Applications of Mobile Activity Recognition*

Jeffrey W. Lockhart
Fordham University
lockhart@cis.fordham.edu

Tony Pulickal
Fordham University
pulickal@cis.fordham.edu

Gary M. Weiss
Fordham University
gweiss@cis.fordham.edu

ABSTRACT
Activity Recognition (AR), which identifies the activity that a user performs, is attracting a tremendous amount of attention, especially with the recent explosion of smart mobile devices. These ubiquitous mobile devices, most notably but not exclusively smartphones, provide the sensors, processing, and communication capabilities that enable the development of diverse and innovative activity recognition-based applications. However, although there has been a great deal of research into activity recognition, surprisingly little practical work has been done in the area of applications in mobile devices. In this paper we describe and categorize a variety of activity recognition-based applications. Our hope is that this work will encourage the development of such applications and also influence the direction of activity recognition research.

Author Keywords activity recognition, mobile computing, ubiquitous computing, context-aware, applications

ACM Classification Keywords H.4.0 [Information Systems Applications]: General.

General Terms Human Factors; Management.

INTRODUCTION
Activity Recognition (AR), which identifies the activity (e.g., walking, sitting, reading) that a user performs, has generated a great deal of interest within ubiquitous and mobile computing. In particular, the recent explosion of smart mobile devices with sensing, processing, and network capacity have opened up a huge range of possibilities for activity recognition. Given the large numbers of researchers and companies working in this area, one might expect that there would be many deployed activity recognition applications. Surprisingly, however, relatively little practical work has been done on the uses and applications of activity recognition with mobile devices. In this paper we describe and categorize possible activity recognition applications, with the hope that it will encourage the development of such applications and influence the direction of current activity recognition research.

To illustrate the general process for sensor-based mobile activity recognition, we use our Actitracker [2] application which is built upon the Wireless Sensor Data Mining (WISDM) Platform [12]. Actitracker runs on smart phones and collects readings from the accelerometer sensor. These readings are then transformed into examples which summarize short (10-second) periods with simple features such as average acceleration and frequency. Predictive models are then built using lightweight classification algorithms, such as neural networks and J48 decision trees, which can predict activities like walking or jogging with accuracies above 98% [9, 25]. These models are used to predict a user's activities throughout the day and the results are made available to the user via a web interface.

Many other activity recognition projects and platforms are being developed [10], but generally, there are few practical, deployed applications proposed for AR. Meta-applications such as "Code in the Air" are even being developed to make application development and AR easy for developers, as well as for end users who want to develop their own activity context-aware applications [20]. However, these efforts leave the question of what to do with AR up to the developers and users, and at this point such platforms are underutilized. With the availability of these platforms, it is even more important to find useful applications for AR.

Activity recognition applications fall across a broad range of disciplines. Generally speaking, there are three major types of applications: those that benefit end users, those that benefit developers or third parties, and those that benefit crowds and groups using the application. These types of applications are not mutually exclusive; an application that benefits groups will invariably benefit end users. In the following section we describe a number of activity recognition applications and organize them by categories.

APPLICATIONS FOR END USERS

Fitness Tracking
In recent years the personal fitness industry has seen an explosion of consumer devices with sensors for monitoring activity. These devices offer a variety of features including those offered by traditional pedometers, such as distance traveled and estimated calories burned, but newer devices are beginning to utilize accelerometers and location sensors to provide other data, such as stairs climbed, intensity and duration of physical activity, activity logging with context, and even how well one sleeps. One popular device, the Fitbit Ultra, claims to "motivate users to make small changes that add up to big results" [5]. Fitbit and similar fitness trackers, such as the Nike+ Fuelband [17], feature a combined score using the data they collect from accelerometers in order to provide users with a single value that can be used for comparative assessments of activity levels.

*This material is based upon work supported by the National Science Foundation under Grant No. 1116124 and two Fordham Summer Science Research Internship Grants.

UbiComp'12, September 5-8, 2012, Pittsburgh, USA.
ACM 978-1-4503-1224-0/12/09.

Most of these commercial devices have only rudimentary activity recognition capabilities. With Actitracker, we expand upon the existing functionality by detecting a greater variety of physical activities with a greater level of detail to provide more meaningful data to the end user. With our online reporting system, users will be able to track their detailed activity history and summaries to get a better picture of their activity and fitness levels on a daily basis but also over a span of weeks and months. Such detailed information can motivate more healthy activity by allowing users to set goals and see their progress. By combining the user's activity with location awareness, we believe that we can provide users with more effective motivation. For example, if the user is at work and the application detects a long period of time with no physical activity, it can encourage the user to stretch or walk for a few minutes. Most importantly, Actitracker will provide detailed and highly accurate activity histories for users to base judgments about their fitness. This avoids the problems of unreliable memory and wishful-thinking which can otherwise distort a user's perception of her activity. We anticipate that an initial version of Actitracker will be deployed for general (free) use by the public in September 2012 and that similar smartphone-based activity recognition fitness-tracking systems will eventually become commercially available. More information on Actitracker is available in our recent publications [9, 12, 25].

Health Monitoring

Smartphone-based activity monitoring can also provide clinicians and medical researchers with the ability to monitor and diagnose patients using continuously generated data, rather than data from a single short medical appointment. Certain patients, such as those facing Parkinson's Disorder, need longer periods of examination to accurately identify the symptoms and monitor the course of the condition [13]. Moreover, some symptoms may not present themselves during a medical examination. Activity recognition can be a tool to help doctors diagnose such conditions as they monitor daily activities in order to detect deviations from a typical routine or deterioration of a patient's current physical status. Patients undergoing physical therapy, in particular, gain the benefit of being able to have their statuses monitored in finer detail and therapists are able to ensure their patients improved quality of care with more accurate evaluations [11].

Additionally, much work has been done using actigraphs to measure the overall energy of a person's motion and these studies could benefit from more detailed activity recognition. One study found that children with ADHD reduced their activity levels significantly when faced with positive reinforcement from actigraphs worn on their waist [24]. Finer grained activity recognition could enable psychologists to target more specific behavior and the use of highly-available smart mobile devices can make these tools accessible to a wider audience.

Fall Detection

Falling, especially for the elderly, can cause serious injury. Mobile activity recognition systems, especially smartphone-based systems, can autonomously send an alert for help if the user has fallen. While commercial applications are available in the form of wearable sensors, they have significant deficiencies which dramatically reduce their effectiveness including: limitations on distance that can be traveled from the base, a limited number of response methods, and cost [3]. Utilizing smartphones instead of specialized wearable sensors allows the user greater flexibility in travel as they need not be tied to a receiver at home and a variety of response methods can be established within the application (e.g., text messaging, email, calling a relative, and calling emergency response if unable to reach a primary contact). Most importantly, as smartphones become more commonplace, such a system can operate with minimal or even no cost. This technology has been successfully demonstrated [3, 6] and such applications would be unobtrusive and inexpensive compared to having assisted-living staff monitor the individual or having extensive camera systems observing the residence for falls [15]. Mobile sensing devices could provide greater independence for at-risk-persons who would otherwise be confined to supervised environments. Other systems that attempt to do this detect a problem only when people miss activities in their daily routines [19], so they cannot respond as rapidly to falls and injuries.

Context-Aware Behavior

Higher level activity recognition (such as "in a meeting," "at lunch," or "with family") may be able to customize the device's behavior to benefit the user by disabling incoming calls or setting the device to silent mode. Or, if a user is determined to be exercising, the device can play music from a workout playlist. Other activities, such as "arrived at work" may trigger the device to check for new email. Research [8] has found that messages and alerts delivered when users are transitioning between activities are better received than those delivered at random times because the users do not feel that they are being interrupted. Thus, activity-context aware mobile devices would do best to hold notifications and messages until activity transitions.

Context-aware devices can also adapt their presentation for a user's activity. When one is walking, for instance, it is more difficult to read the screen of the device because it naturally moves with one's hand. Mashita *et al.* [14] suggest that context-aware devices adapt their presentation by using larger text when a user is active. Similarly, the noise caused by motion can make it harder to hear rings and alerts from the device, so we envision these types of applications going beyond the visual display to adjust volume and other settings. It is clear that these kinds of applications are close to the system's core functionality and we believe that AR will soon be embedded within mobile systems.

Home and Work Automation

Much work has been done in the field of smart homes and workplace automation using stationary sensors to detect the locations and activities of people in the room, but these systems can be costly to implement on a large scale. Tapia *et al.* [23] suggested using accelerometers attached to a user to perform AR and integrate that information with a smart home system. Modern smartphone technology makes this kind of work feasible for real-world applications where users are not willing to tape multiple devices to their extremities or fill their buildings with expensive sensors.

Self-Managing Systems

Activity recognition presents possibilities for self-managing systems that stretch human-computer interaction to its most unobtrusive limits. These applications are similar to the context-aware applications that we recently discussed—and in fact are context-aware applications—but the focus here is not on *directly* improving the user's experience, but rather to better manage the internal resources of the system. For example, when a user is mobile (e.g. walking, running, riding), it may make sense to turn off short range radios like Bluetooth and WiFi to save power (although in the case of driving it may be desirable to enable Bluetooth for interaction with the vehicle or one's hands-free headset). Similarly, when a user is stationary, it may make sense to disable location sensors (e.g. GPS), or reduce the frequency with which they update their location. Limiting radio use can also increase a user's privacy by not connecting to or broadcasting her presence to local radio receivers. These steps can also save substantial battery life for mobile devices, a very valuable feature given the demands placed on mobile devices. Herrmann *et al.* [7] demonstrated that activity context-aware tuning of mobile sensing applications can extend the battery life of a device up to 5x by reducing the rate at which sensor readings are taken when a user is inactive. Other research [27] has found more moderate, yet still impressive, power savings of 20-50% using this technique. This highlights the importance of self-management for mobile sensing systems.

APPLICATIONS FOR THIRD PARTIES

Activity recognition is a valuable source of information for businesses and organizations. While providing the basic front-end for visitors or employees, such as exhibit information or work-related communication systems, the application can also utilize activity recognition to provide an experience tailored to the individual.

Targeted Advertising

Advertising has already taken the next step in targeting users based on their current locations and past purchasing history. However, if an ad service knows a user's current or frequent activities, it can display ads that are even more relevant to that user [18]. These ads not only generate more revenue since the user is more likely to respond to the product or service advertised, but also make advertising less obtrusive because the ads are relevant to the user's activities and interests.

Research Platforms for Data Collection

Researchers in many disciplines, from marketing to healthcare, will always be interested in collecting activity data. Some applications, like our SensorCollector, exist solely to collect data from mobile devices for researcher evaluation; we need raw data to work with so that we can develop accurate and efficient modeling and prediction techniques [26]. Other applications will provide data to researchers working in a broad range of clinical fields; activity data can be useful for people studying physical therapy, psychology, fitness, and much more. Some platforms are being developed as generic systems that researchers can use to conduct studies and facilitate data collection [1]. These platforms are generally integrated with surveys and other traditional participant monitoring and feedback, but could benefit dramatically from the highly accurate activity log that AR can provide. Human observation is costly and participant reporting and recall are unreliable. Applications for mobile devices also offer the possibility of unprecedented scale because they can collect and aggregate data autonomously.

Corporate Management and Accounting

Activity recognition can also assist with employee management and accounting for employee time. Favela [4] demonstrates that mobile activity recognition systems can be used to build pervasive, context- and activity-aware networks for the monitoring of hospital staff, whose whereabouts and activities are important information for colleagues. In other contexts, where the location and activities of employees is not critical to health care, this technology certainly raises questions regarding employer ethics. However, voluntary AR has proven successful in other contexts. For example, the Progressive® insurance company runs a voluntary program where drivers are given a device that monitors the acceleration and duration of their driving. If a driver is deemed to drive safely, she is given a discount on her insurance payments [22]. Similarly, a company or organization could utilize AR on a voluntary basis, perhaps in combination with incentives.

APPLICATIONS FOR CROWDS AND GROUPS

Crowd-sourcing applications are not new, but activity recognition opens new opportunities for group, time, and place analysis and networking.

Traditional Social Networking

Social networking sites, such as Facebook and Twitter, provide users a medium to communicate their daily activities and thoughts with a broad audience. Applications like Foursquare allow users to advertise their current location, so broadcasting a user's current activity may be a valuable social tool. While fitness tracking technologies exist that provide daily or weekly fitness updates to social networking profiles [5], they do not provide responses to a user's current activity. In fact, nearly all facets of social networking rely on manual user operations. With AR, an automated aspect to social networking can be introduced allowing users to automatically update their profiles with information such as “jogging along Broadway with the track team” based on their current activity, location, and proximity to other users [16].

Activity-Based Social Networking

Beyond the simple posting of current activities, activity recognition can be used to *generate* social networks. After identifying users who are in proximity of one another and share activity patterns, an application, like Actitracker, can suggest a sphere of friends who fit similar activity qualities [2, 16]. This methodology creates a new basis on which to build a network, one that is based on more than just proximity or interests and instead links individuals based on their physical actions.

Activity-based Crowdsourcing

Some areas will naturally have high concentrations of specific activities (such as running at a track or sitting in a stadium). By analyzing the activity of many people in the same area, applications can learn and tag places and times as popular for biking or other recreation. Once these patterns are discovered, applications can detect abnormal behavior. If, for instance, a large number of people are suddenly running where they would normally sit or walk, the system could generate a notification of a possible emergency or disaster.

CONCLUSIONS AND FUTURE DIRECTIONS

Mobile activity recognition is a thriving research area. However, as this research area begins to mature, it is essential that high quality and innovative commercial applications are developed to exploit this research, and that the research anticipates and supports these applications. At the moment, such applications are rare, although we do believe that they will arrive shortly due to the ubiquity and increasing power of mobile devices such as smartphones, as well as the many advantages these applications can provide. This paper outlined a number of applications and application categories in the hope of raising awareness of these potential applications and accelerating their development. This paper also serves as a brief literature survey of existing activity recognition applications.

Our view is that most of the categories listed in Table 1 will yield commercially viable applications in the next 2-5 years. We believe that the health and fitness tracking applications will provide the biggest benefit and have the potential to assist with critical societal problems, such as childhood obesity, by monitoring activities and activity levels. This view is certainly supported by our National Science Foundation (NSF) grant for activity recognition, which was awarded as part of the NSF’s cross-cutting “Smart Health and Wellbeing” program. But we also believe that activity recognition will become embedded within the mobile devices to allow them to respond better to their users’ needs and to better conserve battery-life, which will remain a critical resource for the foreseeable future.

End User Applications	
Fitness tracking	Actitracker provides online activity history
Health monitoring	Evaluate patients over time rather than single session
Fall detection	Detect falls and take action
Context-aware behavior	Disable calls while jogging
Home/work automation	Smart homes that anticipate user's needs
Self-managing systems	Save battery by turning off WiFi while jogging
Third Party Applications	
Targeted advertising	Provide users with relevant ads
Research platform	Provide platform for collecting activity data
Corporate management & accounting	Track employee time and ensure spent appropriately
Crowd and Social Network (SN) Applications	
Traditional SNs	Share activity information with friends and followers
Activity-based SNs	Connect people based on their activity profiles
Place & event detection	Identify popular areas for exercise and recreation

Table 1: Summary of activity recognition applications (with example applications)

There are some changes that should occur on mobile devices in order to facilitate activity recognition-based applications. Mobile devices have to be optimized for the continuous monitoring and processing of sensor data. On Google's Android devices, the sensors are not active unless the processor is active, requiring developers to implement a “wake lock” in order to prevent the system from going to sleep while a background sensing application is running. Wake locks, however, keep the device fully functional and prevent it from conserving battery life by sleeping when it is not in use [12]. Moreover, hardware and software differences between manufacturers of Android mobile devices hinder development and integration of activity recognition applications. Other mobile platforms, like Apple's iOS, Windows Mobile, and Blackberry, provide even less support because they restrict access to sensors and the ability to run applications in the background, making it very difficult to develop smart sensing applications that do not interfere with the device’s regular use [21]. Mobile operating systems will need to adapt to support the emerging fields of smart mobile sensing and activity recognition before they become widely used.

REFERENCES

1. Actipal. http://actipal.com
2. Actitracker. https://actitracker.com
3. Dai, J.T., Bai, X.T., Yang, Z.T., Shen, Z., Xuan, D.T. PerFallD: a pervasive fall detection system using mobile phones. *Proc. IEEE International Conference on Pervasive Computing and Communications Workshops* (March-April, 2010): 292-297.
4. Favela, J., Tentori, M., Castro, L.A., Gonzalez, V.M., Moran, E.B., Martinez-Garcia, A.I. Activity Recognition for context-aware hospital applications: issues and opportunities for the deployment of pervasive networks. *Mobile Networks and Applications* 12, 2-3 (March, 2007): 155-171.
5. Fitbit. http://www.fitbit.com
6. He, Y. Fall Detection by built-in tri-accelerometer of smartphone. *Proc. IEEE-EBMS International Confrence on Biomedical and Health Informatics*, Hong Kong, China (Jan 2012): 184-187
7. Herrmann, R., Zappi, P., Rosing,T.S. Context aware management of mobile systems for sensing applications. *Proc. 2nd International Workshop on Mobile Sensing,* Beiging, China (April, 2012).
8. Ho, J., Intille, S.S. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. *Proc. SIGCHI conference on Human factors in computing systems*, New York, NY (2005): 909-918.
9. Kwapisz, J., Weiss, G.M., and Moore, S.A. Activity recognition using cell phone accelerometers. *ACM SIGKDD Explorations*, 12(2):74-82.
10. Lane, N.D., Miluzzo, E., Hong L., Peebles, D., Choudhury, T., and Campbell, A.T. A survey of mobile phone sensing. *Communications Magazine, IEEE* Vol.48, 9 (Sept. 2010): 140-150.
11. Lau, S.L. Supporting patient monitoring using activity recognition with a smartphone. *Proc. 7th International Symposium on Wireless Communication Systems* (Sept. 2010): 810-814.
12. Lockhart, J.W., Weiss, G.M., Xue, J.C., Gallagher, S.T. Grosner, A.B., and Pulickal, T.T. Design considerations for the WISDM smart phone-based sensor mining architecture. *Proc. 5th International Workshop on Knowledge Discovery from Sensor Data,* San Diego, CA (at KDD-2011): 25-33.
13. Loseu, V., Ghasemzadeh, H., Ostadabbas, S., Raveendranathan, N., Malan, J., and Jafari R. Applications of sensing platforms with wearable computers. *Proc. 3rd International Conference on Pervasive Technologies Related to Assistive Environments*, New York, NY, (2010).
14. Mashita, T., Shimatani, K., Iwata, M., Miyamoto, H., Komaki, D., Hara, T., Kiyokawa, K., Takemura, H., Nishio, S. Human activity recognition for a content search system considering situations of smartphone users *Proc. IEEE Virtual Reality*, Costa Mesa, CA (March, 2012): 1-2.
15. Miaou, S., Sung, P., and Huang, C., A customized human fall detection system using omni-camera images and personal information. *Proc. 1st Transdiciplinary Conference on Distributed Diagnosis and Home Healthcare* (2006): 39-42.
16. Miluzzo, E., Lane, N.D., Eisenman, S.B., and Campbell, A.T. CenceMe – injecting sensing presence into social networking applications. *Smart Sensing and Context* (2007): 1-28.
17. Nike+ Fuelband. http://www.nike.com/fuelband
18. Partridge, K. and Begole, B. Activity-based advertising techniques and challenges. *Proc. Workshop on Pervasive Advertising*. PARC, Nara, Japan (2009).
19. QuietCare, Care Innovations. http://www.careinnovations.com/products/quietcare-assisted-living-technology
20. Ravindranath, L., Thiagarajan, A., Balakrishnan, H., and Madden S. Code in the air: simplifying sensing and coordination tasks on smartphones. *Proc. HotMobile'12* San Diego, CA.
21. Sieling, J.D., Moon, J.K. Performance of smartphone on-board accelerometers for recording activity. *Obesity Society Annual Conference*, Orlando, FL (2011). www.actipal.com/MEI-TOS2011-Onboard_Accelerometer.pdf
22. Snapshot, Progressive. http://www.progressive.com/auto/snapshot-privacy-statement.aspx
23. Tapia, E.M., Intille, S.S., and Larson, K. Activity recognition in the home using simple and ubiquitous sensors. *Pervasive Computing* (2004): 158-175.
24. Tryon, W.W., Tryon, G.S., Kazlausky, T., Gruen, W., and Swanson, J.M. Reducing hyperactivity with a feedback actigraph: initial findings. *Clin Child Psychol Psychiatry* 11 (2006): 607-617.
25. Weiss, G.M., and Lockhart, J.W. The impact of personalization on smartphone-based activity recognition. *Proc. Activity Context Representation: Techniques and Languages Workshop*, Toronto, CN (at AAAI-2012).
26. WIreless Sensor Data Mining (WISDM) Lab. http://www.cis.fordham.edu/wisdm
27. Yan, Z., Subbaraju, V., Chakraborty, D., Misra A., Aberer, K. Energy-efficent continuous activity recognition on mobile phones: an activity-adaptive approach. *Proc. 16th International Symposium on Wearable Computers*, Newcastle, UK (June 2012).