

Modelling of Behavioural Patterns for Abnormality Detection in the Context of Lifestyle Reassurance

Fabien Cardinaux¹, Simon Brownsell¹, Mark Hawley¹, and David Bradley²

¹ School of Health and Related Research, University of Sheffield, UK
{F.Cardinaux, S.Brownsell, Mark.Hawley}@shef.ac.uk

² School of Science & Engineering, University of Abertay Dundee, UK
D.Bradley@abertay.ac.uk

Abstract. As a consequence of the growing number of older and vulnerable people, health and care providers are increasingly considering new approaches to support people in their own homes. In this context, lifestyle reassurance analyses data collected from a range of sensors to determine a person's 'routine' and highlights any important changes. This paper proposes a new approach for detection of individual deviation from normal behaviour focusing on building probabilistic models of behaviour based on a set of activity attributes. Models are trained using only normal behaviour. Variations from the models are considered as abnormal behaviours and these can be highlighted for subsequent review or intervention. Case study experiments with real life data suggest that some users' activities follow regular patterns and that these patterns can be learned with probabilistic models.

Keywords: Lifestyle reassurance, probabilistic models, GMM, telecare.

1 Introduction

Faced throughout much of the world with aging populations, health and care providers are considering new approaches to health and care provision, including the use of telecare to support growing numbers of older or vulnerable people. Telecare can be defined as the delivery of health and social care services to people in their own home by means of telecommunications and computer-based systems [2]. Lifestyle reassurance (sometimes called lifestyle monitoring) can then be considered to be subset of telecare where deviations from 'normal' patterns of behaviour are sought to determine that an intervention may be beneficial for the user; for example, following a change in sleep pattern. When applied in a home setting, activity is recorded by a range of sensors for motion detection, space and storage utilisation, and appliance use, located throughout the home. Data from these sources enables the establishment of what can be considered a normal profile of behaviour. Deviations and variations from this established norm can then be detected so that early signs of deteriorating health or quality of life can be investigated and an appropriate intervention provided.

One of the key issues in building such a system is to establish an accurate representation of normal behaviour. This paper explores a way of building probabilistic models of normal behaviour using historical lifestyle reassurance data. It also shows how these models could be used to detect variations from the norm, and the subsequent possibility that additional care and support may be required.

Various publications have addressed the problem of lifestyle activity or behaviour recognition [1,9,12], yet relatively little effort has focused on abnormality detection in the context of lifestyle reassurance, with some noticeable exceptions [8,10,11]. The methods presented in these three papers can fall into the broad category of histogram based approaches. Virone *et. al.* [10,11] presented an approach where, for each hour of the day, the average proportion of time spent within a specific room is estimated. Large deviations from this proportion are considered as abnormal. The system proposed by Ohta *et. al.* [8] is based on a similar concept, except that the average time spent in each room is estimated for each day (not each hour) and they also propose to monitor other variables such as the daily distance moved by a person.

The approach suggested in this paper is based on, what we define as a Probabilistic Model of Behaviour (PMB). A set of attributes or features is associated with each occurrence of this type of activity and a probabilistic model of this type activity is learned from historical data. Deviation from the model is then considered as abnormal.

In the proposed approach, daily activities are detected using rule based algorithms, each activity being associated with a set of attributes. Using automatically detected activity occurrences in historical data (the training set), the parameters of the Gaussian mixture model (GMM) are learned for each type of activity. Using the test set, when an activity is detected, the activity likelihood given the model of this type of activity can be estimated, providing an opinion on whether this can be considered to be normal behaviour. Comparing likelihoods to a pre-set threshold can then be used to generate alerts when interventions for a specific user may be beneficial.

A previous 12 month evaluation of a commercial system conducted by the authors provides a significant source of lifestyle reassurance data [3] on which to test the suggested approach.

2 Proposed Approach

2.1 Abnormal Behaviour Detection

Automatic detection of abnormal behaviour is challenging due to the nature of such abnormal behaviour, which is, by definition, rare and unexpected. It has attracted attention in a variety of fields, such as computer vision [13], fault detection in engine turbines [7] and telephone fraud detection [6]. A groundwork on activity monitoring is presented in [6] where the authors differentiate two main groups of methods:

- Profiling methods that build a model of normality. It is then possible to detect any deviation from this model. The approaches presented by Tarassenko *et. al.* [7] and Virone *et. al.* [10,11] fall into this category.
- Discriminating methods that build a model which is trained to discriminate an abnormal activity from a normal activity. It is assumed that abnormal events have occurred before (in the training set). The approaches presented by Zhang *et. al.* [13] and Chan *et. al.* [4] fall into this category.

While it might be possible to build discriminative models of abnormality for some applications (*eg.* mobile phone fraud detection [6]), in a lifestyle reassurance scenario, abnormal behaviour can occur in countless forms and it therefore appears problematic to be able to build a good generic model of abnormality.

Consequently we have chosen to build statistical descriptions of normality by including historical data which can be regarded as representing normal behaviour and then detecting any deviation from these models (profiling approach). The primary advantage of this approach lies in being able to detect abnormality, even when this had not previously been seen or expected.

Fawcett *et. al.* [6] also differentiate uniform profiling that builds a general profile of normal activity, assuming it is uniform across the users and individual profiling that builds a profile of normal activity of each user. We show in Section 4.1 that for our task, behaviours are not uniform across users and therefore opt for an individual profiling strategy.

In the context of lifestyle reassurance, it can be expected that deviations from normal behaviours is an indicator of the need for care or assistance, however we acknowledge that it is not necessarily always the case that an abnormal behaviour requires a health or care intervention.

2.2 Probabilistic Formulation

This section sets out a probabilistic formulation of abnormal behaviour detection. Where such abnormal behaviour is defined as a behaviour that deviates from normality. Let us denote the parameter set λ describing a normal behaviour. Given a new observed behaviour \mathbf{x} defined by a set of I attributes a_i , we find the likelihood of this behaviour to be a normal behaviour $P(\mathbf{x}|\lambda)$. In practice, likelihoods are usually compared in log domain in order to avoid computer underflow problems. An opinion on the behaviour to be a normal behaviour is thus estimated with the log-likelihood:

$$A(\mathbf{x}) = \log(P(\mathbf{x}|\lambda)) \quad (1)$$

A decision can then be taken whether to raise an alert which requires an intervention based on the $A(\mathbf{x})$ value. Given a pre-set threshold τ , the behaviour is considered as a normal behaviour if $A(\mathbf{x}) > \tau$ and as an abnormal behaviour if $A(\mathbf{x}) \leq \tau$. The thresholds can be set as a function of mean and standard deviation of the log-likelihoods resulting from the training set such as:

$$\tau = m(A_{train}) - \alpha \cdot std(A_{train}) \quad (2)$$

where $m(\Lambda_{train})$ and $std(\Lambda_{train})$ are respectively the mean and the standard deviation of the log-likelihoods from the training set Λ_{train} and α is a value that sets the sensitivity of the system.

Note that Virone *et. al.* [10,11] propose to set several thresholds values corresponding to different levels - indicating different priorities and/or responses to the end user. Another possibility would be to raise an alert only when several behaviours have exceeded the threshold more than a pre-set number of times during a specific period of time. Also the thresholds value can be re-set over time using recently collected data. However, this paper focuses on the modelling of behaviour and other aspects are outside the scope of this paper.

2.3 Gaussian Mixture Models

The main challenge of this approach is to find a good probabilistic model of behaviour to describe normality. In this paper it is proposed to model normal behaviour with Gaussian mixture model (GMM). Assuming that a set of behaviour can be modelled as a mixture of N_g gaussian distributions, the likelihood of an individual behaviour can be estimated as

$$P(\mathbf{x}|\lambda) = \sum_{k=1}^{N_G} w_k \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \Sigma_k) \quad (3)$$

$$\lambda = \{w_k, \boldsymbol{\mu}_k, \Sigma_k\}_{k=1}^{N_G} \quad (4)$$

Here, $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \Sigma_k)$ is a D-dimensional Gaussian density function [5] with mean $\boldsymbol{\mu}_k$ and diagonal covariance matrix Σ_k . N_G is the number of Gaussians and w_k is the weight for Gaussian k . The parameters for the generic model are estimated using maximum likelihood criterion. In other words we find the model parameters that fit the best normal behaviour data used to train the model. This is achieved by Expectation Maximization algorithm [5].

GMMs have been previously used in the context of lifestyle reassurance by Barger *et. al.* [1], however, while their approach presents some similarity with that being proposed, the objectives are different. Indeed, while our main objective is to find deviation from normal behaviour, in [1], the authors use a GMM as an unsupervised clustering technique in order to recognise activities.

3 Lifestyle Reassurance Data

3.1 Setup

After the appropriate ethical approval and consents, a lifestyle reassurance system was installed in a UK sheltered housing scheme with 24 users. The system consisted of bed and chair occupancy sensors, passive infra-red movement detectors, door contact monitors, and electrical usage sensors. The sensor activations were sent via a telephone line to a central server [3]. The resulting data were split into two sets, namely a training set and a test set. The training set is composed

of 26 weeks of data and is used for preliminary data analysis and models training. The test set comprises 8 weeks of data and is used to observe abnormalities by detecting variations from the trained model.

3.2 Observations

An observation represents a lifestyle activity occurrence that can be identified with series of sensor activations. Activities can be, for example, movement in the kitchen, TV usage, usage of another electrical appliance, and so on. Consequently, an observation is a first level of interpretation of the raw data (sensor activations). A set of attributes can be associated with each observation. Among the possible attributes are:

- **Start time:** Time of day the observation took place.
- **Duration:** Amount of time the behaviour was observed for.
- **Weekend or Weekday:** Which day type it occurred on.
- **Activity Level:** Number of sensors triggered during the observation.

The determination of the list of behaviours to monitor, as well as their associated attributes, in order to maximise the quality of life and health of the user is the subject of ongoing studies by the authors.

Another parameter often regarded as a good indicator for lifestyle reassurance is the frequency of observation. The frequency of an activity is defined by the number of times the behaviour was observed during a specific time period.

While this paper intends to present a general approach for lifestyle reassurance data modelling, the results presented in Section 4 focus on the modelling of sleep pattern and television usage. These activities are defined by their start time and duration.

Rule-based algorithms are used to detect these activities. For example, a sleeping period is detected when a bed sensor “in” event is observed followed by a sensor “out” event and the duration (*ie.* time lapse between these two events) is longer than 30 minutes. The start time is then defined as the “in” event and is stored as a real value representing the number of hours since midnight. The duration is stored as a real value representing the number of minutes between “in” and “out” events. It is assumed that the user is sleeping when the bed is occupied for more than 30 minutes. As a result of this algorithm, we can build a dataset X composed of a list of activities \mathbf{x}_n :

$$X = \{\mathbf{x}_n\}_1^N \quad (5)$$

where N is the number of detected activities and \mathbf{x}_n are the activity vectors defined by a set of attribute a_i .

$$\mathbf{x}_n = [a_1, a_2, \dots, a_I] \quad (6)$$

Figure 1 shows a graphical representation of an example of detected activities for one user over a single week. As an example, the first detected sleeping period on the graph starts at 01:55 ($a_1 = 1.91$) and last for 3.5 hours ($a_2 = 210.7$).




Fig. 1. Detected TV Usage and Sleeping activities for one user over a single week. Each horizontal stripe represents a single activity occurrence. Several sleeping periods can be detected in a single night if the user gets out of bed during night.

4 Results

The case studies present results for Sleeping and TV Usage activities. Each Feature vector x_n (see equation 6) representing a detected activity is defined with start time and duration attributes. The case study is presented for two users randomly chosen among those where a TV electrical appliance sensor and a bed sensor were installed.

4.1 Data Analysis

Figure 2 shows plots of sleeping and TV Usage activities represented by start time and duration for two users. It can be seen that for both of these activities and for both users some regular behaviour patterns are evident as well as a dependency between start time and duration. For example, the user A appears to have two main sleeping periods: the first tends to start between 10pm-11pm




Fig. 2. Activity Patterns. Plots of sleeping (top) and TV usage (bottom) activities represented by start time and duration for two users. Each point represents a detected activity occurrence.

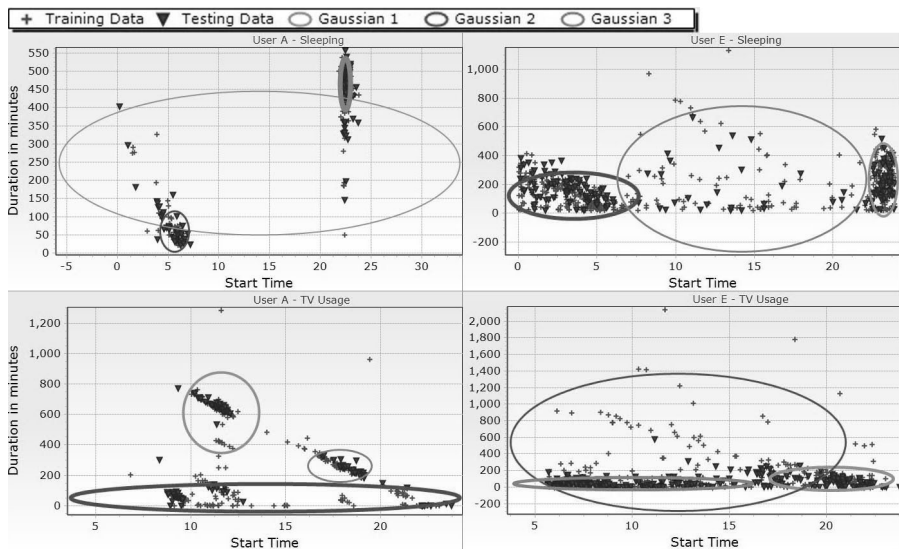


Fig. 3. Gaussian mixtures along with train and testing data plots of sleeping (top) and TV usage (bottom). Activities from the training set are represented with a cross symbol while activities from the test set are represented with triangles. Gaussians are represented by ellipses. The ellipse is centred at the mean μ of the Gaussian, the size of the ellipse in each dimension represents the variances Σ and the width of the line represents the weight w of the Gaussian k .

and last for about 7.5 hours (450 minutes) and the second starts at around 6am and last for about 50 minutes. This suggests the user gets up for a short while and then goes back to bed.

The observed dependency between attributes advocates that multidimensional probability density functions to model the behaviour is more adapted than modelling independently the different attributes with one-dimensional models. We also notice that different users have different behaviour patterns which suggests training user-specific model as opposed to user independent models.

4.2 Modelling Experiments

Following the approach presented in Section 2.3, GMM parameters (means, variances and weights) are learned using the 26 weeks of training data. We heuristically choose to train models with three Gaussians. Figure 3 shows the data plots where Gaussians are represented by ellipses.

Table 1 reports the alert detection results on the test set where thresholds are set with $\alpha = 3$ (see equation 2). Only four alerts have been detected using the chosen thresholds. By looking closely at these alerts we notice that the two alerts on sleeping activity for user A and user E are both long sleeping periods starting later than usual. The TV usage alert for user E is generated by a usage of the TV

Table 1. Detected alerts on the test set

Activity User	A_{train}		τ	Number of Alerts	Alerts		
	mean	std			Time Start	Duration (min)	$A(x)$
Sleep	A	-10.32 \pm 2.13	-16.7	2	0 : 12	402	-17.85
	E	-12.93 \pm 1.52	-17.4	1	23 : 43	357	-16.75
TV	A	-12.90 \pm 2.29	-19.8	1	0 : 09	381	-18.00
	E	-13.24 \pm 1.62	-18.1	0	8 : 21	302	-19.97

occurring earlier than usual and that lasts for approximately five hours which is not usual at that time of the day for this user.

This work is part of an ongoing project which is intended to form the basis of a new approach to lifestyle reassurance. While experiments show examples of generated alerts in real life situations, we were not able to assess the pertinence of these alerts. Ongoing data collection with the real-time running of the algorithms as well as feedback from the users will allow us to perform further evaluation of the approach.

5 Conclusions

This paper proposed a novel framework for abnormality detection in the context of lifestyle reassurance. The focus was in building probabilistic models of behaviours (PMB) based on a set of activity attributes. The PMB are trained using only normal behavioural data and variations from the models can be detected to generate alerts.

Experiments based on real life case studies show that user activities follow regular patterns that can be learned by the PMB and that individual deviation from normality can be detected.

Compared with previously published works [10,11,8], where histogram based approaches are used, the proposed approach has potential advantages as it builds a model for each type of activity to monitor. Moreover, the use of multi-dimensional GMMs assures that dependency between different attributes can be taken into consideration (time and duration as presented in the case studies). As an example, for a user, watching the TV for five hours and starting to watch TV at 8am can each be normal if we consider these two attributes independently, however it might be unusual to watch TV for five hours starting at 8am.

A limitation of the proposed method, where alerts are generated for a specific activity occurrence, is that no alert is generated in case of absence of activity occurrence. This problem can be tackled by generating alerts on activity frequency which would lead to a combination between PMB and histogram based approaches.

In future work, we will investigate the addition of more activity attributes such as the activity level (*ie.* number of sensor activations during activity) and activity frequency (which is the number of occurrences of an activity within a given period of time). We will also explore the inclusion of a categorical or

binary variable in an activity's set of attributes such as week/week-end. Further evaluation of the generated alerts and relationship to subsequent health or care interventions, and how such information can be fed back into the system, is also under consideration.

Acknowledgement

We are grateful to project participants and Department of Health (HTD313) for funding this work.

References

1. Barger, T.S., Brown, D.E., Alwan, M.: Health-status monitoring through analysis of behavioral patterns. *IEEE Trans. on Sys., Man and Cybernetics* 35, 22–27 (2005)
2. Barnes, N.M., Edwards, H.M., Rose, D.A.D., Garner, P.: Lifestyle monitoring-technology for supported independence. *Comp. & Control Eng. J.* 9, 169–174 (1998)
3. Brownsell, S., Blackburn, S., Hawley, M.: An evaluation of second and third generation telecare services in older people's housing. *J. of Telemedicine and Telecare* 14, 8–12 (2008)
4. Chan, M.T., Hoogs, A., Schmiederer, J., Petersen, M.: Detecting rare events in video using semantic primitives with HMM. In: 17th International Conference on Pattern Recognition, vol. 4, pp. 150–154 (2004)
5. Duda, R., Hart, P., Stork, G.: *Pattern Classification*. Wiley, Chichester (2001)
6. Fawcett, T., Provost, F.: Activity Monitoring: Noticing interesting changes in behavior. In: *The conference on Knowledge Discovery in Data*, pp. 53–62 (1999)
7. Tarassenko, L., Nairac, A., Townsend, N., Buxton, I., Cowley, P.: Novelty detection for the identification of abnormalities. *International Journal of Systems Science* 31, 1427–1439 (2000)
8. Ohta, S., Nakamoto, H., Shinagawa, Y., Kishimoto, T.: Home Telehealth: Connecting Care Within the Community. *Medical telematics*, 198–209 (2006)
9. Osmani, V., Balasubramaniam, S., Botvich, D.: Human activity recognition in pervasive health-care: Supporting efficient remote collaboration. *J. of Network and Computer Applications* (in press)
10. Virone, G., Noury, N., Demongeot, J.: A system for automatic measurement of circadian activity deviations in telemedicine. *IEEE Trans. on Biomedical Engineering* 49, 1463–1469 (2002)
11. Virone, G., Alwan, M., Dalal, S., Kell, S., Turner, B., Stankovic, J.A., Felder, R.: Behavioral Patterns of Older Adults in Assisted Living. *IEEE Trans. on Information Technology in Biomedicine* (in press)
12. Wyatt, D., Philipose, M., Choudhury, T.: Unsupervised Activity Recognition Using Automatically Mined Common Sense. In: *Proceedings of AAAI 2005*, pp. 21–27 (2005)
13. Zhang, D., Gatica-Perez, D., Bengio, S., McCowan, I.: Semi-Supervised Adapted HMMs for Unusual Event Detection. *IEEE Computer Vision and Pattern Recognition* 1, 611–618 (2005)