

Ambulation: a tool for monitoring mobility patterns over time using mobile phones

Jason Ryder, Brent Longstaff, Sasank Reddy, Deborah Estrin
Center for Embedded Networked Sensing, UCLA

jryder@cs.ucla.edu, blongstaff@ucla.edu, sasank@ee.ucla.edu, destrin@cs.ucla.edu
UCLA 3563 Boelter Hall, Los Angeles, CA 90095-1596

Abstract

An important tool for evaluating the health of patients who suffer from mobility-affecting chronic diseases such as MS, Parkinson's, and Muscular Dystrophy is assessment of how much they walk. Ambulation is a mobility monitoring system that uses Android and Nokia N95 mobile phones to automatically detect the user's mobility mode. The user's only required interaction with the phone is turning it on and keeping it with him/her throughout the day, with the intention that it could be used as his/her everyday mobile phone for voice, data, and other applications, while Ambulation runs in the background. The phone uploads the collected mobility and location information to a server and a secure, intuitive web-based visualization of the data is available to the user and any family, friends or caregivers whom they authorize, allowing them to identify trends in their mobility and measure progress over time and in response to varying treatments.

1. Introduction

Many personal health monitoring systems exist that require expensive, customized hardware [1,2]. This can make these systems cost-prohibitive and complicated for a patient to use. A less expensive and simpler alternative is possible using modern mobile phones, which have built-in accelerometers and GPS capabilities and are ubiquitous in today's society [6,7].

This paper describes an algorithm and accompanying implementation that has been developed for such phones to infer a user's transportation mode (walking, motorized, or stationary) with data from these sensors. The application samples periodically (on the order of seconds) and classifies the user's activity with the data from each sample [3,9,12]. Ambulation starts automatically when the phone is turned on and requires no user interaction (it is possible for the user enable/disable the application manually, if desired).

We also address issues related to energy usage. Data sampling (especially GPS) and uploading over the cellular data network put a heavy strain on the battery and result in phone battery life of approximately 4-6 hours. This is unacceptable for the average user, as it would mean having to charge the phone several times daily in order to fully participate. Thus, it is necessary to actively manage the power consumed by the on-device application and for it to be aware of the current power state (e.g. running on a very low battery, plugged into an external power source) so as to take full advantage of the situational opportunity. With algorithms that make intelligent use of the GPS, the application consumes less power. The user can continue to use the phone as normal (for calls, messages, web browsing, etc.) while the activity classifier is running in the background. The phone automatically uploads the classifications and measurements to a private account on the web server.

Uploaded data is processed on the server to provide further inference of activity classification. A user is then able to securely login to a front-end web server and view data and classifications through meaningful visualizations that deliver complex and dense data in an intuitive and simple manner. This will allow the user to view trends and progress made with respect to their mobility over days, weeks, or months (broken down into the number of ambulatory trips for a given range of time periods).

The experimental design can be varied to suit needs as they arise. An example protocol may have chronic disease patients run this system for a full week (during all waking hours), several times a year. The week long duration would allow for a more natural capture of participants' physical status than would be possible to observe during individual measures within a clinical setting. This also provides a practical augmentation to current procedures (i.e. clinical visits), which would allow the status to be assessed several times a year instead of once or twice [3].

2. System architecture

In this section we describe the system architecture and flow of data displayed in Figure 1.




Figure 1. System architecture

- **Mobile Sensing Client** – Carried by the user to gather data pertaining to location and activity.
- **Data Storage Service** – The implementation uploads to SensorBase [4] as the data store, but is flexible to use other repositories.
- **Processing Module** – Used to process data collected by the mobile phone.
- **Visualization Platform** – Presents a human-readable view into large amounts data.
- **Notification Generator** – Provides the ability to notify caretakers based on the automatic detection of significant pattern deviations.

3. Mobile Sensing Client

The ambulation application tracks the user's location and mobility state. It determines the location by periodically sampling the GPS. To ascertain the current mobility state of the user (e.g. walking, running, or stationary), the phone records acceleration and speed information. It then runs a classifier on these data to classify the user's activity. There are two classifiers used: one that uses both accelerometer and GPS for use outdoors and an indoor classifier that relies solely on information from the accelerometer.

3.1. Classification

From the user's speed, the algorithm can discern patterns in the changes in acceleration, allowing it to characterize the different transportation modes. The indoor classifier collects accelerometer data and uses a decision tree classifier to distinguish the user's transport activity. The accelerometer data are recorded similarly to the method described in [1]. The variance, mean, and FFT coefficients from 1-10 Hz are used to characterize the movement of the device with the magnitude of the acceleration. It can detect whether the

user is stationary, walking or running. The motion caused by these activities differs enough that the accelerometer can effectively distinguish among them even without knowing the user's speed. Thus, GPS, which does not perform well indoors, is unnecessary for classifying these indoor activities. The outdoor classifier can determine whether the user is still, walking, running, biking, or driving because in addition to using a subset of the accelerometer features used by the indoor classifier, it takes advantage of GPS to determine the user's speed, which helps to distinguish among the various mobility modes.

3.2. GPS Power optimization

GPS is power-intensive, and can significantly decrease mobile phone battery life if left on. Since GPS is only useful outdoors, sampling GPS while a user is indoors is a waste of energy. A naive solution would have users deactivate the GPS when going indoors and reactivate it when they leave a building. However, not only is this inconvenient, but it can lead to missing data when he/she forgets to reactivate the GPS (or a dead battery if he/she forgets to turn it off). The application automates this with an adaptive GPS sampling strategy. It detects when it should activate or turn off the GPS by using motion as an indicator of when the user might be indoors or outdoors. When the device is still for five minutes, it infers that the user is probably indoors and does not need GPS. It switches it off and starts using the indoor classifier. If the user begins moving consistently (for the better part of the previous 90 seconds), the accelerometer senses it and starts the GPS. Whenever the phone is trying to use GPS but is unable (due to being indoors or near tall buildings), it continues to use the indoor classifier. When the GPS successfully obtains a satellite connection, the phone switches to the outdoor classifier. Figure 2 summarizes this process.




Figure 2. Flowchart

This optimization increases battery life by keeping the GPS off when it is unnecessary. The energy savings depend on how mobile the user is. The GPS will never activate if the user is stationary all the time, whereas it will always be on if the user travels nonstop. Without the optimization, the Android battery lasts

nearly 6.3 hours, whereas with the adaptive GPS it lasts up to 9 hours in a typical scenario – a 43% gain.

3.4. Data upload Power optimization

In addition to energy consumed by data collection, there is also an energy cost due to data upload. This is very expensive and takes a toll on battery life. In order for the implementation to remain practical to a wide audience of users, this requires an adaptive approach to decide when the best time for uploading is. The most naïve approach has the user activate and deactivate upload manually when conditions are good for uploading. This approach is victim to the same problems as for the manual de/activation of GPS – the user may forget to en/dis-able it. In the case of upload, it results in potentially long data upload delays (possibly filling the phone memory to capacity) or a quickly depleted battery. Fortunately, the proliferation of 3G networks and Wi-Fi enabled smartphones, allows for high data-rate transfers (with slower speed networks there is an issue of not being able to upload at the same pace as collection, resulting in a growing backlog of data that consumes phone memory – another limited resource). The high-data rate transfers allow for uploading a large time period's worth of data in a relatively short period of time. These properties allow for a simple and efficient upload strategy: upload any backlogged data upon application start up and subsequently upload only when plugged into an external power source. This strategy allows keeping up (with respect to data upload versus collection) with the user that does not run the application when the phone is plugged in, while still reaping the benefits of not using the upload radios for most of the time when the phone is running solely on battery power.

4. Server system and user interface

Once the data have been collected and uploaded, the next challenge is to create meaningful content for patients and caretakers in the form of data visualizations and statistical processing to reveal significant features of the data that are not immediately apparent. It is also extremely important to protect a users' privacy. With privacy in mind, a traditional role-based permissions framework is augmented with data-access controls at the level of user *relationships* (i.e. only users with explicitly defined relationships can view another's data).

An application of this system is as a chronic health monitoring system (this particular application was suggested by Dr. Florence Haseltine, NIH, and Kyle Brown of Innolyst.com). While the patient has the phone with them, their transport mode and location are being monitored. This information can be used to build

a baseline profile for the user. In the event of significant changes, the system can algorithmically detect significant pattern variations and take appropriate action. The goal/purpose of this system is not to diagnose medical problems, but to provide a tool for individuals to observe how their health status is changing over time and in response to various treatments and conditions, as well as to provide information to authorized medical professionals and caretakers to assess the patient's progress.

4.1. Data Storage Service

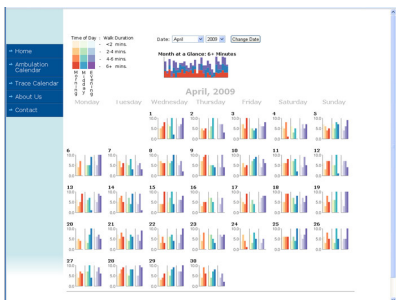
All data for an individual is private by default. The user must explicitly share their information with a caretaker in order for the data to be available to them. This will ensure that their information is only viewed by those they have granted permission to. To enforce these access rules, there must be guards in place at the data access level. Typical role-based permissions systems grant permissions based on the role(s) a user is assigned, rather than using finer-grained constraints on data access [7]. An example of such a role-based system would be a hospital healthcare system where any user with the *Hospital Staff* role would be able to access any patients' data. This can lead to problems such as leaks of information [7] or other inappropriate use of the information. Enforcing data access at the *relationship* level allows for much tighter control over who has access to what data. The finer-grained access controls also allow for sharing amongst other patients and family members that do not necessarily hold the role of healthcare professional.

4.2. Processing module and visualization

The system currently supports two modes of data visualization and is in the process of supporting automated anomaly and pattern deviation detection. With automation, it will be possible to generate a notification to concerned parties when there are significant changes in a patient's habits or patterns.

4.2.1. Ambulation summary visualization. A six minute ambulatory test is used by many health professionals to aid in diagnosis and progress assessment for chronic health problems in patients [3]. The ambulation summary, shown in figure 3, provides a daily snapshot of the user's ambulatory activity. Using the classified data, a post-processing step is used to chunk ambulation episodes into duration categories for each day. Furthermore, each day is broken down into the morning, midday and evening to give a more detailed view of ambulation activity. Each duration and time of day pair is then color-coded to assist in human visual processing. Finally, a complete summary for the

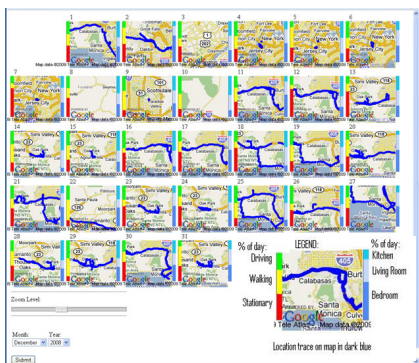
month is provided in a sparkline-style bar chart at the top of the screen. This overall summary provides a quick way to visually identify deviations from the norm for a patient, while the daily summaries allow for closer inspection of the days' details.



http://jryder.cens.ucla.edu/aipw/view/ambulation_sample

Figure 3. Ambulation summary

4.2.2. Daily trace calendar. The daily trace calendar visualization displays the user's location traces for each day in a monthly calendar. The expectation is that from this visual summary that spans an entire month, differences in daily patterns and habits with respect to location can be identified at a glance, whereas a day to day inspection may not reveal significant enough changes. For example, if a patient walks a specific route every Monday, but discontinues this towards the end of the month, this deviation would be visually identifiable from such a representation of the data.



http://jryder.cens.ucla.edu/aipw/view/trace_calendar_sample

Figure 4. Day trace calendar

5. Conclusion

Easily used applications such as the Ambulation system could help patients and caregivers track important trends relating to the health of patients with chronic ambulatory problems. Our on-phone application implements indoor and outdoor mobility classifiers that determine the user's transportation mode. An adaptive GPS strategy allows the phone to save power. Users simply keep the mobile device with them and their mobility information is collected and

uploaded to a server for analysis. An intuitive and powerful web-based front end then makes these data available to the individuals authorized by the patient.

One of the more interesting areas for future work is automated anomaly detection. The ability for a human to recognize interesting data characteristics is limited in the volume of data and the subtlety/complexity of the characteristics. While there are limitations in automated reasoning (which is one motivating factor behind providing visual representations), there are many mature statistical algorithms for detecting anomalous and other interesting characteristics of large datasets. One such algorithm that seems to be a good fit and is currently being implemented is Eigen Behavior through the use of singular value decomposition (SVD). Using the Eigen Behavior approach it will be possible to create a similarity score for different datasets. For example, comparing a week of data from May to a week of data in June could determine if there are significant differences in the activities of the patient during that time period [8].

6. References

- [1] <http://www.caregivertech.com/>.
- [2] <http://www.quietcaresystems.com/>.
- [3] L. Bao and S. Intille. Activity recognition from userannotated acceleration data. Lecture Notes in Computer Science, pages 1–17, 2004.
- [4] K. Chang, N. Yau, M. Hansen, and D. Estrin. Sensorbase.org-a centralized repository to slog sensor network data. Center for Embedded Network Sensing, page 116, 2006.
- [5] J. Fischer, D. Marino, R. Majumdar, and T. Millstein. Fine-Grained Access Control with Object-Sensitive Roles. In ECOOP, 2009.
- [6] A. Holzinger, G. Searle, and A. Nischelwitzer. On some aspects of improving mobile applications for the elderly. Lecture Notes In Computer Science, 4554:923, 2007.
- [7] S. Intille. A new research challenge: persuasive technology to motivate healthy aging. IEEE Transactions on information technology in Biomedicine, 8(3):235–237, 2004.
- [8] D. Lederer, S. Arcasoy, J. Wilt, F. D'Ovidio, J. Sonett, and S. Kawut. Six-minute-walk distance predicts waiting list survival in idiopathic pulmonary fibrosis. American journal of respiratory and critical care medicine, 174(6):659, 2006.
- [9] J. Lester, T. Choudhury and G. Borriello. A practical approach to recognizing physical activities. Lecture Notes in Computer Science, 3968:1–16, 2006.
- [10] S. Reddy, J. Burke, D. Estrin, M. Hansen, and M. Srivastava. Determining Transportation Mode On Mobile Phones. In ISWC, 2008.
- [11] S. Reddy, K. Shilton, J. Burke, D. Estrin, M. Hansen, and M. Srivastava. Using Context Annotated Mobility Profiles to Recruit Data Collectors in Participatory Sensing. In LoCA, pages 52–69, 2009.
- [12] Y. Zheng, L. Liu, L. Wang, and X. Xie. Learning transportation mode from raw GPS data for geographic applications on the Web. 2008.