Activity Recognition With Mobile Phones

Jordan Frank[1], Shie Mannor[2], and Doina Precup[1]

[1] School of Computer Science, McGill University, Montreal, Canada,
{jordan.frank,dprecup}@cs.mcgill.ca
[2] Department of Electrical Engineering, Technion, Israel
shie@ee.technion.ac.il

Abstract. Our demonstration consists of a working activity and gait recognition system, implemented on a consumer smartphone. The activity recognition feature allows participants to train various activities, such as running, walking, or jumping, into the phone, and the system can then identify when those activities are performed. The gait recognition feature learns particular characteristics of how participants walk, allowing the phone to identify the carrier.

1 Overview

We have implemented a system for activity and gait recognition for the Google Android OS ©, which runs on smartphones from a variety of manufacturers. During our demonstration, we will allow participants to try out the various features, either demonstrating the phone's ability to recognize a pre-programmed set of activities (running, jogging, jumping, or walking), or the phone's ability to learn the characteristics of a participants' gait, or style of walking, and showing that the phone can later recognize that participant. Activity and gait recognition using wearable sensors are active research topics [see 8, 1, 9, for example]. Unlike previous work, we do not require specialized wearable sensors or careful placement of the phone on the participant's body. Our application works when the phone is simply placed in a trouser pocket. Additionally, the computation requirements for training these systems are prohibitive for real-time demonstrations, and in many cases even the classification requires more computation that can be afforded on a mobile device such as a smartphone. Our system was designed with computational efficiency in mind, and can both learn classifiers and perform classification in real-time on a low-powered mobile device.

The demonstration will be interactive, and participants will be asked to perform different activities, with the classification results displayed on a separate monitor for everyone to see.

2 Background

The activity and gait recognition systems are based on an algorithm called geometric template matching (GTM) [5] for comparing segments of time series.

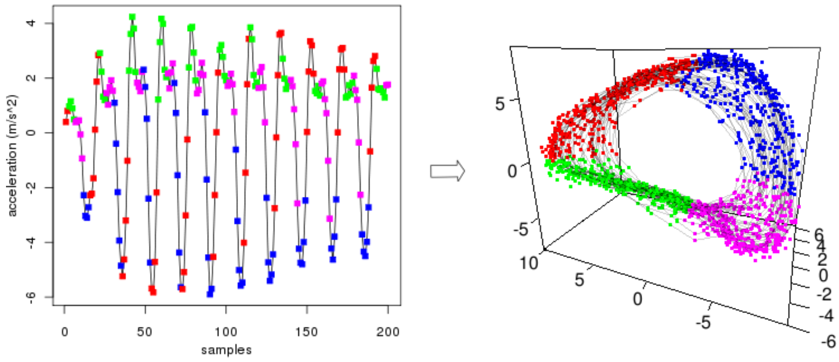

Fig. 1. An example of a time-delay embedding for a sequence of data collected from an accelerometer while a subject was riding a bicycle. The colours are purely for illustrative purposes and show that similar-looking segments of time series appear in similar regions of the embedded model.

GTM operates by first building models, or templates, from short segments of labeled univariate time series, such as that collected by accelerometer sensors. These models are constructed using a technique called time-delay embedding, which is an approach developed in the nonlinear dynamical systems community for reconstructing dynamical models from measurements of some latent nonlinear system [6]. To compute the time-delay embedding for a sequence of measurements $o_1, \ldots, o_N$, one choses a lag parameter τ and a reconstruction dimension m, and then constructs the sequence of m-dimensional points:

$$x_i = (o_i, o_{i+\tau}, o_{i+2\tau}, \ldots, o_{i+(m-1)\tau}),$$

for $i = 1, \ldots, N-(m-1)\tau$. The model, then, consists of points, each composed of m τ-lagged samples of the data. The choice of m and τ are important, and while there are heuristics for choosing these values [3, 7], we have found that standard machine-learning approaches to parameter tuning, such as cross-validation and grid search are most effective. In addition, GTM is fairly robust with respect to the parameters, and values that work well for one problem (identifying running, for example), tend to work well for other problems (identifying biking or walking, for example). Figure 1 depicts the process of time-delay embedding for a segment of accelerometer data collected from an accelerometer sensor while a subject rode a stationary bicycle. It is clear from the embedding that, although the data is nonstationary, the periodic element is captured, and regions in the time-series that appear qualitatively similar are mapped to nearby points in the embedding.

The second step of GTM is to compare a new segment of time series to existing models. To do this, the new time series is projected into the embedding space and then a measure of similarity is computed by considering pairs of subsequent points as vectors and then computing a measure of similarity between the vectors representing the new time series and their nearest neighbours in the model. As a result, GTM computes a measure of similarity between the new time series and each of the models, and these similarity scores can be used to find the model that most closely matches the new time series, or as features for a more complex classifier. We have used these features to train an SVM [4] for classifying activities, and used the model with the highest similarity score for gait recognition [5]. In both cases, we achieve accuracies that exceed the performance of state-of-the art systems at a lower computational cost.

The advantages of GTM over other approaches for comparing segments of time series, such as dynamic time warping (DTW) [2] include its computational efficiency ($O(n \log n)$ to compute each similarity score, as opposed to $O(n^2)$ for DTW), and the fact that it can compare segments of different lengths.

3 Implementation

We have implemented a system for building models and performing classification for accelerometer data collected on mobile phones that run the Google Android OS ©. Due to the efficiency of the GTM algorithm, our system is able to build new models and perform classification in real-time on the phone without drastically reducing the battery life of the phone. The application that we will demonstrate provides an interface for collecting labeled training data from subjects, and for sending live classification results to a separate laptop for visualisation. The phone does not need to be affixed in a special way to the body of the subject, and can be placed in a trouser pocket. We intend to demonstrate two systems, one for recognizing particular activities, and one for recognizing individuals based on their gait[3].

For activity recognition, we will pretrain the phone with a number of activities that can be safely performed indoors without additional equipment, such as running, walking, jogging, and jumping up and down. Participants will then be asked to perform these activities, one at a time, while carrying the phone in their pocket. The classification will be performed on the phone, but to allow other participants to view the results, we will stream them wirelessly from the phone to a separate laptop that will display the results. In the event that participants come up with activities that we haven't thought to pretrain the phone with, the participants will be able to demonstrate the new activity while carrying the phone, and the phone will learn the new activity. Then, this new activity will be included in the set of activities that the phone can recognize.

For gait recognition, participants will be asked to walk with the phone for approximately 15 seconds, while the phone builds a model of their gait. The phone will be pretrained with the gait from 4 other subjects. The participant will be asked to walk again with the phone, and the phone will attempt to classify

[3] Given the nature of the data that the phone will be collecting, participants will be asked to sign a consent form. This demonstration will be performed in accordance with the McGill University Research Ethics Board, and our Certificate of Ethical Acceptability of Research Involving Humans will be clearly displayed. Any data that is collected will be anonymous, and will not be associated with any personally identifying information.

the particular person that is carrying it from the 5 current models. As with the activity recognition, the results will be streamed to a laptop, allowing others to view the results.

4 Conclusion

We propose to demonstrate a state of the art activity and gait recognition system. The demonstration will be interactive, allowing participants to try out the systems, and the results will be displayed in real-time for other to see. These types of systems are typically evaluated in controlled environments, using special-purpose sensors, and often the processing is computationally expensive and done offline at a later time. This demo is interesting in that it uses consumer-grade mobile phones, does not require them to be carefully placed on a participant, and the processing is all done in real-time on the phone. There are many industrial applications for these types of systems, such as health-care monitoring, fitness, and biometric security, making this demonstration appealing to both academic researchers and industry practitioners.

Additionally, the software is open-source and publicly available[4], and so participants that are interested in evaluating the system on their own can install the software on their own phones.

References

[1] Ailisto, H., Lindholm, M., Mantyjarvi, J., Vildjiounaite, E., Makela, S.: Identifying people from gait pattern with accelerometers. In: Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. vol. 5779, pp. 7–14 (2005)
[2] Berndt, D., Clifford, J.: Using dynamic time warping to find patterns in time series. In: AAAI-94 workshop on knowledge discovery in databases. pp. 229–248 (1994)
[3] Buzug, T., Pfister, G.: Optimal delay time and embedding dimension for delay-time coordinates by analysis of the global static and local dynamical behavior of strange attractors. Phys. Rev. A 45(10), 7073–7084 (1992)
[4] Cristianini, N., Shawe-Taylor, J.: An introduction to Support Vector Machines: and other kernel-based learning methods. Cambridge University Press (2000)
[5] Frank, J., Mannor, S., Precup, D.: Activity and gait recognition with time-delay embeddings. In: Proc. of the AAAI Conference on Artificial Intelligence (2010)
[6] Kantz, H., Schreiber, T.: Nonlinear time series analysis. Cambridge University Press (2004)
[7] Kennel, M., Brown, R., Abarbanel, H.: Determining embedding dimension for phase space reconstruction using the method of false nearest neighbors. Phys. Rev. A 45(6), 3403–3411 (1992)
[8] Lester, J., Choudhury, T., Borriello, G.: A practical approach to recognizing physical activities. Lecture Notes in Computer Science 3968, 1–16 (2006)
[9] Subramanya, A., Raj, A., Bilmes, J., Fox, D.: Recognizing activities and spatial context using wearable sensors. In: Proc. of Uncertainty in Artificial Intelligence (2006)

[4] `http://jwf.github.com/Humansense-Android-App/`