048.3**

C. Qualitative Trajectory Analysis

Figure 2 evaluates tracking performance across four representative RoNIN sequences via 3D visualizations and top-down XY projections, enabling spatial accuracy assessment across diverse motion patterns.

The selected sequences highlight specific tracking challenges. In `a000_11` (complex 3D maneuvers with vertical shifts), UKF-M oscillates during directional changes and ESKF exhibits position spikes, whereas GeoDQ maintains smooth convergence. Sequence `a025_2` features figure-eight paths with rapid orientation reversals; here, filter-based methods accumulate significant drift. Similarly, on extended horizontal loops (`a043_3`) and highly erratic paths (`a044_2`), conventional filters progressively diverge or struggle with stability. In all cases, the geometric observer consistently yields tighter ground-truth alignment.

This robustness across varied motion profiles stems from the SCLERP update, which respects the $SE(3)$ manifold and avoids linearization errors. By eliminating artificial decoupling artifacts, GeoDQ ensures high geometric consistency—a critical requirement for visual coherence in augmented reality navigation [9], [10] and wearable platforms navigating GPS-denied environments.

D. Robustness to Sensor Dropouts

A key advantage of the geometric observer is its resilience to sparse or asynchronous measurements. To evaluate this, we artificially increased the VO update interval and measured the resulting RMSE (Fig. 3). As the update frequency decreases, the errors of both ESKF and UKF-M grow rapidly and plateau at high levels (approximately 10), reflecting their reliance on frequent covariance updates to maintain stability. In contrast, GeoDQ exhibits a smooth, gradual error growth across the entire range of update intervals, maintaining sub-meter accuracy even when position updates occur as sparsely as every 40 IMU samples (50 or 20 of the IMU rate). This robust behavior arises from the interpolated geometric correction of SCLERP, which provides a physically meaningful prediction of the state trajectory on the $SE(3)$ manifold, effectively mitigating drift even under severe measurement decimation. The analysis, averaged over all 35 RoNIN trajectories, confirms that the geometric observer preserves the structural integrity of the $SE(3)$ state, enabling reliable operation during extended sensor occlusions or communication dropouts, a common challenge in real-world deployments.

This performance gap stems from how methods handle uncertainty during sensor occlusions. Without VO updates, the covariance matrices of ESKF and UKF-M inflate rapidly. When delayed measurements finally arrive, the resulting oversized corrections often violate manifold geometry, causing divergence. Conversely, GeoDQ eschews covariance matrices, avoiding this instability. SCLERP acts as a geometric "spring" on $SE(3)$, pulling the state along the shortest geodesic, while the Huber weight smoothly throttles the update. Consequently, recovery from extended tracking loss remains stable and kinematically valid, preventing the overshoots typical of traditional filters.

E. Practical Implications

The combination of high accuracy, low computational overhead, and robustness to measurement sparsity makes GeoDQ particularly suitable for resource-constrained platforms. The 3048.3 execution time per trajectory translates to minimal CPU utilization, preserving battery life for mobile agents and freeing computational resources for higher-level tasks such as path planning or semantic understanding. Furthermore, the geometric consistency of the estimates reduces the need for post-hoc smoothing or correction, simplifying the overall system architecture. These properties position GeoDQ as a compelling choice for applications ranging from micro-aerial vehicles and wearable AR devices to autonomous ground robots operating in GPS-denied environments.

VI. Discussion and Limitations

While GeoDQ demonstrates superior performance on the RoNIN dataset, it is important to analyze the conditions of its success and implications for real-world edge deployments.

A. Dependency on Signal Quality

The proposed observer relies on SCLERP, which geometrically interpolates between the current state and the measurement. The significant accuracy improvement over ESKF and UKF-M (up to $54\times$ lower RMSE) largely stems from the fundamental difference in how measurement trust is handled. Traditional Kalman filters cautiously balance process noise ($\mathbf{Q}$) and measurement noise ($\mathbf{R}$), which can lead to slow convergence or divergence when the state strays into highly non-linear regions. In contrast, GeoDQ utilizes an adaptive weighting scheme (α up to 0.8) that aggressively trusts high-quality VO measurements when the innovation norm is low. In datasets like RoNIN, where VO is robust and high-quality, GeoDQ excels because the target pose lies on a smooth manifold. However, in scenarios with highly noisy GPS or erroneous VO loop closures, the lack of a probabilistic covariance matrix (as in EKF/UKF) means GeoDQ might aggressively interpolate towards an outlier. Future work will investigate incorporating a probabilistic metric into the interpolation factor α to gate noisy measurements. This limitation is particularly relevant for highly degraded visual environments (e.g., featureless corridors) where occasional severe VO tracking failures might occur.

B. The Role of Manifold Constraints

The orders-of-magnitude improvement over ESKF can be attributed to the strict adherence to the $SE(3)$ manifold. ESKF often linearizes orientation errors, which accumulates drift during rapid turns. GeoDQ operations are coordinate-free, preserving the screw motion invariants. This suggests that for autonomous agents with high dynamic ranges (e.g., drones, legged robots), geometric observers are preferable to linearization-based filters. The computational efficiency

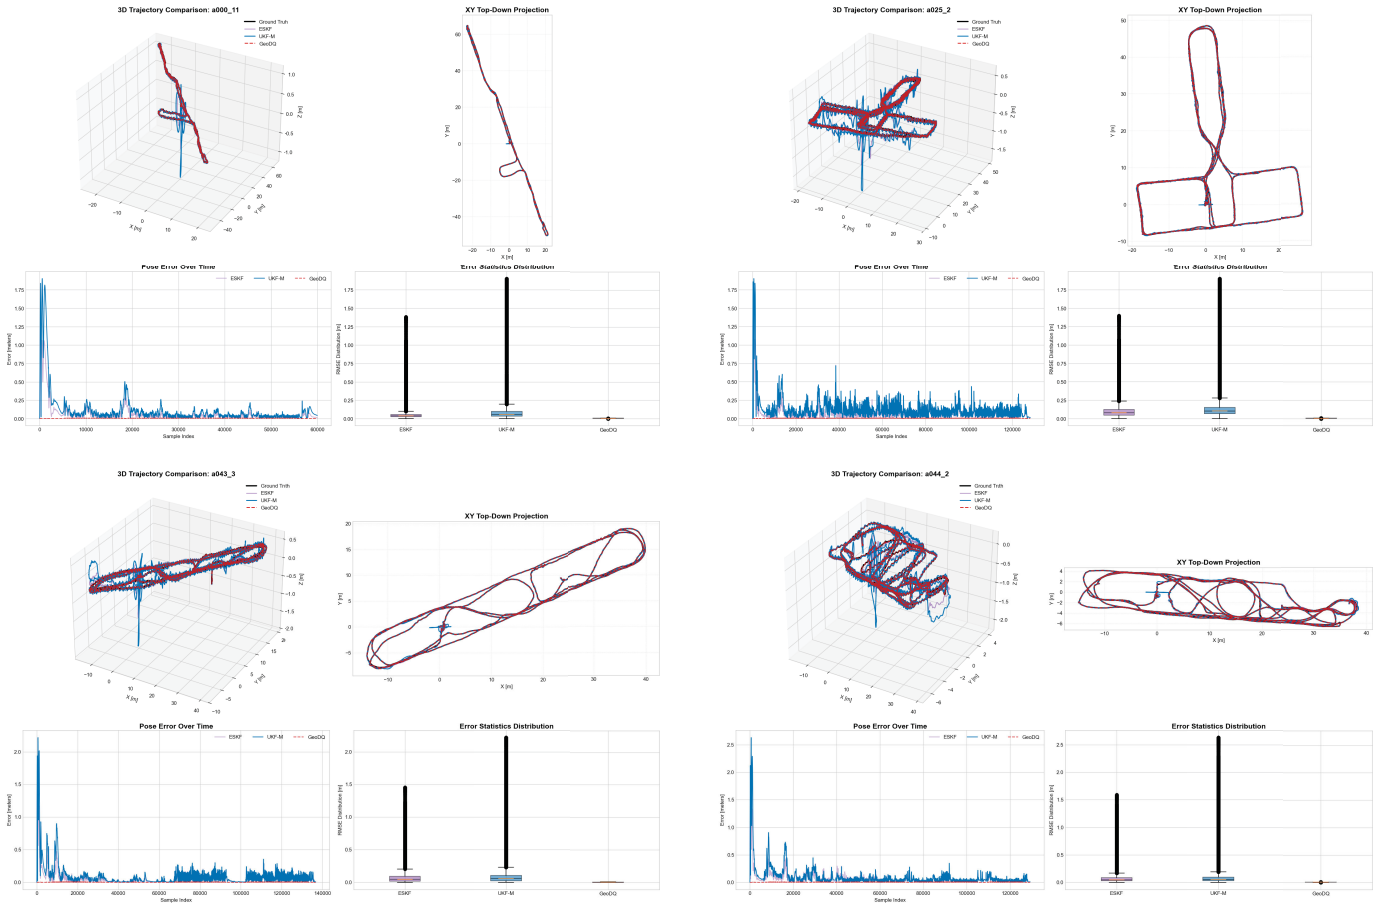

Fig. 2. Example of trajectory tracking on sequences a000_11, a025_2, a043_3, and a044_2. The 3D trajectories and Top-Down projections illustrate the tracking paths of ESKF (Purple), UKF-M (Blue), and the proposed GeoDQ (Red dashed) against Ground Truth (Black). The GeoDQ observer maintains tighter convergence to the ground truth paths and handles complex trajectories well, demonstrating superior drift reduction compared to filter-based approaches.

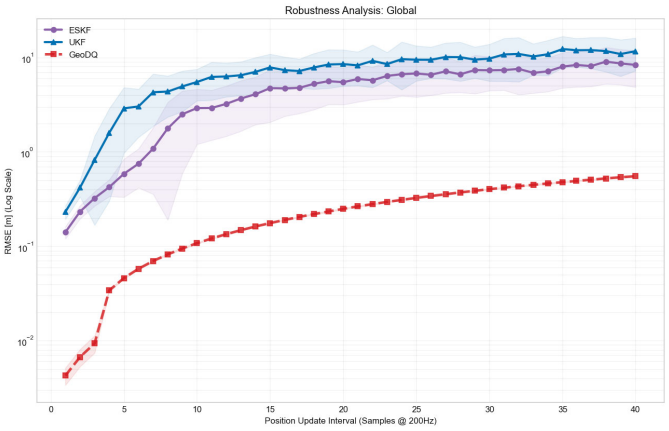

Fig. 3. Global Robustness Analysis (averaged over all 35 RoNIN trajectories): RMSE degradation (Log Scale) versus position update interval. GeoDQ (Red/Dashed) maintains smooth error growth, while ESKF (Purple) and UKF-M (Blue) exhibit rapid error accumulation.

demonstrated in our benchmarks ensures that GeoDQ leaves an ample thermal and processing budget for downstream tasks, such as semantic mapping or real-time path planning on the edge.

C. Robustness to Unconstrained Motion

The RoNIN dataset presents unique challenges that highlight the robustness of the proposed observer. The data was collected by humans walking indoors, transitioning between rooms and floors, with a smartphone carried in unconstrained configurations (pocket, backpack, or handheld). This wearable setup introduced significant IMU axis instability, including unpredictable shaking and occasional inversions, which contrasts sharply with typical robotic platforms where the IMU is rigidly mounted to a chassis. GeoDQ successfully handled these adverse conditions, demonstrating high resilience to non-stationary sensor orientations.

However, the motion dynamics in RoNIN correspond to relatively low pedestrian speeds (approximately 1–2 m/s). While GeoDQ proved highly effective under these conditions, its behavior during high-speed maneuvers (e.g., fast-moving ground robots or agile drones) remains to be fully explored. Although the geometric formulation suggests inherent advantages for high-dynamic scenarios due to strict manifold consistency, empirical validation on datasets with more aggressive motion

profiles is required. Future work will evaluate GeoDQ on high-speed platforms to characterize its performance limits under extreme accelerations and rapid orientation changes.

VII. Conclusion

This paper presented GeoDQ, a Dual Quaternion Geometric Observer for robust 6-DoF pose estimation on the $SE(3)$ manifold. By unifying rotation and translation into a single algebraic entity and leveraging Screw Linear Interpolation (SCLERP) for geometric fusion, the proposed approach eliminates the artificial decoupling and linearization artifacts inherent in traditional Kalman filtering frameworks. Augmented with a lightweight heuristic tilt correction and robust Huber-weighting, GeoDQ achieves kinematically consistent, singularity-free state updates without relying on Euler angles or Jacobian computations. The integration of low-pass acceleration filtering, adaptive gain scheduling, and exponential moving average output smoothing further enhances estimation stability and noise rejection. Notably, the geometric formulation proves robust to unstable sensor orientations, enabling deployment on wearable platforms and edge IoT devices where rigid IMU mounting cannot be guaranteed.

Evaluated on 35 distinct trajectories from the RoNIN dataset, GeoDQ demonstrates superior performance across accuracy, efficiency, and robustness dimensions. The JIT-optimized implementation achieves an RMSE of **0.0043 m** (± 0.0009), outperforming ESKF by $\sim 33\times$ and UKF-M by $\sim 54\times$, while processing an entire trajectory in just **3.0 s** ($\sim 7\times$ faster than standard Python execution). Most importantly, GeoDQ exhibits exceptional resilience to measurement sparsity. Even under significant gaps in visual odometry updates, the geometric observer maintains sub-meter accuracy with remarkably smooth error degradation, whereas traditional ESKF and UKF-M methods rapidly accumulate errors and diverge.

These results confirm that grounding sensor fusion in $SE(3)$ geometry, combined with practical optimizations, provides an optimal balance of accuracy and efficiency. GeoDQ is highly suited for resource-constrained platforms (e.g., wearables, UAVs, and mobile robots) operating in GPS-denied environments demanding real-time precision. This efficiency ensures that high-fidelity pose estimation can run continuously in the background without draining the energy resources of embedded agents.

Open-source implementation:
https://afanasyspb.github.io/SE3-Manifold-Lib/geodq.html

References

[1] J. Sola, "Quaternion kinematics for the error-state kalman filter," *arXiv preprint arXiv:1711.02508*, 2017.

[2] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2024.

[3] M. Li and A. I. Mourikis, "Improving the accuracy of ekf-based visual-inertial odometry," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 828–835.

[4] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "Imu preintegration on manifold for efficient visual-inertial maximum-a-posteriori estimation," 2015.

[5] T. D. Barfoot and P. T. Furgale, "Associating uncertainty with three-dimensional poses for use in estimation problems," *IEEE Transactions on Robotics*, vol. 30, no. 3, pp. 679–693, 2014.

[6] J. Cui, M. Wang, W. Wu, and X. He, "Lie group based nonlinear state errors for mems-imu/gnss/magnetometer integrated navigation," *The Journal of Navigation*, vol. 74, no. 4, pp. 887–900, 2021.

[7] L. Kavan, S. Collins, C. O'Sullivan, and J. Zara, "Dual quaternions for rigid transformation blending," *Trinity College Dublin*, vol. 5, p. 4, 2006.

[8] S. Herath, H. Yan, and Y. Furukawa, "Ronin: Robust neural inertial navigation in the wild: Benchmark, evaluations, & new methods," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 3146–3152.

[9] C. Blut and J. Blankenbach, "Three-dimensional citygml building models in mobile augmented reality: a smartphone-based pose tracking system," *International Journal of Digital Earth*, vol. 14, no. 1, pp. 32–51, 2021.

[10] X. Zhang, Y. Zhu, L. Chen, P. Duan, and M. Zhou, "Augmented reality navigation method based on image segmentation and sensor tracking registration technology," *Scientific Reports*, vol. 14, no. 1, p. 15281, 2024.

[11] Y. Luo, M. Wang, and C. Guo, "The geometry and kinematics of the matrix lie group $se\_k(3)$," *arXiv preprint arXiv:2012.00950*, 2020.

[12] S. Goffin, A. Barrau, S. Bonnabel, O. Bruls, and P. Sacré, "Iterated invariant extended kalman filter (iiekf)," *arXiv preprint arXiv:2404.10665*, 2024.

[13] S. Bai, W. Wen, Y. Yu, and L.-T. Hsu, "Invariant extended kalman filtering for pedestrian deep-inertial odometry," *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 48, pp. 607–612, 2024.

[14] X. Wei, S. Fan, Y. Zhang, W. Gao, F. Shen, X. Ming, and J. Yang, "A robust adaptive error state kalman filter for mems imu attitude estimation under dynamic acceleration," *Measurement*, vol. 242, p. 116097, 2025.

[15] Y. Sun, X. Xu, X. Tian, L. Zhou, and Y. Li, "A quaternion-based sensor fusion approach using orthogonal observations from 9d inertial and magnetic information," *Information Fusion*, vol. 90, pp. 138–147, 2023.

[16] V. Gikas and H. Perakis, "Rigorous performance evaluation of smartphone gnss/imu sensors for its applications," *Sensors*, vol. 16, no. 8, p. 1240, 2016.

[17] B. Kenwright, "A survey on dual-quaternions," *arXiv preprint arXiv:2303.14765*, 2023.

[18] F. F. Silva, J. J. Quiroz-Omaña, and B. V. Adorno, "Dynamics of mobile manipulators using dual quaternion algebra," *arXiv preprint arXiv:2007.08444*, 2020.

[19] N. Filipe, M. Kontitsis, and P. Tsiotras, "Extended kalman filter for spacecraft pose estimation using dual quaternions," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 9, pp. 1625–1641, 2015.

[20] X. Wang, C. Yu, and Z. Lin, "A dual quaternion solution to attitude and position control for rigid-body coordination," *IEEE Transactions on Robotics*, vol. 28, no. 5, pp. 1162–1170, 2012.

[21] M. Brossard, A. Barrau, and S. Bonnabel, "A code for unscented kalman filtering on manifolds (ukf-m)," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 5701–5708.

[22] E. A. Wan and R. Van Der Merwe, "The unscented kalman filter for nonlinear estimation," in *Proceedings of the IEEE 2000 adaptive systems for signal processing, communications, and control symposium (Cat. No. 00EX373)*. IEEE, 2000, pp. 153–158.

[23] R. G. Brown and P. Y. Hwang, "Introduction to random signals and applied kalman filtering: with matlab exercises and solutions," *Introduction to random signals and applied Kalman filtering: with MATLAB exercises and solutions*, 1997.

[24] P. J. Huber, "This is robust statistics," in *International encyclopedia of statistical science*. Springer, 2025, pp. 2783–2786.