Can Magnetic Field Sensors Replace Gyroscopes in Wearable Sensing Applications?

Kai Kunze, Gernot Bahle, Paul Lukowicz
Embedded Systems Lab
University Passau, Passau, Germany
first.lastname@uni-passau.de

Kurt Partridge
Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304
kurt@parc.com

Abstract

We investigate if and how magnetic sensors can be used to replace gyroscopes in wearable activity recognition. The work is motivated by (1) sensor configurations typically found in smart phones where magnetic sensors are used to complement GPS position with orientation and (2) the fact that gyroscopes are an important source of information for activity recognition. We propose a method to compute angular velocity from 3D magnetic sensor data and discuss its fundamental limitations. We present an elaborate evaluation of the accuracy on 5 previously published data set with a total of nearly 20 hours of data from 15 users with activities ranging from bicycle repair, through homemaking to gym exercise.

1. Introduction

After accelerometers, gyroscopes are the second most popular type of wearable motion sensors. To a degree both sensor types provide similar information. However, there are some significant differences. The accelerometer signal reflects a mixture of earth gravity, change in linear motion speed (linear acceleration) and forces related to rotational motions (centripetal acceleration, Coriolis acceleration). The gyroscope, on the other hand, is insensitive to gravity and linear acceleration, providing information about angular velocity only. In general, accelerometers tend to be more useful when complex motions or motion sequences need to be recognized with a single sensor type, since their signal contains information that the gyroscope misses entirely. This comes at the price of ambiguity since the different contributions to the signal can not be separated. Gyroscope are more appropriate whenever exact, non ambiguous information related to rotational motions is relevant. Thus, they are for example extensively used in gait analysis (limb motion is essentially a rotation around the corresponding joint). It has been repeatedly demonstrated that a combination of accelerometers and gyroscopes can lead to better system performance than each sensor on its own [10]. In previous work, we have also shown that a combination of an accelerometer and gyroscope is well suited to compensate for sensor displacement effects [7]. Displacement means a system is trained with a sensor at a particular body location and tested with the sensor shifted (but remaining on the same body part).

Our work is motivated by the observation that mobile phones are often equipped with magnetic field sensors, yet lack gyroscopes. On the other hand, smart phones are increasing being used as activity recognition and motion monitoring devices. As phones are more often loosely placed in a pocket or back than firmly fixed to a given body location, displacement effects can be relevant. Thus, the question arises, if and how a magnetic field sensor can be used to measure angular velocity and enhance activity recognition and motion monitoring applications just like a gyroscope. Note that this is relevant not just to compensate for a specific sensor choice currently made in mobile phone design. Being able to detect absolute orientation, magnetic field sensors are valuable for many applications. Showing that at the same time the sensor can be used as a gyroscope replacement would allow many wearable systems to get ”two sensors for the price of one”, saving cost, size and energy.

Paper Contributions. Magnetic sensors are mostly used to determine orientation with respect to magnetic ”North”. Since angular velocity can be defined as the rate of orientation change with respect to a fixed point, in theory, magnetic field sensors should be able to replace gyroscopes[1]. Unfortunately, in practice, issues such as inhomogeneity of the earth magnetic field, magnetic disturbances (e.g. due to electrical appliances) and numerical stability issues mean that it is not obvious if and how magnetic field sensors can be used in place of gyroscopes. To answer this question, this paper makes the following contributions:

1) We describe a method for deriving 3D angular velocity information from a 3-axis magnetic field sensor.
2) We elaborate the sources of error and fundamental limitations on the approximation of a gyroscope signal with a magnetic field sensor.
3) We evaluate the estimation accuracy for the angular velocity on five previously published data sets.
4) On a previously published problem we investigate how much the accuracy differences affect activity inference.

1. Except when the axis of rotation is parallel to the field lines, as explained later on.

Related Work. So far, we are unaware of other work using a 3D magnetic field sensor to estimate angular velocity for activity recognition. Quite a number of researchers combine gyroscopes, accelerometers and magnetic field sensors to deduce context in a variety of situations (e.g. [8], [5], [1]). Some work specifically deals with placement-indifferent, more robust inferences: Blanke et. al. deduce the location of a person using a gyro in a pocket [2], and Förster et. al. show how to use a clustering approach to gain robustness against displacement for motion sensor based systems [4].

2. Approach

In this section we describe an approach to estimating angular velocity using only magnetometer data.

Magnetic Sensor Signals and Rotation. A naive notion of a magnetic field sensor's functionality is that, like an analog compass, it points straight north. However, what the needle of an analog compass really does, is to orient itself parallel to the tangential of the magnetic field line at the specific location. Similarly, the three field strength components that a 3D magnetic field sensor outputs represent a vector $\vec{B}$ that is tangential to the magnetic field line at sensor location. This vector is given in the local coordinate system of the sensor with the vector length representing the scalar field strength (norm of the field vector at the location). Thus, if we orient the sensor in such a way that one sensor axis (e.g. x-axis) points in the direction of the magnetic field (is tangential to the field line) then the sensor reading will be $\vec{B(t)} = (b, 0, 0)$ with $b = \|\vec{B(t)}\|$ being the magnetic field strength at the location. If we orient the sensor with the (x, y) plane being tangential to the field line than the output will be $(b * cos(\varphi), b * (sin(\varphi), 0)$. Generalizing to arbitrary orientations of the sensor with respect to the field line we have:[2] $B_i(t) = \|\vec{B(t)}\| \cdot cos(\varphi_i)$ The angle $\varphi_i(\vec{B}(t))$ between the i-th axis and the magnetic field strength vector $\vec{B}(t)$ measured at time t is then given by $\varphi_i(\vec{B}(t)) = \arccos \frac{B_i(t)}{\|\vec{B}(t)\|}$ where $B_i(t)$ is the i-th component of $\vec{B}(t)$.
Angular velocity then equals the first derivative of the angle:

$$\varphi_i(\vec{B}(t))' = -\frac{B_i(t)' + \frac{B_i(t) \cdot \|\vec{B}(t)\|'}{\|\vec{B}(t)\|}}{\sqrt{\|\vec{B}(t)\|^2 - B_i(t)^2}} \tag{1}$$

Since measurements happen at discrete points in time, a continuous differential is not available. Instead, differentials of $\vec{B}(t)$ and $\|\vec{B(t)}\|$ have to be approximated by difference quotients. One possible definition is given by taking the average of the differences between measurements at times $(t+1, t)$ and $(t, t-1)$, divided by the actual time that elapsed between those measurements. If the sample rate is completely uniform, then the difference in timestamps can be shortened to $\frac{1}{f}$, where f denotes the sampling frequency:

$$\vec{B}(t)' = \frac{f}{2} \cdot (\vec{B}(t+1) - \vec{B}(t-1)) \tag{2}$$

Problems and Limits. An obvious concern is that magnetic disturbances caused by electrical appliances and metallic objects in the environment. Such disturbances are known to cause significant problems whenever the magnetic field is used to estimate orientation (e.g. for inertial navigation or in MARG motion tracking systems). Interestingly they have a much smaller effect on the estimation of angular velocity from magnetic sensor signals. This is because, as described above, the estimation does not involve absolute orientation. Instead it relies on the orientation of the sensor with *respect to the local magnetic field lines*. Thus, for the angular speed calculation, it does not matter whether the local magnetic field corresponds to the true earth field. However, this assumes that the direction of the local magnetic lines remains constant between the two field measurements that are used to compute the angular velocity. Unfortunately, this is not always the case. First of all, many disturbances are not constant (e.g. fields caused by electric motors). In addition, we have to take into account the fact that magnetic field lines are curved. Thus, linear displacement of the sensor (with no rotational components) can lead to a change of the angle between the sensor and the field lines, since at different locations the magnetic field lines point in different directions. As long as we are dealing only with the earth magnetic field this effect can be neglected due to extremely small curvature of the field lines. However in the presence of environmental fields with stronger curvatures, significant errors can be caused by this effect.

A second source of problems stems from the fact that rotations around an axis parallel to the field line will not lead to a change of angle between the sensor and the field line. Furthermore, rotations around axes that are close to parallel to the field line will lead to small angle changes only and are likely to produce noisy estimations.

3. Signal Level Evaluation

DataSets. The above discussion shows that the ability to accurately estimate angular velocity from 3D magnetic field sensors depends on a broad range of environmental factors as well as on the types of motion that are being performed. To better understand the practical implications for wearable systems we have thus evaluated our computation method on five previously published data sets that span different application domains and activity types. Overall, the evaluation set contains nearly 20 hours of data from 15 users and activities ranging from bicycle repair, car inspection, homemaking, having breakfast, video gaming, and gym exercises. All data sets were based on XSENSE MTX inertial sensor modules containing both a gyroscope and a magnetic field sensor.

2. This assumes the magnetic sensor delivers data in a standard basis coordinate system; it is, however, possible to perform a basis transformation if that is not the case for a given device.

Description	Placement	Citation
office and home setting, 3 subjects, 9 hours	head, wrist, torso, lower leg	[6]
food intake and desk work, 4 subjects, 3 hours	upper arm, wrist	[3]
bicycle repair 4 subjects, 3 hours	upper arm, wrist	[10]
opportunity project data set, 4 subjects, 5 hours	upper arm, wrist, back.	[9]
gym exercises, 2 subjects, 3 hours	upper arm, wrist, lower leg	[7]

Table 1. The data sets used for the signal level evaluation.

locomotion exercises			
8 classes: walk, run, run uphill, bike, rowing, stairs, ski, crosstrain			
Modality	Same	Trained on 1	Trained on 2
Acceleration	100 %	63%	65%
Gyro	80 %	72%	75%
Magnetic	79%	68%	71%
Gyro + Acceleration	-	78%	90%
Magnetic + Acceleration	-	72%	79%
gym arm exercises			
8 classes: lat, pectorial, shoulder press, upper back, arm extension/curl, pull down, chest press			
Modality	Same	Trained on 1	Trained on 2
Acceleration	97 %	24%	31%
Gyro	79 %	55%	61%
Magnetic	66%	41%	53%
Gyro + Acceleration	-	74%	82%
Magnetic + Acceleration	-	57%	69%

Table 2. The effectiveness of gyroscopes and magnetic sensors in compensating sensor shift with the method [7].

The sets include a variety of different body positions such as upper leg, arm, wrist, torso, head, back. The sampling rates used are either 50 or 100 Hz. An overview of the sets is given in Table 1. For details see the cited publications.

Evaluation. For the evaluation we first band-pass filter (1-25 Hz) the magnetic field signal. Then, we calculate the angular velocity based on the formula given in Equation 1. As error measure we use the mean of the absolute percentage difference between the true gyro signal and our estimation. It is calculated over a 1.5 second sliding window. We smooth the error to account for differences arising from temporal jitter between the two measurements.

Results. The results are summarized in Figure 2. The vast majority of data points have error below 20%, and there are only few with an error of more than 40%. The mean, median and standard deviation of the error is given by body location of the sensor in Table 3. It can be shown that statistically significant differences in the error distribution exist in three distinct groups of locations: (1) the back, torso and head, (2) upper armand lower leg, and (3) the wrist (p-values between 0.02 and 0.04). These differences can be explained by the way the respective body parts are moved and the probability of coming near metallic objects (or other sources of disturbance). The head, torso and back all move slowly, and are less likely to come very close to objects and devices. In contrast, errors can arise through the wrist's frequent and rapid displacement through motion within curved magnetic

Placement	Mean Error	Median Error	Standard Deviation
Head	17%	11%	90%
Torso	20.1%	12.5%	110%
Back	23.5%	19.5%	150%
Wrist	53.2%	44.2%	214%
Lower leg	34%	23.4%	162%
Upper arm	32.0%	25%	173%

Table 3. Mean, median error and standard deviation for different placements.

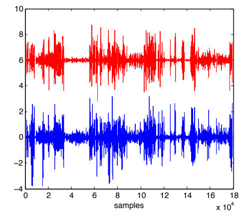

Figure 1. Example traces of signal level estimation of angular velocity using the magnetic field sensor.

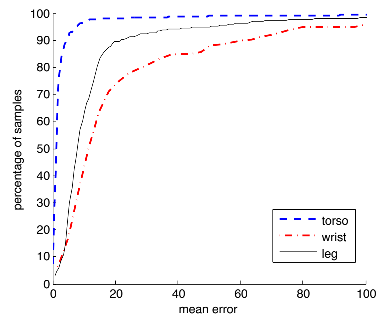

Figure 2. Cumulative distribution of errors. The x-axis gives the percent error between true and estimated angular velocity, and the y-axis gives the percentage of measurements with that error or less.

field lines, as described in the previous section. Errors from sensors on the arm and the leg lie in between.

There are also variances between data sets, however they are not statistically significant.

4. Implications on context recognition

Clearly, deviations in the range of 20%-40% are unacceptable for precise tracking. However, for activity recognition and applications that perform more qualitative analysis (e.g. recognizing general motion types) the accuracy requirements are less strict. To investigate to how well a magnetic field sensor can replace a gyroscope in such systems, we consider the gym exercise data set from [7]. We choose this data set because it was originally used to demonstrate how adding a gyroscope to an acceleration sensor can help compensate for sensor displacements[3] . This is an important application

3. The type of displacement considered was within body part shift. Thus, for example a sensor originally placed close to the wrist may be shifted up towards the elbow. Another common example would be a mobile phone in a trousers pocket which can be anywhere from deep down on one side to just barely in the pocket on the other side. In general, for non negligible shifts, a motion sensor based system that was trained on one location will perform poorly if tested on the other (see results below). In [7] we had shown how a adding a gyroscope to an acceleration sensor based system can make recognition more robust against such displacements. The general idea is based on the observations that for motions dominated by translations, the acceleration signal is actually insensitive to within body part shifts. It is only during rotations that shifted accelerometers produce different signals. The gyroscope signal, on the other hand, is insensitive to such shifts. In short, we use the ratio of acceleration to rotation to determine if a motion is translation or rotation dominated, and dynamically select either acceleration or gyroscope features.

locomotion exercises			
Modality	Trained only on Gyro	Trained on Gyro (normalized)	Individually Trained
Gyroscope	80%	78%	80%
Magnetic	42%	71%	79%
gym arm exercises			
Modality	Trained only on Gyro	Trained on Gyro (normalized)	Individually Trained
Gyroscope	79%	68%	79%
Magnetic	15%	21%	66%

Table 4. Classification for the gym leg and arm exercises.

that is particularly relevant to mobile phones, which are often loosely placed in pockets where they can shift around.

The Data Set. The data set consists of two groups of exercises: leg related and arm related. For both groups we had two test subjects each executing each locomotion activity for around 10 min. and each arm exercise for 20-25 repetitions. Both locomotion and gym arm exercises have 8 distinct activities (Table 4). For the leg exercises the subjects upper leg is equipped with 6 MTx Sensors: three mounted on the front and two on the back. For the arm exercises, there are four sensors placed on the forearm. We use the features and classifiers described in [7].

Results. The results are summarized in Tables 2 and 4. In Table 4 the recognition rates are compared for a gyro and a magnetic sensor based system. We differentiate three cases. First, the system is trained on features derived the raw gyro signals and tested with the magnetic field sensor derived angular velocities. This leads to a significant performance drop (from 80% to 41% for leg and 68% to 21% for arm exercises). Second we consider a system trained on normalized gyro signals. Normalize in this context means to map the gyro and magnetic field sensor values are mapped in a value space between -1 and 1. It can be seen that the gyro has system nearly the same accuracy (78%). At the same time, the magnetic sensor based recognition dramatically improves for the leg (to 71%) and somewhat improves for the arm. Finally it can be seen that if we train the system directly with the magnetic field derived values, the performance is nearly identical (79%) to the gyro case for the leg and close (66% vs. 79%) for the arm. This confirms the suspicion that the method works much better for the legs than for the arms, although when directly training on the magnetic field the arms results are still reasonable.

Table shows the results of using the magnetic field sensor to compensate for sensor shifts. The results for Gyro and accelerometer are taken from [7]. For the magnetic field we have used the same algorithms, substituting the gyroscope signal for angular velocity values computed from the magnetic sensor using our method. In most cases, the magnetic sensor performs about 10% to 15% worse than the gyro. Nonetheless, when it comes to compensating for sensor shifts, it is still useful and achieves improvements of up to over 100% (from 24% to 57%).

5. Conclusion

The results presented in this paper indicate that in many cases angular velocity information derived from magnetic field sensors can approximate and replace gyroscope signals. Two observations are particularly noteworthy. First, the magnetic field sensor can be used to compensate sensor shifts using the same method that has been developed for gyroscopes. This is particularly interesting for mobile phone based applications where the device is loosely carried in a pocket where it can shift around. Second, in some recognition tasks we can feed data produced by the magnetic sensor into a system that has been trained on gyroscope data. This is relevant e.g. for increasing system robustness against changing sensor configurations and sensor failures.

While our approximation method does not require the local magnetic field to reliably point north, fast linear motions in close proximity to strong magnetic disturbances can lead to significant errors. As a consequence the method works much better on the torso then on arms and wrists.

We are currently investigating whether combining an accelerometer and magnetometer can improve the angular velocity estimations. In particular, the estimated gravity vector is subject to different errors (translation instead of magnetic interference), and generally points in a different direction than magnetic north, thereby allowing estimation of angular velocity around the magnetic north vector.

References

[1] A. Benbasat and J. Paradiso. An inertial measurement framework for gesture recognition and applications. *Gesture and Sign Language in Human- ...*, Jan 2002.

[2] U. Blanke and B. Schiele. Sensing location in the pocket. *Ubicomp Poser Session*, page 2, Aug 2008.

[3] J. Cheng, O. Amft, and P. Lukowicz. Active capacitive sensing: Exploring a new wearable sensing modality for activity recognition. *Pervasive Computing*, Jan 2010.

[4] K. Förster, D. Roggen, and G. Tröster. Unsupervised classifier self-calibration through repeated context occurences: is there robustness against sensor displacement to gain? *2009 International Symposium on ...*, Jan 2009.

[5] H. Junker, O. Amft, P. Lukowicz, and G. Tröster. Gesture spotting with body-worn inertial sensors to detect user activities. *Pattern Recognition*, Jan 2008.

[6] K. Kunze and P. Lukowicz. Using acceleration signatures from everyday activities for on-body device location. *Wearable Computers, 2007 11th IEEE International Symposium on*, pages 115 – 116, Sep 2007.

[7] K. Kunze and P. Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing*, Sep 2008.

[8] J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. *Proceedings of Pervasive*, Jan 2006.

[9] P. Lukowicz, G. Pirkl, D. Bannach, and F. Wagner. Recording a complex, multi modal activity data set for context recogntion. *Workshop on Context ...*, Jan 2010.

[10] T. Stiefmeier, G. Ogris, H. Junker, P. Lukowicz, and G. Troster. Combining motion sensors and ultrasonic hands tracking for continuous activity recognition in a maintenance scenario. *Wearable Computers, IEEE International Symposium*, 0:97–104, 2006.