

Maximum likelihood calibration of a magnetometer using inertial sensors

Manon Kok* Thomas B. Schön**

* *Division of Automatic Control, Linköping University, SE-581 83
Linköping, Sweden (email: manko@isy.liu.se)*

** *Department of Information Technology, Uppsala University,
SE-751 05 Uppsala, Sweden (e-mail: thomas.schon@it.uu.se)*

Abstract: Magnetometers and inertial sensors (accelerometers and gyroscopes) are widely used to estimate 3D orientation. For the orientation estimates to be accurate, the sensor axes need to be aligned and the magnetometer needs to be calibrated for sensor errors and for the presence of magnetic disturbances. In this work we use a grey-box system identification approach to compute maximum likelihood estimates of the calibration parameters. An experiment where the magnetometer data is highly disturbed shows that the algorithm works well on real data, providing good calibration results and improved heading estimates. We also provide an identifiability analysis to understand how much rotation is needed to be able to solve the calibration problem.

Keywords: Magnetometers, calibration, inertial sensors, maximum likelihood, grey-box system identification, sensor fusion.

1. INTRODUCTION

Inertial sensors (3D accelerometers and 3D gyroscopes) in combination with 3D magnetometers are in many applications used to obtain orientation estimates using an extended Kalman filter (EKF). When the sensor is subject to low accelerations, the accelerometer measurements are dominated by the gravity component from which it is possible to deduce information about the inclination. The magnetometer measures the local magnetic field vector and its horizontal component can be used to obtain heading information. The orientation estimates are only accurate when the sensor axes of the inertial sensors and the magnetometers are aligned and when the magnetometer is properly calibrated. This calibration consists of two parts. First, the magnetometer needs to be calibrated for errors inherent in the sensor, for instance non-orthogonality of the magnetometer sensor axes. Second, it needs to be calibrated for the presence of magnetic disturbances. These magnetic disturbances are frequently present due to the mounting of a magnetometer and cause a constant disturbance that needs to be calibrated for. When using a magnetometer it is therefore always advisable to calibrate it before use.

Our main contribution is a new practical algorithm for calibration of a magnetometer when we also have access to measurements from inertial sensors rigidly attached to the magnetometer. The algorithm implements a maximum likelihood (ML) estimator to find the calibration parameters. These parameters account for magnetometer sensor errors, the presence of constant magnetic disturbances caused by mounting the magnetometer close to magnetic

* This work is supported by MC Impulse, a European Commission, FP7 research project and CADICS, a Linnaeus Center funded by the Swedish Research Council (VR).




Fig. 1. Calibration results for experimental data where the original magnetometer data is plotted in red and the calibrated magnetometer data is plotted in blue.

objects and misalignments between the magnetometer and the inertial sensors. An illustration of the results obtained from calibrating a magnetometer using the proposed algorithm is available in Fig. 1.

In many practical applications, mounting the magnetometer onto for instance a car or a boat severely limits the rotational freedom of the sensor. A secondary contribution of this work is a quantification of how much rotation is needed to be able to solve the calibration problem. This is done via an identifiability analysis, deriving how much rotation is needed in the case of perfect measurements.

Many recent magnetometer calibration approaches are based on ellipse fitting. These calibration algorithms use the fact that when rotating a magnetometer, its measurements should lie on a sphere if the magnetometer measures a constant local magnetic field vector, i.e. the norm of the magnetic field should be constant. The presence of magnetic disturbances or magnetometer sensor errors leads to an ellipsoid of data instead. Many calibration algorithms are used to obtain improved orientation estimates. Hence, the radius of the sphere, i.e. the actual magnitude of the local magnetic field vector, is irrelevant and we can without loss of generality assume that its norm is equal to 1. The calibration has therefore been successful when the calibrated measurements can be seen to lie on a (unit) sphere (recall Fig. 1). The problem of fitting an ellipsoid of data to a sphere was considered for example by Gander et al. [1994]. As shown by Markovsky et al. [2004] the ordinary least squares estimate is inconsistent when using measurements corrupted by noise. Different calibration approaches have been developed to overcome this problem, see e.g. Gebre-Egziabher et al. [2006], Renaudin et al. [2010].

An ellipse fitting method can only solve the calibration up to an unknown rotation. Combining the magnetometer with inertial sensors requires the sensor axes to be aligned, i.e. the sphere needs to be oriented such that the magnetometer sensor axes are aligned with the inertial sensor axes. A few recent approaches add a second step to the calibration procedure to estimate the misalignment between the magnetometer and the inertial sensor axes [Vasconcelos et al., 2011, Li and Li, 2012, Salehi et al., 2012, Bonnet et al., 2009]. They first use an ellipse fitting approach and determine the misalignment between the magnetometer and the inertial sensor axes in a second step. These approaches generally use the accelerometer measurements for estimating the misalignment, but discard the gyroscope measurements. Troni and Whitcomb [2013] focus on using the gyroscope measurements to obtain an estimate of the misalignment. Our work follows a similar approach, but makes use of both the gyroscope and the accelerometer measurements and aims at obtaining a maximum likelihood estimate by combining both steps into a nonlinear optimization problem.

2. PROBLEM FORMULATION

Our solution to the magnetometer calibration problem makes use of a nonlinear state space model describing the sensor's orientation q_t and its measurements,

$$q_{t+1} = f_t(q_t, u_t, \theta) + B(q_t)v_t(\theta), \quad (1a)$$

$$y_t = h_t(q_t, \theta) + e_t(\theta). \quad (1b)$$

Here, q_t denotes the state variable representing the sensor orientation encoded using a unit quaternion. Furthermore, u_t and y_t denote observed input and output variables, respectively. The dynamics is denoted by $f_t(\cdot)$ and the measurement model is denoted by $h_t(\cdot)$. Finally, v_t and e_t represent mutually independent i.i.d. process and measurement noise, respectively and $B(q_t)$ is a matrix describing how the noise v_t affects the state q_t . The model is introduced in detail in Section 3.

The calibration problem is formulated as a maximum likelihood grey-box system identification problem. Hence, based on N measurements of the observed input $u_{1:N}$ =

$\{u_1, \dots, u_N\}$ and output $y_{1:N} = \{y_1, \dots, y_N\}$ variables, find the parameters θ that maximize the likelihood function,

$$\hat{\theta}^{\text{ML}} = \arg \max_{\theta \in \Theta} p_\theta(y_{1:N}), \quad (2)$$

where $\Theta \subseteq \mathbb{R}^{n_\theta}$. Using conditional probabilities and the fact that the logarithm is a monotonic function we have the following equivalent formulation of (2),

$$\hat{\theta}^{\text{ML}} = \arg \min_{\theta \in \Theta} - \sum_{t=1}^N \log p_\theta(y_t | y_{1:t-1}), \quad (3)$$

where we have used the convention that $y_{1:0} \triangleq \emptyset$. The ML estimator (3) enjoys well-understood theoretical properties including strong consistency, asymptotic normality, and asymptotic efficiency [Ljung, 1999]. The state space model (1) is nonlinear, implying that there is no closed form solution available for the one step ahead predictor $p_\theta(y_t | y_{1:t-1})$. This can be handled using sequential Monte Carlo methods (e.g. particle filters and particle smoothers), see e.g. Schön et al. [2011], Lindsten and Schön [2013]. However, for the magnetometer calibration problem under study it is sufficient to make use of a more pragmatic approach; we simply approximate the one step ahead predictor using an extended Kalman filter (EKF). The result is

$$p_\theta(y_t | y_{1:t-1}) \approx \mathcal{N}(y_t | \hat{y}_{t|t-1}(\theta), S_t(\theta)), \quad (4)$$

where $\mathcal{N}(y_t | \hat{y}_{t|t-1}(\theta), S_t(\theta))$ denotes the probability density function for the Gaussian random variable y_t with mean value $\hat{y}_{t|t-1}(\theta)$ and covariance $S_t(\theta)$. Here, $S_t(\theta)$ is the residual covariance from the EKF [Gustafsson, 2012]. Inserting (4) into (3) results in the following optimization problem,

$$\min_{\theta \in \Theta} \frac{1}{2} \sum_{t=1}^N \|y_t - \hat{y}_{t|t-1}(\theta)\|_{S_t^{-1}(\theta)}^2 + \log \det S_t(\theta), \quad (5)$$

which we can solve for the unknown parameters θ . The problem (5) is non-convex, implying that a good initial value for θ is required.

3. MODELS

3.1 Dynamic model

We model the orientation q_t as the sensor's orientation from the body frame b to the navigation frame n at time t , expressed as q_t^{nb} . The body frame b is the coordinate frame of the inertial sensor with its origin in the center of the accelerometer triad and its axes aligned with the inertial sensor axes. The navigation frame n is aligned with the earth gravity and the local magnetic field vector. The dynamic model of the orientation (1a) takes the gyroscope measurements as an input, which are modeled as

$$y_{\omega,t} = \omega_t + \delta_\omega + v_{\omega,t}, \quad (6)$$

where ω_t denotes the angular velocity, δ_ω denotes the gyroscope bias and $v_{\omega,t} \sim \mathcal{N}(0, \Sigma_\omega)$. The dynamic equation for the orientation states is then given by [Gustafsson, 2012]

$$q_{t+1}^{\text{nb}} = (\mathcal{I}_4 + \frac{T}{2}L(y_{\omega,t} - \delta_\omega)) q_t^{\text{nb}} + \frac{T}{2}\bar{L}(q_t^{\text{nb}})v_{\omega,t}. \quad (7)$$

Here, \mathcal{I}_4 denotes a 4×4 identity matrix, T denotes the sampling time, the matrices $\bar{L}(q_t^{\text{nb}})$ and $L(y_{\omega,t} - \delta_\omega)$ are given by

$$\bar{L}(q) = \begin{pmatrix} -q_1 & -q_2 & -q_3 \\ q_0 & -q_3 & q_2 \\ q_3 & q_0 & -q_1 \\ -q_2 & q_1 & q_0 \end{pmatrix}, L(\omega) = \begin{pmatrix} 0 & -\omega_1 & -\omega_2 & -\omega_3 \\ \omega_1 & 0 & \omega_3 & -\omega_2 \\ \omega_2 & -\omega_3 & 0 & \omega_1 \\ \omega_3 & \omega_2 & -\omega_1 & 0 \end{pmatrix}, \quad (8)$$

where the subscripts denote the different components of the respective vectors.

3.2 Measurement model

The measurements in (1b) consist of the accelerometer measurements $y_{a,t}$ and the magnetometer measurements $y_{m,t}$. The former are modeled as

$$y_{a,t} = R_t^{\text{bn}}(a_t^n - g^n) + e_{a,t} \approx -R_t^{\text{bn}}g^n + e_{a,t}, \quad (9)$$

where a_t^n denotes the sensor's acceleration which is assumed to be approximately zero, g^n denotes the gravity vector in navigation frame and $e_{a,t} \sim \mathcal{N}(0, \Sigma_a)$. The matrix R_t^{bn} is the rotation matrix representation of the quaternion $q_t^{\text{bn}} = (q_i^{\text{nb}})^c$, where c denotes the quaternion conjugate [Hol, 2011].

The magnetometer measurements are modeled as

$$y_{m,t} = DR_t^{\text{bn}}m^n + o + e_{m,t}, \quad (10)$$

where the local magnetic field is denoted by m^n and $e_{m,t} \sim \mathcal{N}(0, \Sigma_m)$. Note that m^n is assumed to be constant. The calibration matrix D and offset vector o contain corrections for both the magnetometer sensor errors, the presence of magnetic disturbances and the misalignment between the magnetometer and inertial sensor axes. The magnetometer and inertial sensor axes are assumed not to be mirrored with respect to each other. Kok et al. [2012] provide a more detailed derivation and discussion of this model.

3.3 Parameter vector

The parameter vector θ to be estimated consists first of all of the calibration matrix D , the offset vector o and the local earth magnetic field m^n as discussed in Section 3.2. Furthermore, the gyroscope bias δ_ω introduced in Section 3.1 and the noise covariances Σ_ω , Σ_a and Σ_m of the three sensors are treated as unknown parameters. The different components of the unknown parameters θ therefore consist of

$$D \in \mathbb{R}^{3 \times 3}, \quad (11a)$$

$$o \in \mathbb{R}^3, \quad (11b)$$

$$m^n \in \{\mathbb{R}^3 : \|m^n\|_2^2 = 1, m_x^n > 0, m_y^n = 0\}, \quad (11c)$$

$$\delta_\omega \in \mathbb{R}^3, \quad (11d)$$

$$\Sigma_\omega \in \{\mathbb{R}^{3 \times 3} : \Sigma_\omega \succeq 0, \Sigma_\omega = \Sigma_\omega^T\}, \quad (11e)$$

$$\Sigma_a \in \{\mathbb{R}^{3 \times 3} : \Sigma_a \succeq 0, \Sigma_a = \Sigma_a^T\}, \quad (11f)$$

$$\Sigma_m \in \{\mathbb{R}^{3 \times 3} : \Sigma_m \succeq 0, \Sigma_m = \Sigma_m^T\}, \quad (11g)$$

where m_x^n and m_y^n denote the x - and y - component of m^n , respectively.

Although (11c) and (11e) - (11g) suggest that constrained optimization is needed, it is possible to circumvent this via a suitable reparametrization. The covariance matrices can be parametrized in terms of their Cholesky factorization, leading to only 6 parameters for each 3×3 covariance matrix. The local magnetic field can be parametrized using only one parameter denoted by d [Kok et al., 2012].

4. NEW CALIBRATION ALGORITHM

The resulting calibration algorithm is given in Algorithm 1. The optimization problem is solved using a Broyden-Fletcher-Goldfarb-Shanno (BFGS) method with damped BFGS updating [Nocedal and Wright, 2006].

Algorithm 1 Magnetometer and inertial calibration

- (1) Set $i = 0$ and initialize $\hat{D}_0, \hat{o}_0, \hat{m}_0^n, \hat{\delta}_{\omega,0}, \hat{\Sigma}_{\omega,0}, \hat{\Sigma}_{a,0}, \hat{\Sigma}_{m,0}$.
 - (2) Repeat,
 - (a) Run an EKF using the current estimates $\hat{D}_i, \hat{o}_i, \hat{m}_i^n, \hat{\delta}_{\omega,i}, \hat{Q}_i = \hat{\Sigma}_{\omega,i}, \hat{R}_i = \begin{pmatrix} \hat{\Sigma}_{a,i} & 0_{3 \times 3} \\ 0_{3 \times 3} & \hat{\Sigma}_{m,i} \end{pmatrix}$ and obtain $\{\hat{y}_{t|t-1}\}_{t=1}^N, S_{1:N}$.
 - (b) Determine $\hat{\theta}_{i+1}$ using the numerical gradient of the objective function in (5), its approximate Hessian and a line search algorithm.
 - (c) Obtain $\hat{D}_{i+1}, \hat{o}_{i+1}, \hat{m}_{i+1}^n, \hat{\delta}_{\omega,i+1}, \hat{\Sigma}_{\omega,i+1}, \hat{\Sigma}_{a,i+1}, \hat{\Sigma}_{m,i+1}$ from $\hat{\theta}_{i+1}$.
 - (d) Set $i := i + 1$ and repeat from 2(a) until convergence.
-

Since the optimization problem (5) is a non-convex problem, it is important to start Algorithm 1 with a good initialization of the parameters. We will do this using a two-step approach.

In a first step the ellipsoid of magnetometer data is mapped onto a unit sphere, using an ellipse fitting approach [Gander et al., 1994]. The second step is based on the invariance of the inner product of two vectors under rotation and is similar to approaches in Vasconcelos et al. [2011], Li and Li [2012], Salehi et al. [2012], Bonnet et al. [2009]. This requires estimates of the inclination which can be obtained by running the EKF using only the accelerometer and gyroscope data. These two steps result in initial estimates \hat{D}_0, \hat{o}_0 and \hat{m}_0^n and are described in more detail in Kok et al. [2012]. The gyroscope bias and the noise covariance matrices can be initialized based on a small period of stationary data.

5. MINIMUM ROTATION NEEDED

For magnetometer calibration, the sensor needs to be rotated in all possible orientations so the magnetometer measurements describe (a part of) an ellipsoid. Work on magnetometer calibration generally assumes that the sensor can be "sufficiently" rotated for the calibration parameters to be identified. Often, however, magnetometers are mounted in such a way that their movement is more or less constrained to the 2D plane [Wu et al., 2013a,b]. Although some approaches do consider the case of limited rotation, to the best of the authors' knowledge there is no study on how much rotation is sufficient. We will now answer this question via an identifiability analysis. Note that our analysis assumes noise-free measurements and it therefore gives a lower bound on the number of measurements needed to identify the calibration parameters. In many applications, rotations around the z -axis are possible, while rotations around the other axes are constrained. We therefore ask the question of how much rotation around the x - and y -axes is needed in addition to rotation around the z -axis.

Identifiability analysis of the calibration parameters $\theta_c = (\text{vec } D \ o^\top \ d)^\top$, where d is used to parametrize m^n as described in Section 3.3, can be performed by studying local observability of an extended system [Walter, 1982]. This extended system consists for our problem of the state q_t (see (1a)) and the parameters θ . Assuming noise-free and bias-free measurements, the orientation at different time steps is known from the inertial measurements, save for the initial heading. The analysis can therefore be limited to the study of θ_c , where it needs to be kept in mind that the initial heading, denoted by α , is unknown.

Discrete-time observability analysis considers the system of equations

$$y_{m,1} = DR_1^{\text{bn}}(\alpha)m^n + o = h_m(R_1^{\text{bn}}, \alpha, \theta_c), \quad (12a)$$

$$y_{m,2} = DR_2^{\text{bn}}(\alpha)m^n + o = h_m(R_2^{\text{bn}}, \alpha, \theta_c), \quad (12b)$$

⋮

$$y_{m,K} = DR_K^{\text{bn}}(\alpha)m^n + o = h_m(R_K^{\text{bn}}, \alpha, \theta_c). \quad (12c)$$

We will analyze the identifiability of the parameters θ_c assuming that the initial heading is known. The identifiability of the initial heading will be discussed separately afterwards. The parameters θ_c are said to be *locally identifiable* if there exists an $R_{1:K}^{\text{bn}}$ such that we can solve for θ_c in (12).

Due to the nonlinear nature of (12) it is typically hard to analyze it directly. Instead, observability is often studied by considering a linearized version of (12), making use of the Jacobian

$$\mathcal{J} = \begin{pmatrix} \nabla h_m(R_1^{\text{bn}}, \alpha, \theta_c) \\ \nabla h_m(R_2^{\text{bn}}, \alpha, \theta_c) \\ \vdots \\ \nabla h_m(R_K^{\text{bn}}, \alpha, \theta_c) \end{pmatrix}, \quad (13)$$

where ∇ denotes the derivative with respect to θ_c . If (13) is full rank for all θ_c , the parameters are said to be *locally weakly identifiable* [Nijmeijer and van der Schaft, 1990].

The Jacobian \mathcal{J} can be built by stacking different measurements on top of each other. For notational simplicity, we consider 90° -rotations around the axes. Note that the analysis is equally valid for any other amount of rotation. For rotations of 0° , 90° and -90° around the z -axis, 90° and -90° around the x -axis and 90° and -90° around the y -axis, the Jacobian becomes

$$\mathcal{J}(\theta_c) = \begin{pmatrix} m_x \mathcal{I}_3 & 0_{3 \times 3} & m_z \mathcal{I}_3 & \mathcal{I}_3 & D_3 \frac{\partial m_z}{\partial d} + D_1 \frac{\partial m_x}{\partial d} \\ 0_{3 \times 3} & m_x \mathcal{I}_3 & m_z \mathcal{I}_3 & \mathcal{I}_3 & D_3 \frac{\partial m_z}{\partial d} + D_2 \frac{\partial m_x}{\partial d} \\ 0_{3 \times 3} & -m_x \mathcal{I}_3 & m_z \mathcal{I}_3 & \mathcal{I}_3 & D_3 \frac{\partial m_z}{\partial d} - D_2 \frac{\partial m_x}{\partial d} \\ \hline m_x \mathcal{I}_3 & -m_z \mathcal{I}_3 & 0_{3 \times 3} & \mathcal{I}_3 & D_1 \frac{\partial m_x}{\partial d} - D_2 \frac{\partial m_z}{\partial d} \\ m_x \mathcal{I}_3 & m_z \mathcal{I}_3 & 0_{3 \times 3} & \mathcal{I}_3 & D_2 \frac{\partial m_x}{\partial d} + D_1 \frac{\partial m_z}{\partial d} \\ \hline m_z \mathcal{I}_3 & 0_{3 \times 3} & -m_x \mathcal{I}_3 & \mathcal{I}_3 & D_1 \frac{\partial m_z}{\partial d} - D_3 \frac{\partial m_x}{\partial d} \\ -m_z \mathcal{I}_3 & 0_{3 \times 3} & m_x \mathcal{I}_3 & \mathcal{I}_3 & D_3 \frac{\partial m_z}{\partial d} - D_1 \frac{\partial m_x}{\partial d} \end{pmatrix}, \quad (14)$$

$\underbrace{\hspace{10em}}_{\text{vec } D} \quad \underbrace{\hspace{2em}}_o \quad \underbrace{\hspace{2em}}_d$

where D_i denotes the i 'th column of the matrix D . For clarity, the contributions of the different rotation axes are separated by dashed lines. For each column it is explicitly indicated which part of the derivative with respect to θ_c it represents.

The rank of (14) can be studied for subsets of its rows. Considering only rotations around the z -axis, i.e. the first

$\mathcal{J}(\theta_c)$ full rank	Pole		Equator		Elsewhere		
#meas z -axis	3	3	3	3	3	3	3
#meas x -axis	2	1	0	1	1	0	2
#meas y -axis	1	2	2	1	1	2	0

Table 1. Summary of how many measurements (i.e. lower bound) around which axes are needed for identifiability of the calibration parameters θ_c .

three rows in (14), rank $\mathcal{J}(\theta_c) = 9$ as long as $m_x \neq 0$ (the calibration is not performed on the magnetic north or south pole). The matrix will not gain any rank from adding more measurements around the z -axis. Adding two measurements around the x - and/or y -axis can be shown to lead to rank $\mathcal{J}(\theta_c) = 13$ under very mild conditions on the calibration matrix D in case the calibration is not performed on the equator or on a magnetic pole. A sufficient condition for this is that $\det D \neq 0$. On the equator and the magnetic poles it is also possible to obtain identifiability of all calibration parameters. In these cases, however, stricter requirements on the measurements are needed. The results on how many measurements around which axes are needed are summarized in Table 1. It is straightforward to show that also α is identifiable based on these measurements, provided that the calibration is not performed on one of the magnetic poles. In this case, the initial heading is not identifiable for any number of measurements since the horizontal component of the magnetic field is zero.

The identifiability analysis above was done assuming 90° rotations between each measurement. The same result holds, however, for any other difference in rotations. Although the amount of rotation does not influence the identifiability of the calibration parameters, it will indeed influence the quality of the estimates. This can be understood in terms of the condition number [Kailath et al., 2000], i.e. the ratio between the maximum and minimum singular value of the Jacobian (14). Any singular values being zero implies that θ_c is not identifiable, but a smaller condition number also implies that certain parameters are more difficult to estimate. Fig. 2 shows the singular values of the Jacobian in (14) for different amounts of rotation. Five measurements are considered, corresponding to three measurements around the z -axis, one around the x -axis and one around the y -axis for $D = \mathcal{I}_3$, $o = (0 \ 0 \ 0)^\top$ and m^n equal to that in Linköping, Sweden. The amount of rotation between the measurements is assumed to be equal for all four differences and is varied between 0° (red) and 90° (blue). As can be seen, the larger the amount of rotation, the smaller the condition number of (14).

It can therefore be concluded that rotation around the z -axis only is never enough to identify the calibration parameters θ_c . The parameters will, however, become identifiable already with very little excitation in the other directions as summarized in Table 1. More excitation around the different axes will always lead to better estimates as can be concluded from Fig. 2.

6. EXPERIMENTS

6.1 Experimental setup

An experiment has been conducted with an Xsens MTi-100 inertial measurement unit (IMU), collecting inertial and




Fig. 2. Singular values of the Jacobian in (14) for small rotation ranges (red) to large rotation ranges (blue).



Fig. 3. Experimental setup where a calibration experiment is performed outdoors. An Xsens MTi-100 IMU (orange box) together with a magnetic disturbance is placed in a aluminum block.

magnetometer data using a sampling frequency of 100 Hz. The experimental setup is shown in Fig. 3. The data has been collected outdoors to ensure that the local magnetic field is as homogeneous as possible. A magnetic disturbance has been rigidly attached to the sensor, causing severe magnetic disturbances. Both the block in which the IMU is placed and the table shown in Fig. 3 are made of non-magnetic material.

The sensor is slowly rotated around all three axes for roughly 2 minutes. The collected magnetometer data is shown in red in Fig. 1. Note that for plotting purposes this data has been downsampled to 1 Hz.

To aid the convergence of the optimization problem (5), a few iterations of step 2 in Algorithm 1 were run for parts of the parameter vector, first estimating only D , o , m^n and δ_ω and afterwards only estimating Σ_ω , Σ_a and Σ_m . The results from these steps were used as an initial estimate to obtain the final maximum likelihood estimate of the parameters θ .

6.2 Experimental results

The calibration resulted in the following calibration matrix \hat{D} and offset vector \hat{o}

$$\hat{D} = \begin{pmatrix} 0.74 & -0.14 & 0.02 \\ -0.12 & 0.68 & 0.01 \\ -0.04 & 0.43 & 1.00 \end{pmatrix}, \quad \hat{o} = \begin{pmatrix} 1.37 \\ 1.22 \\ -0.94 \end{pmatrix}. \quad (15)$$

Correcting the magnetometer measurements with the obtained calibration matrix and offset vector leads to the corrected measurements shown in blue in Fig. 1, which

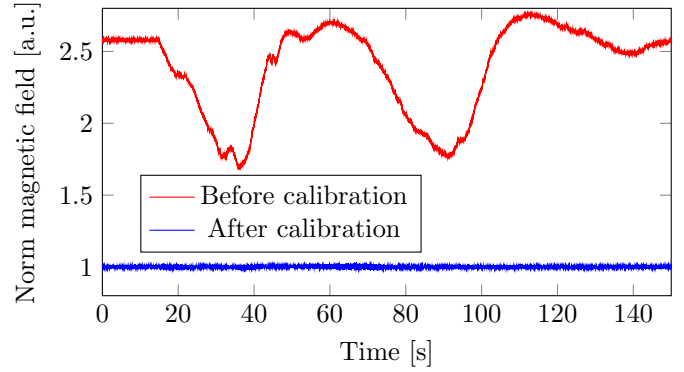


Fig. 4. Norm of the magnetic field measurements before (red) and after (blue) calibration.




Fig. 5. Normalized residuals $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ from the EKF after calibration for the original data set (left) and the validation data set (right). A Gaussian distribution (red) is fitted to the data.

can be seen to lie on a unit sphere. To better visualize the deviation from the unit norm, the norm of the magnetic field before and after calibration is depicted in Fig. 4. The norm after calibration can be seen to lie around 1.

A histogram of the normalized residuals after calibration from the EKF $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ is shown in Fig. 5. A Gaussian distribution has been fitted to the normalized residuals showing that the residuals are indeed more or less Gaussian distributed.

6.3 Validation

The calibration results have been applied to a second data set recorded with the same experimental setup as in Fig. 3. The results are very similar to the results displayed in Fig. 1 and Fig. 4, implying that the calibration parameters can well be used to calibrate the magnetometer measurements of the other data set. As can be seen from Fig. 5, the resulting residuals still look fairly Gaussian, save for the presence of a small number of outliers. Closer inspection of the residuals and the data shows that these outliers are the result of high accelerometer measurements due to a small impact during the collection of the validation data set.

6.4 Heading estimation

An important goal of magnetometer calibration is to facilitate good heading estimates. To check the quality of




Fig. 6. Calibrated magnetometer data of an experiment rotating the sensor into 24 different sensor orientations where the blue, green and red lines represent the data from the x -, y - and z -axis of the magnetometer, respectively.

z-axis		x-axis		y-axis
z up	z down	x up	x down	y down
-0.23	-0.84	0.08	0.98	-0.31
0.21	-2.70	-0.02	1.66	0.35
-0.44	1.81	-0.82	-0.71	-0.07
0.42	2.00	0.36	-1.89	0.45

Table 2. Difference in estimated orientation between two subsequent rotations around the sensor axes using calibrated magnetometer data. The values represent the deviation in degrees from 90° .

the heading estimates after calibration, the block in which the sensor is placed (shown in Fig. 3) is rotated around all axes. This block has right angles and it can therefore be placed in 24 orientations that differ from each other by 90° . The calibrated magnetometer data of this experiment is shown in Fig. 6. Orientation estimates are determined by taking the mean value of 500 magnetometer and accelerometer samples in each orientation and using the accelerometer to estimate the sensor's inclination and the magnetometer data to estimate its heading. After calibration, we expect the difference of the estimated heading between each subsequent rotation to be 90° . Note that when rotating around an axis, the sensor is always rotated back to its initial position, enabling the computation of 4 orientation differences per rotation axis. Table 2 enlists the deviation from 90° between two subsequent rotations. As can be seen, the deviation is very small, indicating that good heading estimates are obtained after calibration. Note that the metal object causing the magnetic disturbance as shown in Fig. 3 physically prevents the setup to properly be placed in all orientations around the y -axis. Rotation around the y -axis with the y -axis pointing upwards has therefore not been included in Table 2.

7. CONCLUSIONS

A grey-box system identification approach has been used to compute maximum likelihood estimates of magnetometer calibration parameters for a magnetometer in combination with inertial sensors. The results show that the algorithm works well on real data and leads to improved heading estimates.

ACKNOWLEDGEMENTS

The authors would like to thank Laurens Slot, Dr. Henk Luinge and Dr. Jeroen Hol from Xsens Technologies for their support in collecting the data sets and for interesting discussions.

REFERENCES

- S. Bonnet, C. Bassompierre, C. Godin, S. Lesecq, and A. Barraud. Calibration methods for inertial and magnetic sensors. *Sensors and Actuators A: Physical*, 156(2):302–311, 2009.
- W. Gander, G.H. Golub, and R. Strebler. Least-squares fitting of circles and ellipses. *BIT Numerical Mathematics*, 34(4):558–578, 1994.
- D. Gebre-Egziabher, G.H. Elkaim, J.D. Powell, and B.W. Parkinson. Calibration of strapdown magnetometers in magnetic field domain. *Journal of Aerospace Engineering*, 19(2):87–102, April 2006.
- F. Gustafsson. *Statistical Sensor Fusion*. Studentlitteratur, 2012.
- J.D. Hol. *Sensor Fusion and Calibration of Inertial Sensors, Vision, Ultra-Wideband and GPS*. PhD thesis, Linköping University, 2011.
- T. Kailath, A.H. Sayed, and B. Hassibi. *Linear estimation*. Prentice Hall, 2000.
- M. Kok, J.D. Hol, T.B. Schön, F. Gustafsson, and H. Luinge. Calibration of a magnetometer in combination with inertial sensors. In *Proceedings of the 15th International Conference on Information Fusion*, pages 787 – 793, Singapore, July 2012.
- X. Li and Z. Li. A new calibration method for tri-axial field sensors in strap-down navigation systems. *Measurement Science and Technology*, 23(10), October 2012.
- F. Lindsten and T.B. Schön. Backward simulation methods for Monte Carlo statistical inference. *Foundations and Trends in Machine Learning*, 6(1):1–143, 2013.
- L. Ljung. *System Identification, Theory for the User*. Prentice Hall PTR, 2nd edition, 1999.
- I. Markovsky, A. Kukush, and S. Van Huffel. Consistent least squares fitting of ellipsoids. *Numerische Mathematik*, 98(1):177–194, 2004.
- H. Nijmeijer and A.J. van der Schaft. *Nonlinear dynamical control systems*. Springer-Verlag, 1990.
- J. Nocedal and S.J. Wright. *Numerical Optimization*. Springer Series in Operations Research, 2nd edition, 2006.
- V. Renaudin, M.H. Afzal, and G. Lachapelle. Complete triaxis magnetometer calibration in the magnetic domain. *Journal of Sensors*, 2010.
- S. Salehi, N. Mostofi, and G. Bleser. A practical in-field magnetometer calibration method for IMUs. In *Proceedings of the IROS Workshop on Cognitive Assistive Systems: Closing the Action-Perception Loop*, pages 39–44, Vila Moura, Portugal, October 2012.
- T.B. Schön, A. Wills, and B. Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.
- G. Troni and L.L. Whitcomb. Adaptive estimation of measurement bias in three-dimensional field sensors with angular-rate sensors: Theory and comparative experimental evaluation. In *Proceedings of Robotics: Science and Systems (RSS)*, Berlin, Germany, June 2013.
- J.F. Vasconcelos, G. Elkaim, C. Silvestre, P. Oliveira, and B. Cardeira. Geometric approach to strapdown magnetometer calibration in sensor frame. *IEEE Transactions on Aerospace and Electronic Systems*, 47(2):1293–1306, April 2011.
- E. Walter. *Identifiability of state space models*. Springer-Verlag, 1982.
- Z. Wu, X. Hu, M. Wu, and J. Cao. Attitude-independent magnetometer calibration for marine magnetic surveys: regularization issue. *Journal of Geophysics and Engineering*, 10(4), June 2013a.
- Z. Wu, X. Hu, M. Wu, and J. Cao. Constrained total least-squares calibration of three-axis magnetometer for vehicular applications. *Measurement Science and Technology*, 24(9), July 2013b.