

Improving Attitude and Heading Reference Systems Performance via Machine Learning-driven Parameters Fine-Tuning

Tommaso C. Ferrari¹ and Felipe O. Silva², *Senior Member, IEEE*

Abstract—Attitude and Heading Reference Systems (AHRs) based on Error-State Extended Kalman Filters (ES-EKF) require careful tuning of stochastic parameters to achieve optimal performance. Traditional approaches rely on Allan Variance (AV) analysis for parameter identification, which, while physically grounded, does not guarantee optimal navigation performance. This paper presents three advanced optimization methodologies—Gaussian Processes (GP), Nondominated Sorting Genetic Algorithm III (NSGA-III), and Multi-Objective Tree-structured Parzen Estimator (MO-TPE)—to systematically fine-tune critical ES-EKF parameters directly against flight performance metrics. The optimization targets accelerometer/magnetometer noise/bias characteristics, vehicle dynamics/Earth magnetic field uncertainty, and Innovation Filter (IF) thresholds across a multi-dimensional parameter space. Experimental validation using 19 real Unmanned Aerial Vehicle (UAV) flights demonstrates substantial improvements over traditional AV-based tuning. The proposed methods achieve 81-91% reductions in Root Mean Square (RMS) attitude errors, 82-91% improvements in Mean Absolute Errors (MAE), and 41-93% enhancements in estimation consistency across roll, pitch, and yaw axes. MO-TPE consistently delivers the best overall performance, achieving optimal results in 7 out of 9 metric-axis combinations, followed closely by NSGA-III, while GP provides competitive single-objective optimization. The largest improvements are observed in yaw estimation, where traditional approaches struggle most with magnetic disturbances. These results demonstrate that machine learning-driven parameter optimization can significantly enhance AHRs accuracy and robustness without modifying the underlying ES-EKF structure, offering a practical path for improving navigation performance in real-world UAV applications.

I. INTRODUCTION

Navigation is the area of science that deals with the determination of position, velocity, and angular orientation, a.k.a. attitude of a body, regardless of its type (vehicles, for example) [1]. Attitude and Heading Reference Systems (AHRs) are specialized navigation systems that provide information about vehicle attitude only [2]. They are considered to be “integrated navigation systems” in the sense that they combine information from different navigation sensors, each conveying its particular error characteristics. Traditionally, AHRs’ sensors comprise a triad of angular

rate gyros, which provide full attitude increments for the vehicle based on the digital integration of their outputs, plus a triad of accelerometers and magnetometers, which provide absolute information about the vehicle roll/pitch and yaw, respectively [3].

While the rate gyros-derived attitude solution accumulates errors over time, as expected for any dead-reckoning system, the absolute attitude from the accelerometers and magnetometers suffers from high short-term noise. Notwithstanding, the latter are only deemed to be reliable, provided the vehicle is not accelerating and spurious magnetic fields (e.g., those created by the vehicle electronics/motors and/or navigation equipment itself) have been systematically mitigated via an online calibration procedure [4]. Such requirements represent encumbrances to AHRs in the sense that, during their regular operation: (a) the vehicle is susceptible to accelerations, turns and/or vibration; and (b) effective magnetometer calibration procedures require: (i) the vehicle to be rotated, online, around all its three axes; as well as (ii) the knowledge of the Earth magnetic flux density magnitude, which is location-dependent (remember that AHRs only provide us with attitude information, not position).

As per the integration part, AHRs have been fused/combined via different strategies. Complementary filters are probably the simplest of the latter, allowing the rate gyros-derived full attitude of the vehicle to be combined with the accelerometers’ roll/pitch and magnetometers’ yaw via gains that are constant and complementary (i.e., they sum to unity) [5], [6]. As an alternative, Mahony *et al.* [7] proposed an integration (filtered) approach based on a proportional-integral controller that acts on the attitude error estimated from the sensor data. Aiming at optimal solutions, variants of the Kalman Filter (KF) [8], such as the Error-State Extended KF (ES-EKF), are probably the most employed tools for sensor fusion within AHRs. In an ES-EKF, the rate gyro-derived attitude error is modeled and predicted forward in time, using appropriate dynamic/evolution models, being in sequence, updated and compensated for whenever reliable measurements from the accelerometers and magnetometers are available [2].

Despite providing good performance overall, ES-EKF-based AHRs are still suboptimal, as they are unable to comply with a basic assumption KFs should be based on, namely, the whiteness and Gaussian distribution of the process and measurement noises [1]. In general, such an issue is fruit of: (a) the existence of stochastic error components on the sensors outputs whose identification/parametrization in terms of state space models is either impossible or only a rough

*This study was financed in part by the Coordination for the Improvement of Higher Education Personnel (CAPES), in part by the Brazilian National Council for Scientific and Technological Development (CNPq), under grant 312194/2022-6, and in part by the Minas Gerais State Agency for Research and Development (FAPEMIG), under grant APQ-04659-22.

¹Tommaso C. Ferrari is with the Autonomous Robotics Research Center, Technology Innovation Institute, Masdar City, Abu Dhabi 20015, UAE tommaso.castiglione@tii.ae

²Felipe O. Silva is with Department of Automatics, Engineering School, Federal University of Lavras, Lavras, MG 37200-900, Brazil felipe.oliveira@ufla.br

approximation [9]; and (b) the fact that, in addition to the aforementioned error components, the ES-EKF measurement noise is also a function of the vehicle's acceleration/vibration level and residual internal/external magnetic fields, which are time varying, and hence, hard to characterize in terms of time correlation and statistical distribution [2].

To mitigate problem (b) above, one option is to tune empirically the ES-EKF-based AHRS measurement covariance matrix using historical data collected during a benign navigation phase of the vehicle, and then reject spurious accelerometers/magnetometers measurements/outliers via, e.g. an Innovation (a.k.a. spike) Filter (IF) [1]. Suitably setting the IF threshold, though, is not intuitive, especially for highly non-Gaussian noise distributions. An alternative solution to accommodate outliers in nonlinear state estimation problems has been suggested by Aghapour *et al.* [10], via their Risk-Averse Performance-Specified (RAPS) approach [10].

Regarding the aforementioned problem (a), stochastic noise parameter identification tools, such as the Allan Variance (AV), accompanied by suitable ES-EKF state augmentation techniques, have proven to be effective [11], [12]. For the particular problem of barometer-aided inertial navigation for instance, optimized approaches based on Weighted Least Squares (WLS) [9] and Genetic Algorithms (GA) [13] have allowed the intricate stochastic error components existing in the outputs of navigation sensors to be accurately identified and modeled. However, as assessed by Cavalcanti *et al.* [14], optimization of the sensors' stochastic error models, at least from the AV standpoint, does not guarantee optimality in terms of navigation performance, which is the ultimate goal of the ES-EKF designer. In this sense, the use of optimization approaches (driven, e.g., by machine learning) aimed at finding suitable ES-EKF tunings and/or IF thresholds that would maximize AHRS performance seems to be a relevant topic of research, which, to the best of the authors' knowledge, has not been explored in depth in the literature yet.

Given the aforementioned scenario, this paper revisits the AHRS sensor integration problem and compares the navigation performance of the traditional ES-EKF-based implementation with AV parameter identification against three optimization-based approaches that systematically fine-tune critical filter parameters. The proposed methodologies employ: (a) Gaussian Processes (GP)-based Bayesian optimization for single-objective parameter tuning; (b) Nondominated Sorting Genetic Algorithm III (NSGA-III) and (c) Multi-Objective Tree-structured Parzen Estimator (MO-TPE) for Multi-Objective (MO) optimization that simultaneously minimizes roll, pitch, and yaw estimation errors. Unlike traditional approaches that rely solely on AV characterization and empirical tuning, these methods optimize parameters directly against flight performance metrics, potentially discovering superior configurations that transcend the limitations of conventional noise analysis. As the main contribution of this paper, we show that the proposed optimization techniques significantly improve AHRS attitude estimation accuracy w.r.t. traditional AV-based parameter selection, with quantified performance improvements across diverse Unmanned

Aerial Vehicle (UAV) flight scenarios.

The rest of this paper is organized as follows: Sections II to II-D review the fundamental principles of ES-EKF-based AHRS implementation and its critical tuning parameters. Section III introduces the three advanced optimization frameworks (GP, NSGA-III, and MO-TPE) proposed to systematically optimize AHRS navigation performance. Section IV details the stochastic parameter fine-tuning methodology, including parameter space definition, objective function formulation, and optimization workflow implementation. Subsequently, Sections V and VI present comprehensive experimental results from real UAV flight tests and conclude the paper, respectively.

II. ES-EKF-BASED AHRS

This Section is devoted to briefly summarizing the main calculation/estimation steps that take place within an ES-EKF-based AHRS. As suggested in [2], [15], a Tightly-Coupled (TC) version of the latter is adopted.

A. Propagation/Prediction Phase

The first step within an ES-EKF-based AHRS consists of propagating the rate gyro-based full attitude solution via the digital integration of their measurements, namely, the angular rate vector of the vehicle frame¹ w.r.t. the inertial frame², resolved on vehicle frame axes, ω_{ib}^b . In order to avoid Euler angle singularities, the vehicle attitude is parametrized here as a Direction Cosine Matrix (DCM) encoding the rotation of the vehicle frame w.r.t. the navigation frame³, \mathbf{C}_b^n . The rate gyro-derived attitude can be propagated using [1], [16]:

$$\hat{\mathbf{C}}_b^n(+)=\hat{\mathbf{C}}_b^n(-)\hat{\mathbf{C}}_{b+}^b-, \quad (1)$$

with

$$\hat{\mathbf{C}}_{b+}^b-=\mathbf{I}_3+\frac{\sin|\hat{\alpha}_{ib}^b|}{|\hat{\alpha}_{ib}^b|}(\hat{\alpha}_{ib}^b\times)+\frac{1-\cos|\hat{\alpha}_{ib}^b|}{|\hat{\alpha}_{ib}^b|^2}(\hat{\alpha}_{ib}^b\times)^2, \quad (2)$$

$$\hat{\alpha}_{ib}^b=\hat{\omega}_{ib}^b\tau_s, \quad (3)$$

where symbol $\hat{(\)}$ represents an estimated variable, i.e., a variable that contains errors; (+) and (-) denote variables in the current and immediately preceding epochs, respectively; \mathbf{I}_3 is a 3-by-3 identity matrix; \times is the cross-product operator; and τ_s is the propagation time (usually the rate gyro sampling time).

As the rate gyro measurements contain errors (both systematic and stochastic), the digital integration of the latter leads to error accumulation in the estimated attitude DCM,

¹The vehicle frame is represented in this paper by sub/superscript b , and is defined as having its origin at the vehicle center-of-mass, and x , y , and z -axes pointing forward, rightward, and downward, respectively.

²The inertial frame is represented in this paper by subscript i , and is defined as having its origin at the Earth center-of-mass, and z , x , and y -axes pointing along the Earth rotation axis, crossing the Equatorial plane with the Greenwich meridian, and completing the right-hand orthogonal set, respectively [1].

³The navigation frame is represented in this paper by sub/superscript n , and is defined as having its origin at the vehicle center-of-mass, and x , y , and z -axes pointing to the North, East, and Down (NED) directions, respectively.

$\hat{\mathbf{C}}_b^n$. The latter can be modeled (to first order) as a function of the true attitude DCM, \mathbf{C}_b^n , as [1]:

$$\hat{\mathbf{C}}_b^n = \mathbf{C}_b^n + \delta\mathbf{C}_b^n \approx (\mathbf{I}_3 + \delta\boldsymbol{\Psi}_{nb}^n \times) \mathbf{C}_b^n. \quad (4)$$

where $\delta\boldsymbol{\Psi}_{nb}^n$ is the small-angle rotation vector that encompasses the misalignment between the estimated and true navigation frames, and which can also be seen as the error in the Euler (i.e., roll, pitch, and yaw) angles [1].

What the EKF attempts to do in an Error-State (ES) framework is exactly to estimate (and then correct) $\delta\boldsymbol{\Psi}_{nb}^n$, which is known to evolve in time approximately as [2]:

$$\delta\dot{\boldsymbol{\Psi}}_{nb}^n \approx \mathbf{C}_b^n \delta\boldsymbol{\omega}_{ib}^b, \quad (5)$$

where $\delta\boldsymbol{\omega}_{ib}^b \equiv \mathbf{b}_g + \boldsymbol{\eta}_g$ are the rate gyro errors, which are assumed to be corrupted by an instability bias vector, \mathbf{b}_g , and an additive white and Gaussian noise vector, $\boldsymbol{\eta}_g$, with Power Spectral Density (PSD), \mathbf{S}_g .

As suggested in [9], [14], and to comply with basic KF assumptions (whiteness and Gaussian distribution of the process noise), the rate gyro bias vector can be augmented to the ES-EKF state vector and predicted forward in time as:

$$\dot{\mathbf{b}}_g = -\frac{1}{\tau_{bg}} \mathbf{b}_g + \boldsymbol{\eta}_{bg}, \quad (6)$$

where τ_{bg} is the gyro bias correlation time, and $\boldsymbol{\eta}_{bg}$ is its driving noise, whose PSD, \mathbf{S}_{bg} , can be approximated as [9]:

$$\mathbf{S}_{bg} \approx \frac{2\sigma_{bg}^2 \ln 2}{\pi 0.4365^2 \tau_{bg}} \mathbf{I}_3, \quad (7)$$

where σ_{bg} is the rate gyro bias Standard Deviation (STD).

The prediction performance of the ES-EKF-based AHRS is highly dependent on the proper tuning of the parameters \mathbf{S}_g , τ_{bg} , and σ_{bg} . The AV method [11], [17], [18] has proven effective in helping designers determine appropriate values for these parameters, particularly through the use of optimization techniques such as WLS [9] and GA [13].

B. Update/Correction Phase

In order to make $\delta\boldsymbol{\Psi}_{nb}^n$ and \mathbf{b}_g separately observable/estimable in the ES-EKF-based AHRS, (reliable) measurements from accelerometers and magnetometers must be made available at a regular rate. In a TC implementation of such AHRS, what the accelerometers and magnetometers input to the ES-EKF are measurements (symbol $\tilde{(\cdot)}$) of the specific force acceleration, \mathbf{f}_{ib}^b , and magnetic field density, \mathbf{m}_m^b , vectors, respectively, which can be modeled, without loss of generality, as [2], [19]:

$$\tilde{\mathbf{f}}_{ib}^b = \mathbf{f}_{ib}^b + \mathbf{b}_a + \boldsymbol{\eta}_a, \quad (8)$$

$$\tilde{\mathbf{m}}_m^b = \mathbf{m}_m^b + \mathbf{b}_m + \boldsymbol{\eta}_m, \quad (9)$$

where \mathbf{b}_a and \mathbf{b}_m are instability bias/residual hard-iron magnetism vectors, respectively, and $\boldsymbol{\eta}_a$ and $\boldsymbol{\eta}_m$ are additive white and Gaussian noise vectors with PSDs \mathbf{S}_a and \mathbf{S}_m , respectively, both corrupting the outputs of the accelerometers and magnetometers.

The core of the ES-EKF-based AHRS update phase consists of forming a measurement innovation vector, $\delta\mathbf{z}$, that is linearly related to the error states. This can be done by subtracting the aiding sensor measurements from the corresponding AHRS-derived estimates, or vice-versa:

$$\delta\mathbf{z} = \begin{bmatrix} \tilde{\mathbf{f}}_{ib}^b - \hat{\mathbf{f}}_{ib}^b \\ \tilde{\mathbf{m}}_m^b - \hat{\mathbf{m}}_m^b \end{bmatrix}. \quad (10)$$

By modeling such measurement estimates as $\hat{\mathbf{f}}_{ib}^b = \mathbf{f}_{ib}^b + \delta\mathbf{f}_{ib}^b$ and $\hat{\mathbf{m}}_m^b = \mathbf{m}_m^b + \delta\mathbf{m}_m^b$, where $\delta\mathbf{f}_{ib}^b$ and $\delta\mathbf{m}_m^b$ are the corresponding estimation errors, and plugging (8) and (9) into (10) yields, after simplification:

$$\delta\mathbf{z} \equiv \begin{bmatrix} -\delta\mathbf{f}_{ib}^b + \mathbf{b}_a + \boldsymbol{\eta}_a \\ -\delta\mathbf{m}_m^b + \mathbf{b}_m + \boldsymbol{\eta}_m \end{bmatrix}. \quad (11)$$

KF theory [20] states that the measurement innovation vector should be a (linear) function of the state vector and white, Gaussian noises only. To accommodate (11) within such requirements, a common approach [2] is to augment the ES-EKF vector with the accelerometer and magnetometer biases, which could be propagated in time (within the ES-EKF prediction phase) according to dynamic models similar to (6) [9]. In doing so, it is obvious that the corresponding parameters τ_{ba} , τ_{bm} , \mathbf{S}_{ba} , and \mathbf{S}_{bm} become now part of the ES-EKF tuning problem. While AV-based approaches can generally perform well in helping the designer to determine suitable values for the accelerometer bias parameters, the latter is not exactly true for magnetometers, as their sensor bias stochastic characteristics may vary significantly on-the-fly, as a function of residual hard-iron magnetisms [4].

1) *Accelerometer Measurement Model:* In order to derive a suitable model that connect $\delta\mathbf{f}_{ib}^b$ in (11) with the attitude error state $\delta\boldsymbol{\Psi}_{nb}^n$, let us first recall how the rate of change of the vehicle velocity w.r.t. the Earth (subscript e) resolved on navigation frame axes, \mathbf{v}_{eb}^n , relate to the accelerometer specific force [1], [2]:

$$\dot{\mathbf{v}}_{eb}^n = \mathbf{C}_b^n \mathbf{f}_{ib}^b + \mathbf{g}_b^n - (\boldsymbol{\omega}_{en}^n \times + 2\boldsymbol{\omega}_{ie}^n \times) \mathbf{v}_{eb}^n, \quad (12)$$

where \mathbf{g}_b^n is the local plumb-bob gravity vector, $\boldsymbol{\omega}_{en}^n$ is the transport rate vector, and $\boldsymbol{\omega}_{ie}^n$ is the Earth rate vector, all resolved on the navigation frame axes [1].

Rearranging (12) yields:

$$\mathbf{f}_{ib}^b = (\mathbf{C}_b^n)^T [-\mathbf{g}_b^n + \dot{\mathbf{v}}_{eb}^n + (\boldsymbol{\omega}_{en}^n \times + 2\boldsymbol{\omega}_{ie}^n \times) \mathbf{v}_{eb}^n], \quad (13)$$

Equation (13) shows that, for an accurate estimation of the specific force measurement (needed in (10)), the velocity and rate of the vehicle must be known. However, and as previously stated, AHRSs are not aware of such information; hence, generally, the best approximation for $\hat{\mathbf{f}}_{ib}^b$ will be:

$$\hat{\mathbf{f}}_{ib}^b \approx -\hat{\mathbf{g}}_b^n, \quad (14)$$

which, of course, is accurate only when the vehicle is stationary or has a small velocity. The computation of the gravity vector in (14) is generally accomplished (via empirical models [21]) using an approximate location of the vehicle, which again, is not provided by the AHRS nor expected to vary significantly.

Back to the ES-EKF measurement model, linear perturbation [21] over (13) allow us to write:

$$\begin{aligned} \delta \mathbf{f}_{ib}^b &= (\delta \mathbf{C}_b^n)^T [-\mathbf{g}_b^n + \dot{\mathbf{v}}_{eb}^n + (\omega_{en}^n \times + 2\omega_{ie}^n \times) \mathbf{v}_{eb}^n] \\ &+ (\mathbf{C}_b^n)^T [-\delta \mathbf{g}_b^n + \delta \dot{\mathbf{v}}_{eb}^n + (\omega_{en}^n \times + 2\omega_{ie}^n \times) \delta \mathbf{v}_{eb}^n \\ &+ \mathbf{v}_{eb}^n (\delta \omega_{en}^n \times + 2\delta \omega_{ie}^n \times)], \end{aligned} \quad (15)$$

where δ represents an error in the subsequent variable.

From (4), one has, after rearrangement:

$$\delta \mathbf{f}_{ib}^b = (\mathbf{C}_b^n)^T [-\mathbf{g}_b^n + \dot{\mathbf{v}}_{eb}^n + (\omega_{en}^n \times + 2\omega_{ie}^n \times) \mathbf{v}_{eb}^n] \times \delta \Psi_{nb}^n + \eta_D, \quad (16)$$

with

$$\eta_D = (\mathbf{C}_b^n)^T [-\delta \mathbf{g}_b^n + \delta \dot{\mathbf{v}}_{eb}^n + (\omega_{en}^n \times + 2\omega_{ie}^n \times) \delta \mathbf{v}_{eb}^n + \mathbf{v}_{eb}^n (\delta \omega_{en}^n \times + 2\delta \omega_{ie}^n \times)]. \quad (17)$$

2) *Magnetometer Measurement Model*: To derive a model that connects $\delta \mathbf{m}_m^b$ in (11) with $\delta \Psi_{nb}^n$, let us consider now the scenario in which the only external magnetic field acting on the vehicle is the one created by the Earth, \mathbf{m}_e^n . In this case, it is straightforward to write:

$$\mathbf{m}_m^b = (\mathbf{C}_b^n)^T \mathbf{m}_e^n, \quad (18)$$

with

$$\mathbf{m}_e^n = \begin{bmatrix} \cos \alpha \cos \beta \\ \sin \alpha \cos \beta \\ \sin \beta \end{bmatrix} B, \quad (19)$$

where α , β , and B are the Earth's magnetic field declination angle, inclination angle, and total intensity, respectively, which can be accurately estimated from empirical models as a function of the vehicle location and time [22]. Again, as AHRSs do not provide us with location information, the best we can do is to employ an approximate location of the vehicle, which is not expected to vary significantly.

Applying linear perturbation [21] to (19) gives us

$$\delta \mathbf{m}_m^b = (\delta \mathbf{C}_b^n)^T \mathbf{m}_e^n + (\mathbf{C}_b^n)^T \delta \mathbf{m}_e^n, \quad (20)$$

which, from (4) and after rearrangement, yields:

$$\delta \mathbf{m}_m^b = [(\mathbf{C}_b^n)^T \mathbf{m}_e^n \times] \delta \Psi_{nb}^n + \eta_e \quad (21)$$

with $\eta_e = (\mathbf{C}_b^n)^T \delta \mathbf{m}_e^n$.

C. Complete State-Space Model

Once the TC ES-EKF-based AHRS prediction and measurement models have been established, substitution of (16) and (21) in (11) allows us to write, in state-space format:

$$\delta \dot{\mathbf{x}} = \begin{bmatrix} \mathbf{0}_3 & \mathbf{C}_b^n & \mathbf{0}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & -\mathbf{I}_3 & \mathbf{0}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & \mathbf{0}_3 & \frac{-\mathbf{I}_3}{\tau_{bg}} & \mathbf{0}_3 \\ \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{0}_3 & \frac{-\mathbf{I}_3}{\tau_{bm}} \end{bmatrix} \delta \mathbf{x} + \mathbf{G} \begin{bmatrix} \eta_g \\ \eta_{bg} \\ \eta_{ba} \\ \eta_{bm} \end{bmatrix} \quad (22)$$

$$\delta \mathbf{z} = \begin{bmatrix} \mathbf{H}_{11} & \mathbf{0}_3 & \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{H}_{21} & \mathbf{0}_3 & \mathbf{0}_3 & \mathbf{I}_3 \end{bmatrix} \delta \mathbf{x} + \begin{bmatrix} \eta_a - \eta_D \\ \eta_m - \eta_e \end{bmatrix}, \quad (23)$$

with:

$$\delta \mathbf{x} = [(\delta \Psi_{nb}^n)^T \quad \mathbf{b}_g^T \quad \mathbf{b}_a^T \quad \mathbf{b}_m^T]^T, \quad (24)$$

$$\mathbf{G} = \text{diag}(\mathbf{C}_b^n, \mathbf{I}_3, \mathbf{I}_3, \mathbf{I}_3), \quad (25)$$

$$\mathbf{H}_{11} = (\mathbf{C}_b^n)^T [-\mathbf{g}_b^n + \dot{\mathbf{v}}_{eb}^n + (\omega_{en}^n \times + 2\omega_{ie}^n \times) \mathbf{v}_{eb}^n] \times \quad (26)$$

$$\mathbf{H}_{21} = [(\mathbf{C}_b^n)^T \mathbf{m}_e^n \times] \quad (27)$$

where $\delta \mathbf{x}$ is the error state vector, and $\mathbf{0}_3$ are 3-by-3 zero matrices.

In addition to the sensor bias correlation times, the main ES-EKF-based ARHS tuning parameters are the process noise density matrix, \mathbf{S} , and measurement noise covariance matrix, \mathbf{R} , defined as follows:

$$\mathbf{S} = \text{diag}(\mathbf{S}_g, \mathbf{S}_{bg}, \mathbf{S}_{ba}, \mathbf{S}_{bm}), \quad (28)$$

$$\mathbf{R} = \text{diag}(\mathbf{R}_{ad}, \mathbf{R}_{me}), \quad (29)$$

where \mathbf{R}_{ad} and \mathbf{R}_{me} stand for the covariance matrices associated with the combined effects of accelerometer measurement noise (η_a)/vehicle dynamics uncertainty (η_D) and magnetometer measurement noise (η_m)/Earth magnetic field uncertainty (η_e), respectively.

Again, when information about the velocity and rate of the vehicle is not available, which is generally the case for AHRSs, an approximation for $\mathbf{H}_{11} \approx -(\mathbf{C}_b^n)^T \mathbf{g}_b^n \times$ has to be done based on an approximate information of the vehicle location. Moreover, as should be obvious from (17) and (29), whether the latter information is not available, nor is the information about the vehicle dynamics uncertainty, making the assessment of \mathbf{R}_{ad} very cumbersome. The net effect of making \mathbf{R}_{ad} too small, for instance, is having a noisy and jumpy AHRS attitude solution, especially for roll and pitch. On the contrary, making \mathbf{R}_{ad} too big may lead the latter to drift unbounded, as a consequence of the ES-EKF trusting too much on the rate gyros-propagated attitude.

D. Innovation Filtering

A solution that is sometimes adopted for addressing the latter issue (suitable tuning of the ES-EKF-based AHRS measurement covariance matrix) is to deploy innovation (or spike) filtering in parallel with the regular update phase of the ES-EKF [1]. Innovation filtering, in principle, determines whether the aiding sensor measurements are consistent with previous information. This is done by comparing the magnitude of each normalized measurement innovation, y_j , against a threshold (T_h) and rejecting the measurement if the latter is exceeded. In equations:

$$y_j = \frac{\delta z_j}{\sqrt{C_{\delta z_j, j}}} \quad (30)$$

where δz_j is the j -component of the measurement innovation vector $\delta \mathbf{z}$, and $C_{\delta z_j, j}$ is its variance, which is extracted from row/column j of the measurement innovation covariance matrix $\mathbf{C}_{\delta z}$, to be computed as [1]:

$$\mathbf{C}_{\delta z} = \mathbf{H} \mathbf{P} \mathbf{H}^T + \mathbf{R}, \quad (31)$$

where \mathbf{H} should be intuitive from (23), and \mathbf{P}^- is the ES-EKF predicted error state covariance matrix.

Where the normalized innovations have a zero-mean unit-variance Gaussian distribution, 99.72% of genuine measurements are passed by the IF set with a threshold of 3 [1]. In this sense, and back to the AHRS tuning problem, if one inputs data to the AHRS that correspond to a benign navigation phase of the vehicle, i.e., when the latter is not accelerating, not subject to external spurious magnetic fields, and at most, possesses small velocity, the analysis of the percentage of measurements that are rejected by a threshold of 3 would allow us to ascertain how far the tuning of our \mathbf{R} matrix is from the expected behavior.

Whether the experimental STD of the normalized measurement innovations is too below 1, which means that our \mathbf{R} was set above the expected, i.e., too high (please see (30) and (31)), it is likely that non-genuine measurements will not be rejected by the IF. On the contrary, if the normalized measurement innovation experimental STD is far above 1, that means that our \mathbf{R} is too small (refer to (30) and (31) again), and in this case, it is likely that the IF will reject genuine measurements inadvertently.

Hence, a judicious analysis of the normalized measurement innovation experimental STDs (collected during a benign navigation phase of the vehicle) can be useful for tuning properly either \mathbf{R} and/or T_h . Of course, the latter rationale is to be effective when the normalized measurement innovations have a Gaussian distribution, which is not always the case, particularly for AHRSs (please see (17) and (23) for details). Hence, proper tuning of the IF threshold becomes an additional encumbrance for the AHRS designer to handle. A proposed solution to this issue is discussed in sequence.

III. OPTIMIZATION APPROACHES FOR AHRS TUNING

Effective tuning of AHRS parameters significantly influences navigation performance. Given the complexity and multidimensionality of AHRS parameter spaces, traditional tuning techniques, such as empirical or manual adjustments, often fall short. In this paper, we explore three optimization methodologies designed to address this challenge: GP, NSGA-III, and MO-TPE. Each method provides unique strengths in handling complex search spaces and optimizing multiple, often conflicting, objectives.

A. Gaussian Processes

GPs are used as probabilistic surrogate models within Bayesian optimization to efficiently explore and exploit complex search spaces, aiming at identifying optimal hyperparameters for nonlinear filtering algorithms, such as the ES-EKF-based AHRS [23]. Given a set of candidate ES-EKF tuning parameters, GP models the objective function as a stochastic process characterized by a mean function and a covariance (kernel) function. Formally,

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')), \quad (32)$$

where $m(\mathbf{x})$ is the mean function (often taken as zero) and $k(\mathbf{x}, \mathbf{x}')$ is the covariance kernel between two inputs \mathbf{x}

and \mathbf{x}' , commonly chosen as the squared-exponential (RBF) kernel:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|^2}{2\ell^2}\right). \quad (33)$$

where σ_f^2 is the signal variance (kernel amplitude) and ℓ is the length-scale controlling the smoothness (correlation range) of the functions. Bayesian optimization iteratively updates the GP posterior and selects new evaluations by maximizing an acquisition function; in this work, we use Expected Improvement (EI) to balance exploration and exploitation:

$$EI(\mathbf{x}) = \mathbb{E} \left[\max(0, f_{\text{best}} - f(\mathbf{x})) \right], \quad (34)$$

where f_{best} denotes the best observed objective value so far.

B. Nondominated Sorting Genetic Algorithm III

NSGA-III is an evolutionary MO optimization algorithm particularly well-suited for complex nonlinear problems like tuning ES-EKF-based AHRS [24]. NSGA-III evolves a population of candidate solutions using genetic operators such as selection, crossover, and mutation. Solutions are ranked according to dominance into fronts based on Pareto optimality, where a solution \mathbf{x}_1 dominates another solution \mathbf{x}_2 if:

$$\mathbf{x}_1 \succ \mathbf{x}_2 \Leftrightarrow \forall i, f_i(\mathbf{x}_1) \leq f_i(\mathbf{x}_2) \wedge \exists j, f_j(\mathbf{x}_1) < f_j(\mathbf{x}_2). \quad (35)$$

C. Multi-Objective Tree-structured Parzen Estimator

MO-TPE is an extension of the standard TPE method tailored for MO optimization [25], [26]. MO-TPE uses probabilistic surrogate models to optimize parameters by partitioning the parameter space based on multiple performance objectives. Separate Parzen estimators model these partitions, creating distinct densities for each objective:

$$p(\mathbf{x}|y) = \begin{cases} l(\mathbf{x}), & \text{if } y < y^* \\ g(\mathbf{x}), & \text{if } y \geq y^* \end{cases}, \quad (36)$$

where y^* denotes quantile thresholds for each objective, separating high-performing observations from others, and $l(\mathbf{x}), g(\mathbf{x})$ are kernel density estimates. Candidate solutions are chosen based on an MO expected improvement function, simultaneously balancing improvements across all objectives.

IV. STOCHASTIC PARAMETERS FINE-TUNING

The stochastic parameter fine-tuning framework presented in this work addresses the challenge of optimally configuring ES-EKF-based AHRSs by searching through the parameter space to minimize attitude estimation errors. This Section outlines the practical implementation of the optimization process, including parameter selection, objective function formulation, and the iterative refinement methodology.

A. Optimization Parameter Space

The fine-tuning process focuses on the most critical stochastic parameters that directly influence AHRS performance, as identified from the ES-EKF state-space formulation presented in Section II. The optimization targets two main categories: (a) accelerometers/vehicle dynamics; and (b) magnetometers/Earth magnetic field, with their associated noise characteristics and IF thresholds.

For the first group, the optimization space encompasses:

- **Accelerometer/vehicle dynamics noise standard deviation (σ_{aD}):** Represents the white noise characteristics of the accelerometer measurements and vehicle dynamics uncertainty, with search bounds of $[0.01, 1.0]$ m/s² per axis.
- **Accelerometer bias standard deviation (σ_{ba}):** Captures the instability bias magnitude, bounded within $[1 \times 10^{-4}, 1 \times 10^{-2}]$ m/s² per axis.
- **Accelerometer bias correlation time (τ_{ba}):** Defines the time constant for bias evolution, ranging from $[10, 300]$ seconds per axis.
- **Accelerometer/vehicle dynamics IF threshold (T_{had}):** Sets the rejection threshold for spurious measurements, with bounds $[0.5, 15.0]$.

For the second group of parameters, one has:

- **Magnetometer/Earth magnetic field noise standard deviation (σ_{mE}):** Characterizes measurement noise, bounded within $[1 \times 10^{-8}, 1 \times 10^{-3}]$ Tesla per axis.
- **Magnetometer bias PSD (S_{bm}):** Represents random walk driving noise⁴, ranging from $[1 \times 10^{-18}, 1 \times 10^{-14}]$ T²/s per axis.
- **Magnetometer/Earth magnetic field threshold (T_{hme}):** Controls outlier rejection sensitivity, bounded within $[0.5, 15.0]$.

B. UAV Multi-Flight Objective Function Formulation

The optimization framework employs a robust UAV multi-flight objective function that aggregates performance metrics across diverse flight scenarios to ensure generalizability. For each parameter configuration $\theta = \{\sigma_{aD}, \sigma_{ba}, \tau_{ba}, T_{had}, \sigma_{mE}, S_{bm}, T_{hme}\}$, the AHRS estimator processes sensor data from multiple UAV flights, generating attitude estimates that are compared against high-precision Ground Truth (GT) measurements.

The individual attitude errors are computed using angular Root Mean Square (RMS) error metrics that properly handle angular wrapping. For each flight, individual error metrics are computed, and in sequence, the multi-flight aggregation strategy computes the mean error across all flights ($\bar{E}_{roll}(\theta)$, $\bar{E}_{pitch}(\theta)$, and $\bar{E}_{yaw}(\theta)$).

C. Algorithm-Specific Objective Formulations

The three optimization algorithms employ different objective function formulations to leverage their respective strengths:

⁴In this work, we opted for modeling the magnetometer biases as random walk processes, which is a particular case of (6), with $\tau_{bm} \rightarrow \infty$.

1) *GP Optimization:* For GP-based optimization, a single composite objective function is constructed using a weighted combination of the individual error metrics:

$$J_{GP}(\theta) = 0.33 \cdot \bar{E}_{roll}(\theta) + 0.33 \cdot \bar{E}_{pitch}(\theta) + 0.33 \cdot \bar{E}_{yaw}(\theta) \quad (37)$$

The equal weighting strategy ensures balanced optimization across all attitude components, preventing the algorithm from over-optimizing any single axis at the expense of others.

2) *MO Optimization:* Both NSGA-III and MO-TPE algorithms utilize the individual error components as separate objectives:

$$J_{MO}(\theta) = \begin{bmatrix} \bar{E}_{roll}(\theta) \\ \bar{E}_{pitch}(\theta) \\ \bar{E}_{yaw}(\theta) \end{bmatrix} \quad (38)$$

This formulation enables the algorithms to identify Pareto-optimal solutions that represent different trade-offs between roll, pitch, and yaw accuracy, providing the system designer with multiple optimal configurations to choose from based on application-specific requirements.

D. Optimization Workflow and Implementation

The parameter fine-tuning process follows a structured workflow designed for efficiency and robustness:

- 1) **Iterative parameter search:** Each optimization algorithm iterates through the parameter space according to its specific sampling strategy:
- 2) **AHRS estimation:** After having chosen a new set of parameters, the AHRS is ran to compute again the roll, pitch, and yaw errors, averaging them for every flight in the fine-tuning dataset.
- 3) **Robustness handling:** Failed AHRS estimations (due to numerical instability or convergence issues) are penalized with high error values (1000 deg), encouraging the algorithms to avoid unstable parameter regions.

This is repeated for the specified number of steps, or until the search algorithms reach an objective error lower bound.

V. EXPERIMENTS

This Section presents results based on data collected from 19 real flights conducted in May 2025. Out of these 19 flights, 13 were used to fine-tune the stochastic parameters, and 6 were used to validate the different performances. The experiments employed a Foxtech AYK-250 vertical take-off and landing UAV, equipped with an Orientus inertial measurement unit module from Advanced Navigation Pty Ltd., a three-axial RM3100 magneto-inductive sensor from Drotek company, and controlled with the autopilot PX4 version 1.14.3 [27]. Each flight lasted 45 minutes. Subsequently, each optimizer was run for 750 epochs on the collected data to ensure convergence and comprehensive exploration of the parameter space.

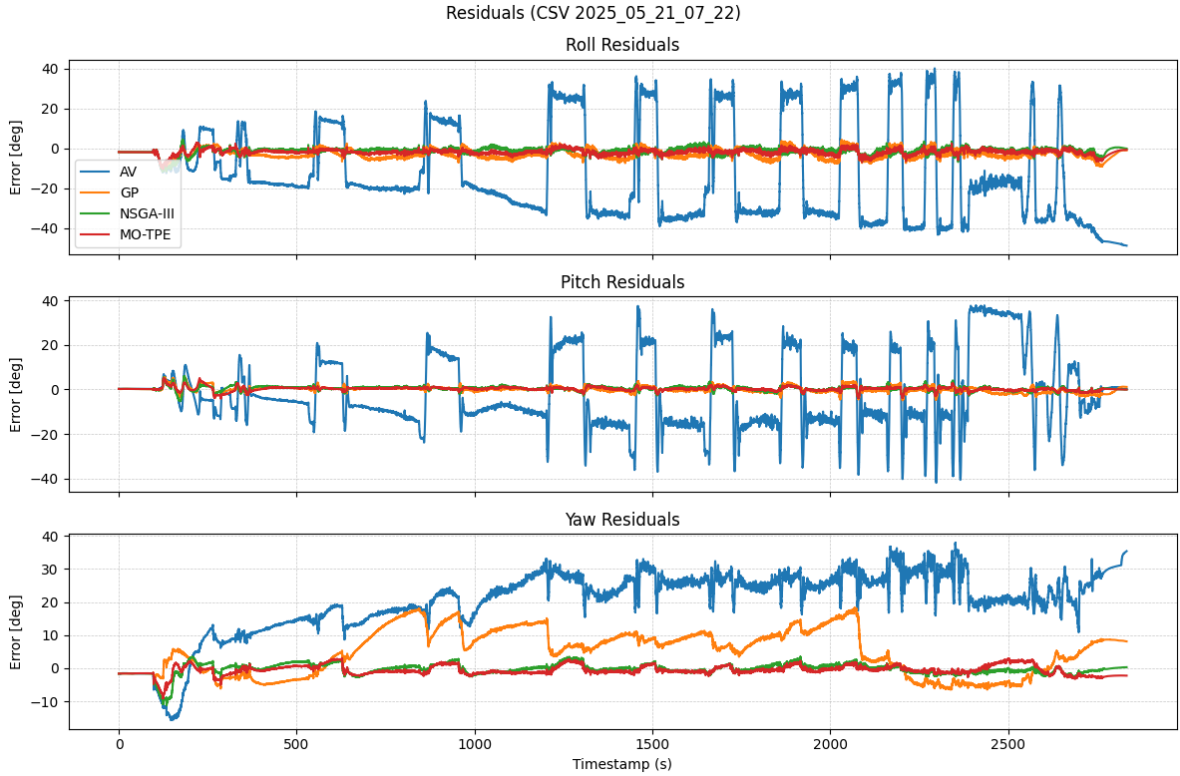


Fig. 1. Attitude estimation residuals for roll, pitch, and yaw angles comparing AV baseline with optimized parameter sets from GP, NSGA-III, and MO-TPE algorithms.

A. Performance Evaluation Methodology

The performance evaluation compared four approaches:

- **AV**: Traditional parameter identification using AV analysis as the baseline approach.
- **GP**: Single-objective Bayesian optimization targeting combined attitude error minimization.
- **NSGA-III**: Multi-objective evolutionary algorithm with reference-point-based selection.
- **MO-TPE**: Multi-objective optimization using probabilistic surrogate models.

The performance metrics were computed using three complementary error measures across roll, pitch, and yaw axes: RMS error, to emphasize large errors; Mean Absolute Errors (MAE), to provide robust central tendency measures; and STD, to quantify estimation consistency.

B. Attitude Estimation Results

Figure 1 illustrates the time-series attitude estimation residuals for a representative flight (CSV 2025_05_21_07_22), comparing all four optimization approaches against GT measurements. The residual plots demonstrate the temporal behavior of estimation errors across the roll, pitch, and yaw channels, revealing the dynamic performance characteristics of each optimization method.

The residual analysis reveals that the optimization-based approaches generally exhibit reduced amplitude variations compared to the AV baseline, particularly evident in the roll

and pitch channels. The yaw channel demonstrates the most significant improvements, with substantial reduction of error magnitudes and improved temporal consistency.

C. Multi-Flight Aggregated Performance

To assess the generalizability and robustness of the optimization approaches, Table I presents performance metrics averaged across all evaluation flights, providing a comprehensive assessment of each method's effectiveness across diverse flight conditions.

The aggregated results reveal significant performance improvements across all optimization methods compared to traditional AV tuning. The baseline AV approach, was consistently outperformed by the optimization-based methods with 80-90% error reductions across most metrics. Pitch

TABLE I
AVERAGE ATTITUDE ERRORS ACROSS ALL EVALUATION FLIGHTS, IN DEGREES.

Metric	Angle	AV	GP	NSGA-III	MO-TPE
RMS	Roll	10.94	2.15	1.83	1.90
	Pitch	10.19	1.02	0.97	0.95
	Yaw	11.69	4.77	2.51	2.26
MAE	Roll	9.19	1.89	1.61	1.66
	Pitch	8.65	0.83	0.8	0.79
	Yaw	10.41	4.00	1.89	1.78
STD	Roll	9.7	1.14	1.00	0.98
	Pitch	10.12	0.83	0.81	0.74
	Yaw	6.42	3.82	2.29	1.97

estimation showed the most remarkable gains with over 90% improvement, achieving sub-degree accuracy across all optimizers. Roll estimation similarly benefited from 80-90% error reductions, while yaw—historically the most challenging axis due to magnetic disturbances—demonstrated substantial 70-80% improvements despite remaining the most error-prone channel.

Among the optimization approaches, MO-TPE emerged as the clear leader, delivering superior performance across most metrics through its ability to simultaneously balance trade-offs between all attitude axes. NSGA-III proved particularly effective for roll estimation, while GP provided consistent, well-balanced improvements across all channels.

VI. CONCLUSION

This paper showed that optimization-based fine-tuning of stochastic parameters delivers substantial gains with ES-EKF-based AHRS over traditional AV identification. Using multi-flight evaluation on real UAV data, all three optimizers—GP, NSGA-III, and MO-TPE—consistently reduced attitude errors across roll, pitch, and yaw. On average, MO-TPE achieved the best overall performance, NSGA-III was a close second, and GP provided competitive single-objective results. Besides lowering RMS and MAE, the methods reduced error variability, indicating improved robustness to vehicle dynamics and residual magnetic disturbances.

Limitations include the dependence on representative training data and the offline nature of the present workflow, as well as modeling approximations (e.g., approximate gravity/magnetic field and non-Gaussian innovations).

Future work will address (a) on-board, compute-aware online adaptation; (b) risk-aware objectives and robust likelihoods for heavy-tailed innovations; (c) joint tuning of filter parameters and IF thresholds in closed loop; (d) generalizability of the approach to other state estimation methods (e.g., particle filter); (e) investigation and comparison of other optimizers; and (f) the generalizability to other autonomous vehicle families.

REFERENCES

- [1] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed. Artech House Publishers, 2013.
- [2] J. A. Farrell, *Aided Navigation: GPS with High Rate Sensors*. McGraw-Hill, Inc., 2008.
- [3] J. Guerrero-Castellanos, H. Madrigal-Sastre, S. Durand, N. Marchand, W. Guerrero-Sánchez, and B. Salmerón, "Design and implementation of an attitude and heading reference system (AHRS)," in *2011 8th International Conference on Electrical Engineering, Computing Science and Automatic Control*, 2011, pp. 1–5.
- [4] R. Paes Menezes Filho, F. Oliveira e Silva, G. Esau Villalobos Hernandez, A. Wakode, I. Campos, and C. de Souza, "Two-dimensional solution for the analytical extended two-step magnetometer calibration," *IEEE Sensors Journal*, vol. 25, no. 7, pp. 10 879–10 893, 2025.
- [5] M. Euston, P. Coote, R. Mahony, J. Kim, and T. Hamel, "A Complementary Filter for Attitude Estimation of a Fixed-Wing UAV," in *2008 IEEE/RSJ International Conf. on Intelligent Robots and Systems*, Nice, France, September 2008, pp. 340–345.
- [6] M. Blachuta, R. Grygiel, R. Czyba, and G. Szafranski, "Attitude and heading reference system based on 3D complementary filter," in *2014 19th International Conference on Methods and Models in Automation and Robotics (MMAR)*, 2014, pp. 851–856.
- [7] R. Mahony, T. Hamel, and J.-M. Pfimlin, "Nonlinear complementary filters on the special orthogonal group," *IEEE Transactions on Automatic Control*, vol. 53, no. 5, pp. 1203–1218, 2008.
- [8] L. Chang, F. Zha, and F. Qin, "Indirect Kalman filtering based attitude estimation for low-cost attitude and heading reference systems," *IEEE/ASME Transactions on Mechatronics*, vol. 22, no. 4, pp. 1850–1858, 2017.
- [9] J. A. Farrell, F. O. Silva, F. Rahman, and J. Wendel, "Inertial measurement unit error modeling tutorial: Inertial navigation system state estimation with real-time sensor calibration," *IEEE Contr. Syst. Mag.*, vol. 42, no. 6, pp. 40–66, 2022.
- [10] E. Aghapour, F. Rahman, and J. A. Farrell, "Outlier accommodation in nonlinear state estimation: A risk-averse performance-specified approach," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 2, pp. 567–579, 2021.
- [11] D. Allan, "Statistics of atomic frequency standards," *Proc. of the IEEE*, vol. 54, no. 2, pp. 221–230, 1966.
- [12] J. Hidalgo-Carrió, S. Arnold, and P. Poulakis, "On the design of attitude-heading reference systems using the Allan variance," *IEEE T. Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 63, no. 4, pp. 656–665, 2016.
- [13] V. M. G. B. Cavalcanti, F. O. Silva, and D. A. d. Lima, "Genetic algorithm-based tunings for baro-fused inertial navigation systems," in *2024 27th International Conference on Information Fusion (FUSION)*, 2024, pp. 1–8.
- [14] V. M. G. B. Cavalcanti, F. O. Silva, C. R. C. Durão, G. E. V. Hernandez, T. C. Ferrari, C. Souza Jr., and J. A. Farrell, "On the impact of optimized stochastic error modeling on the performance of EKF-based baro-aided INS," in *2025 64th IEEE Conference on Decision and Control (CDC)*, vol. accepted, 2025, pp. 1–8.
- [15] J. Yang, T. Du, X. Liu, B. Niu, and L. Guo, "Method and implementation of a bioinspired polarization-based attitude and heading reference system by integration of polarization compass and inertial sensors," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 11, pp. 9802–9812, 2020.
- [16] P. G. Savage, "A unified mathematical framework for strapdown algorithm design," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 2, pp. 237–249, 2006.
- [17] "IEEE standard specification format guide and test procedure for linear single-axis, nongyroscopic accelerometers," *IEEE Std 1293-2018 (Revision of IEEE Std 1293-1998)*, pp. 1–271, 2019.
- [18] "IEEE standard for specifying and testing single-axis interferometric fiber optic gyros," *IEEE Std 952-2020 (Revision of IEEE Std 952-1997)*, pp. 1–93, 2021.
- [19] Á. H. A. Maia, F. O. Silva, D. A. de Lima, and R. P. Menezes Filho, "Identificação e modelagem de erros estocásticos em sensores inerciais via variância de Allan e otimização," in *SBAI 2023*, 2022, pp. 1–8.
- [20] M. S. Grewal and A. P. Andrews, *Kalman Filtering: Theory and Practice Using MATLAB*, 4th ed. Wiley-IEEE Press, 2014.
- [21] P. G. Savage, *Strapdown Analytics*. Strapdown Associates, 2007.
- [22] A. Chulliat, W. Brown, P. Alken, C. Beggan, M. Nair, G. Cox, A. Woods, S. Macmillan, B. Meyer, and M. Paniccia, "The US/UK world magnetic model for 2020–2025: Technical report," National Centers for Environmental Information, NOAA and British Geological Survey, Technical Report, 2020. [Online]. Available: <https://doi.org/10.25923/ytk1-yx35>
- [23] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [24] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints," *IEEE Transactions on Evolutionary Computation*, vol. 18, no. 4, pp. 577–601, 2014.
- [25] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 24, 2011, pp. 2546–2554.
- [26] Y. Ozaki, Y. Tanigaki, S. Watanabe, and M. Onishi, "Multiobjective tree-structured parzen estimator for computationally expensive optimization problems," in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, ser. GECCO '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 533–541. [Online]. Available: <https://doi.org/10.1145/3377930.3389817>
- [27] L. Meier, D. Honegger, and M. Pollefeys, "PX4: A node-based multithreaded open source robotics framework for deeply embedded platforms," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 6235–6240.