

The Impact of Personalization on Smartphone-Based Activity Recognition

Gary M. Weiss and Jeffrey W. Lockhart

Department of Computer and Information Science, Fordham University
441 East Fordham Road, Bronx NY, 10458
{gweiss, lockhart}@cis.fordham.edu

Abstract

Smartphones incorporate many diverse and powerful sensors, which creates exciting new opportunities for data mining and human-computer interaction. In this paper we show how standard classification algorithms can use labeled smartphone-based accelerometer data to identify the physical activity a user is performing. Our main focus is on evaluating the relative performance of impersonal and personal activity recognition models. Our impersonal (i.e., universal) models are built using training data from a panel of users and are then applied to new users, while our personal models are built with data from each user and then applied only to new data from that user. Our results indicate that the personal models perform dramatically better than the impersonal models—even when trained from only a few minutes worth of data. These personal models typically even outperform hybrid models that utilize both personal and impersonal data. These results strongly argue for the construction of personal models whenever possible. Our research means that we can unobtrusively gain useful knowledge about the habits of potentially millions of users. It also means that we can facilitate human computer interaction by enabling the smartphone to consider context and this can lead to new and more effective applications.

Introduction

Smartphones and other mobile devices now contain diverse and powerful sensors. These sensors include GPS sensors, audio sensors (microphones), image sensors (cameras), light sensors, direction sensors (compasses), proximity sensors, and acceleration sensors (accelerometers). Because of the small size of these “smart” mobile devices, their substantial computing power, their ability to send and receive data, and their nearly ubiquitous use in our society, these devices open up exciting new areas for research in data mining and human-computer interaction. The goal of our WISDM (Wireless Sensor Data Mining) project (Weiss 2012a) is to explore the research and application issues associated with mining sensor data from these powerful mobile devices. In this paper we explore the use of the smartphone accelerometer sensor to identify the activity a user is performing—a task known as activity recognition.

We employ a supervised learning approach for addressing the activity recognition task. We collect accelerometer data from 59 users as they walk, jog, climb stairs, sit, stand, and lie down, and then aggregate this raw time series data into examples that cover 10 seconds of activity. Each example is labeled with the activity that occurred during the interval and activity recognition models are induced using several standard classification algorithms. We utilize Android-based smartphones because the Android operating system is free, open-source, easy to program, and is the most popular mobile operating system. But the tri-axial accelerometers present in Android cell phones are found in virtually all new smartphones and smart music players, including the iPhone and iPod Touch (Apple 2009), and thus the research we describe in this paper could easily be applied to other mobile platforms.

Accelerometers were initially included in smartphones to support advanced game play and to enable automatic screen rotation. However, they can also support applications that exploit activity recognition, such as health applications that monitor a user’s daily activities. Such applications can help address the health concerns that arise due to inactivity (e.g., cardiovascular disease, hypertension, and osteoporosis) and can help combat the critical public health threat of childhood obesity (Koplan et al. 2005). Given that the World Health Organization (2002) maintains that approximately two million deaths per year can be attributed to physical inactivity, such applications are sorely needed. The WISDM project is pursuing such a health application by developing the Actitracker application (Weiss 2012b).

Activity recognition can also enable a smartphone to tailor its behavior based on what the user is doing; such context-sensitive applications can automatically forward calls to voicemail when a user is exercising or play “upbeat” music when a user begins to slow during a daily jog. Applications like these make human-computer interaction transparent by responding to a user’s natural activities rather than requiring conscious interaction with an interface.

Accelerometer-based activity recognition is not new. In fact, numerous activity recognition systems have been developed, but virtually all of these rely on the use of multiple accelerometers strapped to the subject’s extremities. Our work differs from this in that we rely on a commercial mass-marketed device rather than specialized hardware,

and we use a single device, conveniently located in the user's pocket, rather than multiple devices distributed across the body. Thus our activity recognition work can readily be deployed on a wide scale. But in this paper we make another key contribution, by comparing and analyzing the performance of personal and impersonal activity recognition models. Personal models are generated for specific users by having a sample of labeled activity data for the user, while universal/impersonal models are built from a collection of users and then applied to new users. As our results will show, personal models dramatically outperform impersonal models—even though they are built using vastly less data. No existing smartphone-based system performs this type of comparison in a comprehensive manner. The accuracy and portability of activity recognition models are critical concerns for activity representation systems that are built using these models.

The Activity Recognition Task

The activity recognition task involves mapping time-series accelerometer data to a single physical user activity. In this paper we formulate this into a standard classification problem by aggregating the time-series data into examples. We consider six common activities that collectively cover much of the time in a typical user's day: walking, jogging, stair climbing (up and down), sitting, standing and lying down. We assume the smartphone is in the user's pocket, but in the future will consider belt-worn smartphones. The axes associated with the smartphone are aligned as indicated in Fig. 1. The accelerometer measures the acceleration due to gravity, about 9.8m/s^2 , and this is incorporated into the y-axis values for activities where the user (or at least the phone) is upright.

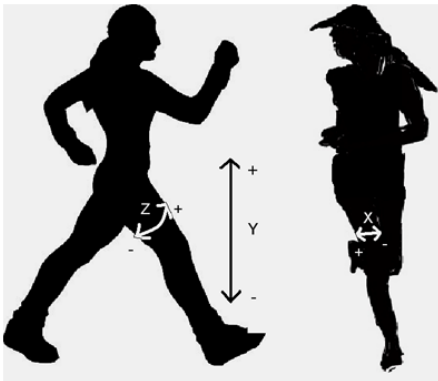


Figure 1. Modern smartphone accelerometers are tri-axial. The figure shows the axes of motion relative to the user, assuming the phone is in the user's pocket.

The six activities that we recognize include three static activities (standing, sitting, and lying down) and three dynamic activities (walking, jogging and stair climbing). Fig. 2 displays the accelerometer graph for the standing activity. Due to space limitations, the graphs for sitting and lying down are not shown, but for those two static activities

gravity no longer aligns with the y-axis and so the y-axis values no longer dominate.






Figure 2. The accelerometer graph for standing shows that gravity aligns with the y-axis.


Fig. 3 displays sample accelerometer data for walking, jogging, and stair climbing. All exhibit periodic behavior.



(a) Walking



(b) Jogging



(c) Stairs (Ascending)

Figure 3. Accelerometer graphs for three dynamic activities.

Experiment Methodology

In this section we describe the methodology for generating the activity recognition models. We discuss the data collection procedures, the method for transforming the accelerometer data into examples, and the model induction process. We also describe the methodology for generating and evaluating the personal, impersonal, and hybrid models.

Data Collection

We collected data by having 59 users carry an Android-based smartphone in their pocket while performing the six everyday activities mentioned previously. The data collection process is controlled by our sensor collection “app,” which is available from the Android Marketplace. Our research team members directed the participants to perform the various activities and input the activity labels into the app. The sensor data is stored on the phone and also transmitted to our server. For this study we sampled the accelerometer 20 times per second. Fifteen different Android smartphone models were used to collect the data, running Android OS versions from 1.5 to 2.3. The accelerometers on all of these phones were observed to operate similarly.

Data Transformation

Standard classification algorithms cannot directly handle time-series data, so we first transform the raw accelerometer data into examples (Weiss and Hirsh 1998). To accomplish this each example summarizes 10 seconds of data (this time is sufficient to capture several repetitions of periodic motions and was empirically shown to perform well). Given 20 samples per second and 3 axes, this yields 600 accelerometer values per example. We then use the following 6 basic features to generate 43 features from these raw values (the number of features generated for each feature-type is noted in brackets):

- Average[3]: Average acceleration (per axis)
- Standard Deviation[3]: Standard deviation (per axis)
- Average Absolute Difference[3]: Average absolute difference between the value of each of the 200 values in an example and their mean (per axis)
- Average Resultant Acceleration[1]: Average of the square roots of the sum of the values of each axis squared $\sqrt{(x_i^2 + y_i^2 + z_i^2)}$ over the example
- Binned Distribution[30]: The fraction of the 200 values that fall within each of 10 equally spaced bins, spanning the range of values (per axis).
- Frequency[3]: the frequency of the periodic wave associated with repetitive activities (per axis)

The frequency feature uses a heuristic method to identify all of the clearly distinct peaks in the wave and then calculates the average time between successive peaks. For sam-

ples where at least three peaks cannot be found, a special null value is used.

Table 1 shows the number and distribution of the transformed examples, per activity. Walking is the most common activity. The time spent jogging and stair climbing had to be limited because they are so strenuous, and we limited the time spent on the static activities because they were found to be easy to identify.

	Walk	Jog	Stair	Sit	Stand	Lie	Total
Total	3397	1948	1549	1143	689	565	9291
%	36.6	21.0	16.7	12.3	7.4	6.1	100

Table 1. The number of examples per activity over the entire data set, collected from 59 users.

Model Induction and Experiments

Our activity recognition models are induced from the labeled examples using the following WEKA (Witten and Frank 2005) classification algorithms: decision trees (J48 and Random Forest, RF), instance-based learning (IBk), neural networks (Multilayer Perceptron, NN), rule induction (J-Rip), Naïve Bayes (NB), Voting Feature Intervals (VFI), and Logistic Regression (LR). Default settings are used for all learning methods except NB, where kernel estimation is enabled, and IBk, where we set $k=3$ (IB3) so we use 3 nearest neighbors, rather than the default of 1.

In our research we induce three types of models: impersonal, personal, and hybrid. Each model addresses a slightly different learning problem and makes different assumptions about how the model will be applied. The type of model impacts how we partition the data into training and test data. The different models are described below:

- Impersonal Models: these models are generated using training data from a panel of users that will not subsequently use the model (thus the training and test sets will have no common users). These models can be applied to a new user without requiring additional labeled training data or model regeneration.
- Personal Models: these models are generated using labeled training data from *only* the user for whom the model is intended. These models require a training phase to collect labeled data from each user. The training and test data come from the same (and only) user, but will contain distinct examples.
- Hybrid Models: these models are a mixture of the impersonal and personal models. The training data will have data from the test subject and from other users. Thus the training and test set will have overlapping sets of users (but distinct examples).

Impersonal models have the advantage that they can be built once for all users and can include data from many users for training purposes. These models can be viewed as *universal* models, although technically they should only be

used for users not in the training set. Personal models have the advantage that they may match the idiosyncrasies of the intended user, but require each user to provide training data; they will also have to rely on limited data from a single user. The hybrid model, like the personal model, requires training data and model generation for each user, but can potentially outperform the personal model because it makes additional training data available (from other users).

The experiments associated with each model vary in how they are set up. For impersonal models data from 58 users is placed into the training set and data from 1 user is placed into the test set. This process is repeated 59 times, which allows us to generate reliable performance metrics and also easily characterize the performance on a per-user basis. For personal models, 10-fold cross validation is applied to each user’s data and thus 590 (59×10) personal models are evaluated. Since each user has a very limited amount of data (on average 160 examples), 10-fold cross validation is essential. The confusion matrices that are generated from both of these types of models are created by summing the counts in each cell over all 59 runs. The setup for the hybrid models is much simpler, since we just place all of the user data into a single file and then use 10-fold cross validation. Thus, in this case the training and test set have overlapping sets of users.

Results

The predictive accuracy associated with the personal, hybrid, and impersonal models, for each of the eight classification algorithms, is displayed in Table 2. Those activities that occur more frequently have a greater impact on performance. These results make it quite clear that for every classification algorithm the personal models perform best, the hybrid models perform second best, and the impersonal models perform worst. Furthermore, the personal models always achieve a very high level of accuracy and perform *dramatically* better than the impersonal models. While this result is easy to justify, since different people may move differently, the result is far from obvious, since the personal models are trained from dramatically less data—and, in an absolute sense, from a *small* amount of data (i.e., an average of $0.9 \times 160 = 144$ examples).

	RF	NB	LR	IB3	NN	J48	VFI	J-Rip
Personal	98.4	97.6	97.7	98.3	<u>98.7</u>	96.5	96.6	95.1
Hybrid	95.0	82.8	84.6	<u>96.5</u>	92.1	91.8	76.0	91.1
Impersonal	<u>75.9</u>	74.5	72.7	68.4	67.8	69.1	68.3	70.2
Average	<u>89.8</u>	85.0	85.0	87.7	86.2	85.8	80.3	85.5

Table 2. Predictive accuracy of the activity recognition task.

The hybrid models typically perform closer to the personal models than the impersonal models. We were greatly surprised by this because even though the hybrid models evaluate performance on users that are in the training set, we did not think it could really exploit the data from just

one of the 59 users. These results clearly show that this is not true, and that the good performance of the hybrid models is due to the ability to focus on and exploit a small fraction of the training data (on average 1/59). This means that the classification algorithms can effectively identify the movement patterns of a particular user from among a host of users. In retrospect this is not surprising, since recent work by our research group demonstrated that biometric models induced from accelerometer data can identify a user from a set of users with near perfect accuracy (Kwapisz, Weiss, and Moore 2010a). Because the hybrid model performs more poorly than the personal model, but still requires the acquisition of labeled training data from each user, we see little reason to use the hybrid model—except possibly in cases where one has only a *tiny* amount of personal data (we will evaluate this in the future).

Our main focus is on the comparative performance of the three types of models, but our results suggest which classification methods may be best suited to activity recognition, given our formulation of the problem. For personal models, NN does best, although RF and IB3 also perform competitively. For impersonal models, RF does best.

Table 3 shows the activity recognition performance for the personal and impersonal models for each activity, and for three best-performing classification algorithms and a “baseline” strategy. The baseline strategy always predicts the specified activity, or, when assessing overall performance, the most common activity. The personal models outperform the impersonal models for every activity, usually by a substantial amount, although impersonal models still outperform the baseline strategy.

	% of Records Correctly Classified						
	Personal			Impersonal			Base-line
	IB3	RF	NN	IB3	RF	NN	
Walking	<u>99.1</u>	98.9	99.0	65.2	<u>73.0</u>	56.8	36.6
Jogging	99.5	99.6	<u>99.8</u>	89.0	<u>95.2</u>	92.1	21.0
Stairs	96.4	96.8	<u>98.0</u>	65.1	61.5	<u>68.0</u>	16.7
Sitting	98.2	<u>98.7</u>	98.1	67.6	<u>81.5</u>	66.7	12.3
Standing	96.4	<u>97.8</u>	97.5	75.2	<u>91.9</u>	88.0	7.4
Lying Down	95.9	<u>95.0</u>	<u>97.5</u>	34.0	45.1	<u>45.5</u>	6.1
Overall	98.3	98.4	98.7	68.4	75.9	67.8	36.6

Table 3. Predictive accuracy on a per-activity basis.

Table 4 provides the confusion matrices associated with the Random Forest learner for the impersonal and personal models. We begin our analysis with the impersonal models since they have the most errors. The results in Table 4a indicate that {walking, stairs} and {lying down, sitting} are the two sets of activities most often confused. The confusion between walking and stairs may be due to the similar time between steps and exacerbated by the differences that people have when performing each of these activities. It is easy to see how lying down and sitting can be confused, since the orientation of one’s pocket will be at a similar angle in both cases. While the results for personal

models in Table 4b show that these two sets of activities are still confused most, the frequency of such errors is reduced by more than a factor of 10. This indicates that it is possible to learn the user-specific differences in these two sets of activities—and that these differences are not the same for all people. This is a key lesson for activity recognition and an argument for the use of personal models.

		Predicted Class					
		Walk	Jog	Stairs	Sit	Stand	Lie
Actual Class	Walking	2480	66	819	22	8	2
	Jogging	51	1854	41	1	0	1
	Stairs	518	69	953	2	4	3
	Sitting	7	5	3	931	19	178
	Standing	3	0	12	19	633	22
	Lying down	7	0	5	284	14	255

(a) Impersonal

		Predicted Class					
		Walk	Jog	Stairs	Sit	Stand	Lie
Actual Class	Walking	3359	3	30	1	3	1
	Jogging	5	1940	3	0	0	0
	Stairs	40	5	1500	2	2	0
	Sitting	3	0	1	1128	2	9
	Standing	5	0	8	2	674	0
	Lying down	3	2	4	18	1	537

(b) Personal

Table 4. Confusion matrices for random forest (RF) algorithm.

The results presented thus far are averages over all users. However, it is informative to see how activity recognition performance varies between users. Fig. 4 provides this information for the personal models and shows that these models perform consistently well for almost all users. The minor outliers that do show poor performance are primarily due to the high levels of class imbalance in those user’s data. For example, the user with the second highest error rate has 59 examples of walking data but only between 5 and 8 examples each for stairs, sitting, standing, and lying down. The user with the worst accuracy had a similar class distribution, but also had a leg injury. Thus, the few problems that do occur for the personal models appear to be due to high levels of class imbalance or due to an injury.




Figure 4. Per-user accuracy for personal models.

Fig. 5 shows a much broader distribution of performance for the impersonal models. There are still some users with

classification accuracies in the 95-100% range, but there are many users with *extremely* poor performance. Our analysis showed that most of these poor-performing users performed quite well when using personal models. This provides clear evidence that there are many users who move differently from other users, which “confuses” the impersonal models—but not the personal models, which can learn how a specific user moves.




Figure 5. Per-user accuracy for impersonal models

As part of our data collection protocol, we collect information about the physical characteristics of each user (height, weight, sex, shoe size, etc.). We analyzed this information to determine if people with particular or extreme characteristics are especially hard to predict when using the impersonal models, but, partly due to the limited number of users, we could find only suggestive patterns. For example, of the ten users that were hardest to predict using the RF impersonal models, three were among the oldest users in the study. In the future we plan to collect data from substantially more users—and diverse users—so that we can better assess the impact of such factors.

The success of the personal models encourages us to provide a self-training interface so that users can generate labeled training data themselves. We have already done this to a great extent and this has greatly reduced the workload for our researchers. In the near future we will also allow users to automatically generate their own personal models. But this is not to say that the performance of the impersonal models is inadequate for all applications. For example, although the predictions for each individual 10-second interval may be relatively poor, we still may be able to accurately quantify the time spent on each activity over a long period of time.

Related Work

Activity recognition has recently garnered significant attention due to the availability of accelerometers in consumer products. But activity recognition has been an active research topic for many years. Early work employed multiple accelerometers positioned across the user’s body. One well-known study used five biaxial accelerometers to identify 20 activities (Bao and Intille 2004). That study determined that the accelerometer placed on the thigh was most powerful for distinguishing between activities—a fortunate

finding given that smartphones are often stored in a user's pocket. Another study (Krishnan et al. 2008), used two accelerometers and claimed that a single thigh accelerometer is insufficient for classifying activities such as sitting, lying down, walking, and running—a claim our research refutes. A study by Tapia et al. (2007) used five accelerometers to identify thirty gymnasium activities while yet another used three accelerometers to identify seven lower body activities (Krishnan and Panchanathan 2008). While the systems described in these papers are capable of identifying a wide range of activities with good accuracy, they cannot be used for mainstream applications due to their reliance on multiple special-purpose accelerometers. These studies also involved relatively few participants.

Some studies have used other sensors in conjunction with an accelerometer to perform activity recognition, including: a heart monitor (Tapia et al. 2007), a light sensor (Maurer et al. 2006), and an image sensor worn at the waist (Cho et al. 2008). Smartphones can support such multi-sensor approaches, and we plan to exploit this in the future.

Several activity recognition systems have incorporated smartphones, but only as a storage device (Gyorbiro et al. 2008; Ravi et al. 2005). Other systems have used commercial mobile devices as the primary component of an activity recognition system. One effort explored the use of a variety of smartphone sensors (microphone, accelerometer, GPS, and camera) for activity recognition and mobile social networking applications (Miluzzo et al. 2008). An activity recognition system using the Nokia N95 phone to distinguish between sitting, standing, walking, running, driving, and bicycling was able to achieve relatively high activity recognition accuracy, but did not consider stair climbing and only involved four users (Yang 2009). Another effort also used the Nokia N95 phone to recognize six user activities, but only evaluated personalized models (Brezmes et al. 2009), while another smartphone-based system only evaluated hybrid models (Kwapisz, Weiss, and Moore, 2010b). Khan et al. (2010) achieved 96% accuracy on the same activities evaluating smartphone data from diverse body locations, but the study only included 6 users and the methodology does not provide enough information to determine the model type.

There has been relatively little comparative analysis of the different types of activity recognition models. Most studies, like the following, analyze only one type of model: impersonal models (Brezmes et al. 2009; Gyorbiro et al. 2008; Ravi et al. 2005), personal models (Miluzzo et al. 2008; Yang 2009), and hybrid models (Kwapisz, Weiss, and Moore, 2010b). Two studies did compare personal and impersonal models, but neither of these used a smartphone and both employed five accelerometers—thus any conclusions would not necessarily apply to a smartphone-based system. The first of these studies concluded that impersonal models always outperform personal ones, due to the additional training data; it further showed that when the impersonal and personal training set sizes are equalized, per-

sonal models only slightly outperform impersonal ones (Bao and Intille 2004). Our results clearly and dramatically contradict this result (but for smartphone-based systems). In the second study the personal models outperform the impersonal models—but there is virtually no analysis or discussion of this—as the paper focuses on other issues (Tapia et al. 2007). Thus our paper is the most comprehensive study on the impact of model-type on activity recognition—especially as it relates to smartphone-based systems. Furthermore, we additionally evaluate hybrid models and include many more users in our study than prior studies, leading to more general and reliable results.

Activity recognition is a key field in human-computer interaction, particularly when mobile devices are involved. Even before smartphones came equipped with sensors, researchers strapped sensors to these mobile devices to support context awareness (Hinckley et al. 2000) and to allow them to respond to basic context, such as orientation, by rotating the display. Schmidt (2000) proposed that implicit interaction would be the next big shift in human-computer interaction because it would further reduce the human overhead of handling the interface. Activity recognition, including the type described in this paper, completely eliminates this human overhead by allowing devices to respond to a user without any explicit input.

Conclusion and Future Work

In this paper we describe and evaluate a machine learning approach for implementing activity recognition, in an unobtrusive manner, using only a smartphone. We demonstrate that nearly perfect results can be achieved if a personalized model is constructed, even using only a small amount of user-specific training data. We further show that impersonal models perform much more poorly than personal models. An analysis of the data shows that impersonal models cannot effectively distinguish between certain activities, and further, that this is largely due to impersonal models performing horrendously on some users. Personal models can easily handle the problematic users in almost every case. We also evaluate hybrid models and show that their performance is closer to that of personal models than impersonal models—but since hybrid models require user-specific training data, one may as well just use personal models instead. This is the first such study to provide a careful analysis of these different models/learning scenarios, and we view this as a key contribution of the paper. This work should greatly influence the design of future activity recognition systems and the higher level activity representation systems which rely on them.

We plan to extend our activity recognition work in several ways. We plan to cover additional activities and to utilize additional sensors (e.g., GPS). We continue to collect data and thus will add many more users to our study. We also plan to incorporate physical characteristics (e.g., height, weight, sex) into our activity recognition models

and expect that this will improve the performance of all of our models—but especially the impersonal models by “personalizing” them with relatively easy to obtain information. We also plan to analyze those users who perform very poorly with impersonal models to gain a better understanding of why they fare so poorly—and if this can be remedied. We also plan to collect data from users over a period of days, weeks, and months in order to determine how their movement patterns change over time (currently personal models are evaluated on data from within the same training session).

One of our main goals is to make our activity recognition work available to smartphone users and researchers, via a downloadable app. We have been working on this for over a year and this effort is now progressing quickly with the support of NSF grant 1116124. The Actitracker system we are building will track a user’s activities and provide reports via a secure account and web interface. This application will help people ensure that they and their children are sufficiently active to maintain good health, and, as noted in the introduction, this can help improve and save lives, given the number of conditions and diseases associated with physical inactivity. The system will also include the ability to collect one’s own labeled training data and generate a personalized activity recognition model. Such a system will hopefully stimulate further advances in activity recognition and highlight the importance of personal activity recognition models. An early version of this system is currently being tested for deployment.

References

- Apple iPhone and Apple iPod Touch. 2009. Apple Inc. www.apple.com.
- Bao, L. and Intille, S. 2004. Activity recognition from User-Annotated Acceleration Data. *Lecture Notes Computer Science* 3001:1-17.
- Brezmes, T., Gorricho, J.L., and Cotrina, J. 2009. Activity Recognition from Accelerometer Data on Mobile Phones. In *Proceedings of the 10th International Work-Conference on Artificial Neural Networks*, 796-799.
- Cho, Y., Nam, Y., Choi, Y.-J., and Cho, W.-D. 2008. Smart-Buckle: Human Activity Recognition using a 3-axis Accelerometer and a Wearable Camera. *HealthNet*.
- Gyorbiro, N., Fabian, A., and Homanyi, G. 2008. An Activity Recognition System for Mobile Phones. *Mobile Networks and Applications* 14(1):82-91.
- Hinckley, K., Pierce, J., Sinclair, M., and Horvitz, E. 2000. Sensing Techniques for Mobile Interaction. In *Proceedings of the 13th annual ACM symposium on User Interface Software and Technology*, New York, NY, 91-100.
- Khan, A.M., Lee, Y.K., Lee, S.Y., and Kim, T.S. 2010. Human Activity Recognition via an Accelerometer-Enabled-Smartphone Using Kernel Discriminant Analysis. In *Proceedings of 5th International Conference on Future Information Technology*. 1-6.
- Koplan, J.P., Liverman, C.T., and Kraak, V.I. 2005. Preventing Childhood Obesity: Health in Balance. National Academies Press, Washington DC.
- Kwapisz, J.R., Weiss, G.M., and Moore, S.A. 2010a. Cell Phone-Based Biometric Identification. In *Proceedings of the IEEE Fourth International Conference on Biometrics: Theory, Applications and Systems*, Washington DC.
- Kwapisz, J.R., Weiss, G.M., and Moore, S.A. 2010b. Activity Recognition using Cell Phone Accelerometers, In *Proceedings of the 4th International Workshop on Knowledge Discovery from Sensor Data*, Washington DC, 10-18.
- Maurer, U., Smailagic, A., Siewiorek, D., and Deisher, M. 2006. Activity Recognition and Monitoring using Multiple Sensors on Different Body Positions. In *Proceedings of the IEEE International Workshop on Wearable and Implantable Sensor Networks*, vol. 3 no.5.
- Miluzzo, E., et al., 2008. Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application. The *6th ACM Conference on Embedded Networked Sensor Systems*, 337-350.
- World Health Organization. 2002. Physical inactivity a leading cause of disease and disability, warns WHO. [<http://www.who.int/mediacentre/news/releases/release23>].
- Ravi, N., Dandekar, N., et al., 2005. Activity Recognition from Accelerometer Data. In *Proceedings of the 17th Conference on Innovative Applications of Artificial Intelligence*, 1541-1546.
- Schmidt, A. 2000. Implicit Human Computer Interaction Through Context. *Personal Technologies* 4(2):191-199.
- Tapia, E.M., Intille, S.S., et al.. 2007. Real-Time Recognition of Physical Activities and Their Intensities using Wireless Accelerometers and a Heart Rate Monitor. In *Proceedings of the 2007 11th IEEE International Symposium on Wearable Computers*.
- Weiss, G.M., and Hirsh, H. 1998. Learning to Predict Rare Events in Event Sequences. In *Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining*. AAAI Press, Menlo Park, CA, 359-363.
- Weiss, G.M. 2012a. WISDM (Wireless Sensor Data Mining) Project. Fordham University, Department of Computer and Information Science, <http://www.cis.fordham.edu/wisdm>.
- Weiss, G.M. 2012b. Actitracker. <http://www.actitracker.com>.
- Witten, I.H., and Frank, E. 2005. *Data Mining: Practical Machine Learning Tools and Techniques*, 2nd ed. Morgan Kaufmann.
- Yang, J. 2009. Toward Physical Activity Diary: Motion Recognition using Simple Acceleration Features with Mobile Phones. In *1st International Workshop on Interactive Multimedia for Consumer Electronics at ACM Multimedia*.