



Research article

Adaptive PID-based gain control for robust vision-inertial attitude estimation on SO(3)

Abdelwahhab Khaldi, Choukri Bensalah and Amal Choukchou Braham*

Laboratoire de Tlemcen (LAT), University of Tlemcen, Algeria

* **Correspondence:** Email: amal.choukchou@gmail.com.

Abstract: Existing solutions to attitude estimation for unmanned aerial vehicles (UAVs) in GPS-denied environments face a persistent trade-off: classical filtering approaches such as the extended Kalman filter (EKF) suffer from linearization errors, while geometric nonlinear observers on the special orthogonal (SO(3)) Lie group avoid this issue but rely on a manually tuned, fixed correction gain. A fixed gain cannot simultaneously deliver fast convergence during nominal operation and robust noise rejection when inertial sensors degrade, and existing adaptive or gain-scheduled alternatives either require manual heuristics or lack a real-time responsiveness mechanism, leaving gain tuning an open and recognized limitation in this line of work. This paper proposes a robust vision-inertial fusion method for UAV attitude estimation using a nonlinear observer on the SO(3) Lie group with adaptive gain control that directly addresses this gap. We introduce a PID-based gain adaptation mechanism that dynamically modulates the observer's correction term based on real-time gyroscope variance metrics. When inertial data deteriorates due to noise or bias, the observer automatically relies more on visual corrections; during stable periods, it prioritizes inertial data for faster convergence. MATLAB/Simulink simulation results demonstrate that the adaptive observer reduces attitude estimation error by approximately 65 percent during disturbances and achieves reconvergence over 10 times faster after disturbances compared to fixed-gain baselines, while maintaining bounded errors during severe gyroscope perturbations where fixed gain observers diverge. This work provides a computationally efficient solution to the gain-tuning challenge in geometric attitude estimation for autonomous UAVs operating in GPS-denied environments. Validation is performed through comprehensive simulation studies.

Keywords: SO(3) observer; attitude estimation; proportional-integral-derivative (PID); visual-inertial data fusion; optical flow

1. Introduction

Reliable attitude estimation is essential for the autonomous control of unmanned aerial vehicles (UAVs), especially in complex and uncertain environments. The fusion of inertial measurements from gyroscopes and accelerometers with exteroceptive sensors such as cameras has become fundamental to achieving robust state estimation [1, 2]. This multimodal approach helps mitigate the effects of sensor noise, drift, and temporary loss of individual sensing modalities. The demand for robust sensor fusion methods has intensified particularly in GPS-denied settings such as indoor environments, underground facilities, and urban canyons where satellite-based navigation is unavailable or unreliable [3–5].

Modern UAV applications impose three critical requirements on attitude estimation systems: real-time computational performance, resilience to sensor degradation, and adaptability to dynamically changing environmental conditions [6, 7]. Traditional filtering approaches, such as the extended Kalman filter (EKF), have been widely applied but often suffer from linearization errors and can lose the geometric constraints inherent in rotational dynamics. To address these limitations, recent research has focused on observer designs based on geometric principles, particularly those formulated directly on Lie groups such as $SO(3)$ [8, 9].

Alongside these observer-based approaches, a parallel body of work has pursued UAV attitude *control* (rather than estimation) using PID and PID-derived architectures directly on the rotation group. Goodarzi et al. [10] formulated a geometric nonlinear PID controller on Special Euclidean group ($SE(3)$) that preserves the manifold structure of the configuration space, while Noordin et al. [11] designed a sliding-mode-based auto-tuning adaptive PID (APIDC) for quadrotor attitude and position stabilization under parameter uncertainty and wind disturbance, and Chen et al. [12] used a time-varying-gain extended state observer within the attitude tracking loop itself. These control-oriented studies confirm that gain/parameter adaptation is likewise an active concern on the control side of UAV attitude problems, but they act on the control input given a state estimate, whereas the present work adapts the correction gain of the state *estimator* itself; the two problems are complementary rather than interchangeable.

Nonlinear observers on $SO(3)$ preserve the structure of rotation matrices without suffering from Euler angle singularities or quaternion ambiguities, offering local exponential stability guarantees under appropriate conditions, with global convergence remaining an open problem for most constructions [7, 9]. These geometric approaches have proven especially effective in vision-inertial fusion systems where the coupling between rotational and translational motion must be carefully modeled [13–15]. The seminal work of Rehbinder and Ghosh [1] introduced a geometric observer that combines gyroscope measurements with visual directional cues from 3D line features projected onto the image plane, guaranteeing local exponential convergence under ideal sensor conditions. This foundational approach inspired subsequent research into vision-aided observers on $SO(3)$, often augmented by gain scheduling and adaptive mechanisms [2, 12, 16].

Visual information in these systems often comes from line features found in structured environments like indoor corridors or city facades. Traditional line detection methods, such as the Hough transform [17] and edge-based techniques [7], are commonly used due to their simplicity and reliability under ideal conditions. However, in cluttered or dynamic scenes, these classical approaches struggle with occlusion, lighting variations, or texture loss. To address these challenges, deep learning methods have emerged, including Line Segment Detection Network (LSDNet) [18], DeepLSD [19], and Holistically-

Attracted Wireframe Parsing (HAWP) [20], which offer greater precision and reliability.

For tracking the motion of these line features, optical flow techniques are employed. Traditional methods such as Lucas-Kanade [21], Horn-Schunck [22], and Farnebäck [23] are computationally efficient and perform well under small motion and good illumination. However, their reliability decreases with high-speed motion or in low-texture regions, prompting the adoption of deep learning-based methods like Recurrent All-Pairs Field Transforms for Optical Flow (RAFT) [24] and FlowFormer [25], and learned visual-odometry pipelines such as Depth and Flow for Visual Odometry (DF-VO) [26], which provide dense, accurate optical flow and perform well in challenging visual conditions.

Despite their theoretical elegance and practical success, geometric observers on $SO(3)$ face a persistent challenge: the selection and tuning of correction gains. The correction gain directly governs the trade-off between convergence speed and noise sensitivity. High gains enable rapid convergence and strong disturbance rejection but may amplify measurement noise and cause instability. Conversely, low gains provide smooth filtering but result in slow convergence and inadequate response to sensor degradation. This fundamental trade-off becomes particularly problematic in real-world UAV operations where inertial sensor quality fluctuates due to vibrations, temperature variations, magnetic interference, or mechanical wear. Existing adaptive observers either rely on heuristic gain scheduling that requires manual parameter adjustment or lack real-time responsiveness to changing sensor conditions [2, 12, 16]. Several researchers have acknowledged gain tuning as an unresolved issue requiring future investigation, with some explicitly recommending it as a direction for future work to address performance trade-offs [2, 12, 16]. While state-dependent and time-varying gain mechanisms have been proposed [2, 12, 27], and very recently a fuzzy adaptive-gain complementary filter based on a vector-norm judgment metric was reported for UAV attitude estimation [28], to the best of our knowledge no prior $SO(3)$ -observer work has used gyroscope variance specifically as a direct, real-time feedback metric for adaptive gain control in $SO(3)$ observers. The variance of angular velocity measurements provides an intuitive and computationally efficient indicator of inertial sensor reliability, yet this metric has remained underutilized in geometric observer design, in contrast to the norm-based or covariance-based metrics used in [28] and [16], respectively. Furthermore, the application of classical PID control principles to regulate observer gains on manifolds represents an unexplored intersection between control theory and geometric estimation.

It is worth distinguishing this proposal from the broader landscape of nonlinear/adaptive PID *controller* variants developed outside the attitude-estimation literature, since several such variants could, in principle, be repurposed as alternative gain-adaptation laws. These include the sigmoid PID controller [29], the neuroendocrine-PID controller [30], the brain-emotional-learning-based intelligent controller (BELBIC) combined with PID [31], the Gudermannian-function-based PID controller [32], the softsign fractional-order PID (softsign-FOPID) controller [33], and sliding-mode-based auto-tuning adaptive PID schemes such as the APIDC of [11]. All of these variants share a common design philosophy: they replace the linear error term of a classical PID controller with a bounded, smooth, often biologically- or metaheuristically-inspired nonlinear mapping of the *tracking error* (or its derivatives), and their gains are typically tuned offline by a metaheuristic optimizer (e.g., starfish optimization, sine-cosine algorithm, safe experimentation dynamics) against a fixed reference scenario such as an Automatic Voltage Regulator (AVR) step response or a fixed trajectory. The proposed gain-adaptation law differs from this family in three respects: (i) its input is not a tracking/output

error but an independently computed *sensor-reliability metric* (windowed gyroscope variance), so it does not require a reference trajectory to be defined; (ii) it modulates the correction *gain* of a state observer evolving on the $SO(3)$ manifold rather than the control effort of a plant, and therefore must be shown—as we do in Section 3.4—to preserve the Lyapunov stability properties of the underlying geometric observer, a requirement absent from the AVR/DC-motor/manipulator settings in which the above PID variants were validated; and (iii) its nonlinearity (the PID law of Eq. (3.3)) is fixed by construction rather than being an additional tunable nonlinear mapping, which keeps the number of free parameters at three (K_P, K_I, K_D) rather than introducing further shape parameters as in the Gudermannian or softsign constructions. We view these PID variants as complementary rather than competing: replacing the linear PID law in Eq. (3.3) with one of the bounded nonlinear mappings above is a natural direction for future refinement of the gain-adaptation law itself (see Section 5.9).

This paper addresses the gain-tuning challenge through the following key contributions:

- **Novel adaptive mechanism:** To the best of our knowledge, we introduce the first PID-based adaptive gain control mechanism for $SO(3)$ observers driven specifically by a real-time gyroscope-variance reliability metric, enabling autonomous adjustment of correction gains without manual intervention or a reference trajectory; Section 1 situates this contribution relative to related time-varying-gain observers and to nonlinear PID controller variants from the broader control literature.
- **Theoretical framework:** We establish that local exponential stability of the original Rehbinder-Ghosh observer is preserved under bounded, slowly-varying adaptive gains, extending classical Lyapunov-based proofs to the time-varying gain case.
- **Quantitative simulation evidence:** The theoretical claims above are supported by MATLAB/Simulink simulations incorporating realistic sensor disturbances (bias and high-frequency noise), which show an approximately 65% root mean square error (RMSE) reduction, more than 10× faster reconvergence relative to a fixed-gain baseline, and bounded errors under perturbations that cause the fixed-gain observer to diverge; these results should be read as simulation-based evidence supporting the mechanism above, not as an independent contribution in themselves.
- **Computational efficiency:** The variance-based feedback mechanism requires minimal computational overhead ($\approx 5\%$ increase in execution time, Section 5.4), a property that distinguishes it from covariance-based adaptive filters [16] and learning-based alternatives [34, 35], and makes it suitable for embedded UAV systems with limited processing resources.

The remainder of this paper is structured as follows. Section 2 formulates the attitude estimation problem and reviews the baseline Rehbinder-Ghosh observer on $SO(3)$. Section 3 presents the adaptive gain control design, including the PID-based adaptation law and stability considerations. Section 4 describes the simulation framework and implementation details. Section 5 presents comprehensive results and discusses their implications in comparison with fixed-gain approaches and related work. Section 6 concludes the paper and outlines future research directions.

2. Problem statement

This section establishes the mathematical foundation for UAV attitude estimation using vision-inertial fusion on $SO(3)$. We first present the rigid body attitude dynamics and measurement models, then review the geometric observer proposed by Rehbinder and Ghosh [1], which serves as the baseline for our adaptive extension.

2.1. Notation and preliminaries

Table 1. Summary of mathematical notations.

Symbol	Description	Domain
$R(t)$	True rotation matrix (body to inertial frame)	$SO(3)$
$\hat{R}(t)$	Estimated rotation matrix	$SO(3)$
$\tilde{R}(t)$	Estimation error: $\tilde{R} = \hat{R}^\top R$	$SO(3)$
$\Omega(t)$	Angular velocity in body frame	$\mathbb{R}^3$
$\Omega^\wedge$	Skew-symmetric matrix of Ω	$\mathfrak{so}(3)$
d_i	Known inertial direction vectors	$\mathbb{R}^3, \ d_i\ = 1$
$y_i(t)$	Visual measurement normals	$\mathbb{R}^3, \ y_i\ = 1$
$\sigma(t)$	Visual correction term (skew-symmetric)	$\mathfrak{so}(3)$
$k(t)$	Time-varying observer gain	$\mathbb{R}_+$
$v(t)$	Normalized gyroscope variance metric	$[0, 1]$
K_P, K_I, K_D	PID controller gains	$\mathbb{R}_+$

The special orthogonal group is defined as

$$SO(3) = \left\{ R \in \mathbb{R}^{3 \times 3} \mid R^\top R = I, \det(R) = 1 \right\},$$

which represents the set of all rotation matrices in three dimensions. Its associated Lie algebra,

$$\mathfrak{so}(3) = \left\{ A \in \mathbb{R}^{3 \times 3} \mid A^\top = -A \right\},$$

consists of all 3×3 skew-symmetric matrices. For any angular velocity vector $\omega = [\omega_x, \omega_y, \omega_z]^\top \in \mathbb{R}^3$, the skew-symmetric operator $(\cdot)^\wedge : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ is defined as

$$\omega^\wedge = \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix},$$

as shown in Eq. (2) below.

2.2. Rigid body attitude dynamics

Let $R(t) \in SO(3)$ represent the true rotation matrix describing the UAV's orientation at time t . This matrix maps vectors from the body-fixed frame to the inertial frame. The attitude kinematics evolve according to

$$\dot{R}(t) = R(t)\Omega^\wedge(t), \quad (2.1)$$

where $\Omega(t) = [\omega_x(t), \omega_y(t), \omega_z(t)]^\top \in \mathbb{R}^3$ denotes the angular velocity expressed in the body frame, and the operator $(\cdot)^\wedge : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ maps a vector to its corresponding skew-symmetric matrix. Explicitly,

$$\Omega(t)^\wedge = \begin{bmatrix} 0 & -\omega_z(t) & \omega_y(t) \\ \omega_z(t) & 0 & -\omega_x(t) \\ -\omega_y(t) & \omega_x(t) & 0 \end{bmatrix}. \quad (2.2)$$

The gyroscope provides noisy measurements of the true angular velocity, modeled as

$$\Omega_{\text{measured}}(t) = \Omega(t) + b(t) + \eta(t), \quad (2.3)$$

where $\eta(t)$ denotes zero-mean Gaussian noise and $b(t)$ represents a time-varying bias that may arise from thermal drift, calibration imperfections, or electromagnetic interference.

2.3. Visual measurements from line features

The UAV's monocular camera observes 3D line segments in the environment. These lines are projected onto the image plane, and normal vectors $\mathbf{y}_i(t) \in \mathbb{R}^3$ corresponding to known inertial directions $\mathbf{d}_i \in \mathbb{R}^3$ are estimated using optical flow-based tracking [36]. Specifically, points belonging to detected line segments are tracked over time, and their apparent motion is used to compute projected normal vectors through the following geometric relationship.

Each visual measurement $\mathbf{y}_i(t)$ satisfies the orthogonality constraint:

$$(R(t)\mathbf{d}_i)^\top \cdot \mathbf{y}_i(t) = 0 \quad (2.4)$$

provided that feature tracking and the projection model are accurate. This constraint states that the inertial direction $\mathbf{d}_i$, when rotated into the body frame by $R(t)$, must be perpendicular to the observed normal $\mathbf{y}_i(t)$.

Assumptions: We assume the availability of at least three non-collinear inertial directions $\{\mathbf{d}_1, \mathbf{d}_2, \mathbf{d}_3\}$, which are known and fixed in the inertial frame. These typically correspond to orthogonal line features in structured environments, such as building edges, corridor boundaries, or urban infrastructure elements.

The optical flow computation relies on Lucas-Kanade tracking [21] with a pyramid implementation (3 levels, window size 15×15 pixels) to handle moderate inter-frame motion. For detailed derivations of the visual measurement model and normal vector computation from optical flow, readers are referred to [1, 36].

2.4. Baseline Rehbinder-Ghosh observer

The rigid-body kinematics of Eqs. (2.1)–(2.3) and the visual measurement model of Eq. (2.4) above are standard formulations reproduced from the existing literature (namely [1] and our own prior conference work [36]), rather than models newly derived or experimentally identified in this paper; they are restated here only to keep the paper self-contained and to fix the notation used in Section 3. Section 4 onward describes what *is* original to this work: the adaptive gain law and its simulation-based evaluation. Our approach builds upon the nonlinear observer introduced by Rehbinder and Ghosh [1], which fuses inertial and visual measurements directly on $\text{SO}(3)$ to estimate attitude. The observer dynamics are governed by

$$\dot{\hat{R}}(t) = \hat{R}(t)(\Omega^\wedge(t) - k \cdot \sigma(t)), \quad (2.5)$$

where $\hat{R}(t)$ is the estimated rotation matrix, $k > 0$ is a scalar gain, and $\sigma(t) \in \mathbb{R}^{3 \times 3}$ is the visual correction term calculated as follows:

$$\sigma(t) = \sum_{i=1}^m V_i(t), \quad (2.6)$$

where m is the lines number detected and tracked in the scene, and each individual correction component $V_i(t)$ is defined as

$$V_i(t) = -S(y_i(t) \times \hat{R}(t)d_i)(\hat{R}(t)d_i)^T y_i(t). \quad (2.7)$$

Here, $S(\cdot)$ denotes the skew-symmetric matrix operator. The term $(y_i \times R\mathbf{d}_i)$ measures the angular error between the predicted and observed normal directions, while the quadratic form $(R\mathbf{d}_i)^T y_i$ quantifies the violation of the orthogonality constraint. The correction $\sigma(t)$ is skew-symmetric by construction, ensuring that $R(t)$ remains on $SO(3)$ throughout the evolution.

Stability Properties: Under the assumptions of continuous visual feature availability, accurate inertial sensor readings, and properly selected constant gain k , the observer guarantees local exponential convergence to the true attitude. The proof employs a Lyapunov function

$$V(\tilde{R}) = \frac{1}{2} \text{trace}(I - \tilde{R}) \quad (2.8)$$

based on the estimation error $\tilde{R}(t) = \hat{R}^T(t)R(t)$. Near the identity (small initial error), the time derivative $\dot{V}$ is negative definite, ensuring convergence with a rate proportional to k .

$$V(\tilde{R}) = \frac{1}{2} \text{trace}(I - \tilde{R})$$

based on the estimation error

$$\tilde{R}(t) = \hat{R}^T(t)R(t).$$

Near the identity (small initial error), the time derivative $\dot{V}$ is negative definite, ensuring convergence with rate proportional to k .

Limitations: In practice, these assumptions are frequently violated. Visual features may be intermittently lost due to occlusion or motion blur, and inertial data inevitably suffers from noise and bias. Most critically, the observer's performance is highly sensitive to the choice of gain k . A fixed gain represents a compromise: it cannot simultaneously achieve fast convergence during stable operation and robust noise rejection during sensor degradation. This motivates our adaptive gain approach described in the following section.

3. Adaptive gain control design

This section presents the core contribution of this paper: a PID-based adaptive gain mechanism that dynamically adjusts the observer's correction strength based on real-time inertial sensor reliability. We first introduce the modified observer structure, then derive the adaptation law, and finally discuss stability considerations under time-varying gains.




Figure 1. Control-system block diagram of the adaptive gain-control loop, with input/output variables labeled on each block.

3.1. Modified observer with adaptive gain

The proposed adaptive observer extends Eq. (2.5) by replacing the constant gain k with a time-varying function $k(t)$:

$$\dot{\hat{R}}(t) = \hat{R}(t)(\Omega^\wedge(t) - k(t) \cdot \sigma(t)) \quad (3.1)$$

where $k(t) \in \mathbb{R}_+$ is the time-varying observer gain, regulated by a feedback control mechanism described below.

The key principle is that $k(t)$ should increase when inertial measurements are unreliable (amplifying visual corrections) and decrease when inertial data is trustworthy (allowing faster convergence with minimal visual interference).

3.2. Variance-based feedback metric

To quantify inertial sensor reliability in real time, we employ the variance of angular velocity measurements as a feedback signal. Intuitively, high variance in $\Omega(t)$ indicates the presence of noise, bias drift, or external disturbances, while low variance suggests stable, reliable gyroscope operation.

To reliably characterize the time-varying quality of inertial measurements, the gyroscope noise level is estimated online over a sliding temporal window of duration $T_w = 1.0$ s, following the adaptive noise-estimation strategy commonly adopted in recent navigation systems [37]. Instead of assuming a fixed noise profile, the estimator computes an incremental update of the sample covariance using Welford's numerically stable online algorithm, which maintains constant memory and computation per sample. Given the measurement history $\{\Omega(t - T_w), \dots, \Omega(t)\}$, the instantaneous reliability metric $\text{Var}(\Omega(t))$ is defined as the trace of the 3×3 sample covariance matrix, averaged across the three angular velocity components.

To ensure bounded operation of the PID controller, we normalize this variance into a dimensionless metric:

$$v(t) = \min(1, \text{Var}(\Omega(t))) \quad (3.2)$$

This saturation maps $v(t) \in [0, 1]$, where $v(t) \approx 0$ indicates stable inertial sensing and $v(t) \rightarrow 1$ signals severe degradation. The choice of a unit saturation threshold was determined empirically through preliminary simulations; alternative normalizations (e.g., sensor datasheet-based scaling) could be employed depending on the application.

3.3. PID-based gain adaptation law

The time-varying gain $k(t)$ is generated by a PID controller that tracks the variance metric $v(t)$:

$$k(t) = k_0 \cdot \left(1 + K_P v(t) + K_I \int_0^t v(\tau) d\tau + K_D \frac{dv(t)}{dt} \right) \quad (3.3)$$

where:

- $k_0 > 0$ is the nominal baseline gain, representing the desired gain under ideal conditions ($v \approx 0$),
- $K_P \geq 0$ is the proportional gain, providing immediate response to current variance levels,
- $K_I \geq 0$ is the integral gain, accumulating persistent variance deviations to eliminate steady-state error,
- $K_D \geq 0$ is the derivative gain, anticipating rapid changes in variance (e.g., sudden disturbance onset).

Operational Principle: When the inertial signal is stable ($v \approx 0$), the PID terms vanish and $k(t) \rightarrow k_0$, allowing the observer to trust gyroscope measurements and converge rapidly. When variance increases ($v \rightarrow 1$), the proportional term $K_P v(t)$ immediately raises the gain, strengthening the influence of visual corrections. The integral term addresses sustained degradation (e.g., slow bias drift), while the derivative term enables preemptive adjustment during transient disturbances.

Parameter Selection: The PID gains were tuned empirically through iterative simulation trials. We selected $k_0 = 10$, $K_P = 1.5$, $K_I = 0.01$, and $K_D = 0.05$ to balance responsiveness and stability. The proportional gain K_P dominates the response, providing the primary adaptation mechanism. The integral gain K_I is kept small ($K_I \ll K_P$) to prevent excessive windup during prolonged disturbances. The derivative gain K_D is moderate to smooth transient responses without introducing high-frequency oscillations from noise in dv/dt .

3.4. Stability analysis under adaptive gain

The introduction of the time-varying gain $k(t)$ requires careful stability analysis to ensure that the adaptive observer inherits the convergence properties of the fixed-gain baseline. We now establish conditions under which local exponential stability is preserved.

Proposition 1 (Local Exponential Stability with Adaptive Gain). Consider the adaptive observer (Eq. 8) with gain adaptation law (Eq. 10). Assume the following:

- (A1) The PID controller ensures $k(t)$ remains bounded: $k_{\min} \leq k(t) \leq k_{\max}$ for all $t \geq 0$, with $0 < k_{\min} < k_{\max} < \infty$.
- (A2) The gain varies slowly: $\left| \frac{dk(t)}{dt} \right| \leq L$ for some Lipschitz constant $L > 0$.

(A3) Visual measurements satisfy persistent excitation: at least three non-collinear inertial directions $\{d_1, d_2, d_3\}$ are continuously tracked.

(A4) The initial estimation error is small: $\|\tilde{R}(0) - I\| < \delta$ for sufficiently small $\delta > 0$.

Then the adaptive observer is locally exponentially stable, with the estimation error $\tilde{R}(t)$ converging to the identity at an exponential rate.

Proof Sketch. The proof extends the Lyapunov-based analysis from [1] to the time-varying gain case. Consider the same Lyapunov candidate function:

$$V(\tilde{R}) = \frac{1}{2} \text{trace}(I - \tilde{R}). \quad (3.4)$$

This function is positive definite near the identity and vanishes when $\tilde{R} = I$. Computing its time derivative along the error dynamics and using the skew-symmetry of $\sigma(t)$ yields:

$$\dot{V} = -k(t) \text{trace}(\sigma(\tilde{R})(\tilde{R} - I)). \quad (3.5)$$

Under assumptions (A1)–(A4), the term $\text{trace}(\sigma(\tilde{R})(\tilde{R} - I))$ is positive definite near the identity due to persistent excitation (A3). The boundedness of $k(t)$ (A1) ensures that $\dot{V}$ remains negative definite, while the slow variation (A2) guarantees that the Lyapunov function does not increase due to rapid gain fluctuations. Therefore, $V(t)$ decreases exponentially, implying exponential convergence of $\tilde{R}(t)$ to I .

This theoretical framework, combined with the empirical evidence presented in Section 5, provides strong support for the practical viability of the proposed adaptive observer.

4. Simulation framework and implementation

This section describes the comprehensive simulation environment developed to validate the proposed adaptive observer. We detail the MATLAB/Simulink architecture, parameter settings, sensor models, and evaluation scenarios designed to assess robustness under realistic disturbances.

4.1. Simulation architecture

The simulation platform was implemented in MATLAB R2023a using Simulink for the modular block-based design. Figure 2 illustrates the overall system architecture, which integrates multiple subsystems for sensor simulation, observer computation, and gain adaptation.




Figure 2. Block scheme for SO(3) observer with fixed and adaptive gain.

Figure 2 presents the complete block diagram of the SO(3) observer with both fixed-gain and adaptive-gain configurations. The architecture enables direct performance comparison under identical sensor conditions. The primary components include visual line processing modules, observer blocks, disturbance generators, and the PID-based gain controller, each of which is described in detail below.

Visual Line Modules (y_1, y_2, y_3): Three parallel blocks compute the measured normal vectors for tracked line features in the image plane using simulated optical flow data [36]. Each module processes one inertial direction $\mathbf{d}_i$ and outputs the corresponding visual measurement $\mathbf{y}_i(t)$, computed according to Eq. (2.4). The optical flow is synthesized from known attitude trajectories using the Lucas–Kanade algorithm [21] with pyramid refinement (3 levels, 15×15 pixel windows). This provides realistic feature motion patterns without requiring actual video input, enabling controlled evaluation of observer performance.

Rehbinder–Ghosh Observer (Fixed Gain): One instantiation of the observer implements Eq. (2.5) with constant gain $k = 12$. This value was selected through manual tuning to achieve satisfactory performance under nominal conditions, representing a typical engineering compromise between convergence speed and noise sensitivity. This fixed-gain observer serves as the performance baseline for comparison.

Rehbinder–Ghosh Observer with Adaptive Gain: The second observer implements Eq. (3.1) with PID-controlled gain $k(t)$ from Eq. (3.3). This block receives the variance metric $v(t)$ as an additional input and dynamically adjusts its correction strength in real time. Both observers operate in parallel, processing identical sensor inputs to ensure a fair comparison.

Gyroscope Degradation Generator: A programmable disturbance module injects controlled noise

and bias into the angular velocity signal $\Omega(t)$ to simulate sensor malfunction. The disturbance is activated by a time-based flag and includes both constant bias (simulating thermal drift) and high-frequency additive Gaussian noise (simulating vibration or electromagnetic interference). This enables systematic evaluation of observer robustness under known perturbation profiles.

PID Controller and Variance Estimator: This subsystem computes the real-time variance $\text{Var}(\Omega(t))$ over a sliding 1.0-second window. The normalized metric $v(t) = \min(1, \text{Var}(\Omega(t)))$ is then fed into a discrete-time PID controller implemented using the trapezoidal integration rule with anti-windup. The PID block outputs the adaptive gain $k(t)$, which modulates the observer's correction term. The controller operates at 100 Hz, matching the observer's sampling rate.

4.2. Simulation parameters

Table 2 summarizes the key simulation parameters used in all experiments. These settings were chosen to represent realistic UAV operating conditions while maintaining computational tractability.

Table 2. Simulation parameters.

Parameter	Value / Description
Simulation duration	30 seconds
Time step (sampling rate)	0.01 s (100 Hz)
Initial attitude	Identity matrix / Euler(0, $\frac{\pi}{8}$, 0)
Inertial line directions d_1, d_2, d_3	Orthogonal unit vectors $[1, 0, 0]^T$, $[0, 1, 0]^T$, $[0, 0, 1]^T$
Nominal observer gain (k_0)	10
Fixed baseline gain (k_{fixed})	12
PID gains	$K_P = 1.5$, $K_I = 0.01$, $K_D = 0.05$
Variance window duration (T_w)	1.0 s
Gyroscope noise (nominal)	$\sigma_\eta = 0.01$ rad/s (white Gaussian)
Gyroscope perturbation interval	5–10 seconds
Perturbation: constant bias	$[0.5, 0.3, 0.4]^T$ rad/s
Perturbation: additive noise	$\sigma_\eta = 0.5$ rad/s (white Gaussian)
Optical flow algorithm	Lucas–Kanade pyramidal (3 levels, 15×15 window)
Camera frame rate	30 Hz
Visual measurement delay	0.03 s (1 frame latency)

Synthetic Attitude Trajectory: To generate ground-truth attitude data and corresponding visual measurements, we simulate a smooth periodic motion representative of UAV maneuvering. The true Euler angles (roll ϕ , pitch θ , yaw ψ) evolve as:

$$\begin{aligned}
 \phi(t) &= \frac{\pi}{4} \sin\left(\frac{\pi}{2}t\right), \\
 \theta(t) &= \frac{\pi}{4} \sin\left(\frac{\pi}{2}t + \frac{\pi}{4}\right), \\
 \psi(t) &= \frac{\pi}{4} \sin\left(\frac{\pi}{2}t + \frac{\pi}{2}\right).
 \end{aligned} \tag{4.1}$$

This sinusoidal profile provides continuous excitation across all three axes without singularities,

ensuring persistent observability throughout the simulation. The corresponding rotation matrix $R(t)$ is computed via the standard ZYX Euler angle composition, and angular velocities $\Omega(t)$ are derived analytically by differentiating $R(t)$ according to Eq. (1).

4.3. Perturbation scenarios and performance metrics

Disturbance Protocol: To rigorously evaluate robustness, we inject a severe gyroscope perturbation between $t = 5$ s and $t = 10$ s. This five-second interval contains:

1. **Constant bias:** A fixed offset $\mathbf{b} = [0.5, 0.3, 0.4]^T$ rad/s applied to all three axes, simulating sensor calibration error or thermal drift.
2. **High-frequency noise:** White Gaussian noise with standard deviation $\sigma_\eta = 0.5$ rad/s (50× nominal), representing vibration-induced disturbances or electromagnetic interference.

This combined perturbation significantly degrades inertial measurement quality, challenging the observer's ability to maintain accurate attitude estimates. We analyze the system's response during three phases: pre-disturbance (0–5 s), active disturbance (5–10 s), and recovery (10–30 s).

Quantitative Evaluation Metrics: To enable objective comparison, we compute the following performance metrics for both fixed-gain and adaptive-gain observers:

1. Root Mean Square Error (RMSE):

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{k=1}^N \|\mathbf{e}(t_k)\|_2^2}, \quad (4.2)$$

where $\mathbf{e}(t) = [\phi(t) - \hat{\phi}(t), \theta(t) - \hat{\theta}(t), \psi(t) - \hat{\psi}(t)]^T$ is the Euler angle error vector, and N is the number of samples.

2. Maximum Absolute Error (MAE):

$$\text{MAE} = \max_{1 \leq k \leq N} \|\mathbf{e}(t_k)\|_\infty. \quad (4.3)$$

3. **Recovery Time:** Duration required for the estimation error to return within 5% of pre-disturbance levels after the perturbation ends at $t = 10$ s.
4. **Gain Response Time:** Time delay between disturbance onset and the peak gain value, characterizing the PID controller's responsiveness.

These metrics are computed separately for the disturbance phase (5–10 s) and recovery phase (10–30 s) to isolate transient vs. steady-state behavior.

5. Results and discussion

This section presents comprehensive simulation results comparing the proposed adaptive observer with the classical fixed-gain baseline. We analyze performance under sensor disturbances, examine the gain adaptation mechanism, and discuss implications for practical UAV applications. The results are organized to address three key questions:

1. How does the adaptive gain respond to inertial degradation?
2. What impact does adaptation have on estimation accuracy?
3. How does the proposed method compare to existing approaches?

5.1. Adaptive gain behavior under perturbation

Figure 3 illustrates the evolution of the adaptive gain $k(t)$ over the 30-second simulation. The gain exhibits three distinct phases corresponding to the perturbation protocol described in Section 4.3. Figure 3 shows the time-varying observer gain $k(t)$ generated by the PID controller. The horizontal axis represents time in seconds, while the vertical axis shows the gain magnitude. During the initial stable phase (0–5 s), the gain remains near the nominal value $k_0 = 10$, with minor fluctuations due to residual measurement noise. At $t = 5$ s, when the gyroscope perturbation begins, the variance metric $v(t)$ increases sharply (see Figure 6), causing the PID controller to rapidly elevate the gain. The gain peaks at approximately $k(t) \approx 32$ around $t = 6.5$ s, representing a $3.2\times$ increase relative to the nominal value. This amplification strengthens the visual correction term $\sigma(t)$, effectively compensating for the unreliable inertial data.




Figure 3. Gain of adaptive observer vs. time.

Throughout the disturbance interval (5–10 s), the gain remains elevated, fluctuating between 25 and 35 in response to time-varying noise characteristics. Once the perturbation ends at $t = 10$ s, the variance $v(t)$ decreases rapidly, and the PID controller smoothly reduces the gain back toward k_0 . By $t = 12$ s, the gain has stabilized near its nominal value, demonstrating effective self-regulation. This adaptive behavior contrasts sharply with the fixed-gain observer, which maintains $k = 12$ throughout all phases, regardless of sensor conditions.

Key Observation: The gain adaptation is both responsive and stable. The proportional term K_P provides immediate reaction to variance increases, while the derivative term K_D introduces damping that prevents oscillatory overshoot. The integral term K_I accumulates persistent errors, ensuring that sustained disturbances (e.g., slow bias drift) trigger appropriate gain adjustments. The smooth return to nominal values after disturbance cessation confirms that the PID controller does not exhibit excessive windup or instability.

5.2. Attitude estimation performance: fixed vs. adaptive gain

5.2.1. Fixed-gain observer performance

Figure 4 presents the estimated and true Euler angles for the classical Rehbinder-Ghosh observer with constant gain $k = 12$. The three subplots show pitch θ (top), yaw ψ (middle), and roll ϕ (bottom) over the 30-second simulation. Orange curves represent ground truth, while blue curves show observer estimates.

During the initial stable phase (0–5 s), the fixed-gain observer tracks the true attitude accurately, with estimation errors remaining below 0.05 rad ($\approx 3^\circ$) across all axes. However, when the gyroscope perturbation activates at $t = 5$ s, the estimates begin to diverge noticeably. The pitch and roll estimates exhibit increasing oscillations and steady-state offsets, with peak errors reaching approximately 4 rad ($\approx 230^\circ$) for pitch around $t = 8$ s. These large deviations indicate that the fixed gain $k = 12$ is insufficient to reject the combined effects of bias and high-frequency noise injected into the angular velocity measurements.




Figure 4. Estimated and true yaw, pitch, and roll angles using classical observer.

Most critically, the oscillations persist well beyond the end of the disturbance at $t = 10$ s. Even by $t = 20$ s, the pitch estimate has not fully reconverged to the true trajectory, exhibiting residual errors

of ≈ 0.5 rad. This slow recovery demonstrates the fundamental limitation of fixed-gain designs: low gains provide smooth filtering but cannot adequately respond to severe transient disturbances, resulting in prolonged estimation errors.

Quantitative Analysis (Fixed-Gain):

- RMSE during disturbance (5–10 s): 1.82 rad ($\approx 104^\circ$)
- Maximum absolute error: 4.13 rad ($\approx 237^\circ$) at $t = 8.2$ s
- Recovery time (to 5% of pre-disturbance error): > 15 s (does not fully recover within observation window)

These metrics confirm that the fixed-gain observer experiences severe degradation under the test conditions, failing to maintain bounded errors during the disturbance and exhibiting unacceptably slow reconvergence afterward.

5.2.2. Adaptive-gain observer performance

Figure 5 shows the corresponding results for the proposed adaptive observer. The format is identical to Figure 4, enabling direct visual comparison.

In stark contrast to the fixed-gain case, the adaptive observer successfully maintains tracking throughout the entire simulation. During the disturbance interval (5–10 s), the estimates deviate temporarily from ground truth, with peak errors of approximately 0.8 rad ($\approx 46^\circ$) for pitch at $t = 7$ s. However, these deviations remain bounded and significantly smaller than those of the fixed-gain observer. More importantly, the adaptive observer demonstrates rapid reconvergence: by $t = 11.5$ s—just 1.5 seconds after the disturbance ends—the estimation errors return to pre-disturbance levels (< 0.05 rad).




Figure 5. Estimated and true yaw, pitch, and roll angles using tuned gain observer.

The improved performance is directly attributable to the adaptive gain mechanism. During the high-variance period (5–10 s), the elevated gain $k(t) \approx 30$ amplifies the visual correction term $\sigma(t)$, effectively overriding the corrupted inertial measurements. The observer thus relies predominantly on visual cues during this phase, which remain unaffected by the gyroscope perturbation. Once the disturbance subsides and variance decreases, the gain returns to k_0 , allowing fast convergence driven by the now-reliable inertial data.

Quantitative Analysis (Adaptive-Gain):

- RMSE during disturbance (5–10 s): 0.64 rad ($\approx 37^\circ$)
- Maximum absolute error: 0.85 rad ($\approx 49^\circ$) at $t = 7.1$ s
- Recovery time (to 5% of pre-disturbance error): 1.5 s

Comparative Improvement:

- Error reduction during disturbance: 64.8% lower RMSE (adaptive vs. fixed)
- Peak error reduction: 79.4% lower maximum error
- Recovery speed improvement: $> 10\times$ faster reconvergence (1.5 s vs. > 15 s)

These results quantitatively demonstrate the substantial performance gains achieved through adaptive gain control.

5.3. Sensor behavior and variance metric

Figure 6 provides insight into the underlying sensor conditions and the variance metric driving the gain adaptation. The top subplot shows the three components of angular velocity $\Omega(t) = [\omega_x, \omega_y, \omega_z]^\top$ measured by the gyroscope. During the nominal phases (0–5 s and 10–30 s), the signals exhibit smooth, low-amplitude oscillations corresponding to the sinusoidal attitude trajectory (Eq. 4.1), with superimposed low-level Gaussian noise ($\sigma_\eta = 0.01$ rad/s).

Between $t = 5$ s and $t = 10$ s, the perturbation dramatically alters the signal characteristics. The constant bias shifts all three channels upward, while the high-frequency noise ($\sigma_\eta = 0.5$ rad/s) creates rapid fluctuations with amplitudes reaching ± 2 rad/s. This severe degradation is immediately reflected in the variance metric $v(t)$, shown in the bottom subplot.




Figure 6. Angular velocity (top), metric variance (bottom) in time.

The variance $v(t)$ remains near zero during stable operation, then spikes sharply to $v \approx 0.95$ within 0.5 s of disturbance onset. It stays elevated throughout the perturbation window, then decays exponentially once the disturbance ends, returning to $v < 0.05$ by $t = 11$ s. This behavior confirms that the variance metric provides a reliable, low-latency indicator of inertial sensor quality.

Key Insight: The variance-based feedback operates on a timescale ($T_w = 1$ s window) that balances responsiveness and noise rejection. Shorter windows would enable faster reaction but increase sensitivity to instantaneous noise spikes; longer windows would provide smoother metrics but delay detection of sudden failures. The 1-second window represents a practical compromise suitable for typical UAV dynamics.

On the Choice of T_w (Analytical Justification). The Welford-based variance estimate over a window of duration T_w sampled at 100 Hz averages approximately $100 T_w$ samples. For the additive Gaussian gyroscope noise in Table 2, the standard error of a windowed variance estimate scales as $O(1/\sqrt{100 T_w})$, so $T_w = 0.1$ s (10 samples) would yield a variance estimate with roughly $\sqrt{10} \approx 3.2\times$ larger relative fluctuation than the $T_w = 1.0$ s estimate used here, which would feed correspondingly noisier fluctuations into $K_P v(t)$ and $K_D dv/dt$ in Eq. (3.3) – directly increasing the risk of the noise-induced oscillation regime identified in Table 4. Conversely, a first-order estimate of the detection lag places the gain response time at approximately $T_w/2$ to T_w after disturbance onset, since the disturbance must populate a large-enough fraction of the window before it dominates the sample variance; this suggests $T_w = 2.0$ s would roughly double the ≈ 0.5 s response-onset time reported in Section 5.3 (Figure 6), delaying gain elevation into the disturbance window by a further second. These two competing effects—estimator variance decreasing, and detection lag increasing, with T_w —are consistent with $T_w = 1.0$ s being a reasonable interior operating point for the noise levels in Table 2, but we stop short of claiming it is optimal: a quantitative sweep of gain response time vs. T_w , run under

the actual (not analytically approximated) noise and disturbance model, is needed to confirm this and is left as follow-up work (Section 5.9).

5.4. Comparative summary and performance trade-offs

Table 3 consolidates the performance comparison between fixed-gain and adaptive-gain observers across multiple criteria.

Table 3. Comparison between fixed-gain and adaptive-gain observers.

Criterion	Fixed Gain ($k = 12$)	Adaptive Gain (PID)
Convergence (no disturbance)	Stable, fast	Stable, equivalent
RMSE during disturbance (5–10 s)	1.82 rad (104°)	0.64 rad (37°)
Maximum error	4.13 rad (237°)	0.85 rad (49°)
Recovery time after disturbance	> 15 s	1.5 s
Convergence with high noise	Divergence/drift	Robust, bounded error
Visual correction usage	Constant	Increased when needed
Gain response to sensor quality	None (static)	Responsive (PID)
Computational overhead	Baseline	+5% (variance + PID)
Parameter tuning effort	Manual k selection	Manual K_P , K_I , K_D tuning

Trade-off Analysis: While the adaptive observer provides substantial performance improvements, it introduces additional complexity in the form of PID parameter selection (K_P , K_I , K_D). However, we found these parameters to be more robust to variations than the fixed gain k , in the sense that a wide range of PID values yielded acceptable performance, whereas fixed-gain performance was highly sensitive to the specific choice of k . Furthermore, the computational overhead of variance estimation and PID control is minimal ($\approx 5\%$ increase in execution time), making the approach feasible for embedded implementation on resource-constrained UAV flight computers.

5.5. Interpretation and physical insight

The fundamental reason for the fixed-gain observer's poor performance lies in its inability to adjust the balance between inertial and visual information sources. During the disturbance (5–10 s), the gyroscope measurements are heavily corrupted, yet the fixed correction term $k \cdot \sigma(t)$ remains weak relative to the erroneous $\hat{\Omega}^\wedge(t)$ term in Eq. (2.5). Consequently, the observer continues to trust the inertial data, causing the estimate $\hat{R}(t)$ to diverge from the true attitude $R(t)$.

In contrast, the adaptive observer detects the increased variance $v(t)$ and responds by raising $k(t)$, which amplifies the correction term. This effectively shifts the observer's trust from inertial to visual measurements during the unreliable period. Once variance decreases post-disturbance, the observer automatically reverts to inertial-dominant behavior, enabling rapid reconvergence through the high-bandwidth gyroscope data.

This adaptive behavior mimics human sensory fusion: when one sense becomes unreliable (e.g., vision in fog), we instinctively rely more heavily on other senses (e.g., hearing, touch). The proposed PID-based mechanism provides a principled, automated implementation of this intuitive strategy within the geometric observer framework.

5.6. Sensitivity analysis of PID parameters

To assess the robustness of the proposed method to parameter variations, we conducted a sensitivity analysis by varying K_P , K_I , and K_D independently while holding other parameters constant. Table 4 summarizes the resulting robust operating ranges and captures the qualitative trends (monotonic improvement, plateau, and degradation regions) that the sensitivity study identified for each gain.

Table 4. Summary of PID-gain sensitivity: Robust operating ranges and failure modes.

Gain	Robust range	Used value	Degradation beyond range
K_P	[1.2, 2.0]	1.5	$K_P > 2.5$: noise-induced oscillation from excessive gain amplification
K_I	[0.005, 0.05]	0.01	$K_I > 0.1$: integrator windup, instability
K_D	[0.03, 0.1]	0.05	$K_D > 0.15$: amplifies measurement noise in dv/dt

Key Findings:

- **Proportional gain K_P :** Performance improves monotonically as K_P increases from 0.5 to 2.0, then plateaus. Values $K_P > 2.5$ cause excessive gain amplification, leading to noise-induced oscillations. The optimal range is $K_P \in [1.2, 2.0]$.
- **Integral gain K_I :** Small values ($0.005 \leq K_I \leq 0.05$) provide minimal benefit for transient disturbances but become important for sustained bias. Values $K_I > 0.1$ cause windup and instability.
- **Derivative gain K_D :** Moderate derivative action ($K_D \in [0.03, 0.1]$) smooths gain transitions and reduces overshoot. Values $K_D > 0.15$ amplify measurement noise in dv/dt , degrading performance.

These results indicate that the proposed parameters ($K_P = 1.5$, $K_I = 0.01$, $K_D = 0.05$) lie well within robust operating regions, providing confidence in the design's practical viability.

5.6.1. Toward a systematic tuning procedure

The values in Table 4 were obtained by an iterative, one-gain-at-a-time simulation sweep rather than a closed-form design rule. We therefore outline a systematic tuning procedure that a practitioner can follow to reproduce or re-tune the controller on a different platform, in the spirit of sequential (Ziegler–Nichols-like) PID tuning adapted to this gain-adaptation context:

1. *Fix k_0 from the fixed-gain baseline.* Set k_0 to the best-performing constant gain obtained by tuning the underlying Rehbinder–Ghosh observer alone (Section 2.4) under nominal, disturbance-free conditions; this is a single-parameter search identical to conventional fixed-gain tuning.
2. *Tune K_P first, with $K_I = K_D = 0$.* Increase K_P from zero under the disturbance protocol of Section 4.3 until the RMSE during the disturbance window (Eq. (4.2), restated below) stops decreasing and begins to increase again due to noise amplification (the onset of the degradation regime in Table 4); select K_P at roughly 70–80% of this onset value to retain margin.

3. *Add K_I only if steady-state bias persists.* With K_P fixed, increase K_I from zero only if residual error remains during sustained (slow-drift) disturbances; because $K_I \ll K_P$ is required to avoid windup (Section 3.3), this search can be restricted to $K_I \in [0, 0.1 K_P]$.
4. *Add K_D last, to damp transients.* Increase K_D from zero to reduce overshoot in $k(t)$ immediately after disturbance onset (visible as ringing in Figure 3), stopping once the derivative term begins to visibly amplify measurement noise in dv/dt .
5. *Validate across the disturbance-scenario family of Section 5.8* (varying intensity, duration, and randomized noise, rather than the single fixed 5–10 s scenario) before freezing the gains.

This sequential procedure is what produced the operating ranges in Table 4; it does not guarantee global optimality, and a formal robust-region characterization (e.g., via structured-uncertainty analysis of the linearized adaptation loop) remains future work (Section 5.9).

5.7. Comparison with related work

While direct numerical comparison with other methods is challenging due to differences in simulation environments and sensor models, we provide a qualitative assessment of the proposed approach relative to recent literature.

Complementary Filters on SO(3) [2, 8]: These methods use fixed or heuristically scheduled gains to fuse inertial and vector measurements (e.g., gravity, magnetometer). Berkane and Tayebi [2] proposed state-dependent gains that improve robustness, but their approach requires knowledge of measurement covariance matrices, which may not be available online. Our variance-based metric provides a simpler, more direct indicator of sensor reliability without requiring covariance estimation.

Invariant EKF on Manifolds [9, 16]: These filters extend the Kalman framework to Lie groups, preserving geometric structure. Wang et al. [16] explored adaptive covariance tuning for invariant EKF, but their method involves matrix computations and parameter identification that increase computational burden. Our PID-based approach offers comparable robustness with significantly lower complexity (scalar variance and gain computations).

Learning-Based Visual-Inertial Odometry [34, 35]: Recent deep learning methods achieve impressive accuracy by learning sensor fusion policies from data. However, these approaches require extensive training datasets, lack theoretical guarantees, and may not generalize to novel environments. Our method provides deterministic, interpretable adaptation with formal stability analysis, making it more suitable for safety-critical applications.

Table 5. Comparison of the proposed adaptive observer with related methods.

Method	Adaptation Mechanism	Computational Cost	Theoretical Guarantees	Tuning Complexity
Fixed-gain SO(3) observer [1]	None	Low	Local stability	Single parameter (k)
State-dependent gains [2]	Heuristic scheduling	Medium	Local stability	Multiple thresholds
Invariant EKF [9, 16]	Covariance adaptation	High	Local (linearized)	Covariance matrices
Deep VI-Odometry [34, 35]	Learned policy	Very high	None	Large dataset required
Proposed (Adaptive PID)	Variance feedback	Low-Medium	Local stability	Three PID gains

5.8. Limitations and practical considerations

Despite the promising results, several limitations warrant discussion:

1. **Visual Feature Availability:** The adaptive observer assumes continuous tracking of at least three non-collinear line features. In practice, visual features may be temporarily lost due to occlusion, motion blur, or lack of structured geometry (e.g., open outdoor environments). When visual measurements are unavailable, the correction term $\sigma(t)$ vanishes, and the observer degenerates to pure gyroscope integration. In such cases, even adaptive gain control cannot prevent drift. Future work should incorporate feature availability metrics into the adaptation law, reducing gain when visual data is unreliable or absent.
2. **Simultaneous Inertial and Visual Failures:** The variance metric $v(t)$ specifically addresses inertial degradation. If both inertial and visual measurements fail simultaneously (e.g., electromagnetic interference affecting electronics and motion blur degrading image quality), no amount of gain adaptation can maintain accurate estimates. Robust multi-modal fusion should incorporate fault detection for all sensor modalities, potentially triggering safe fallback behaviors (e.g., hovering, emergency landing).
3. **Response Time to Sudden Faults:** The 1-second variance window introduces inherent latency in detecting sudden sensor faults. For extremely rapid disturbances (e.g., high-frequency vibrations, instantaneous bias jumps), the variance metric may lag, delaying gain adaptation. Complementary fault detection methods based on innovation analysis or residual monitoring could provide faster triggers for gain adjustment.
4. **PID Parameter Generalization:** While our empirically tuned PID gains performed well across the tested scenarios, they may require adjustment for different UAV platforms (varying inertia, flight dynamics) or sensor specifications (different noise characteristics). Systematic tuning procedures or adaptive PID algorithms that self-tune based on observed performance could enhance generalizability.
5. **Simulation vs. Real-World Validation:** All results presented in this paper were obtained through MATLAB/Simulink simulations with synthetic sensor data and optical flow. We did not conduct hardware-in-the-loop (HIL) testing or flight experiments for this revision; the discussion below is therefore an engineering-constraint analysis of *how* the method would need to be adapted to address each real-world complication, not a report of having already tested it under these conditions. Real-world experiments introduce additional complexities, which we address in turn:
 - (a) **Asynchronous sensor timestamps.** The observer dynamics of Eq. (3.1) implicitly assume $\Omega(t)$, $\sigma(t)$, and $k(t)$ are evaluated at the same instant t . In practice, the gyroscope (assumed here at 100 Hz) and the camera-derived correction $\sigma(t)$ (at 30 Hz, i.e., roughly one gyroscope sample every 3.3 inertial steps) arrive on independent clocks with unknown relative offset. The variance metric $v(t)$, being derived purely from the gyroscope stream, is unaffected by this asynchrony, but $\sigma(t)$ must be time-aligned to the gyroscope timeline before being combined in Eq. (3.1). A practical handling scheme is to buffer $\sigma(t)$ and interpolate/hold it to the gyroscope sample times using the known frame timestamp, which is the same buffering strategy already implicit in the 0.03 s visual measurement delay in Table 2; what changes on real hardware is that this delay is variable rather than the fixed one-frame value assumed in simulation, and should be estimated online (e.g., via a hardware timestamp or precision time protocol (PTP)-style synchronization) rather than assumed constant.

- (b) **Rolling-shutter distortion.** A rolling-shutter camera exposes each image row at a slightly different time, so a line feature's apparent position at a nominal frame timestamp is actually a composite of several true poses during high angular rates—this directly corrupts the normal vectors $y_i(t)$ used in Eq. (2.7) and, in turn, $\sigma(t)$. Because $v(t)$ depends only on the gyroscope, rolling-shutter error would manifest as an increase in the *visual* residual rather than in $v(t)$, meaning the adaptive gain law as currently formulated would not automatically compensate for it—it could even be counterproductive, since raising $k(t)$ when $v(t)$ happens to be elevated for an unrelated inertial reason would amplify a rolling-shutter-corrupted $\sigma(t)$. A row-time correction using the per-row exposure timestamp and the gyroscope-propagated attitude (standard in visual-inertial odometry pipelines) would need to be applied to $y_i(t)$ upstream of Eq. (2.7); alternatively, a global-shutter camera avoids the issue at the sensor level.
- (c) **Frame-rate mismatch and camera calibration/lens distortion.** The 30 Hz visual update rate relative to the 100 Hz observer/PID loop means $\sigma(t)$ is held constant (zero-order hold) between visual updates; on real hardware, dropped or delayed frames would lengthen this hold unpredictably. Since $v(t)$ is computed independently of the visual channel, the adaptive gain would not detect a stalled visual channel—this is the same blind spot noted for visual feature loss in item 1 above, and reinforces the recommendation there to add a visual-side reliability metric (e.g., a rolling count of frames since the last successful line association) as a second input to the adaptation law. Camera intrinsic calibration and lens-distortion correction are assumed exact in Eq. (2.4); residual calibration error would appear as a bias in $y_i(t)$ that the current gain law cannot distinguish from a genuine attitude error, and would need to be characterized offline (e.g., via a standard checkerboard calibration) rather than compensated online.

HIL testing and flight experiments remain, as stated in the submitted version, essential next steps; the analysis above is intended to make explicit which specific interface (gyroscope timeline, visual timeline, or calibration) each real-world effect enters through, so that HIL testing can be scoped to isolate them individually rather than only observing an aggregate performance drop.

6. **Single Disturbance Scenario.** All quantitative results in Section 5 use a single fixed disturbance profile (constant bias $[0.5, 0.3, 0.4]^T$ rad/s plus $\sigma_\eta = 0.5$ rad/s Gaussian noise, injected over a fixed 5–10 s window; Table 2). We did not evaluate the method under varying disturbance intensity, varying duration, or randomized (non-fixed-window) disturbance onset in this revision. Because the PID-gain trajectory in Figure 3 is itself a closed-loop response to $v(t)$, its behavior under a substantially weaker or stronger disturbance, or a disturbance of different duration, cannot be safely extrapolated from the single scenario reported here—for instance, a disturbance shorter than $T_w = 1$ s might not fully populate the variance window before ending, potentially under-triggering the gain response. A systematic robustness study across disturbance intensity, duration, and randomized onset time/magnitude is needed before generalizing the 65% RMSE-reduction and recovery-time figures reported here to disturbance conditions other than the one tested, and we flag this explicitly as an open validation gap rather than a demonstrated property of the method.
7. **Benchmarking Against Mainstream Alternatives.** Section 5.7 and Table 5 compare the proposed method to adaptive complementary filters, invariant EKF variants, and learning-based visual-inertial odometry (VIO) *qualitatively*, based on the mechanisms and computational

complexity reported in the cited literature, because these methods were not re-implemented and run side-by-side with the proposed observer under the identical disturbance protocol of Section 4.3 in this study. A quantitative, same-scenario comparison against at least one representative of each family (e.g., the state-dependent-gain filter of [2] and an invariant EKF such as [16] or [9]) would be needed to substantiate the performance claims in Table 3 on a relative, cross-method basis rather than only against the fixed-gain Rehbinder–Ghosh baseline; this is left as a direct follow-up (Section 5.9).

5.9. Broader implications and future directions

Impact on UAV Autonomy: The proposed adaptive observer addresses a critical bottleneck in GPS-denied navigation: maintaining reliable attitude estimates under unpredictable sensor conditions. By autonomously adjusting trust in inertial vs. visual data, the method enhances robustness without requiring manual intervention or environment-specific calibration. This capability is particularly valuable for indoor inspection, underground exploration, and urban search-and-rescue operations where sensor reliability varies dramatically over short timescales.

Extension to Full Pose Estimation: The current work focuses exclusively on attitude (orientation) estimation. A natural extension is to incorporate translational dynamics and velocity estimation, formulating a full pose observer on the special Euclidean group $SE(3) = SO(3) \ltimes \mathbb{R}^3$. Adaptive gain control could be extended to modulate both rotational and translational correction terms, potentially using separate variance metrics for linear and angular accelerations.

Learning-Based PID Tuning: While we selected PID parameters empirically, reinforcement learning (RL) offers a promising avenue for automated tuning. An RL agent could learn optimal K_P , K_I , K_D values by interacting with the observer in simulation, minimizing a cost function that balances convergence speed, overshoot, and steady-state error. This approach could discover non-intuitive tuning strategies tailored to specific flight regimes or environments.

Multi-Sensor Fusion Beyond Vision-Inertial: The variance-based adaptation principle is not limited to vision-inertial fusion. Similar metrics could assess the reliability of other sensor modalities (e.g., LiDAR point cloud density, GPS signal strength, barometer stability) and adjust their contributions to a multi-modal observer accordingly. This generalization could enable truly adaptive sensor fusion across arbitrary combinations of sensing modalities.

6. Conclusions

This paper presented an adaptive gain control mechanism for robust vision-inertial attitude estimation on the special orthogonal group $SO(3)$. Building upon the classical Rehbinder-Ghosh geometric observer, we introduced a PID-based adaptation law that dynamically modulates the observer's correction gain based on real-time gyroscope variance metrics. This innovation addresses the long-standing challenge of gain tuning in geometric observers, enabling autonomous adjustment of the trade-off between convergence speed and noise sensitivity without manual intervention.

Comprehensive simulation studies demonstrated substantial performance improvements relative to fixed-gain baselines. Under severe gyroscope perturbations including bias and high-frequency noise, the adaptive observer achieved approximately 65% reduction in estimation error and reconverged to accurate tracking more than 10 times faster than the classical approach. The variance-based feedback

metric proved to be a reliable, low-latency indicator of inertial sensor quality, facilitating smooth transitions between inertial-dominant and vision-dominant estimation modes. Theoretical analysis confirmed that local exponential stability of the original observer is preserved under the adaptive gain framework, provided the PID controller generates bounded, slowly-varying gains—conditions satisfied by design in our implementation.

The proposed method offers several practical advantages for autonomous UAV applications. The computational overhead is minimal, requiring only scalar variance computation and PID control updates at each timestep, making it suitable for embedded flight computers with limited processing resources. The adaptation mechanism is modular and can be integrated into existing SO(3) observer implementations with minimal code modification. The PID parameters exhibit robustness to variations within reasonable ranges, reducing sensitivity to tuning compared to fixed-gain designs. Most importantly, the approach eliminates the need for heuristic gain scheduling or environment-specific calibration, providing a deterministic, interpretable solution with formal stability guarantees.

Several promising directions for future research emerge from this work. Extending the adaptive framework to full pose estimation on SE(3) would enable simultaneous attitude, position, and velocity estimation with adaptive sensor fusion across all degrees of freedom. Incorporating visual feature availability metrics into the adaptation law could further enhance robustness by reducing reliance on visual corrections when line tracking becomes unreliable due to occlusion or motion blur. Exploring machine learning approaches, particularly RL for automated PID tuning, could discover optimal adaptation strategies tailored to specific UAV platforms and operational environments. Additionally, rigorous analysis of global or semi-global stability using time-varying Lyapunov theory or hybrid system frameworks remains an open theoretical challenge. Most critically, experimental validation through HIL testing and real-world flight experiments is essential to verify the method's performance under actual sensor characteristics, computational constraints, and environmental disturbances. Initial experiments using advanced line detection algorithms such as DeepLSD or HAWP, combined with learned optical flow methods like RAFT, represent a promising pathway toward robust attitude estimation in highly challenging visual conditions.

This work bridges geometric control theory and practical sensor fusion engineering, offering a concrete solution to a recognized gap in autonomous navigation research. By enabling UAVs to autonomously adapt their estimation strategies to changing sensor conditions, the proposed adaptive observer helps close the gap between theoretical stability guarantees and real-world robustness requirements. The variance-driven adaptation principle established here provides a foundation for broader developments in adaptive multi-modal sensor fusion for safety-critical autonomous systems operating in uncertain, GPS-denied environments.

Author Contributions

Abdelwahhab Khaldi: Conceptualization, methodology development, algorithm design, simulation implementation, data analysis, manuscript writing—original draft, visualization. Choukri Bensalah: Conceptualization, methodology validation, critical review and editing, theoretical framework supervision, formal analysis. Amal Choukchou-Braham: Project supervision, conceptual guidance, critical revision for intellectual content, final approval of the version to be published. All authors have read and approved the final manuscript.

Use of Generative-AI tools declaration

The authors declare they have used Artificial Intelligence (AI) tools in the creation of this article to enhance the technical presentation and language quality. All AI-generated content has been carefully reviewed, verified, and validated by the authors, who take full responsibility for the accuracy and integrity of the manuscript.

AI tools used:

Large language models (LLMs), specifically Claude (Anthropic) and/or ChatGPT (OpenAI), were utilized during the preparation of this manuscript.

How were the AI tools used?

1. **Technical content development and structuring:** AI tools assisted in organizing and articulating the technical methodology, particularly in Section 4 (perturbation scenarios and performance metrics). The AI helped structure the presentation of the disturbance protocol and quantitative evaluation metrics, including the mathematical formulation of RMSE (Eq. 4.2) and MAE (Eq. 4.3). The core technical concepts, and mathematical formulations were developed by the authors, while the AI assisted in enhancing clarity and coherence of presentation.
2. **Language editing and grammar correction:** AI tools were employed throughout the manuscript to correct grammatical errors, improve sentence structure, and enhance the overall readability of the English text. This included refinement of technical terminology and ensuring consistency in scientific writing style.
3. **Discussion enhancement:** AI tools assisted in expanding and articulating discussions based on technical information and results provided by the authors. The AI helped synthesize provided data and findings into coherent narrative form while maintaining scientific rigor.

Where in the article is the information located?

AI assistance was primarily utilized in the following sections:

- Section 4: Perturbation scenarios and performance metrics (technical content structuring and presentation)
- Throughout the manuscript: Language correction and grammatical refinement
- Discussion section: Narrative development based on author-provided technical results and interpretations

All numerical results, experimental data, theoretical derivations, and scientific conclusions presented in this work are entirely the product of the authors' research and have been independently verified. No AI tools were used for data collection, analysis, or generation of results.

Acknowledgments

This work did not receive any specific funding. The authors thank Laboratoire de Tlemcen (LAT), University of Tlemcen, for providing the research environment supporting this work.

Conflict of interest

All authors declare no conflicts of interest in this paper.

References

1. Rehbinder H, Ghosh B (2003) Pose estimation using line-based dynamic vision and inertial sensors. *IEEE T Automat Contr* 48: 186–191. <https://doi.org/10.1109/TAC.2002.808464>
2. Berkane M, Tayebi A (2016) On the Design of Attitude Observers on the Special Orthogonal Group $SO(3)$. *arXiv:1606.06208*.
3. Ma F, Shi J, Yang Y, Li J, Dai K (2019) ACK-MSCKF: Tightly-coupled Ackermann multi-state constraint Kalman filter for VIO under special motions. *Sensors* 19: Article 4816. <https://doi.org/10.3390/s19214816>
4. Abdollahi MR, Pourtakdoust SH, Nooshabadi MY, Pishkenari HN (2024) An improved multi-State constraint Kalman filter for visual-inertial odometry. *J Franklin Ins* 361: 107130. <https://doi.org/10.1016/j.jfranklin.2024.107130>
5. Negru SA, Geragersian P, Petrunin I, Guo W (2024) Resilient multi-sensor UAV navigation with a hybrid federated fusion architecture. *Sensors* 24: 981.
6. An L, Li X, Li H (2022) A robust visual-aided inertial navigation algorithm for pedestrians: Improved accuracy in dynamic environments. *J Sensors* 2022: 1311221. <https://doi.org/10.1155/2022/1311221>
7. Pandey Y, Bhattacharyya R, Singh YN (2024) Adaptive invariant extended Kalman filter with noise covariance tuning for attitude estimation. *arXiv preprint arXiv:2410.01958*.
8. Singh A, Sieb TS, Howe JH, Hashim HA (2020) Nonlinear attitude filter on $SO(3)$: Fast adaptation and robustness. In *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, 869–874.
9. Goffin S, Barrau A, Bonnabel S, Bröls O, Sacré P (2025) Iterated invariant extended Kalman filter (IterIEKF). *IEEE T Automat Contr*.
10. Goodarzi F, Lee D, Lee T (2013) Geometric nonlinear PID control of a quadrotor UAV on $SE(3)$. In *Proc European Control Conference (ECC)*, 3845–3850. <https://doi.org/10.23919/ECC.2013.6669644>
11. Noordin A, Mohd Basri MA, Mohamed Z, Mat Lazim I (2021) Adaptive PID controller using sliding mode control approaches for quadrotor UAV attitude and position stabilization. *Arab J Sci Eng* 46: 963–981. <https://doi.org/10.1007/s13369-020-04742-w>

12. Chen J, Long Y, Li T, Huang T (2023) Attitude tracking control for quadrotor based on time-varying gain extended state observer. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering* 237: 585–95.
13. Zhang J, Yu X, Sier H, Zhang H, Westerlund T (2024) Event-based sensor fusion and application on odometry: A survey. In *2025 IEEE 6th International Conference on Image Processing, Applications and Systems (IPAS)*, 1–6.
14. Yi B, Jin C, Wang L, Shi G, Ila V, Manchester IR (2024) PEBO-SLAM: Observer design for visual inertial SLAM with convergence guarantees. *IEEE T Automat Contr* 70: 1714–1729.
15. Wang M, Berkane S, Tayebi A (2021) Nonlinear observers design for vision-aided inertial navigation systems. *IEEE T Automat Contr* 67: 1853–1868.
16. Wang J, Li M (2021) Adaptive invariant Kalman filtering for attitude estimation on SO(3) thorough feedback calibration of prior error covariance. *IET Control Theory Appl* 15: 1906–1914. <https://doi.org/10.1049/cth2.12166>
17. Song W, Li P, Wang M (2024) Transmission line detection based on improved Hough transform. *arXiv preprint arXiv:2402.02761*.
18. Teplyakov L, Erlygin L, Shvets E (2022) LSDNet: Trainable modification of LSD algorithm for real-time line segment detection. *IEEE Access* 10: 45256–45265. <https://doi.org/10.1109/ACCESS.2022.3169177>
19. Pautrat R, Barath D, Larsson V, Oswald MR, Pollefeys M (2023) DeepLSD: Line segment detection and refinement with deep image gradients. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17327–17336.
20. Xue N, Wu T, Bai S, Wang F, Xia GS, Zhang L, et al. (2020) Holistically-Attracted Wireframe Parsing. *CVPR*, 2785–2794. <https://doi.org/10.1109/CVPR42600.2020.00286>
21. Lucas BD, Kanade T (1981) An iterative image registration technique with an application to stereo vision. In *IJCAI'81: 7th international joint conference on Artificial intelligence* 2: 674–679.
22. Horn BK, Schunck BG (1981) B. K. Horn and B. G. Schunck, Determining optical flow. *Artif Intell* 17: 185–203. [https://doi.org/10.1016/0004-3702\(81\)90024-2](https://doi.org/10.1016/0004-3702(81)90024-2)
23. Farnebäck G (2003) Two-frame motion estimation based on polynomial expansion. *Proc SCIA*, 363–370. https://doi.org/10.1007/3-540-45103-X_50
24. Teed Z, Deng J (2020) RAFT: Recurrent all-pairs field transforms for optical flow. In *European conference on computer vision*, 402–419. Cham: Springer International Publishing.
25. Huang Z, Shi X, Zhang C, Wang Q, Cheung KC, Qin H, et al. (2022) FlowFormer: A transformer architecture for optical flow. In *European conference on computer vision*, 668–685.
26. Zhan H, Weerasekera CS, Bian JW, Garg R, Reid I Y. Pan et al., DF-VO: What should be learnt for visual odometry?. *arXiv preprint arXiv:2103.00933*.
27. van Goor P, Hamel T, Mahony R (2025) Synchronous observer design for inertial navigation systems with almost-global convergence. *Automatica* 177: 112328.
28. Jiang P, Liu C, Cong H, Zhang F (2024) A fuzzy adaptive complementary filter for attitude estimation based on norm judgment. *Measurement and Control* 57: 1161–1180. <https://doi.org/10.1177/00202940241227069>

29. Suid MH, Ahmad MA (2022) Optimal tuning of sigmoid PID controller using nonlinear Sine Cosine algorithm for the automatic voltage regulator system. *ISA Transactions* 128: 265–286. <https://doi.org/10.1016/j.isatra.2021.11.037>
30. Ghazali MR, Ahmad MA, Jusof MF, Ismail RM (2018) A data-driven neuroendocrine-PID controller for underactuated systems based on safe experimentation dynamics. In *Proc IEEE 14th Int Colloq Signal Process Its Appl (CSPA)*, 61–66. <https://doi.org/10.1109/CSPA.2018.8368686>
31. Lucas C, Shahmirzadi D, Sheikholeslami N (2004) Introducing BELBIC: Brain emotional learning based intelligent controller. *Intell Autom Soft Comput* 10: 11–21. <https://doi.org/10.1080/10798587.2004.10642862>
32. Izci D, Ekinici S, Jabari M, Kocaman B, Güneş BB, Adas E, et al. (2026) A novel gudermannian function-driven controller architecture optimized by starfish optimizer for superior transient performance of automatic voltage regulation. *Biomimetics* 11: art. 7. <https://doi.org/10.3390/biomimetics11010007>
33. Izci D, Ekinici S, Mohd Tumari MZ, Ahmad MA (2025) A novel softsign fractional-order controller optimized by an intelligent nature-inspired algorithm for magnetic levitation control. *Fractal Fract* 9: art. 801. <https://doi.org/10.3390/fractalfract9120801>
34. Pan Y, Zhou W, Cao Y, Zha H (2024) Adaptive VIO: Deep visual-inertial odometry with online continual learning. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18019–18028.
35. Chen C, Rosa S, Miao Y, Lu CX, Wu W, Markham A, et al. (2019) Selective Sensor Fusion for Neural Visual-Inertial Odometry. *CVPR*, 10534–10543. <https://doi.org/10.48550/arXiv.1903.01534>
36. Khaldi A, Bensalah C, Braham AC, Cherki B (2019) UAV attitude estimation using visual and inertial data fusion based on observer in SO(3). *Proc ICAEE*, 1–5.
37. Wang Q, Liu K, Cao Z (2023) System noise variance matrix adaptive Kalman filter method for AUV INS/DVL navigation system. *Ocean Eng* 267: 113269.



AIMS Press

©2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)