

Accelerometer-based Activity Recognition on Smartphone

Xing Su
City University of New York
New York, NY 10016
sunnyxing@gmail.com

Hanghang Tong
Arizona State University
Tempe, AZ 85281
Htong6@asu.edu

Ping Ji
City University of New York
New York, NY 10016
pj.ping@gmail.com

ABSTRACT

Smartphones are ubiquitous and becoming more and more sophisticated, with ever-growing computing, networking and sensing powers. How can we help the users form a healthy habit by sending a reminder if s/he is sitting too long? How can we localize where we are inside a building and/or find the reception desk? Recognizing the physical activities (e.g., sitting, walking, jogging, etc) is a core building block to answer these questions and many more.

We present ACRE, a human activity recognition application on smartphone. ACRE takes the motion data from different sensors on smartphones as inputs (e.g., accelerometer, compass, etc), and predicts a user's motion activities (e.g., walking upstairs, standing, sitting, etc) in real-time. It provides some additional functionalities, such as incorporating a user's feedback, daily activity summerization, etc. The application is built on iOS 7.0 and will be released soon in Apple's App Store. We will invite the audience to experiment with our ACRE in terms of its effectiveness, efficiency and applicability to various domains and the potential for further improvements.

1. INTRODUCTION

Smartphones are ubiquitous and becoming an important part of people's daily life. Moreover, they are becoming more and more sophisticated, with ever-growing computing, networking, and sensing powers. Many main stream smartphones are equipped with various sensors, including accelerometers, GPS, light sensors, temperature sensors, gyroscopes, barometers, etc. This has opened the door for many interesting data mining applications, ranging from health and fitness monitoring, personal biometric signature, urban computing, assistive technology and elderly-care, to indoor localization and navigation, etc.

A common, key building block (often as the very first step) behind these applications is to recognize a user's physical activities, such as walking, sitting, walking upstairs, jogging, etc. For example, in order to help the user form a

healthy fitness habit, we can send a reminder if s/he has been sitting too long. Several recent popular fitness trackers are built upon wearable sensors and activity recognition techniques. They track people's steps taken, stairs climbed, calorie burned, distance travelled, quality of sleep, etc, and provide online services for living statistics.

We develop ACRE, a human activity recognition application on smartphones. ACRE takes the motion data from different sensors on smartphones as inputs (e.g., accelerometer, compass, etc) and predicts a user's motion activities (e.g., walking upstairs, standing, sitting, etc). Compared to other mobile activity detection applications, our work has the following key features. First (*Effectiveness and Efficiency*), by carefully designing the features and classification algorithms, ACRE is able to detect *more types* of activities, with a *higher detection accuracy* (up to 98.7%) in cross validation. Second (*Personalization*), ACRE provides the functionality to incorporate user-feedback so as to personalize the detection model to a specific user. Third (*Summarization*), ACRE provides a daily activity summerization, such as the statistics of each type of activities as well as the comparison of moving speed at different time ticks, etc. Fourth (*Deployment*), our application is built on iOS 7.0 combining the latest motion recognition API[1] and will soon be released in Apple's App Store (currently in the review process).

The rest of the paper is organized as follows. Section 2 demonstrates the main functionalities of ACRE. Section 3 presents the technical details. Section 4 reviews the related works and Section 5 concludes the paper.

2. APPLICATION DEMONSTRATION

In this section, we demonstrate the main functionalities of ACRE. Fig. 1 shows the three major parts of ACRE. Fig. 1(a) is the main screen is for *raw data reading*, including the line chart of accelerometion in X, Y and Z axis, followed by accelerometer data and the compass reading. The second part of the main screen is for *recognition*. The recognized activity will be shown with bigger and bold font, such as "Walking" in Fig. 1(a), a picture indicating current activity would appear on the left also. The *Turn* field is for turn detection. The button *View Activity* will lead to the second screen Fig. 1(b), the button *View Stats* will lead to the third screen Fig. 1(c), and the button *Send Data* will send current raw data to the server.

The second screen (Fig. 1(b)) is for activity recognition and personalization. ACRE displays the three most probable activities of a user at the current time tick. For each type of recognized activity, it also shows the confidence level. The

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

CIKM'14, November 3–7, 2014, Shanghai, China.

Copyright 2014 ACM 978-1-4503-2598-1/14/11

<http://dx.doi.org/10.1145/2661829.2661836>

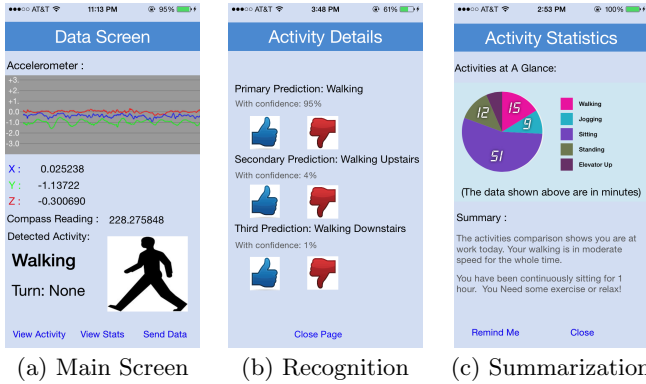


Figure 1: Activity Recognition and Summarization

thumb up and thumb down buttons provide the functionality to collect a user’s feedback. If there’s a hit in the three most possible activities, user could click the “*Thumb Up*” button under the correct prediction. The user can always click the “*Thumb Down*” as a negative feedback for the wrong predictions. The feedback information will be sent to the backend server to refine the prediction model so that it fits better with the current user’s specific moving patterns.

The third screen (Fig. 1(c)) is for the daily activity summarization and comparison. It has two parts. The first part is a pie chart showing a comparison of the total durations of the different activities within a certain period (e.g., one day). The second part is a summary of the daily overview in terms of his motion status. This could help reveal the user’s status of the day, e.g., s/he was busy rushing here and there in the morning, or s/he had a healthy day with lots of exercise, or it was just a long day of work with the dominant “sitting” activity for a long period.

Engaging the Audience. We expect that our demo will mainly attract two types of audience, including (1) the practitioners who are interested in developing novel mobile applications based on the activity recognition, and (2) database and data mining researchers who are interested in developing new algorithms and tools to address the unique challenges posed by human activity recognition.

3. TECHNICAL DETAILS

In this section, we present the algorithm details of our ACRE, followed by some implementation and evaluation details. Fig. 2 shows the overall flowchart of ACRE. It consists of two main parts. The *activity recognition* part shows the recognition process. ACRE will then collect data and train a general classification model in the backend. In the real-time recognition stage, by incorporating users’ feedback, it refines the classification model to be tailored to the current individual. The applications module lists several potential web services based on activity recognition results. These services include (but are not limited to) building users’ biometric database for individuals, displaying the user’s daily activity statistics, providing a life style logger and activity recognition augmented indoor localization service.

3.1 Method Details

Data Description and Feature Extraction. We use two types of raw data as the inputs of our ACRE application,

including the acceleration and the compass data. For the accelerometer data, we record acceleration along the x , y and z axis respectively, together with the corresponding time stamps. A vector A is used to represent one such reading in the following format: $A_i = \langle accX_i, accY_i, accZ_i, t_i \rangle$, where $-2.0 \leq accX_i, accY_i, accZ_i \leq 2.0$ in the unit of g -force and t_i is the time stamp. Sampling rate is set to 64Hz.

Based on the raw accelerometer reading data, we extract features in both time and frequency domains. For each of the three axis (x , y and z), we segment the data every four seconds. In the time domain, we compute the basic statistics of each axis (x , y and z), including the mean, the average, the maximum, the minimum and the standard deviation. In the frequency domain, we first perform Fast Fourier Transform (FFT) and compute an energy/magnitude array of the length 256 as $E_i = \sqrt{T_real_i^2 + T_image_i^2}$, $i = 0, 1, 2, \dots, 255$, where T_real_i and T_image_i are the real and imaginary parts of the FFT of a time series T . Different activities have distinctive frequency domain features. We compute the offset, the maximum, the standard deviation of the energy array. We also generate a histogram for the energy array, and calculate the top four peak frequencies.

Another type of raw data is the compass reading. In the iOS platform, the compass data is represented as float numbers between 0 and 360 and indicates the angle between current smartphone’s heading direction and the absolute north in the clockwise. We use a vector C to represent one compass reading: $C_i = \langle comp_i, t_i \rangle$, where $0 \leq comp_i \leq 360$ is the smartphone’s heading angle, and t_i is time stamp.

In our ACRE application, we use the compass reading for *turn detection*. To be specific, we created a direction dictionary: $Direction_Dict = [N : (330, 30), NE : (30, 60), E : (60, 120), ES : (120, 150), S : (120, 210), SW : (210, 240), W : (240, 300), WN : (300, 330)]$. Each key value entry in this dictionary defines the range of a specific direction. A significant change in the compass reading indicates the potential turn. For example, if the last average reading is 317.88 and the current average reading is 42.34, we will detect a right turn as the direction changes from WN to NE .

Model Building and Refinement. The basic activity recognition model is trained off-line using the data recorded from different users at different time inside a building. If we treat the activity recognition at every time window (4 seconds) separately, it can be naturally casted as a classification task with the actual activity (e.g., walking, jogging, etc) as the class label. Many classic (supervised) data mining algorithms (e.g., logistic regression, decision tree, multi-player perception, etc) can be plugged into ACRE. In the current deployment of ACRE, we chose decision tree as the basic classification model.

In order to maximally boost the recognition performance, we further impose a *smoothing* phase. Here, our key observation is that the adjacent activities of a user often follow a specific pattern. For example, almost every other activity happens after the *walking* activity. In order to capture such a temporal dependency, we apply hidden Markov model (HMM) on top of the classification model.

Finally (*personalization*), the motion patterns might differ from user to user. In order to make the model more personalized to a special user, we provide the feedback functionality in ACRE. If the user gives a positive feedback on certain prediction, the application rewards the correspond-

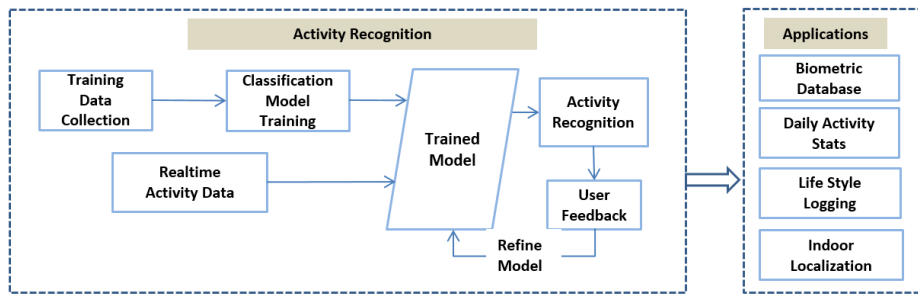


Figure 2: Application Work Flow Chart

ing learning parameter. In contrast, a penalty is imposed when a negative feedback is received.

3.2 Evaluation

We perform a systematic cross-validation evaluation on a benchmark data set [4]. The result is summarized in Fig. 3. Compared with the existing method in [4], our method is consistently better in terms of the recognition accuracy.

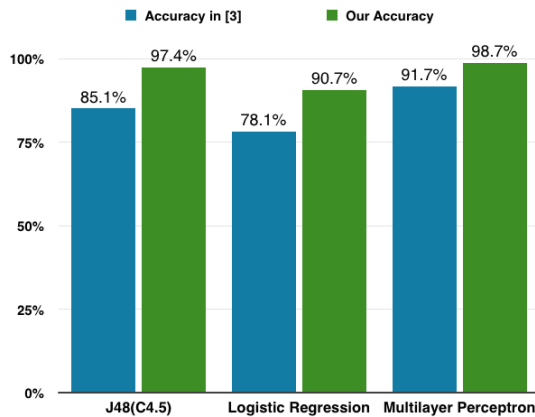


Figure 3: Cross Validation Result Comparison

4. RELATED WORK

Human activity recognition can be naturally casted as a (supervised) data mining problem. Kwapisz et al [4] used the standard *WEKA* tools for activity recognition with the accelerometer readings. Longstaff et al [5] leveraged active and semi-supervised learning to make the pre-trained recognition model adaptive to a specific user. Human activity recognition is a very powerful building block for many applications, including health and fitness monitoring [7], personal biometric signature [3], assistive technology [6], elderly-care [2], indoor localization and navigation[8], etc.

5. CONCLUSIONS

We present ACRE, a human activity recognition application on smartphone. ACRE takes the motion data from different sensors on smartphones as inputs and predicts a user's motion activities in real-time. It is able to recognize *eight* types of different activities in total with a much higher detection accuracy. It can be naturally personalized to a user's specific motion pattern by incorporating his/her feedback.

It also provides a daily activity summarization functionality. The application will be released soon in Apple's App Store. We will invite the audience to experiment with our ACRE in terms of its effectiveness, efficiency, applicability to various domains (e.g., smart-phone based navigation, traffic mode detection, etc) and the potential for further improvements (e.g., adaptive algorithms to accommodate different individuals as well as different body parts; multi-modality classification algorithm to incorporate other types of sensors, etc).

6. ACKNOWLEDGEMENT

This material is supported by the National Science Foundation under Grant No. IIS1017415, by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053, by Defense Advanced Research Projects Agency (DARPA) under Contract Number W911NF-11-C-0200 and W911NF-12-C-0028, and by Region II University Transportation Center under the project number 49997-33 25.

7. REFERENCES

- [1] Apple Document: <https://developer.apple.com>.
- [2] C. N. Doukas and I. Maglogiannis. Emergency fall incidents detection in assisted living environments utilizing motion, sound, and visual perceptual components. *Information Technology in Biomedicine, IEEE Transactions on*, 2011
- [3] J. R. Kwapisz, G. M. Weiss, and S. A. Moore. Cell phone-based biometric identification. In *BATS*, 2010.
- [4] J. R. Kwapisz, G. M. Weiss, and S. A. Moore. Activity recognition using cell phone accelerometers. *ACM SIGKDD Explorations Newsletter*, 2011.
- [5] B. Longstaff, S. Reddy, and D. Estrin. Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In *4th PervasiveHealth*, 2010.
- [6] D. Lymberopoulos, A. Bamis, and A. Savvides. Extracting spatiotemporal human activity patterns in assisted living using a home sensor network. *Universal Access in the Information Society*, 2011.
- [7] C. Seeger, A. Buchmann, and K. Van Laerhoven. myhealthassistant: a phone-based body sensor network that captures the wearer's exercises throughout the day. In *ICST*, 2011.
- [8] H. Ye, T. Gu, X. Zhu, J. Xu, X. Tao, J. Lu, and N. Jin. Ftrack: Infrastructure-free floor localization via mobile phone sensing. In *PerCom*, 2012.