Modeling Visit Behaviour in Smart Homes using Unsupervised Learning

Ahmed Nait Aicha
Dept of Computer Science
Amsterdam University of
Applied Sciences
Amsterdam, The Netherlands
a.nait.aicha@hva.nl

Gwenn Englebienne
Dept of Computer Science
University of Amsterdam
Amsterdam, The Netherlands
g.englebienne@uva.nl

Ben Kröse
Dept of Computer Science
University of Amsterdam
Amsterdam, The Netherlands
b.j.a.krose@uva.nl

UbiComp '14, September 13 - 17 2014, Seattle, WA, USA
Copyright 2014 ACM 978-1-4503-3047-3/14/09...$15.00.
http://dx.doi.org/10.1145/2638728.2638809

Abstract
Many algorithms on health monitoring from ambient sensor networks assume that only a single person is present in the home. We present an unsupervised method that models visit behaviour. A Markov modulated multidimensional non-homogeneous Poisson process (M3P2) is described that allows us to model weekly and daily variations and to combine multiple data streams, namely the front-door sensor transitions and the general sensor transitions. The results from nine months of sensor data collected in the apartment of an elderly person show that our model outperforms the standard Markov modulated Poisson process (MMPP).

Author Keywords
Ambient Assisted Living (AAL), Smart Homes, Markov Modulated Poisson Process (MMPP)

ACM Classification Keywords
I.2.6 [ARTIFICIAL INTELLIGENCE]: Learning; I.5.4 [PATTERN RECOGNITION]: Application Signal Processing; G.3 [PROBABILITY AND STATISTICS]: Probabilistic algorithms

Introduction
In recent years there has been a lot of work on ambient sensor networks for monitoring human activities in a home

environment, see [1, 2] for surveys. Most studies on activity recognition, anomaly detection or trend detection assume that there is only a single resident in the home. However, using supervised learning, multi-person activity recognition methods in smart homes have been presented, either based on joint location uncertainty [8], naive Bayes [7], activity models trained on individuals [3] or joint activity modeling [11].

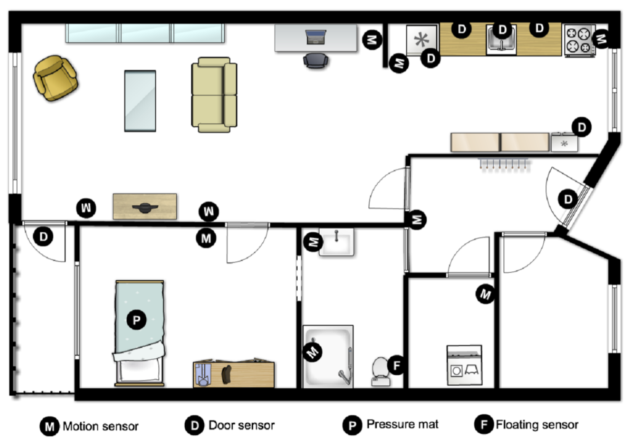

Figure 1: A map of the volunteer's apartments equipped with a wireless sensor network. The orientation of the motion sensor is indicated by the orientation of the letter M.

Currently our group monitors about 20 elderly living alone using sensor networks. For a correct health assessment it is important that we are sure that the data originates from the activities of the resident only and that we know whether there are visitors. Supervised methods for visitor detection have been presented [6], but our future projects will involve even more people, and we can not rely on supervised methods. For that reason, the objective of our research is to develop and test an unsupervised method that is able to detect visits from ambient sensor data.

We consider visits as an abnormal activity that will increase sensor counts. Unsupervised methods for the detection of abnormal behaviour in smart homes have been presented for applications like fall detection or wandering [5] but not for visit detection. A model that has been successfully applied for the detection of anomalous events in sensor counts is the Markov modulated Poisson process (MMPP) [10, 4]. However, MMPPs are univariate, and as such cannot deal with the richer datasets that are common to AAL. In this paper we present the Markov Modulated Multidimensional heterogeneous Poisson Process (M3P2) for the detection of visits. The contribution of our research is twofold: a) the novel Multidimensional Markov Modulated Poisson Process (M3P2) model can deal with multiple data streams, and b) we show that the heterogeneous model allows us to model weekly and daily variations and to distinguish between normal and abnormal visits. In order to measure the performance of our model we carried out experiments with real-life sensor data. We use a data set, that will be made public, consisting of nine months of sensor data collected in the apartment of an elderly person. The results show that our model outperforms the standard MMPPs.

Sensor Data

We have continuously collected data for more than a year in the apartment of an elderly person, in a care centre in The Netherlands, which is equipped with a sensor network. Our sensor network uses the Z-Wave protocol and consists of off-the-shelf binary sensors that measure motion, pressure on the bed, toilet flushing and the opening and closing of cabinets and doors. An overview of the location of the sensors in the apartment of one resident is shown in

Figure 1. The elderly are living their routine life and are not told to modify their behaviour in any way.

A set of 9 months sensor data, collected from one apartment between March 31 and December 31, 2013, is used to conduct the experiments in this paper. Although our method is unsupervised, we need ground truth data for evaluation. This is done by visually inspecting the raw sensor data. To simplify the search, we relied on the information given by the resident during several interviews. To find the unusual visits, we focused on the times at which we know the resident occasionally receives visits from his children. To find the unusual *absence* of visits, we focused our search on the daily visits from a caregiver every day around 8:30 in the morning and 9 in the evening and the weekly visits of the cleaner every Friday between 9 and 12 in the morning.

Feature Extraction

Similar to [6], sensor-transitions, defined as a tuple of two consecutive sensor readings, are used as features. However, we particularly select features that deal with the presence of multiple people (*i.e.* visits). Define $N^M(t)$ to be the number of transitions during time slice t, between sensors that are *not* topologically connected. We define two sensors to be topologically connected if one person can activate the sensors consecutively without activating another sensor. For example, the sensors from the bedroom and the bathroom are topologically connected, while the sensors from the kitchen and the bathroom are not topologically connected. The only way the resident can move from the kitchen to the bathroom is through the living-room and the bedroom, as shown in Figure 1. In addition to these, we define $N^D(t)$ to be the number of sensor-transitions during time slice t, for which one of the sensor readings originates from the front-door sensor.

Both features, $N^M(t)$ and $N^D(t)$, are informative and should be taken into account in the modelling of visits behaviour.

Multidimensional MMPP

The Poisson process is a widely used stochastic process for modelling counts of random events that occur during a time interval. A Markov Modulated Poisson Process (MMPP, [9]) is a non-homogeneous Poisson process, where the rate $\lambda(t)$ of the count data $N(t)$ varies over time and follows a Markov chain. We follow [4], where both the periodic and non-periodic influences on the count data are modelled. The periodic aspects (*i.e.*, daily and weekly cycles) are modelled by decomposing the rate $\lambda(t)$ as follows:

$$\lambda(t) = \lambda_0 \cdot \delta_{d(t)} \cdot \eta_{d(t),h(t)} \tag{1}$$

where λ_0 represents the average rate over a full week, δ_j represents the effect of the day j of the week and $\eta_{j,i}$ represents the effect of the time i of day j of the week.

An important limitation the MMPP is that the model is restricted to one-dimensional observations. In our application, both the counts $N^M(t)$ and $N^D(t)$ are informative and should be taken into account. Although multivariate MMPP have been analysed [12], this does not extend to non-homogeneous MMPP. We therefore extend the model to multiple simultaneous count features, resulting in M3P2. A graphical model of the M3P2 is given in Figure 2.

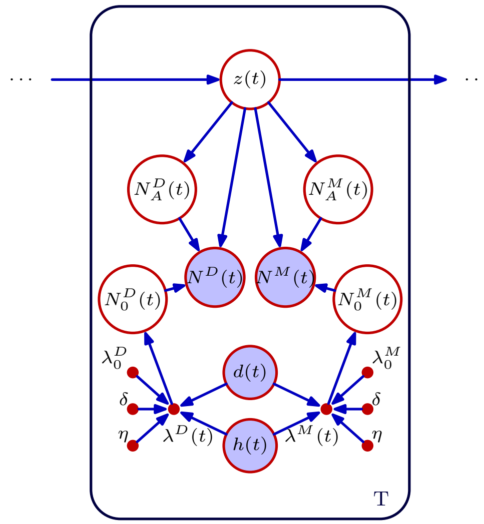

Figure 2: Graphical model for different distributions and variables defining the M3P2. The shaded nodes are observed and the small solid nodes indicate deterministic parameters. The directed edges indicate conditional dependence of the nodes. $\lambda^M(t)$ depends on the parameters $\lambda_0^M, \delta, \eta, d(t)$ and $h(t)$, while $N^M(t)$ depends on $N_0^M(t), N_A^M(t)$ and $z(t)$.

Let $N^i(t)$ denote the i-th observation stream. Similar to the MMPP, we model each observed count $N^i(t)$ as the sum of the base counts $N_0^i(t)$ and the variation (positive or negative) $N_A^i(t)$ due to a visit (presence of absence). Assume that these counts are independent. As with MMPP, $z(t)$ denotes the latent variable with the states -1 (absence of visit), 0 (normal) and 1 (visit). In our application, it is advantageous to explicitly model the usage of the front-door by introducing a new state $z(t) = 2$. This state acts as a gating state without which it is impossible to transition from the absence of a visit to its presence, and vice-versa.

The calculation of the posterior probability $p(z(t)|N^M(t), N^D(t))$ in case of two data stream is done using the Markov Chain Monte Carlo (MCMC) sampling method following [4].

Experiments

A set of three experiments is conducted to study the performance of the model. In the first experiment we studied the effect of the time discretization: the duration of the time-slice in which we count the transitions. If this is too small we may have too few counts to make a good model, if it is too large we lose resolution and will miss information: the duration of a visit may vary from few minutes to several hours. A second experiment was carried out to investigate the performance of the model in capturing temporal variations. We model the effect of the daily and weekly periodicity by modulating the rate λ with parameters δ and η. Apart from this periodicity we expect also a yearly (seasonal) periodicity, which we did not model as our data spans only 9 months. To investigate these effects, we trained the model with a sequence of 13 weeks and the full sequence of 39 weeks. In the third experiment we compared our model M3P2 with the baseline, the standard MMPP model.

To measure the performance of the model we used the annotated data as described in Section 2. In order to map the ground truth into discrete time slices, we decided to label all the time slices that are fully or partially covered by an unusual visit as a positive class ($z = 1$). The same procedure holds for the unusual absence of a visit ($z = -1$). The precision, recall and the F-value are computed separately for the two states ($z = 1$ and $z = -1$). An unusual visit lasting more than one time slice is defined to be correctly detected if at least one time slice of this visit is correctly detected by the classifier.

Results

Effect of time discretization

We varied the number of time slices in a day $D \in \{24, 12, 8\}$, corresponding to time slice lengths of $1, 2, 3$ hours. The reason for limiting these values to only 3 is because the majority of the unusual visits has a duration of 3 hours or less. We performed 3-fold cross validation and report the average results. The results, listed in Table 1, show that the precision increases with D while the recall decreases, resulting in little change in the F-values.

	slice length	unusual visits ($z = 1$)		
D	(hours)	precision	recall	F-value
8	3	0.67	0.75	0.71
12	2	0.63	0.83	0.72
24	1	0.54	0.95	0.69

Table 1: Precision, recall and F-value obtained when MMPP is applied on $N(t)$. The standard deviation is less than 0.1 in all cases.

These results may be explained by the fact that decreasing the time slice length, which implies an increase of D, facilitates the discovery of (short) unusual visits. As a consequence, the recall will increase. On the other hand, decreasing the time slice length will decrease the number of observation per time slice, which will increase the variance. As a consequence, the number of false positives will increase, which decreases the precision. The results of the unusual absence of visits ($z = -1$) are not reported because of lack of data, resulting in a very large standard deviation of the F-value. The first fold (calendar weeks 14-27) and the third fold (calendar weeks 40-53) have only one or two 'unusual absences of visit', while the other 'unusual absences' lie in the second fold. In our case, the recall was almost always equal to 1 (respectively equal to 0) when testing with the first fold (respectively the third fold). We chose for $D = 24$ because it gives the best compromise between the performance and the practical usage.

Temporal variations

An experiment where W, *i.e.* the period in which we assume that behaviour does not change, is set to different values between 4 weeks and 39 weeks (approximately 9 months) is conducted. Setting $W = 39$ means that we assume there are no seasonal influences on the behaviour, while setting $W = 4$ means that we assume there are 'monthly' influences on the behaviour. The results given in Table 2 show that the best performance is obtained when using a period of 13 weeks long. This is remarkably close to the duration of a meteorological season, and it seems very plausible that there is a seasonal effect in the data. It will be interesting, in future work, to incorporate this in the model and evaluate it on multi-year data. The results for the absence of regular visit ($z = -1$) are not reported for the same reason as the previous experiment.

W	unusual visits ($z = 1$)			std dev
(weeks)	precision	recall	F-value	F-value
4	0.58	0.84	0.66	0.12
6	0.57	0.87	0.67	0.10
8	0.58	0.87	0.68	0.10
13	0.56	0.96	0.71	0.03
39	0.48	0.97	0.64	0.01

Table 2: Precision, recall and F-value obtained when MMPP is applied on $N(t)$ using $D = 24$ and different values for the period W (*i.e.* different total number of time slices T).

Comparison of the models M3P2 and MMPP

For the application of M3P2 on the collected data and the comparison with MMPP, the results of the first two experiments applied on MMPP are used. Hence, the priors that resulted in the best performance of MMPP are used as start point to find the best priors for the first data stream $N^M(t)$ of M3P2.

The results of this experiment, listed in Table 3 show that the M3P2 results in higher F-values than MMPP. The better precisions of M3P2 compared to MMPP reflect the lower false positives obtained when M3P2 is used. This can be explained by the fact that M3P2 gives the possibility to separately set the priors for each data stream, which influences the precision in a positive way.

Model (MMPP/M3P2)	unusual visits $(z = 1)$ precision	recall	F-value
MMPP	0.56	0.96	0.71
M3P2	0.64	0.87	0.74

Table 3: Precision, recall and F-value obtained when M3P2 is applied on $N^M(t)$ and $N^D(t)$ using one hour as time slice length (*i.e.* $D = 24$) and a season for the total time period T (*i.e.* $W = 13$). The variance is less than 0.01 in all cases.

Most of the MMPP's false positives are caused by slight temporal shifts of the nurse's daily visits. An earlier (respectively later) visit of the cleaner results in a higher value of $N^M(t)$ in the preceding (respectively following) time slice than the one in which the visit normally takes place. An example of such visits, which last few minutes, is shown in Figure 3(a). Another reason for false positives, in both models, is the lack of an accurate annotation, it was in some cases difficult to determine whether a visit took place or not by only using the raw sensor data. We chose in such case to not label that time slice as a visit, which may result in false positives. An example of such a case in shown in Figure 3(b).

Conclusions

In this paper we presented an unsupervised method for the detection of abnormal visits in the home of an elderly. The method is based on a MMPP. Our method, referred as M3P2, solves the limitation of MMPP that can only deal with one-dimensional observation stream. The M3P2 is tested to nine months of sensor data, collected in the apartment of an elderly living alone person. The study has shown that M3P2 is able to detect abnormal visits with a clearly higher f-measure than the MMPP. In particular, the reduced amount of false positives reflected in the much higher precision is of great practical importance in care environments. In addition, the approach is able to model daily and weekly characteristics, and can be used to distinguish between recurrent and irregular visits.

The findings of this study have a number of important implications for future practice. Detecting visits and analysing them gives an insight in the social life of the resident. In future work, applying the M3P2 to other features, such as the mobility at home, may result in an insight in the functional health status of the resident. These insights can be used to generate different kinds of alarms.

Acknowledgments

This work is part of the research programs SIA-raak Smart Systems for Smart Services, Health-lab and COMMIT/. The authors would like to thank the participants at Vivium Zorggroep Naarderheem.

References

[1] Acampora, G., Cook, D. J., Rashidi, P., and Vasilakos, A. A survey on ambient intelligence in healthcare. *Proceedings of the IEEE 101*, 12 (2013), 2470 – 2494.

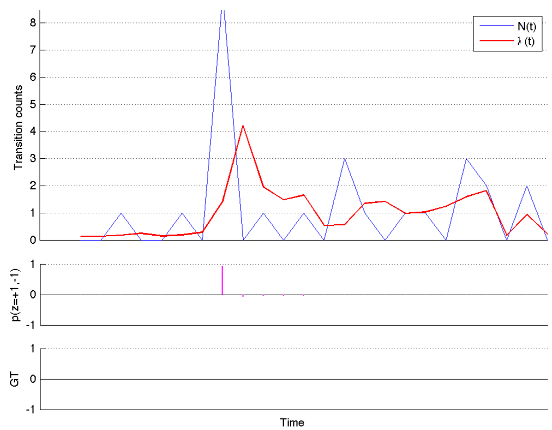

(a) Tuesday April 9. The daily visit of the nurse in the morning is slightly earlier than usual.

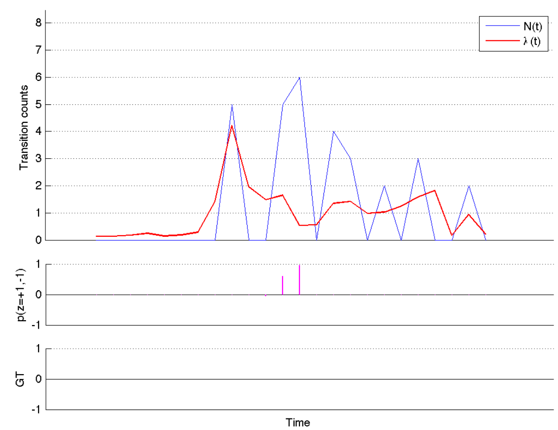

(b) Wednesday October 30. It is difficult to determine from the raw data if the detected false positives were truly false positives.

Figure 3: MP-transition counts $N(t)$ along with $\lambda(t)$, the corresponding posterior probabilities (p(z)) and the ground truth.

[2] Chen, L., Hoey, J., Nugent, C., Cook, D., and Yu, Z. Sensor-based activity recognition. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on 42*, 6 (2012), 790–808.

[3] Crandall, A. S., and Cook, D. J. Coping with multiple residents in a smart environment. *Journal of Ambient Intelligence and Smart Environments 1*, 4 (2009), 323–334.

[4] Ihler, A., Hutchins, J., and Smyth, P. Learning to detect events with markov-modulated poisson processes. *ACM Transactions on Knowledge Discovery from Data (TKDD) 1*, 3 (2007), 13.

[5] Lotfi, A., Langensiepen, C., Mahmoud, S., and Akhlaghinia, M. Smart homes for the elderly dementia sufferers: identification and prediction of abnormal behaviour. *Journal of Ambient Intelligence and Humanized Computing 3*, 3 (2012), 205–218.

[6] Petersen, J., Larimer, N., Kaye, J. A., Pavel, M., and Hayes, T. L. Svm to detect the presence of visitors in a smart home environment. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, IEEE (2012), 5850–5853.

[7] Phua, C., Sim, K., and Biswas, J. Multiple people activity recognition using simple sensors. In *PECCS* (2011), 224–231.

[8] Roy, N., Roy, A., and Das, S. K. Context-aware resource management in multi-inhabitant smart homes a nash h-learning based approach. In *Pervasive Computing and Communications, 2006. PerCom 2006. Fourth Annual IEEE International Conference on*, IEEE (2006), 11–pp.

[9] Scott, S. L. Detecting network intrusion using a markov modulated nonhomogeneous poisson process. *Submitted to the Journal of the American Statistical Association* (2001).

[10] Scott, S. L., and Smyth, P. The markov modulated poisson process and markov poisson cascade with applications to web traffic data. *Bayesian Statistics 7* (2003), 671–680.

[11] Singla, G., Cook, D. J., and Schmitter-Edgecombe, M. Recognizing independent and joint activities among multiple residents in smart environments. *Journal of ambient intelligence and humanized computing 1*, 1 (2010), 57–63.

[12] Sumita, U., and Masuda, Y. Analysis of multivariate markov modulated poisson processes. *Operations research letters 12*, 1 (1992), 37–45.