A Novel Similarity Measure for Time Series Data with Applications to Gait and Activity Recognition

Jordan Frank
McGill University
jordan.frank@cs.mcgill.ca

Shie Mannor
McGill University
shie@ee.technion.ac.il

Doina Precup
McGill University
dprecup@cs.mcgill.ca

ABSTRACT

In this abstract, we propose a novel approach to modeling time-series for the purpose of comparing segments of data in order to classify activities based on accelerometer sensor data. Our approach consists of producing an ensemble of simple classifiers that can be built and can classify new data efficiently. We present empirical results from an implementation of our algorithm running on a mobile phone, demonstrating the efficiency and performance of our technique on real-world data. Our algorithm is able to identify individuals based on their gait, and can be used in a semi-supervised setting to label large data sets using a small number of labeled examples. Our method can also be used in an unsupervised setting to visualize time-series data, for example, to identify the number of different activities that occur in an unlabeled data set.

Author Keywords Activity recognition, gait recognition, supervised learning, semi-supervised learning, unsupervised learning, wearable computing.

ACM Classification Keywords I.5.1 [Computing Methodologies]: Pattern Recognition, Models.

General Terms Algorithms, Design, Experimentation, Human Factors

INTRODUCTION

Mobile devices are ubiquitous. Of the over 174 million smart-phones sold worldwide in 2009 [1], the majority contained sensors such as accelerometers and GPS that can provide rich contextual information about the user's daily activities. This creates the potential for intelligent applications that recognize the user's activity and take appropriate actions. For example, a phone that can recognize the gait of the person carrying it can lock important private information if someone other than the owner were to pick it up and walk off with it. A phone that can recognize physical activities can assist with fitness routines, or provide health-care professionals with a simple way to monitor patients' activity levels during rehabilitation.

Our focus is on detecting primitive physical activities from data collected using accelerometers built into standard off-the-shelf mobile phones. We propose an novel method for constructing features from time-series data that is better suited for the task of representing measurements from a nonlinear dynamical system (the human body) than the features that are traditionally considered [2,3], such as those used by the linear signal-processing community. One significant advantage of our approach is that robust models can be built from short segments of data.

We demonstrate our approach on a number of real-world data sets collected by ourselves and others. We demonstrate that by comparing data to models built from data collected while a person is walking, we can accurately recognize the individual carrying the phone. We also demonstrate that our technique can discriminate between different activities such as walking, running, etc. Finally, we demonstrate the use of our technique for building a similarity matrix between short segments of accelerometer data, which allows us to use clustering techniques to visualize the data, and to perform semi-supervised labeling, where unlabeled data is used in the learning algorithm to improve the overall performance. Due to lack of space, only some of these results are presented in this abstract.

OUR APPROACH

Our approach to modeling and comparing segments of time-series data is composed of two steps, a modeling step, and a comparison step. The modeling step uses a technique called time-delay embedding [4]. A time-delay model of a sequence of data is computed by taking as points in model-space a set of univariate measurements from a dynamical system evenly spaced in time. For example, a 3-dimensional model with the delay parameter set to 3 for the sequence *(a,b,c,d,e,f,g,h,i,j)* would consist of the points *(a,d,g)*, *(b,e,h)*, *(c,f,i)*, and *(d,g,j)*. We call this set of points the reconstruction of the original data set in model-space. It is well known in the nonlinear dynamical systems community that if the dimension and delay parameter are chosen appropriately, then the reconstruction will be diffeomorphic to the original state space of the dynamical system and its dynamics [4]. While this statement holds only when there is no noise in the measurements, we have found that these time-delay embeddings of noisy sensor data preserve enough of the dynamics of the original system such that we can accurately differentiate not only between different activities, but between different people performing the same activity. The models are compact, comprising of only the raw data points with no pre-processing of the signal, and efficient to construct. We store the points in a KD-Tree for efficient nearest-neighbour lookups.

UbiComp'10, September 26–29, 2010, Copenhagen, Denmark.
ACM 978-1-4503- 0283-8/10/09.

Once models of the sensor data of interest have been computed, we can compare new data to each of the models. This is done by first projecting the data into the model-space by computing a time-delay embedding of the data. Next, the nearest-neighbours in the model can be found efficiently by searching the KD-tree. We then compare the new data to the nearest-neighbours in the model using an algorithm that we call the Geometric Template Matching algorithm [5]. We omit the details of this algorithm for lack of space. This algorithm produces a measure of similarity between the data and each of the models. These scores can be used to classify the new data by choosing the model that is most similar, can be used as inputs to supervised learning algorithms, or a similarity matrix between pairs of time series segments can be computed and used to visualize the data.

GAIT RECOGNITION

We employed our technique to build models for 40 individuals. Each individual was asked to place a mobile phone in their trouser pocket and walk for approximately 20 seconds. Data was collected at 32Hz from the accelerometer in the phone. The data was divided into 5 segments. One segment was used to build a model for each user, and then the other four segments were compared against the 3-dimensional time-delay embedding models for all of the users. This process was repeated 5 times using each of the segments to build the models. The segments were classified according to the model that gave the highest average similarity score, and remarkably the algorithm chose the correct model with 100% accuracy.

ACTIVITY CLUSTERING

We collected data from four individuals while performing a sequence of fitness activities: riding a stationary bike, using an elliptical machine, using a stationary rowing machine, using a stair-climbing machine, and running on a treadmill. The data was partitioned into 597 five-second segments, and each pair of segments was compared against each other to compute a similarity matrix. To highlight the efficiency of our approach, we note that this 597x597 similarity matrix was computed in approximately 10 minutes on a 3.0GHz Quad-Core desktop computer. As is typically done in spectral clustering [6], an eigendecomposition of the similarity matrix was performed, and we used the first 4 components (the number of people) of the eigenvectors as representatives for each point. We use the Kamada-Kawai algorithm [7] to visualize the 4-dimensional representation in 2 dimensions, and the result is shown in Figure 1. For clarity, the data for only two of the subjects is shown, and it is clear that the individual activities form clusters, and while the separation between the activities is clear, the data for each individual also form sub-clusters. Sequential points are connected by dotted lines, and so we can see that the majority of the activities that appear to be placed in the wrong cluster actually represent transitions between activities.

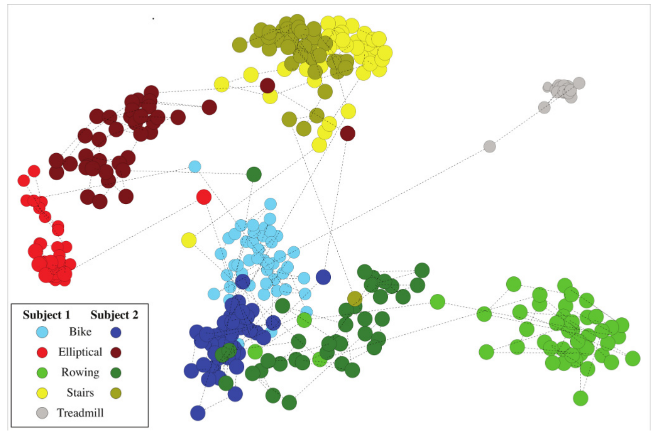

Figure 1. Visualization of Activity Dataset.

CONCLUSION AND DISCUSSION

We present a method for modeling and comparing sensor data that takes into account the inherent nonlinearities in the dynamical system being measured. The approach that we present is efficient both in terms of the computation and the space required to store the models. The algorithm has been implemented and runs in real-time on a mobile phone. We have evaluated these methods on a number of real-world data sets and shown that they are able to accurately compute similarity measures between segments of data, which can be used for classification and data visualization.

REFERENCES

1. IDC Worldwide Quarterly Mobile Phone Tracker, February 4, 2010.
2. J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. Lecture Notes in Computer Science, 3968:1–16, 2006.
3. A. Subramanya, A. Raj, J. Bilmes, and D. Fox. Recognizing activities and spatial context using wearable sensors. In Proceedings of Uncertainty in Artificial Intelligence, 2006.
4. T. Sauer, J. Yorke, and M. Casdagli. Embedology. Journal of Statistical Physics, 65(3):579–616, 1991.
5. J. Frank, S. Mannor, and D. Precup. Activity and Gait Recognition with Time-Delay Embeddings. In Proceedings of AAAI Conference on Artificial Intelligence, 2010.
6. U. von Luxburg. A tutorial on spectral clustering. Statistics and Computing, 17(4):395–416, 2007.
7. T. Kamada and S. Kawai. An algorithm for drawing general undirected graphs. Information Processing Letters, 31(1):7–15, 19.