Efficient Simultaneous Calibration of a Magnetometer and an Accelerometer

Conrado S. Miranda, Janito V. Ferreira

***Abstract*—This paper describes a calibration algorithm to simultaneously calibrate a magnetometer and an accelerometer without any information besides the sensors readings. Using a linear sensor model and maximum likelihood cost, the algorithm is able to estimate both sensors' biases, gains, and covariances, besides sensor orientations and Earth's fields. Results show errors of less than 0.1 standard deviations in simulation, and high-quality estimates with real sensors even when the algorithm's assumptions are violated.**

***Index Terms*—Maximum likelihood; Parameter estimation; Sensor calibration.**

I. Introduction

SENSOR reading is an important step in a robot's control loop. However, the values obtained may be inaccurate because of incorrect sensor calibration. In the field of aerial vehicles, accelerometers and magnetometers are frequently used to estimate attitude [1]. As both these types of sensors are calibrated using Earth's gravitational or magnetic field as reference, most research into calibration of these devices focuses on magnetometer calibration, which can also be used to calibrate accelerometers in stationary settings. This work focuses on batch algorithms, that is, algorithms that collect a data set and then calibrate the sensors, which can be used to provide the initial conditions for online methods that solve the calibration problem by filtering [2] or iterating [3], [4].

A common approach to calibration is to assume that the magnetic field's norm is known or unitary, and then adjust the gain and bias so that the readings match this value. [5] extended the TWOSTEP algorithm, which tries to minimize the difference between the field's norm and the sensed norm, to calibrate both the bias and gain. [6] proved that noiseless magnetometer readings, including soft- and hard-iron effects, occur on the surface of an ellipsoid manifold, and then developed an algorithm to calibrate the bias and gain.

Although both methods can be used to calibrate accelerometers as well as magnetometers, the calibration must be performed independently, so the magnetometer and accelerometer are calibrated in different frames, and the rotation between them is unknown. Furthermore, neither method is able to estimate the direction of the magnetic field, which may be unknown and is essential for attitude estimation using magnetometers [2], [7], [8].

[9] partially solves this issue by calibrating the magnetometer while also capturing a reference for the z direction, provided by some independent filter. This allows the direction of the Earth's magnetic field to be estimated, and the calibration to be performed in some previously known reference frame, while also estimating the magnetometer's bias and gain. However, a reference orientation, whether reliable or not, is not always available, thus limiting the use of algorithm.

[10] tried to solve both problems by simultaneously calibrating the magnetometer and accelerometer, allowing the magnetic field, the attitude, and all the sensors' parameters to be estimated. This eliminates the need for a reference so that the calibration can be performed from sensor readings alone, that is, no external apparatus or knowledge besides the sensor readings is required. However, the problem is replete with low-quality local minima, which sometimes significantly degrade the calibration performance, and the solution has a very high computational cost, making the algorithm impractically slow.

The novelties in this paper are the improvement over [10] made possible by the use of better initial estimates of the parameters, leading to better minima, and the approximation for the rotation estimation, which also helps to avoid poor minima and reduce computational time because of its closed-form solution, making it suitable for use in practice. The initial conditions are given by an adaptation of [9], where the z reference is replaced by an estimate computed from the accelerometer. To our knowledge, the algorithms described in the present work and in [10] are the only accelerometer and magnetometer calibration algorithms that are able to perform the complete calibration from collected data alone. Moreover, the algorithm proposed here is also the most statistically correct and generic solution based on linear models since all the parameters (gains, biases, Earth's fields, and orientations) must be estimated at the same time.

The algorithm is tested in simulated and real sensors. The simulations show that the algorithm correctly calibrates different sets of sensor parameters and allows comparison between the estimates and the ground truth. Several variations of the algorithm are tested, and the best-performing one is also used to calibrate the real sensors. The data capture was designed to violate the basic assumptions underlying the algorithm so that its robustness to adverse real-world conditions could be tested. It is shown that even in this imperfect setting the proposed algorithm is able to achieve high-quality calibration.

The paper is organized as follows. Sections II and III describe the sensor model used and the collection and preprocessing of the data, respectively. Based on the sensor model, a maximum likelihood cost function is described in Section IV. The initial parameters' estimates are given in Section V, and the optimization steps are described in Section VI. The results of simulated and real-world experiments are reported in Section VII and are used to evaluate the proposed method.

C. S. Miranda is with the School of Electrical and Computer Engineering, University of Campinas, Brazil. E-mail: conrado@dca.fee.unicamp.br.

J. V. Ferreira is with the School of Mechanical Engineering, University of Campinas, Brazil. E-mail: janito@fem.unicamp.br.

The concluding remarks are presented in Section VIII.

II. Sensor Model

The linear sensor model with Gaussian noise, where the correct value is scaled by a matrix and added to a bias and a noise, is frequently used in the literature [5], [6] and is also used in this paper. The value $v_s \in \mathbb{R}^3$ read by the sensor s, where $s = a$ and $s = m$ for the accelerometer and magnetometer, respectively, is given by:

$$v_s = K_s v_s^* + b_s + \epsilon_s, \quad \epsilon_s \sim \mathcal{N}(0, \Sigma_s), \tag{1}$$

where $K_s \in \mathbb{R}^{3\times3}$ and $b_s \in \mathbb{R}^3$ are the gain matrix and bias vector associated with the sensor, $v_s^* \in \mathbb{R}^3$ is the nominal value to be read, and $\Sigma_s \in \mathcal{S}^3_{++}$ is the noise covariance, which is a positive-definite matrix. It is important to note that the linear sensor model for the magnetometer is also able to handle soft- and hard-iron effects, which can be embedded in the sensor's gain and bias [6].

Conversely, given the parameters and a value read by the sensor, the nominal value can be estimated as:

$$v_s^* = K_s^{-1}(v_s - b_s) + \epsilon_s^*, \quad \epsilon_s^* \sim \mathcal{N}(0, K_s^{-1}\Sigma_s K_s^{-T}),$$

where it is assumed that the gain matrix K_s is invertible. If it is not, then its columns are not linear independent and the measured values v_s must lie in a reduced subspace spanned by these columns. Therefore, only the case where K_s is invertible is considered, and the nominal measurement v_s^* can be recovered.

The nominal value for a sensor is given by the Earth's respective field rotated by the current orientation of the sensor frame, so that $v_s^* = R v_s^{\mathcal{I}}$, where $v_s^{\mathcal{I}}$ is the field value for a given sensor in the fixed inertial frame $\mathcal{I}$ and $R \in SO(3)$ is the rotation representing the current sensor orientation. Since the choice of $\mathcal{I}$ is arbitrary, one can choose it such that the gravity g only has a negative component in $\vec{z}$ and the magnetic field h has a positive component in $\vec{x}$, such that

$$g = \begin{bmatrix}0 & 0 & g_z\end{bmatrix}^T, \quad h = \begin{bmatrix}h_x & 0 & h_z\end{bmatrix}^T, \tag{2a}$$

$$g_z < 0, \quad h_x > 0, \tag{2b}$$

$$s_a^{\mathcal{I}} = g, \quad s_m^{\mathcal{I}} = h. \tag{2c}$$

It should be noted that the assumption that the magnetic field component h_x is greater than zero can be satisfied almost everywhere except at the magnetic poles, where the magnetic field is perpendicular to the Earth's surface and the magnetometer does not provide more orientation information than the accelerometer does for the gravitational field. In this particular case, one cannot use this combination to perform the calibration, but one can do so in all other contexts, as will be shown.

As discussed in [10] and highlighted later, the sensor frame also has multiple possible definitions, since the gain matrix can be written as $K_s = RK_s'$, where R is a rotation matrix. Therefore, some constraints must be fixed to make the solution unique. This paper assumes that K_a is triangular, which is a sufficient assumption to define a single solution. Moreover, K_m is allowed to have any value, so that it can incorporate the misalignment between the accelerometer and magnetometer.

Since during the calibration the algorithm does not use information about the body frame, and this may not be available if the calibration is being performed outside the desired body, other algorithms such as [11] can be used to align the sensor and body frame.

III. Data Preprocessing

Assuming the linear model introduced in Sec. II and that the sensors can be held relatively still when readings are being taken so that a set of sensor measurements have the same nominal value, the measured values can be written as:

$$v_s[i,j] \sim \mathcal{N}(\mu_s[i], \Sigma_s), \quad \mu_s[i] = K_s R_i v_s^{\mathcal{I}} + b_s, \tag{3a}$$

$$j \in \{1, 2, \ldots, \Delta_s[i]\}, \quad i \in \{1, 2, \ldots, N\}, \tag{3b}$$

where N is the number of measurement sets, $\Delta_s[i]$ is the number of data points collected for the i-th set, $\mu_s[i]$ is the mean value shared by the measurements in the i-th set, and R_i is the orientation at which the i-th set was collected.

Since the measurements in each set are assumed to have sampled the same information, the mean values $\mu_s[i]$, which are different for each set, and the covariance matrix Σ_s, which is shared among all sets, can be estimated.

The minimum variance unbiased estimator for $\mu_s[i]$ is the sample mean [12], given by:

$$\hat{\mu}_s[i] = \frac{1}{\Delta_s[i]} \sum_{j=1}^{\Delta_s[i]} v_s[i,j], \tag{4}$$

and has a distribution described by:

$$\hat{\mu}_s[i] \sim \mathcal{N}\left(\mu_s[i], \frac{\Sigma_s}{\Delta_s[i]}\right). \tag{5}$$

Since the sample mean $\hat{\mu}_s[i]$ characterizes the values for the i-th set, it can be used instead of them as a single measurement with appropriate weight, associated with the number of measurements $\Delta_s[i]$, during optimization to decrease the computing time.

The minimum variance unbiased estimator for the noise covariance is given by:

$$\hat{\Sigma}_s = \frac{\sum_{i=1}^{N}\sum_{j=1}^{\Delta_s[i]} \left(v_s[i,j] - \hat{\mu}_s[i]\right)\left(v_s[i,j] - \hat{\mu}_s[i]\right)^T}{\left(\sum_{i=1}^{N} \Delta_s[i]\right) - N}, \tag{6}$$

where the subtraction of N from the denominator is a generalization of the Bessel's correction [13] for N measurement sets, instead of 1 as in the original correction.

IV. Cost Function

Given that it is assumed that only the measured values are known, the full parameter set is given by all the sensors' and fields' parameters and all orientations. This set can be written as:

$$\Theta = \{\Sigma_a, K_a, b_a, g_z, \Sigma_m, K_m, b_m, h_x, h_z, R_1, R_2, \ldots, R_N\}.$$

Note that this is the most generic setting possible for this problem with the assumption of affine measurements, since all parameters are considered unknown and any previously known parameter can be fixed to the known value.

Using the model defined in Eq. (3), the negative log-likelihood can be used to create a cost function [14], which is given by:

$$J(\Theta; s_a, s_m) = \sum_{s\in\{a,m\}} \sum_{i=1}^{N} \sum_{j=1}^{\Delta_s[i]} \left(\ln(|\Sigma_s|) + \tilde{v}_s[i,j]^T \Sigma_s^{-1} \tilde{v}_s[i,j]\right) \tag{7a}$$

$$\tilde{v}_s[i,j] = K_s R_i v_s^{\mathcal{I}} + b_s - v_s[i,j]. \tag{7b}$$

In this format, the calibration problem becomes a traditional minimization problem over all the parameters in Θ, where the restrictions over Σ_s, g_z, h_x, and R_i must be satisfied.

A simplified cost function can be defined on a smaller parameter set $\Theta' = \Theta \setminus \{\Sigma_a, \Sigma_m\}$, where the covariances are replaced by their estimates computed using Eq. (6). In this case, the maximum likelihood cost can be simplified to

$$J'(\Theta'; \hat{\mu}_a, \hat{\mu}_m, \hat{\Sigma}_a, \hat{\Sigma}_m) = \sum_{i=1}^{N} \sum_{s\in\{a,m\}} \Delta_s[i] \tilde{\mu}_s[i]^T \hat{\Sigma}_s^{-1} \tilde{\mu}_s[i], \tag{8a}$$

$$\tilde{\mu}_s[i] = K_s R_i v_s^{\mathcal{I}} + b_s - \hat{\mu}_s[i], \tag{8b}$$

in which the real measurements $v_s[i,j]$ can be replaced by their sample means $\hat{\mu}_s[i]$ with weight $\Delta_s[i]$, which greatly reduces the computational cost.

Since both cost functions are subject to many local minima, which can lead to poor estimates as shown in [10], a good initial estimate must be provided. Then, instead of optimizing all the parameters at once, it will be shown that optimizing subsets one at a time greatly simplifies the problem. The next two sections deal with these problems.

V. Initial Estimate

This section is an adaptation of the work presented in [9], where a magnetometer was calibrated using a known correct reference measure of the z direction. Since it is assumed that no such z is available, this reference is replaced here by the estimated accelerometer's field value, which only has a negative component in the z direction.

Furthermore, the constraint that the magnetometer cannot have mirrored axis relative to the inertial axis is relaxed, which was not possible in the original reference. The reader is referred to the original paper for a more detailed derivation of the equations and implementation details.

Assuming that the field $v_s^{\mathcal{I}}$ is unitary and that the noise ϵ_s can be neglected, the following approximation can be performed:

$$0 \approx \hat{\mu}_s[i]^T A \hat{\mu}_s[i] + b^T \hat{\mu}_s[i] + c \tag{9a}$$

$$A = K_s^{-T} K_s^{-1} \tag{9b}$$

$$b = -2 b_s^T K_s^{-T} K_s^{-1} \tag{9c}$$

$$c = b_s^T K_s^{-T} K_s^{-1} b_s - 1, \tag{9d}$$

which follows directly from taking the norm of both sides of Eq. (1) and replacing the single measurements by their mean given by Eq. (4), which has smaller variance and thus is a better estimate.

By separating the known terms in Eq. (9a) from the unknown terms and considering all N sample sets, the approximation can be written as:

$$J\eta = \begin{bmatrix} J_1 \\ \vdots \\ J_N \end{bmatrix} \eta \approx 0 \tag{10}$$

$$J_i = \begin{bmatrix} \hat{\mu}_s[i]^T \otimes \hat{\mu}_s[i]^T & \hat{\mu}_s[i]^T & 1 \end{bmatrix}$$

$$\eta = \begin{bmatrix} \text{vec}A \\ b \\ c \end{bmatrix}$$

where $\otimes$ denotes the Kronecker product and vec denotes the vectorization operator. Assuming that A is symmetric, the nontrivial solution $\hat{\eta}$ that minimizes the approximation error is given by the right eigenvector corresponding to the smallest singular value of J [9].

Since any $\alpha\hat{\eta}$ is also a solution to the approximation in Eq. (10), one can find the correct α by manipulating the elements in Eq. (9), which provides the correct scale. Hence the correct value is given by:

$$\alpha = \left(\frac{1}{4}\hat{b}^T \hat{A}^{-1} \hat{b} - \hat{c}\right)^{-1}.$$

From the definitions in Eq. (9) and the estimates $\hat{\eta}$ and α, the gain and bias must satisfy

$$\hat{K}_s^{-T} \hat{K}_s^{-1} = \alpha \hat{A} \tag{11}$$

$$\hat{b}_s = -\frac{1}{2} \hat{A}^{-1} \hat{b},$$

which gives the bias estimate for each sensor directly.

The gain $\hat{K}_s$ can not be uniquely determined, as any $\hat{K}_s R$ with $RR^T = I_3$ is also a solution to Eq. (11). If the upper triangular solution obtained by the Cholesky decomposition is denoted by $\tilde{K}_s$, the estimates can be written as:

$$\hat{K}_m = \tilde{K}_m R, \quad \hat{K}_a = \tilde{K}_a \tag{12}$$

where it is assumed that only $\hat{K}_m$ is subject to rotation as the rotation between the two gains is fixed and there would be multiple solutions if both had full unknown rotations.

So far, the solution described in [9], which provides estimates for the sensor bias and gain, has been used to compute the initial estimates. This solution could be applied here since the accelerometer and magnetometer could be considered decoupled problems. However, in Eq. (12), $\hat{K}_a$ can be estimated directly while $\hat{K}_m$ depends on the unknown rotation R. Since only the magnetometer was calibrated in [9], this was not an issue as the ground-truth z direction was available. Moreover, the magnetic field component h_z could also be estimated directly. To solve this problem, the ground truth z was replaced by the estimate of the nominal value v_a^* from the accelerometer, since this is already calibrated.

Therefore, estimates for the rotation R between the accelerometer and the magnetometer and the component h_z of

the magnetic field can be found simultaneously by solving the following optimization problem:

$$\min_{R,h_z} \quad \frac{1}{2}\sum_{i=1}^{N}\left(\Delta_a[i]+\Delta_m[i]\right)\left\|h_z + z_a[i]^T R^T z_m[i]\right\|_2^2 \tag{13a}$$

$$\text{s.t.} \quad z_s[i] = \tilde{K}_s^{-1}\left(\hat{\mu}_s[i] - \hat{b}_s\right) \tag{13b}$$

$$R \in SO(3) \tag{13c}$$

where Δ_a and Δ_m are used to weight the number of samples in each data set. As stated earlier, this equation is similar to the one in [9], but the ground truth z is replaced by the estimates $z_a[i]$.

Test showed that, unlike in [9], this problem is subject to a few low-quality local minima because of the replacement of precise knowledge of the z direction by the accelerometer estimate and that the initial condition $R = I_3$ and $h_z = 0$ sometimes leads to poor estimates. Hence, an optimization with random restarts was performed and the initial conditions were given by a random rotation as described in [15] and $h_z \sim \mathcal{U}([-1,1])$. Tests showed that at most 100 iterations were needed to find good initial estimates.

Since it was assumed that the fields were unitary, it follows that

$$\hat{h}_x = \sqrt{1-\hat{h}_z^2}.$$

As discussed in [10], having the fields' and gains' scales as variables leads to multiple solutions, since the gain can be multiplied by a constant while the field is divided by the same value. Therefore, in this paper the value $g_z = -1$ and $h_x = 1$ are adopted, and the magnetometer estimates must be adjusted as:

$$\hat{K}'_m = \hat{K}_m \hat{h}_x, \quad \hat{h}'_z = \hat{h}_z/\hat{h}_x,$$

while the accelerometer does not have to be adjusted because it already has a unitary field.

Finally, the rotation R_i associated with each sample set is given by the eigenvector corresponding to the minimal eigenvalue of B_i [16], where B_i is defined as:

$$B_i = w_a A_{a,i}^T A_{a,i} + w_m A_{m,i}^T A_{m,i} \tag{14a}$$

$$A_{s,i} = f(\hat{s}^*_{s,i} + \hat{s}^{\mathcal{I}}_s, \hat{s}^*_{s,i} - \hat{s}^{\mathcal{I}}_s) \tag{14b}$$

$$\hat{s}^*_{s,i} = \hat{K}_s^{-1}\left(\hat{\mu}_s[i] - b_s\right), \; w_s = |\hat{K}_s^T \hat{\Sigma}_s^{-1} \hat{K}_s|^{\frac{1}{3}} \tag{14c}$$

$$f(x,y) = \begin{bmatrix} 0 & -y_1 & -y_2 & -y_3 \\ y_1 & 0 & -x_3 & x_2 \\ y_2 & x_3 & 0 & -x_1 \\ y_3 & -x_2 & x_1 & 0 \end{bmatrix} \tag{14d}$$

and the weights w_a and w_m introduce uncertainty weighting.

All initial estimates therefore have closed-form efficient solutions except for R and h_z. However, Eq. (13) has few decision variables and can be efficiently optimized by random restarts and gradient descent.

VI. Optimization Steps

As the cost functions in Eqs. (7) and (8) have many joint parameters, the parameter set is split into four disjoint subsets that can be solved efficiently one at a time. The optimization starts with the estimates provided in Sec. V and then iterates each subset optimization while keeping the other parameters fixed until convergence is achieved. Convergence is defined by $J_{k-1} - J_k < \gamma$, where J_k is the cost after the k-th iteration and γ is the stop parameter. Since the true values for any parameter in Θ are not available, the hat notation is dropped where convenient to simplify the expressions.

This section is an extension of the optimization presented in [10]. The main differences are, firstly, the approximation for the rotation estimation described in Sec. VI-A. This uses Eq. (14) to compute an efficient approximation to the real orientation during the optimization, which is not present in [10], leading to a significant slow-down of the algorithm and worse results. Secondly, the format of the gain matrix K_a computed in Sec. VI-C is assumed to be triangular here because its initial condition is given by a Cholesky decomposition while in [10] it is assumed to be symmetric. The main difference between the method proposed here and the one described in [9], apart from the fact that two sensors are used and no ground-truth is available in the former, is that the optimization is performed in steps with efficient solutions, while in [9] the initial estimate is computed and then all parameters are optimized at the same time, preventing the inherent structure of the problem from being exploited.

A. Rotation Estimation

From Eq. (8), it is clear that the cost for the i-th rotation is given by:

$$J'_{R_i} = \sum_{s\in\{a,m\}} \Delta_s[i]\tilde{\mu}_s[i]^T \Sigma_s^{-1}\tilde{\mu}_s[i], \tag{15}$$

$$\tilde{\mu}_s[i] = K_s R_i s_s^{\mathcal{I}} + b_s - \hat{\mu}_s[i].$$

As a closed-form solution could not be found, the problem is optimized directly through gradient descent using a quaternion to represent the rotation. However, the approximation used to compute the initial estimates, which is given by the eigenvector associated with the minimum eigenvalue of B_i from Eq. (14), can be used to speed-up the initial steps of the optimization, when the estimates are still far from the true values. The performance of this approximation approximation is evaluated in the experimental section.

B. Bias and Field Estimation

Since the sensor readings are assumed linear in the biases and fields and have Gaussian noise, the problem can be posed as a generalized least squares (GLS) problem [17] given by finding β to approximate

$$y = X\beta + \epsilon, \; \mathrm{E}[\epsilon|X] = 0, \; \mathrm{Var}[\epsilon|X] = \Omega,$$

whose solution is given by:

$$\hat{\beta} = \left(X^T\Omega^{-1}X\right)^{-1} X^T\Omega^{-1}y.$$

Writing $\hat{\mu}_a[i]$ in the GLS format using Eqs. (3) and (5) gives

$$\begin{bmatrix}\hat{\mu}_a[1]\\ \vdots \\ \hat{\mu}_a[N]\end{bmatrix} = \begin{bmatrix} K_a R_1 \vec{z} & I_3 \\ \vdots & \vdots \\ K_a R_N \vec{z} & I_3 \end{bmatrix} \begin{bmatrix} g_z \\ b_a \end{bmatrix} + \epsilon,$$

$$\Omega = \mathrm{diag}\left(\frac{\Sigma_a}{\Delta_a[1]}, \ldots, \frac{\Sigma_a}{\Delta_a[N]}\right),$$

where $\vec{z}$ is the unitary vector in the z direction. As Sec. V established that $g_z = -1$, the gain matrix is adjusted to ensure that this equality is still valid. The estimate is therefore adjusted so that $\hat{K}'_a = -\hat{K}_a \hat{g}_z$, where $\hat{g}_z$ is the estimated gravity.

The magnetometer estimate $\hat{\mu}_m[i]$ can be written in a similar fashion, allowing $\hat{h}_x$, $\hat{h}_z$ and $\hat{b}_m$ to be estimated. As $h_x = 1$, it follows that $\hat{K}'_m = \hat{K}_m \hat{h}_x$ and $\hat{h}'_z = \hat{h}_z / \hat{h}_x$, which maintains the scale.

C. Gain Estimation

The gains can also be written in the GLS format using Eqs. (3) and (5). Assuming that K_a is upper triangular, which is valid for the initial estimate since it comes from a Cholesky decomposition, one has that:

$$\begin{bmatrix} \hat{\mu}_a[1] - b_a \\ \vdots \\ \hat{\mu}_a[N] - b_a \end{bmatrix} = \begin{bmatrix} G_1 \\ \vdots \\ G_N \end{bmatrix} \begin{bmatrix} K_{a,11} \\ K_{a,12} \\ K_{a,13} \\ K_{a,22} \\ K_{a,23} \\ K_{a,33} \end{bmatrix} + \epsilon,$$

$$G_i = \begin{bmatrix} g_{R_i,1} & g_{R_i,2} & g_{R_i,3} & 0 & 0 & 0 \\ 0 & 0 & 0 & g_{R_i,2} & g_{R_i,3} & 0 \\ 0 & 0 & 0 & 0 & 0 & g_{R_i,3} \end{bmatrix},$$

$$g_{R,j} = (Rg)_j, \ \Omega = \text{diag}\left(\frac{\Sigma_a}{\Delta_a[1]}, \ldots, \frac{\Sigma_a}{\Delta_a[N]} \right),$$

where $(v)_i$ represents the i-th component of the vector v.

Again, a similar formulation can be used for the magnetometer, but the gain K_m is considered full, as it is the product of a rotation with an upper triangular matrix.

D. Covariance Estimation

The maximum likelihood covariance estimator is similar that given in Eq. (6), but the reference $\hat{\mu}_s[i]$ is replaced by an estimate of $\mu_s[i]$ using parameters in Θ, and Bessel's correction is not used, since it lowers the likelihood of the data in order to provide an unbiased estimate. The new estimate is then given by:

$$\hat{\Sigma}_s = \frac{\sum_{i=1}^{N} \sum_{j=1}^{\Delta_s[i]} \tilde{s}_s[i,j] \tilde{s}_s[i,j]^T}{\sum_{i=1}^{N} \Delta_s[i]}, \tag{16}$$

$$\tilde{s}_s[i,j] = s_s[i,j] - \left(K_s R_i s_s^{\mathcal{I}} + b_s \right).$$

VII. Experiments

Two experiments are performed to validate the proposed algorithm: calibration of simulated sensors and calibration of real sensors. The simulation is used to evaluate the effects of changing the stop condition, the number of sampled intervals, and the algorithm hypothesis. The performance is analyzed in terms of errors in relation to the ground-truth values, which are available in the simulation, and the total time required for the simulation. Using a Monte Carlo approach, many combinations of sensors can be simulated to evaluate the robustness of the proposed algorithm.

The performance with test data, which is not used for training, can be evaluated to test the generalization of the algorithm [14].

In the real experiment, a magnetometer and an accelerometer had their readings measured while being held by hand inside a building, which slightly violates the hypothesis that all samples measure the same value. It will be shown that the algorithm is able to fit the sampled data correctly.

It is important to note that only the algorithm derived in this paper is used for calibration since the authors do not know of any other work, apart from their own previous contribution [10] upon which this paper is based, that describes an algorithm that can perform this kind of calibration without external references. If an algorithm such as the one proposed in [9] were used, additional information would have to be provided, making comparison of the results impossible since the calibration would be performed on different data. Thus, no comparison with any previous work is performed.

A. Simulated calibration

To evaluate the calibration algorithm, 100 Monte Carlo simulations with new sensor parameters on each run were performed, allowing evaluation of both common and possible exceptional calibration behavior. This analysis is important to show that the method described handles different sensor characteristics well. Table I presents a summary of the parameters used in the experiment; the symbols in the table are those used throughout the paper.

For each Monte Carlo run, a total of $N = 15$ sets of sensor measurements were considered, with the same number of samples $\Delta[i] \sim \mathcal{U}(\{400, 401, \ldots, 600\})$ for both sensors. The rotations for each set were randomly created using the method described in [15], and the stop condition was set to $\gamma = 10^{-4}$.

When varying the stop condition γ, the same sample sets were used for each value tested, so that the performance can be viewed as the same algorithm stopping at different times. This avoids differences caused by random fluctuations and allows better comparison. For a varying number of sets, runs with $N + 1$ sample sets merged a new set to the previous N already used, simulating collection of an increasing number of sets and also preventing random fluctuations from significantly affecting the comparison. Alternatively, the training with N sets can be viewed as training on a subset of the data set

Table I: Experimental parameters

	Description	Value
N	# of measurement sets	15
$\Delta[i]$	# of samples for each set i	$\mathcal{U}(\{400, 401, \ldots, 600\})$
R_i	Rotation for set i	See [15]
γ	Stop condition	10^{-4}
g_z	Gravity's z component	$\mathcal{U}([-1.5, -0.5])$
h_x	Magnetic field's x component	$\mathcal{U}([0.5, 1.5])$
h_z	Magnetic field's z component	$\mathcal{U}([-1.5, 1.5])$
$b_{s,i}$	Bias' component	$\mathcal{U}([-1, 1])$
K_s	Gain matrix	See Eq. (17)
Σ_s	Covariance matrix	See Eq. (18)
	# of Monte Carlo runs	100

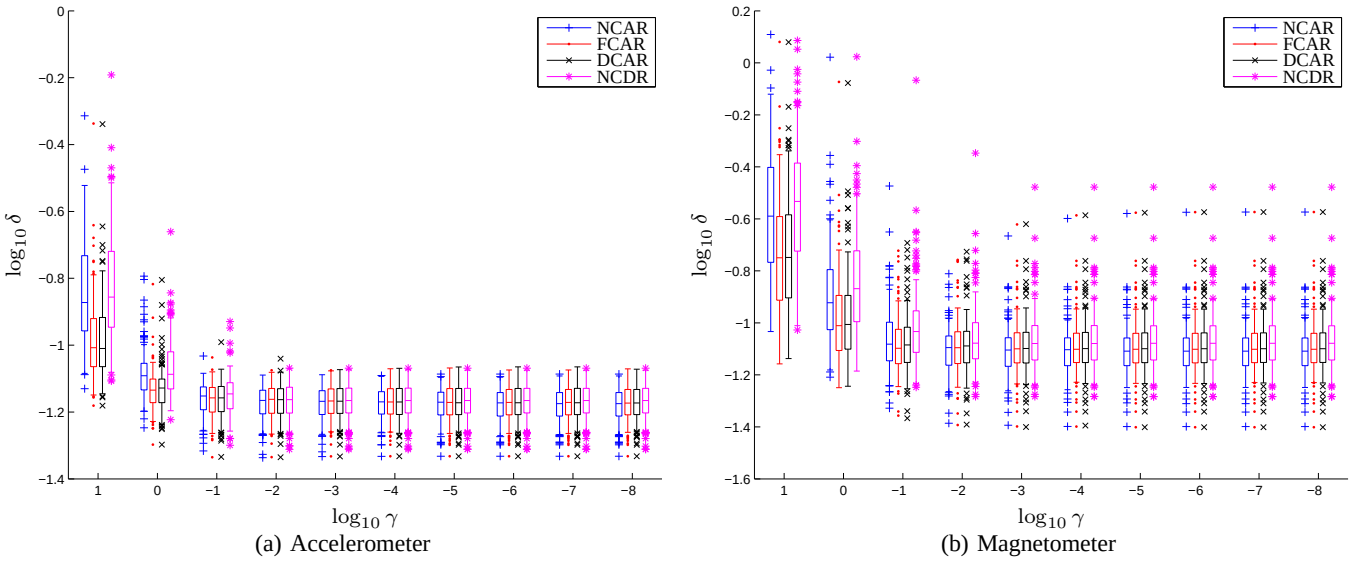

(a) Accelerometer (b) Magnetometer

Figure 1: Reconstruction error for training data for different stop conditions. Best viewed in color.

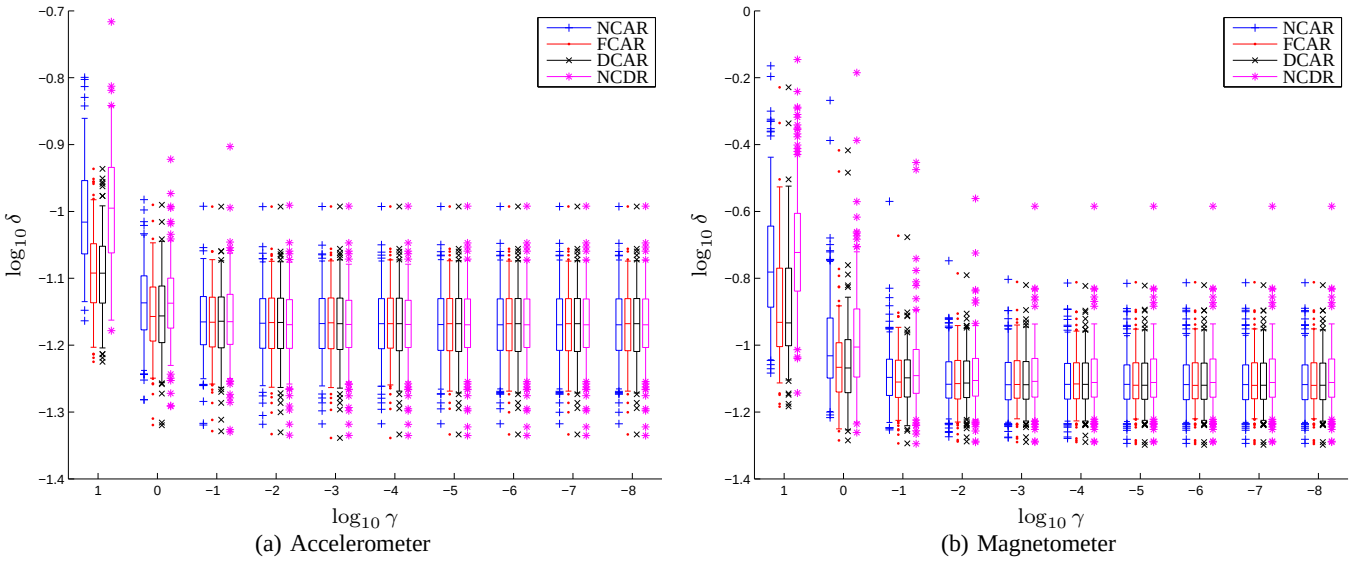

(a) Accelerometer (b) Magnetometer

Figure 2: Reconstruction error for test data for different stop conditions. Best viewed in color.

provided to the algorithm that trained with $N+1$ sets, thus maintaining consistency.

The field components g_z, h_x and h_z are sampled uniformly over $[-1.5,-0.5]$, $[0.5,1.5]$ and $[-1.5,1.5]$, respectively. Parameters for the accelerometer and magnetometer are created in a similar fashion. Each component $b_{s,i}$ of the biases has a random value sampled uniformly in $[-1,1]$. The gain matrices are given by

$$K_s = I_3 + \delta K_s,\ \ \delta K_s \sim \mathcal{U}([-0.1,0.1]), \tag{17}$$

and the covariances are given by

$$\Sigma_s = \alpha_s \begin{cases} \Sigma_{s,ij} \sim \mathcal{U}([0.5,2]), & \text{if } i=j \\ \Sigma_{s,ij} \sim \mathcal{U}([-0.2,0.2]), & \text{otherwise,} \end{cases} \tag{18}$$

where $\alpha_s = 10^{\epsilon},\ \epsilon \sim \mathcal{U}([-2,-4])$, which guarantees that the covariance matrices are positive definite, but allows high condition numbers. Instead of using the nominal value K_m, the magnetometer is rotated by R' and mirrored by M, so that the resulting gain matrix is given by $K'_m = MR'K_m$, showing that the algorithm described is able to handle any configuration between the accelerometer and magnetometer, unlike [9]. The rotation R' is chosen randomly [15], and $M = \text{diag}(\delta_1,\delta_2,\delta_3),\ \delta_i \sim \mathcal{U}(\{-1,1\})$, is a random diagonal matrix that indicates the mirroring of each component.

After calibration, the estimated value of Θ was used to rebuild estimates $\breve{\mu}_s[i]$ for $\mu_s[i]$, as in Eq. (3). To compare the quality of the reconstructions, the following measurement was used:

$$\delta_s = \sqrt{\frac{\sum_{i=1}^{N} \Delta[i]\,(\mu_s[i]-\breve{\mu}_s[i])^T\,\Sigma_s^{-1}\,(\mu_s[i]-\breve{\mu}_s[i])}{\sum_{i=1}^{N}\Delta[i]}}, \tag{19}$$

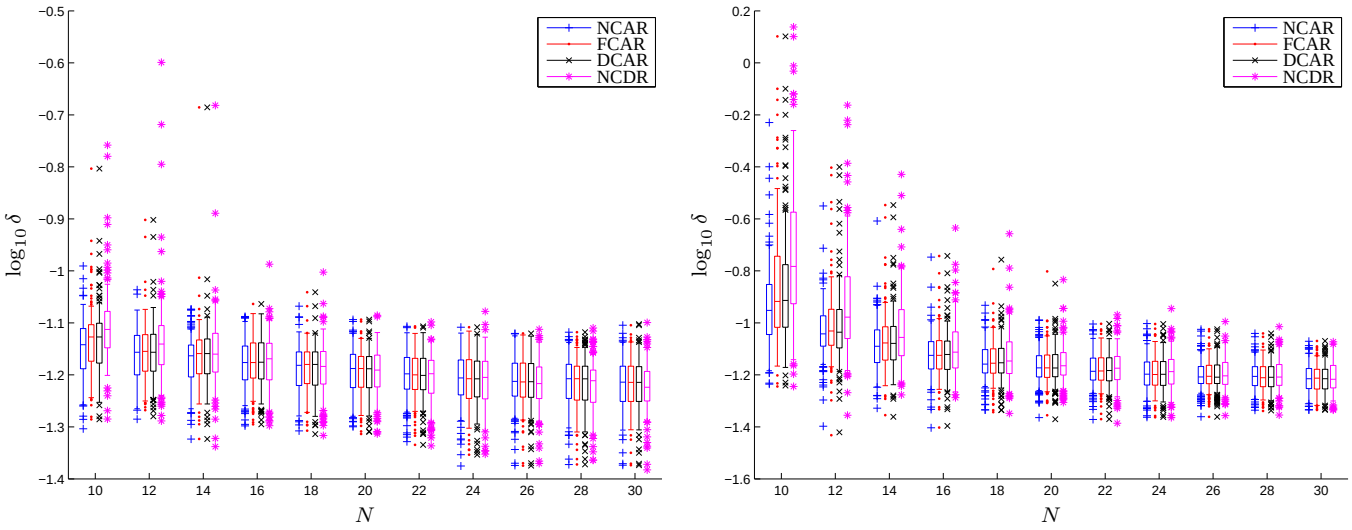

(a) Accelerometer (b) Magnetometer

Figure 3: Reconstruction error for training data for different numbers of sample sets. Best viewed in color.

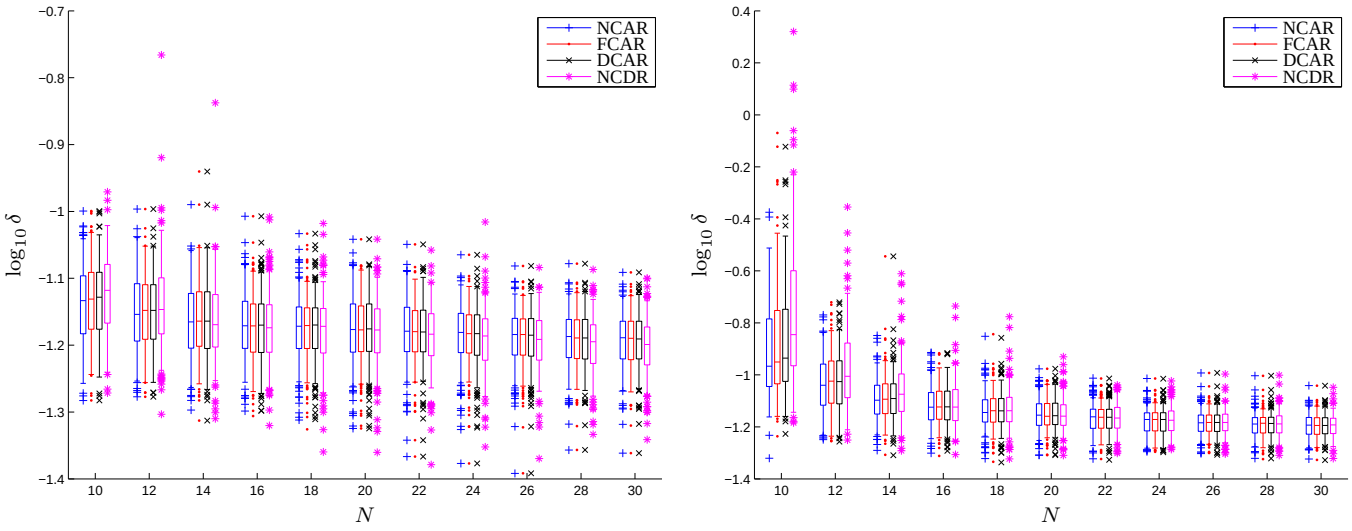

(a) Accelerometer (b) Magnetometer

Figure 4: Reconstruction error for test data for different numbers of sample sets. Best viewed in color.

where Σ_s is the real sensor covariance. This measurement is derived from the Mahalanobis distance [18] and represents the average error in standard deviations.

The rotation optimization described in Sec. VI-A can be performed in one of two ways: always using gradient descent in Eq. (15) or using the approximation in Eq. (14) until $J_k > J_{k-1}$ and then changing to gradient descent in subsequent iterations. As the approximation does not minimize the cost directly, eventually the change to gradient descent will be made. Another variation involves re-estimating the covariance using Eq. (16) or only using the initial estimate given by Eq. (6). As some approaches consider the covariance matrix to be diagonal [6], this restriction is also evaluated.

Hence there are four algorithms to compare: not refitting covariance and optimizing the rotation directly using gradient descent (NCDR); not refitting covariance and approximating the rotation (NCAR); and fitting full covariance and diagonal covariance and approximating the rotation (FCAR and DCAR, respectively). The results obtained while readjusting the covariance with direct rotation optimization are not reported as these were considerably worse.

For each sample set, another set with different samples and rotations for the same sensor parameters was created to evaluate the generalization of the fitted parameters, that is, if the parameters are able to fit the data not present during calibration well. These new samples follow the description in Sec. III and provide new values $\hat{\mu}'[i]$, which are used to compute the new rotations. The rotations were first approximated and then optimized directly to reduce the cost function, as described in Sec. VI-A.

Figure 1 shows the reconstruction error for different stop conditions. For $\gamma < 10^{-4}$, the error does not change signifi-

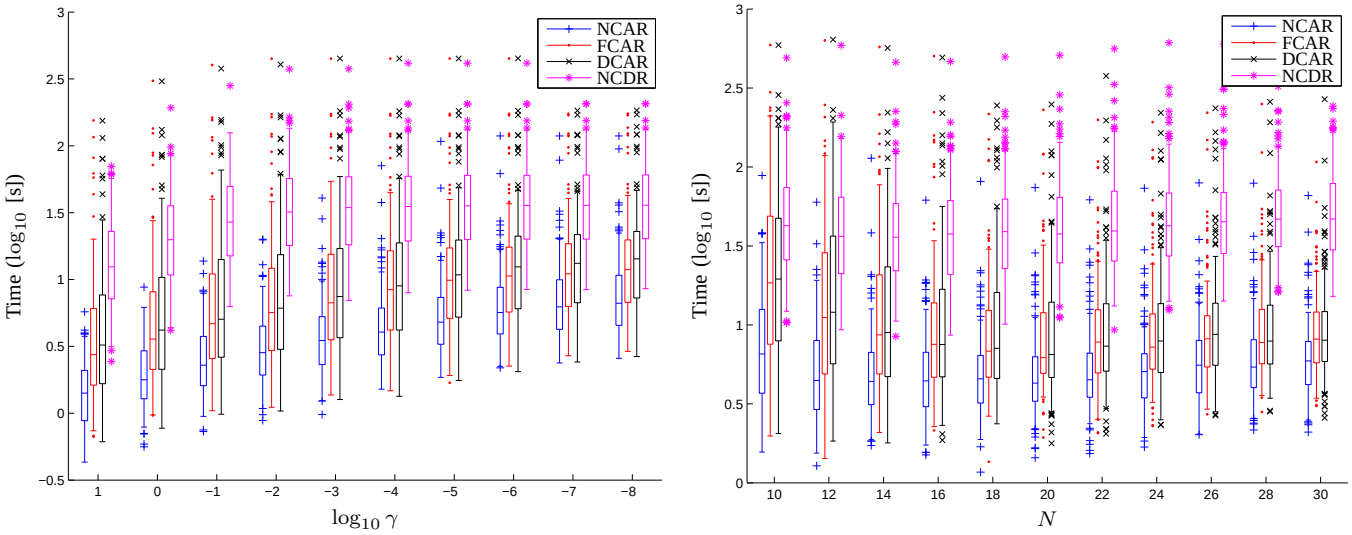

(a) Varying stop condition

(b) Varying number of samples sets

Figure 5: Calibration time. Best viewed in color.

cantly, justifying the choice of this value as a reference when evaluating different number of samples. Except for NCDR, all the algorithms have similar behavior for lower values of γ, while adjusting the covariance was beneficial for larger stop conditions. However, as larger stop conditions can only be justified if the computational cost is very high, it will be shown that this difference is not relevant. NCDR has slightly worse magnetometer calibration because of the failure of the gradient descent method to escape local minima.

It is interesting to note the apparent magnetometer overfitting for $\gamma < 10^{-2}$, as the reconstruction error for the training set increases even though the cost function decreases. Nonetheless, Fig. 2 shows that the reconstruction error for the new data not used during training generally decreases. This suggests that, although the reconstruction worsens, the parameters improve the generalization. The significantly higher error for one case with NCDR indicates that it is not as robust as the other methods using approximations. For all the algorithms, about 75% of all cases have errors of less than 0.1 standard deviations from the correct value for both sensors, showing that the proposed method is able to fit the parameters well.

Figure 3 shows that all the algorithms eventually converge to solutions that are close to each other as the number of samples increases, with errors on every run lower than 0.1 standard deviations for $N > 25$ for all methods, except for one run with NCDR. However, a diminishing return on the number of samples is observed for $N > 20$ as the error decreases very little if more data is collected. The NCAR algorithm appears to be the most robust, with a lower error spread for a smaller number of samples, so that less user data is needed to achieve acceptable performance for most situations. The NCDR algorithm usually has the largest errors, which, as in the previous case, are the result of poor local minima. Both algorithms with readjusted covariance have similar performance, indicating that the diagonal simplification is a good one. Figure 4 shows slightly better performance overall, indicating that the estimates obtained are able to generalize and fit new unseen samples well.

While the NCAR, FCAR and DCAR algorithms had similar performances for reasonable stop conditions and numbers of samples, the NCDR algorithm not only had the worst performance for all the conditions simulated but also required more computational time, as shown in Fig. 5a. It should be noted that, although DCAR has a simpler covariance to fit than FCAR, it has a higher computational cost because of implementation details. Hence, the calibration time should be about the same for these algorithms in a final implementation. As stated earlier, using higher values of γ and covariance re-estimation would be justified if the computational cost was high. However, the cost for NCAR is significantly lower, with $\gamma = 10^{-2}$ having the same cost as $\gamma = 10$ for FCAR. Therefore re-estimating the covariance using the rotation approximation does not improve performance and instead slows the processing considerably.

As shown in Figs. 1 and 2, γ values of less than 10^{-4} do not improve the performance significantly but cause the computing time for the approximation methods to increase. This occurs because rotation optimization is still taking place, which, although having little effect on the calibration, is time consuming. This also explains why the time taken by the NCDR algorithm stabilizes.

The time for NCAR is almost constant for all the numbers of samples tested, as shown in Fig. 5b. This occurs despite the fact that the cost function in Eq. (8) is a sum that increases as the number of sample sets increases and the stop condition compares two iterations of the algorithm. As more terms are used in the sum, the difference $J_{k-1} - J_k$ would also be expected to increase, so that it takes longer to satisfy the stop condition. Moreover, the optimization steps also increase in complexity, requiring more time for each step in Sec. VI. Nonetheless, the presence of more samples increases the speed at which the estimates converge, requiring less iterations to

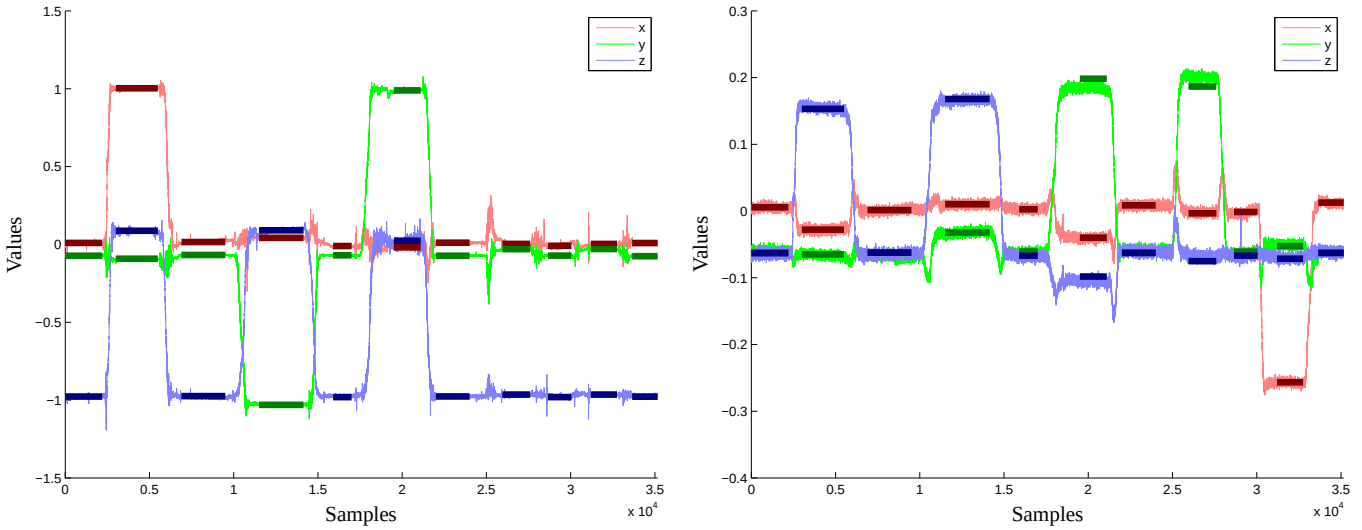

(a) Accelerometer (b) Magnetometer

Figure 6: Calibrated real sensors. The light lines show the values read and the thick, darker straight lines show the estimated reading for each set. Best viewed in color.

achieve smaller values of J.

B. Real calibration

Once it has been shown that the proposed algorithm works correctly with different simulated sensors, calibration of real sensors must be tested. The main difference between the real and simulation experiments is that, although the simulation allowed a wide variety of sensor parameters to be tested, it also simulated a perfectly still system during the capture of each dataset, which was an assumption made in Sec. III. Since the NCAR algorithm showed the best performance across all simulations, it was chosen to calibrate the real system.

In real-world applications the system may not be perfectly still, and the robustness of the algorithm in applications where this is the case must be tested. Another assumption was that the magnetic field is constant throughout the measurements, since it is assumed that it can be written as in Eq. (2). This also may not be valid in the real world because of the influence of power cables and any metal structures nearby.

To test the algorithm's robustness under these conditions, a RoboVero board was held by hand to introduce vibration and measurements were taken inside a large room so that magnetic disturbances were present but not overwhelming. Both the accelerometer and the magnetometer were configured to have the same sample rate so that both sensors took the same number of measurements for each set. The set length was determined by using time slices in which the norms of the values read were kept approximately constant. This also introduces errors as the underlying values can vary without changing the norm considerably.

The results are shown in Fig. 6, where the darker lines are the estimated mean sensor readings for each measurement set. These are computed from the estimated parameters using Eq. (3), with length equal to the number of measurements $\Delta_s[i]$ in each set. The light lines are the real sampled values.

Fig. 6a shows that the estimated values for the accelerometer were very close to the real ones despite the sensors not being held still during each set. This confirms the robustness of the algorithm when the first assumption is violated and supports its use when the sensors cannot be kept stationary.

Fig. 6b shows that the magnetometer is also calibrated correctly but that its measurements vary a lot more than the accelerometer measurements. As a result, some measurement sets have slight offsets, as is the case of the sets around 20000 and 26000 samples. These errors are associated with distortion of the magnetic field by the building's metallic structure. However, the small size of these offsets show that this does not affect significantly the calibration, confirming the algorithm's robustness in the presence of this kind of distortion.

It will be recalled that the calibration was performed using only the measured values shown in Fig. 6 without any external references or sensors. Therefore, the proposed algorithm's robustness to violation of the assumptions about the only measurements available, together with the high-quality calibration that it provides, should make it a good candidate for any calibration of these kinds of sensors.

VIII. Conclusion

This paper has described an algorithm to calibrate the parameters of an accelerometer and a magnetometer using only sensor measurements without any external information. The algorithm is designed to estimate the gain, bias, and covariance of each sensor, as well as the orientation of each measurement and the direction of the Earth's magnetic field. Hence, the calibration can be performed almost anywhere by anyone as only the sensor readings are required. This is the most generic setting for this kind of problem, since all the parameters are calibrated and any external information can be considered a constraint on the parameters and does not make the calibration harder.

A comparison was made of the base method and its variants, which use an approximation of the cost function being minimized. Simulated results show that the simplest variant is also the fastest and most robust, with the smallest worst-case errors. All the algorithms are able to achieve an error of less than 0.1 standard deviations for the data used to calibrate the devices, indicating that the algorithms were correctly trained, and also for new data not used in the calibration, indicating that the parameters obtained are suitable for use when taking sensor readings and that the learning algorithm was able to generalize. Test with real sensors and adverse conditions that violate the assumptions underlying the derivation of the algorithms showed that high-quality calibration can be achieved in non-controlled settings, confirming the robustness of the algorithm.

Future studies should investigate use of the estimated sensor rotations to compute the rotation between body and sensor frame, as described in [11], since the sensor frame used was specific to this study and may be different from the desired frame. Another possibility for further study is to integrate external references, such as vision systems [19], so that other sources of information can be used during calibration to reduce the errors even further.

ACKNOWLEDGMENT

The authors would like to thank CNPq and FAPESP for the financial support.

REFERENCES

[1] M. I. Lizarraga, G. H. Elkaim, and R. Curry, "SLUGS UAV: A Flexible and Versatile Hardware/Software Platform for Guidance Navigation and Control Research," *Airborne Experimental Test Platforms: From Theory to Flight, Washington DC*, 2013.

[2] A. M. Sabatini, "Quaternion-based extended Kalman filter for determining orientation by inertial and magnetic sensing," *Biomedical Engineering, IEEE Transactions on*, vol. 53, no. 7, pp. 1346–1356, Jul. 2006.

[3] T. Pylvanainen, "Automatic and adaptive calibration of 3D field sensors," *Applied Mathematical Modelling*, vol. 32, no. 4, pp. 575–587, Apr. 2008.

[4] J. L. Crassidis, K. Lai, and R. R. Harman, "Real-time attitude-independent three-axis magnetometer calibration," *Journal of Guidance Control and Dynamics*, vol. 28, no. 1, pp. 115–120, Jan. 2005.

[5] R. Alonso and M. D. Shuster, "Complete linear attitude-independent magnetometer calibration," *Journal of the Astronautical Sciences*, vol. 50, no. 4, pp. 477–490, 2002.

[6] J. F. Vasconcelos, G. Elkaim, C. Silvestre, P. Oliveira, and B. Cardeira, "Geometric approach to strapdown magnetometer calibration in sensor frame," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 47, no. 2, pp. 1293–1306, 2011.

[7] D. Gebre-Egziabher, G. H. Elkaim, J. D. Powell, and B. W. Parkinson, "A gyro-free quaternion-based attitude determination system suitable for implementation using low cost sensors," in *Position Location and Navigation Symposium, IEEE 2000*. Ieee, 2000, pp. 185–192.

[8] S. O. H. Madgwick, A. J. L. Harrison, and R. Vaidyanathan, "Estimation of IMU and MARG orientation using a gradient descent algorithm," in *Rehabilitation Robotics (ICORR), 2011 IEEE International Conference on*, vol. 2011, Jun. 2011, pp. 1–7.

[9] M. Kok, J. Hol, T. B. Schön, F. Gustafsson, and H. Luinge, "Calibration of a magnetometer in combination with inertial sensors," in *Proceedings of the IEEE International Conference on Information Fusion*. IEEE, 2012, pp. 787–793.

[10] C. S. Miranda and J. V. Ferreira, "Greedy statistically correct simultaneous calibration of magnetometers and accelerometers," in *Proceedings of the ABCM International Congress of Mechanical Engineering*. ABCM, 2013.

[11] ——, "Sensor and Body Frames Rotation Calibration Through Attitude Restriction," in *World Congress of the International Federation of Automatic Control*, vol. 19, no. 1, 2014, pp. 7579–7584.

[12] K. Krishnamoorthy, *Handbook of Statistical Distributions with Applications*. CRC Press, 2006.

[13] J. Kenney and E. Keeping, *Mathematics of statistics*, ser. Mathematics of Statistics. Van Nostrand, 1957, no. v. 2.

[14] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2006.

[15] J. J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *Proceedings of the IEEE International Conference on Robotics and Automation*, no. April. IEEE, 2004, pp. 3993–3998, vol. 4.

[16] C. F. F. Karney, "Quaternions in molecular modeling," *Journal of Molecular Graphics and Modelling*, vol. 25, no. 5, pp. 595–604, Jun. 2005.

[17] T. Kariya and H. Kurata, *Generalized least squares*. Wiley, 2004.

[18] P. C. Mahalanobis, "On the generalized distance in statistics," in *Proceedings of the National Institute of Sciences of India*, vol. 2, no. 1. New Delhi, 1936, pp. 49–55.

[19] Q. V. Le and A. Y. Ng, "Joint calibration of multiple sensors," in *International Conference on Intelligent Robots and Systems*. IEEE, 2009, pp. 3651–3658.

Conrado S. Miranda received his M.S. degree on Mechanical Engineering and his B.S. in Control and Automation Engineering from the University of Campinas (Unicamp), Brazil, in 2014 and 2011, respectively. He is currently a Ph.D. student at the School of Electrical and Computer Engineering, Unicamp. His main research interests are machine learning, multi-objective optimization, neural networks, and statistical models.

Janito V. Ferreira received his Ph.D. in Mechanical Engineering from the Imperial College, London, in 1999, and his M.S. and B.S. in Mechanical Engineering from the University of Campinas (Unicamp), Campinas, SP, Brazil, in 1989 and 1983, respectively. Since 1984 he is with the Department of Computational Mechanics at Unicamp, where he is presently the head of the Autonomous Mobility Laboratory. His main interests are in the areas of autonomous mobility, analysis of dynamic systems, applied and experimental mechanics, project of integrated system, and signal and image processing.