How Lonely is your Grandma? Detecting the Visits to Assisted Living Elderly from Wireless Sensor Network Data

Ahmed Nait Aicha
Dept of Computer Science
HvA 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

Abstract
Existing research on the recognition of Activities of Daily Living (ADL) from simple sensor networks assumes that only a single person is present in the home. In reality, the resident receives visits from family members or professional health care givers. In such cases activity recognition must take into account the presence of multiple persons. Here we investigate the problem of detecting multiple persons in a home environment equipped with a sensor network consisting of 13 binary sensors. We collected data during more than one year in our living labs and used Hidden Markov Model (HMM) for a visitor detection. A cross validation method was used to determine the best set of features from the binary data. Using this set of features the detection rate is approximately 85%.

Author Keywords
Ambient assisted living and health monitoring, Sensor networks for pervasive health care, Hidden Markov Models

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

A large number of countries face severe population ageing in the near future. As a consequence of this, the cost of health care is expected to grow enormously the coming years. One way to keep these costs limited is to introduce technology that offers physical help, cognitive help or social help. For this, it is essential that the health state of the user is monitored continuously. Many monitoring systems focus on recognizing ADL: activities performed on a daily basis such as sleeping, toileting and cooking [10, 1, 11]. Networks of simple sensors such as motion detectors and door switches are becoming popular for measuring these activities. All presented research on the recognition of ADL from simple sensors assumes that only a single person is present in the home.

In our group we carry out long term experiments where we infer ADL from sensor network data. We monitor a number of elderly for more than a year, and study trends and correlations with ADL measurements from professional nurses. For that reason it is important to be sure that the data we are collecting indeed originates from a single person, and to discard data that is caused by multiple persons in the home. In this paper we present our system that is able to detect whether there is a single person in the home or multiple persons. We used a HMM with two states 'visits' and 'no visits'. We describe the data and annotation, the feature selection and the experiments on real data from a senior user.

Related work

Much research has been carried out on detecting, counting and tracking multiple people and monitoring their activities using video cameras [7]. However, there is much less research on the detection and tracking of multiple people with sensor networks, especially those with simple binary sensors.

A lot of research has focused on recognizing the ADLs of a single person. Different kinds of methods like Hidden Markov Models (HMM) [11], hierarchical hidden Markov model (HHMM) [6], Switching hidden semi-Markov models (SHSMM) [2], Rao-Blackwellised particle filter [9] have been used for modelling and inference of the ADLs of a single person.

In [13] the problem of simultaneous tracking and activity recognition (STAR) was introduced. The goal of STAR is to track multiple persons in a home setting and to recognize their activities. From an experimental setting where mainly motion sensors and contact switches are installed, it is shown that the use of Particle Filtering method has potential to solve the STAR problem. However, the accuracy of the tracking decreases as more people are in the house.

In [1] two HMM's (an activity model and a person model) have been used to detect a human social interaction in an office. Unfortunately a deeper insight into the methodology and accuracy is not given in the paper. In [4], [12] and [3] an emerging pattern based multi-user recognizer (epMAR) to recognize both single-user and multi-user activities is proposed. The sensor platform used for this experiment mainly consists of wearable sensors. The sensors are not only simple binary sensors, but also sophisticated sensors like acceleration, location, sound and voice sensors.

In [10] a video sensor network is used together with a wireless sensor network to distinguish between different people in a multi-person setting, e.g., the resident vs. a visitor. Unfortunately, the use of video camera's is not a possibility in our living labs due to privacy issues.

The focus of this paper is detecting the presence of one or more persons in an environment using a simple sensor network consisting of binary sensors.

Sensor data

In this section, we describe the way the sensor data has been collected, the tools used to visualise the sensor data and how the sensor data is annotated.

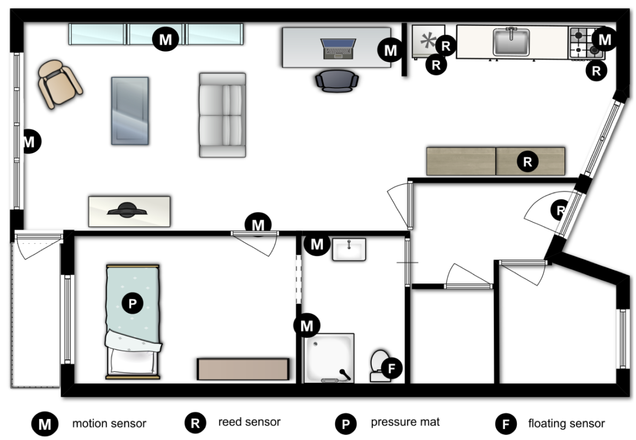

Figure 1: A map of the volunteer's apartments equipped with a wireless sensor network. The number of used sensors, their types and their position in the apartment do not differ a lot between the different apartments.

Sensor id	Sensor name	Sensor type	Room (number)
4dd	toilet	floating	bathroom (1)
3ac	front door	switch	hall (3)
d23	sink	motion	bathroom (1)
3d7	freezer	switch	kitchen (4)
460	microwave	switch	kitchen (4)
73e	fridge	switch	kitchen (4)
3b5	bedroom door	motion	bedroom (2)
717	shower	motion	bathroom (1)
3a6	bed	pressure	bedroom (2)
d12	stove	motion	kitchen (4)
d14	couch	motion	livingroom (5)
d0d	livingroom	motion	livingroom (5)
d22	desk	motion	livingroom (5)

Table 1: A list of the sensors (id, name, type and room) installed in the apartment of resident 1. as shown in Figure 1.

Data collection

Four apartments in the assisted living department of a care centre in the Netherlands were equipped with a wireless sensor network during a period of more than a year. The wireless sensor network consist of binary sensors described in [11]. The elderly are living their routine life and not told to perform a specific ADL in a specific way at some specific time. An overview of the location of the sensors in the apartment of the resident 1 is shown in Figure 1. The location of the sensors is chosen so that the most important rooms in the apartment are covered and in a way the network does not affect the elderly daily life. For instance, the pressure sensor for the bed is installed under the mattress, motion sensors in the living-room are installed under the TV-cabinet and in the bookshelf and switch sensors in the kitchen are installed above the stove, under the freezer, etc. Around 15 sensors are installed in each apartment. From these sensors there are at least two motion sensors in the living room, three motion sensors in the bathroom and at least 5 sensors are installed in the kitchen.

Sensor events, signals fired by sensors, are stored on a local computer at the resident's apartment. The sensor data computer can be remotely accessed, so that the sensor data is available every moment for analysis. Remote access has also the advantage of monitoring the (mal)functioning of the sensor network. For example, during the first three days of July 2012, the microwave sensor in apartment 1 fired a lot of consecutive events in a short time even during the night. These (mal)function events are not taken into account in the data analysis. Figure 2 illustrates all different types of events.

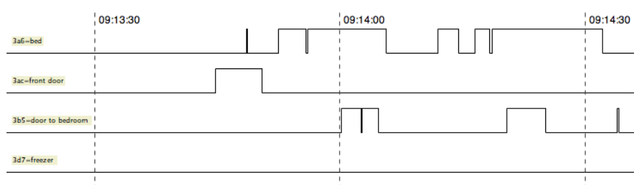

Figure 2: A sample of the sensor data generated by the binary sensors (bed, front door, door to bedroom and freezer). The front door sensor has fired an 'OPEN' event at time-stamp 09:13:44 and a 'CLOSE' event at time-stamp 09:13:50. The bed sensor has fired 6 'YES' and 6 'NO' events between 09:13:48 (first 'YES' event) and and 09:14:32 (last 'NO' event). The door to bedroom sensor has fired 4 'ON' and 4 'OFF' events between 09:14:00 (first 'ON' event) and 09:14:34 (last 'OFF' event). The freezer sensor did not fire any event.

Annotation of data

In our previous work [8], a video camera is used to record the activities of an office user and the visits to his office. Annotation of the sensor data is then done in an accurate way by viewing the video recordings. In our living labs located in an elderly care centre, there is no possibility to use video camera's to record the ADLs due to privacy issues. The volunteers are asked to register some information about the visits they received during the first two weeks of July 2012. A special form was designed to make the registration easy for them. Some elderly do find this a difficult job due to their physical disability or just forgot to fill in the form. During these two weeks, the researcher occasionally phoned the elderly to make sure the registration process is going well.

After the period of annotation, an interview with the elderly took place where the sensor data is compared to the annotation the elderly has made. Unusual patterns found in the sensor data were clarified by the elderly during this interview. For example the bed sensor of resident 4 produced daily events every morning after a sleep. This pattern corresponds to the daily activity 'putting stockings'.

The start and end time of the visits registered by the elderly were approximate times. For this reason the exact time generated by the front- and back-door sensors are used for the annotation in stead of the times filled in by the elderly. An overview of the types and the duration of the visits of the resident 1 is given in Figure 3. This resident receives a visit of a nurse twice a day. The visits in the morning last approximately 10 minutes to help the resident out of his bed. In case the resident also needs to be showered, these visits last approximately 30 minutes. The visits in the evening last just few minutes to help the resident wear the compression stockings. The visits of the cleaner are every Thursday between 13h and 15h. Private visits of family members and friends are incidental. The visits received for the sensor network maintenance are filtered out as all the sensors are triggered various times by the maintenance team to check their functionality.

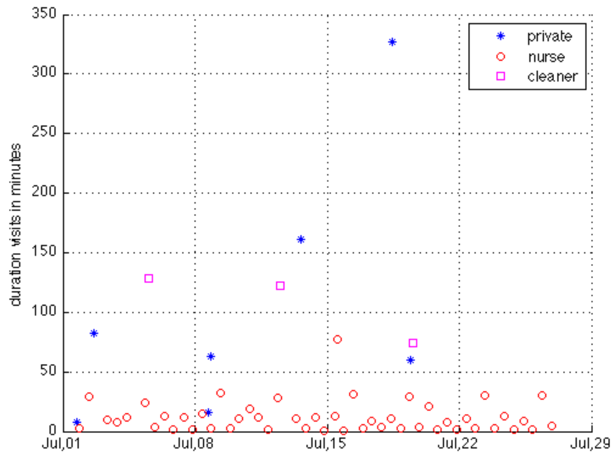

Figure 3: The type versus the duration of the visits to resident 1 during the month of July 2012.

Feature Selection

The residents of our living labs receive in general three types of visits: daily visits of a health care professional like a nurse, weekly visits of a cleaner and visits of family members, neighbours or friends. We believe that when a resident has a visit the order in which different sensors are activated will be different than when a resident is alone at home. For example, an event fired by the front door sensor followed by an event of the bed sensor within two seconds is most unlikely to be generated by the resident alone. We also believe that the density of some sensors events in some rooms like the living-room will be different when the resident has a visit. Therefore, we chose the room events transition τ_{ij} defined in equations 1 and 2 as features for visits.

To this end we define:

- The set $S = \{s_1, ..., s_{|S|}\}$ as the set of sensors installed in the apartment of the elderly, where $|S|$ is the total number of sensors.
- The set $R = \{r_1, ..., r_{|R|}\}$ as the set of rooms monitored in the apartment of the elderly, where $|R|$ is the total number of rooms.
- The set $P = \{t_1, \ldots, t_{|P|}\}$ as the set of event's time stamps.
- The observation $\boldsymbol{o}^{(t)} = (o_1^{(t)}, ..., o_{|S|}^{(t)})$ as the vector of sensor events at time stamp $t \in P$ generated by the $|S|$ sensors where

$$o_i^{(t)} = \begin{cases} 1 & \text{if } s_i \text{ fires an ON/OPEN/YES-event at } t \\ -1 & \text{if } s_i \text{ fires an OFF/CLOSE/NO-event at } t \\ 0 & \text{otherwise} \end{cases}$$

  Note that as only sensor event time stamps are considered ($t \in P$), there will be no null observation vector $\boldsymbol{o}^{(t)} = (0, 0, \ldots, 0)$. For notational simplicity, the superscript (t) will be discarded in some equations.

- The event of sensor s_i defined by $e_i^{(t)} := |o_i^{(t)}|$. $e_i^{(t)} = 1$ if sensor s_i fires an event at time stamp t, $e_i^{(t)} = 0$ otherwise.
- The room event of room j defined by $f_j^{(t)} := \vee_{s_i \in r_j} e_i^{(t)}$. $f_j^{(t)} = 1$ if at least one sensor in room j fires an event at time-stamp t. $f_j^{(t)} = 0$ otherwise.

For the detection of multiple persons we used the feature vector $\boldsymbol{\tau}^{(t)} = (\tau_{11}^{(t)}, \tau_{12}^{(t)}, \cdots, \tau_{|R||R|}^{(t)})$, where $\tau_{ij}^{(t)}$ is a room event transition from room i to room j at time stamp $t = t_n$ defined by:

$$\tau_{ij}^{(t)} = f_i^{(t_n)} \cdot f_j^{(t_{n+1})} \quad \text{if } i \neq j \tag{1}$$

and

$$\tau_{jj}^{(t)} = \vee_{k,l \in r_j, k \neq l} e_k^{(t_n)} \cdot e_l^{(t_{n+1})} \quad \forall j \in R \tag{2}$$

This means that the transition $\tau_{ij}^{(t)} = 1$ only if two consecutive room events occured in different rooms i and j. The self transitions $\tau_{jj}^{(t)} = 1$ only if different sensors of room r_j generate two consecutive events. Note that as the hall consists of only one sensor, the hall self transition is always equal to 0.

The feature vector $\boldsymbol{\tau}^{(t)}$ or a subset of it is used to construct the emission matrix B of the HMM described in the next section.

Approach/classifier model

The problem of detecting visits can be tackled using HMM. The observation vector at time t denoted by x_t is represented by the feature vector $\boldsymbol{\tau}^{(t)}$ or a subset it as described in the previous section. The hidden state at time t denoted by z_t is represented by the visits to the resident. $z_t = 1$ if the number of persons in the apartment is exactly one i.e. there is no visit. $z_t = 2$ if the number of persons in the apartment is two or more i.e. there is a visit. The HMM is mathematically represented by equation 3.

$$\lambda = (A, B, \pi) \tag{3}$$

where $A = \{a_{ij} : 1 \leq i, j \leq N\}$ defines the transition probability matrix, $B = \{b_j(k) : 1 \leq j \leq N, 1 \leq k \leq M\}$ defines the emission probability matrix and $\pi = \{\pi_i : 1 \leq i \leq N\}$ defines the initial state probability. a_{ij} is the transition probability of taking the transition from state s_i to state s_j. $b_j(k)$ is the conditional probability of emitting the $k-th$ observation given the state s_j.

The parameters (A, B, π) of the HMM used in equation 3 are estimated using the maximum likelihood estimates.

Given a sequence of features $x_{1:T}$ with length T, we want to predict whether the resident has a visit during this period of time. For this reason we calculate for each time stamp t $(1 \leq t \leq T)$ the posterior state probability. This posterior is, for the visit's state, defined as the conditional probability of being at state s_2 at time stamp t, given the observed sequence of features $x_{1:T}$. Calculating this posterior probability for every t $(1 \leq t \leq T)$, will result in a vector of posterior state probabilities denoted by $p(z_t = s_2 | x_{1:T})$. Counting the number of posteriors with values bigger than 0.5, the probability of a visit is then equal to number of these counts divided by the length of the posteriors vector, which is equal to T. The general posterior state probability denoted by $p(z_t | x_{1:T})$ is calculated using the Forward and Backward procedures.

For the implementation of the estimation of the parameters and the inference we used the HMM functions `hmmestimate` and `hmmdecode` of the numerical computing environment Matlab.

Experiments

In this section we describe the objective, the setup and the results of the conducted experiments.

Objective
We have conducted two experiments. In the first experiment we determined the best set of features using the two weeks of annotated data. In the second experiment we determined the performance of the HMM model with the best features on an independent data set. This data set, referred as the second data set, comprises the last two weeks of July 2012. The performance is calculated using the sensor data of only resident 1 because of lack of annotated test data for that period in the other apartments

Experiments Setup
Since we have five rooms, the feature vector is a 25-dimensional vector. This dimension is too high for the training of all possible combinations of the subsets of $\boldsymbol{\tau}^{(t)}$. In order to find the best subset of features we therefore followed a stochastic procedure. Randomly a subset (between 1-25) of features was drawn and the HMM was trained. The followed steps were:

1. Construct the feature matrix X by extracting the room event transition vectors $\boldsymbol{\tau}^{(t)}$ at each time stamp t from the first data set.
2. Select a random subset of $\boldsymbol{\tau}^{(t)}$ with a random dimension h and construct the feature matrix X_h with this subset.
3. Apply a 3-fold cross validation using HMM on X_h and calculate the average accuracy Acc_h, defined in Equation 4, over the 3 rounds.
4. Repeat the steps 2 and 3 at least 50000 times.
5. Calculate the biggest Acc_h and select the corresponding subset of $\boldsymbol{\tau}^{(t)}$.

In the second experiment we used the best feature τ_{best} found in the first experiment and used the first data set to construct the best HMM. Using this HMM on the second data set, the accuracy is calculated to determine the performance of the model.

Classifier performance
In the conducted experiments the percentage of the generated sequences belonging to the 'visit' class may be too small compared to the percentage of the 'no visits' class. To this end we used the Geometric Mean [5] given in equation 4 which maximizes the accuracy of each class while keeping these accuracies balanced.

$$Acc_{gm} \quad = \quad \sqrt{\frac{TP}{TP+FN} \times \frac{TN}{TN+FP}} \qquad (4)$$

Note that $Acc^{+} := TP/(TP+FN)$ is the same as the recall of the positive class and $Acc^{-} := TN/(TN+FP)$ the recall of the negative class. Acc^{+} is also called the true positive rate and Acc^{-} is called the true negative rate.

Results
The total number of sensor events collected in the apartment of resident 1 during the first two weeks of July 2012 is 80862. From this amount of data we distilled 61732 room transition vectors $\boldsymbol{\tau}^{(t)}$, where a percentage of 11% belongs to the 'visits' class and a percentage of 89% belongs to the 'no visit' class.
The length of the feature sequence $x_{1:T}$ is set to $T = 10$ because a sequence with length 10 has a duration which is approximately equal to the duration of a short visit.
Running a 3-fold cross validation (step 3 of the first experiment) takes approximately three CPU-second on a

regular linux virtual machine (2,5GHz single core and 4Gb RAM). This means that the running of a 3-fold cross validation of all the possible subsets of X_h will take $\sum_{h=1}^{25} \binom{25}{h}$ CPU-seconds which is equal to about 97 months.
Conducting the first experiment during few days showed that the obtained maximum average Acc_{gm} is comparable for all the HMMs λ_h except λ_1 and λ_{25} which have a much smaller Acc_{gm}. Also, the features with a dimension $h = 12$ often have high accuracies. The top three of the best features with dimension $h = 12$ are given in Table 2.

The total number of sensor events collected in the apartment of resident 1 during the last two weeks of July 2012 is 60322 which is comparable with the size of the first data set. The best features found in the first experiment are used to train HMMs which are applied to the second data set. Calculating the geometric mean of these HMMs resulted in a performance of 85%. The visits posterior probability for three days is shown in Figure 4. These figures show that all the types of visits are detected, but there are also some false positives.
The frequently occurred transitions in the best features justify our presumptions. For example the transitions τ_{15} and τ_{51} explain the presence of the cleaner in the bathroom while the resident is in the living room. The self transition in the living room (τ_{55}) explains the presence of a (private) visit in the living room.

rank	best feature combinations
1	$(\tau_{11}, \tau_{13}, \tau_{15}, \tau_{22}, \tau_{24}, \tau_{32}, \tau_{35}, \tau_{41}, \tau_{44}, \tau_{52}, \tau_{53}, \tau_{55})$
2	$(\tau_{11}, \tau_{13}, \tau_{22}, \tau_{23}, \tau_{33}, \tau_{34}, \tau_{41}, \tau_{44}, \tau_{51}, \tau_{52}, \tau_{53}, \tau_{55})$
3	$(\tau_{12}, \tau_{15}, \tau_{22}, \tau_{23}, \tau_{33}, \tau_{34}, \tau_{35}, \tau_{44}, \tau_{51}, \tau_{52}, \tau_{53}, \tau_{55})$

Table 2: Best feature combinations with dimension $h = 12$

Conclusions and Future Work
Our research group is interested in recognizing ADLs of elderly people using data from a simple wireless sensor network. It is important to know that the collected data is originated from a single person, the resident. In our pervious work [8], we studied the problem of detecting the presence of multiple persons in an environment. The experimental setting used is limited to one room, the office of the supervisor, and the number of sensors is small. Here, the experimental setting is more realistic. There are more than one monitored rooms in the apartment and the number of the sensors is high. Also the type of visits is diverge. Using HMM with room event transition as features, we achieved a reasonable high accuracy of visits detection.
At the moment of writing, we are in the implementation phase of a Markov Modulated Poisson Process (MMPP) as an unsupervised approach to detect visits. It is interesting to compare the MMPP method with the supervised HMM method used in this paper.
The sensor network used does not monitor all the rooms in the apartment. There are even some unmonitored areas in the monitored rooms like the kitchen. A visitor in these unmonitored areas may not be detected. We need to monitor every area in the apartment to be able to detect visits and recognise ADLs. A new wireless sensor network with more sensors leading to a bigger monitoring range is installed in our living labs since April, 2013. Using this new wireless sensor network we hope to achieve better results in the future.

Acknowledgements
This work is part of the research programs Health-lab, COMMIT and Smart Systems for Smart Services that has been supported by SIA. The authors would like to thank the participants at Vivium Zorggroep Naarderheem,

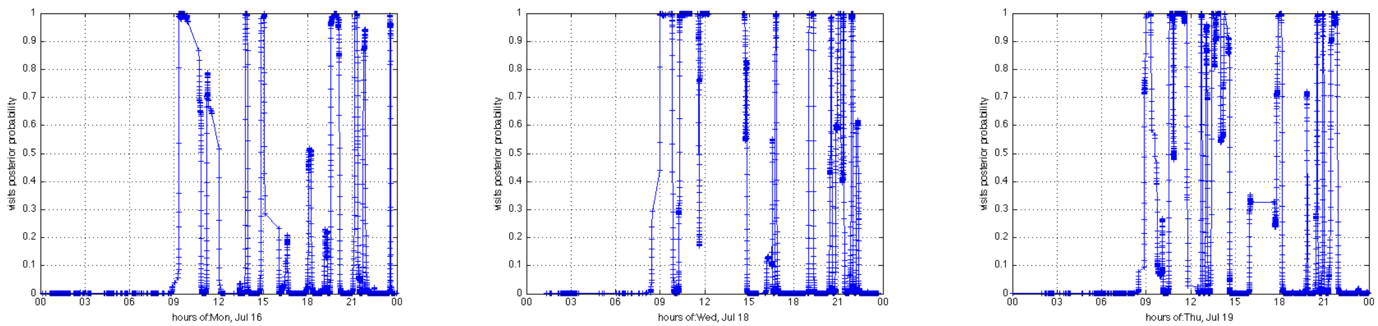

Figure 4: The visits posterior probability for three typical days in de second data set. (a) a day with the daily two visits of the nurse in the morning at 09:20 and in the evening at 21:10. (b) a day with a private visit between 1:10 and 15:37. (c) a day with a private visit between 10:30 and 11:30 and the visit of the cleaner between 13:00 and 14:15.

References

[1] Cook, D. J., Crandall, A., Singla, G., and Thomas, B. Detection of social interaction in smart spaces. *Cybernetics and Systems 41*, 2 (2011/10/26 2010), 90–104.

[2] Duong, T., Bui, H., Phung, D., and Venkatesh, S. Activity recognition and abnormality detection with the switching hidden semi-markov model. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 1 (june 2005), 838 – 845 vol. 1.

[3] Gu, T., Wu, Z., Tao, X., Pung, H. K., and Lu, J. epsicar: An emerging patterns based approach to sequential, interleaved and concurrent activity recognition. In *Pervasive Computing and Communications, 2009. PerCom 2009. IEEE International Conference on* (march 2009), 1 –9.

[4] Gu, T., Wu, Z., Wang, L., Tao, X., and Lu, J. Mining emerging patterns for recognizing activities of multiple users in pervasive computing. In *Mobile and Ubiquitous Systems: Networking Services, MobiQuitous, 2009. MobiQuitous '09. 6th Annual International* (july 2009), 1 –2.

[5] Kubat, M., Holte, R., and Matwin, S. Learning when negative examples abound. In *Machine Learning: ECML-97*, M. van Someren and G. Widmer, Eds., vol. 1224 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 1997, 146–153. note.

[6] Luhr, S., Bui, H., Venkatesh, S., and West, G. Recognition of human activity through hierarchical stochastic learning. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on* (march 2003), 416 – 422.

[7] Moeslund, T. B., Hilton, A., and Krüger, V. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding 104*, 2 (2006), 90–126.

[8] Nait Aicha, A., Englebienne, G., and Kröse, B. How busy is my supervisor. *PETRAE 2012: submitted to the 5th international conference on Pervasize Technologies Related to Assistive Environments* (2012), 1–8.

[9] Nguyen, N., Phung, D., Venkatesh, S., and Bui, H. Learning and detecting activities from movement trajectories using the hierarchical hidden markov model. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 2 (june 2005), 955 – 960 vol. 2.

[10] Skubic, M., Alexander, G., Popescu, M., Rantz, M., and Keller, J. A smart home application to eldercare: Current status and lessons learned. *Technology And Health Care 17*, 3 (2009), 183 – 201.

[11] van Kasteren, T., Noulas, A., Englebienne, G., and Kröse, B. Accurate activity recognition in a home setting. In *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing*, ACM (New York, NY, USA, 2008), 1–9.

[12] Wang, L., Gu, T., Tao, X., and Lu, J. Sensor-based human activity recognition in a multi-user scenario. In *Ambient Intelligence*, M. Tscheligi, B. de Ruyter, P. Markopoulus, R. Wichert, T. Mirlacher, A. Meschterjakov, and W. Reitberger, Eds., vol. 5859 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2009, 78–87.

[13] Wilson, D., and Atkeson, C. Simultaneous tracking and activity recognition (star) using many anonymous, binary sensors. In *Pervasive Computing*, H. Gellersen, R. Want, and A. Schmidt, Eds., vol. 3468 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2005, 329–334.