

Human Activity Recognition: Accuracy across Common Locations for Wearable Sensors

Daniel Olguín Olguín, Alex (Sandy) Pentland
MIT Media Laboratory, Human Dynamics Group
20 Ames St. Room E15-383
Cambridge, MA 02139
{dolguin, sandy}@media.mit.edu

Abstract

In recent years much work has been done on human activity recognition using wearable sensors. As we begin to deploy hundreds and even thousands of wearable sensors on regular workers, hospital patients, and army soldiers, the question shifts more toward a balance between what information can be gained and their broad immediate user acceptance. In this paper we compare the activity classification accuracy of four different configurations of accelerometer placement on the human body using hidden Markov models (HMMs). We find the classification accuracy of a single accelerometer placed in three different parts of the body and evaluate whether there is a significant improvement in recognition accuracy by adding multiple accelerometers or not. We also find the number of hidden states that best models each activity by achieving the lowest test error using K-fold cross-validation.

1 Introduction

Being able to automatically recognize human motion patterns using unobtrusive wearable sensors can be useful in monitoring the elderly in their homes and keep track of their daily activities and behavioral changes. This could lead to a better understanding of numerous medical conditions and treatments. Other applications of human activity recognition range from context aware computing to physical training, physical rehabilitation, and military applications such as intelligent outfit design for soldiers.

In this paper we study different configurations of accelerometer placement to classify human activities that are frequent in at least one of these application areas. Eight different activities were modeled using HMMs and continuous Gaussian observation vectors. Three wireless accelerometers (*MITes: MIT Environmental Sensors*) [6] were placed

on different parts of the body: right wrist (A1), left hip (A2), and chest (A3). We selected these locations on the basis of lengthy discussions with three potential user groups (office workers, hospital patients, and army soldiers) concerning their general acceptance of wearable sensors. We found that there was broad consensus about the acceptability of the chest and hip locations, and so we are developing a wearable electronic badge that will be worn on the chest. This badge will be able to share information with a Bluetooth-enabled cellular phone that could be worn on the hip and might have a second accelerometer.

2 Related Work

Previous work in human activity recognition using accelerometers has shown that it is possible to classify several postures and activities in real time. In [3], the authors developed a two-layer model that combined a Gaussian mixture model with first-order Markov models to classify a range of activities including: sitting, walking, biking, and riding the subway. A single 3-axis accelerometer placed on the torso was used. In [4], the authors combined data from three accelerometers and two microphones placed on different body locations to classify activities performed in a wood shop with 84.4% accuracy. They modeled most of the activities using single Gaussian hidden Markov models. The number of hidden states to model each activity was selected through visual inspection. In [1], several algorithms to classify twenty different physical activities from data acquired using five 2-axis accelerometers were evaluated with an overall recognition rate of 84%.

In [5], the authors used decision trees to classify six different activities with a single accelerometer, placed on six different body positions commonly used for wearing electronic devices, with accuracies ranging from 16.7% to 92.8% depending on the position of the accelerometer and the features used.

3 Data Collection and Processing

Three different subjects were asked to perform the following sequence of activities: 1) Sit down, 2) Run, 3) Squat, 4) Walk, 5) Stand, 6) Crawl, 7) Lay down (on the chest), and 8) Hand movements (while standing). The data collection process was repeated three times for each subject. 90 seconds of data were collected for each activity and labeled according to the start and stop times of each activity performed. The data were divided into nine datasets, each of them containing 80 observation sequences of one-second duration. An HMM was trained for each activity using eight datasets (640 observations) and tested on the ninth dataset (80 observations). This process was repeated nine times, each time using different training and test sets to obtain the K-fold cross-validation classification accuracy.

We resampled each dataset at $f'_s = 50$ Hz. The mean and variance of the acceleration in the x , y , and z axes from each of the three accelerometers were calculated over 200-millisecond time slices.

4 Activity Classification Using HMMs

Given a training set of labeled observation sequences (features extracted from the acceleration readings in the x , y , and z axes from three accelerometers placed on different parts of the body), corresponding to each of the activities that we want to classify, we first want to estimate the model parameters $\lambda = (\mathbf{A}, \mathbf{B}, \pi)$, where $\mathbf{A} = \{a_{ij}\}$ is the state transition probability distribution, $\mathbf{B} = \{b_j(k)\}$ is the observation symbol probability distribution in state j , and $\pi = \{\pi_i\}$ is the initial state distribution for each activity. Given a new set of observations we would like to classify each sequence according to the model that gives the maximum likelihood for that particular sequence.

We modeled each observation sequence as a 5-state left-to-right HMM with continuous Gaussian observation vectors and two hidden states. Each observation vector was formed by combining the mean and variance in the x , y , and z axes from each accelerometer. These features were previously used in [2]. An HMM was trained for each class $(\lambda_1, \lambda_2, \dots, \lambda_C)$, where λ_c indicates the learned HMM model for class c , and $C = 8$ is the total number of classes, using the labeled data from eight datasets as training dataset \mathcal{T}_k . The ninth data set was used as the validation dataset \mathcal{V}_k .

The HMM toolbox for Matlab developed by [7] was used to train and test the different models. The log-likelihood of each model was calculated for each observation sequence in the ninth dataset. Each observation sequence $O^l = \{O_1^l O_2^l \dots O_T^l\}$ (with $T = 5$ time slices) in the validation dataset $\mathcal{V}_k = \{O^l\}_{l=1}^L$ was classified according to the model that gave the maximum likelihood.

The final classification was obtained as

$$\hat{G}(O^l) = \arg \max_c \mathcal{L}(\lambda_c). \quad (1)$$

The classifier \hat{G} takes values in the class set $G = \{1, 2, \dots, C\}$.

This process was repeated $K = 9$ times using k-fold cross-validation. The average cross-validation classification accuracy per class is compared for the four possible configurations of accelerometer placement shown in table 1.

	Configuration			
	C1	C2	C3	C4
Right wrist	YES	YES		YES
Left hip	YES		YES	YES
Chest		YES	YES	YES

Table 1. Accelerometer configurations

5 Results

Figure 1 shows a comparison of classification accuracy when a single accelerometer was used for activity classification. We are able to discern activities such as walking (65.68%) and performing hand movements (56.30%) using only accelerometer A1 (right wrist). Accelerometer A2 (left hip) played the most important role when classifying activities such as sitting (66.05%), running (97.78%), crawling (69.26%), and lying down (87.04%). Accelerometer A3 (chest) was best for classifying activities such as squatting (75.8%) and standing (77.78%).




Figure 1. K-fold cross validation accuracy comparison using a single accelerometer.

Figure 2 shows the average classification accuracy per activity when combinations of two and three accelerometers were placed in the four different configurations described in table 1. Table 2 shows a comparison between the classification accuracy obtained when a single accelerometer (A1 to A3) was used, and the classification accuracy obtained when multiple accelerometers were used (C1 to C4).




Figure 2. K-fold cross validation accuracy comparison for the four different configurations of accelerometer placement.

Class	A1	A2	A3	C1	C2	C3	C4
1	58.40%	66.05%	63.21%	88.52%	94.44%	95.56%	97.78%
2	80.86%	97.78%	96.54%	98.89%	95.43%	100.0%	98.89%
3	23.46%	42.72%	75.80%	72.10%	67.90%	92.72%	82.72%
4	65.68%	41.73%	54.44%	84.32%	77.28%	61.11%	84.20%
5	69.75%	68.77%	77.78%	90.00%	94.44%	97.65%	95.56%
6	45.93%	69.26%	57.65%	84.20%	69.26%	73.58%	91.98%
7	43.21%	87.04%	41.98%	81.11%	76.42%	86.05%	87.16%
8	56.30%	53.51%	32.22%	98.77%	79.85%	55.28%	98.77%
Global	55.45%	65.86%	62.45%	87.24%	81.88%	82.74%	92.13%

Table 2. Classification accuracy comparison

Our results show that it is possible to recognize some of the most common activities using a single accelerometer on the chest (with 62.45% average accuracy). Adding a second accelerometer on the hip or the wrist improved our classification accuracy by approximately 20%. Adding a third accelerometer improved the global classification accuracy by an additional 10%.

The results presented so far were obtained by modeling each activity with an HMM having two hidden states. However, we think that classification results could be improved by modeling each activity with a different number of hidden states. In some cases, two hidden states might not be enough for capturing the different stages of a particular activity, especially when the activity involves different movements and body positions. Previous studies have not taken this into consideration and have modeled all activities using the same number of hidden states.

Table 3 shows the K-fold cross-validation accuracy when modeling each activity with two-hidden-state HMMs and when using the number of hidden states, Q_{min} , that gives the minimum test error for configuration C4 of accelerometer placement. Based on these results, we select the number of hidden states, Q_{opt} , that best models each of the activities studied in this paper.

Class	$Q = 2$	Q_{min}	Var	Q_{opt}
1	97.78%	93.33%	-4.45%	2
2	98.89%	98.89%	0%	2
3	82.72%	82.72%	0%	2
4	84.20%	88.64%	+4.44%	4
5	95.56%	96.67%	+1.11%	4
6	91.98%	86.29%	-5.69%	2
7	87.16%	87.26%	+0.1%	4
8	98.77%	99.01%	+0.24%	5

Table 3. Classification accuracy variation and optimal number of hidden states.

6 Conclusions

We found that the best global classification performance was achieved when using configuration C4 (92.13%), although it might be possible to obtain similar results using only two accelerometers. We showed that the global classification accuracy that can be achieved using a single accelerometer is around 60%, and determined the activities that are best classified with each accelerometer placed on three different parts of the body. Modeling each activity with a different number of hidden states improved the results. We are convinced that a minimum system formed by a wearable badge and a cellular phone can achieve fairly good results in daily activity recognition (80%).

References

- [1] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *Proceedings of the 2nd International Conference on Pervasive Computing*, pages 1–17, Apr 2004.
- [2] M. L. Blum. Real time context recognition. Master’s thesis, Swiss Federal Institute of Technology Zurich (ETH) & MIT Media Lab, Oct 2005.
- [3] R. W. DeVaul and S. Dunn. Real-time motion classification for wearable computing applications. Technical report, MIT Media Lab, Dec 2001.
- [4] P. Lukowicz, J. A. Ward, H. Junker, M. Stager, G. Troster, A. Atrash, and T. Starner. Recognizing workshop activity using body worn microphones and accelerometers. In *Proceedings of the 2nd International Conference on Pervasive Computing*, pages 18–22, Apr 2004.
- [5] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *International Workshop on Wearable and Implantable Body Sensor Networks*, Apr 2006.
- [6] E. Munguia Tapia, N. Marmasse, S. S. Intille, and K. Larson. Mites: Wireless portable sensors for studying behavior. In *Proceedings of the 6th International Conference on Ubiquitous Computing*, Sep 2004.
- [7] K. Murphy. <http://www.cs.ubc.ca/murphyk/Software/HMM/>.