Towards the Detection of Unusual Temporal Events during Activities Using HMMs

Shehroz S. Khan[1], **Michelle E. Karg**[1,2], **Jesse Hoey**[1], **Dana Kulić**[1,2]
[1]David R. Cheriton School of Computer Science
[2]Department of Electrical and Computer Engineering
University of Waterloo, 200, University Ave. West, Waterloo, Ontario, Canada N2L 3G1, Canada {s255khan,mekarg,jhoey,dkulic}@uwaterloo.ca

ABSTRACT
Most of the systems for recognition of activities aim to identify a set of normal human activities. Data is either recorded by computer vision or sensor based networks. These systems may not work properly if an unusual event or abnormal activity occurs, especially ones that have not been encountered in the past. By definition, unusual events are mostly rare and unexpected, and therefore very little or no data may be available for training. In this paper, we focus on the challenging problem of detecting unusual temporal events in a sensor network and present three Hidden Markov Models (HMM) based approaches to tackle this problem. The first approach models each normal activity separately as an HMM and the second approach models all the normal activities together as one common HMM. If the likelihood is lower than a threshold, an unusual event is identified. The third approach models all normal activities together in one HMM and approximates an HMM for the the unusual events. All the methods train HMM models on data of the usual events and do not require training data from the unusual events. We perform our experiments on a Locomotion Analysis dataset that contains gyroscope, force sensor, and accelerometer readings. To test the performance of our approaches, we generate five types of unusual events that represent random activity, extremely unusual events, unusual events similar to specific normal activities, no or little motion and normal activity followed by no or little motion. Our experiments suggest that for a moderately sized time frame window, these approaches can identify all the five types of unusual events with high confidence.

Author Keywords
Unusual Event, Abnormal Activity, Hidden Markov Models, Sensor Networks

INTRODUCTION
The world's population of older adults is increasing rapidly due to the improvements in medical science and health-care in the past century. A lot of recent research has focused on enabling older adults (especially those with cognitive diseases) to live independently by developing assistive technologies that provide aid in their activities of daily living (ADL) [6, 1, 27]. Some of the research on Assistive Technologies [17] is centered around activity recognition, studying the actions, behaviours and goals of an individual attempting to recognize them and provide the desired assistance. A central focus of many of these studies is the detection of usual daily human activities e.g. walking, hand washing, making breakfast etc. However, in many scenarios detection of unusual activities is of more importance as it may render an older adult at risk and vulnerable. Consider an activity monitoring system where the normal activities such as walking, sitting, or standing are important to identify, but the more challenging and useful thing to identify is when the person deviates from these normal or relatively safe activities. In this context these *unusual* activities include incurring a fall or suffering a stroke. A typical activity recognition system may misclassify 'fall' as one of the already existing normal activities because 'fall' may not have occurred earlier. An alternative strategy is to detect specific unusual activities such as fall [33, 12]. However, this may require extensive domain knowledge, an understanding of the activities that may be encountered and data collection for the type of unusual activity to be modelled. These algorithms would only be able to detect the specific unusual activity on which they are trained and cannot be generalized to other types of unusual activities. Moreover, in emergency situations it is important to first identify if an unusual event has occurred and later on efforts can be expanded to find their specific details.

Zhang et al. [34] defines *unusual events* for the audio-visual steams as the ones that are *rare*, *unexpected* and hold *relevance* for a particular task. The rarity of unusual events yet to be observed leads to a lack of sufficient data for training the model. More than one type of unusual events may also occur in a data sequence and the unexpectedness of unusual events makes it difficult to model them in advance. Yin et al. [32] provides a similar definition for *abnormal activities* in sensor-based human activity detection. They mention that due to the scarcity of such activities, it is a challenging problem to design a detection system that can reduce both the false positives and false negatives. Collecting abnormal activities data can be cumbersome because it may require the person to actually undergo such unusual events which may

UbiComp '12, Sep 5-Sep 8, 2012, Pittsburgh, USA.

be harmful. In addition to very few or no labelled data, the diversity and types of unusual events further makes it difficult to model them efficiently.

Most of the research in unusual activity recognition is based on computer vision systems [30, 5, 34]. Vision based systems work well in an indoor setting, however when an older adult goes out of the sensing range these systems cannot provide much help. Though methods for specific abnormal activities such as fall detection exist, not much research work has been carried out for general unusual activity detection in the sensor network area. In this paper, we focus on detecting unusual activities from sensor-based systems which can be worn easily in both indoor and outdoor settings. The paper presents three Hidden Markov Model (HMM) based methods for unusual event detection; the first method trains itself on individual normal activities where as the second method combines all normal activities into a single class. For both of these methods, based on a predefined threshold it raises an alarm for unusual activity. In the third method, an HMM is trained for usual activities and the model for unusual activities is approximated by varying the covariances of the observations of the usual events. As training data for the unusual events is typically sparse or unavailable, all three approaches for detecting unusual events use observations from the usual events during the training phase; unusual events are only used during the testing phase. It is interesting to note that the unusual activities may be scarce in number and modelling different unusual activities as a single class can be challenging. In all the three methods, we model only the usual events together / separately (see Figure 1 and 2) based on the fact that different unusual activities are sufficiently different from the normal activities. While some unusual activities may be quite similar to a usual activity (e.g. falling onto or sitting down upon the couch), our hypothesis is that we will be able to train a classifier to detect this difference. That is, there is some separation between the usual and unusual events, and we can use a generative model to distinguish them. The inter-difference among various unusual activities is out of scope of the present paper as we only focus on detecting an unusual event irrespective of its type.

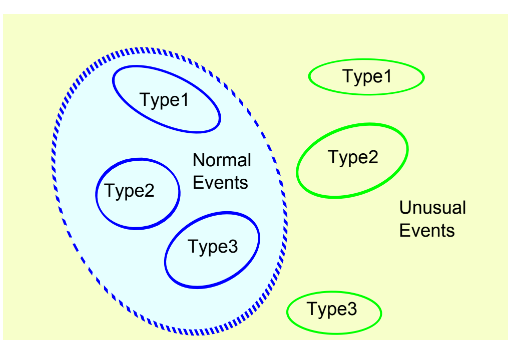

Figure 1. Conceptual Diagram of Modelling Usual and Unusual events for Approach I and II.

The rest of the paper is structured as follows. In the following section, we present a short survey of related work on various unusual events and abnormal activity recognition both in vision and sensor-based systems. Then we present three approaches to identify unusual activities in sensor-based systems. After that we present the experimental results, followed by a discussion on future extensions and conclusions.

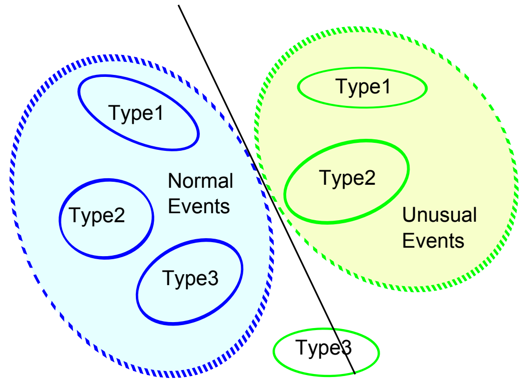

Figure 2. Conceptual Diagram of Modelling Usual and Unusual events for Approach III.

RELATED WORK

In the past decade, there has been considerable research carried out in the field of activity recognition using HMMs [15, 21, 10]. Most of these works are focused on modelling normal ADL using different types of sensors (such as accelerometers, GPS, WLAN, video camera etc) and employing supervised learning for recognizing the activities. However, these methods require a large amount of labelled data to train these supervised classification models, which may not be available for modelling unusual events because there is always a scarcity of labelled training data. Therefore, these methods cannot be directly adapted to detect unusual events or abnormal activities and new solutions are desired.

Several approaches have been proposed for abnormal activity recognition in the past, both in the field of computer vision and sensor based networks. Xiang and Gong [31] propose a Dynamic Bayesian Network approach to model each normal video pattern and use a threshold to detect an abnormal activity. This approach is simple, however choosing a threshold remains challenging. Duong et al. [5] introduce Switching Hidden Semi-Markov Model for modelling normal activities and identifying abnormal activities using multiple camera tracking. However, they only focus on a specific type of abnormality that corresponds to spending too much or too little time at a location and can be of interest in an elder care application. Zhang et al. [34] propose a semi-supervised adapted HMM framework for audio-visual data streams which comprises of supervised learning of normal data and unsupervised learning of unusual events using Bayesian adaptation. Their method has an iterative structure, where each iteration corresponds to a new detected unusual event. However, it is not clear from their work how many iterations are needed to terminate the process of outlier detection. Their model assumes that the usual data contains unusual events and guarantees one outlier per iteration. In cases where the usual data contains no unusual events, their method would still finds an outlier per iteration, which is highly undesirable. Pruteanu-Malinici and Carin [22] propose infinite HMM modelling to train normal video sequences; unusual events are detected if a low likelihood is observed. The infinite HMM modelling retains the full posterior density function as well as the underlying HMM states. Zhang et al. [35] propose an abnormal event detection algorithm from video sequences using a three-phased approach. Firstly, they build a set of weak classifiers using Hierarchical Dirichlet Process Hidden Markov Model (HDP-HMM) and then use ensemble learning to identify abnormal events. Finally, they extract abnormal events from the normal ones in unsupervised manner to reduce the FP rates. Hu et al. [7] propose a refinement of the HDP-HMM method by incorporating Fisher Kernel into One-Class Support Vector Machines (OSVM) instead of ensemble learning and using sensor data instead of video data that can be discrete or continuous. The advantage of their method relies on using the HDP-HMM models that can decide the appropriate number of states of the underlying HMM automatically.

Several recent works focus on using the sensor networks to detect abnormal activities. Yin et al. [32] propose a two-stage abnormal activity detection, in which an OSVM is first trained on normal activities and the abnormal activities are filtered out and passed on to a kernel nonlinear regression routine to derive abnormal activity models from a general normal activity model in an unsupervised manner. The method iteratively detects different types of abnormal activities based on a threshold. They claim that this method provides a good trade-off between false alarm and abnormal activity detection without collecting and labelling the abnormal data. The data is collected by using wearable sensors attached to a user and abnormal instances were collected by simulating 'falls' and 'slipping' in different positions. Therefore, in their problem formulation, the types of abnormal activities are known. Another limitation is that they model a specific type of abnormality, however in practice a user can carry out complex and interleaved set of activities. Rivera-Illingworth et al. [25] present an adaptive neural network architecture that can grow in size and add nodes to the hidden layers in an online manner whenever an unseen example is encountered. They add an additional memory layer to their neural network architecture to capture temporal information, collect data from various sensors fitted in a smart-home set up and build a general model of normal activities. They proposed two methods for abnormality detection; the first method applies a threshold in the hidden layer whereas the second method takes the entire output pattern into account and matches it against a threshold. However, the model complexity remains high and the significance of additional memory layer is not substantiated. Quin et al. [23] present a general framework of Switched Linear Dynamical Systems (SLDS) for condition monitoring of a premature baby receiving intensive care. They introduce the 'X-factor' to deal with unmodelled variation from the normal events that may not have been seen previously. The general principle to identify an unusual event is to vary the covariance of the mode of normal events to determine the interval with the highest likelihood where events can be classified as 'not normal'. To model dynamic detection of unusual events, they add a new factor to the existing SLDS model by inflating the system noise covariance of the normal dynamics. The sensor data is collected using various probes connected to each baby. The computation of the factor related to increasing the covariance remains challenging and is critical in this application. Lotfi et al. [18] describe a solution for supporting independent living of the elderly by equipping their homes with various sensors to monitor their behaviour. They use start-time and duration and employ various clustering techniques to detect abnormal behaviours in a smart-home set up using a threshold and on the placement of different data points within a cluster and the size of clusters.

Some of work discussed above [34, 22, 7] use thresholds based on log-likelihood to detect unusual events using variants of HMMs, mostly in computer vision based systems. In this paper, we present and compare three variations of unusual event detection methods that are designed for sensor based networks. Our methods are (1) classifying different normal activities separately and detecting unusual events, (2) classifying all normal activities as a single class and detecting unusual events using log-likelihood based thresholds, and (3) training an HMM on normal activities and detecting the abnormal activities by approximating an unusual event model by modifying the covariances of the normal activities model. The focus of the present paper is on identifying unusual temporal events or abnormal activities using non-intrusive sensors that can be conveniently employed in both indoor or outdoor settings. We intend not to model any specific type of unusual activity or use any prior domain knowledge during the training phase, rather we attempt to detect deviations from the normal activities during the testing phase. Keeping this view in mind we analyze three approaches based on HMM that are described in the next section.

APPROACHES

HMMs are applied to model actions [10, 9]. Time series recordings of an action are modelled by a cyclic left-to-right HMM which evolves through a number of k states. The observations $o_j(t)$ in state j are modelled by Gausian distributions. Each model i is described by the set of parameters $\lambda_i = \{\pi, A, \mathcal{N}_j(\mu, \Sigma)\}$ containing the prior π, the transition matrix A, and the description of the observation probabilities $P(o_j|j)$ for each state j. The parameters are trained by the Baum-Welch algorithm [24]. A set of n_a different actions is modelled by n_a HMMs. We apply this method to model the usual activities.

The following two properties of unusual events are taken into account for developing a detection algorithm:

- Rare occurrence of unusual events: In comparison to usual actions, only a small amount of training data may be available for unusual events. Considering fall detection in gerontechnology, no training data may be available or only a rough description of sensor measurements during a fall. Hence, it is desired that the approach does not require training data for the unusual events.
- Unexpected and versatile characteristic of unusual events:

Unusual events are unexpected and may appear in a various number of forms, e.g. a fall may happen in various ways. Hence, the method should be generalizable to a set of different unusual events.

Hence, the investigated methods do not aim to model unusual events based on training data. Instead the following alternative approaches are considered:

1. Approach I models each usual action i by an HMM. It estimates the probability that the observed sequence has been generated by each of the n_a models of usual events. If this probability falls below a threshold T_i, an unusual event is detected. One way to choose the threshold T_i for each action i is to define it as the lowest observed probability $log(P(O_i|\lambda_i))$ for all training sequences O_i. This approach assumes that the training data is labelled and that the sets of training sequences for each action i do not contain outliers. Outliers in the training set would result in a too low threshold T_i and unusual event detection will be biased towards missing unusual events.

2. Approach II models all usual actions by a common HMM. It estimates the probability that the observed sequence has been generated by this common model for all usual actions. Similar to approach I, if this probability falls below a threshold T, an unusual event is detected. One way to choose this threshold is to define it as the lowest observed probability $log(P(O|\lambda))$ for all training sequences O. Also, this approach assumes that the training data does not contain outliers.

3. Approach III estimates a model for the unusual events by modifying the model for the usual events by varying the covariances of their observations. It applies the X-factor introduced in [23] for SLDS to HMMs. The X-factor is a model for unusual events $\lambda = \{\pi, A, \mathcal{N}(\mu^{(x)}, \Sigma^{(x)})\}$ which is gained by alternating the model of the usual events. The new model is constructed by

$$\Sigma^{(x)} = \xi\Sigma\ , \quad \mu^{(x)} = \mu \tag{1}$$

   and increasing the covariance of the observations by choosing $\xi > 1$. Hence, the HMMs for the usual and the unusual events differ only in the covariances of the observation distributions.

It is to be noted that in all the three approaches discussed above, only the training data from usual events are needed to build the respective HMMs and no training data of the unusual events is required. The data from unusual events are only used when testing the detection rate of these models. These three approaches are based on the assumption that all unusual events have in common that their observations differ sufficiently from the observations of the usual events even though the unusual events may differ from each other.

EXPERIMENTAL RESULTS

The three approaches are evaluated on a locomotion dataset and unusual events are simulated by five types of artificially generated sensor data. Unusual events are simulated to test the detection rate on a range of unusual events and can provide insights into the behaviour of these approaches for different types of unusual events.

Dataset

We used the Locomotion Analysis dataset [8] in our experiment as normal activities. The data contains 10-dimensional numeric readings from different sensors attached to the body of users and contains three sets of activities while walking on a predefined path – level walking, walking up, and walking down the stairs. The following sensors are used: accelerometers, air pressure sensor, gyroscopes, and force sensitive resistors. The accelerometers may be affixed with different orientation, so that only the absolute value of the accelerometer at the knee and the hip is used later on. Table 1 lists the different recorded features f_i and the investigated feature vector y is

$$y = [f_1, f_3, \sqrt{f_5^2 + f_6^2 + f_7^2}, \sqrt{f_8^2 + f_9^2 + f_{10}^2}\,]^T\ . \tag{2}$$

Feature 2 is excluded since we consider it as irrelevant for the current activity recognition problem. The sampling rates of all sensors is 100Hz except for the sampling rate of the air pressure sensor (1Hz). For applying the HMM, the data is down-sampled to 33Hz so that calculated logP values stay within a reasonable range for longer sequences.
The dataset consists of two separate observations for the 1^{st}, 2^{nd} and 3^{rd} subject, and seven observations for the 4^{th} subject. For the following analysis, the sequences of the four subjects are combined. The dataset is segmented in windows ranging from 2 -15s. The windows do not overlap and segments do not contain transitions between walking and going upstairs or downstairs. The number of segments decreases with increasing window size. For 5s windows, there are 210 segments for walking, 211 for going upstairs, and 78 for going downstairs. For 10s windows, there are 92 segments for walking and 71 segments for going upstairs. Segments of going downstairs and upstairs are shorter than 7s for subject 1-3. Hence, all segments of going upstairs come from subject 4 for window sizes larger than 7s.

Simulating Unusual Events

As discussed earlier, in this paper we are not modelling specific unusual temporal event(s) rather we want to model any abnormal activity that might occur. This may be related to either a user encountering an unusual event, behaving in a manner which is not normal, or may be due to sensors working improperly. Our primary concern is to investigate the properties of HMM-based methods for detecting unusual events in general, and so we simulate unusual events by artificially generating the sensor data and use them as test cases to validate our results. The following is a list of five methods to generate different kinds of unusual events. These are not meant to be an exhaustive set, but rather a range of possibilities that are reasonable and will allow us to delve into the properties of the classifiers we are investigating.

1. *Unusual1 (U1)* – Combine the data from all three normal activities together and artificially generate uniformly distributed random sensor readings within the range of

Features f_i	Details
1	Data from Gyroscope mounted above right knee (a.U)
2	Data from Air Pressure Sensor (mbar)
3	Data from Force Sensitive Resistor mounted under ball of right foot (a.U.)
4	Data from Force Sensitive Resistor mounted under hell of right foot (a.U.)
5	Data from Accelerometer mounted above right knee; Axis: top to bottom (9.81m/s/s)
6	Data from Accelerometer mounted above right knee; Axis: left to right (9.81m/s/s)
7	Data from Accelerometer mounted above right knee; Axis: back to front (9.81m/s/s)
8	Data from Accelerometer mounted on back of body on a belt; Axis: top to bottom (9.81m/s/s)
9	Data from Accelerometer mounted on back of body on a belt; Axis: left to right (9.81m/s/s)
10	Data from Accelerometer mounted on back of body on a belt; Axis: back to front (9.81m/s/s)

Table 1. Details of the sensor data

each sensor. This unusual event represents random activity which is still within the sensor readings encountered before, for e.g. unusual event occurring during either level walking, ascending or descending the stairs.

2. *Unusual2 (U2)* – Combine the data from all three normal activities together and artificially generate uniformly distributed random sensor readings outside (twice) the range of each sensors. This unusual event represents an unusual event that has not been captured by the sensors before and therefore it is out of their present captured intervals. This may also represent a case when the sensor has malfunctioned.

3. *Unusual3 (U3)* – Artificially generate uniformly distributed random sensor readings within the range of each sensor for all the three normal activities separately. This unusual event represents unusual activity occurring while performing a specific task, for e.g. unusual event occurring during descending the stairs. The way this *Unusual3* is generated may resemble the actual activity in 'some ways' as the values are are within the sensors' ranges but distributed randomly.

4. *Unusual4 (U4)* – Artificially generate uniformly distributed random sensor readings with much less variance and mean close to zero. This scenario may represent a 'lying' state or no or little motion.

5. *Unusual5 (U5)* – Keep half of the sensor readings in a given time frame the same as normal activity and for the remaining half, artificially generate a sudden drop in sensor readings, with a mean close to zero and much lower variance. This unusual event may represent a sudden 'fall' following a normal activity and then stay in a state of no or little motion for the rest of the time frame.

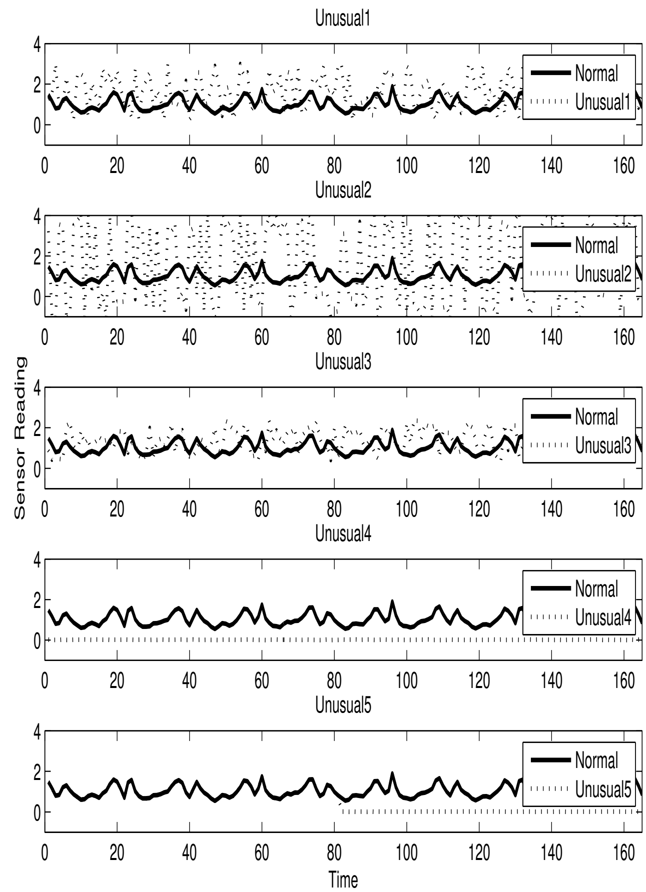

Figure 3. Normal, U1, U2, U3, U4 and U5 sensor readings for feature 4 for first 165 seconds

Figure 3 shows the readings for the first time frame for Normal, U1, U2, U3, U4 and U5 sensor data for feature 4. The five simulations of artificially generated sensor data represents different types of *unexpectedness* of the unusual events. All other unusual sensor readings for the other features are generated the same way.

These simulated unusual events are used only for testing, and are not included in training the detection algorithms. The number of unusual events lie in a similar range as the test samples of the usual event. 10,000 samples are generated. The length of each sequence is aligned with the segment length of the usual actions. This results in either 30 sequences containing 330 frames for 10s, 60 sequences containing 165 frames for 5s or 151 sequences containing 66 frames for 2s. In online scenarios, the occurrence of the unusual events would be less frequent.

Model for Usual Activities

In the following, 3-fold cross validation is applied to evaluate the HMM based approaches. Each HMM is modelled with 6 states and the variances are constrained between 0.01 and 3. A larger number of states did not improve the accuracy significantly and increased the computation time. A maximum number of 5 iterations is used during training.

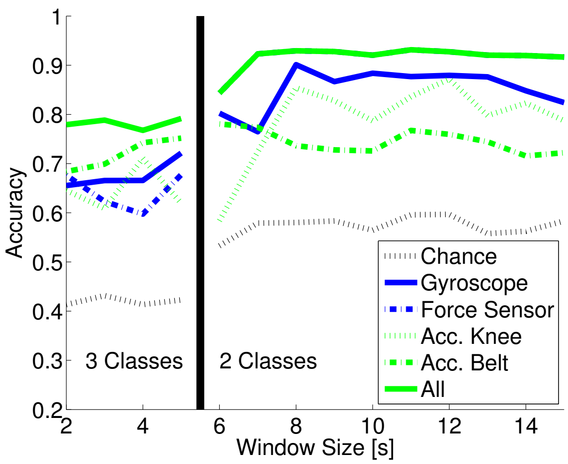

Figure 4. Accuracy for the recognition of walking and climbing stairs up for increasing window size.

Figure 4 shows the accuracy over the window size ranging from 2s to 15s. Up to a window size of 5s, segments are available for walking and going up and down the stairs. For this 3-class problem, an accuracy of 79% is achieved for a window size of 5s and using 4 features. Acceleration measured at the belt is the best single feature. Using only the studied sensors, recognition above chance is achieved, but confusions still occur. This result is in accordance with Lester et al. [14], who report that going up/down stairs is most commonly confused with walking in comparison to other actions such as sitting, riding an evaluator, or brushing teeth.

If the window size is larger than 5s, segments are only available for walking and going up stairs. For this 2-class problem, an accuracy of 92% is achieved for a window size of 5s and using 4 features. The best single feature is the force sensor and its accuracy matches exactly the accuracy of using all 4 features. Furthermore, it is noted that the increase in accuracy, which occurs for window sizes above 7s, is explained by a change in the dataset. Segments for climbing up stairs with a duration longer than 6s are obtained only for subject 4 and no data is available for subject 1, 2, and 3 for this subset. This results in less inter-subject variability in the data.

Performance Metrics for Unusual Event Detection

To evaluate our results, we employed the following performance metrics.

1. Precision for unusual events:

$$Precision = \frac{TP_{Unusual}}{TP_{Unusual} + FP_{Unusual}}$$

2. Recall for unusual events:

$$Recall = \frac{TP_{Unusual}}{TP_{Unusual} + FN_{Unsual}}$$

3. Overall accuracy:

$$Accuracy = \frac{True\ Unusual + True\ Usual}{All\ Unusual + All\ Usual} \quad (3)$$

where $TP_{Unusual}$ are the correctly identified unusual events, $FP_{Unusual}$ are the usual events falsely classified as unusual and FN_{Unsual} are the unusual events falsely classified as usual events. For Approach I, $True\ Usual$ is

$$True\ Usual = \sum_{i=1}^{\#NormalActivities} True\ Positives_i \quad (4)$$

The reason to only compute precision and recall values for unusual events is that we are only focused on identifying unusual events and not the inter-difference between the usual events. The accuracy will give an overall measure of performance of the detection algorithm.

Approach I

Tables 2 and 3 shows the precision / recall values for different unusual events and the overall accuracy when using Approach I. It can be inferred that Approach I can detect the five types of unusual events with high precision for 10 and 2 seconds, absolute precision for 5 seconds and absolute recall for both 10, 5 and 2 seconds time frames. The accuracy of detecting different normal actions and unusual events U1, U2, U3, U4 and U5 is very high for 10, 5 and 2 seconds time frame.

Unusual Event	Precision			Recall		
	10s	5s	2s	10s	5s	2s
U1, U2, U3, U4, U5	0.9792	1	0.9978	1	1	1

Table 2. Precision and Recall of unusual events using Approach I

Unusual Event	10s	5s	2s
U1, U2, U3, U4, U5	0.9406	0.9584	0.9636

Table 3. Accuracy using Approach I

Approach II

Approach II considers all the normal activities (walking, ascending, descending) as a single class. Table 4 shows the precision and recall values for the 10, 5 and 2 seconds time frame. We observe that Approach II gives absolute recall for all the unusual events for 10, 5 and 2 seconds time frame. The precision is absolute for 10 seconds and close to unity for all the unusual events for 5 and 2 seconds. From Table 5, we observe that for the 10 seconds time frame, unusual events U1, U2, U3, U4 and U5 give 100% accuracy whereas for the 5 and 2 seconds time frame it is very close to unity. An important difference between Approach I and Approach II is that the accuracy at 10, 5 and 2 seconds time frame is worse using Approach I due to misclassification among normal activities. That problem does not occur in Approach II, because all normal activities are considered as one class and their intermixing is not relevant.

Unusual Event	Precision			Recall		
	10 s	5 s	2 s	10 s	5 s	2 s
U1, U2, U3, U4, U5	1	0.9892	0.9913	1	1	1

Table 4. Precision and Recall of unusual events using Approach II

Unusual Event	10 secs	5 secs	2 secs
U1, U2, U3, U4, U5	1	0.9971	0.9982

Table 5. Accuracy using Approach II

Approach III

In Approach III, all the normal activities are combined together and modelled as a single HMMs. The value of parameter ξ is varied as 10, 100 and 1000. From Table 6, we observe that the precision and recall values for unusual events U1, U2 and U4 remain very high for all values of ξ and for both 10, 5 and 2 seconds time frames. For U3, as the value of ξ is increased to 1000 for the 10 seconds time frame, the recall drops due to misclassifying unusual event as normal activity. The reason is that U3 is generated such that it resembles each of the individual normal activities and increasing the ξ leads to improper inferences. For unusual event U5, for ξ=100, for the 5 seconds time frame we observe a drop in recall and when ξ=1000, it is classified as normal and thereby recall is zero for 5 and 2 seconds and very low at 10 seconds. The precision for U5 is undefined – due to zero $TP_{Unusual}$ and zero $FP_{Unusual}$ in one of the folds of the cross validation for 5 and 2 seconds and in all three folds of the cross validation for 10 seconds time frame. Similarly for accuracy values from Table 7, we observe that for all values of ξ for 10, 5 and 2 seconds absolute or very high accuracy is obtained for unusual events U1, U2 and U4. For unusual event U3 the accuracy drops with an increase in ξ for 10 seconds and for U5 accuracy drops for 10, 5 and 2 seconds due to high level of misclassification of unusual events as normal activities.

Unusual Event for $\xi = 10, 100, 100$	Precision			Recall		
	10s	5s	2s	10s	5s	2s
$\xi = 10$						
U1, U2, U3, U4, U5	1	0.9841	0.9661	1	1	1
$\xi = 100$						
U1, U2, U3, U4	1	1	1	1	1	1
U5	1	1	1	1	0.9	1
$\xi = 1000$						
U1, U2, U4	1	1	1	1	1	1
U3	1	1	1	0.5222	0.9778	0.9978
U5	-	-	-	0.3889	0	0

Table 6. Precision and Recall of unusual events using Approach III

Unusual Event for $\xi = 10, 100, 100$	Accuracy		
	10 s	5 s	2 s
$\xi = 10$			
U1, U2, U3, U4, U5	1	0.9956	0.9927
$\xi = 100$			
U1, U2, U3, U4	1	1	1
U5	1	0.9736	1
$\xi = 1000$			
U1, U2, U4	1	1	1
U3	0.8301	0.9941	0.9995
U5	0.7818	0.7349	0.7956

Table 7. Accuracy using Approach III

CONCLUSIONS AND DISCUSSION

In this paper, we present three HMM based methods that can raise an alarm whenever an unusual event occurs or any event that deviates from the usual ADL. We hypothesize that the different normal ADL carried out by a person are different from unusual events, and that these unusual events may be different among each other. In this work, we do not aim to model specific unusual events, rather we focus on a general concept of unusual event's occurrence and its subsequent detection. We present three different approaches to achieve our goal. The first approach (Approach I) models all the normal ADL as separate HMMs and identifies the unusual events based on a threshold. The second approach (Approach II) combines all the normal activities in one class, models a single HMM to represent usual ADL, and attempts to identify unusual events based on a threshold. The last approach (Approach III) creates HMM model for all normal ADL considered together as single class and vary the covariances to approximate the model of unusual events. We tested our methods on the Locomotion Analysis data that captures normal activities like walking, ascending and descending the stairs based on a number of wearable sensors attached to a person at points on the body. To simulate abnormal activities, we generate five types of unusual events that represents random activity (*Unusual1*), extremely unusual activity (*Unusual2*), slight deviations from specific normal activities (*Unusual3*), lying or little motion (*Unusual4*) and sudden change from specific normal activity to no or little motion (*Unusual5*). We segment the data in time frames of 10, 5 and 2 seconds to simulate a scenario that we are observing these many segments at a time. As a pre-processing step we find the magnitude of acceleration values in x, y and z direction because the considered activities detection does not require direction of acceleration, and to handle the case where accelerometers may have been mounted with differing orientations.

Approach I was able to identify all normal activities and the unusual events except for the case of U5, where half the time frame is identical to normal activity and half contains little or no motion. Moreover, in this approach misclassification among different normal activities led to reduced accuracy of the overall method. Approach II circumvents this problem by considering all the normal activities as a single class and therefore the intermixing of various normal activities (such as classifying walking up as level walking) is not relevant. Due to this formulation, Approach II successfully detects all the different unusual events with high precision and recall and the accuracy of the method remains very high for 10, 5 and 2 seconds time frame. A possible reason for the effectiveness of Approach II is that all the three normal activities are very similar to each other. The efficiency of this method for the case when large number of diverse normal activities are considered, needs to be investigated as a future work. Approach III trains an HMM on usual activities and builds the model for unusual activities by varying the covariances (using a parameter ξ) of the normal activities model. At high values of ξ, this method fails to identify unusual events U3 and U5 correctly. Our experiments suggests that on the Locomotion Analysis dataset, a value of ξ between 10 and 100 can be a good choice. A disadvantage of this method is to find the best value of ξ to obtain good classification. The unusual events U3 and U5 were difficult to classify using Approach III in comparison to other unusual events because they bear similarity with the normal activities. An interesting point to note is that at 5 and 2 seconds time frame length, Approach II is able to identify unusual events with very high precision, recall and accuracy. This can represent a real scenario, where observing a person for very few seconds, an unusual event can be identified and an alarm can be raised and further steps can be taken. An alternative to observe time frame in seconds is to segment the data in single steps and train and model each step by a non-cyclic HMM.

There are several challenges to take the problem of temporal unusual events detection forward. The present paper focuses on identifying different types of unusual events separately, however it remains challenging to build models to further discriminate these unusual events. We believe that approaches based on non-parametric and hierarchical clustering could be useful to cluster different unusual events. The main assumption behind our approaches is that, to start with all the data recorded so far contains only normal activities, however in real situations the data to start with may contain some unusual events along with the normal activities. We would like to adapt our approaches to handle such scenarios and by benefiting from the unusual event data as they are always scarce. To model specific abnormal activities which have either very few data to begin with or data collection is difficult, approaches based on over-sampling of minority class [3], and synthetic data generation for activity recognition [20, 19] can be explored to simulate more data and generative classification approaches be used. In such cases, investigating the role of discriminative temporal classifiers such as Conditional Random Fields [28, 16] will be interesting to explore. The number of hidden states of HMMs in the current implementation is chosen by testing the accuracy rates at various states, however advanced methods such as those based on HDP-HMM [7] and Infinite HMMs [22] can find the optimum number of hidden states.

A drawback with sensor based methods is their intrusive nature, a person has to wear the sensors all the time which may be uncomfortable and lead to a refusal to wear them [26]. The dataset used in our analysis comes from body worn sensors and the results suggest that accelerometers are adequate for detecting unusual events with the ongoing normal activities. There is recent interest in performing activity recognition using smartphones [13, 4, 11], which have become a common commodity across the globe and are easy to use and carry. Most activity recognition based on smartphones uses information captured from accelerometers. We plan to adapt our unusual event detection approaches to use smartphones to capture accelerometer data, which do not have the hassles introduced by wearable sensors. Smartphones can also be used to collect other types of sensor readings such as GPS, gyroscope, WLAN etc that can help in modelling the detection of specific unusual events such as wandering or falling. Another advantage of using smartphones for unusual event detection is that it can work in both indoor and outdoor settings where the traditional computer vision or camera based systems may fail. A possible limitation might be that recorded data of some actions, e.g. sitting down and falling, is very similar for acceleration sensors integrated in a smart phone. To increase accuracy further, additional wearable sensors, e.g. pressure sensors, might be necessary.

The problems and methods discussed in the paper are based on offline supervised classification or some of its variants. The real utility of a temporal unusual event detection mechanism can be achieved if new methods can be developed and adapted to support online identification of unusual events. In this direction, we believe that approaches based on incremental HMM learning [29, 2] can be helpful.

ACKNOWLEDGEMENTS

This research was funded in part by the Natural Sciences and Engineering Research Council of Canada. We are grateful for support from the Toronto Rehabilitation Institute, Canada.

REFERENCES

1. T. Adlam, B. E. Carey-Smith, N. Evans, R. Orpwood, J. Boger, and A. Mihailidis. Implementing monitoring and technological interventions in smart homes for people with dementia - case studies. In B. Gottfried and H. K. Aghajan, editors, *BMI Book*, volume 3 of *Ambient Intelligence and Smart Environments*, pages 159–182. IOS Press, 2009.

2. R. V. Andreão, S. M. T. Muller, J. Boudy, B. Dorizzi, T. F. Bastos-Filho, and M. Sarcinelli-Filho. Incremental HMM training applied to ECG signal analysis. *Comput. Biol. Med.*, 38(6):659–667, June 2008.

3. K. W. Bowyer, N. V. Chawla, L. O. Hall, and W. P. Kegelmeyer. SMOTE: Synthetic minority over-sampling technique. *CoRR*, abs/1106.1813, 2011.

4. T. Brezmes, J.-L. Gorricho, and J. Cotrina. Activity recognition from accelerometer data on a mobile phone. In S. Omatu, M. Rocha, J. Bravo, F. F. Riverola, E. Corchado, A. Bustillo, and J. M. Corchado, editors, *IWANN (2)*, volume 5518 of *Lecture Notes in Computer Science*, pages 796–799. Springer, 2009.

5. T. V. Duong, H. H. Bui, D. Q. Phung, and S. Venkatesh. Activity recognition and abnormality detection with the switching hidden semi-markov model. In *CVPR (1)*, pages 838–845, 2005.

6. J. Hoey, P. Poupart, A. von Bertoldi, T. Craig, C. Boutilier, and A. Mihailidis. Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process. *Computer Vision and Image Understanding*, 114(5):503–519, 2010.

7. D. H. Hu, X.-X. Zhang, J. Yin, V. W. Zheng, and Q. Yang. Abnormal activity recognition based on HDP-HMM models. In C. Boutilier, editor, *IJCAI*, pages 1715–1720, 2009.

8. H. Junker. Modes of Locomotion. `http://www.pervasive.jku.at/Research/Context_Database`. Online Accessed 19^{th} June 2012.

9. A. Kale, A. Sundaresan, A. Rajagopalan, N. Cuntoor, A. Roy-Chowdhury, V. Krueger, and R. Chellappa. Identification of humans using gait. *IEEE Transactions on Image Processing*, 13(9):1163–1173, 2004.

10. D. Kulić, W. Takano, and Y. Nakamura. Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains. *The International Journal of Robotics Research*, 27(7):761–784, 2008.

11. J. R. Kwapisz, G. M. Weiss, and S. A. Moore. Activity recognition using cell phone accelerometers. *SIGKDD Explor. Newsl.*, 12(2):74–82, Mar. 2011.

12. T. Lee and A. Mihailidis. An intelligent emergency response system: preliminary development and testing of automated fall detection. *Journal of Telemedicine and Telecare*, 11(4):194–198, 2005.

13. Y. Lee and S.-B. Cho. Activity recognition using hierarchical hidden markov models on a smartphone with 3D accelerometer. In E. Corchado, M. Kurzynski, and M. Wozniak, editors, *HAIS (1)*, volume 6678 of *Lecture Notes in Computer Science*, pages 460–467. Springer, 2011.

14. J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. In *In Proc. of Pervasive*, pages 1–16, 2006.

15. J. Lester, T. Choudhury, N. Kern, G. Borriello, and B. Hannaford. A hybrid discriminative/generative approach for modeling human activities. In L. P. Kaelbling and A. Saffiotti, editors, *IJCAI*, pages 766–772. Professional Book Center, 2005.

16. L. Liao, D. Fox, and H. A. Kautz. Extracting places and activities from GPS traces using hierarchical conditional random fields. *I. J. Robotic Res.*, 26(1):119–134, 2007.

17. E. F. LoPresti, A. Mihailidis, and N. Kirsch. Assistive technology for cognitive rehabilitation:state of the art. *Neuropsychological Rehabilitation*, 14(1/2):5–39, 2004.

18. A. Lotfi, C. Langensiepen, S. M. Mahmoud, and M. J. Akhlaghinia. Smart homes for the elderly dementia sufferers: identification and prediction of abnormal behaviour. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–14–14, 2011.

19. A. Mendez-Vazquez, A. Helal, and D. Cook. Simulating events to generate synthetic data for pervasive spaces. In *Workshop on Developing Shared Home Behavior Datasets to Advance HCI and Ubiquitous Computing Research*, 2009.

20. D. Monekosso and P. Remagnino. Synthetic training data generation for activity monitoring and behavior analysis. In *Proceedings of the European Conference on Ambient Intelligence*, AmI '09, pages 267–275, Berlin, Heidelberg, 2009. Springer-Verlag.

21. N. Oliver, E. Horvitz, and A. Garg. Layered representations for human activity recognition. In *ICMI*, pages 3–8. IEEE Computer Society, 2002.

22. I. Pruteanu-Malinici and L. Carin. Infinite hidden markov models for unusual-event detection in video. *IEEE Transactions on Image Processing*, 17(5):811–822, 2008.

23. J. A. Quinn, C. K. Williams, and N. McIntosh. Factorial switching linear dynamical systems applied to physiological condition monitoring. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(9):1537–1551, 2009.

24. L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

25. F. Rivera-Illingworth, V. Callaghan, and H. Hagras. "A neural network agent based approach to activity detection in AmI environments". In *IEE International Workshop, Intelligent Environments 2005 (IE05), Colchester, UK, 28-29th June 05*, 2005.

26. B. Schulze, M. Floeck, and L. Litz. Concept and design of a video monitoring system for activity recognition and fall detection. In M. Mokhtari, I. Khalil, J. Bauchet, D. Zhang, and C. D. Nugent, editors, *ICOST*, volume 5597 of *Lecture Notes in Computer Science*, pages 182–189. Springer, 2009.

27. M. Sinn and P. Poupart. Smart walkers!: enhancing the mobility of the elderly. In L. Sonenberg, P. Stone, K. Tumer, and P. Yolum, editors, *AAMAS*, pages 1133–1134. IFAAMAS, 2011.

28. D. L. Vail, M. M. Veloso, and J. D. Lafferty. Conditional random fields for activity recognition. In E. H. Durfee, M. Yokoo, M. N. Huhns, and O. Shehory, editors, *AAMAS*, page 235. IFAAMAS, 2007.

29. K. Wakabayashi and T. Miura. Data stream prediction using incremental hidden markov models. In *Proceedings of the 11th International Conference on Data Warehousing and Knowledge Discovery*, DaWaK '09, pages 63–74, Berlin, Heidelberg, 2009. Springer-Verlag.

30. T. Xiang and S. Gong. Video behaviour profiling and abnormality detection without manual labelling. In *Proceedings of the Tenth IEEE International Conference on Computer Vision - Volume 2*, ICCV '05, pages 1238–1245, Washington, DC, USA, 2005. IEEE Computer Society.

31. T. Xiang and S. Gong. Video behaviour profiling and abnormality detection without manual labelling. In *ICCV*, pages 1238–1245. IEEE Computer Society, 2005.

32. J. Yin, Q. Yang, and J. J. Pan. Sensor-based abnormal human-activity detection. *IEEE Trans. Knowl. Data Eng.*, 20(8):1082–1090, 2008.

33. X. Yu. Approaches and principles of fall detection for elderly and patient. In *10th International Conference on ehealth Networking Applications and Services, HealthCom'08*, pages 42–47. IEEE, 2008.

34. D. Zhang, D. Gatica-Perez, S. Bengio, and I. McCowan. Semi-supervised adapted HMMs for unusual event detection. In *CVPR (1)*, pages 611–618, 2005.

35. X.-X. Zhang, H. Liu, Y. Gao, and D. H. Hu. Detecting abnormal events via hierarchical dirichlet processes. In T. Theeramunkong, B. Kijsirikul, N. Cercone, and T. B. Ho, editors, *PAKDD*, volume 5476 of *Lecture Notes in Computer Science*, pages 278–289. Springer, 2009.