

Robust Heading Estimation Indoors

Jonas Callmer, David Törnqvist, Fredrik Gustafsson

Division of Automatic Control

E-mail: callmer@isy.liu.se, tornqvist@isy.liu.se,
fredrik@isy.liu.se

23rd April 2013

Report no.: LiTH-ISY-R-3060

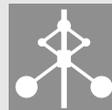
Submitted to 16th International Conference on Information Fusion:
Fusion 2013

Address:

Department of Electrical Engineering
Linköpings universitet
SE-581 83 Linköping, Sweden

WWW: <http://www.control.isy.liu.se>

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET



Abstract

The problem of estimating heading is central in the indoor positioning problem based on measurements from inertial measurement and magnetic units. Integrating rate of turn angular rate gives the heading with unknown initial condition and a linear drift over time, while the magnetometer gives absolute heading, but where long segments of data are useless in practice because of magnetic disturbances. A basic Kalman filter approach with outlier rejection has turned out to be difficult to use with high integrity. Here, we propose an approach based on convex optimization, where segments of good magnetometer data are separated from disturbed data and jointly fused with the yaw rate measurements. The optimization framework is flexible with many degrees of freedom in the modeling phase, and we outline one design. A recursive solution to the optimization is derived, which has a computational complexity comparable to the simplest possible Kalman filter. The performance is evaluated using data from a handheld smartphone for a large amount of indoor trajectories, and the result demonstrates that the method effectively resolves the magnetic disturbances.

Keywords: Heading estimation, magnetometer, gyro, disturbances, optimization

Robust Heading Estimation Indoors using Convex Optimization

Jonas Callmer, David Törnqvist, Fredrik Gustafsson
Department of Electrical Engineering
Linköping University, Linköping, Sweden
callmer@isy.liu.se

Abstract—The problem of estimating heading is central in the indoor positioning problem based on measurements from inertial measurement and magnetic units. Integrating rate of turn angular rate gives the heading with unknown initial condition and a linear drift over time, while the magnetometer gives absolute heading, but where long segments of data are useless in practice because of magnetic disturbances. A basic Kalman filter approach with outlier rejection has turned out to be difficult to use with high integrity. Here, we propose an approach based on convex optimization, where segments of good magnetometer data are separated from disturbed data and jointly fused with the yaw rate measurements. The optimization framework is flexible with many degrees of freedom in the modeling phase, and we outline one design. A recursive solution to the optimization is derived, which has a computational complexity comparable to the simplest possible Kalman filter. The performance is evaluated using data from a handheld smartphone for a large amount of indoor trajectories, and the result demonstrates that the method effectively resolves the magnetic disturbances.

I. INTRODUCTION

Dead reckoning systems are based on a movement and a direction in which the unit is moving. To determine the direction of a user, a magnetometer is often used. It works as a compass, measuring the earth magnetic field and thereby giving a measurement of the unit heading. Indoors, magnetic disturbances from steel structures, electrical wiring, cabinets etc. are often present making magnetometer readings unreliable.

A gyro measures rate of turn, which means that it measures how the heading is changing. It can be used to stabilize the heading estimation, but when used alone the integrated heading estimate is prone to drift off.

In this work we study the problem of accurately estimating heading indoors using an inexact gyro and a magnetometer which is sometimes heavily disturbed.

The problem is illustrated in Figure 1. It shows the cumulative sum of gyro signals with the drift clearly present, the significantly disturbed magnetic heading and the actual ground truth heading. Note that sometimes the magnetic field is undisturbed, providing information of user heading that can and should be used. In this work we aim to support the heading estimation by detecting and utilizing these undisturbed magnetic heading sections.

The problem of how to handle magnetic disturbances for usage in positioning has not received a lot of attention. Often

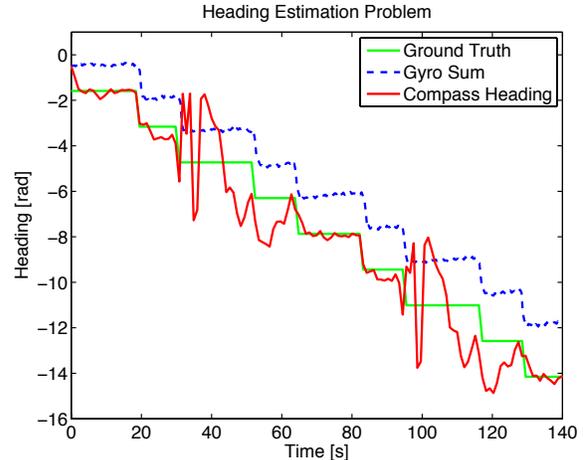


Fig. 1. How do we fuse the disturbed magnetometer signal (solid red) with the biased gyro measurements (dashed blue) to reproduce the ground truth (solid light green)?

the signal has been discarded as unreliable indoors and simply ignored [1, 2].

The most common approach to handle the magnetic disturbances is to study the signal norm and/or the dip angle [3–6]. If either the signal norm or the dip angle deviates from a narrow interval, the magnetic signal is deemed completely unreliable and is subsequently discarded. In Roetenberg *et al* [7] this is handled a bit less abrupt by modelling the magnetometer measurement noise as a function of the signal norm and the dip angle.

Based on our large scale experiments in public venues, our conclusion is that the norm or the dip angle of the magnetic signal cannot be used to find good data segments in a robust way, or even to find just some of the bad ones. The magnetometer signal simply fluctuates too much even during good data segments to be useful.

A slightly different approach was taken in [8] where the magnetic disturbances are detected using a hidden Markov model where the modes represent disturbed and undisturbed data, respectively. The resulting mode estimate affects the parameter settings of a Kalman filter that estimated the attitude.

All in all, ways to detect and handle magnetic disturbances have not received a great deal of attention and the previously suggested solutions almost all focus on features in the mag-

netic signal itself.

Our suggested heading estimation system is based on the assumption that the errors in the gyro measurements are from insufficient calibration. The angular velocity signal is hence seen as primarily corrupted by sensor deficiencies such as a slightly incorrect gain and a small bias. By estimating bias and gain error correction terms and also an initial heading, the current heading can be estimated robustly.

The bias and gain correction parameters are estimated using convex optimization. They are estimated by altering the shape of the vector of the cumulative sum of angular velocity measurements, see the dashed blue signal in Figure 1. The aim is to tweak it until it matches the shape of the vector of magnetometer heading measurements as closely as possible. In the end, the problem is formulated as a weighted least squares problem.

A more standard approach to estimate heading would be to use a Kalman filter. The angular velocity measurements are used as an input in the time update while the magnetometer measurements are used in a measurement update. The disturbances could for example be handled by studying the innovation and discarding the perceived bad measurements.

There are two main shortcomings of such a Kalman filter solution:

- Disturbances are detected using only the latest magnetometer measurement.
- It needs to be reinitialized if the filter diverges.

The proposed method can handle both of these issues. Since it studies the whole trajectory of data, it becomes easier to detect and discard the disturbed sections of the magnetic signal. This makes the proposed method more robust. It will also recover automatically if diverged once more data is available, given that some of that data is undisturbed. Finally, we have derived a computationally cheap solution to the problem, that compares well with a Kalman filter.

If one still wants to use a Kalman filter to estimate heading, a more elaborate scheme than the standard Kalman filter is needed to handle the disturbances and to reinitiate if diverged.

In this work we have restricted the problem to yaw estimation in the 2D plane. The sensor is from a smartphone that is handheld and the method is evaluated on a large number of data sets. Strengths and weaknesses of the method are also discussed.

II. PROBLEM FUNDAMENTALS

The true heading $h(t)$ is a function of the true but unknown angular velocities $\omega(t)$ as

$$h(t) = \int_0^t \omega(\tau) d\tau + h(0). \quad (1)$$

Now we sample $\omega(t)$ using a perfect sensor with sampling time T and zero order hold. That means $\omega_k = \omega(t_k)$. Assuming that the bandwidth of the user movements is small compared to the

sampling frequency, the true time discrete heading h_k is

$$h_k = T \sum_{i=1}^N \omega_i + h_0 \quad (2)$$

where $k = NT$.

Unfortunately, sensors are not perfect, and cheap MEMS gyros are prone to be poorly calibrated why their measurements have slight errors. In this work we have assumed that these errors are a minor constant change in gain \bar{g} , a small constant additive bias \bar{b} and measurement noise $\bar{e}_k \sim \mathcal{N}(0, \bar{\sigma}_\omega^2)$ making the gyro measurements

$$y_k^\omega = \bar{g}\omega_k + \bar{b} + \bar{e}_k. \quad (3)$$

This can be rewritten as

$$\omega_k = g y_k^\omega + b + e_k \quad (4)$$

where $g = \frac{1}{\bar{g}}$ and $b = -\frac{\bar{b}}{\bar{g}}$ which are constant, and $e_k \sim \mathcal{N}(0, \sigma_\omega^2)$. Consequently, we have the fundamental assumption that

$$h_k = T \sum_{i=1}^N (g y_i^\omega + b + e_i) + h_0 \quad (5)$$

where the variance of h_k is $NT^2\sigma_\omega^2$.

Now, g , b and h_0 are estimated as \hat{g} , \hat{b} and \hat{h}_0 and we assume that σ_ω is very small. The estimated heading is then

$$\begin{aligned} \hat{h}_k &= \sum_{i=1}^N T(\hat{g}y_i^\omega + \hat{b}) + \hat{h}_0 \\ &= \hat{g} \sum_{i=1}^N T y_i^\omega + \hat{b} \sum_{i=1}^N T + \hat{h}_0. \end{aligned} \quad (6)$$

Introduce $m_k = \sum_{i=1}^N T y_i^\omega$ and $n_k = \sum_{i=1}^N T$. Since they are given by the sensor and therefore are known, $m_k = m_{k-1} + T y_k^\omega$ and $n_k = n_{k-1} + T$. The heading estimation model is now

$$\hat{h}_k = m_k \hat{g} + n_k \hat{b} + \hat{h}_0. \quad (7)$$

Putting the sought after states in a vector θ

$$\theta = (g \quad b \quad h_0)^T, \quad (8)$$

the estimated heading at time k is

$$\begin{aligned} \hat{h}_k &= (m_k \quad n_k \quad 1) \hat{\theta} \\ &= \Omega_k \hat{\theta}. \end{aligned} \quad (9)$$

Since θ is assumed constant, the vector of all heading estimates $\hat{H}_{0:k} = (\hat{h}_0 \quad \hat{h}_1 \quad \dots \quad \hat{h}_k)^T$ can be written as

$$\begin{aligned} \hat{H}_{0:k} &= \begin{pmatrix} 0 & 0 & 1 \\ m_1 & n_1 & 1 \\ \vdots & \vdots & \vdots \\ m_k & n_k & 1 \end{pmatrix} \hat{\theta} \\ &= \Omega_{0:k} \hat{\theta} \end{aligned} \quad (10)$$

III. HEADING ESTIMATION

To estimate θ we use all the heading measurements acquired by the magnetometer from time 0 to k , $\bar{Y}_{0:k}^h = (y_0^h \ y_1^h \ \dots \ y_k^h)^T$. Some of these will be greatly disturbed while some will be correct.

The approach is based on optimization and estimates θ by altering the shape of $\hat{H}_{0:k}$. The sought after shape is the one that matches $\hat{H}_{0:k}$ to the seemingly accurate sections of $\bar{Y}_{0:k}^h$ as closely as possible. To focus the attention only to those sections of the data, the shape matching is performed using weighted least squares.

A. Detect Disturbed Magnetometer Readings

We therefore need to detect which parts of the magnetometer data that are disturbed and which ones that are not. One measure of the trustworthiness of the magnetometer data is to compare how well the magnetometer and the gyro measurements agree. The residual

$$\lambda_k = y_k^h - y_{k-1}^h - T y_{k-1}^\omega \quad (11)$$

is only close to 0 if the magnetometer and the gyro measure the same change in heading over that time interval. The interval is only one sample here but can be generalized to more samples by summing up measurements if the sampling times differ.

Sometimes disturbed magnetometer data and gyro data coincide over a short interval just by chance, making λ_k small when we do not want it to. This will make disturbed measurements seem trustworthy. To avoid the problems this can introduce and also to capture the fact that disturbances are location dependent and therefore time dependent, λ_k is averaged using a sliding window.

$$\bar{\lambda}_k = \frac{1}{N} \sum_{i=k-N+1}^k |\lambda_i| \quad (12)$$

The length of the window N is a tradeoff between having false positives and false negatives. We averaged the absolute value of λ_k to avoid problems when λ_k was changing sign in the interval. Since λ_0 is unknown it is given a value reflecting a poor measurement.

The actual weight function used in the vector shape matching is

$$w(\bar{\lambda}_k) = e^{-\bar{\lambda}_k^2/s}. \quad (13)$$

It is close to 1 when $\bar{\lambda}_k$ is small and to 0 when $\bar{\lambda}_k$ is large. s determines the shape of $w(\bar{\lambda}_k)$, thereby stating what a "small" or a "large" value of $\bar{\lambda}_k$ is.

In the end, the weight function is quite ad hoc. It makes sense when studied step by step, but we cannot in any way prove that it is the best choice of weight. It is reasonable that the weights should be large when the magnetometer and gyro seem to agree. It is also reasonable that we should smooth out the measured difference between the sensors a bit over time, since disturbances are often location dependent. $\bar{\lambda}_k$ should then be transformed so that if $\bar{\lambda}_k$ is large $w(\bar{\lambda}_k) \sim 0$ and if $\bar{\lambda}_k$ is small $w(\bar{\lambda}_k) \sim 1$. The actual values used on N and

s have been determined by studying a number of challenging datasets. At the end of the day, what is important is getting a weight vector that looks reasonable.

The weight function is included in the optimization problem as a diagonal matrix $W(\bar{\lambda}_{0:k})$ with $\{w(\bar{\lambda}_i)\}_{i=0}^k$ as its entries.

B. Parameter Estimation

The parameter estimation can be solved as a weighted least squares problem. We know that the approximate true value of the parameters are $\hat{g} \approx 1$ and $\hat{b} \approx 0$, so we should put a restriction on their deviation from these values. Since there are no hard boundaries on these values, giving fixed intervals for the values as constraints is not a good idea.

The parameter values should instead be kept in order using regularization. That means that the deviation of the parameter value from the assumed one is associated with an additional cost in the cost function. It is not obvious though how to add regularization terms to the cost function, since the dimension of $\hat{H}_{0:k}$ is constantly growing and therefore also the cost is growing.

There are two basic ways to solve this problem: normalizing the regularization terms using the actual cost or using the dimension of $\hat{H}_{0:k}$. In the first we solve the optimization problem with no restrictions on \hat{b} and \hat{g} as

$$C_k = \text{minimize}_{\hat{b}_k, \hat{g}_k, \hat{h}_{0,k}} \|\bar{Y}_{0:k}^h - \hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})\|_{W(\bar{\lambda}_{0:k})}^2 \quad (14)$$

and note the final cost C_k . We then resolve the problem with the additional regularization terms that depend on C_k and the parameter values.

$$\text{minimize}_{\hat{b}_k, \hat{g}_k, \hat{h}_{0,k}} \|\bar{Y}_{0:t}^h - \hat{H}_{0:t}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})\|_{W(\bar{\lambda}_{0:k})}^2 + f_g^{C_k}(C_k, \hat{g}_k) + f_b^{C_k}(C_k, \hat{b}_k) \quad (15)$$

An alternative would be to normalize the regularization terms using the dimension of $\bar{Y}_{0:k}$. Unfortunately, how the cost C_k grows with the dimension will be different for each dataset, why such a normalization will be very approximative.

In this work we therefore normalized using the cost C_k and used the regularization terms

$$f_g^{C_k}(C_k, \hat{g}_k) = \beta C_k (\hat{g}_k - 1)^2 \quad (16)$$

$$f_b^{C_k}(C_k, \hat{b}_k) = \alpha C_k \hat{b}_k^2 \quad (17)$$

Both have a quadratic cost for deviating from the nominal values of \hat{b}_k and \hat{g}_k . This means that we allow small changes of \hat{b}_k and \hat{g}_k , while large changes are associated with great costs. Also, the values of α and β are chosen so that these costs relate well to each other. Since the deviation in \hat{b}_k is likely to be larger than in \hat{g}_k , we use $\beta \gg \alpha$. Their actual values have in the end been chosen by studying the experimental results of a couple of datasets.

C. Magnetic Heading Vector Unwrapping

Before one can solve (15) there is a crucial preprocessing step. In order to match $\bar{Y}_{0:k}^h$ and $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$ they have to look like the gyro sum and the magnetic heading in Figure 1. When measured, the raw magnetometer heading measurements

$Y_{0:k}^h \in [-\pi \pi]$ but such a vector cannot be straightforwardly compared to $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$. First it has to be unwrapped.

Unfortunately, it is extremely important that the unwrapping is done correctly. For the estimation to work, it is crucial that the magnetic heading is unwrapped so that it is centered around the ground truth at all times, like in Figure 1. If it is not, $\hat{H}_{0:k}(\hat{b}_k, \hat{g}_k, \hat{h}_{0,k})$ will be tweaked to match a vector that has the wrong shape, ruining the estimation.

Of course, we do not have the ground truth to do the unwrapping around. Instead we have to use the gyro sum for this. We add multiples of 2π to the magnetic vector as $\bar{Y}_i^h = Y_i^h + j \cdot 2\pi$, $j = \{\dots, -1, 0, 1, \dots\}$, so that $-\pi \leq \bar{Y}_i^h - \hat{H}_i(0, 1, Y_0^h) \leq \pi$ for all i . This works fine if the dataset is short and/or the gyro bias is small. Else, eventually the gyro drift will have caused $\hat{H}_i(0, 1, Y_0^h)$ to drift off, resulting in an incorrectly unwrapped magnetic heading vector and therefore poor heading estimation.

The alternative of unwrapping around $\hat{H}_i(0, 1, Y_0^h)$ is to unwrap using $\hat{H}_i(\hat{b}_{k-1}, \hat{g}_{k-1}, \hat{h}_{0,k-1})$. This works perfectly fine as long as $\hat{\theta}_{k-1}$ is estimated correctly. But, if there is an error in the estimation this will result in $Y_{0:k}^h$ being incorrectly unwrapped. When later used in estimation, the incorrectly unwrapped $\bar{Y}_{0:k}^h$ will make $\hat{\theta}_k$ even more wrong. The system hence becomes unstable. Primarily this can happen in datasets where large magnetic disturbances occur early.

Since unwrapping using $\hat{H}_i(0, 1, Y_0^h)$ works fine for short datasets and using $\hat{H}_i(\hat{b}_{k-1}, \hat{g}_{k-1}, \hat{h}_{0,k-1})$ works well for long datasets, we use them like that. For the first 30 seconds of data, the raw gyro sum measurements are used for unwrapping. When more than 30 seconds of data is available, the estimated heading vector is used. This has turned out to be a stable way to unwrap the magnetic heading vector as later experiments will show.

D. Solver Outline

Solving (15) once $\bar{Y}_{0:k}^h$ is correctly unwrapped is straightforward since it is just a weighted least squares problem. Even though the dimensions of $\bar{Y}_{0:k}^h$, $\hat{H}_{0:k}$ and $W(\bar{\lambda}_{0:k})$ grow with time, the solving time is constant using Cholesky decomposition.

Derivating the cost function and putting it equal to zero results in (14) being solved as

$$A_k \hat{\theta}_k = b_k \quad (18)$$

where A_k is a symmetric 3×3 matrix and b_k is 3×1 . Using Cholesky decomposition, $A_k = L_k L_k^T$, (18) can be solved very cheaply. Over time, A_k also evolves in a structured way as $A_k = A_{k-1} + x_k x_k^T$ where

$$x_k = \begin{pmatrix} m_k w_k^{1/2} & n_k w_k^{1/2} & w_k^{1/2} \end{pmatrix}^T \quad (19)$$

where $w_k = w(\bar{\lambda}_k)$. Therefore, the Cholesky decomposition only has to be done once and then it is updated with each new measurement using rank-one updates, $L_k L_k^T = L_{k-1} L_{k-1}^T + x_k x_k^T$

Calculating the actual cost C_k is also cheap since the terms in the cost function evolve in a structured way.

$$\begin{aligned} C_k &= (\bar{Y}_{0:k} - \Omega_{0:k} \hat{\theta}_k)^T W_{0:k} (\bar{Y}_{0:k} - \Omega_{0:k} \hat{\theta}_k) \\ &= \bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k} - 2 \hat{\theta}_k^T \Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k} + \hat{\theta}_k^T \Omega_{0:k}^T W_{0:k} \Omega_{0:k} \hat{\theta}_k \\ &= \underbrace{\bar{Y}_{0:k-1}^T W_{0:k-1} \bar{Y}_{0:k-1} + \bar{y}_k^2 w_k}_{\bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k}} \\ &\quad - 2 \hat{\theta}_k^T \underbrace{(\Omega_{0:k-1}^T W_{0:k-1} \bar{Y}_{0:k-1} + \bar{y}_k w_k \Omega_k^T)}_{\Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k}} \\ &\quad + \hat{\theta}_k^T \underbrace{(\Omega_{0:k-1} W_{0:k-1} \Omega_{0:k-1} + x_{k+1} x_{k+1}^T)}_{\Omega_{0:k}^T W_{0:k} \Omega_{0:k}} \hat{\theta}_k \end{aligned} \quad (20)$$

where $W_{0:k} = W(\bar{\lambda}_{0:k})$ and $w_k = w(\bar{\lambda}_k)$. So if $\bar{Y}_{0:k}^T W_{0:k} \bar{Y}_{0:k}$, $\Omega_{0:k}^T W_{0:k} \bar{Y}_{0:k}$ and $\Omega_{0:k}^T W_{0:k} \Omega_{0:k}$ are saved for the next iteration, updating the cost function will then only require multiplication of 3-dimensional vectors and 3×3 dimensional matrices.

Now, when (15) is derived one gets

$$(A_k + \bar{C}_k) \hat{\theta}_k = \bar{b}_k \quad (21)$$

where

$$\bar{C}_k = \begin{pmatrix} 2C_k \beta & 0 & 0 \\ 0 & 2C_k \alpha & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (22)$$

$A_k + \bar{C}_k$ can be subsequently be Cholesky decomposed as $L_k L_k^T + c_{\alpha,k} c_{\alpha,k}^T + c_{\beta,k} c_{\beta,k}^T$ using two more rank-one updates where

$$\begin{aligned} c_{\beta,k} &= (\sqrt{2\beta C_k} \quad 0 \quad 0)^T \\ c_{\alpha,k} &= (0 \quad \sqrt{2\alpha C_k} \quad 0)^T \end{aligned} \quad (23)$$

Thus, (15) can be solved very cheaply no matter the dimension of $\bar{Y}_{0:k}^h$, $\hat{H}_{0:k}$ and $W(\bar{\lambda}_{0:k})$.

IV. EXPERIMENTAL RESULTS

The devices used for data collection are inertial magnetic measurement units from hand held smartphones. Sensors used in smartphones are not so well calibrated, since calibration is an expensive procedure. Therefore, gyro bias and scale errors are common problems in these low grade sensors. We will present the detailed results from one experiment and then mass evaluations using a large number of datasets.

When a new magnetometer heading measurement Y_k is available, $\hat{\theta}_k$ is reestimated using all data. \hat{h}_k is then computed as (9) and stored in a vector $\bar{H}_{0:k} = [\bar{H}_{0:k-1} \quad \hat{h}_k]$. $\bar{H}_{0:k}$ is referred to as the filtering estimate since it for each time instant contains the estimate that was given at that time.

Later, when more data is available, improved estimates of earlier headings $\hat{H}_{0:k-1}$ are readily available from (10), since θ is assumed constant throughout the experiment. $\hat{H}_{0:\bar{k}}(\hat{\theta}_{\bar{k}})$ where \bar{k} is the final time instant, will be referred to as the smoothing estimate.

In almost all datasets, $\hat{H}_{0:\bar{k}}(\hat{\theta}_{\bar{k}})$ will be a better estimate than $\bar{H}_{0:\bar{k}}$. In a few rare cases, the final magnetometer measurements pull the estimate away from the true values. This

causes the smoothing trajectory to be off during a larger part of the trajectory, making it slightly worse than the filtering version.

A. Detailed Experiment

In the detailed dataset an HTC Sensation XE smartphone was used for data collection. A gyro sampling rate of 25 Hz and a magnetometer sampling rate of 1 Hz were used. In (11), this difference in sampling time has been handled by summing up all the angular velocity measurements between the magnetometer measurements.

The route used for experimental validation is the same as in Figure 1. It has four turns of 90 degrees and each lap takes about 65 seconds to walk. A major magnetic disturbance is present each lap for about 45 seconds. That means, a majority of the trajectory is constituted of a severe magnetic disturbance. There is also a significant bias in the angular velocity measurements.

The filtering results are shown in solid black in Figure 2. In the top plot it is clear that two segments of good data, one in the beginning and one starting around $t = 70$ seconds, are needed to stabilize the results. Before the second segment the parameter estimates, bottom plot, fluctuate quite a lot, but after around 70 seconds they become more stable. The magnetometer disturbances starting around $t = 95$ therefore have a much smaller influence on the filter estimates than the disturbances starting around $t = 30$.

Also included are the smoothing results in dashed magenta. The main difference between these and the filtering results is that also the early disturbances can now be handled making the smoothing estimate close to ground truth for the entire experiment. The final estimated parameter values are $\hat{g} = 1.005$, $\hat{b} = -0.020$ and $\hat{h}_0 = -1.668$. The mean absolute error is 0.924 rad for the magnetic heading, 0.476 rad for the filtering result and 0.243 for the smoothing result.

The middle plot in Figure 2 shows the weights calculated as in (13). It is clear that the weights are the largest and most consistent for the good magnetometer data sections around $t = 20$, $t = 80$ and $t = 140$. During the disturbed periods the weights are most often low. There are some spikes in the weights in sections where the magnetometer disturbances happen to look like the gyro data, like at $t = 120$, but this is more or less unavoidable. Their influence on the result is small in the end.

B. Mass Experiments

We have also conducted mass experiments, 651 datasets from more than 50 venues in 5 continents, to evaluate the performance of the heading estimation system. For each dataset, the estimated heading is compared to the ground truth that has been created based on knowledge about how the user was walking. The datasets are between 15 seconds and 8 minutes long.

The results are shown as mean absolute error for each dataset. In Figure 3, each experiment is indicated as a dot. The value on the x axis is the mean absolute error of the filter

result while the value on the y axis is the mean absolute error of the magnetic heading. The red line indicate whether an improvement has occurred or not. All dots above the line are experiments where the estimate is better than the magnetic heading, while all dots below the line mean that the result has worsened. The green diamond is the mean result which is 0.334 rad for the magnetic heading, 0.266 for the filtering and 0.208 for the smoothing.

Figure 3 show the filtering results. That means, for each time instant only information up until that time instant is used to estimate the heading. In Figure 4, the smoothing estimate is shown. That means that all data has been used to estimate all the headings. The smoothing results are better since there are fewer dots under the red line, and for most of the experiments the dots are moved to the left compared to the filtering results, indicating a reduction in mean absolute estimation error.

In practice, the main difference between the filtering results and the smoothing results is in the beginning of the dataset. After the estimation has converged, the filtering and smoothing results are more or less identical.

The experiments show that the system can produce reliable estimates of heading even under heavy magnetic disturbances and that it can recover from poor early estimates.

V. DISCUSSION

We have presented a method to estimate heading indoors. It is based on the assumption that the gyro measurements are correct down to a slight multiplicative error and an additive error. These errors are estimated from the magnetic heading measurements using weighted least squares. The experimental results indicate that the method can produce robust accurate estimates of user heading. We will now discuss the pros and cons of the system and how easy it is to work with.

The method uses the entire data set to reestimate the parameters at each iteration. An advantage with this approach is that earlier mistakes can be corrected once more data is available. Computationally this method is also extremely cheap and the number of computations for each iteration is constant throughout the dataset, Section III-D.

Convex optimization can be a powerful tool in solving large scale signal processing problems due to the impressive solvers available, but sometimes fitting the problem into the solver framework requires a bit of ad hoc tampering. This led to, in our case, that the weighted optimization solution using batched data was not that easy to trim. Some terms like the weights and the regularization terms are chosen quite ad hoc.

In the end, the solver is only looking for the cheapest solution given the cost function and it can sometimes be quite hard to figure out why an occasional poor solution can be so cheap and what can be done to adjust the cost function to avoid it. Even though all steps seem reasonable when they are made, a system where they are all put together can turn out to be hard to trim. Therefore care has to be taken when designing a convex optimization based solution to a signal processing problem, but in the end the rewards can be great.

A. Conclusions

The fundamental assumption that the summation of the angular velocities is correct down to a few parameters, seems viable. The errors in estimated heading acquired in the experiments are commonly small and the system can handle significant magnetic disturbances.

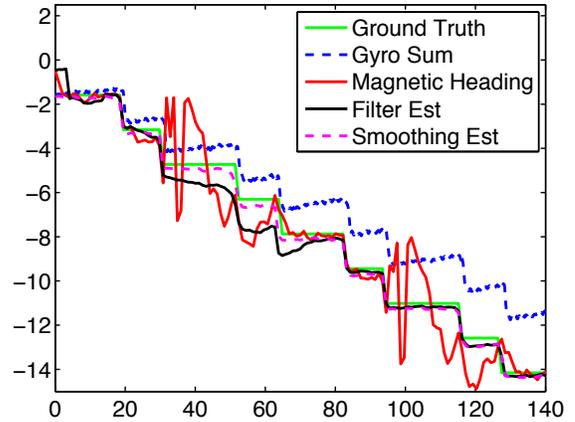
Convex optimization offers powerful solutions to signal processing problems but one has to be careful when designing the problem formulation. The method is somewhat less forgiving than for example the Kalman filter but its ability to, among others, handle large sets of data compensates for that kind of limitations.

ACKNOWLEDGMENTS

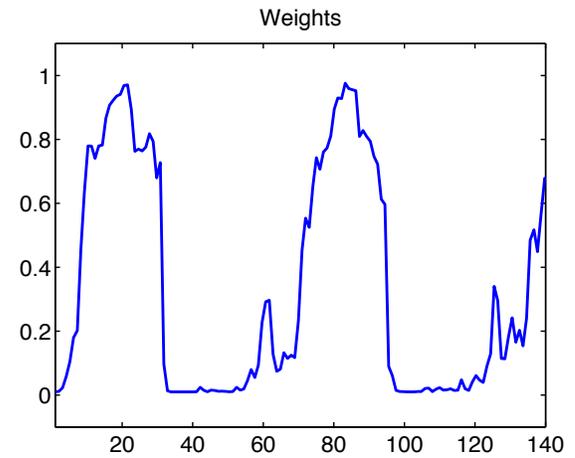
The authors would like to thank prof Stephen Boyd at Stanford University and Daniel Petersson at Automatic Control, Linköping University for helpful discussions.

REFERENCES

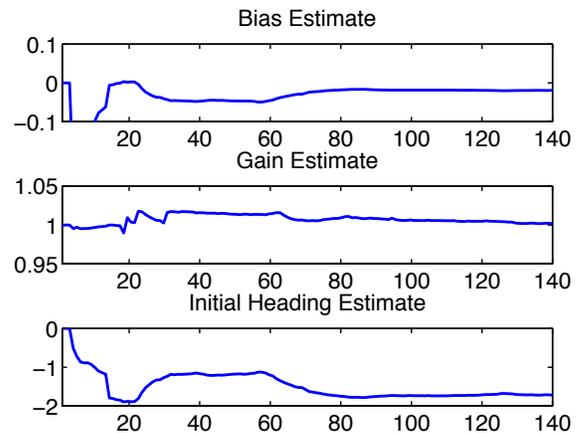
- [1] R. Stirling, J. Collin, K. Fyfe, and G. Lachapelle, "An innovative shoe-mounted pedestrian navigation system," in *Proceedings of European Navigation Conference GNSS*, 2003.
- [2] J. Borenstein and L. Ojeda, "Heuristic drift elimination for personnel tracking systems," *Journal of Navigation*, vol. 63, no. 3, pp. 591 – 606, 2010.
- [3] A. M. Sabatini, "Quaternion-based extended kalman filter for determining orientation by inertial and magnetic sensing," *IEEE Transactions on Biomedical Engineering*, vol. 53, no. 7, pp. 1346 – 1356, 2006.
- [4] J. K. Lee and E. J. Park, "Minimum-order kalman filter with vector selector for accurate estimation of human body orientation," *IEEE Transactions on Robotics*, vol. 25, no. 5, pp. 1196 – 1201, 2009.
- [5] T. Harada, H. Uchino, T. Mori, and T. Sato, "Portable absolute orientation estimation device with wireless network under acceleration situation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.
- [6] A. M. Sabatini, "Estimating three-dimensional orientation of human body parts by inertial/magnetic sensing," *Sensors*, vol. 11, no. 2, pp. 1489 – 1525, 2011.
- [7] D. Roetenberg, H. J. Luinge, C. T. M. Baten, and P. H. Veltink, "Compensation of magnetic disturbances improves inertial and magnetic sensing of human body segment orientation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 13, no. 3, pp. 395 – 405, 2005.
- [8] C. W. Kang and C. G. Park, "Improvement of attitude estimation using hidden markov model classification," in *AIAA Infotech@Aerospace 2010 Conference and Exhibit*, 2010.



(a) Magnetic heading in solid red, ground truth in solid light green, cumulative summation of gyro measurements in dashed blue, the filter heading estimate in solid black and the smoothing heading estimate in dashed magenta. The filter estimate becomes stable after 75 seconds while the smoothing estimate is stable for the entire data set.



(b) The weights assigned to the heading measurements used for heading estimation. For the most time the good magnetometer heading sections have a high weight and the disturbed ones have a low.



(c) The estimated bias, gain and initial heading parameters. After the second good magnetometer data segment is found around $t = 80$, the parameters converge to their final values: $\hat{b} = -0.020$, $\hat{g} = 1.005$ and $\hat{h}_0 = -1.668$.

Fig. 2. [top] Experimental results, [middle] weights given to the magnetometer headings, [bottom] estimated parameter values.

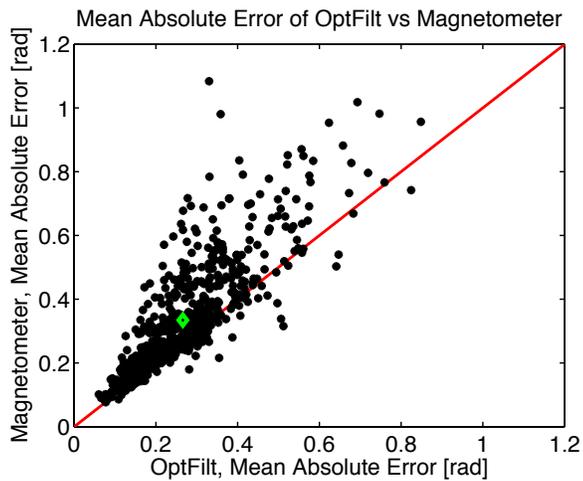


Fig. 3. Mean absolute error of the filtering result plotted versus the mean absolute error of the magnetic heading. Each dot represent one dataset. All dots above the line indicate the result has improved and the further to the left the smaller the error.

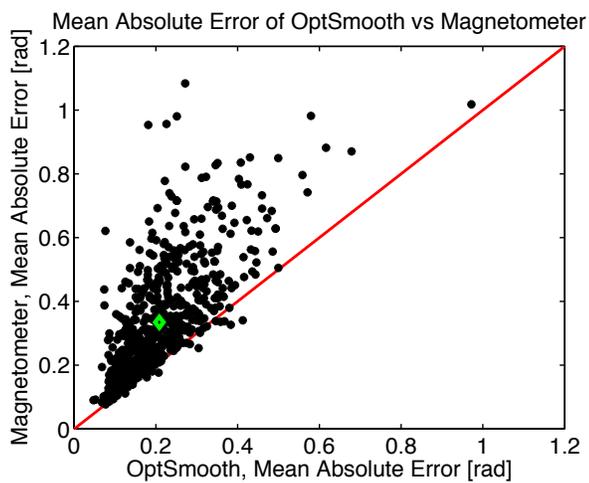


Fig. 4. Mean absolute error of the smoothing result and the magnetic heading. Compared to Figure 3 most of the dots have moved to the left, since the smoothing results are better than the filtering results.

	Avdelning, Institution Division, Department Division of Automatic Control Department of Electrical Engineering	Datum Date 2013-04-23
	Språk Language <input type="checkbox"/> Svenska/Swedish <input checked="" type="checkbox"/> Engelska/English <input type="checkbox"/> _____	Rapporttyp Report category <input type="checkbox"/> Licentiatavhandling <input type="checkbox"/> Examensarbete <input type="checkbox"/> C-uppsats <input type="checkbox"/> D-uppsats <input checked="" type="checkbox"/> Övrig rapport <input type="checkbox"/> _____
URL för elektronisk version http://www.control.isy.liu.se		LiTH-ISY-R-3060
Titel Robust Heading Estimation Indoors Title		
Författare Jonas Callmer, David Törnqvist, Fredrik Gustafsson Author		
Sammanfattning Abstract <p>The problem of estimating heading is central in the indoor positioning problem based on measurements from inertial measurement and magnetic units. Integrating rate of turn angular rate gives the heading with unknown initial condition and a linear drift over time, while the magnetometer gives absolute heading, but where long segments of data are useless in practice because of magnetic disturbances. A basic Kalman filter approach with outlier rejection has turned out to be difficult to use with high integrity. Here, we propose an approach based on convex optimization, where segments of good magnetometer data are separated from disturbed data and jointly fused with the yaw rate measurements. The optimization framework is flexible with many degrees of freedom in the modeling phase, and we outline one design. A recursive solution to the optimization is derived, which has a computational complexity comparable to the simplest possible Kalman filter. The performance is evaluated using data from a handheld smartphone for a large amount of indoor trajectories, and the result demonstrates that the method effectively resolves the magnetic disturbances.</p>		
Nyckelord Keywords Heading estimation, magnetometer, gyro, disturbances, optimization		