Activity Recognition for Personal Time Management

Zoltán Prekopcsák, Sugárka Soha, Tamás Henk, and Csaba Gáspár-Papanek

Budapest University of Technology and Economics, Hungary
prekopcsak@tmit.bme.hu, sugarka.soha@gmail.com,
henk@tmit.bme.hu, gaspar@tmit.bme.hu
http://prekopcsak.hu

Abstract. We describe an accelerometer based activity recognition system for mobile phones with a special focus on personal time management. We compare several data mining algorithms for the automatic recognition task in the case of single user and multiuser scenario, and improve accuracy with heuristics and advanced data mining methods. The results show that daily activities can be recognized with high accuracy and the integration with the RescueTime software can give good insights for personal time management.

1 Introduction

Activity recognition is a popular research area since the 1980s, when camera based methods have been developed, but due to the fast advancement in small sensor technology, it is gaining even more attention in the past few years. Since the first sensor based results [11], activity recognition has been used for the surveillance of medical patients [7], providing context-aware information [2,9] and many more in military and public domains [3]. The algorithmic results are very promising and activity information could be used in many ways, but so far there is no widespread application of this recognition technology.
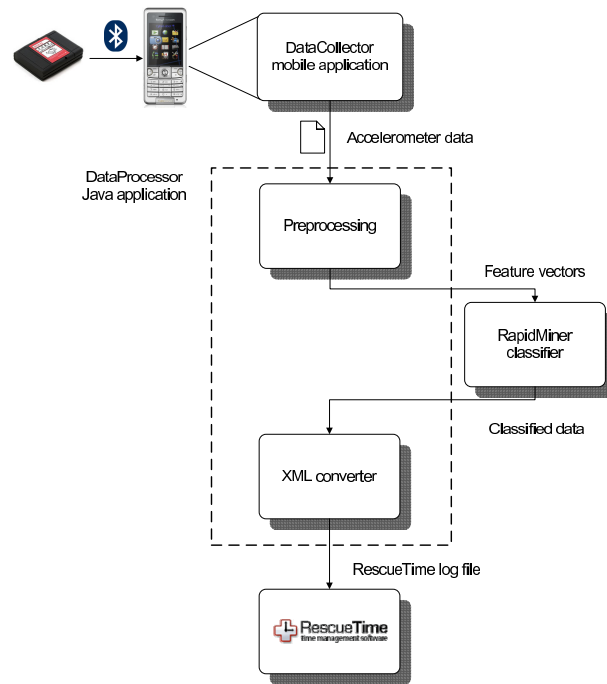
Furthermore, our life is getting faster and information workers try to be as efficient as possible which made a hype for time management theories, like GTD (Getting Things Done [1]) and gave birth to many software applications which track the time that we spend in front of our computers (e.g. RescueTime). They can give detailed statistics about the time spent with specific applications and documents, and they can even provide suggestions on managing our time better.

As computer usage is continuously shifting towards mobile, we have no information about the context in which we are using our devices, and of course there is a significant amount of time that we spend without using them. Mobile activity recognition solutions can provide additional information about our context and activities even when we are not using electronic devices.

2 System Description

Our goal was to design an activity recognition system that can be used in real life for personal time management. We have used a regular mobile phone and a

small external accelerometer sensor (SparkFun WiTilt v3.0) which was accessed via Bluetooth connection and placed in a belt case, so we could easily track its orientation which makes the recognition easier. It is not a serious limitation as many people uses belt cases, but it is possible to extend our recognition tests in the future for other mobile wearing habits.

Fig. 1. Building blocks of the system

For the data collection period, we have created a mobile application that helps to collect annotated data about our lives that can serve as training data for the recognizer algorithms. We have decided to pre-define 10 daily (or weekly) activities: walking, running, working (while sitting), watching television, cooking, vacuuming, taking stairs, using the elevator, traveling on a bus, lying. These activities have been selected because they are done frequently and they might be of interest from a personal time management perspective. On the mobile screen, the user can select from the list of these activities, which starts the data collection to the memory card or the phone's internal memory. The user can always change the activity in the list or stop the data collection. The application is written in mobile Java, so it can run in the background and seamlessly track longer activities too.

The accelerometer data files are synchronized to a computer where the recognition software operates. After some data preprocessing, we use the RapidMiner open source data mining framework [8] for the recognizer algorithm, which produces time stamps and activity labels as its output. This is then converted to a special file format for the RescueTime time management software, which adds these activities to its statistics. The whole system can be seen on Figure 1.

3   Activity Recognition

3.1   Data Preprocessing

To be able to use the continuous 3-dimensional accelerometer and the one dimensional gyroscope data for recognition, we had to preprocess it first. We have decided to take two seconds long chunks from the data which resulted 40 data points for each dimension, as the sensor was working on 20Hz. The optimal window length was investigated in earlier papers and they concluded that different activities have different optimal window lengths, but the global optimum is around 2 seconds [5].

From these 4x40 data points, we have calculated many statistical features for the recognition. Previous works suggested minimum, maximum, average and variance values for each dimension [6,9,5], energy measures from the Fourier transformation [2,9,5], and Pearson correlation between the different dimensions [2,9,5]. We believe that these features should well describe the underlying activities, and the classification performance proved our assumptions.

3.2   Classification

With these features, we have tested three different classifier algorithms from the RapidMiner framework: decision tree (DT), support vector machine (SVM), k-nearest neighbor (k-NN). These standard machine learning and data mining algorithms are part of the framework, so we only needed to optimize their parameters. The cross-validation results show that SVM has the highest accuracy, but all methods have the recognition rate around 90%. It is an excellent result considering that there are 10 different activities to be recognized, so random guessing would produce 10% accuracy. We have to note, that for an application that runs only on mobile, decision tree can be the best choice, as it is easy to program on mobile devices and needs low resources.

We have also tried some heuristics to enhance our results. The most successful from these was the usage of a 10 seconds long moving window, which included five 2-second pieces and calculated the most frequently predicted class of these 5 classifications. As it is impossible to have a 2 seconds long cooking activity in minutes of walking, this moving window could correct these errors. This simple heuristic improved the accuracy and sometimes corrected more than half of the errors. Results can be seen on Table 1.

Table 1. Recognition accuracy

Classifier	Accuracy	Accuracy with heuristics
Decision Tree	87.42%	91.97%
Support Vector Machine	93.28%	96.57%
k-Nearest Neighbor	91.27%	95.77%

3.3   Advanced Data Mining Methods

The above mentioned methods are standard data mining algorithms usually used for activity recognition. However, in the past decade, there has been a lot of development in classification algorithms. One of the most important trends is boosting, which is a combination method for a set of weak decision rules to form a highly accurate classifier. A comprehensive overview of boosting methods is written by Schapire [10] and the LogitBoost algorithm is described by Friedman et al. [4]. We have used the latter algorithm with iteratively training small decision trees and combining them for the final model.

Cross-validation was used to optimize the iteration number of LogitBoost to avoid overfitting. The optimal iteration number was 60, which resulted an accuracy of 95.53% on its own and 97.59% when using the heuristics described above. The iteration number means that the final model includes 60 small decision trees. It seems complex, but these 60 trees can be programmed to a mobile phone with code generation.

In case of 97.59% accuracy, the error rate is only 2.41%, so we have analyzed these errors to get some ideas about further improvements. It turned out that most of these errors come from using the elevator being misclassified as cooking which makes sense, as these are very similar activities. We do not consider using the elevator as a very important activity, so by removing it from the study, the accuracy exceeds 99%.

3.4   Multiuser Scenario

These results apply for the case of a single user, when training and test data was collected from the same person. To test the generalization performance of the system, we have made a measurement with training on one user and testing on another. As expected, accuracy drops significantly to around 58–63% for all four algorithms, but we achieved 68.39% with combining them with simple majority voting. This result can be improved if we collect more data from different users.

4   Conclusion

We have shown an activity recognition system that can be used with everyday mobile phones. We have compared standard data mining algorithms and boosting methods, and presented some heuristics to improve the results. The recognition accuracy is over 99% for a single user and 68% for a new user, but there is

a huge potential of improving the latter result by additional data collected from more users. Similar results can be achieved with a mobile implementation using low resources.

This system is mainly intended for personal time management, so there is no need for perfect recognition, mistakes are not critical, because its goal is only to track important trends and provide overall statistics. We have found that it can give interesting insights to our personal lives and help us manage better the everyday tasks that we face.

Personal time management is a popular application area, so we believe that it will help spreading activity recognition components for mobile phones in the next few years. If activity recognition becomes a basic service in the mobile, all applications will know much more about our context and they can adapt to us.

References

1. Allen, D.: Getting Things Done: The Art of Stress-Free Productivity. Viking, New York (2001)
2. Bao, L., Intille, S.S.: Activity Recognition from User-Annotated Acceleration Data. In: Ferscha, A., Mattern, F. (eds.) PERVASIVE 2004. LNCS, vol. 3001, pp. 1–17. Springer, Heidelberg (2004)
3. Choudhury, T., Borriello, G., et al.: The Mobile Sensing Platform: An Embedded System for Activity Recognition. In: IEEE Pervasive Magazine - Special Issue on Activity-Based Computing (2008)
4. Friedman, J., Hastie, T., Tibshirani, R.: Additive Logistic Regression: a Statistical View of Boosting. Annals of Statistics 28 (1998)
5. Huynh, T., Schiele, B.: Analyzing Features for Activity Recognition. In: Proceedings of Smart Objects & Ambient Intelligence Conference (2005)
6. Kern, N., Schiele, B., Schmidt, A.: Multi-Sensor Activity Context Detection for Wearable Computing. In: Aarts, E., Collier, R.W., van Loenen, E., de Ruyter, B. (eds.) EUSAI 2003. LNCS, vol. 2875, pp. 220–232. Springer, Heidelberg (2003)
7. Mathie, M.J., Celler, B.G., Lovell, N.H., Coster, A.C.: Classification of basic daily movements using a triaxial accelerometer. Medical & Biological Engineering & Computing 42, 679–687 (2004)
8. Mierswa, I., Wurst, M., Klinkenberg, R., Scholz, M., Euler, T.: YALE: Rapid Proto-typing for Complex Data Mining Tasks. In: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2006 (2006)
9. Ravi, N., Dandekar, N., Mysore, P., Littman, M.L.: Activity recognition from accelerometer data. In: Proceedings of IAAI 2005 (2005)
10. Schapire, R.: The Boosting Approach to Machine Learning: An Overview. In: MSRI Workshop on Nonlinear Estimation and Classification (2003)
11. Uiterwaal, M., Glerum, E.B.C., Busser, H.J., Van Lummel, R.C.: Ambulatory monitoring of physical activity in working situations, a validation study. Journal of Medical Engineering & Technology 22 (1998)