Magnetometer Calibration Using Inertial Sensors

Manon Kok and Thomas B. Schon

Linköping University Post Print

N.B.: When citing this work, cite the original article.

©2016 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Manon Kok and Thomas B. Schon, Magnetometer Calibration Using Inertial Sensors, 2016, IEEE Sensors Journal, (16), 14, 5679-5689.

<http://dx.doi.org/10.1109/JSEN.2016.2569160>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-130378>

Magnetometer calibration using inertial sensors

Manon Kok, *Student Member, IEEE*, and Thomas B. Schön, *Senior Member, IEEE*

Abstract—In this work we present a practical algorithm for calibrating a magnetometer for the presence of magnetic disturbances and for magnetometer sensor errors. To allow for combining the magnetometer measurements with inertial measurements for orientation estimation, the algorithm also corrects for misalignment between the magnetometer and the inertial sensor axes. The calibration algorithm is formulated as the solution to a maximum likelihood problem and the computations are performed offline. The algorithm is shown to give good results using data from two different commercially available sensor units. Using the calibrated magnetometer measurements in combination with the inertial sensors to determine the sensor's orientation is shown to lead to significantly improved heading estimates.

Index Terms—Magnetometers, calibration, inertial sensors, maximum likelihood, grey-box system identification, sensor fusion.

I. INTRODUCTION

NOWADAYS, magnetometers and inertial sensors (gyroscopes and accelerometers) are widely available, for instance in dedicated sensor units and in smartphones. Magnetometers measure the local magnetic field. When no magnetic disturbances are present, the magnetometer measures a constant local magnetic field vector. This vector points to the local magnetic north and can hence be used for heading estimation. Gyroscopes measure the angular velocity of the sensor. Integration of the gyroscope measurements gives information about the change in orientation. However, it does not provide absolute orientation estimates. Furthermore, the orientation estimates suffer from integration drift. Accelerometers measure the sensor's acceleration in combination with the earth's gravity. In the case of small or zero acceleration, the measurements are dominated by the gravity component. Hence, they can be used to estimate the inclination of the sensor.

Inertial sensors and magnetometers have successfully been used to obtain accurate 3D orientation estimates for a wide range of applications. For this, however, it is imperative that the sensors are properly calibrated and that the sensor axes are aligned. Calibration is specifically of concern for the magnetometer, which needs recalibration whenever it is placed in a (magnetically) different environment. When the magnetic disturbance is a result of the mounting of the magnetometer onto a magnetic object, the magnetometer can be calibrated to compensate for the presence of this disturbance. This is the focus of this work.

Manon Kok is with the Department of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden e-mail: (see <http://users.isy.liu.se/en/rt/manko/>).

Thomas B. Schön is with the Department of Information Technology, Uppsala University, SE-751 05 Uppsala, Sweden e-mail: (see <http://user.it.uu.se/~thosc112/>).

Fig. 1. Example calibration results with an ellipsoid of magnetometer data before calibration (red) and a unit sphere of data after calibration (blue).

Our main contribution is a practical magnetometer calibration algorithm that is designed to improve orientation estimates when combining calibrated magnetometer data with inertial data. The word *practical* refers to the fact that the calibration does not require specialized additional equipment and can therefore be performed by any user. More specifically, this means that the orientation of the sensor is not assumed to be known. Instead, the calibration problem is formulated as an orientation estimation problem in the presence of unknown parameters and is posed as a maximum likelihood (ML) problem. The algorithm calibrates the magnetometer for the presence of magnetic disturbances, for magnetometer sensor errors and for misalignment between the magnetometer and the inertial sensor axes. Using the calibrated magnetometer measurements to estimate the sensor's orientation is experimentally shown to lead to significantly improved heading estimates. We aggregate and extend the work from [1] and [2] with improvements on the implementation of the algorithm. Furthermore, we include a more complete description and analysis, more experimental results and a simulation study illustrating the heading accuracy that can be obtained with a properly calibrated sensor.

To perform the calibration, the sensor needs to be rotated in all possible orientations. A perfectly calibrated magnetometer would in that case measure rotated versions of the local magnetic field vector. Hence, the magnetometer data would lie on a sphere. In practice, however, the magnetometer will often measure an ellipsoid of data instead. The calibration maps the ellipsoid of data to a sphere as illustrated in Fig. 1. The alignment of the inertial and magnetometer sensor axes determines the orientation of the sphere. Since we are interested in improving the heading estimates, the actual magnitude of the local magnetic field is of no concern. Hence, we assume without loss of generality that the norm is equal to 1, i.e. the sphere in Fig. 1 is a unit sphere.

II. RELATED WORK

Traditional magnetometer calibration approaches assume that a reference sensor is available which is able to provide accurate heading information. A well-known example of this is compass swinging [3]. To allow for any user to perform the calibration, however, a large number of approaches have been developed that remove the need for a source of orientation information. One class of these magnetometer calibration algorithms focuses on minimizing the difference between the magnitude of the measured magnetic field and that of the local magnetic field, see e.g. [4]. This approach is also referred to as scalar checking [5]. Another class formulates the calibration problem as an ellipsoid fitting problem, i.e. as the problem of mapping an ellipsoid of data to a sphere, see e.g. [6]–[8]. The benefit of using this formulation, is that there is a vast literature on solving ellipsoid fitting problems, see e.g. [9], [10]. Outside of these two classes, a large number of other calibration approaches is also available, for instance [11], where different formulations of the calibration problem in terms of an ML problem are considered.

The benefit of the approaches discussed above is that they can be used with data from a magnetometer only. Our interest, however, lies in calibrating a magnetometer for improved heading estimation in combination with inertial sensors. Alignment of the sensor axes of the inertial sensors and the magnetometer is in this case crucial. This alignment can be seen as determining the orientation of the blue sphere of calibrated magnetometer data in Fig. 1. Algorithms that only use magnetometer data can map the red ellipsoid of data to a sphere, but without additional information, the rotation of this sphere remains unknown.

A number of recent approaches include a second step in the calibration algorithm to determine the misalignment [6], [12]–[14] between different sensor axes. A common choice to align the magnetometer and inertial sensor axes, is to use accelerometer measurements from periods of fairly small accelerations [12], [13]. The downside of this approach is that a threshold for using accelerometer measurements needs to be determined. Furthermore, data from the gyroscope is hereby omitted. In [15] on the other hand, the problem is reformulated in terms of the change in orientation, allowing for direct use of the gyroscope data.

In our algorithm we instead formulate the magnetometer calibration problem as a problem of estimating the sensor's orientation in the presence of unknown (calibration) parameters. This formulation naturally follows from the fact that the problem of orientation estimation and that of magnetometer calibration are inherently connected: If the magnetometer is properly calibrated, good orientation estimates can be obtained. Reversely, if the orientation of the sensor is known accurately, the rotation of the sphere in Fig. 1 can accurately be determined, resulting in a good magnetometer calibration. In this formulation, data from the accelerometer *and* the gyroscope is used to aid the magnetometer calibration.

Our formulation of the calibration problem requires solving a non-convex optimization problem to obtain ML estimates of the calibration parameters. To obtain good initial values of

the parameters, an ellipsoid fitting problem and a misalignment estimation problem are solved. Solving the calibration problem as a two-step procedure is similar to the approaches in [12], [13]. We analyze the quality of the initial estimates and of the ML estimates in terms of their heading accuracy, both for experimental and simulated data. Based on this analysis, we show that significant heading accuracy improvements can be obtained by using the ML estimates of the parameters.

III. PROBLEM FORMULATION

Our magnetometer calibration algorithm is formulated as a problem of determining the sensor's orientation in the presence of unknown model parameters θ . It can hence be considered to be a grey-box system identification problem. A nonlinear state space model on the following form is used

$$x_{t+1} = f_t(x_t, \omega_t, e_{\omega,t}, \theta), \quad (1a)$$

$$y_t = \begin{pmatrix} y_{a,t} \\ y_{m,t} \end{pmatrix} = \begin{pmatrix} h_{a,t}(x_t) \\ h_{m,t}(x_t, \theta) \end{pmatrix} + e_t(\theta), \quad (1b)$$

where the state x_t represents the sensor's orientation at time t . We use the change in orientation, i.e. the angular velocity ω_t , as an input to the dynamic model $f_t(\cdot)$. The angular velocity is measured by the gyroscope. However, the measurements $y_{\omega,t}$ are corrupted by a constant bias δ_ω and Gaussian i.i.d. measurement noise with zero mean and covariance Σ_ω , i.e. $e_{\omega,t} \sim \mathcal{N}(0_{3 \times 1}, \Sigma_\omega)$.

The measurement models $h_{a,t}(\cdot)$ and $h_{m,t}(\cdot)$ in (1b) describe the accelerometer measurements $y_{a,t}$ and the magnetometer measurements $y_{m,t}$, respectively. The accelerometer measurement model assumes that the acceleration of the sensor is small compared to the earth gravity. Since the magnetometer is not assumed to be properly calibrated, the magnetometer measurement model $h_{m,t}(\cdot)$ depends on the parameter vector θ . The exact details of the magnetometer measurement model will be introduced in Section IV. The accelerometer and magnetometer measurements are corrupted by Gaussian i.i.d. measurement noise

$$e_t = \begin{pmatrix} e_{a,t} \\ e_{m,t} \end{pmatrix} \sim \mathcal{N}\left(0_{6 \times 1}, \begin{pmatrix} \Sigma_a & 0_{3 \times 3} \\ 0_{3 \times 3} & \Sigma_m \end{pmatrix}\right). \quad (2)$$

The calibration problem is formulated as an ML problem. Hence, the parameters θ in (1) are found by maximizing the likelihood function $p_\theta(y_{1:N})$,

$$\hat{\theta}^{\text{ML}} = \arg \max_{\theta \in \Theta} p_\theta(y_{1:N}), \quad (3)$$

where $y_{1:N} = \{y_1, \dots, y_N\}$ and $\Theta \subseteq \mathbb{R}^{n_\theta}$. Using conditional probabilities and the fact that the logarithm is a monotonic function we have the following equivalent formulation of (3),

$$\hat{\theta}^{\text{ML}} = \arg \min_{\theta \in \Theta} - \sum_{t=1}^N \log p_\theta(y_t | y_{1:t-1}), \quad (4)$$

where we use the convention that $y_{1:0} \triangleq \emptyset$. The ML estimator (4) enjoys well-understood theoretical properties including strong consistency, asymptotic normality, and asymptotic efficiency [16].

The state space model (1) is nonlinear, implying that there is no closed form solution available for the one step ahead

predictor $p_\theta(y_t \mid y_{1:t-1})$ in (4). This can systematically be handled using sequential Monte Carlo methods (e.g. particle filters and particle smoothers), see e.g. [17], [18]. However, for the magnetometer calibration problem 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

$$p_\theta(y_t \mid y_{1:t-1}) \approx \mathcal{N}(y_t; \hat{y}_{t|t-1}(\theta), S_t(\theta)), \quad (5)$$

where the mean value $\hat{y}_{t|t-1}(\theta)$ and the covariance $S_t(\theta)$ are obtained from the EKF [19]. Inserting (5) into (4) and neglecting all constants not depending on θ results in the following optimization problem,

$$\min_{\theta \in \Theta} \frac{1}{2} \sum_{t=1}^N \|y_t - \hat{y}_{t|t-1}(\theta)\|_{S_t^{-1}(\theta)}^2 + \log \det S_t(\theta), \quad (6)$$

which we can solve for the unknown parameters θ . The problem (6) is non-convex, implying that a good initial value for θ is required.

IV. MAGNETOMETER MEASUREMENT MODEL

In the case of perfect calibration, a magnetometer measures the local magnetic field and its measurements will therefore lie on a sphere with a radius equal to the local magnetic field. Since we are interested in using the magnetometer measurements to improve the orientation estimates from the state space model (1), the actual magnitude of the local magnetic field is of no concern. Hence, we assume without loss of generality that its norm is equal to one. We denote the normalized local magnetic field by m^n . Ideally, the magnetometer measurements then lie on a sphere with radius equal to one as

$$h_{m,t} = m_t^b = R_t^{bn} m^n, \quad (7)$$

where $h_{m,t}$ is defined in (1b). The explicit dependence on x_t and θ has been omitted for notational simplicity. The matrix R_t^{bn} is the rotation matrix representation of the orientation at time t . The superscript bn denotes that the rotation is from the *navigation frame* n to the *body frame* b . The body frame b is aligned with the sensor axes. The navigation frame n is aligned with the earth's gravity and the local magnetic field. In case the coordinate frame in which a vector is defined can be ambiguous, we explicitly indicate in which coordinate frame the vector is expressed by adding a superscript b or n . Hence, m^n denotes the normalized local magnetic field in the navigation frame n while m_t^b denotes the normalized local magnetic field in the body frame b . The latter is time-dependent and therefore also has a subscript t . Note that the rotation from navigation frame to body frame is denoted R_t^{nb} and $R_t^{bn} = (R_t^{nb})^\top$.

In outdoor environments, the local magnetic field is equal to the local *earth* magnetic field. Its horizontal component points towards the earth's magnetic north pole. The ratio between the horizontal and vertical component depends on the location on the earth and can be expressed in terms of the *dip angle* δ . In indoor environments, the magnetic field can locally be assumed to be constant and points towards a *local*

magnetic north. This is not necessarily the earth's magnetic north pole. Choosing the navigation frame n such that the x -axis is pointing towards the local magnetic north, m^n can be parametrized in terms of its vertical component m_z^n

$$m^n = \begin{pmatrix} \sqrt{1 - (m_z^n)^2} & 0 & m_z^n \end{pmatrix}^\top, \quad (8a)$$

or in terms of the dip angle δ

$$m^n = \begin{pmatrix} \cos \delta & 0 & -\sin \delta \end{pmatrix}^\top. \quad (8b)$$

Note that the two parametrizations do not encode exactly the same knowledge about the magnetic field; the first component of m^n in (8a) is positive by construction while this is not true for (8b). However, both parametrizations will be used in the remainder. It will be argued that no information is lost by using (8b) if the parameter estimates are properly initialized.

The main need for magnetometer calibration arises from the fact that a magnetometer needs recalibration each time it is placed in a magnetically different environment. Specifically, a magnetometer measures a superposition of the local magnetic field and of the magnetic field due to the presence of magnetic material in the vicinity of the sensor. In case this magnetic material is rigidly attached to the magnetometer, it is possible to calibrate the magnetometer measurements for this. The magnetic material can give rise to both hard and soft iron contributions to the magnetic field. Hard iron effects are due to permanent magnetization of the magnetic material and lead to a constant 3×1 offset vector o_{hi} . Soft iron effects are due to magnetization of the material as a result of an external magnetic field and therefore depend on the orientation of the material with respect to the local magnetic field. We model this in terms of a 3×3 matrix C_{si} . Hence, the magnetometer measurements do not lie on a sphere as in (7), but instead, they lie on a translated ellipsoid as

$$h_{m,t} = C_{si} R_t^{bn} m^n + o_{hi}. \quad (9)$$

As discussed in Section II, when calibrating the magnetometer to obtain better orientation estimates, it is important that the magnetometer and the inertial sensor axes are aligned. Let us now be more specific about the definition of the body frame b and define it to be located in the center of the accelerometer triad and aligned with the accelerometer sensor axes. Furthermore, let us assume that the accelerometer and gyroscope axes are aligned. Defining the rotation between the body frame b and the magnetometer sensor frame b_m as R^{bm} , the model (9) can be extended to

$$h_{m,t} = C_{si} R^{bm} R_t^{bn} m^n + o_{hi}. \quad (10)$$

Finally, the magnetometer calibration can also correct for the presence of sensor errors in the magnetometer. These errors are sensor-specific and can differ for each individual magnetometer. They can be subdivided into three components, see e.g. [6]–[8]:

- 1) Non-orthogonality of the magnetometer axes, represented by a matrix C_{no} .
- 2) Presence of a zero bias or null shift, implying that the magnetometer will measure a non-zero magnetic field even if the magnetic field is zero, defined by o_{zb} .

Algorithm 1 Magnetometer and inertial calibration

- 1) Determine an initial parameter estimate $\hat{D}_0, \hat{o}_0, \hat{m}_0^n, \hat{\delta}_{\omega,0}, \hat{\Sigma}_{\omega,0}, \hat{\Sigma}_{a,0}, \hat{\Sigma}_{m,0}$ using three steps
 - a) Initialize $\hat{\delta}_{\omega,0}, \hat{\Sigma}_{\omega,0}, \hat{\Sigma}_{a,0}, \hat{\Sigma}_{m,0}$.
 - b) Obtain an initial $\hat{D}_0$ and $\hat{o}_0$ based on ellipsoid fitting (see Section VI-A).
 - c) Obtain initial $\hat{D}_0, \hat{o}_0$ and $\hat{m}_0^n$ by initial determination of the sensor axis misalignment (see Section VI-B).
- 2) Set $i = 0$ and repeat,
 - a) Run the EKF using the current estimates $\hat{D}_i, \hat{o}_i, \hat{m}_i^n, \hat{\delta}_{\omega,i}, \hat{\Sigma}_{\omega,i}, \hat{\Sigma}_{a,i}, \hat{\Sigma}_{m,i}$ to obtain $\{\hat{y}_{t|t-1}(\theta_i), S_t(\theta_i)\}_{t=1}^N$ and evaluate the cost function in (6).
 - b) Determine $\hat{\theta}_{i+1}$ using the numerical gradient of the cost function in (6), its approximate Hessian and a backtracking 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 Step 2a until convergence.
- 3) Difference in sensitivity of the three magnetometer axes, represented by a diagonal matrix C_{sc} .

We can therefore extend the model (10) to also include the magnetometer sensor errors as

$$h_{m,t} = C_{sc}C_{no} (C_{si}R_t^{bmb}R_t^{bn}m^n + o_{hi}) + o_{zb}. \quad (11)$$

To obtain a correct calibration, it is fortunately not necessary to identify all individual contributions of the different components in (11). Instead, they can be combined into a 3×3 distortion matrix D and a 3×1 offset vector o where

$$D = C_{sc}C_{no}C_{si}R_t^{bmb}, \quad o = C_{sc}C_{no}o_{hi} + o_{zb}. \quad (12)$$

The resulting magnetometer measurement model in (1b) can be written as

$$y_{m,t} = DR_t^{bn}m^n + o + e_{m,t}. \quad (13)$$

In deriving the model we have made two important assumptions:

Assumption 1. *The calibration matrix D and offset vector o in (12) are assumed to be time-independent. This implies that we assume that the magnetic distortions are constant and rigidly attached to the sensor. Also, the inertial and the magnetometer sensor axes are assumed to be rigidly attached to each other, i.e. their misalignment is represented by a constant rotation matrix. Additionally, in our algorithm we will assume that their misalignment can be described by a rotation matrix, i.e. that their axes are not mirrored with respect to each other.*

Assumption 2. *The local magnetic field m^n is assumed to be constant. In outdoor environments, this is typically a physically reasonable assumption. In indoor environments, however, the local magnetic field can differ in different locations in the building and care should be taken to fulfill the assumption.*

V. CALIBRATION ALGORITHM

In our magnetometer calibration algorithm we solve the optimization problem (6) to estimate the parameter vector θ . In this section we introduce the resulting calibration algorithm which is summarized in Algorithm 1. In Section V-A, we first discuss our optimization strategy. A crucial part of this optimization strategy is the evaluation of the cost function. Some details related to this are discussed in Section V-B. Finally, in Section V-C we introduce the parameter vector θ in more detail.

A. Optimization algorithm

The optimization problem (6) is solved in Step 2 of Algorithm 1. Standard unconstrained minimization techniques are used, which iteratively update the parameter estimates as

$$\theta_{i+1} = \theta_i - \alpha_i [\mathcal{H}(\theta_i)]^{-1} \mathcal{G}(\theta_i), \quad (14)$$

where the *direction* of the parameter update at iteration i is determined by $[\mathcal{H}(\theta_i)]^{-1} \mathcal{G}(\theta_i)$. The step *length* of the update at iteration i is denoted by α_i .

Typical choices for the search direction include choosing $\mathcal{G}(\theta_i)$ to be the gradient of the cost function in (6) and $\mathcal{H}(\theta_i)$ to be its Hessian. This leads to a Newton optimization algorithm. However, computing the gradient and Hessian of (6) is not straightforward. Possible approaches are discussed in [20], [21] for the case of linear models. In the case of nonlinear models, however, they only lead to approximate gradients, see e.g. [22], [23]. For this reason we make use of a numerical approximation of $\mathcal{G}(\theta_i)$ instead and use a Broyden-Fletcher-Goldfarb-Shanno (BFGS) method with damped updating [24] to approximate the Hessian. Hence, the minimization is performed using a quasi-Newton optimization algorithm. A backtracking line search is used to find a good step length α_i .

Proper initialization of the parameters is crucial since the optimization problem (6) is non-convex. Step 1 summarizes the three-step process used to obtain good initial estimates of all parameters.

B. Evaluation of the cost function

An important part of the optimization procedure is the evaluation of the cost function in (6). This requires running an EKF using the state space model (1) to estimate the orientation of the sensor. This EKF uses the angular velocity ω_t as an input to the dynamic model (1a). An estimate of the angular velocity is obtained from the gyroscope measurements $y_{\omega,t}$ which are modeled as

$$y_{\omega,t} = \omega_t + \delta_{\omega} + e_{\omega,t}. \quad (15)$$

The measurement model (1b) entails the accelerometer measurements and the magnetometer measurements. The magnetometer measurement model can be found in (13). The accelerometer measurements $y_{a,t}$ are modeled as

$$y_{a,t} = R_t^{bn}(a_t^n - g^n) + e_{a,t} \approx -R_t^{bn}g^n + e_{a,t}, \quad (16)$$

where a_t^n denotes the sensor's acceleration in the navigation frame and g^n denotes the earth's gravity. The rotation matrix R_t^{bn} has previously been introduced in Section IV.

The state in the EKF, which represents the sensor orientation, can be parametrized in different ways. In previous work we have used a quaternion representation as a 4-dimensional state vector [1]. In this work we instead use an implementation of the EKF, which is sometimes called a multiplicative EKF [25]–[27]. Here, a 3-dimensional state vector represents the *orientation deviation from a linearization point*. More details on this implementation can be found in [28].

The EKF returns the one step ahead predicted measurements $\{\hat{y}_{t|t-1}(\theta)\}_{t=1}^N$ and their covariance $\{S_t(\theta)\}_{t=1}^N$ which can be used to evaluate (6). The cost function needs to be evaluated for the current parameter estimates in Step 2a but also needs to be evaluated once for each component of the parameter vector θ to compute the numerical gradient. Hence, each iteration i requires running the EKF at least $n_\theta + 1$ times. Note that the actual number of evaluations can be higher since the backtracking line search algorithm used to determine α_i can require a varying number of additional evaluations. Since $n_\theta = 34$, computing the numerical gradient is computationally rather expensive. However, it is possible to parallelize the computations.

C. The parameter vector θ

As apparent from Section IV, our main interest lies in determining the calibration matrix D and the offset vector o , which can be used to correct the magnetometer measurements to obtain more accurate orientation estimates. To solve the calibration problem, however, we also estimate a number of other parameters.

First, the local magnetic field m^n introduced in Section IV is in general scenarios unknown and needs to be estimated. In outdoor environments, m^n is equal to the local earth magnetic field and is accurately known from geophysical studies, see e.g. [29]. In indoor environments, however, the local magnetic field can differ quite significantly from the local earth magnetic field. Because of that, we treat m^n as an unknown constant. Second, the gyroscope measurements that are used to describe the change in orientation of the sensor in (1a) are corrupted by a bias δ_ω . This bias is slowly time varying but for our relatively short experiments it can be assumed to be constant. Hence, it is treated as part of the parameter vector θ . Finally, we treat the noise covariance matrices Σ_ω , Σ_a and Σ_m as unknown. In summary, the parameter vector θ consists of

$$D \in \mathbb{R}^{3 \times 3}, \quad (17a)$$

$$o \in \mathbb{R}^3, \quad (17b)$$

$$m^n \in \{\mathbb{R}^3 : \|m^n\|_2^2 = 1, m_x^n > 0, m_y^n = 0\}, \quad (17c)$$

$$\delta_\omega \in \mathbb{R}^3, \quad (17d)$$

$$\Sigma_\omega \in \{\mathbb{R}^{3 \times 3} : \Sigma_\omega \succeq 0, \Sigma_\omega = \Sigma_\omega^T\}, \quad (17e)$$

$$\Sigma_a \in \{\mathbb{R}^{3 \times 3} : \Sigma_a \succeq 0, \Sigma_a = \Sigma_a^T\}, \quad (17f)$$

$$\Sigma_m \in \{\mathbb{R}^{3 \times 3} : \Sigma_m \succeq 0, \Sigma_m = \Sigma_m^T\}, \quad (17g)$$

where m_x^n and m_y^n denote the x - and y - component of m^n , respectively. The notation $\Sigma \succeq 0$ denotes the assumption that the matrix Σ is positive semi-definite.

Although (17c) and (17e) – (17g) suggest that constrained optimization is needed, it is possible to circumvent this via

suitable reparametrizations. 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 as in (8). Note that in our implementation we prefer to use the representation (8b) for the ML problem (6). Although this latter parametrization does not account for the constraint $m_x^n > 0$, this is of no concern due to proper initialization. The procedure to obtain good initial estimates of all parameters is the topic of the next section.

VI. FINDING GOOD INITIAL ESTIMATES

Since the optimization problem is non-convex, the parameter vector θ introduced in Section V needs proper initialization. An initial estimate $\hat{\theta}_0$ is obtained using a three-step method. As a first step, the gyroscope bias δ_ω and the noise covariances of the inertial sensors, Σ_ω , Σ_a , and of the magnetometer, Σ_m , are initialized. This is done using a short batch of stationary data. Alternatively, they can be initialized based on prior sensor knowledge. As a second step, described in Section VI-A, an ellipsoid fitting problem is solved using the magnetometer data. This maps the ellipsoid of data to a sphere but can not determine the rotation of the sphere. The rotation of the sphere is determined in a third step of the initialization procedure. This step also determines an initial estimate of the normalized local magnetic field m^n .

A. Ellipsoid fitting

Using the definition of the normalized local magnetic field m^n , we would expect all calibrated magnetometer measurements to lie on the unit sphere,

$$\begin{aligned} \|m^n\|_2^2 - 1 &= \|R_t^{bn} m^n\|_2^2 - 1 \\ &= \|D^{-1}(y_{m,t} - o - e_{m,t})\|_2^2 - 1 = 0. \end{aligned} \quad (18)$$

In practice, the measurements are corrupted by noise and the equality (18) does not hold exactly. The ellipsoid fitting problem can therefore be written as

$$y_{m,t}^T A y_{m,t} + b^T y_{m,t} + c \approx 0, \quad (19)$$

with

$$A \triangleq D^{-T} D^{-1}, \quad b \triangleq -2o^T A, \quad c \triangleq o^T A o. \quad (20)$$

Assuming that the matrix A is positive definite, this can be recognized as the definition of an ellipsoid with parameters A , b and c (see e.g. [9]). We can rewrite (19) as a linear relation of the parameters as

$$M\xi \approx 0, \quad (21)$$

with

$$M = \begin{pmatrix} y_{m,1} \otimes y_{m,1} & y_{m,1} & 1 \\ y_{m,2} \otimes y_{m,2} & y_{m,2} & 1 \\ \vdots & \vdots & \vdots \\ y_{m,N} \otimes y_{m,N} & y_{m,N} & 1 \end{pmatrix}, \quad \xi = \begin{pmatrix} \text{vec } A \\ b \\ c \end{pmatrix}, \quad (22)$$

where $\otimes$ denotes the Kronecker product and vec denotes the vectorization operator. This problem has infinitely many solutions and without constraining the length of the vector ξ , the

trivial solution $\xi = 0$ would be obtained. A possible approach to solve the ellipsoid fitting problem is to make use of a singular value decomposition [2], [9]. This approach inherently poses a length constraint on the vector ξ , assuming that its norm is equal to 1. It does, however, not guarantee positive definiteness of the matrix A . Although positive definiteness of A is not guaranteed, there are only very few practical scenarios in which the estimated matrix A will not be positive definite. A non-positive definite matrix A can for instance be obtained in cases of very limited rotation of the sensor. The problem of allowing a non-positive definite matrix A can be circumvented by instead solving the ellipsoid fitting problem as a semidefinite program [30], [31]

$$\begin{aligned} \min_{A, b, c} \quad & \frac{1}{2} \left\| M \begin{pmatrix} \text{vec } A \\ b \\ c \end{pmatrix} \right\|_2^2, \\ \text{s.t.} \quad & \text{Tr } A = 1, \quad A \in S_{++}^{3 \times 3}, \end{aligned} \quad (23)$$

where $S_{++}^{3 \times 3}$ denotes the set of 3×3 positive definite symmetric matrices. By constraining the trace of the matrix A , (23) avoids the trivial solution of $\xi = 0$. The problem (23) is a convex optimization problem and therefore has a globally optimal solution and does not require an accurate initial guess of the parameter vector ξ . The optimization problem can easily be formulated and efficiently solved using freely available software packages like YALMIP [32] or CVX [33].

Initial estimates of the calibration matrix D and the offset vector o can be obtained from the estimated $\hat{A}, \hat{b}, \hat{c}$ as

$$\beta = \left(\frac{1}{4} \hat{b}^T \hat{A}^{-1} \hat{b} - \hat{c} \right)^{-1}, \quad (24a)$$

$$\tilde{D}_0^T \tilde{D}_0 = \beta \hat{A}^{-1}, \quad (24b)$$

$$\hat{o}_0 = \frac{1}{2} \hat{A}^{-1} \hat{b}, \quad (24c)$$

where $\hat{o}_0$ denotes the initial estimate of the offset vector o . From (24b) it is not possible to uniquely determine the initial estimate of the calibration matrix D . We determine an initial estimate of the calibration matrix D using a Cholesky decomposition, leading to a lower triangular $\tilde{D}_0$. However, any $\tilde{D}_0 U$ where $U U^T = \mathcal{I}_3$ will also fulfill (24b). As discussed in Assumption 1 in Section IV, we assume that the sensor axes of the inertial sensors and the magnetometers are related by a rotation, implying that we restrict the matrix U to be a rotation matrix. The initial estimate $\hat{D}_0$ can therefore be defined in terms of $\tilde{D}_0$ as

$$\hat{D}_0 = \tilde{D}_0 R_D. \quad (25)$$

The unknown rotation matrix R_D will be determined in Section VI-B.

B. Determine misalignment of the inertial and magnetometer sensor axes

The third step of the initial estimation aims at determining the misalignment between the inertial and the magnetometer sensor axes. It also determines an initial estimate of the normalized local magnetic field $\hat{m}_0^n$. These estimates are obtained by combining the magnetometer measurements with the inertial sensor measurements. The approach is based on

the fact that the inner product of two vectors is invariant under rotation. The two vectors considered here are m^n and the vertical $v^n = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}^T$. Hence, it is assumed that the inner product of the vertical v_t^b in the body frame b ,

$$v_t^b = R_t^{bn} v^n, \quad (26a)$$

and the normalized local magnetic field m_t^b in the body frame,

$$m_t^b = R_D^T \tilde{D}_0^{-1} (y_{m,t} - \hat{o}_0), \quad (26b)$$

is constant. The matrix R_D in (26b) denotes the rotation needed to align the inertial and magnetometer sensor axes. The rotation matrix R_t^{nb} in (26a) is a rotation matrix representation of the orientation estimate at time t obtained from an EKF. This EKF is similar to the one described in Section V-B. It does not use the magnetometer measurements, since they have not properly been calibrated yet and can therefore not result in accurate heading estimates. However, to determine the vertical v_t^b , only the sensor's inclination is of concern, which can be determined using the inertial measurements only.

The inner product between m^n and v^n is equal to m_z^n (see also (8a)). Since this inner product is invariant under rotation, we can formulate the following minimization problem

$$\begin{aligned} \min_{R_D, m_{z,0}^n} \quad & \frac{1}{2} \sum_{t=1}^N \|m_{z,0}^n - (v^n)^T R_t^{nb} R_D^T \tilde{D}_0^{-1} (y_{m,t} - \hat{o}_0)\|_2^2, \\ \text{s.t.} \quad & R_D \in SO(3). \end{aligned} \quad (27)$$

The rotation matrix R_D can be parametrized using an orientation deviation from a linearization point similar to the approach described in Section V-B. Hence, (27) can be solved as an unconstrained optimization problem.

Based on these results and (25) we obtain the following initial estimates

$$\hat{D}_0 = \tilde{D}_0 \hat{R}_D, \quad (28a)$$

$$\hat{m}_0^n = \begin{pmatrix} \sqrt{1 - (\hat{m}_{z,0}^n)^2} & 0 & \hat{m}_{z,0}^n \end{pmatrix}^T. \quad (28b)$$

Hence, we have obtained an initial estimate $\hat{\theta}_0$ of the entire parameter vector θ as introduced in Section V.

VII. EXPERIMENTAL RESULTS

A. Experimental setup

Experiments have been performed using two commercially available inertial measurements units (IMUs), an Xsens MTi-100 [34] and a Trivisio Colibri Wireless IMU [35]. The experimental setup of both experiments can be found in Fig. 2. The experiment with the Xsens IMU was performed outdoors to ensure a homogeneous local magnetic field. The experiment with the Trivisio IMU was performed indoors. However, the experiment was performed relatively far away from any magnetic materials such that the local magnetic field is as homogenous as possible. The Xsens IMU was placed in an aluminum block with right angles which can be used to rotate the sensor 90° to verify the heading results. For both sensors, inertial and magnetometer measurements were collected at 100 Hz.

Fig. 2. Top: experimental setup where a calibration experiment is performed outdoors. An Xsens MTi-100 IMU (orange box) together with a magnetic disturbance is placed in an aluminum block. Bottom: experimental setup using a Trivisio Colibri Wireless IMU (black box). A phone is used as a source of magnetic disturbance. To avoid saturation of the magnetometer, the phone is not attached directly to the IMU.

Fig. 3. Calibration results from the experiment with the Trivisio IMU. The ellipsoid of magnetometer data (red) lies on a unit sphere after calibration (blue).

B. Calibration results

For calibration, the IMU needs to be slowly rotated such that the assumption of zero acceleration is reasonably valid. This leads to an ellipsoid of magnetometer data as depicted in red in Figs. 1 and 3. Note that for plotting purposes the data has been downsampled to 1 Hz. To emphasize the deviation of the norm from 1, the norm of the magnetometer data is depicted in red in Fig. 4 for both experiments.

For the experiment with the Xsens IMU, the following calibration matrix $\hat{D}$ and offset vector $\hat{o}$ are found

$$\hat{D} = \begin{pmatrix} 0.74 & -0.13 & 0.01 \\ -0.12 & 0.68 & 0.01 \\ -0.03 & 0.43 & 1.00 \end{pmatrix}, \quad \hat{o} = \begin{pmatrix} 1.36 \\ 1.22 \\ -0.94 \end{pmatrix} \quad (29)$$

Fig. 4. Norm of the magnetic field measurements before (red) and after (blue) calibration for (top) the experiment with the Xsens IMU and for (bottom) the experiment with the Trivisio IMU.

using Algorithm 1. Applying the calibration result to the magnetometer data leads to the unit sphere of data in blue in Fig. 1. The norm of the magnetometer data after calibration can indeed be seen to lie around 1, as depicted in blue in Fig. 4.

As a measure of the calibration quality, we analyze the normalized residuals $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ after calibration from the EKF. For each time t , this is a vector in $\mathbb{R}^6$. In the case of correctly calibrated parameters that sufficiently model the magnetic disturbances, we expect the stacked normalized residuals $\{S_t^{-1/2}(y_t - \hat{y}_{t|t-1})\}_{t=1}^N \in \mathbb{R}^{6N}$ to be normally distributed with zero mean and standard deviation 1. The histogram and a fitted Gaussian distribution can be found in Fig. 5a. The residuals resemble a $\mathcal{N}(0, 1)$ distribution except for the large peak around zero and – not visible in the plot – a small amount of outliers outside of the plotting interval. This small amount of outliers is due to the fact that there are a few measurement outliers in the accelerometer data. Large accelerations can for instance be measured when the setup is accidentally bumped into something and violate our assumption that the acceleration of the sensor is approximately zero. We believe that the peak around zero is due to the fact that the algorithm compensates for the presence of the large residuals.

To analyze if the calibration is also valid for a different (validation) data set with the same experimental setup, the calibrated parameters have been used on a second data set. Figures of the ellipsoid of magnetometer data and the sphere of calibrated magnetometer data are not included since they look very similar to Figs. 1 and 4. The residuals after calibration of this validation data set can be found in Fig. 5b. The fact that these residuals look very similar to the ones for the original data suggests that the calibration parameters obtained are also valid for this validation data set.

Fig. 5. Histogram of the normalized residuals $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ from the EKF after calibration for the estimation data set (left) and for a validation data set (right) for the experiments performed with the Xsens IMU. A Gaussian distribution (red) is fitted to the data.

Fig. 6. Histogram of the normalized residuals $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ from the EKF after calibration for the estimation data set (left) and for a validation data set (right) for the experiments performed with the Trivision IMU. A Gaussian distribution (red) is fitted to the data.

The Trivision IMU outputs the magnetometer data in microtesla. Since our algorithm scales the calibrated measurements to a unit norm, the obtained $\hat{D}$ and offset vector $\hat{o}$ from Algorithm 1 are in this case of much larger magnitude,

$$\hat{D} = \begin{pmatrix} 61.74 & 0.59 & 0.09 \\ -1.01 & 60.74 & 0.23 \\ -0.39 & 0.06 & 60.80 \end{pmatrix}, \quad \hat{o} = \begin{pmatrix} -19.77 \\ -1.68 \\ -6.98 \end{pmatrix}. \quad (30)$$

The sphere of calibrated data and its norm can be found in blue in Figs. 3 and 4. Note that for plotting purposes, the magnetometer data before calibration is scaled such that its mean lies around 1. The obtained $\hat{D}$ and $\hat{o}$ are scaled accordingly to plot the red ellipsoid in Fig. 3. The normalized residuals $S_t^{-1/2}(y_t - \hat{y}_{t|t-1})$ of the EKF using both the estimation and a validation data set are depicted in Fig. 6. For this data set, the accelerometer data does not contain any outliers and the residuals resemble a $\mathcal{N}(0, 1)$ distribution fairly well.

From these results we can conclude that Algorithm 1 gives good magnetometer calibration results for experimental data from two different commercially available IMUs. A good fit of the ellipsoid of data to a sphere is obtained and the algorithm seems to give good estimates analyzed in terms of its normalized residuals. Since magnetometer calibration is generally done to obtain improved heading estimates, it is important to also interpret the quality of the calibration in

terms of the resulting heading estimates. In Section VII-C this will be done based on experimental results. The heading performance will also be analyzed based on simulations in Section VIII.

C. Heading estimation

An important goal of magnetometer calibration is to facilitate good heading estimates. To check the quality of the heading estimates after calibration, the block in which the Xsens IMU was placed (shown in Fig. 2) 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 degrees. The experiment was conducted in Enschede, the Netherlands. The dip angle δ at this location is approximately 67° [29]. Hence, we expect the calibrated magnetometer measurements to resemble rotations of the normalized magnetic field $m^n = (0.39 \ 0 \ -0.92)^T$ (see also (7) and (8b)). The calibrated magnetometer data from the experiment is shown in Fig. 7 and consists of the following stationary time periods:

z-axis up During the period 0 – 105s, the magnetometer is flat with its z -axis pointing upwards. Hence, the z -axis (red) of the magnetometer measures the vertical component of the local magnetic field m_z^n . During this period, the sensor is rotated by 90° around the z -axis into 4 different orientations and subsequently back to its initial orientation. This results in the 5 steps for measurements in the x - (blue) and y -axis (green) of the magnetometer.

z-axis down A similar rotation sequence is performed with the block upside down at 110–195s, resulting in a similar pattern for measurements in the x - and y -axis of the magnetometer. During this time period, the z -axis of the magnetometer measures $-m_z^n$ instead.

x-axis up The procedure is repeated with the x -axis of the sensor pointing upwards during the period 200 – 255s, rotating around the x -axis into 4 different orientations and back to the initial position. This results in the 5 steps for measurements in the y - and z -axis of the magnetometer.

x-axis down A similar rotation sequence is performed with the x -axis pointing downwards at 265 – 325 seconds.

y-axis down Placing the sensor with the y -axis downwards and rotating around the y -axis results in the data at 350 – 430 seconds. The rotation results in the 5 steps for measurements in the x - and z -axis of the magnetometer.

y-axis up A similar rotation sequence is performed with the y -axis pointing upwards at 460 – 520 seconds.

Since the experimental setup was not placed *exactly* vertical, it is not possible to compare the *absolute orientations*. However, it is possible to compare the *difference in orientation* which is known to be 90° due to the properties of the block in which the sensor was placed. To exclude the effect of measurement noise, for each of the stationary periods in Fig. 7, 500 samples of magnetometer and accelerometer data are selected. Their mean values are used to estimate the orientation of the sensor. Here, the accelerometer data is used to estimate the inclination. The heading is estimated from the horizontal component of the magnetometer data. This procedure makes use of the fact that

Fig. 7. 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.

the orientation of the sensor can be determined from two linearly independent vectors in the navigation frame – the gravity and the direction of the magnetic north – and in the body frame – the mean accelerometer and magnetometer data. It is referred to as the TRIAD algorithm [36]. Table I reports the deviation from 90° between two subsequent rotations. Note that the metal object causing the magnetic disturbance as shown in Fig. 2 physically prevents the setup from being properly 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 I.

Our experiment investigates both the heading errors and the improvement of the heading estimates over the ones obtained after the initial calibration, i.e. Step 1 in Algorithm 1. In Table I we therefore include both the heading errors using the initial parameter estimates $\hat{D}_0$ (28a) and $\hat{o}_0$ (24c) and the heading errors using ML parameter estimates $\hat{D}$ and $\hat{o}$ (29) obtained using Algorithm 1. As can be seen, the deviation from 90° is small, indicating that good heading estimates are obtained after calibration. Also, the heading estimates using the initial parameter estimates are already fairly good. The mean error is reduced from 1.28° for the initial estimate to 0.76° for the ML estimate. The maximum error is reduced from 4.36° for the initial estimate to 2.48° for the ML estimate. Note that the results of the ML estimate from Algorithm 1 are slightly better than the results previously reported by [1]. This can be attributed to the fact that we now use orientation error states instead of the quaternion states in the EKF (see Section V-B). This results in slightly better estimates, but also in a smoother convergence of the optimization problem. The quality of the heading estimates is studied further in Section VIII based on a simulation study.

VIII. SIMULATED HEADING ACCURACY

Magnetometer calibration is typically performed to improve the heading estimates. It is, however, difficult to check the heading accuracy experimentally. In Section VII-C, for instance, we are limited to doing the heading validation on a different data set and we have a limited number of available data points. To get more insight into the orientation accuracy that is gained by executing all of Algorithm 1, compared to just its initialization phase (Step 1 in the algorithm), we

engage in a simulation study. In this study we focus on the root mean square (RMS) heading error for different simulated sensor qualities (in terms of the noise covariances and the gyroscope bias) and different magnetic field disturbances (in terms of different values for the calibration matrix D and offset vector o).

In our simulation study, we assume that the local magnetic field is equal to that in Linköping, Sweden. The calibration matrix D , the offset vector o and the sensor properties in terms of the gyroscope bias and noise covariances are all sampled from a uniform distribution. The parameters of the distributions from which the sensor properties are sampled are chosen as physically reasonable values as considered from the authors' experience. The noise covariance matrices Σ_ω , Σ_a and Σ_m are assumed to be diagonal with three different values on the diagonal. The calibration matrix D is assumed to consist of three parts,

$$D = D_{\text{diag}} D_{\text{skew}} D_{\text{rot}}, \quad (31)$$

where D_{diag} is a diagonal matrix with elements D_{11} , D_{22} , D_{33} and D_{rot} is a rotation matrix around the angles ψ , θ , ϕ . The matrix D_{skew} models the non-orthogonality of the magnetometer axes as

$$D_{\text{skew}} = \begin{pmatrix} 1 & 0 & 0 \\ \sin \zeta & \cos \zeta & 0 \\ -\sin \eta & \cos \eta \sin \rho & \cos \eta \cos \rho \end{pmatrix}, \quad (32)$$

where the angles ζ , η , ρ represent the different non-orthogonality angles. The exact simulation conditions are summarized in Table II.

The simulated data consists of 100 samples of stationary data and subsequently 300 samples for rotation around all three axes. It is assumed that the rotation is exactly around the origin of the accelerometer triad, resulting in zero acceleration during the rotation. The first 100 samples are used to obtain an initial estimate of the gyroscope bias $\hat{\delta}_{\omega,0}$ by computing the mean of the stationary gyroscope samples. The covariance matrices $\hat{\Sigma}_{\omega,0}$, $\hat{\Sigma}_{a,0}$ and $\hat{\Sigma}_{m,0}$ are initialized based on the covariance of these first 100 samples. The initial estimate then consists of these initial estimates $\hat{\delta}_{\omega,0}$, $\hat{\Sigma}_{\omega,0}$, $\hat{\Sigma}_{a,0}$, $\hat{\Sigma}_{m,0}$ and the initial calibration matrix $\hat{D}_0$ (28a), the initial offset vector $\hat{o}_0$ (24c) and the initial estimate of the local magnetic field m_0^n (28b).

To study the heading accuracy, the EKF as described in Section V-B is run with both the initial parameter values $\hat{\theta}_0$ and their ML values $\hat{\theta}^{\text{ML}}$. The orientation errors Δq_t , encoded as a unit quaternion are computed using

$$\Delta q_t = \hat{q}_t^{\text{nb}} \odot (q_{\text{ref},t}^{\text{nb}})^c, \quad (33)$$

where $\odot$ denotes a quaternion multiplication and the superscript c denotes the quaternion conjugate (see e.g. [27]). It is computed from the orientation $\hat{q}_t^{\text{nb}}$ estimated by the EKF and the ground truth orientation $q_{\text{ref},t}^{\text{nb}}$. Computing the orientation errors in this way is equivalent to subtracting Euler angles in the case of small angles. However, it avoids subtraction problems due to ambiguities in the Euler angles representation. To interpret the orientation errors Δq_t , they are converted to Euler angles. We focus our analysis on the heading error, i.e. on the third component of the Euler angles.

TABLE I

DIFFERENCE IN ESTIMATED HEADING BETWEEN TWO SUBSEQUENT ROTATIONS AROUND THE SENSOR AXES USING CALIBRATED MAGNETOMETER DATA. THE VALUES REPRESENT THE DEVIATION IN DEGREES FROM 90° . INCLUDED ARE BOTH THE RESULTS USING THE ML ESTIMATES FROM ALGORITHM 1 AND THE RESULTS USING INITIAL ESTIMATES FROM STEP 1 IN THE ALGORITHM.

z -axis				x -axis				y -axis	
z up		z down		x up		x down		y down	
ML	init	ML	init	ML	init	ML	init	ML	init
0.11	0.36	0.69	1.34	0.22	0.16	0.86	1.01	0.18	1.57
0.22	0.90	2.48	4.36	0.07	0.20	1.57	1.45	0.29	0.76
0.46	1.52	1.53	3.57	0.97	0.94	0.61	0.71	0.20	0.78
0.30	0.94	1.92	2.40	0.29	0.59	1.78	1.70	0.50	0.45

TABLE II

SETTINGS USED IN THE MONTE CARLO SIMULATION.

D_{diag}	D_{skew}	D_{rot}	o	δ_ω	Σ_ω	Σ_a	Σ_m
D_{11}, D_{22}, D_{33} $\sim \mathcal{U}(0.5, 1.5)$	ζ, η, ρ $\sim \mathcal{U}(-30^\circ, 30^\circ)$	ψ, θ, ϕ $\sim \mathcal{U}(-10^\circ, 10^\circ)$	o_1, o_2, o_3 $\sim \mathcal{U}(-1, 1)$	$\delta_{\omega,1}, \delta_{\omega,2}, \delta_{\omega,3}$ $\sim \mathcal{U}(-1, 1)$	$\Sigma_{\omega,1}, \Sigma_{\omega,2}, \Sigma_{\omega,3}$ $\sim \mathcal{U}(10^{-3}, 10^{-2})$	$\Sigma_{a,1}, \Sigma_{a,2}, \Sigma_{a,3}$ $\sim \mathcal{U}(10^{-3}, 10^{-1})$	$\Sigma_{m,1}, \Sigma_{m,2}, \Sigma_{m,3}$ $\sim \mathcal{U}(10^{-3}, 10^{-1})$

Fig. 8. Histogram of the heading RMSE using the ML parameter estimate from Algorithm 1 (left, blue) and the initial parameter estimate from Step 1 in the algorithm (right, red). Note the different scales in the two plots.

The RMS of the heading error is plotted for 150 Monte Carlo simulations in Fig. 8. As can be seen, the heading root mean square error (RMSE) using the estimate of the calibration parameters from Algorithm 1 is consistently small. The heading RMSE based on the initialization phase in Step 1 of the algorithm, however, has a significantly larger spread. This clearly shows that orientation accuracy can be gained by executing all of Algorithm 1. Note that in all simulations, analysis of the norm of the calibrated magnetometer measurements as done in Fig. 4 does not indicate that the ML estimate is to be preferred over the estimate from the initialization phase. Hence, analysis of the norm of the calibrated magnetometer measurements does not seem to be a sufficient analysis to determine the quality of the calibration in the case when the calibration is performed to improve the heading estimates.

IX. CONCLUSIONS

We have developed a practical algorithm to calibrate a magnetometer using inertial sensors. It calibrates the magnetometer for the presence of magnetic disturbances, for magnetometer

sensor errors and for misalignment between the inertial and magnetometer sensor axes. The problem is formulated as an ML problem. The algorithm is shown to perform well on real data collected with two different commercially available inertial measurement units.

In future work the approach can be extended to include GPS measurements. In that case it is not necessary to assume that the acceleration is zero. The algorithm can hence be applied to a wider range of problems, like for instance the flight test example discussed in [2]. The computational cost of the algorithm would, however, increase, since to facilitate the inclusion of the GPS measurements, the state vector in the EKF needs to be extended.

Another interesting direction for future work would be to investigate ways of reducing the computational cost of the algorithm. The computational cost of the initialization steps is very small but actually solving the ML problem in Step 2 of Algorithm 1 is computationally expensive. The algorithm both needs quite a large number of iterations and each iteration is fairly expensive due to the computation of the numerical gradients. Interesting lines of future work would either explore different optimization methods or different ways to obtain gradient estimates.

Finally, it would be interesting to extend the work to online estimation of calibration parameters. This would allow for a slowly time-varying magnetic field and online processing of the data.

X. ACKNOWLEDGEMENTS

This work is supported by CADICS, a Linnaeus Center, and by the project *Probabilistic modeling of dynamical systems* (Contract number: 621-2013-5524), both funded by the Swedish Research Council (VR), and by MC Impulse, a European Commission, FP7 research project. The authors would like to thank Laurens Slot, Dr. Henk Luinge and Dr. Jeroen Hol from Xsens Technologies and Dr. Gustaf Hendeby from Linköping University for their support in collecting the data

sets and for interesting discussions. The authors would also like to thank the reviewers for their constructive comments.

REFERENCES

- [1] M. Kok and T. B. Schön, "Maximum likelihood calibration of a magnetometer using inertial sensors," in *Proceedings of the 19th World Congress of the International Federation of Automatic Control*, Cape Town, South Africa, Aug. 2014, pp. 92–97.
- [2] 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*, Singapore, Jul. 2012, pp. 787 – 793.
- [3] N. Bowditch, *The American Practical Navigator*. United States government, 2002.
- [4] R. Alonso and M. D. Shuster, "Complete linear attitude-independent magnetometer calibration," *The Journal of the Astronautical Sciences*, vol. 50, no. 4, pp. 477–490, 2002.
- [5] G. M. Lerner, "Scalar checking," in *Spacecraft attitude determination and control*, J. R. Wertz, Ed. Dordrecht, Holland: D. Reidel Publishing Company, 1978, pp. 328–334.
- [6] 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*, vol. 47, no. 2, pp. 1293–1306, Apr. 2011.
- [7] V. Renaudin, M. H. Afzal, and G. Lachapelle, "Complete triaxis magnetometer calibration in the magnetic domain," *Journal of Sensors*, 2010.
- [8] 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*, vol. 19, no. 2, pp. 87–102, Apr. 2006.
- [9] W. Gander, G. H. Golub, and R. Strebler, "Least-squares fitting of circles and ellipses," *BIT Numerical Mathematics*, vol. 34, no. 4, pp. 558–578, 1994.
- [10] I. Markovsky, A. Kukush, and S. V. Huffel, "Consistent least squares fitting of ellipsoids," *Numerische Mathematik*, vol. 98, no. 1, pp. 177–194, 2004.
- [11] Y. Wu and W. Shi, "On calibration of three-axis magnetometer," *IEEE Sensors Journal*, vol. 15, no. 11, Nov. 2015.
- [12] X. Li and Z. Li, "A new calibration method for tri-axial field sensors in strap-down navigation systems," *Measurement Science and Technology*, vol. 23, no. 10, Oct. 2012.
- [13] 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*, Vila Moura, Portugal, Oct. 2012, pp. 39–44.
- [14] S. Bonnet, C. Bassompierre, C. Godin, S. Lesecq, and A. Barraud, "Calibration methods for inertial and magnetic sensors," *Sensors and Actuators A: Physical*, vol. 156, no. 2, pp. 302–311, 2009.
- [15] 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, Jun. 2013.
- [16] L. Ljung, *System Identification, Theory for the User*, 2nd ed. Prentice Hall PTR, 1999.
- [17] T. B. Schön, A. Wills, and B. Ninness, "System identification of nonlinear state-space models," *Automatica*, vol. 47, no. 1, pp. 39–49, 2011.
- [18] F. Lindsten and T. B. Schön, "Backward simulation methods for Monte Carlo statistical inference," *Foundations and Trends in Machine Learning*, vol. 6, no. 1, pp. 1–143, 2013.
- [19] F. Gustafsson, *Statistical Sensor Fusion*. Studentlitteratur, 2012.
- [20] K. J. Åström, "Maximum likelihood and prediction error methods," *Automatica*, vol. 16, no. 5, pp. 551–574, 1980.
- [21] M. Segal and E. Weinstein, "A new method for evaluating the log-likelihood gradient, the Hessian, and the Fisher information matrix for linear dynamic systems," *IEEE Transactions on Information Theory*, vol. 35, no. 3, pp. 682–687, 1989.
- [22] M. Kok, J. Dahlin, T. B. Schön, and A. Wills, "Newton-based maximum likelihood estimation in nonlinear state space models," in *Proceedings of the 17th IFAC Symposium on System Identification*, Beijing, China, Oct. 2015, pp. 969 – 974.
- [23] J. Kokkala, A. Solin, and S. Särkkä, "Sigma-point filtering and smoothing based parameter estimation in nonlinear dynamic systems," *Preprint*, 2015, arXiv:1504.06173.
- [24] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer Series in Operations Research, 2006.
- [25] F. L. Markley, "Attitude error representations for Kalman filtering," *Journal of guidance, control, and dynamics*, vol. 26, no. 2, pp. 311–317, 2003.
- [26] J. L. Crassidis, F. L. Markley, and Y. Cheng, "A survey of nonlinear attitude estimation methods," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 1, pp. 12–28, 2007.
- [27] J. D. Hol, "Sensor fusion and calibration of inertial sensors, vision, ultra-wideband and GPS," Dissertation No. 1368, Linköping University, Linköping, Sweden, June 2011.
- [28] M. Kok, "Probabilistic modeling for positioning applications using inertial sensors," Licentiate's Thesis No. 1656, Linköping University, Linköping, Sweden, June 2014.
- [29] National Centers for Environmental Information, <https://www.ngdc.noaa.gov/geomag/geomag.shtml>, Accessed on Oct. 2, 2015.
- [30] G. Calafiore, "Approximation of n-dimensional data using spherical and ellipsoidal primitives," *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, vol. 32, no. 2, pp. 269–278, 2002.
- [31] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [32] J. Löfberg, "YALMIP: A toolbox for modeling and optimization in MATLAB," in *Proceedings of the IEEE International Symposium on Computer Aided Control Systems Design, 2004*, Taipei, Taiwan, Sep. 2004, pp. 284 – 289.
- [33] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.0 beta," <http://cvxr.com/cvx>, Sep. 2013.
- [34] Xsens Technologies B. V., <http://www.xsens.com>, Accessed on Jan. 7, 2016.
- [35] Trivisio Prototyping GmbH, <http://www.trivisio.com>, Accessed on Jan. 7, 2016.
- [36] M. D. Shuster and S. D. Oh, "Three-axis attitude determination from vector observations," *Journal of Guidance, Control, and Dynamics*, vol. 4, no. 1, pp. 70–77, 1981.

Manon Kok is a Ph.D. student in Automatic Control at the Department of Electrical Engineering at Linköping University. She received the Tech. Lic. degree in automatic control in June 2014 from Linköping University (Sweden), the MSc degree in Applied Physics in Jan. 2009 and the MSc degree in Philosophy of Science, Technology and Society in June 2007, both from the University of Twente (the Netherlands). Between Feb. 2009 and July 2011 she worked as a research engineer at Xsens Technologies B.V., the Netherlands. Her main research interests

are in probabilistic modeling for position and orientation estimation using inertial sensors.

Thomas B. Schön is Professor of the Chair of Automatic Control in the Department of Information Technology at Uppsala University. He received the PhD degree in Automatic Control in Feb. 2006, the MSc degree in Applied Physics and Electrical Engineering in Sep. 2001, the BSc degree in Business Administration and Economics in Jan. 2001, all from Linköping University. He has held visiting positions with the University of Cambridge (UK), the University of Newcastle (Australia) and Universidad Técnica Federico Santa María (Valparaíso, Chile).

He is a Senior member of the IEEE and an Associate Editor of *Automatica*. He received the Arnerberg prize by the Royal Swedish Academy of Sciences in 2016. He was awarded the *Automatica* Best Paper Prize in 2014, and in 2013 he received the best PhD thesis award by The European Association for Signal Processing (EURASIP). He received the best teacher award at the Institute of Technology, Linköping University in 2009. Schön's main research interest is nonlinear inference problems, especially within the context of dynamical systems, solved using probabilistic methods.