Applied Soft Computing 55 (2017) 168–177

Contents lists available at ScienceDirect

Applied Soft Computing

journal homepage: www.elsevier.com/locate/asoc

Detecting falls with X-Factor Hidden Markov Models

Shehroz S. Khan [a,*], Michelle E. Karg [b], Dana Kulić [b], Jesse Hoey [a]

[a] David R. Cheriton School of Computer Science, University of Waterloo, Canada
[b] Department of Electrical and Computer Engineering, University of Waterloo, Canada

ARTICLE INFO

Article history:
Received 30 June 2016
Received in revised form
19 December 2016
Accepted 20 January 2017
Available online 25 January 2017

Keywords:
Fall detection
Hidden Markov Models
X-Factor
Outlier detection

ABSTRACT

Identification of falls while performing normal activities of daily living (ADL) is important to ensure personal safety and well-being. However, falling is a short term activity that occurs rarely and infrequently. This poses a challenge for traditional supervised classification algorithms, because there may be very little training data for falls (or none at all) to build generalizable models for falls. This paper proposes an approach for the identification of falls using a wearable device in the absence of training data for falls but with plentiful data for normal ADL. We propose three 'X-Factor' Hidden Markov Model (XHMMs) approaches. The XHMMs have 'inflated' output covariances (observation models). To estimate the inflated covariances, we propose a novel cross validation method to remove 'outliers' from the normal ADL that serves as proxies for the unseen falls and allow learning the XHMMs using only normal activities. We tested the proposed XHMM approaches on two activity recognition datasets and show high detection rates for falls in the absence of fall-specific training data. We show that the traditional method of choosing threshold based on maximum of negative of log-likelihood to identify unseen falls is ill-posed for this problem. We also show that supervised classification methods perform poorly when very limited fall data is available during the training phase.

1. Introduction

Identification of normal Activities of Daily Living (ADL), for e.g., walking, hand washing, making breakfast, etc., is important to understand a person's behaviour, goals and actions [1]. However, in certain situations, a more challenging, useful and interesting research problem is to identify cases when an abnormal activity occurs, as it can have direct implications on the health and safety of an individual. An important abnormal activity is the occurrence of a fall. However, falls occur rarely, infrequently and unexpectedly w.r.t. the other normal ADLs and this leads to either little or no training data for them [2]. The Centers for Disease Control and Prevention, USA [3], suggests that on average, patients incur 2.6 falls per person per year. Recent studies also suggest that even in a long term experimental set up only a few real falls may be captured [4,5]. In these situations with highly skewed fall data, a typical supervised activity recognition system may misclassify 'fall' as one of the already existing normal activity as 'fall' may not be included in the classifier training set. An alternative strategy is to build fall detection specific classifiers that assume abundant training data for falls, which is hard to obtain in practice. Another challenge is the data collection for falls, as it may require a person to actually undergo falling which may be harmful, ethically questionable, and the falling incidences collected in controlled laboratory settings may not be the true representative of falls in naturalistic settings [6].

The research question we address in this paper is: *Can we recognise falls by observing only normal ADL with no training data for falls in a person independent manner?*. We use the HMMs for the present task as they are very well-suited for sequential data and can model human motions with high accuracy [7]. Typically, an HMM can be trained on normal activities and the maximum of negative of log-likelihood on the training data is set as a threshold to identify a fall as an outlier. However, choosing such a threshold may severely effect classifier's performance due to spurious artifacts present in the sensor data and most of the falls may be classified as normal activities. In this paper, we use the outlier detection approach to identify falls and present three X-Factor HMM based approaches for detecting short-term fall events. The first and second method models individual normal activities by separate HMMs or all normal activities together by a single HMM, by explicitly modelling the poses of a movement by each HMM state. An alternative HMM is constructed whose model parameters are the averages of the normal activity models, while the averaged covariance matrix is

* Corresponding author. Present address: 550, University Avenue West, Toronto, Ontario M5G 2A2, Canada.
E-mail addresses: s255khan@uwaterloo.ca (S.S. Khan), mkarg@uwaterloo.ca (M.E. Karg), dkulic@uwaterloo.ca (D. Kulić), jhoey@uwaterloo.ca (J. Hoey).

http://dx.doi.org/10.1016/j.asoc.2017.01.034
1568-4946/

artificially 'inflated' to model unseen falls. In the third method, an HMM is trained to model the transitions between normal activities, where each hidden state represents a normal activity, and adds a single hidden state (for unseen falls) with an inflated covariance based on the average of covariances of all the other states. The inflation parameters of the proposed approaches are estimated using a novel cross-validation approach in which the outliers in the normal data are used as proxies for unseen fall data. We present another method that leverages these outliers to train a separate HMM as a proxy model to detect falls. We also compare the performance of one-class SVM and one-class nearest neighbour approach along with several supervised classification algorithms that use full data for normal activities but the number of falls are gradually increased in the training set. We show that supervised classifiers perform worse when limited data for falls is available during training. This paper is a comprehensive extension of the work of Khan et al. [8] in terms of:

- Proposing two new models to detect unseen falls by (i) modelling transitions among normal activities to train an HMM and adding a new state to model unseen falls, and (ii) training a separate HMM on only the outliers in the normal activities data to model unseen falls.
- Data pre-processing, extraction of signals from raw sensor data, and number and type of features are different from Khan et al.[8].
- Studying the effect of changing the number of states on the proposed HMM methods for fall detection.
- Identifying similarity through experiments between the rejected outliers from the normal activities and the unseen falls.
- Additional experiments evaluating the effect of quantity of fall data available during the training phase on the performance of the supervised versions of the proposed fall detection methods and two other supervised classification methods.

2. Related work

The research in fall detection spans over two decades with several recent papers [2,9,10] that discuss different methodologies, trends and ensuing challenges using body worn, ambient or vision based fall detection techniques. Several research works in fall detection are based on thresholding techniques [11] or supervised classification [2]. One of the major challenges in fall detection is the less availability of fall data [5]; therefore, such techniques are difficult to use in practice. Keeping this view in mind, we survey techniques that attempt to detect falls by employing generative models, outlier/anomaly detection and one-class classification [12] based techniques that only use data from normal activities to build the model and identify a fall as an anomaly or outlier.

Thome et al. [13] present a Hierarchical HMM (HHMM) approach for fall detection in video sequences. The HHMMs first layer has two states, an upright standing pose and lying. They study the relationship between angles in the 3D world and their projection onto the image plane and derive an error angle introduced by the image formation process for a standing posture. Based on this information, they differentiate other poses as 'non-standing' and thus falls can be distinguished from other motions. A two-layer HMM approach, *SensFall* [14], is used to identify falls from other normal activities. In the first layer, the HMM classifies an unknown activity as normal vertical activity or 'other', while in second stage the 'other' activity is classified as either normal horizontal activity or as a fall. Tokumitsu et al. [15] present an adaptive sensor network intrusion detection approach by human activity profiling. They use multiple HMMs for every subject in order to improve the detection accuracy and consider the fact that a single person can have multiple patterns for the same activity. The data is collected using infra-red sensors. A new sequence of activity is fed to all the HMMs and likelihoods are computed. If all the likelihoods calculated from corresponding HMMs are not greater than pre-determined thresholds, then an anomaly is identified. Cheng et al. [16] present a fall detection algorithm based on pattern recognition and human posture analysis. The data is collected through tri-axial accelerometer embedded in the smartphones and several temporal features are computed. HMM is employed to filter out noisy character data and to perform dimensionality reduction. One-class SVM (OSVM) is applied to reduce false positives, followed by a posture analysis to counteract the missed alarms until a desired accuracy is achieved.

Zhang et al. [17] trained an OSVM from positive samples (falls) and outliers from non-fall ADL and show that the falls can be detected effectively. Yu et al. [18] propose to train Fuzzy OSVM on fall activity captured using video cameras and to tune parameters using fall and some non-fall activities. Their method assigns fuzzy membership to different training samples to reflect their importance during classification and is shown to perform better than OSVM. Popescu [19] presents a fall detection technique that uses acoustic signals of normal activities for training and detects fall sounds from it. They train OSVM, one-class nearest neighbour (OCNN) classifier and One-class GMM classifier (that uses a threshold) to train models on normal acoustic signals and find that OSVM performs the best; however, it is outperformed by its supervised counterpart. Medrano et al. [20] propose to identify falls using a smartphone as a novelty from the normal activities and found that OCNN performs better than OSVM but is outperformed by supervised SVM.

The supervised and thresholding techniques for fall detection collect artificial fall data in a laboratory under non-naturalistic settings; however, such fall data may not be true representative of actual falls and learning with them may lead to over-fitting. To overcome the need for a sufficient set of representative 'fall' samples, we propose three 'X-Factor' HMM based approaches to identify falls across different people while learning models only on data from normal activities.

3. Proposed fall detection approaches

The problem we investigate in this paper pertains to activity recognition and the datasets we use capture the temporal activities performed by humans. The Hidden Markov Models (HMM) are effective in modelling the temporal dynamics in data sequences and consider the history of actions when taking a decision on the current sequence. The HMM is a doubly stochastic process for modelling generative sequences that can be characterized by an underlying process generating an observable sequence. Formally, an HMM consists of the following components [21]:

- N – the number of hidden states in the HMM. The hidden states can be connected in several ways, for example in left-to-right manner or fully interconnected (ergodic). the set of states can be denoted as $S = \{S_1, S_2, \ldots, S_N\}$ and the state at time t as q_t.
- M – The number of distinct observation symbols per state that corresponds to the physical output of the system being modelled. The symbols can be denoted as $V = \{v_1, v_2, \ldots, v_M\}$. When the observation is continuous, $M = \infty$, and can be approximated using Gaussian or mixture of Gaussian with mean and covariance corresponding to each hidden state as the underlying parameters.
- A – The state transition probability distribution $A = a_{ij}$, where a_{ij} represents the probability of state j following state i and is expressed as:

$$a_{ij} = P[q_{t+1} = S_j | q_t = S_i] \qquad 1 \le i, j \le N \tag{1}$$