

HHS Public Access

Author manuscript

Proc IEEE Int Conf Pervasive Comput Commun Workshops. Author manuscript; available in PMC 2018 May 04.

Published in final edited form as:

Proc IEEE Int Conf Pervasive Comput Commun Workshops. 2017 March ; 2017: .

ActivityAware: An App for Real-Time Daily Activity Level Monitoring on the Amulet Wrist-Worn Device

George Boateng, John A. Batsis, Ryan Halter, and David Kotz Dartmouth College

Abstract

Physical activity helps reduce the risk of cardiovascular disease, hypertension and obesity. The ability to monitor a person's daily activity level can inform self-management of physical activity and related interventions. For older adults with obesity, the importance of regular, physical activity is critical to reduce the risk of long-term disability. In this work, we present ActivityAware, an application on the Amulet wrist-worn device that measures daily activity levels (sedentary, moderate and vigorous) of individuals, continuously and in real-time. The app implements an activity-level detection model, continuously collects acceleration data on the Amulet, classifies the current activity level, updates the day's accumulated time spent at that activity level, logs the data for later analysis, and displays the results on the screen. We developed an activity-level detection model using a Support Vector Machine (SVM). We trained our classifiers using data from a user study, where subjects performed the following physical activities: sit, stand, lay down, walk and run. With 10-fold cross validation and leave-one-subject-out (LOSO) cross validation, we obtained preliminary results that suggest accuracies up to 98%, for n=14 subjects. Testing the ActivityAware app revealed a projected battery life of up to 4 weeks before needing to recharge. The results are promising, indicating that the app may be used for activity-level monitoring, and eventually for the development of interventions that could improve the health of individuals.

Keywords

physical activity; machine learning; mobile health (mHealth); wearables; support vector machine

I. Introduction

Physical inactivity increases the risk for cardiovascular disease and chronic diseases such as diabetes, hypertension and obesity [1]. Older adults with obesity who are sedentary are at higher risk of long-term disability and physical activity in this population is critical to reducing their risk of functional impairment. The American College of Sports Medicine (ACSM) recommends 150 minutes of moderate intensity activity each week for adults, including older adults [2]. Yet, a manner to unobtrusively track the amount of time spent doing moderate or vigorous activities are needed to enable this population to achieve this important health goal.

In this work, we built *ActivityAware*, an application that continuously monitors the activity level of individuals in real time using acceleration data recorded from an Amulet, a low-power wrist-worn device [3]. We developed a Support Vector Machine (SVM)-based

machine-learning model to detect the activity level of a person using data from a Dartmouth College Institutional Review Board (IRB) approved study. We collected acceleration data from younger, healthy volunteers who wore the Amulet as they performed various activities that could subsequently be adapted to an older adult population. We created an app for the Amulet that implements our activity-level-detection model, continuously records acceleration data, classifies the activity level of an individual, updates the day's accumulated time spent at that activity level, logs the data for later analysis, and then displays the results on its screen.

In the remainder of this work, we describe the Amulet platform on which *ActivityAware* runs and our approach to physical level categorization in Section II. We describe related work and the components of the *ActivityAware* app in Sections III and IV respectively. We describe our approach to developing the *ActivityAware* machine-learning model and the evaluation of the energy efficiency of the *ActivityAware* app in Sections V and VI respectively. We describe the limitations and future work in Section VII, and we conclude in Section VIII.

II. Background

In this section, we describe the Amulet platform on which the *ActivityAware* app runs and why it is suitable for running the app. Then, we describe the categorization of the physical activity levels we use in this work.

A. Amulet Wearable Device Platform

The Amulet is a hardware and software platform for writing energy- and memory-efficient sensing applications that achieve long battery life [3]. The Amulet hardware is a wrist-worn device that has two microcontrollers: an MSP430 running applications, and an nRF51822 for communicating with peripheral Bluetooth Low Energy (BLE) devices such as a heart-rate monitor and a galvanic skin response sensor (Figure 1). It has built-in sensors to measure acceleration, rotation, ambient sound, ambient light, and ambient temperature. The main board has two buttons, capacitive touch sensors, a battery, a haptic buzzer, two LEDs embedded in the case, a secondary storage board that holds a microSD card reader, and a display screen. The energy-efficient Amulet platform is useful for creating and running mHealth applications that monitor the physiological and behavioral health of its wearer, lasting weeks before needing to recharge.

B. Physical Activity Level Categorization

Physical activity levels are defined using the Compendium of Physical Activities, which capture the intensity of activities expressed in metabolic equivalents (METs): 1 MET corresponds the metabolic rate obtained during quiet sitting [4]. According to the Centers for Disease Control and Prevention (CDC) and the ACSM guidelines, activities can be categorized into low, moderate and vigorous based on METs [5]. Low corresponds to activities with METs less than 3 (e.g., sit, stand, lay down), moderate corresponds to activities with METs between 3 and 6 (e.g., walking at a moderate pace), and vigorous

corresponds to activities with METs greater than 6 (e.g., running) [5]. In this work, we use these example activities to characterize our activity levels.

III. Related work

There has been considerable research to classify activities using wrist-worn devices and other devices worn on other parts of the body. Also, there exist commercial fitness trackers that monitor the physical activity of wearers. We describe some of the related work here.

Manini et al. classified four classes of activity (ambulation, cycling, sedentary, and other) using a wrist-worn and ankle-worn device [6]. Using 13 features and an SVM classifier, they achieved accuracies of 95% for the ankle-worn data and 84.7% for the wrist-worn device with leave-one-subject-out cross-validation [6].

Maurer et al. implemented a real-time classifier (decision tree) on a custom-built device called the eWatch, which they placed on various parts of the body including the wrist, belt and pocket [7]. They classified six primary activities: sitting, standing, walking, ascending stairs, descending stairs and running. With six subjects, their classifier had 87% accuracy for the wrist-worn data using 5-fold cross-validation [7].

Fitbit[™] is a wrist-worn device that monitors several fitness parameters such as sleep, steps taken and activity level using data from an accelerometer, a gyroscope, and a heart-rate monitor (for some models). Fitbit calculates 'active minutes' when a person performs activities with METs above 3 – moderate-to-intense activities such as brisk walking, cardio workout and running [8]. Users can use this data to monitor their activity level over days, weeks, and months.

These works demonstrate that various activity groups can be classified using acceleration data from a wrist-worn device. The first two works focus on distinguishing between specific activities rather than activity levels. The real-time classifier developed by Maurer et al. implements a decision tree which tends to overfit the training data set. Although the training phase may reveal high accuracy results, the model might perform poorly on unseen data when deployed in a real-world system. Additionally, the work by Maurer et al. does not focus on tracking and displaying the amount of time spent performing each of the activities, which could be useful for the wearer of the eWatch. Although FitbitTM tracks activity levels and displays the results to users, it is a closed system that uses a proprietary algorithm. As a result it is not clear how active minutes are calculated, and the accuracy of the algorithms being used is unknown.

There is a need for a system that monitors activity levels of users, has high accuracy, uses an approach that can be validated by others, and has long battery life. This system should be able to be modified and used by researchers who might be interested in monitoring the physical activity of various populations. Also, this system needs to have a battery life measured in weeks rather than hours and days. We present a system that satisfy's these constraints and addresses the shortcomings of earlier work.

IV. Solution: ActivityAware

ActivtyAware is an Amulet application that measures the daily activity levels of individuals (sedentary, moderate and vigorous). The app continuously collects acceleration data, classifies the activity level, updates the day's accumulated time spent at that activity level, logs the data for later analysis, and displays the results on the screen as feedback to the wearer. The app consists of four components: data collector, activity-level detector, activity-level monitor, and activity-level display (Figure 2).

A. Data Collector

The data collector samples data from a 3-axis accelerometer at a frequency of 20 Hz, and parses the data stream into 5-second windows. Previous studies have shown that a frequency of 20Hz is sufficient for capturing the frequency range of physical human activities for classifying activities [7].

B. Activity-Level Detector

The activity-level detector determines the activity level of the user. It computes a vector of features from each 5-second window of accelerometer data: mean and standard deviation of four values (individual *x*, *y*, *z* accelerations, and magnitude of the acceleration). This 8-feature vector is then fed to the activity-level classifier, which is an implementation of the decision function of a Linear SVM:

y = wx + b

Here, y is the vector that holds the result of the evaluation for the three activity levels, x is the computed feature vector of size: *number of features*, w is the coefficient matrix of size: *number of classes* × *number of features* and b is the intercept vector of size: *number of classes*. The values for w and b are obtained from the linear model that we train offline using the scikit-learn library (we describe the training of this model in Section V). Because this is a multi-class classification, we implemented the "one-vs-the-rest" approach for multi-class classification since the scikit-learn Linear SVM function uses this method [9]. In this approach, one classifier is trained for each of the classes that correspond to each row in the matrix w. The result of solving the equation is a vector y that contains a value for each of the three classes. The class with the maximum value is the predicted class.

C. Activity-Level Monitor

The activity-level monitor is responsible for keeping track of the number of minutes spent per day, for each of the three activity-level categories. This component tracks three data points for each activity level: total minutes today, mean minutes over all days, and total minutes over all days. (Here, all days refers to the set of days since the app was started.) The value for each of these data points is updated after each classification result, and the total minutes today is reset at midnight each day.

D. Activity-Level Display

The activity-level display component displays information about the amount of time spent for each activity on the Amulet screen computed by the activity-level monitor. The display presents the amount of time in minutes spent for each of the activity levels for that day (first row) and the total minutes over the past days (second row) (Figure 3). In the future, we plan to display graphs of the activity levels to highlight trends over the past days and explore other ways the data should be displayed based on user feedback.

V. Activity-Level Detection Model - Machine Learning Training Offline

SVM is a classifier that constructs a high-dimensional hyper-plane and uses it to perform classification [9]. SVM chooses a hyper-plane that maximizes distance to the nearest points on the either side of the plane for the binary classification case (Figure 4). We use SVM because it uses a subset of the training set – support vectors – for its prediction function. Models like k-nearest neighbor (kNN), on the other hand, need to store all the data points in memory for prediction. SVM is more memory efficient and thus well suited for low-memory platforms like the Amulet. We trained three SVM models: Linear SVM, Polynomial (with degree 1) SVM and Radial Basis Function (RBF) SVM, using the scikit-learn library [9] to distinguish low, moderate, and vigorous activity levels. We use scikit-learn's default parameters for the SVM models.

A. Data Collection

We collected data from volunteer subjects under a study protocol approved by Dartmouth's IRB. All individuals completed a basic baseline demographic questionnaire that assessed age, gender, race, height, weight and handedness (left or right). All data was collected online via Research Electronic Data Capture software (REDCap) into a centralized, HIPAA compliant repository.

1) Activity Data Collection App—We developed an app similar to *ActivityAware* for the purpose of collecting data from the study. The app has three states: Ecological Momentary Assessment (EMA), Data Collection, and Data Logging (Figure 5).

The app begins in the EMA state. Within this state, the user selects which activity they are about to perform from a list of activities using capacitive touch sensors on the Amulet (Figure 6). After the user selects the specific activity and presses the button on the Amulet, the app switches to the data-collection state.

In this state, the app collects and stores acceleration data from a 3-axis accelerometer with range $\pm 2g$ at a frequency of 20 Hz. We discard the first 5 seconds of data. After 2 minutes, the app switches to the data-logging state in which it logs the collected acceleration data along with the activity level onto a micro-SD card on the Amulet.

The app then switches to the EMA mode to allow the user to select the next activity to perform. We accompanied the subject when they performed the activities so we could ensure they completed all activities correctly and the appropriate number of times.

2) Study Protocol—We collected acceleration data from 14 subjects (n=14) as they performed various physical activities. The subjects were college students 18–23 years old. Subjects wore the Amulet on their left wrist and performed each of the following activities for 10 minutes as the Amulet ran the Activity Data Collection App: sit, stand, lay down, walk at a regular pace, and run (Figure 7).

We had 50 minutes of data from each subject resulting in 700 minutes of data total. We categorized the data from these 5 activities into the following classes: low (sit, stand and lay down); moderate (walk); and vigorous (run). We then split the data into 5-second non-overlapping time windows that previous studies have shown to be suitable for activity classification [6].

B. Feature Extraction

From each 5-second window of each subject's data, we extracted eight time-based features that previous studies have shown to be relevant for activity detection: mean and standard deviation of each of four values: acceleration (x, y, z axes), and magnitude of the acceleration vector [6][7]. The result was a training dataset containing a total of 8,362 feature vectors.

C. Training /Classification

We trained three SVM models – Linear, Polynomial and RBF SVM – each of which classified the data into three activity levels: low, moderate and vigorous. We ran experiments to test these three classifiers.

D. Testing Results

We ran 10-fold cross-validation and leave-one-subject-out (LOSO) cross-validation on the data from the study using Linear, Polynomial, and RBF SVM.

Overall, Polynomial SVM performed the best with accuracies of 98.3% and 98.1% for 10fold and LOSO cross-validation respectively (Table 1). Linear SVM performed second best, with accuracies of 95.7% and 95.6% for 10-fold and LOSO cross-validation respectively, and RBF performed worst with accuracies of 59.4% and 59.4% for 10-fold and LOSO crossvalidation respectively. Tuning the hyper-parameters for RBF could improve its accuracy.

VI. Energy Efficiency of ActivityAware

We tested the energy efficiency of the *ActivityAware* app by running it for 7 days (168 hours) as it computed activity levels continuously throughout the duration. We recorded the battery voltage level each hour over the 7-day period.

The graph of the battery level shows battery percentage as the y-axis and time (hours) as the x-axis. The battery level dropped linearly from 100% to 73% over the 7-day period (Figure 8). From a linear extrapolation of this battery-life data, we forecast that the app can run for approximately 4 weeks (26 days) before the Amulet needs to be recharged, which should be

convenient for most users. This result demonstrates that *ActivityAware* is sufficiently energy efficient.

VII. Limitations and Future Work

Our experiments have several limitations that imply the need for additional investigation. We used only a limited number of time-based features and also no frequency-based features for training the classifiers, because the latter are computationally intensive to derive. Computing additional temporal and frequency-based features could produce better models. We additionally could test various subsets of the features to obtain a feature set which is small, less computationally intensive, and yet produces high accuracy.

We implemented a Linear SVM model in the *ActivityAware* app rather than a Polynomial SVM model. While Polynomial SVM had the highest accuracy, it is more computationally intensive and requires more memory to store all the support vectors. Implementing a Polynomial SVM model could lead to more accurate real-time predictions. Also, since we only used the default setting while training the Linear SVM, we could tune the hyper-parameters with a grid search to obtain accuracies comparable to or even better than the Polynomial SVM, in which case there will be no need to implement a Polynomial SVM model in the *ActivityAware* app.

We did not explore the best way to display the activity level information to users. We need to design varied interfaces and get user feedback to determine the best interface to use, and to select which display makes the most impact with reference to personal physical activity level monitoring.

Finally, we have not yet trained or tested our model with subjects in our target population, obese elderly people; we need to validate our preliminary results with data from obese elderly people and with larger populations before final conclusions can be drawn.

VIII. Conclusion

In this work, we present *ActivityAware*, an application on the Amulet wearable platform to measure the activity levels of individuals continuously and in real time. The app continuously collects acceleration data on the Amulet, classifies the activity level of an individual, updates the day's accumulated time spent at that activity level, logs the data for later analysis, and displays the results on the screen.

Our results show accuracies of 59.4% for RBF SVM, 95.7% for Linear SVM and 98.3% for Polynomial SVM with 10-fold cross validation, and accuracies of 59.4% for RBF SVM, 95.6% for Linear SVM and 98.1% for Polynomial SVM with LOSO cross-validation. Testing the *ActivityAware* app revealed a projected battery life of up to 4 weeks when running the Linear SVM classifier.

The apparent accuracy and energy efficiency of our app show that *ActivityAware* has the potential to be used for activity-level measurement and monitoring and could eventually

inform the development of intervention and personal physical activity management that could improve the health of elderly individuals.

Acknowledgments

We are grateful to Dr. Summer Cook for her input on the activity-level categorization. We thank Alexandra Zagaria for assisting with setting up the user study. We thank the members of the Amulet team, including Ron Peterson, Taylor Hardin, Josiah Hester, Joseph Skinner, and Steven Hearndon. This research is supported by the National Science Foundation under award numbers CNS-1314281, CNS-1314342, CNS-1619970, and CNS-1619950. Dr. Batsis' effort was supported in part by the National Institute On Aging under NIH Award Number K23AG051681. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation or the National Institutes of Health.

References

- 1. Warburton, Darren ER, Nicol Crystal Whitney, Bredin Shannon SD. Health benefits of physical activity: the evidence. Canadian Medical Association Journal. 2006; 174:6. 801–809.
- 2. American College of Sports Medicine (ACSM). News Releases. Web. 16 Nov. 2016 http:// www.acsm.org/about-acsm/media-room/news-releases/2011/08/01/acsm-issues-newrecommendations-on-quantity-and-quality-of-exercise
- 3. Hester, Josiah, Peters, Travis, Yun, Tianlong, Peterson, Ronald, Skinner, Joseph, Golla, Bhargav, Storer, Kevin, Hearndon, Steven, Freeman, Kevin, Lord, Sarah, Halter, Ryan, Kotz, David, Sorber, Jacob. Proceedings of the ACM Conference on Embedded Networked Sensor Systems (SenSys). ACM Press; Nov. 2016 Amulet: An energy-efficient, multi-application wearable platform.
- Ainsworth, Barbara E., et al. Compendium of physical activities: an update of activity codes and MET intensities. Medicine and Science in Sports and Exercise. 2000; 32.9(SUPP/1):S498–S504. [PubMed: 10993420]
- 5. U.S. Department of Health and Human Services, Public Health Service, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Division of Nutrition and Physical Activity. Promoting physical activity: a guide for community action. Champaign, IL: Human Kinetics; 1999.
- Mannini, Andrea, Stephen S, Intille, Mary Rosenberger, Angelo M, Sabatini, William Haskell. Activity recognition using a single accelerometer placed at the wrist or ankle. Medicine and Science in Sports and Exercise. 2013; 45:11. 2193.
- Maurer, Uwe, Asim, Smailagic, Daniel, P., Siewiorek, Michael, Deisher. International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06). IEEE; 2006. Activity recognition and monitoring using multiple sensors on different body positions.
- 8. Fitbit Official Site for Activity Trackers & More. Web. 12 Jan. 2017. http://www.fitbit.com/
- 9. Pedregosa, Fabian, et al. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research. Oct.2011 12:2825–2830.

Author Manuscript



Figure 1.

Fully assembled Amulet device (left), Internal Amulet peripherals (middle), custom Amulet circuit board (right)



Figure 2. Components of *ActivityAware* App







Figure 4. Hyper-pane separating two classes in SVM



Figure 5.

States of Activity Data Collection App









Plots of acceleration data from one subject for activities: standing, laying down, sitting, walking and running



Figure 8. Graph of Battery Life over 7 days



Table 1

Accuracy of SVM models using 10-fold and LOSO Cross Validation

	Polynomial SVM	Linear SVM	RBF SVM
10-fold	98.3%	95.7%	59.4%
LOSO	98.1%	95.6%	59.4%