Proceedings of the ASME 2007 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference
IDETC/CIE 2007
September 4-7, 2007, Las Vegas, Nevada, USA

DETC2007-35844

MOTION CLASSIFICATION USING WIRELESS SENSORS FOR ACTIVITY MONITORING IN FIREFIGHTING

Nyugen Nyugen
Department of Computer Science
University of California
Riverside, CA 92521

V. Sundararajan*
Department of Mechanical Engineering
University of California
Riverside, CA 92521

Victor B. Zordan
Department of Computer Science
University of California
Riverside, CA 92521

ABSTRACT

Firefighter' chief reason for entering a burning structure is to search for and rescue potential victims. Currently, their primary method for communication is an often-congested two-way radio which the firefighters use from within a burning building to relay their activities (and other information) to an external battalion chief. In response to discussions with firefighters in the field, we introduce an approach for automatically segmenting and classifying a select set of activities using wireless accelerometers attached to the human body. The activities we focus on are the ones that are most commonly conducted by firefighters and that are important to the battalion chief for understanding the ongoing search and rescue. In our implementation, sensors continuously measure the acceleration of a small number of body segments and transmit data back to a central base station. At runtime, our system classifies data for short intervals, relying on training examples of the activities of interest. We show that our approach can appropriately detect motions in real-time without significant latency using as few as two accelerometers.

1 INTRODUCTION

When a team of firefighters enter a blazing building, a central hub of information is the battalion chief who assesses the progress of the team and makes supervisory decisions regarding the progress of the search and rescue effort. After our discussions with local firefighters, we assessed that it would be valuable to the battalion chief to monitor not only location but also the types of movements of her/his team members so that s/he can make better informed decisions. With their low cost and availability, wireless sensors, such as accelerometers can be used to measure each of the firefighter's movement for specific body parts and with this information we can automatically provide details that will relieve the congestion on the current communication channel, a shared two-way radio. Data collected from different body segments of a firefighter within a time frame can be used to determine his/ her behavior in real time (Figure 1). Once an activity is classified, a graphical system can represent information to the team leader corresponding to that activity.

Our goal in this project is finding a mechanism to correctly predict activities in real-time using minimum number of sensors. There are three main components in this process: data acquisition, data processing (discretization, normalization, analysis) and classification. The first step is done on the sensors and the third step is done at the base station. Though the second step can also be done at the base station, in order to reduce the data transmission rate, we experimented with performing some processing on the sensor node so that only the feature patterns computed over the collected data is transmitted back. This reduction in bandwidth allows for more sensor nodes to use the channel and also reduces the energy used by the sensor nodes.

In this paper, we introduce a solution to track human activities using accelerometers. The acceleration signals received from

*Address all correspondence to this author.

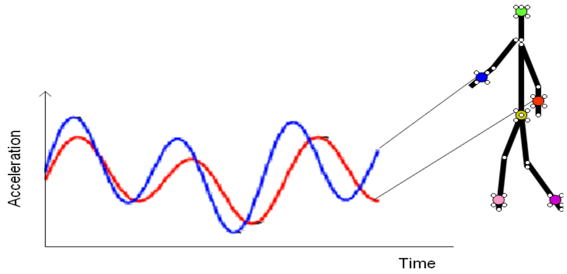

Figure 1. Sensor data representation for human motion.

each sensor represent movement of the body segments on which it is attached. Within a small periodic time-span, the collection of data is compared against the training data to classify the current movement of the object. The specific aims of our project are to: develop the real-time algorithm for classification; determine location and number of sensors; and discover an appropriate frequency of data collection so that for each time period, there is enough information collected and transmitted back to the base station to classify an activity while remaining responsive in light of the anticipated emergency time critical response.

We use two different methods to solve our classification problem. The first uses dynamic time warping (DTW) to gauge similarity between testing and training time series. The other finds feature patterns for classification from training and clusters them using a support vector machine (SVM). Our experiments reveal that the SVM approach results in more reliable classification and that as few as two sensors are good enough to classify the critical behaviors of interest.

This paper is divided into 6 sections. The next section provides some information about current related research and challenges. Section 3 describes our experimental set-up and data recording method used in this project. Section 4 describes the processing of the data collected and classification for behavior. Section 5 is about the experiments and results given by the classification methods. Section 6 concludes the project and points out possible future research directions.

2 Related Work

Sensing and augmentation with wireless technology is not a new idea applied firefighting and several efforts have been set forth to challenge this specific goal. Two notable examples are the Fire Information and Rescue Equipment (FIRE) effort between UC Berkeley's Mechanical Engineering department and the Chicago fire dept [1] and also the so-called "Precision Personnel Locator" (PPL) project at Worchester PolyTechnic Institute [2]. The FIRE project sets it aims in two areas, first providing visual information such as floor layouts to the firefighters via small head-mounted displays and second by sensing for the location of the firefighter and dangers, such as the fire hot spots [3,4]. PPL is focused on the indoor position tracking of individuals (firefighters) and has conducted a thorough investigation of their proposed technique [5, 6]. Work on indoor tracking in general has also been the focus of several other reseach thrusts [7, 8]. In contrast, these projects largely focus on the position of the firefighter while we complement this work by additionally sensing and classifying the type of activity that a firefighter is performing.

2.1 Human Motion Classification

Many researchers have been working on classifying human activities from video sequences [9] [10]. Chai and Hodgin proposed an approach to classify human motion from low control signals by employing video cameras and a small set of retro-reflective markers to create a low-cost system [11]. Joo Hong et al. [12] used accelerometers in hospitals to remotely track patients' movement. In this work, a single accelerometer is attached on elderly patients or movement impaired people to help physicians track if they are active or not. The methods used was principal component analysis. Noel Kaijsers et al. [13] used statistic neural networks to classify human movements. The neural network was used to assess the severity of levodopa-induced dyskinesia from scores achieved by accelerometers' signals. Human motion classification is an important research topic with applications in surveillance. Many researchers have looked at analyzing video data for this purpose. In a different direction, this project focuses on the problem when the vision on the object is occluded and limit data bandwidth can be transferred in the environment. In this case, the accelerometers are used to transmit data via FM radio. In addition, we target this project for application in firefighting which limits the domain of expected behaviors.

3 System Overview

The system layout is shown schematically in Figure 2. Battery driven wireless accelerometers transmit data to a base station during run-time. The base station processes the data and performs classification of the activity. After our discussion with the firefighting team, we selected five different activities recorded:

1. Walking
2. Running
3. Crawling
4. Pulling/dragging a heavy object (simulation of saving an unconscious victim)
5. Idling

The time frame should be defined so that the amount of data collected within that frame is large enough to classify different activities. However, the data should be short enough to provide a

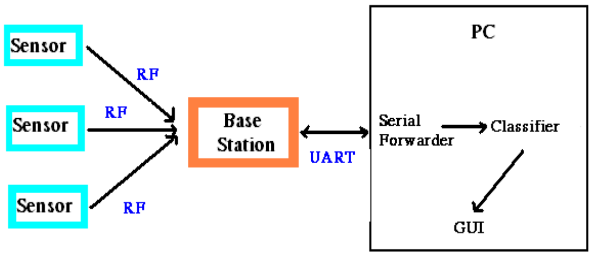

Figure 2. Sensors collect data and send it to the base station via RF signals. The base station is connected with the PC with serial port. The application classifies the data and animation is displayed on the screen.

reasonable update to the battalion chief. Thus, for each activity to be predicted, we decided to record continuous data for 1.5 seconds - intuitively this period is plenty fast enough for the chief to take action.

3.1 Wireless sensors

The accelerometers are mounted on the mica2dot mote (wireless sensor node) (Figure 3) available from Crossbow Technologies [**?**]. The node possesses processing, storage and communication abilities and is programmable using the NesC language and the TinyOS operation system. TinyOS is an embedded operating system developed by the University of California, Berkeley. It is written in NesC, a new language for programming structured component-based applications. The motes are programmed to pack 5 data samples (in both direction x and y, 2 bytes each) into a packet and transmit it to the base station.

Each sensor is programmed with its own identification number. One spare sensor is programmed and used as a controller. The controller sends synchronous signals to all sensors to command the sensors to start measuring accelerations. Figure 4 shows the locations on the human body where we attach the accelerometers. Six sensors are numbered from 1 to 6. There are two sensors on the two wrists, one on each ankle, one on the forehead and one on the waist.

3.2 Data Acquisition

The training person performs each of the activities continuously for 30 seconds. Data is simultaneously recorded from different body segments of the training person. Data sampling rate is set to 20 Hz. Accelerations are measured continuously for 1.5 seconds for each testing data set. Thus, each time series has 30 data points. Each packet sent by the sensors is packed with 5 data samples. When there are lost packets, the data is made up by averaging out with their successors and predecessors packets. Raw samples of the data appear in Figure 5 showing the potential variation for a single behavior for a single subject.

Figure 3. (a) MTS510 accelerometers designed by Crossbow can measure acceleration in two dimensions ranging from -2g to +2g. With the size of a thumb, they can easily be attached on the body. (b) MIB510CA board was developed by UC Berkeley. This board serves as a base station where information is gathered and then transmitted to the processing unit via RS-232 serial interface. The acceleration data is transmitted from the sensors to the base station by RF communication.

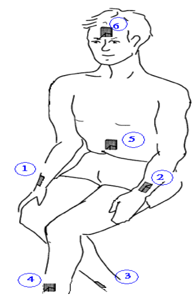

Figure 4. Sensor Attachment

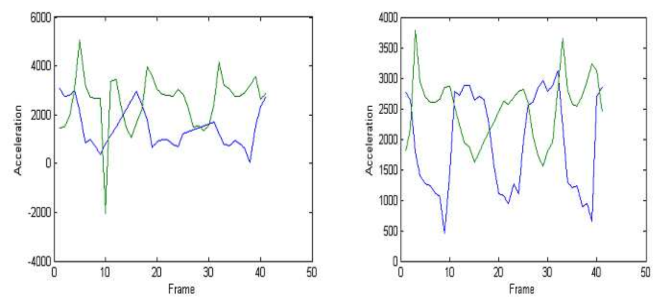

Figure 5. Two data samples measured for x direction from a running activity. The blue and green are for the accelerators on left and right hand respectively (performed by the same person.)

4 Classification

Motion data of the object, both training and testing, are represented as time series of observed acceleration. In this section, we discuss two different classification methods: dynamic time warping (DTW) and support vector machine (SVM) on which

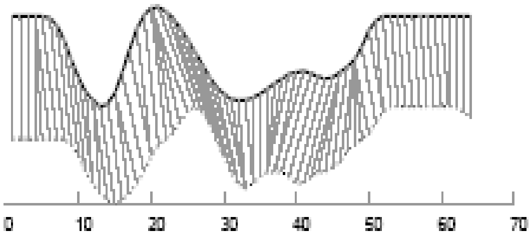

Figure 6. This figure shows an example of 2 similar sequence but different in time. DWT dynamically warp the time axis to achieve best match.

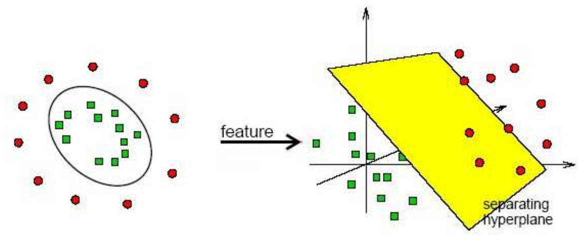

Figure 7. A classification problem which may be complex in a low dimensional space becomes easier when using separating hyperplane in higher dimension.

our solution depends. In addition to the classification method, since we are targeting an online tool, both the speed of the algorithm is a factor as well as its input requirements. The latter adds an additional constraint that an important need of the project is for the sensor to have the capability of analyzing the data required before transferring its signal to the base station. We include these factors in our discussion as well as the specifics of each approach.

4.1 Classification by Dynamic Time Warping (DTW).

The first method we use to classify the activity is lining up the testing time series with the training ones to find the one with minimum distance. The distance between the testing and training time series is the similarity between trends of movements along time axis. The method that we use is called dynamic time warping (DTW) [14]. DTW has been widely used in video, audio, and graphics where streaming data can be represented in time series and the problem is solved in polynomial time. DTW allows more sophisticated distance measurement by warping the time axis on the time series to achieve better alignment. For this reason, DTW can match the similar activity patterns which vary in time and speed. For example a person may walk slow or fast but his activity is still considered walking. In Figure 6, an example of two time series and their best time-warped alignment illustrates DTW.

Acceleration data is normalized both in the training and testing steps. Euclidean distance between the time-warped algorithm is the metric function. The distance between a test example and the training data is computed. The nearest neighbor is classified as the result activity. Since all of the data sample is required for a single classification, the information for each sensor must be transmitted to the base station. However, as long as the training set was kept small, we were able to achieve real-time performance for DTW.

4.2 Classification by SVM Model

Support Vector Machine (SVM) is a useful technique for data classification. SVM operates by finding a hypersurface in space of input data. We take a set of training data to create a model which contains attribute information to predict the target class of testing data. Each training data sample includes a class label and several attributes. The created SVM model contains a set of linear separating hyperplanes or hypersurfaces with maximal margin. Figure 7 demonstrates SVM in a pictorial fashion.

Formally, let T(x_i,y_i) (i = 1, ... , l) be a set of training samples where $x_i \in R_n$ are the samples attributes and y_i are the sample labels, the training vectors x_i are then created by solving the following optimization problem [15].

$$\begin{aligned} &min_{w,b,\xi} \quad \tfrac{1}{2}w^T w + C\textstyle\sum_{i=1}^{l}\xi_i \\ &\text{subject to} \quad y_i(w^T\phi(x_i)+b) \geq 1-\xi_i\ , \\ &\qquad \xi \geq 0. \end{aligned}$$

Training vectors x_i are mapped to multi-dimensional space by function and separating hyperplanes are found with maximal margin. $K(x_i, xj) = \phi(x_i)^T\phi(x_j)$ is called kernel function. There are four different basic kernels: linear, polynomial, radial basis function and sigmoid. In this project, we use polynomial kernel where $K(x_i,x_j) = (\sigma x_i^T x_j + r)^d, \sigma_i > 0$. We use two accelerometers, each testing sample (x_1,y_1,x_2,y_2) contains four instance vectors x_1,y_1,x_2,y_2. The feature vector that feeds into the SVM machine is f = [mean(x_1),std(x_1),mean(x_2),std(x_2)]. Here, we only need to send the feature vector of the sample after computing it on the sensor. Also, the training set can be larger than described for DTW.

5 Implementation and Results

We collect training data by continuously recording each of the defined activities for 30 seconds from a single person (see Figure 8). In the first experiment, we use complete set of six sensors for classification. Testing sequences are recorded both

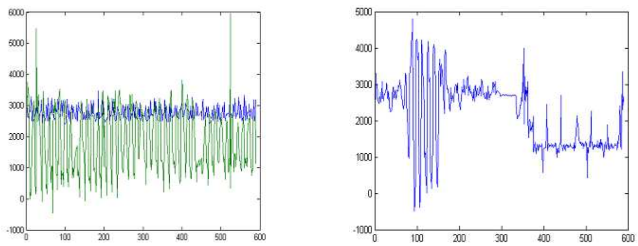

Figure 8. (8a) shows acceleration data recorded on training person A left hand for walking (blue) and running (green) activities. (8b) shows the testing sequence performed by the same person with the order walking, running, walking, idle and crawling.

each subject. We first test short sequence of single activity with DTW method and then on continuous data.

Sensor Weighting. Empirically, we assign different weights to different accelerometers (the closer to the root the accelerometer is, the higher the weight). We consider the position on the torso where we attach the sensor as the root and it is inferred that the sensors on the limbs have lowest weight values. After performing several experiments we can approximate the optimal sliding window size at the data discretization steps. The window size is considered optimal when it yields fairly high distances among the classified classes of activity. This step is only used with DTW.

With the SVM approach, for each of the activities, we segment continuous training data into a number of training samples. From each of training samples, the features are extracted and we use standard deviation and mean values as the feature variables. Testing data is naively segmented into samples of thirty data points. Mean and standard deviation are computed for each sample. Offline, an SVM model is built from the entire collection of training data and is ready to be used for classification steps. For the method using DTW to compare two time series, we use sequential approximation (SAX) [16] to convert the recorded data to a symbolic time series of discrete symbols. The advantage of using SAX is it allows lower bounding of Euclidian distance and dimensionality reduction.

5.1 Experimental Results

Our first grouping of experiments was performed only using DTW to assess a few key aspects of our data. That is, is there enough information in the accelerometer data set to classify the motions of interest? And, second, what is the minimum number of sensors that can be used and what should their placement be?

5.1.1 Experiment 1: Single Subject, Complete Set of Sensors (Base Case) Data from all 6 sensors are used with DTW radius of 1 (equivalent to non time-warped Euclidean distance) to calculate the distance between the testing time series and the training data for a single subject. For each class of activity, we recorded three data sets for testing. The result appear in Table 1. In this case, the total accuracy is around 70 percent. Running and walking are not well distinguished.

	crawling	idle	pulling	running	walking
crawling 1	102	396	374	728	460
crawling 2	107	408	382	739	470
crawling 3	73	395	373	735	451
idle 1	400	18	77	687	190
idle 2	402	25	86	705	194
idle 3	405	16	81	692	189
pulling 1	380	86	68	632	197
pulling 2	388	83	53	635	199
pulling 3	377	87	71	628	202
running 1	695	690	628	375	605
running 2	774	708	659	844	636
running 3	633	586	527	627	549
walking 1	465	205	209	618	175
walking 2	460	181	189	609	231
walking 3	448	180	184	600	105

Table 1. This table shows the one-to-one distance between the testing and training behaviors. Testing time series on left column are lined up with training time series on the top row.

5.1.2 Experiment 2: Tuning for improved classification In order to improve the performance, we allow a shift in the testing of the time series within a small time window and choose the position returning smallest distance. We also increase the DTW radius to 3 to gain better result. This improves the classification quality but increases cost in term of running time. With some tuning based on trial and error, we were able to make the result improve to 100 percent when we assigned different weights to different sensors so that the inner body part such as waist and forehead have higher values than the ones on four limbs.

5.1.3 Experiment 3: Reducing Number of Sensors An important goal for us was to assess the minimum number of sensors and their placement to make our system as practical as possible. In this experiment, we ran various subsets of

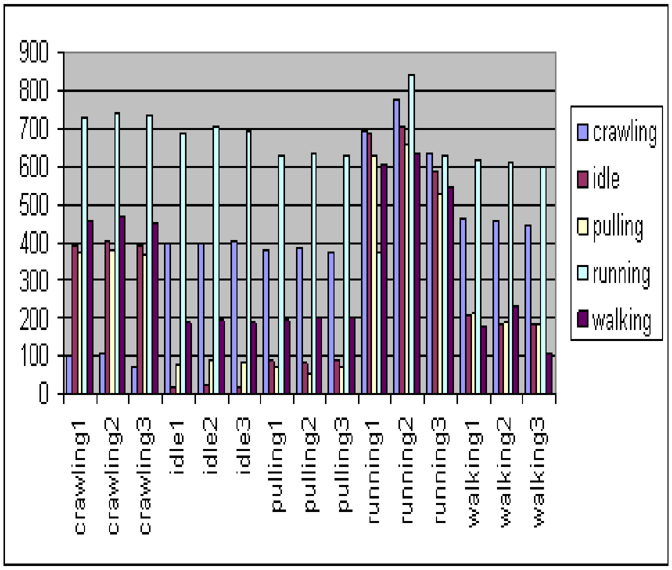

Figure 9. Distance between testing and training behaviors

Sensor included	num of misses w/o DTW	num of misses w/ DTW
1,2,3,4,5,6	1	0
1,2,3,4,6	1	0
1,2,3,4,5	1	0
2,3,4,6	2	0
1,3,6	1	0
2,6	3	0
1,2	2	1
1,3	4	0
2	3	1
4	7	5

Table 2. Different combinations of sensors give different set of misses. The result degrades with the number of sensors. From the result table we can see that with the set of five behaviors as defined in this project, using two sensors is good enough for classification process.

the sensors to see if there is redundancy and we can minimize the number of sensors required for the classification processes. We remove sensor 6. The set of remaining 5 sensors still performed well and gave similar results as the complete one except the dragging was classified as pulling. Intuitively, both sensor 6 (forehead) and sensor 5 (waist) measure the movement of the whole body and there is redundancy in their data. We perform test on different combinations shown in Table 2.

5.1.4 Experiment 4: Testing on continuous, non-segmented data

To bring our experimental set-up closer to the real world case, we recorded subjects performing continuous combinations of the activities of interest over 30-second intervals. These tests are performed on 3 different subjects, labelled persons A, B and C. The training data was recorded on person A previously. Each activity sequence is recorded continuously for 30 seconds with 5 different activities recorded for testing: running, walking, crawling, pulling and idle. These experiments only use two sensors 1 and 3 (one on right hand and the other one on right ankle). Results appear in Figure 9.

From our results, we find that the DTW method is reasonably good only when testing data is comes from A (training person). However, there are still several misses that we can notice at the transitions between the activities. It is because these are mixtures of two defined activities and theoretically they are the activities that have not been defined during training processes. The classification quality degrades when testing on person B and C, i.e. when testing activities are performed on subjects other than the subject that performs training activities.

5.2 Comparing classification methods

Once our initial testing phase was complete, we set out to compare DTW to SVM. In brief, DTW is thorough but slow. Its speed forced us to classify our data in the last experiment with a limited training set, although we believe if we compared with a larger dataset it would do very well based on our 100 percent accuracy in our experiments from Phase 1. In the second phase, our goal was to contrast implementations of both methods on the continuous data.

To start, SVM is run on a large test data set (represented by each point in Figure 10. Looking at the acceleration signals generated by the right hand for all 5 different activities (Figure 10), we can visually distinguish the activities and segment the data sequence by looking at the frequencies and amplitudes which are represented by the means and standard-deviations of those segments.

Next we perform all six previous tests from Experiment 4 with SVM and the results improve significantly (Figures 11 and 12), especially at the transition points. Intuitively, this is because the transition between activities A1 and A2 lies on or close to the boundaries created by the hyperplane in SVM model, therefore the activity is classified as either A1 or A2. (Note, we feel a small transition error is acceptable since it is very short and a simple filtering process could remove such artifacts in a commercial implementation where a firefighter is not changing activities as quickly as was done in our experiments.)

In addition, we experimented with adding different data sets

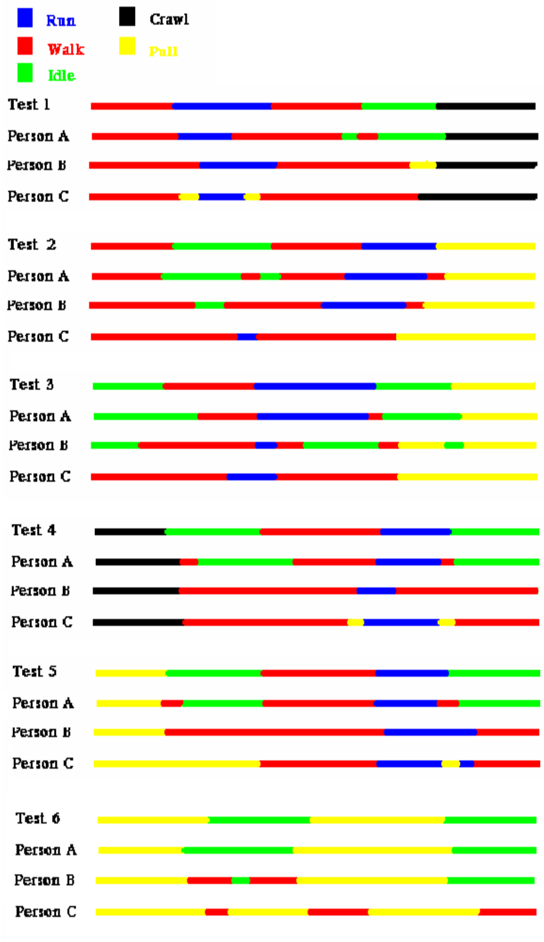

Figure 10. This image shows the outcome of performing classification using DTW on continuous data and on different persons.

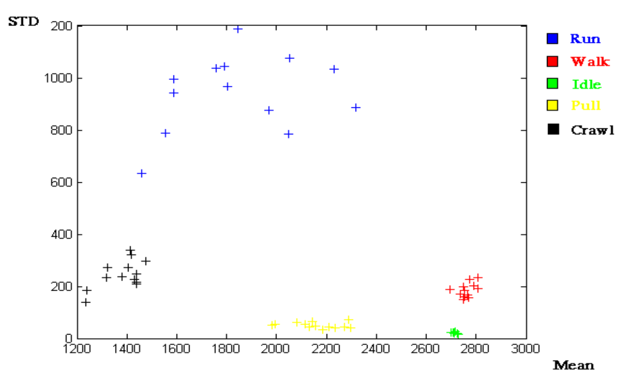

Figure 11. This figure shows the clusters of training data on 2-dimensional space with mean and standard deviation.

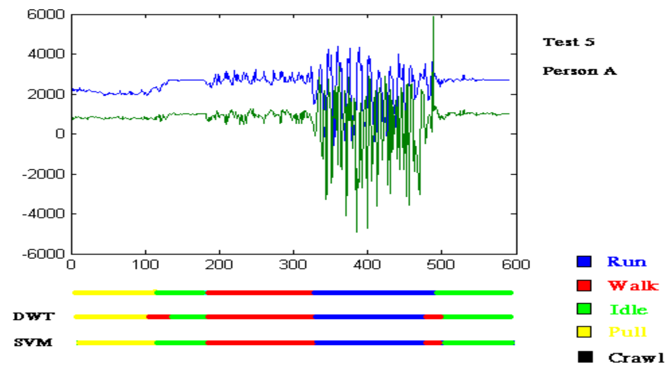

Figure 12. This figure shows the activity sequence 5 (pull, idle, walk, run, idle) performed by the same person whose training data is recorded. SVM outperforms DWT by correctly predicting idling after the pulling activity. The transition between running and walking is reasonable because the subject needs to slow down before idling.

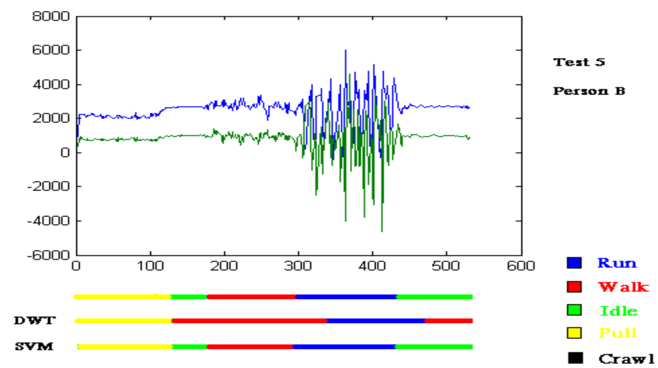

Figure 13. This figure shows the activity sequence 5 performed by person B using training data of person A. DWT miss-classifies idling with walking while SVM makes perfect prediction.

to the training of the SVM model. A set of results from these examples appear in Figure 13.

Finally, we also consider the features that carry the relative information between the sensors by looking at the cross-correlation between the body segments in computing the SVM (Figure 14.) One feature that we extract is that the maximum value of the normalized cross-correlation appears to be between Sensor 1 and 2.

However, when we include this feature to generate the SVM model, the result does not improve. However, we may decide to use cross-correlation if we were to include more a complex set of activities. For the five defined activities, means and standard-deviations of two sensors appears good enough for our classification problem.

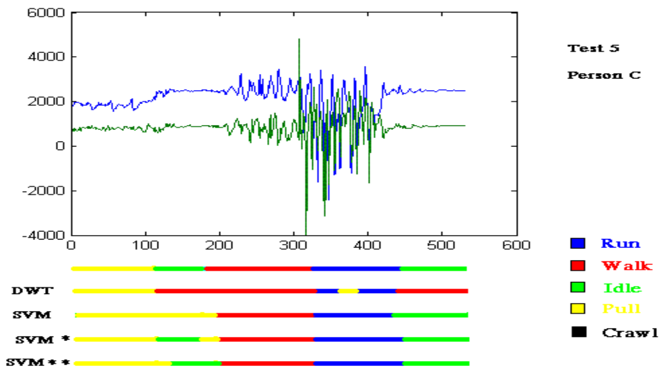

Figure 14. This figure shows improvement made by SVM with ordered activities for person C. Notice idling is classified as walking by DTW. SVM trained on person A does a little better job, however the first idling sequence is still classified as pulling. When we include a small amount of training data by C to the model, the result shown by SVM* is close to perfect except at the transition between idling and walking, the classified activity is pulling. SVM** is tested with training data recorded by person A and B. Its even better than SVM* because its model was created by twice the training samples (i.e. from person A and B).

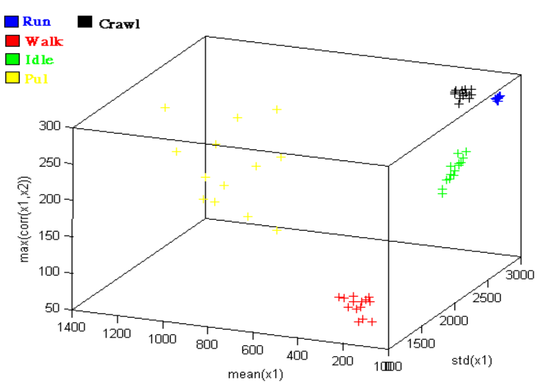

Figure 15. This figure shows the clusters of training data in 3-dimensional space with mean value and standard deviation of x-acceleration read on Sensor 1 and the Z-axis as the maximum value of the cross-correlation between readings from Sensor 1 and Sensor 2.

6 Conclusions

We perform classification on a set of specific behaviors identified from discussion of activities performed during search and rescue in firefighting. Our goal is a real-time system which utilized the smallest number of sensors and took advantage of onboard processing of sensors to minimize message passing (saving battery power) while performing good classification. In addition, sensor placement was another unknown in our investigation. Our method uses wireless sensors which measure acceleration and have limited processing capabilities. In our experiments, we distinguished between two classification approaches. We show results with a high-level of accuracy and conclude that two sensors on hand and on ankle are sufficient for this classification.

Several challenges make this problem difficult. Data can be very noisy because of the complexity of environment. Data is transmitted via FM radio; therefore, the rate of packet loss can be high. The accelerometers used measure only two dimensions. Moreover, the classification must be done quickly (within each time frame) so that the system can work in real-time. Finally, the sensor memory is limited, therefore training data and code must be represented compactly.

We compare DTW and SVM classification method in several scenarios. However, we find that classification using SVM is more advantageous because of the following reasons:

1. SVM gives better results at transitions between two define activities.
2. SVM reduces the data transmission since only the features of the collected data are sent. They can be easily computed on the sensors.
3. SVM can use the features that carry relative information between the sensors while DTW handles sensors separately.

For the best results, we found that training on multiple persons was helpful. Possibly, in the field, we might recommend a test stage with individual firefighters (though this was not required in our experimental set-up.) It is probably better to use 3-dimensional accelerometers so that we can use the absolute acceleration by taking the values recorded on the three dimensional space. In the real application, we may need more than 5 activities and adding more dimensions to the data space by considering cross-correlation information (or other features) may be required. Latency is on the order of a second and would be acceptable depending upon the application and situation.

In the future, we would prefer improvement by performing all classification process on the sensors and the only information sent to the base station is the classified result. This is desirable because when the number of objects is high, the higher rate of data transmission may overload the base station. Another reason is that data is sent over FM radio which is easily lost due to environment. In term of capacity, on-sensor classification may be trivial because the training data does not require much memory. However, the challenge is the processing speed and the fault tolerance of the technique. Also, a robust system must handle the case when a processing sensor may be dead. This leads to a new problem, fault tolerance, in sensor network and provides exciting directions for future extenstions.

REFERENCES

[1] HTTP://FIRE.ME.BERKELEY.EDU/, 2007. UCBerkeley FIRE project, WWW website.

[2] HTTP://WWW.ECE.WPI.EDU/RESEARCH/PPL/, 2007. WPI PPL project, WWW website.

[3] Wilson, J., Steingart, D., Romero, R., Reynolds, J., Mellers, E., Redfern, A., Lim, L., Watts, W., Patton, C., Baker, J., and Wright, P., 2005. "Design of monocular head-mounted displays for increased indoor firefighting safety and efficiency,". In SPIE Volume 5800, pp. 103- 114.

[4] Steingart, D., Wilson, J., Redfern, A., Wright, P., Romero, R., and Lim, L., 2005. "Augmented cognition for fire emergency response: An iterative user study". In Proceeding of the 11th International Conference on Human-Computer Interaction (HCI).

[5] Parikh, H. K., Michalson, W. R., and Duckworth, R. J., 2004. "Performance evaluation of the rf receiver for precision positioning system". In ION NTM.

[6] Cyganski, D., Duckworth, J., Makarov, S., Michalson, W., and Orr, J., 2007. "Wpi precision personnel locator system". In Institute of Navigation, National Technical Meeting.

[7] Werb, J., and Lanzl, C., 1998. "Designing a positioning system for finding things and people indoors". In IEEE Spectrum, Vol. 35, No. 9.

[8] Krishnamurthy, P., Pahlavan, K., and Beneat, J., 1998. "Radio propagation modeling for indoor geolocation applications". In Proc. IEEE PIMRC '98.

[9] INTERNATIONAL WORKSHOP ON HUMAN ACTIVITY RECOGNITION AND MODELING, 2005. Segmentation and Classification of Human Activities, HAREM.

[10] Rittscher, J., and Blake, A., 1999. "Classification of human body". *ICCV*, pp. 634–639.

[11] Chai, J., and Hodgins, K., 2005. "Performance animation from low-dimensional control signals". *ACM Trans. Graph,* **24**, pp. 686–696.

[12] Joo Hong, e. a., 2005. "Classification technique of human motion context based on wireless sensor network". In Proceedings of the 2005 IEEE, Engineering in Medicine and Biology 27th Annual Conference.

[13] Noel Kaijsers, e. a., 2002. "Automatic assesment of levodopa-induced dyskinesia in daily life by neural networks". *Movement Disorder,* **18**, pp. 70–80.

[14] FIRST SIAM INTERNATIONAL CONFERENCE ON DATA MINING, 2001. Dynamic Time Warping with Higher Order Features, SDM.

[15] Boser, B., Guyon, I., and Vapnik, V., 2005. "A training algorithm for optimal margin classifiers". In In Proceedings of the Fifth Annual Workshop on Computational Learning Theory, pp. 144-152. ACM Press.

[16] Lin, J., Keogh, E., Lonardi, S., and Chiu, B., 2003. "A symbolic representation of time series, with implications for streaming algorithms". In proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery.