Activity Detection Using Time-Delay Embedding in Multi-modal Sensor System

Ferdaus Kawsar[1], Md Kamrul Hasan[1], Tanvir Roushan[1],
Sheikh Iqbal Ahamed[1], William C. Chu[2], and Richard Love[3]

[1] Marquette University, USA
[2] Tunghai University, Taiwan
[3] Amader Gram
{ferdaus.kawsar,mdkamrul.hasan,
tanvir.roushan,sheikh.ahamed}@mu.edu
cchu@thu.edu.tw,richardibcrf@gmail.com

Abstract. About two billion people in this world are using smart devices where significant computational power, storage, connectivity, and built-in sensors are carried by them as part of their life style. In health telematics, smart phone based innovative solutions are motivated by rising health care cost in both the developed and developing countries. In this paper, systems and algorithms are developed for remote monitoring of human activities using smart phone devices. For this work, time-delay embedding with expectation-maximization for Gaussian Mixture Model is explored as a way of developing activity detection system. In this system, we have developed lower computational cost algorithm by reducing the number of sensors.

Keywords: Human activity detection, remote monitoring, time delay embedding.

1 Introduction

Accurate information regarding human physical activity and ability to access that information in real time has far-reaching significance. Activity information is important to doctors who wants to monitor their patients remotely. This technology can be used for monitoring elderly people who wants to maintain their independence. However, such monitoring systems usually require complex devices and significant involvement from the participants. Complex devices can be expensive whereas intrusive systems greatly discourages the usage in real life. Consequently, we focused on developing monitoring systems using smart phones since smart phones are ideal candidate for numerous innovation. Smart devices has significant computational power, storage and communication capability and is conveniently carried out by mass people. Developing a system centered around smart phones will most likely remove the necessity of carrying other extra devices. Even if it is required to use other sensors, it is possible to connect with those sensors using Bluetooth connectivity. As, by now 2 billion people worldwide are using smart phones, we now have a unique phenomenon where significant computational power, storage, connectivity, and built-in sensors are carried by mass people willingly as part of their life style. This unique phenomenon provides a great opportunity in terms of research and innovation. A realistic smart phone based activity monitoring system can help to reduce the cost of health care. The initial results of these works are illustrated in [Kawsar et al., 2015].

2 Background

Gaussian Mixture Model (GMM) provides unique opportunity to analyze time series data from multiple sensors. For this reason, we have used this GMM technique for activity detection with time-delay embedding.

2.1 Time-Delay Embedding with Gaussian Mixture Model

Time-delay embedding theorem gives the conditions under which a chaotic dynamical system can be reconstructed from sequence of observations of the state of dynamical system. The reconstruction preserves the properties of dynamical system that do not change under smooth coordinate changes. Taken's theorem [Takens, 1981] provides the conditions under which a smooth attractor can be reconstructed from observations. This theorem essentially provides approaches for reconstructing the essential dynamics of the underlying system using a sequence of observations. The assumption is that the dynamics of the underlying system are significantly different for different activities of a person. In our case, we observed accelerometer data along X and Y axis as well as six pressure sensors from left shoe.

The parameters of time-delay embedding models are learned using a Gaussian Mixture Model. In our experiments, number of mixture models we used are three. In reality, true dimension of phase space is usually unknown. Based on some trial and error, we used a six dimensional phase space with time lag, $\delta = 5$.

2.2 Gaussian Mixture Model

Gaussian mixture models are extension of k-means models. If random variable X is Gaussian, it has the following probability density function(pdf):

$$N(x|\mu, \sigma^2) = p(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}}$$

The two parameters are mean, μ and variance, σ^2. $p(x)$ can be conveniently written as $N(x|\mu, \sigma^2)$. If we have independent and identically distributed observations X_1^n from a Gaussian distribution with unknown mean μ, maximum likelihood estimation for μ will be $\frac{1}{N}\sum_i x_i$.

Gaussian mixture model (GMM) is useful for modeling data that comes from one of several groups. The groups may be different from each other. However, data from same group can be modeled using Gaussian distribution. A superposition of K Gaussian densities can be written as

$$p(x) = sum_{k=1}^{K} \pi_k N(x|\mu, \sigma^2)$$

which is called a mixture of Gaussians. Each Gauassian density is called a component of the mixture and has its own mean, μ_k and variance σ. The parameters π_k is called mixing co-efficients. Also, $sum_{k=1}^{K} \pi_k = 1$ and $0 \leq \pi_k \leq 1$ in order to be valid probabilities.

Expectation-maximization (EM) algorithm is an iterative method for finding maximum likelihood estimates of parameters in Gaussian Mixture environment. Maximum likelihood estimation in Gaussian mixture model is the estimation of π_k, μ and σ of the component of Gaussian mixture.

3 Related Works

Shaji et al. designed an innovative BMA classifier in [Shaji et al., 2016] that can classify different physical activities: walking, jogging, sitting, standing, climbing upstairs, coming downstairs, and lying down. They having got the accuracy of 96.66%. The tested results of the BMA classifier is integrated with a complete system where the classification system is able to show the above mentioned task getting data from a context aware system and using their algorithm.

Capela et al. [Capela et al., 2016] has developed smartphone-based HAR classifier where accelerometer, magnetometer, and gyroscope data were used from smartphone. The author used orientation correction matrix to all sensor data that can give appropriate information on human movement activities for both able-bodied and stroke populations.

For human physical activity recognition, Kwapisz et al. [Kwapisz et al., 2011] used phone-based accelerometers in their research to build human physical recognition system. Twenty-nine users performed daily activities such as walking, jogging, climbing stairs, sitting and standing. Labeled accelerometer data were collected from these users. These data was used as training data to build a predictive model for activity recognition. As users always carry cell phones in their pockets, this work can help to collect information about the habits of millions of users.

They have used accelerometer data from android phone to identify several activities. Android was chosen because the OS is free and open-source, easy to program. This architecture has the advantage of using a device since mass people can keep their mobile phone along with themselves. Authors have used the data to extract six features, namely standard deviation, average absolute difference, average resultant acceleration, time between peaks and binned distribution. Now raw time-series accelerometer data must be transformed into examples since standard classification algorithms cannot be directly applied to it. Three classification techniques - decision trees (J48), logistic regression and multilayer neural network from WEKA data mining suite were used. The system is very unobtrusive as the cell phone carried by users work as data collection system. But it requires the users to carry the phone in a certain location.

Indoor location of a person is estimated in Lee and Mase's [Lee and Mase, 2001] work. The system uses a bi-axial accelerometer, a digital compass and an infrared light detector. This work identifies walking and whether the person is walking in level ground, going up or going down. It also counts the number of steps. The strategy adopted by the researchers is hybrid: dead-reckoning for relative measurements and infrared-based beacon method for absolute measurement. Accumulation of error is common in dead-reckoning system and an infrared-based beacon method that detects signals from a transmitter in a fixed place (stairway) helps to correct those errors. By using conventional peak detection algorithm, the system tries to find the peak values at every sampling. If the values of all four peaks follow some specific conditions, step count is incremented. Another feature called cross-correlation function of x(t) and z(t) is used to improve performance. This feature is helpful for discriminating between level and up/down. The classification results show good performance for level and down behaviors but up behavior detection is not satisfactory. One problem with this work is that the connection to central mobile unit is not wireless.

Yang [Yang, 2009] developed an activity recognition system using the built-in accelerometers in Nokia N95 phone. Although the study achieved relatively high accuracies of prediction, stair climbing was not considered and the system was trained and tested using data from only four users. Decision tree performed best among the four classifiers evaluated. Other classifiers that were evaluated are Naïve Bayes (NB), k-Nearest Neighbor (kNN) and Support Vector Machine (SVM). As phone's position on a human body varies from person to person, its orientation cannot be fixed. Orientation-independent features extraction was also explored in this study.

Miluzzo et al. [Miluzzo et al., 2008] exploits various sensors (such as a microphone, accelerometer, GPS and camera) that are available on commercial smart phones for activity recognition and mobile social networking applications. They collected accelerometer data from ten users to build an activity recognition model for walking, running, sitting and standing. Their applications 'CenceMe', collects sensor data of individuals using off-the-shelf, sensor-enabled mobile phones, analyzes these data, detects the activities and share these information through social networking applications such as Facebook and MySpace. To make the system scalable, classification task was shared between cell phones and back-end servers. They also carried a user study on twenty two people who used CenceMe continuously over a three week period.

Both the Symbian operating system and Java Micro Edition (JME) virtual machine which runs on top of the N95 have been designed to use small amounts of memory and computational resources. Designing and implementing 'CenceMe' application on top of this environment was thus resource-constraining. One of the contributions of the paper is the design of lightweight classifiers, running on mobile phones where classification is split between cell phone and servers. Another contribution is the measurement of the RAM, CPU, and energy performance of the classifiers and the whole 'CenceMe' software suite.

4 Experimental Setup

We have illustrated the experimental setup for smart phone based multimodal activity detection system using plantar pressure sensors in this section. In this work, we have proposed a novel architecture for the unobtrusive detection of human physical activity using accelerometer and gyroscope data from smart phones as well as pressure data from shoes. Our architecture was designed to make the system unobtrusive and robust against various human behavior patterns. We developed a prototype of the activity detection system using smart phones and plantar pressure sensors based on our proposed architecture. We identified the various issues that came up while developing the system alongside the caveats and their origins and possible solutions. We analyzed data from four activities and developed an algorithm based on our analysis. Later we tested how our algorithm performs and achieved very good accuracy for the activities in the data analysis stage. Several modifications of the algorithm and the evaluation of their performances were also discussed.

4.1 Proposed Architecture

In this architecture, pressure sensors are placed inside the shoes. These pressure data are transmitted over Bluetooth to the user's smart phone. If the smart phone is in the range of Bluetooth, the data collection is performed smoothly. Since Bluetooth has a range of 5-30 meters, the distance from a shoe to the mobile phone is always within bluetooth range. In this case, we tried to collect pressure data from shoes. Plantar pressure sensor system are used with cell phones for activity detection in this research. In addition, we have also collected accelerometer and gyroscope data from the cell phone. The data collection system collects data from these four sensor systems and stores them in four files in three different folders. We used the data collection system in two ways: learning and recognition stage. Plantar pressure sensor system based on a fabric sensor array shown in Figure 2 is used in this work.
This system was developed by Lin Shu et al. [Shu et al., 2010]. It has 8 pressure sensors in each shoe. There is also a Bluetooth interface to transfer the pressure data to an Android phone.

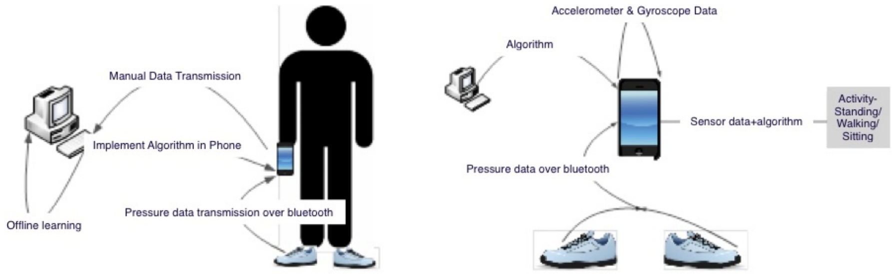

Fig. 1: a) System architecture for learning stage and b) system architecture for activity recognition stage.

Fig. 2: Sensor system inside the sole of shoes.

5 Methodology

In this section we have described activity detection using time-delay embedding with Gaussian Mixture Model. The fundamental idea comes from the work of Takens [Takens, 1981] and Sauer et al. [Sauer et al., 1991]. Their work shows that a time series of observation samples from a system can be used to reconstruct a space topologically equivalent to original system. It is very easy to reconstruct such reconstructive phase space. Time-delay embeddings attempt to reconstruct the state and dynamics of an unknown dynamical from observations of that system taken over time [Frank et al., 2010]. Formulating time series algorithm using multi-dimensional phase space is different than developing algorithms using time or frequency domain features.

Determining the dimension, d, of reconstructive phase space (i.e. how many measurements have to be considered) and determining τ (at what time the measurements should be taken) is a key problem. A row vector is a point in RPS. d must be greater than two times of the box cutting dimension to be topologically equivalent. In our case, we experimented with the parameter $d = 2$ and $d = 3$. We have presented our findings for these two cases. Following is a plot of a time series for pressure data for pressure sensor 1 (PS1) from left shoe.

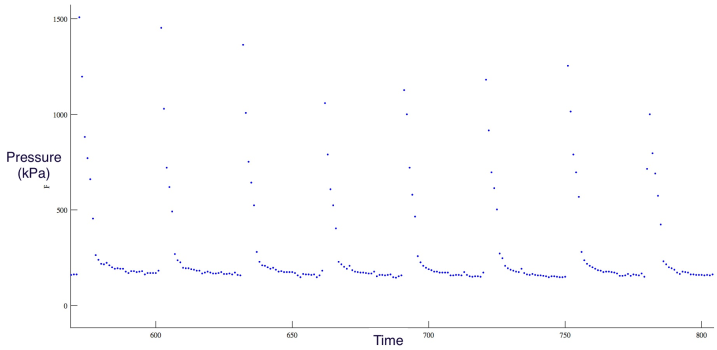

Fig. 3: Data from PS1 for left shoe for running.

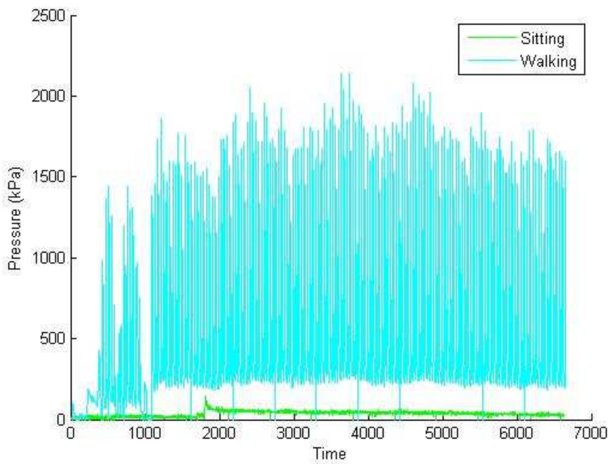

Fig. 4: Data from second pressure sensor (P2) in the left shoe: walking and sitting.

A structure is obvious in Figure 5 where we made a phase plot in 3 dimension for time lag 5 and 10.

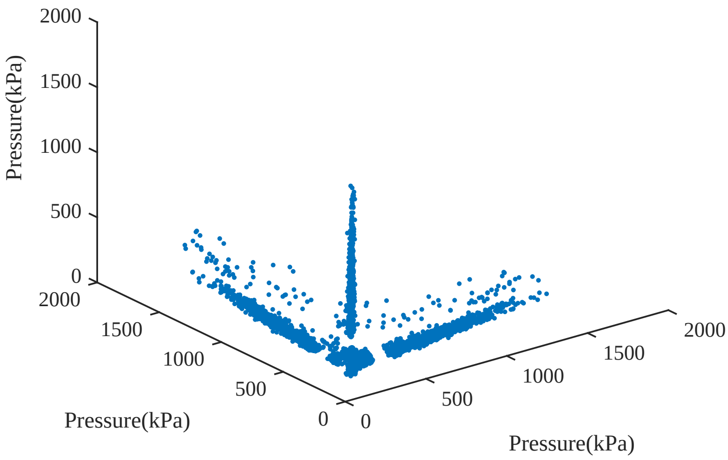

Fig. 5: PhasePlot in 3 dimension with time lag 5 and 10 for running data from pressure sensor 1 of left shoe.

5.1 Our Approach

First we demonstrate our experiment in the case of 2 activities only: running and sitting. We used Gaussian Mixture Model (GMM) with Expectation Maximization(EM) algorithm for classification of embedding features. We build two models. One model for running and other one is for sitting. In both cases, we have collected data from pressure sensor data, PS1 of left shoe. We used 5 mixtures for GMM. When tested, sitting data for sitting model showed higher probability; same was true for running data for running model. That means GMM with time delay embedding can accurately distinguish these two activities of sitting and running.

Symbol	Meaning
R	Running
S	Sitting
W	Walking
St	Standing
Sd	Stair down
Su	Stair up
C	Cycling
D	Driving

Fig. 6: Different symbol for various activites

We define LP_{mn} as the log probability of applying data of n activity on the model of activity m. We also define following symbols for 8 activities which shown in Figure 6. The following table shows the log probabilit for just 1 pressure sensor data, P1 from left shoe. From the above table we see that $LP_{RR}>LP_{sR}$. We also find that $LP_{ss}>LP_{Rs}$. The significance of this numbers is that we can distinguish running and sitting using just 1 pressure sensor P1's data. After a

Case	Log Probability
Testing running data against running model (LP_{RR})	$-3.0050 \times 10^{+04}$
Testing sitting data against sitting model (LP_{ss})	$-1.7031 \times 10^{+04}$
Testing sitting data against running model (LP_{Rs})	$-4.2674 \times 10^{+04}$
Testing running data against sitting model (LP_{sR})	$-1.1068 \times 10^{+06}$

Table 1: Log probabilities for siting and running activity for P1 from left shoe.

running model is developed from running data and sitting model is made from sitting data, applying running and sitting data on these models show that the probability of running data coming from running model is higher than it coming from sitting model.

In the following table, we expand to include standing activity making it a three activity scenario.

Case	Log Probability
Testing running data against running model (LP_{R_R})	$-5.2250 \times 10^{+04}$
Testing sitting data against sitting model (LP_{s_s})	$-3.3829 \times 10^{+04}$
Testing standing data against standing model (LP_{St_St})	$-4.025 \times 10^{+04}$
Testing sitting data against running model (LP_{R_s})	$-1.3754 \times 10^{+05}$
Testing running data against sitting model (LP_{s_R})	$-1.6040 \times 10^{+06}$
Testing standing data against sitting model (LP_{s_St})	$-8.1173 \times 10^{+04}$
Testing sitting data against standing model (LP_{St_s})	$-4.0219 \times 10^{+04}$
Testing standing data against running model (LP_{R_St})	$-1.2025 \times 10^{+05}$
Testing running data against standing model (LP_{St_R})	$-1.1099 \times 10^{+06}$

Table 2: Log probabilities for sitting, running and standing activity for P1 from left shoe.

- Since (LP_{R_R}) >(LP_{s_R}) and (LP_{R_R}) >(LP_{St_R}), running activity can be correctly classified.
- Since (LP_{s_s}) >(LP_{R_s} and (LP_{s_s}) >(LP_{St_s}), sitting activity can be correctly classified.
- Since (LP_{St_St}) >(LP_{s_St}) and (LP_{St_St}) >(LP_{R_St}), standing activity is correctly classified.

We also carried out similar analysis for four activity system and found out that this approach can correctly classify activities in four activity setting. In our case, these four activities are: sitting, standing, walking and running.

6 Result

Based on our preliminary experiments, we expanded our system to 8 activity system. These 8 activities are: cycling, running, climbing stairs down, climbing stairs up, walking, sitting and driving.

For each activity, we worked with 3000 samples and we divided the samples in 20 windows making each window with 150 samples. We applied GMM with EM for training. Here we are working with data of single subject. Parameters are as follows:

Number of Gaussian Mixture: 5
Time Lag, $\tau = 5$
dimension, d =6

Time Lag, τ and dimension, d were empirically obtained. We adopted a grid search approach and observed the values of τ and d for which activity detection accuracy is best. The following Table 3 shows a confusion matrix derived from applying our approach on accelearation along X-axis. We use the following symbols in the tables: C for cycling, R for running, Sd for downstairs, Su for Upstairs, St for standing, W for walking, Si for sitting and D for driving. Miss-classifications are shown in red color. Out of $8 \times 20 = 160$ time segments, 11 time segments are misclassified (93.13% accuracy).

W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12	W13	W14	W15	W16	W17	W18	W19	W20	Actual
C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	**C**
R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	**R**
Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Su	Sd	Sd	Sd	Sd	Sd	Su	Su	**Sd**
Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	**Su**
St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	**St**
W	W	W	W	W	W	W	W	W	W	W	W	W	W	Sd	Sd	Sd	Sd	Sd	Sd	**W**
Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	**Si**
D	D	D	D	D	D	D	D	D	D	D	D	D	D	C	C	D	D	D	D	**D**

Table 3: Confusion matrix using GMM based on accelerometer data along *X-axis*.

A much better accuracy is achieved by using accelerometer along Y-axis as obvious from the following table (Table 4)

W1	W2	W3	W4	W5	W6	W7	W8	W9	W10	W11	W12	W13	W14	W15	W16	W17	W18	W19	W20	Actual
C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C	**C**
R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	**R**
Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	**Sd**
Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	**Su**
St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	**St**
W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	**W**
Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	**Si**
D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	D	**D**

Table 4: Confusion matrix using GMM based on accelerometer data along *Y-axis*.

We generated similar confusion matrix based on data from P1, P2, P3, P4, P5, P7. In each case, there are $20 \times 8 = 160$ classifications. For P1, there are 41 miss-classifications. Similarly, for P2, 31; for P3, 9; for P4, 29; for P5, 9; and for

P7, there were 12 miss-classifications. As an example of performance of pressure sensors, we show the confusion matrix of P3.

W1	**W2**	**W3**	**W4**	**W5**	**W6**	**W7**	**W8**	**W9**	**W10**	**W11**	**W12**	**W13**	**W14**	**W15**	**W16**	**W17**	**W18**	**W19**	**W20**	**Actual**
C	C	C	St	C	Si	C	C	C	C	C	C	C	D	C	C	C	C	C	C	**C**
R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	**R**
Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	Sd	**Sd**
Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	Su	**Su**
St	St	St	St	Si	St	St	St	St	St	St	St	St	St	St	St	St	St	St	St	**St**
W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	W	**W**
St	St	St	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	Si	**Si**
D	D	D	D	D	D	D	D	D	D	D	D	C	D	D	D	D	D	D	C	**D**

Table 5: Confusion matrix using GMM based on pressure sensor data *P3* of left shoe.

7 Conclusion

Most promising aspect about time-delay embedding with GMM is that significantly good accuracy is obtained just from analysis of small number of sensor data. We have not applied this approach on gyroscope data, neither did we apply on pressure data from right shoe. We are now working to develop a fusion of this approach. For example, we can generate decisions from P1, P2 and P3 sensor and obtain the final decision from fusion of multiple sensor for any time segment. Even with one sensor, we have significant accuracy.

Consequently, it will be possible to reduce computational complexity if we use small number of sensors. As a result, a more energy-saving system can be reality. Such systems can reduce energy cost in two ways. First, as fewer sensor data will be transmitted over Bluetooth, energy can be saved by reducing energy for Bluetooth transmission. Second, as there is less data and consequently, less computation, reduced energy will be needed for computation. Memory and computational saving is significant as it is most likely that activity detection applications will run on resource-constraint smart phones. Most activity system demands real-time detection of accuracy. Reduction in computational cost implies extended battery life. Computational power and battery life, both are scarce resource in cell phones and an algorithm that protects these resources are obviously preferable.

8 Future Works

Our experiment with time-delay embedding shows exciting outcome. This prototype is getting ready to use in Taiwan for elderly care research. Dr. William Chu, the prominent researcher in the field of elderly care, is helping us to setting up the prototype there. In the future, we plan to incorporate multimodal fusion approach with time-delay embedding to improve accuracy further.

References

[Capela et al., 2016] Capela, N., Lemaire, E., Baddour, N., Rudolf, M., Goljar, N., and Burger, H. (2016). Evaluation of a smartphone human activity recognition application with able-bodied and stroke participants. *Journal of NeuroEngineering and Rehabilitation*, 13.

[Frank et al., 2010] Frank, J., Mannor, S., and Precup, D. (2010). Activity and gait recognition with time-delay embeddings. In *AAAI*. Citeseer.

[Kawsar et al., 2015] Kawsar, F., Hasan, M. K., Love, R., and Ahamed, S. I. (2015). A novel activity detection system using plantar pressure sensors and smartphone. In *Computer Software and Applications Conference (COMPSAC), 2015 IEEE 39th Annual*, volume 1, pages 44–49. IEEE.

[Kwapisz et al., 2011] Kwapisz, J. R., Weiss, G. M., and Moore, S. A. (2011). Activity recognition using cell phone accelerometers. *ACM SigKDD Explorations Newsletter*, 12(2):74–82.

[Lee and Mase, 2001] Lee, S.-W. and Mase, K. (2001). Recognition of walking behaviors for pedestrian navigation. In *Control Applications, 2001.(CCA'01). Proceedings of the 2001 IEEE International Conference on*, pages 1152–1155. IEEE.

[Miluzzo et al., 2008] Miluzzo, E., Lane, N. D., Fodor, K., Peterson, R., Lu, H., Musolesi, M., Eisenman, S. B., Zheng, X., and Campbell, A. T. (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*, pages 337–350. ACM.

[Sauer et al., 1991] Sauer, T., Yorke, J. A., and Casdagli, M. (1991). Embedology. *Journal of statistical Physics*, 65(3-4):579–616.

[Shaji et al., 2016] Shaji, S., Ramesh, M. V., and Menon, V. N. (2016). Real-time processing and analysis for activity classification to enhance wearable wireless ecg. In *Proceedings of the Second International Conference on Computer and Communication Technologies*, pages 21–35. Springer.

[Shu et al., 2010] Shu, L., Hua, T., Wang, Y., Li, Q., Feng, D. D., and Tao, X. (2010). In-shoe plantar pressure measurement and analysis system based on fabric pressure sensing array. *Information Technology in Biomedicine, IEEE Transactions on*, 14(3):767–775.

[Takens, 1981] Takens, F. (1981). *Detecting strange attractors in turbulence*. Springer.

[Yang, 2009] Yang, J. (2009). Toward physical activity diary: motion recognition using simple acceleration features with mobile phones. In *Proceedings of the 1st international workshop on Interactive multimedia for consumer electronics*, pages 1–10. ACM.