

# Inactivity Recognition: Separating Moving Phones from Stationary Users

James Reinebold  
University of Southern California  
reinebol@usc.edu

Harshvardhan Vathsangam  
University of Southern California  
vathsang@usc.edu

Gaurav S. Sukhatme  
University of Southern California  
gaurav@usc.edu

## ABSTRACT

Accurate methods of detecting whether a person is at rest form an important component in indoor localization and sedentary lifestyle monitoring. The problem of quantifying rest is complicated by the variety of activities and phone configurations that exist even when the user location is stationary. Our study examines whether on-phone kinematic sensors can be used to accurately and consistently detect rest. Rest is defined as a user's absolute positioning with respect to a world coordinate frame not changing significantly over a fixed time interval. An important requirement in our approach is that the algorithm maintains its accuracy independent of orientation and on-body location. The techniques examined show high accuracy classification (>95%) with test participants simulating typical everyday tasks in an office environment. An important contribution of our approach is showing that rest detection accuracy improved when accounting for the orientation of the phone for the activities discussed.

## 1. INTRODUCTION

Detecting when the user is at rest has important applications in localization [1], gaming, and health monitoring [2]. We define "at rest" to be when the user is standing (or sitting) in a fixed position with respect to the world (specifically for this study taken to be not moving more than a meter in a three second window). It is important to note that even when the user is at rest they may be actively using the phone: interacting with it for applications or switching its position relative to their body. An algorithm that claims to accurately detect rest must be robust to these variations.

Work in rest detection has fallen under the broader field of activity recognition - using mobile sensors and machine-learning techniques to recognize aspects of human motion. Previous studies have enabled highly accurate classification of divergent activities such as folding laundry or running a vacuum cleaner [3-6]. However, unlike our approach, these studies assume a constant (or at least known) location for the mobile phone in relation to the body of the user during the course of the experiments.

Similarly, work has been done on localization using sensors typically available on most phones (GPS, Wi-Fi networks, GSM) [7],[8]. These systems rely, at least partially, on signals transmitted over radio waves from known fixed-point locations. However, relying on the constancy of these signals can be problematic. For example, GPS connection

can be lost inside buildings or in the "urban canyons" of major cities [9]. GPS and Wi-Fi are also more power-hungry than other, internal sensors [10]. To avoid these pitfalls, our approach uses only on-phone kinematic sensors.

Although in theory localization could be accomplished by integrating information from accelerometers over time given an initial starting position ("dead reckoning"), in practice sensor noise corrupts the calculations and such methods are inaccurate given a significant length of time [11]. However, solving the simpler problem of determining whether or not the user holding a phone is moving may be valuable information in its own right. If this could be determined with a high degree of accuracy without knowing in advance where the phone is stored on a person's body, it could assist other, more complicated, localization schemes.

Thiagarajan et al. [1] tried similar strategies using acceleration to detect movement in cars but relied on preset thresholds and assumed a constant location for the mobile device. Similarly, Wang et al. [12] detected movement but did not integrate gyroscope data and relied on empirical thresholds. Data driven techniques allow us to operate in non-linear spaces thus permitting flexibility in threshold design.

Our approach to rest detection treats rest as a binary classification problem in the presence of non-rest data. Our study applies the established pattern developed for activity recognition: sampling sensor hardware, extracting features, and using statistical machine learning algorithms to classify unknown data points [3]. However, we expand on these methods with two main contributions: our studies examine in detail which features are most relevant for rest detection and show how correcting for phone orientation improves accuracy. Our techniques do not require a fixed location or orientation of the phone on the user's person. We also pay particular interest to the problem of accurately differentiating rest from the activity of walking (we choose to focus on walking as it is the most typical form of human movement in office environments).

In Section 2 we will cover the design, noting the features used (Section 2.2) and how frame the sensor values reported by the phone into a global frame of reference to provide useful training data for the machine learning algorithms (Section 2.3). We then present the results of a user study in Section 3 and conclude with a summary of our techniques and discuss potential areas of improvement.

## 2. DESIGN

To classify rest, a systematic way of sampling the kinematic sensors, extracting relevant features, and applying these features as inputs to machine learning algorithms is needed. Varying which features are trained on, whether or not rotational correction is performed, and what machine learning algorithms are used can all affect the final performance of the system. Our paper tests along these axes.

### 2.1 Hardware Sensors

For this experiment we used a standard Nexus-S phone equipped with the Android operating system. The custom designed `MovementTracker` App (<https://github.com/mobilesensing-usc/MovementTracker>) was used to record data from two kinds of kinematic sensors: accelerometers reporting triaxial accelerations ( $\text{m}/\text{sec}^2$ ) and triaxial rotational gyroscopes reporting angular speeds ( $\text{radians}/\text{sec}$ ). All sensors recorded these values to text files at the fastest possible sampling frequency permitted by the Android Sensor API [13] (approximately 35 Hz for the accelerometers and 800 Hz for the gyroscopes).

### 2.2 Feature Extraction

Feature extraction replaces raw, potentially noisy data with statistically meaningful aggregations across time intervals. In our study, training features were extracted from the text logs across a sliding window of size three seconds with a one second overlap between consecutive windows. We used the following 16 features to describe phone movement:

- Accelerometer Power (as in [4])
- Accelerometer Means (along X, Y, Z axes)
- Accelerometer Variances (along X, Y, Z axes)
- Gyroscope Means (along X, Y, Z axes)
- Gyroscope Variances (along X, Y, Z axes)
- Covariance between acceleration and gyroscope rotation rates (along X, Y, Z axes)


These features were chosen for their ease of implementation and the fact that they can be computed in  $O(n)$  time and have been used before with success for activity recognition [3],[5]. The features were further grouped as: accelerometer power only (referred to as "Power Only"), accelerometer power and covariance between acceleration and rotation rates (referred to as "Partial"), and the entire set of sixteen (referred to as "Full").

### 2.3 Sensor Data Coordinate Transformation

A distinguishing aspect in our approach is the use of "world rotated features" to describe and characterize rest. Conventional activity recognition algorithms use sensor readings that are normally measured in the local coordinate system of the phone [1], [3-5], [12].

The Android API fuses accelerometer and gyroscope information to return an orientation quaternion of the phone:

$$Q(X,Y,Z, \theta) = [X \cdot \sin(\theta/2), Y \cdot \sin(\theta/2), Z \cdot \sin(\theta/2), \cos(\theta/2)]$$



**Figure 1:** The left side of the diagram shows the orientation of the locally framed axis for the phone. The right side of the diagram shows the orientation of the globally framed axis relative to the earth. Our approach compares the effect of training models in the global and local reference models on detection accuracy (Image source: <http://developer.android.com/>)

Where X, Y and Z are the direction cosines of the axis of rotation and  $\theta$  specifies an angle of rotation about that axis. Knowing the orientation quaternion allows us to rotate the triaxial accelerometer and gyroscope sensor streams from the local coordinate frame of the phone to a global coordinate frame [14]. At this point, sensor readings are said to be "corrected" for phone orientation. Using a global coordinate frame ensures that sensor readings corresponding to a particular axis remain so irrespective of phone orientation. As such, local repetitive movements (as with circular motions of the device) can be distinguished from movements associated with location changes.

### 2.4 Experiment Setup

Data collection was divided into three kinds of trials, grouped by type of movement: constant movement, constant stationary behavior, and a mixture of both movement and stationary behavior.

The median age of the eight test participants was 27.5 years with a standard deviation of  $\pm 6.00$ . Participants had a median body weight of 70.40 kg (standard deviation of  $\pm 11.77$ ) and a median height of 1.79 meters (standard deviation of  $\pm 0.09$ ). All test participants responded that they regularly used mobile phones (although not necessarily Android devices).

#### 2.4.1 Trial 1: Constant Movement

The aim of this trial was to test the accuracy of our algorithm in scenarios where the user is always moving. Test subjects were given the phone and told to walk around the USC campus for five minutes. The subjects were not given explicit instructions on how to carry the phone (during the experiment we observed some subjects holding the phone in their hand and others who kept the phone in a pocket). Ground truth was taken to be the moving state for all data in this trial.

#### 2.4.2 Trial 2: Constant Stationary Behavior

The aim of this trial was to test the accuracy of our algorithm in situations where the user is always at rest. Test subjects were given the phone and told to not move outside

of a one-meter radius for five minutes. While inside the circle, they were instructed to complete various tasks that involved small movements of the phone. Example tasks included using the phone's calculator App to solve a simple math expression, standing up, reorienting the phone towards an object in the room to take a picture, and putting the phone inside (and later removing it from) a drawer. The presence of these tasks ensured that the users would not keep the phone still for the duration of the experiment and produced motions similar to those encountered while using mobile phones for gaming or office work. Ground truth was taken to be the stationary state for all data in this trial.

### 2.4.3 Trial 3: Mixture of Behaviors

The aim of this trial was to test the accuracy of our algorithm in situations involving a mixture of activities typical of daily lifestyles. Test subjects were given the phone and told to complete a list of tasks in five minutes. In this trial, the tasks involved both walking small distances (down a hallway and back) and using the phone to answer questions on the survey as in the second trial. Video recordings made of the test subjects during the trial were used to annotate the ground truth of the data collected as belonging to either the stationary (at rest) or moving (walking) sets. For sliding windows that spanned both classifications (i.e. took place during transitions between the two states), a majority vote of readings taken was used to label the data.

## 3. RESULTS

Data from each of these trials formed the input to classification algorithms. Classification of features as either stationary or moving was implemented the open source machine learning toolkit Weka [15]. Weka includes standard algorithms for k-nearest neighbors, support vector machines, J48 decision trees, and Naive Bayes learning. With the exception of selecting  $k=5$  for kNN, default parameters were used for each of the algorithms. This was because the emphasis was more on finding the right feature spaces for the algorithms and not the algorithms themselves.

Results were evaluated with respect to two categories of user behaviors: classification and training from constant behaviors (the first two trials) and from mixed behaviors (the third trial). Leave-one-out cross validation was used to generate the confusion matrices. Data was collected from a total of eight volunteers for the first two trials and seven volunteers for the third trials (one subject's third trial log had corrupted data and could not be used).

### 3.1 Classifiers Trained on Constant Behaviors

#### 3.1.1 Classification Accuracies


Classification was achieved with a total accuracy of roughly 97.41% for kNN across the subject pool with a sliding window size of length three seconds with a one second overlap. The training was performed with the full set of sixteen globally referenced features. Of the 120 points classified incorrectly, a total of 94 of these occurred in consecutive temporal groups of size  $\geq 2$  (78.33%). A total

of 56 points out of those classified incorrectly occurred in groups of size  $\geq 3$  (46.67%). The next best performing algorithm was SVM with a classification accuracy of 96.68%.

		Truth	Truth
		Stationary	Moving
Prediction	Stationary	2274	73
Prediction	Moving	47	2246

**Table 1: Confusion Matrix for Constant Behaviors (shown for kNN). Each value in the confusion matrix represents one window of features that was assigned as either stationary or moving by the algorithm. All points from Trial 1 were assigned a ground truth of moving and all points from Trial 2 were assigned a ground truth of stationary. Data from both trials was used for this confusion matrix.**

#### 3.1.2 Effect of Different Feature Sets



**Figure 2: The relative accuracy ratings of using only power as a feature compared to a partial feature set of power and covariance between accelerometer and angular rotation speed and using all sixteen features noted in Section 2.2. Using additional features helped the algorithms separate between stationary and moving behaviors.**


Figure 2 illustrates the effect of different feature sets on classification accuracy. Using all sixteen features outperformed using just accelerometer power by as much as 5% for kNN. Using the partial feature set (as defined in Section 2.2) performed somewhere between the full feature set and using just accelerometer power.

One possible explanation for this result is that some movements still associated with rest (i.e. putting the phone in a drawer or moving it around in the air while repositioning it) can actually generate accelerations of sufficient magnitude to confuse them with walking. Adding in covariance between the angular rotation velocities provides additional insight on how the movement is occurring. Surprisingly, adding more features hurt the Naive Bayes classifier's performance (possibly due to overfitting).

#### 3.1.3 Effect of Coordinate Transformation to Global Coordinate Frame

Figure 3 illustrates the effect of rotation to a global frame on classification accuracy. Transforming to a global coordinate

frame resulted in over an 8% increase for some of the algorithms.



**Figure 3: Classification accuracies for the machine learning algorithms when trained on features framed locally versus those framed globally. Training on globally framed features resulted in more accurate classification.**

This implies that while maintaining the orientation of the device with respect to global coordinates may not always work for accurate position estimation, it is still capable of determining if displacement occurs. Rotating to a global frame of reference for the sensors accounts for rotational changes in sensor streams. For example, when a phone is rotated by 90 degrees about an axis, with respect to the local frame of reference, axes of sensor streams will be switched to another axis. Accounting for this rotation ensures that sensors streams always map to the same axis of rotation.

### 3.2 Classifiers Trained on Mixed Behaviors

#### 3.2.1 Classification Accuracies

		Truth	Truth
		Stationary	Moving
Prediction	Stationary	909	33
Prediction	Moving	46	1042

**Table 2: Confusion Matrix for Mixed Behaviors (shown for kNN). Each value of the confusion matrix represents one window of features that was assigned as either stationary or moving by the algorithm. Ground truth was obtained from annotating video recordings of the subjects as they completed the tasks.**


Transitions were handled without much loss of precision, resulting in a total accuracy of roughly 96.11% for kNN when using the full feature set. Of the 79 points classified incorrectly, a total of 59 of these occurred in consecutive temporal groups of

size  $\geq 2$  (74.68%). A total of 43 points out of the those classified incorrectly occurred in groups of size  $\geq 3$  (54.3%). SVM out-performs kNN with a total accuracy of 96.40% when trained on all sixteen features.

Additional features once again improved the performance of algorithms, as illustrated in Figure 4. The improvement was larger than in the constant behavior case (roughly 10% for

decision trees and kNN). As in the previous section, the full feature set performs the best, followed by accelerometer and covariance, with using only accelerometer power performing the worst.

#### 3.2.2 Effect of Different Feature Sets




**Figure 4: The relative accuracy ratings of using only power as a feature compared to a partial feature set of power and covariance between accelerometer and rotation speed and using all sixteen features noted in Section 2.2. The additional features helped the machine learning algorithms overcome the noisy data of mixed behaviors.**

Adding in additional training features helped the algorithms more for mixed behaviors with transitions between the two states than when the behaviors were constant throughout data recording. Data with transitions is noisier and has more periods where the user was at rest as per our definition, but still moving in some way (such as when they are sitting down or standing up). The additions to the feature vector helped to overcome these complications by providing additional descriptive insight on how the motion occurred.

#### 3.2.3 Effect of Coordinate Transformation to Global Coordinate Frame

Figure 5 illustrates the effect of rotation to a global frame on classification accuracy. The difference between locally and globally referenced features is more pronounced for the mixed behavior trial. Rotating to a global frame of reference provides additional insight on whether or not the accelerations are being applied in world space.



**Figure 5: Classification accuracies for the machine learning algorithms when trained on locally versus globally framed features when the data included transitions. Once again using globally framed coordinates aided classification.**

## 4. CONCLUSION

We have shown that recognizing whether or not the user is moving or not can be done with high accuracy (>95%) using only kinematic sensors from a single mobile phone that was not kept in a constant location during the tests. Furthermore, we have demonstrated it in a semi-naturalistic environment (with transitions between rest and non-rest) representative of daily lifestyles of everyday users. By doing so, we have identified the optimal features for high accuracy and also underscored the usefulness of rotating to a global frame of reference in activity recognition.

It should be noted that all approaches used in this study, including using only accelerometer power, give usable classification rates. However, for some applications higher classification rates might be necessary. In particular, if data from the accelerometers and gyroscopes implies that motion is not occurring, then there would be no reason to continually check the GPS or Wi-Fi sensors to determine if the user is moving (thus saving power).

Although in this experiment all analysis was done via post-processing data collected from the mobile phones, the logical next step is to integrate the data collection and machine learning algorithms on the phone hardware itself. The algorithm presented would still function as described in a real-time setting, the only main difference is that instead of writing the sensor readings to file they would instead be stored in memory with classification decisions occurring at the end of every window. Knowing whether or not the person holding a mobile device is at rest will enable richer applications to be developed with diverse goals from documenting sedentary lifestyles [2] to being built into indoor localization schemes [1].

## 5. ACKNOWLEDGEMENTS

We would like to thank Ankit Sharma for his contributions on the MovementTrackr Android App.

This project was funded by NSF (CCR-0120778) as part of the Center for Embedded Network Sensing (CENS). Support for H Vathsangam was provided by the Annenberg Graduate Fellowship Program.

## 6. REFERENCES

[1] Thiagarajan et al. 2011. Accurate, Low-Energy Trajectory Mapping for Mobile Devices. In *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation*.

[2] Berke et al. 2011. Objective Measurement of Sociability and Activity: Mobile Sensing in the Community. *Annals of Family Medicine*. Volume 9. 344-350.

[3] Bao, L., Intille, S. 2004. Activity recognition from user- annotated acceleration data. In *Proceedings of the*

*2nd International Conference on Pervasive Computing*, 1-17.

[4] Lester, J., Choudhury, T. and Borriello, G. 2006. A Practical Approach to Recognizing Physical Activities. In *Proceedings of the Fourth International Conference on Pervasive Computing*.

[5] Ravi et al. 2005. Activity Recognition from Accelerometer Data. In *Proceedings of the Seventeenth Conference on Innovative Applications of Artificial Intelligence*. 1541-1546.

[6] Miluzzo et al. 2008. Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the *CenceMe* Application. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*.

[7] Vathsangam, H., Tulsyan, A., and Sukhatme, G. 2011. A Data-driven Movement Model for Single Cellphone-based Indoor Positioning. In *Body Sensor Networks*.

[8] Constandache, I., Choudhury, R., and Rhee, I. 2010. Towards mobile phone localization without wardriving. In *Proceedings of the 29th conference on Information communications*.

[9] Cui, Y. and Ge. S. 2003. Autonomous Vehicle Position in Urban Canyon Environments. *IEEE Transactions on Robotics and Automation*. Volume 19, Issue 1.

[10] Abdesslem, F., Phillips, A., and Henderson, T. 2009. Less is more: energy-efficient mobile sensing with senseless. In *Proceedings of the 1st ACM workshop on Networking, systems, and applications for mobile handhelds*.

[11] Woodman, O. 2007. An Introduction to Inertial Navigation. Technical Report. University of Cambridge.

[12] Wang, et al. 2009. A Framework of Energy Efficient Mobile Sensing for Automatic Human State Recognition. *ACM Conference on Mobile Systems, Applications, and Services*.

[13] Android Sensor Event: <http://developer.android.com/reference/android/hardware/SensorEvent.html>

[14] Vicci, L. 2001. Quaternions and Rotations in 3-Space: The Algebra and its Geometric Interpretation. Technical Report. University of North Carolina at Chapel Hill.

[15] Hall, et al. 2009. The WEKA Data Mining Software: An Update. *SIGKDD Explorations*. Volume 11, Issue 1.