TwitterJacket: An automated activity and health monitoring solution for the elderly

Shahriyar Amini, Priya Narasimhan

October 18, 2009

CMU-CyLab-10-003

CyLab
Carnegie Mellon University
Pittsburgh, PA 15213

TwitterJacket: An automated activity and health monitoring solution for the elderly

Shahriyar Amini
Carnegie Mellon University
Electrical & Computer Engineering Department
Pittsburgh, PA 15213
shahriyar@cmu.edu

Priya Narasimhan
Carnegie Mellon University
Electrical & Computer Engineering Department
Pittsburgh, PA 15213
priya@cs.cmu.edu

ABSTRACT
This paper explores automation in monitoring the well-being of senior adults. Namely, it introduces TwitterJacket, an automated activity and health monitoring solution through the use of wearable sensors and mobile devices. The system takes advantage of wearable electrocardiography (ECG) monitoring and accelerometer data to recognize user activity and publish on Twitter. In this paper, we discuss the motivation behind such platforms, the TwitterJacket architecture, and the present implementation.

Keywords
Social Networking, Wearable Computing, Activity Recognition, ECG, EKG, Twitter

1. INTRODUCTION
According to the United Nation's Population Division, DESA, the number of persons over the age of 60 has tripled over the past 50 years. Further, it is expected to triple once more over the next 50 years [1]. Based on statistics provided by the U.S. Census Bureau, there are 51.7 million individuals over the age of 60 in the U.S., as of 2008[1]. This figure represents over 17% of the U.S. population. With the world population aging, it is necessary to rely on technology to provide for the elderly in a more efficient and improved manner. In this paper, we focus on an unobtrusive method of informing interested parties regarding the well-being of senior adults.

Many senior adults enjoy the comfort of their homes as opposed to institutional care [15]. With economic reasons as well as personal comfort at hand, many older individuals stay in private homes alone or with family [13]. However, for the relatives of older adults, safety is a major concern when deciding on the issue of institutional care. Considering the busy lifestyle of most individuals, it is not practical for most individuals to maintain a close watch of their love-ones. Improvements in non-invasive monitoring of older individuals affords relatives a greater sense of safety of their loved-ones. As a result of this greater sense of safety, the use or the consideration of institutional care will be postponed by the elderly and their loved-ones. With the average cost of a private room in a nursing home in 2004 at $70,000[2], delaying the move towards institutional care can have cost savings for families. This cost saving is also shared by government programs that provide for the elderly.

Activity and health monitoring is also valuable in the institutional context. Our system will allow for improved information retrieval regarding the type and duration of activities and the health of the patient in nursing homes. The staff will have additional information aiding in the consideration of various aspects of the patients' lives such as diet and exercise. It is impractical to assign a human observer to every patient; therefore, without the proper use of technology to support activity and vital monitoring, the establishment will be less capable in offering personalized care.

It is our goal to create a non-invasive personalized monitoring solution that automatically informs interested stakeholders (family, healthcare professionals, patient care-providers, etc.) of the daily activities and health of senior adults in real-time. Our system utilizes a combination of sensors– wearable wireless accelerometer sensors on the wrist and the thigh to support activity recognition, a wireless wearable ECG solution to monitor the heart, and a mobile device to upload computed features of the data to a back-end server. Refer to Figure 1 for an overview of TwitterJacket. Once the values are uploaded to the server, the server updates the individual's status with an activity status based on the user's setting for update rates. The individual's activity and vital status are automatically updated in real-time on Twitter, allowing the stakeholders an easy way to monitor the individual's health by following his/her tweets. The TwitterJacket, as presented here in its first version, uses a one-to-many feed (one individual observed by many stakeholders). We envision future versions of the TwitterJacket using Twitter's many-to-many feed (a network of individuals observing each other's activities) to support more social-networking applications beyond healthcare.

2. PROBLEM STATEMENT
Our problem statement is divided into three sections describing our goals, non-goals, and assumptions for TwitterJacket.

[1]www.census.gov

[2]www.consumeraffairs.com

Figure 1: TwitterJacket takes advantage of accelerometers on the thigh and the wrist as well as an ECG monitoring t-shirt to observe user activity and health and report the results on Twitter.

Goals: TwitterJacket aims to present an unobtrusive platform to monitor activity and health status for individuals. TwitterJacket should detect a range of coarse-grained activities such as sitting, sleeping, standing, walking, and running. The subject's activity status should be updated on Twitter with reasonable delay, and as near real-time as possible. Considering that activities lasting for less than three minutes are marked as transient and are not reported, the reasonable delay is approximately three minutes from the start of the activity.

Non-Goals: This paper mainly focuses on the architecture of TwitterJacket and does not touch on quality of service aspects such as privacy and security (although Twitter does provide for privacy through its "protected tweets" feature). TwitterJacket does not aim for complete coverage of activities such as ironing clothes, bathing, or loading the dishwasher. Further, it does not aim to perform fine-grained activity recognition, i.e. it does not differentiate between sitting and watching TV.

Assumptions: We assume that there exists a notion of time to provide a chronological order of activities and health status. Further, we assume that the batteries to all of the mobile devices that measure, collect, and analyze data are charged. Finally, we assume that the user cooperates by wearing the TwitterJacket devices during the activity detection and health monitoring phases.

3. DESIGN DECISIONS

3.1 Why Twitter?

Twitter is a publish-and-subscribe (many-to-many follower) enabled service that allows for unobtrusive presentation and retrieval of information. It is convenient for relatives and caregivers to obtain access to Twitter accounts by creating private accounts with protected tweets that are visible only to select users specified by the Twitter account-holder. Furthermore, posting the activity and health information via Twitter allows for utilization of a relatively mature platform for data dissemination, using readily available COTS software and hardware. Twitter posts (tweets) can be remotely accessed through mobile and stationary devices with existing APIs in multiple languages. Because the information provided by TwitterJacket can be used for other applications, the existence of APIs in various programming languages is beneficial.

3.2 Why ECG?

Heart disease is the number one cause of death in the U.S., England, Canada, and Wales[3][4]. As a measure towards monitoring the health of senior adults we have added electrocardiography to our system. We see mobile ECG as a convenient method to detect heart problems early. Further, ECG monitoring is a painless and well-known medical procedure. As a result, senior adults will be more comfortable with providing such vital data.

In addition to the accelerometer data, the ECG features collected can be used in training classifiers and detecting activities. We see ECG as an appropriate measure in differentiating between activities such as sitting and exercising.

4. BACKGROUND

In this section, we discuss the background information regarding the various portions of our activity and health monitoring system, namely, asynchronous messaging, electrocardiography (ECG), and activity recognition.

4.1 Asynchronous Messaging

With the advent of wireless services such as SMS and MMS, and social networking platforms such as Facebook, MySpace, and Twitter[5], wireless networks and the Internet have evolved to address yet another problem which is easier to mitigate in the virtual world. Prior to services such as email, blogs, and micro-blogs, there were few options in updating other individuals with a relatively fresh status of oneself in a non-invasive manner. The more popular solutions in the past include physical mail, fax, and telephone. Considering usage with frequent updates, each technology is associated with limitations.

In the recent years services such as SMS, MMS, and micro-blogs have allowed users to share information with others in a less intrusive manner. The information arrives at the intended destination with relatively little delay and does not require immediate response. Our system's main communication channel in updating activity and health status is built on top of a micro-blogging service, Twitter.

With Twitter's asynchronous messaging model of publish and subscribe (pub/sub), the sender of messages are not required to send information to a particular receiver in hope of a response. Rather, messages are published by the senders and are available to any interested individual who subscribes

[3] www.cdc.gov

[4] www.statistics.gov.uk

[5] www.facebook.com, www.myspace.com, www.twitter.com

to receive them. Furthermore, the messages are stored once published and become available for viewing almost immediately. In the case of Twitter, messages can only be deleted by the publisher.

4.2 Electrocardiography (ECG)

ECG is a simple and painless method of recording the heart's electrical activity. An electrical signal is responsible for the contraction of the heart and the pumping of blood. With each beat, the electrical signal travels from the top to the bottom of the heart. The electrical signals set the rhythm of the heartbeat. ECG shows how fast the heart is beating, whether the heartbeat is steady or irregular, and the strength and timing of electrical signals as they pass through the heart. ECG data can be used to detect and evaluate heart problems such as arrhythmia, heart attack, and heart failure.

In the typical stationary setting, the ECG readings are taken by placing twelve electrodes on the skin of chest, arm, and legs. The electrodes are approximately the size of a quarter. A typical test takes about ten minutes.

4.3 Activity Recognition

Activity Recognition aims to detect the activity performed by a subject or subjects through observations made regarding the activity and the environment. Machine learning techniques such as decision trees, instance-based learning (IBL), and Naïve Bayes are trained on the observations to classify various activities [2]. Previous work has relied on data obtained through sensors such as accelerometers and the subject's location to train classifiers and detect activities [9]. In recent years there has been a growing interest in activity recognition as it allows for more reliable and efficient methods of providing data and services to users.

5. SYSTEM ARCHITECTURE

Our system is based on a three-tier architecture. At the user level, the system consists of sensors and a client front end that allows for the specification of activities. The data collection layer consists of calculating features of the data and sending the computed values to the backend server. The backend server is responsible for storing data for various users, and sending out status updates to Twitter. Figure 2 shows an overview of the TwitterJacket architecture.

5.1 User Level

The user is responsible for wearing two 3-axis accelerometers, one on the thigh and one on the wrist. The thigh and the wrist have been proposed as ideal locations to place accelerometers according to [2]. The accelerometer samples readings at 25Hz, which is in the range of commonly accepted sampling rates for activity recognition. The user wears a mobile ECG solution that samples at 200Hz. The user also carries a GPS and WiFi capable mobile device. The GPS is used to obtain the location of the user in outdoor settings. For indoor localization, a service such as SkyHook[6] or PlaceLab [7] can be used to approximate location based on access point signatures.

[6]www.skyhookwireless.com

In addition to providing user localization, the mobile device will be used as an input device by the user to further specify activities and location. This is an additional feature of TwitterJacket that makes the status updates more informative. The system works even in the cases where the senior adult user is not capable or is not interested in offering additional information. The additional information includes location specifiers: street address, city, state, country, postal code, landmarks; activity specifiers: sitting, standing, sleeping, walking, running, eating, watching, working; object and company specifiers as a list of objects, such as, toothbrush and toothpaste, and list of present groups or individuals; and finally state specifier: at work, at play, and at rest.

Activity recognition signatures are pre-loaded on the server side for coarse activities for a specified set, namely, sitting, standing, sleeping, walking, and running. However, the user can label new sets of activities and modify signatures for pre-set activities by providing labeled samples via the front end. By specifying a new activity, the user, a loved-one, or a caregiver can provide a signature for an additional activity.

The front-end client allows the user to specify privacy settings regarding location updates. Further, the user settings allow for changes in the rate of user status updates. User setting changes are reflected on the backend server.

5.2 Feature Extraction and Data Transfer

Upon collection of the data, certain features of the samples are calculated and sent to the backend server. The data is collected in 128 sample buffers for the accelerometers. Therefore, each buffer contains 5.12 seconds of data. The buffer contains 64 samples of overlap. As a result, each computed value is based on 2.56 seconds of fresh samples, and 2.56 seconds of previously recorded samples. We have selected similar features to [2]. For each 128 set of samples, we calculate the mean, standard deviation, energy, and normalized entropy. The energy is calculated as the sum of the absolute values of the FFT components. The normalized entropy is calculated based on the discrete absolute value of the FFT components. Furthermore, for each pair of axes on the accelerometer, correlation is calculated. The values are uploaded to the server using any available connection on the mobile device, whether WiFi access points or the cellular network.

A similar set of features is calculated for the ECG samples, namely the mean, standard deviation, normalized entropy and the heart rate in beats per minute. The mean of the ECG samples is only used to detect a shift in the baseline of the ECG samples. A first derivative based algorithm was used to calculate the heart rate by detecting peaks in the ECG samples obtained from [5]. Equation (1) is used to calculate a measure for the first derivative, where Y(n) approximates the slope and X refers to the input samples. 0.5 * Max(Y(n)) is used to threshold the values and detect peaks in the ECG samples. Once a peak value is detected, the next 30 samples covering 150ms cannot be flagged as peak values. This puts the ECG samples back near the baseline position and away from the peak.

$$Y(n) = -2X(n-2) - X(n-1) + X(n+1) + 2X(n+2) \quad (1)$$

Once a sample buffer is refilled and features are calculated

Figure 2: Overview of the TwitterJacket system architecture.

the data is transfered to the backend server. The data to be transferred is labeled with all additional inputs from the user, family members, or caregiver.

5.3 Data Storage and Tweeting

TwitterJacket is supported by a backend server, which collects values from the user's mobile device, and updates the user's Twitter status. The samples and corresponding labels are stored in a database for each user. The backend server is responsible for the update in the user's activity and health status. Based on our current architecture, transient activities with a duration of less than three minutes are not considered for activity updates. This limitation prevents the system from posting unnecessary updates for short lived activities that interrupt others. If an activity has continued and remained the same since the previous status update, the system does not update the user's status. The user's average heart rate calculated from the last post until the current post is appended to each status update. In the case of drastic deviations in ECG data, the system will repost a user's status. At the day's end, the server will tweet out duration of all activities performed during the day. This information is helpful in observing the change in user activities, such as exercising, over a period of time.

The backend server will mark the status of the user with GPS coordinates, or approximate location based on the type localization used, i.e. GPS or access point signatures. This only happens if the user has allowed location updates in the privacy settings of the front-end client. A sample tweet takes the form: `vital statistics activity @location`. Figure 3 presents sample tweets offering information activities, location, and data features.

6. IMPLEMENTATION

We use SunSPOTs (Sun Smart Programmable Object Technology) to collect readings and perform feature extraction on accelerometer and ECG samples. SunSPOT is a commercial sensor network platform developed by Sun Microsystems. The SunSPOT nodes operate on a 180MHz 32-bit ARM920T core processor with 512K RAM and 4M Flash. The devices are programmed almost entirely in Java. The nodes communicate using CC2420 chips, which are 802.15.4 compliant. Each roaming node has a 3-axis accelerometer which is run based on a 6G scale for our application. In order to record accelerometer readings, the users have to wear one SunSPOT node on the thigh, and one on the wrist. The node on the thigh is also responsible for collecting ECG data over Bluetooth and computing the features for both the ECG and thigh accelerometer data. The computed values are then sent to the wrist node. The wrist node is responsible for collecting accelerometer data for the wrist as well as feature extraction. The wrist node sends the features of the ECG, and thigh and wrist accelerometers to the mobile device over Bluetooth. Figure 4 shows the wrist SunSPOT setup.

The ECG samples are obtained through the use of a Vital Jacket. The Vital Jacket is a personal mobile ECG solution produced by a Portuguese based company of the same name[7]. The Vital Jacket is a t-shirt worn by the user that connects to three ECG leads. The ECG samples are aggregated on an approximately 2.5"x1.5"x0.75" device that stores the values on an SD card. The aggregation device also streams the ECG samples over bluetooth. In our setup, we use the bluetooth stream to obtain the ECG values. Figure 5 shows a sample ECG recording from the Vital Jacket.

We choose the Nokia N95 as the user mobile device. The user interface is written J2ME which allows for a more complete description of the user activity, location, and surrounding environment. The N95 receives the data from the wrist node over Bluetooth. Upon receiving the data, the packets are labeled with the pre-defined description and sent to the back end server. The values are sent via HTTP POST to server-side PHP scripts.

On the backend server, PHP scripts parse the values obtained from the user's mobile device and place them in a PostgreSQL database. Each user has 4 tables in the database for thigh accelerometer, wrist accelerometer, ECG, and user settings. If the previous user tweet was less than three minutes ago, the server only appends the values to the tables. Once an activity has been performed for longer than three minutes, or the user location has changed for longer than three minutes, the server publishes a new tweet. If the user has provided additional description for the activity, it is used

[7]www.vitaljacket.com

Figure 3: Twitter feed illustrating activities and locations along with values calculated for certain features of the data.

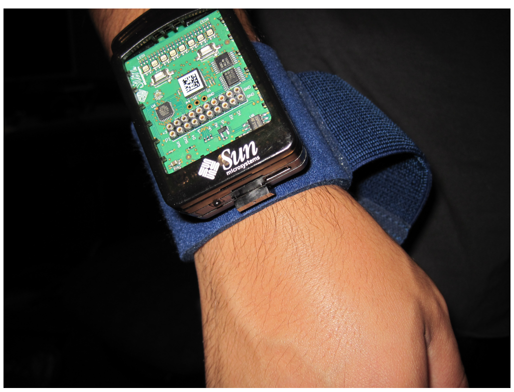

Figure 4: SunSPOTs are worn on the wrist and the thigh. Each node samples the corresponding accelerometer values.

to make the tweets more informative.

Weka[8] data mining software is used to train a decision tree classifier and detect activities. The server is preloaded with signatures for the activities: sitting, sleeping, standing, walking, and running. The user can modify signatures by labeling new samples for stored activities. Further, the user can add new activities by specifying a new activity using the user interface, the mobile device then uploads the samples with the new label. Upon arrival of new activity labels, the decision tree is trained on the new database to allow for detection of new activities. The server stores a tailored copy of the decision tree for each user.

7. RELATED WORK

In [11], Mynatt presents a method to convey information regarding a senior citizen to extended family members through the use of family portraits. Portraits are encircled by icons

[8]www.cs.waikato.ac.nz/ml/weka/

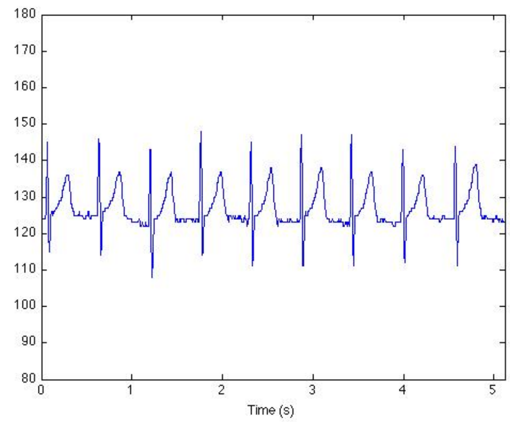

Figure 5: Sample ECG recording obtained through Vital Jacket.

that carry meaning for a single aspect of life such as activities, health, or relationships. The icons are updated daily through variations in size and placement, signifying changes in the elder's daily routine. As a result of the coarse update rate, the freshness of information obtained by the observer is limited. However, the method of presentation is effective as it blends in with the observer's surrounding home environment.

Activity and location recognition using accelerometers has been used to observe and detect activities and context [8]. Video monitoring has been used by [3] to reduce the risks of un-witnessed elopements from a dementia unit of a care facility. Chen trains HMMs to detect a sequence of activities that lead to an elopement. In [4], French proposes virtual coaching for manual wheelchair users through the use of the eWatch [10]. The accelerometer readings are used in helping the users avoid damaging forms of propulsion pattern. In [14], Song proposes an activity recognition system for the elderly using a wearable accelerometer.

In [6], Fulford-Jones et al. present a portable, low-power, wireless ECG system with software components that allow the capture and wireless transmission of heart activity traces. Wireless Medical Sensor Networks have been developed to relay a variety of sensor readings such as GPS, pulse oximetry, and blood pressure over a self-organizing wireless mesh network [12].

Our work differentiates itself by combining a unique form of presentation of activity and health status with accelerometer and ECG readings in a mobile setting.

8. CONCLUSION

TwitterJacket is a wearable system built on readily available sensors enabling automated, non-invasive monitoring of older persons. As a result, relatives and care-givers are better informed regarding their older person of interest, allowing them to provide better personalized care. Better monitoring enables a greater sense of safety for relatives and care-givers alike. Platforms such as TwitterJacket afford cost saving opportunity for families, institutional care

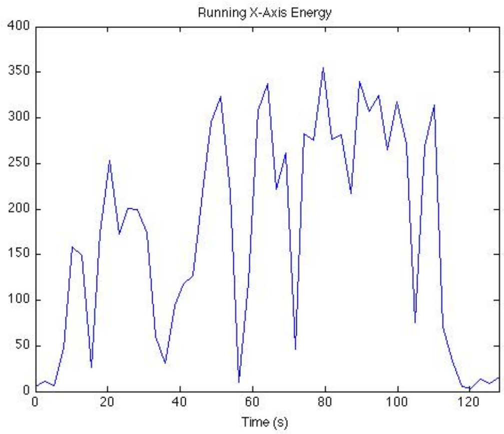

Figure 6: Sample running accelerometer x-axis energy readings.

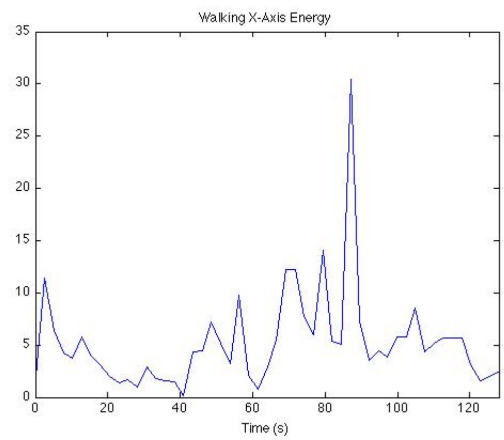

Figure 7: Sample walking accelerometer x-axis energy readings.

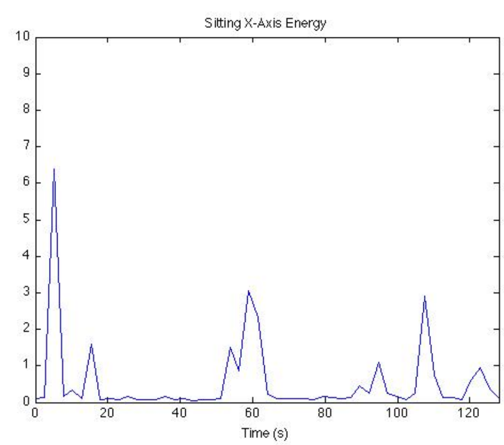

Figure 8: Sample sitting accelerometer x-axis energy readings.

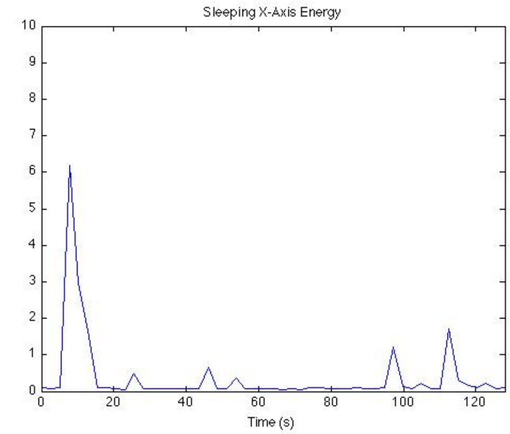

Figure 9: Sample sleeping accelerometer x-axis energy readings.

facilities, and government programs targeted at aiming the elderly.

Clearly, TwitterJacket is a work in progress. In the future, we intend to add to the range of activities that TwitterJacket supports. Specifically, we would like to add recognition support for events such as falling as it leads to earlier detection of emergency events. Further, we intend to explore various levels of information abstraction as to improve the usefulness of content to stakeholders, e.g., presenting information differently to a subject's grand-children vs. a care-giver. We will evaluate activities and health events as to establish categories of information and how each category of information should be presented to various audience. Beyond healthcare, we also intend to explore social-networking opportunities that might be afforded with multiple TwitterJacket wearers.

9. REFERENCES

[1] Magnitude and speed of populationg ageing. In *World Population Ageing 1950-2050*. United Nation's Population Division DESA.

[2] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. pages 1–17. Springer, 2004.

[3] D. Chen, A. Bharucha, and H. Wactlar. Intelligent video monitoring to improve safety of older persons. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pages 3814–3817, August 2007.

[4] B. French, D. Tyamagundlu, D. P. Siewiorek, A. Smailagic, and D. Ding. Towards a virtual coach for manual wheelchair users. In *ISWC '08: Proceedings of the 2008 12th IEEE International Symposium on Wearable Computers*, pages 77–80, Washington, DC, USA, 2008. IEEE Computer Society.

[5] G. Friesen, T. Jannett, M. Jadallah, S. Yates, S. Quint, and H. Nagle. A comparison of the noise sensitivity of nine qrs detection algorithms. *Biomedical Engineering, IEEE Transactions on*, 37(1):85–98, Jan. 1990.

[6] T. R. F. Fulford-jones, G. yeon Wei, and M. Welsh. A portable, low-power, wireless two-lead ekg system. In *In Proceedings of the 26th IEEE EMBS Annual International Conference*, 2004.

[7] J. Hong, G. Borriello, J. Landay, D. Mcdonald, B. Schilit, and D. Tygar. Privacy and security in the location-enhanced world wide web. In *Proceedings of Ubicomp 2003*, 2003.

[8] S.-W. Lee and K. Mase. Activity and location recognition using wearable sensors. *IEEE Pervasive Computing*, 1(3):24–32, 2002.

[9] L. Liao. *Location-based activity recognition*. PhD thesis, Seattle, WA, USA, 2006. Adviser-Fox, Dieter and Adviser-Kautz, Henry.

[10] U. Maurer, A. Rowe, A. Smailagic, and D. P. Siewiorek. ewatch: A wearable sensor and notification platform. In *BSN '06: Proceedings of the International Workshop on Wearable and Implantable Body Sensor Networks*, pages 142–145, Washington, DC, USA, 2006. IEEE Computer Society.

[11] E. D. Mynatt, J. Rowan, S. Craighill, and A. Jacobs. Digital family portraits: supporting peace of mind for extended family members. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 333–340, New York, NY, USA, 2001. ACM.

[12] V. Shnayder, B.-r. Chen, K. Lorincz, T. R. F. F. Jones, and M. Welsh. Sensor networks for medical care. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 314–314, New York, NY, USA, 2005. ACM.

[13] D. B. D. Smith. Human factors and aging: An overview of research needs and application opportunities. In *Human Factors*, pages 509–526, 1990.

[14] S.-K. Song, J. Jang, and S. Park. An efficient method for activity recognition of the elderly using tilt signals of tri-axial acceleration sensor. In *ICOST '08: Proceedings of the 6th international conference on Smart Homes and Health Telematics*, pages 99–104, Berlin, Heidelberg, 2008. Springer-Verlag.

[15] S. L. Willis. Everyday problem solving. In *Handbook of the psychology of aging*, pages 287–307. San Diego: Academic Press, 1996.