On the Use of Magnetic Field Disturbances as Features for Activity Recognition with on Body Sensors

Gernot Bahle, Kai Kunze, and Paul Lukowicz

Embedded Systems Lab
University Passau, Passau, Germany
first.lastname@uni-passau.de

Abstract. We investigate the use of magnetic field disturbances as features for motion based, wearable activity recognition systems. Such disturbances are mostly caused by large metallic objects and electrical appliances, both of which are often involved in human activities. We propose to detect them by subtracting angular velocity values computed from the changes in the magnetic field vector from gyroscope signals. We argue that for activities that are associated with specific objects or devices such features increase system robustness against motion variations, sensor displacement and inter user differences. On a previously published data set of 8 gym exercises we demonstrate that our approach can improve the recognition by up to 31% over gyroscope only and up to 17% over a combination of a gyroscope and 3D accelerometer. Improvements of 9.5% are also demonstrated for user independent training as well as for the case of displaced sensors.

A particularly interesting result is the fact that adding the magnetic disturbance features significantly improves recognition based on the vector norm of accelerometers and gyroscopes. The norm is often used when the orientation of the sensor is not known. This is common when using a mobile phone or other consumer appliance as a sensor.

1 Introduction

Gyroscopes and magnetic field sensors are often used modalities in activity recognition ([7,4,17,15]). In general, magnetic sensors are used to determine orientation (possibly together with an accelerometer) while the gyroscope helps filter out distortions due to disturbances in the magnetic field.

In this paper, we investigate a different way of using the information from a magnetic sensor. Rather than try to filter out the disturbances we specifically use them as features. This approach is aimed at activities that are related to motions in the proximity of fixed metallic objects with each activity being associated with a different object. The example considered in this paper are exercises on gym machines. Such machines are mostly metallic and differ significantly in the shape of metallic parts which users hands come close to.

P. Lukowitz, G. Kortuem, and K. Kunze (Eds.): EuroSSC 2010, LNCS 6446, pp. 71–81, 2010.
© Springer-Verlag Berlin Heidelberg 2010

Metallic objects disturb the earth magnetic field in a way determined by their shape and composition. This means that while the absolute value of the magnetic field signal is predominantly given by position and the specific motion that the user has been performing, the disturbances are more dependent on the object itself (although position and motion also play a role). As will be shown in the paper this reduces the dependence of the recognition performance on sensor placement and motion variations that occur when looking at different users.

We demonstrate this on a data set recorded from gym exercises on different machines that we have previously used to study the effect of sensor placement variations on activity recognition [9].

1.1 Related Work

In general, gyroscopes are used together with magnetic field sensors to filter out magnetic disturbances. Levi et. al. present a patent describing how to improve the magnetic heading information using a gyroscope [12]. Ladetto et. al improve a dead-reckoning system fusing the gyro and compass information [10]. Young et al. use the magnetic field as source of orientation data in motion capturing applications ([16]). The application scenarios for magnetic field sensors range from life recording, over music performances to novel interface designs ([7,3,14,5]). Still the use is as described above.

Merging compass and rotation data for activity recognition is also quite common([11,6,2,1]). Lester et. al. use a combination of accelerometer, sound, gyro and magnetic field sensors embedded in a mobile phone size device [11]. Junker et. al. describe how to combine inertial sensors to spot sporadic user gestures in a data stream [6].

Complementary to the work presented in this paper, Minnen et. al. describe a way to detect unsupervised "Motifs" in sensor data [13]. This might be used to detect "interesting" disturbances in the magnetic field automatically without training.

In previous work, we have shown how to estimate the angular velocity form magnetic field sensor data only [8].

While all of the above research deals with magnetic sensing in activity recognition, our work is fundamentally different as it aims to use magnetic disturbances rather then the pure magnetic signal caused by user motion. As such it is complementary to previous work. Features derived using our method can thus be used and can be applied in situations where magnetic field signals were previously of little use.

2 Approach

To isolate the magnetic disturbances we compare the angular velocity given by the gyroscopes with the value computed from the magnetic field data. To this end we use the approximation of angular velocity from 3D magnetic sensor data described in [8]. It is based on the fact that the three field strength components that a 3D magnetic field sensor outputs represent a vector $\boldsymbol{B}$ 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 $\boldsymbol{B(t)} = (b, 0, 0)$ with $b = \|\boldsymbol{B(t)}\|$ 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:[1] $B_i(t) = \|\boldsymbol{B(t)}\| \cdot cos(\varphi_i)$ The angle $\varphi_i(\boldsymbol{B}(t))$ between the i-th axis and the magnetic field strength vector $\boldsymbol{B}(t)$ measured at time t is then given by $\varphi_i(\boldsymbol{B}(t)) = \arccos \frac{B_i(t)}{\|\boldsymbol{B}(t)\|}$ where $B_i(t)$ is the i-th component of $\boldsymbol{B}(t)$. Angular velocity then equals the first derivative of the angle.

2.1 Analyzing the Distortion

Obviously, the presence of metallic objects and artificial magnetic fields (e.g. caused by electric appliances) has significant influence on the quality of the approximation. Figure 1 shows an example comparing the magnitude of gyro data and the magnitude of estimated angular velocity based on magnetic data. The third graph in 1 highlights the principle we exploit in this work: when angular velocity measured by the gyro and angular velocity estimated by use of the magnetic sensor are equal, their difference is zero. When they diverge, however, the difference takes on non zero values, signifying a disturbance in the magnetic field. In previous work we have tried to filter out this influence and focus on applications where it is negligible. In this paper we do the opposite: we explicitly try to analyze the distortions and use them as a source of information. We focus on scenarios where the user wears a combination of a magnetic field sensor and a gyroscope while performing activities in the proximity of metallic objects and other sources of magnetic disturbance.

As described in [8], the distortion depends on many factors and is not easily described in exact terms. However, in most cases the nature of the source of magnetic distortion is a more significant influence then fine variation in the motions associated with the activity. Thus, in activity recognition tasks where the classes are associated with different objects or devices (=different sources of magnetic field disturbance), the use of magnetic disturbance as a feature can increase the robustness of the system in comparison to purely motion analysis based approaches. This is because the sources of disturbance often remain largely unchanged, whereas the specific motion signals associated with the activities can significantly vary. Such variations are particularly relevant when looking at user independent recognition, but can also be non negligible for a single user. In

[1] 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.

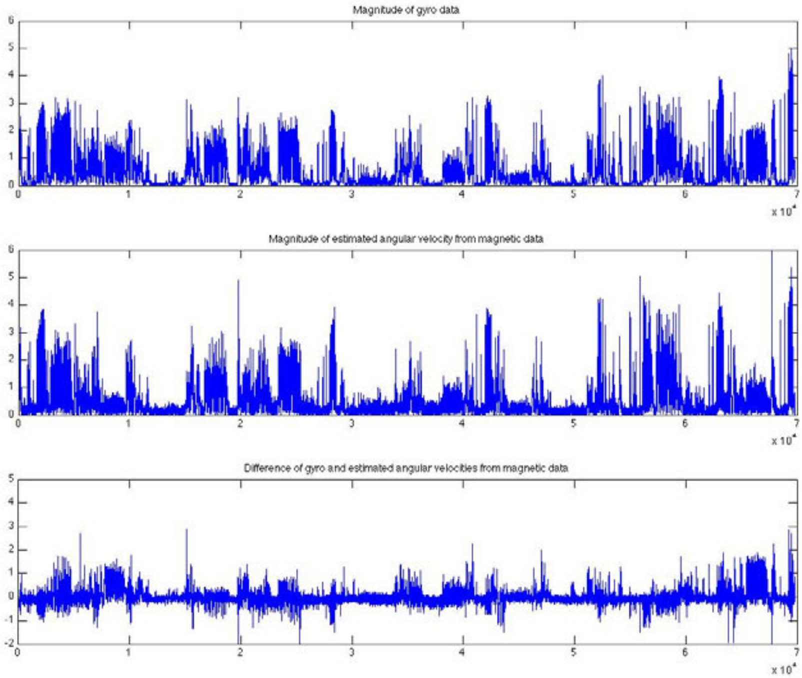

Fig. 1. Magnitudes of (estimated) angular velocities from gyro and magnetic data and difference thereof

addition, as was shown in [9], recognition rates drop dramatically if training data is recorded from a different sensor location than testing data (e.g. because of the sensor being loosely attached and slipping). In all of the above cases using features gained from difference of estimated and real angular velocity, recognition rates can be improved.

2.2 Problems and Limits

The main problem we encountered was that the approximation of angular velocities using magnetic data on individual axes was not good enough to yield useful information when subtracted from the three individual gyroscope signals. As [8] states, the estimate depends on the basis vectors chosen for the projections of the magnetic field strength vector. Calculating the magnitude over all three estimated axes skirts this difficulty, as the magnitude is not dependent on the exact distribution of the velocities between the three axes. Unfortunately, it also destroys potentially useful direction specific information. At the moment, this boils down to using the total amount of distortion as a feature instead of the distortion on each axis. In the future, we will try to further refine the approximation technique to allow for individual axis estimates good enough to use in the way described in this paper.

3 Evaluation and Results

3.1 DataSet

We evaluated our approach on part of the gym data set previously published in [9]. This data set consists of 8 gym exercises (lat machine, butterfly, shoulder press, upper back, arm extension, arm curl, pull down, chest press). The exercises are performed by two subjects with 10 to 15 repetitions on each machine (total of around 200 instances), taking about 5 minutes for each set.

The data set is recorded with XSENSE MTX inertial sensor modules containing both a gyroscope and a magnetic field sensor, as well as an accelerometer. For both test subjects, five modules have been used to record the entire training session; we exploit this to demonstrate our approach for displaced sensors (using data of one sensor for training and then testing on the other 4).

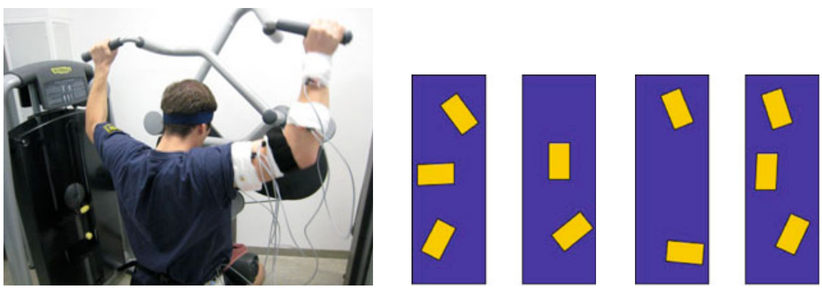

Fig. 2. Sensor placement for gym dataset for subject 1 and 2; blue rectangles represent the rigid bodies (upper / lower arm), yellow boxes the placement of the sensors

The sensors were split into two groups of three and two sensors respectively. Figure 2 shows the positions of the modules).

It is noteworthy that during the recording of subject two, one of the sensors attached to the upper arm shifted position, resulting in a new position between both upper and lower arm sensor groups.

3.2 Evaluation

We calculated features in a sliding window one second long with 0.5 seconds overlap. The features used were *mean, var, amplitude, rms, centre of mass and fluctuation of spectrum and the first five cepstrum coefficients as defined by ift(log(abs(fft(signal))))*.

All features were calculated on the magnitude of acceleration, gyroscope and magnetic 3D data vectors. As stated above, the magnetic field based approximation of angular velocity is not yet sophisticated enough to work on individual axes. Furthermore, using the individual axes for classification performs significantly better only as long as there is no sensor displacement (for reference, we did include the 3D data in the case of non displaced sensors). Once sensor displacement happens, the opposite is actually the case. As translations and rotations change the distribution of force to individual axes far more than the overall amount of force measured by the sensor, using features based on magnitudes of sensor signals is far more robust under displacement. E.g., using gyro and acceleration data, recognition rates are as good as 96.8% using 3 axes but drop down to 24% when the sensor is displaced. On the other hand, resorting to the magnitude features only yields 69.2% without displacement, but only drops to 56.5% in a different sensor location.

A note on computational complexity: All the calculations required for our approach are quite simple; both the estimation of angular velocity from magnetic data and the features used do not require extensive computational power; while we used an already available data set for this work and performed our analysis offline on a workstation running Matlab, it is computationally feasible to implement it in an online setting, e.g. on a mobile device.

For the classification itself, we tested different frame-by-frame classifiers (C4.5 decision tree, kNN, NaiveBayes). Results from these classifiers were within 1% of one another, so without loss of generality, the numbers we present are derived from the NaiveBayes classifier. Training data consisted of half a set (obtained by a 50% percentage split) of full data for non displaced sensors and of a full set of any one sensor in the case of displaced sensors. Testing data likewise was either the other half in the single sensor case or the full data of the remaining four sensors in the displaced case. In all evaluations safe the last one, both training and testing was done on a per subject basis. The results were then averaged. While user independence is a desirable quality in activity recognition, it is also difficult to achieve; in the gym setting, e.g., different users perform motions slightly differently. We address this topic in the last part of the evaluation.

In order to determine whether the difference of gyro signal and magnetic estimated angular velocities contained useful information, in a first step we compared two sets of features. The first set contained only features derived from the individual gyro and magnetic data vectors; the second set was expanded by adding features calculated from the difference of gyro and estimated magnetic angular velocities.

As outlined above, it is not a good idea to use raw magnetic data for activity recognition. In our data set, rotating the sensors (or, just as bad) the gym machines would result in completely different magnetic sensor readings. We therefore performed two more evaluation steps.

In a second evaluation, we compared two sets of data; one set only containing features derived from the gyro readings, with the second set once again including the gyro and magnetic velocity difference features.

In a third step, we added the acceleration data, comparing one set of features calculated on magnitudes of acceleration and gyro data to a set containing those and additionally the information derived from differences in gyro and magnetic angular velocity.

As a final evaluation step, we have tested how well the additional features derived from the angular velocity comparison perform when training data and test data do not originate from the same person. To that end, we used sensor data from subject 1 for training and data from subject 2 for testing; afterwards, we switched persons and repeated the process, averaging over the results.

3.3 Results

Carrying out the evaluation strategy outlined above, we achieved the results shown in Figure 3. The details of each evaluation step are described in the corresponding subsection below.

Comparison: gyro + magnetic data to same + difference in angular velocities. In the case of no sensor displacement, classifying the gym data set with features calculated from gyroscope and magnetic sensor data yields average recognition rates of 90.0% using individual axis data and 74.6% using magnitude of axes data. Extending the feature set by features derived from the difference of gyro angular velocities and estimated magnetic angular velocities can boost these rates to 93.0% in the 3D case and 83.9% in the magnitude case; the gains of 3.0% and 9.3% respectively prove that the features stemming from angular velocity comparison of gyro and magnetic data do contain additional information

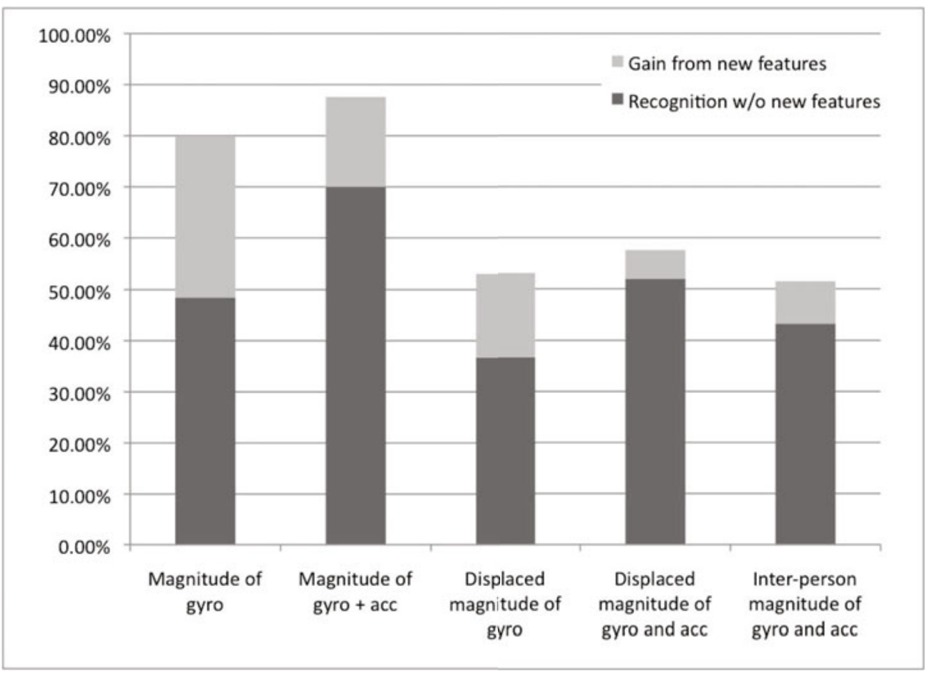

Fig. 3. Overview of base recognition rates and gains for various cases

that can be used to further improve recognition rates. The lower gain in the 3D case is partly due to the much higher base recognition rates, which limit the room for improvement.

In the case of displaced sensors, adding the difference features boosts average recognition rates from 52.4% to 59.7%, resulting in an average gain of 7.3%. Individual gains between each pair of sensors were uniformly distributed even between the different sensor groups of upper and lower arm; this implies that - at the very least - the additional features are less sensor location dependent than e.g. raw gyro or magnetic data. Looking at the individual axis data under displacement underscores the arguments for using the magnitude of sensor data made above, with recognition rates dropping to about 30%. There is, however, still an average gain of 3.7% when adding the additional features.

In conclusion, the results of the first evaluation step show that both in cases with and without sensor displacement, using the magnetic/gyro difference features improves recognition rates. Table 1 summarizes these results.

Table 1. Comparison of recognition rates w/o and with angular velocity comparison features for gyroscope and magnetic data

Modality	Recognition Rate w/o ang. vel. comp. features	Recognition Rate w. ang. vel. comp. features	Gain
3D, undisplaced	90.0%	93.0%	3.0%
magnitude, undisplaced	74.6%	83.9%	9.3%
magnitude, displaced	52.5%	59.7%	7.3%

Comparison: gyro data to gyro + difference in angular velocities. Classifying the gym dataset with features derived solely from gyroscope sensor data yields an average recognition rate of 80.1% in the 3D case and 48.5% in the magnitude case as long as the sensor has not been displaced. In both cases, adding the features resulting from our difference of angular velocities approach net improvements: average recognition rates rise to 90.4% in the 3D case, a gain of 10.3%. Adding the additional features to features calculated from the magnitude gyro signal vector, recognition can be boosted to 79.9%, showing a tremendous gain of 31.4%. This suggests that adding the new features manages to recoup at least part of the information lost by switching from 3D data to magnitude of data. In the case of displaced sensors, recognition rates using only gyroscope data drop by 11.8% to 36.7% using the features calculated on the magnitude of the signal. Once again, adding the information resulting form our comparison of gyro angular velocity and the estimated magnetic one yields significant improvements, raising average recognition rates amongst the four different groups of displaced sensors to 53.1%, a gain of 16.4%.

Table 2 summarizes these results.

Comparison: gyro + acceleration data to same + difference in angular velocities. In the third section of the evaluation, we combined gyroscope and acceleration data for the feature set; these modalities provide a realistic

Table 2. Comparison of recognition rates w/o and with angular velocity comparison features for gyro data

Modality	Recognition Rate w/o ang. vel. comp. features	Recognition Rate w. ang. vel. comp. features	Gain
3D, not displaced	80.1%	90.4%	10.3%
magnitude, not displaced	48.5%	79.9%	31.4%
magnitude, displaced	36.7%	53.1%	16.4%

scenario for activity recognition. In the case of a single, non displaced sensor, recognition rates are very good, especially when using the 3D signal for feature calculation. Averaged over the 5 different sensors and both test persons yields a recognition of 95.7%. This is very close to perfect already and leaves little room for improvement. Even so, however, adding the features derived from the angular velocity comparison raises accuracy slightly to 96.8%, a gain of 1.1%. Using the magnitude of acceleration and gyroscope as basis for feature calculation expectedly drops recognition rates to 70.0% when there is no sensor displacement. Adding the information derived from our approach manages to increase accuracy significantly again to 87.6%, a gain of 17.6%.

Finally, in the case of displaced sensors, recognition rates drop down to an average of 52.0% over all 4 groups of sensors, using just acceleration and gyroscope data. Adding in the gyroscope - magnetic comparison features manages to improve accuracy by 5.6% to 57.6%.

Overall, as in the other cases presented in our evaluation, there is measurable information contained in the features derived from comparing actual and approximated angular velocity. Table 3 summarizes these results.

Table 3. Comparison of recognition rates w/o and with angular velocity comparison features for combined gyro and acceleration data

Modality	Recognition Rate w/o ang. vel. comp. features	Recognition Rate w. ang. vel. comp. features	Gain
3D, not displaced	95.7%	96.8%	1.1%
magnitude, not displaced	70.0%	87.6%	17.6%
magnitude, displaced	52.0%	57.6%	5.6%

Comparison: gyro + acceleration data with and without angular velocity difference between subjects one and two. In activity recognition, an added difficulty may present itself: when classifiers trained on data obtained from one person are used to classify activities performed by a different person, recognition accuracy usually drops, sometimes severely. Thus, while user independence is of course a desirable quality of any recognition system, it is usually hard to achieve. E.g., different users may perform activities slightly differently and sensor placement may diverge slightly. In theory, magnetic distortions should be quite person independent, as they depend mostly on the environment. To test this hypothesis, we trained with acceleration and gyroscope data from subject one and tested on data from subject two and vice versa. We obtained these results: Training with upper arm data without our magnetic distortion features resulted in a recognition rate of 40.05%, while lower arm data performed slightly better at 46.51%. On average, recognition accuracy was 43.28%. Adding the features calculated from angular velocity comparison resulted in noticeable gains; upper arm accuracy increased to 46.86%, a gain of 6.81%; lower arm recognition was boosted to 56.08%, an increase of 9.57%; finally, on average, recognition rose to 51.47%, an improvement of 8.19%. Table 4 summarizes these results.

Table 4. Comparison of recognition rates w/o and with angular velocity comparison features between different subjects

Modality	Recognition Rate w/o ang. vel. comp. features	Recognition Rate w. ang. vel. comp. features	Gain
upper arm	40.05%	46.86%	6.81%
lower arm	46.51%	56.08%	9.57%
both	43.28%	51.47%	8.19%

Since sensor placement on subject one was different from sensor placement on subject two, it is not possible to divide total gains from the additional features into gains from displacement resistance and gains from person independence. Overall, while recognition expectedly performs worse than in the case of same subject training and testing (51.47% compared to 57.6%), the gain from the magnetic distortion features is actually 46% higher (8.19% compared to 5.6%). This indicates that our initial supposition of the added value for achieving a measure of subject independence is correct.

4 Conclusion

The results presented in the previous section are a strong indication that magnetic field disturbances computed from the difference between gyro signals and angular velocity derived from magnetic field vector changes are useful for wearable activity recognition. The inclusion of the corresponding features has lead to an improvement of performance in all experiments. This included cases of shifted sensors and person independent recognition.

Despite the promising results, our evaluation, however, should only be considered a starting point for further work. The technique presented here should be tested on a wider range of datasets. It would also be interesting to apply it to data gathered from mobile devices like the iPhone 4. Another aspect deserving of more research is further refining the estimation approach itself in order to be able to apply it to individual axes with sufficient precision. It would then be possible to characterize distortions along three individual dimensions instead of just their magnitude, possibly making the technique even more valuable. In that context, it might also be worth exploring how much sensor displacement is necessary before the gains from the greater invariance to displacement outweigh the accuracy penalties in case of little displacement when switching from 3D signal data to magnitude of signal data.

References

1. Amft, O., Bannach, D., Pirkl, G., Kreil, M., Lukowicz, P.: Towards wearable sensing-based assessment of fluid intake. In: 8th IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), pp. 298–303 (2010)
2. Benbasat, A., Paradiso, J.: An inertial measurement framework for gesture recognition and applications. Gesture and Sign Language in Human- (January 2002)
3. Fels, S., Lyons, M.: Creating new interfaces for musical expression: introduction to nime. In: SIGGRAPH 2009: SIGGRAPH 2009 Courses (August 2009)
4. Förster, K., Roggen, D., Tröster, G.: Unsupervised classifier self-calibration through repeated context occurences: is there robustness against sensor displacement to gain? In: ISWC 2009 (2009)
5. Hook, J., Taylor, S., Butler, A., Villar, N., Izadi, S.: A reconfigurable ferromagnetic input device. In: Proceedings of the 22nd Annual ACM Symposium on User Interface Software and Technology, UIST 2009 (October 2009)
6. Junker, H., Amft, O., Lukowicz, P., Tröster, G.: Gesture spotting with body-worn inertial sensors to detect user activities. Pattern Recognition (January 2008)
7. Kern, N., Schiele, B., Schmidt, A.: Recognizing context for annotating a live life recording. Personal and Ubiquitous Computing (January 2007)
8. Kunze, K., Bahle, G., Lukowicz, P., Partridge, K.: Can magnetic field sensors replace gyroscopes in wearable sensing applications? To be published in ISWC (October 2010)
9. Kunze, K., Lukowicz, P.: Dealing with sensor displacement in motion-based onbody activity recognition systems. In: Proceedings of the 10th International Conference on Ubiquitous Computing, UbiComp 2008 (September 2008)
10. Ladetto, Q., Gabaglio, V.: Combining gyroscopes, magnetic compass and gps for pedestrian navigation. Systems in Geodesy (January 2001)
11. Lester, J., Choudhury, T., Borriello, G.: A practical approach to recognizing physical activities. In: Proceedings of Pervasive (January 2006)
12. Levi, R., Marshall, R.: Gyro aided magnetic compass. US Patent App. 10/618,549 (January 2003)
13. Minnen, D., Starner, T., Essa, I., Isbell, C.: Improving activity discovery with automatic neighborhood estimation. In: International Joint Conference on . . . (January 2007)
14. Raya, R., Ceres, R., Roa, J., Rocon, E.: Assessment of the involuntary motion of children with motor impairments to improve the accessibility of an inertial interface. In: Proceedings of the 9th International Conference on Interaction Design and Children, IDC 2010 (June 2010)
15. Stiefmeier, T., Roggen, D., Troster, G., Ogris, G.: Wearable activity tracking in car manufacturing. IEEE Pervasive . . . (January 2008)
16. Young, A., Ling, M., Arvind, D.: Distributed estimation of linear acceleration for improved accuracy in wireless inertial motion capture. In: Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN 2010 (April 2010)
17. Zappi, P., Stiefmeier, T., Farella, E., Roggen, D.: Activity recognition from onbody sensors by classifier fusion: sensor scalability and robustness. In: Proceedings of . . . (January 2007)