Activity Recognition Using Cell Phones

Rohit Talreja, Travis Geis, Sanjay Srinivas
CS229 Final Project Fall 2015

Introduction

Modern cell phones have sophisticated onboard positional sensors that provide real-time information about the user's motion. Applying machine learning to these motion data, it is possible to classify the current activity of the user. Accurate activity classification can provide context clues to mobile applications and make them more useful and less prone to interrupting the user.

We experimented with different classification algorithms to assess their accuracy in classifying the user's current activity.

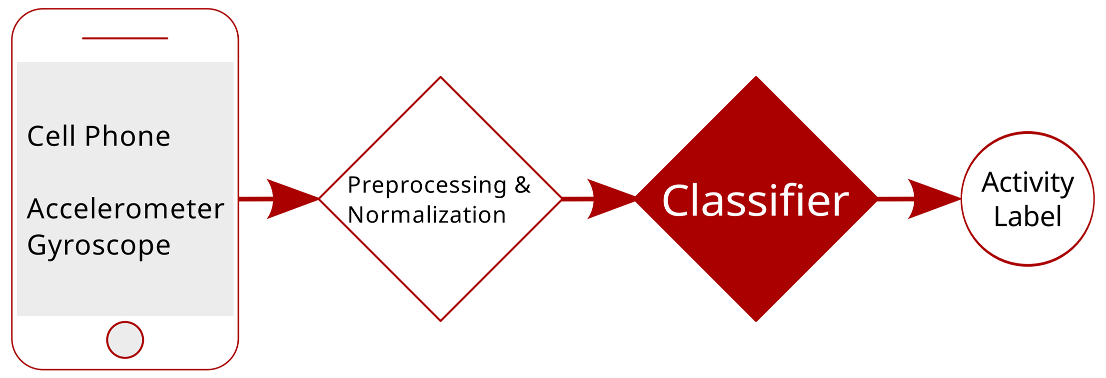

To train and test our classification algorithms, we used an existing dataset containing motion data [1]. The data were collected using a hip-mounted Samsung Galaxy SII phone. 30 participants performed each of 6 activities:

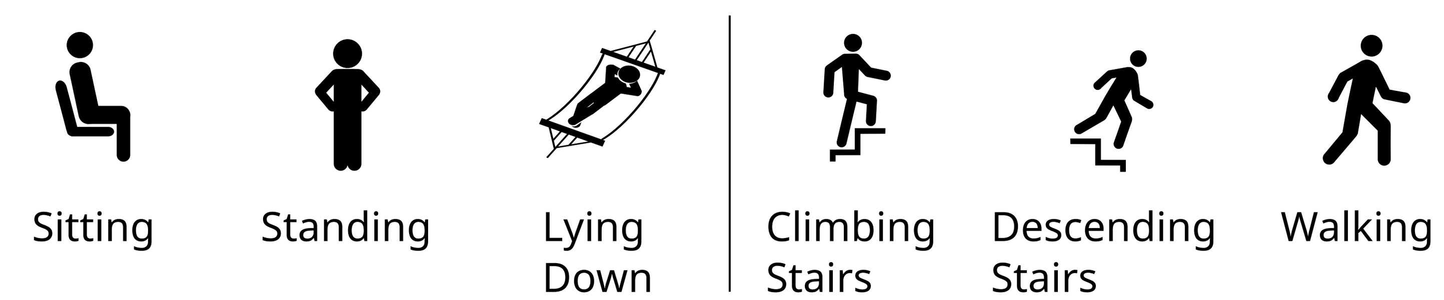

The dataset contains 7352 training examples and 2947 test examples, roughly a 70%-30% split. Each example has 561 features derived from accelerometer and gyroscope data. These include means, peak values, derivatives, standard deviations, and various frequency-space transformations of the original data. Participants were randomly assigned to contribute either to the test or train data set.

Binary SVM Classification

We began our analysis on the dataset by applying a binary SVM classifier [2] to each pair of activities to see which activity pairs were easily distinguishable from each other and which were more difficult to distinguish. This also gave us a baseline accuracy metric for assessing the multi-class SVM algorithm.

As predicted, there was high accuracy when distinguishing between a motion activity and a sedentary activity (eg. walking vs. standing) but a much lower accuracy when distinguishing between two motion activities (walking vs. walking upstairs) or two sedentary activities (sitting vs. standing). In the last example, the accuracy was only marginally above chance.

Binary Classification Accuracies

	Walking Upstairs	Walking Downstairs	Sitting	Standing	Lying Down
Walking	57.9%	77.1%	99.6%	99.6%	99.3%
Walking Upstairs		76.7%	98.4%	98.6%	97.8%
Walking Downstairs			98.8%	98.9%	98.6%
Sitting				53.9%	67.3%
Standing					67%

Multi-Class SVM Automatic Feature Selection

To classify an activity as one of the 6 possible activities, we used a multi-class SVM provided by libSVM [3]. We used a forward feature search over the training set to find the most useful features for the multi-class SVM. Acceleration-based features tended to be more useful, and the learning rate dropped off dramatically after 15 features.

Learning Curve for Multi-Class SVM

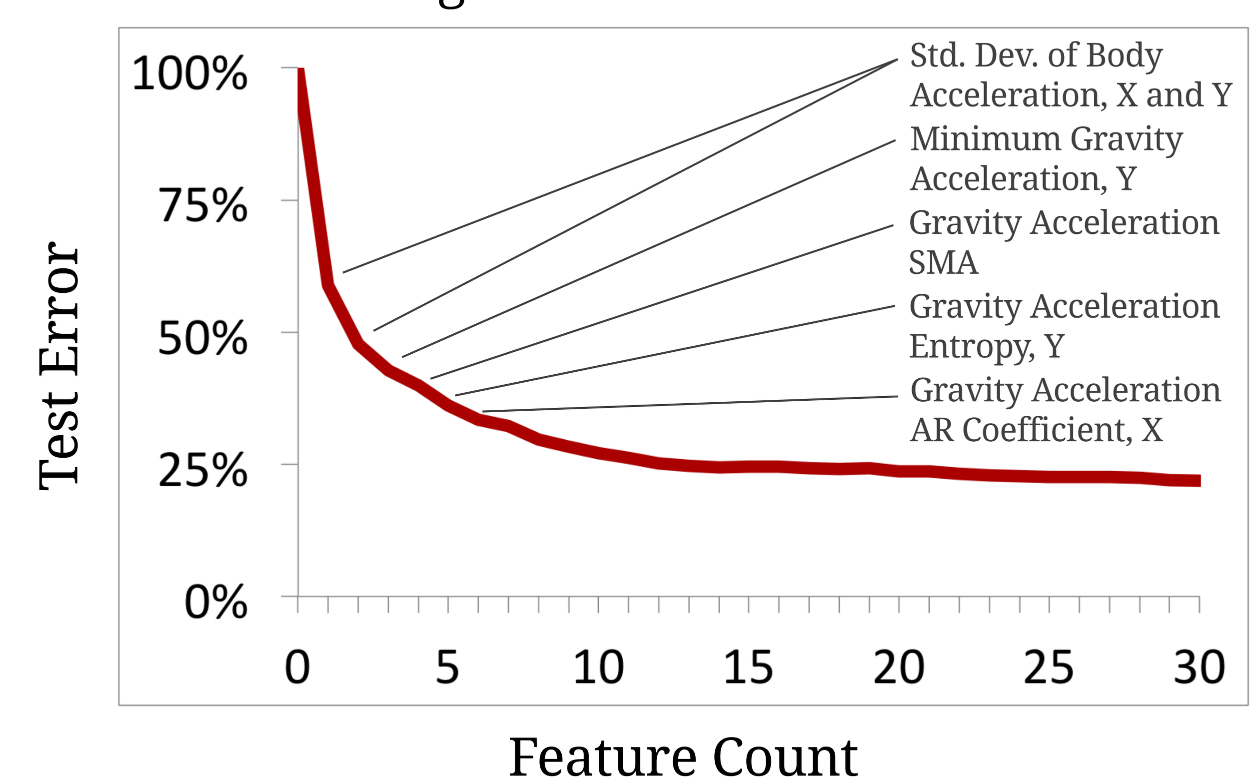

The 6-class SVM achieved a final test set accuracy of 78.1% using all 30 automatically-chosen features.

K-Means Clustering

Since there were 6 distinct activities in the dataset, we were curious how an unsupervised learning algorithm (k-means clustering) would partition the data. We used standard Euclidean distance to measure the similarity between data points consisting of all 563 features. The chart below shows the cluster composition after 23 iterations with 6 randomly-initialized centroids.

We also experimented with initializing one centroid for each activity, but found that the random initialization often converged to one centroid per activity. We found that k-means wasn't ideal for distinguishing between individual activities, but had greater than 98% accuracy distinguishing between the 3 active vs. the 3 passive activities.

Activity Composition of Clusters Found by K-means

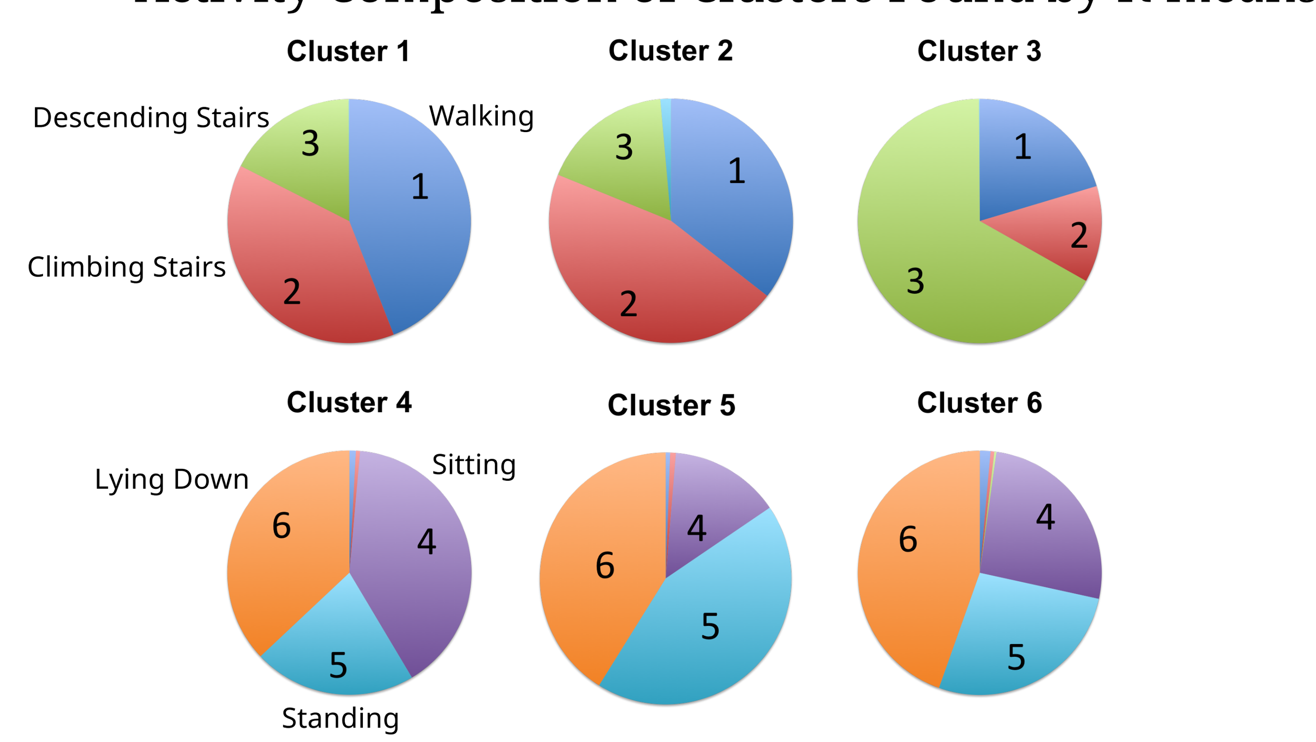

Multi-Layer SVM and Plurality Voting

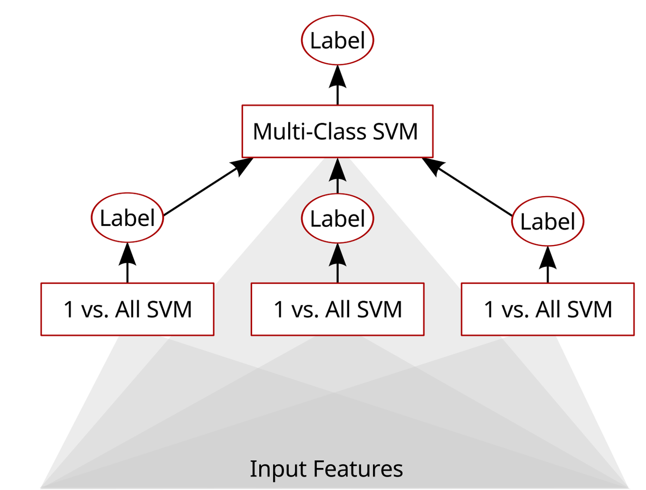

Multi-layer SVM

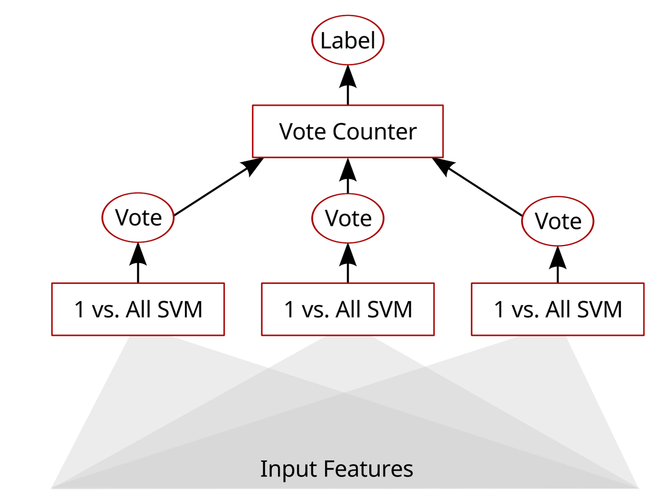

Plurality Voting

In the multi-layer SVM, the ith base-level "1 vs. all" SVM classifies an input as label "i" or "not i". These SVMs use all input features.

The upper-level multi-class SVM then uses the base-level SVM outputs as additional input features, augmenting the 30 automatically-selected features.

Inaccuracies in the predictions of the base-level SVMs degraded the output of the upper-level SVM when compared to the standard multi-class SVM. The multi-layer SVM achieved a test set accuracy of 76.8%.

In plurality voting, the base-level SVMs similarly classify the input features. The vote counter tallies the votes for a given feature to decide which label to apply. It breaks ties randomly.

Using all input features, voting achieved a test-set accuracy of 64.4%.

Conclusions

Classification Accuracies

Algorithm	Accuracy
Multi-Class SVM	78.1%
Multi-Layer SVM	76.8%
Voting	64.4%
K-Means	47.4%

Of the algorithms we implemented, multi-class SVM with automatic feature selection produced the highest test accuracy. The SVM performed better than multi-tiered voting systems or unsupervised algorithms like k-means because of high classification error between two “active” or two “passive” tasks (ie walking upstairs vs. walking).

With a larger training set, we could train and test the SVM on a single user, which would allow better prediction accuracy for that user at the cost of generalizability. A more specific and accurate classifier would be more useful for phone applications that benefit from knowing the current context of the user's activity.

Works Cited

[1] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra and Jorge L. Reyes-Ortiz. A Public Domain Dataset for Human Activity Recognition Using Smartphones. 21th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2013. Bruges, Belgium 24-26 April 2013.

[2] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A Library for Large Linear Classification, Journal of Machine Learning Research 9(2008), 1871-1874. Software available at http://www.csie.ntu.edu.tw/~cjlin/liblinear

[3] Chih-Chung Chang and Chih-Jen Lin, LIBSVM : a library for support vector machines. ACM Transactions on Intelligent Systems and Technology, 2:27:1--27:27, 2011. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm

[4] Ravi, Nishkam, Nikhil Dandekar, Preetham Mysore and Michael L. Littman. "Activity Recognition from Accelerometer Data." 2005. Rutgers University. 6 Dec 2015 <http://paul.rutgers.edu/~nravi/accelerometer.pdf>

[5] Icons by The Noun Project. https://thenounproject.com/.