Multi-Target Data Association in Binary Sensor Networks

Sebastian Matthias Müller
OFFIS – Institute for Information Technology
Oldenburg, Germany
Email: sebastian.mueller@offis.de

Andreas Hein
School of Medicine and Health Sciences
University of Oldenburg
Oldenburg, Germany
Email: andreas.hein@uni-oldenburg.de

***Abstract*—Numerous applications of ambulant medical care, house automation and security use binary sensors such as passive infrared motion sensors or light barriers to monitor activity in the house. Multi-target tracking algorithms allow for at least a partial separation of activity in data from such sensors from multiple persons. While many tracking algorithms demonstrate good performance across various sensing modalities and sensor setups, little research has been done to determine the impact of placement and varying density of sensors for tracking performance. This paper presents the results of an evaluation of a Bayesian multi-hypothesis multi-target tracking algorithm on data of two residents monitored by a network of binary sensors. We evaluate the algorithm on data from sensors of varying quantity and placement. We show that our approach outperforms other approaches in low-resolution setups. While tracking performance naturally decreases with the number of sensors, it also strongly varies by sensor positioning.**

Keywords–Multi-target tracking; Assisted living; Wireless sensor networks.

I. INTRODUCTION

The emergence of research on technical support systems for ambulant care and support for patients and elderly stem from numerous recent societal developments as well as changes in demographic structure.

First, the coincidence of prolonged life expectancy [1] and the atomization of households [2] puts an increasing care demand into the hands of third parties. According to the German Federal Statistical Office, the number of single-households will increase sixfold in relation to the population numbers. At the same time, the ratio between care personnel supply and demand will cut in half [3].

Second, the increasing life expectancy, in combination with improved medical care and "modern lifestyles and behavior" [4] causes an increase in the proportion of population living the chronic diseases, thus further driving demand for ambulant care.

Third, there is a general trend towards outpatient care by hospitals. According to the Avalere Health analysis of American Hospital Association Annual Survey [5], the percentage of revenue for community hospitals in the United States has increased from 25% to 44% between 1992 and 2012.

These developments drive the research on technical support systems in home and care environments. Applications for such include automated assessments [6], activity monitoring [7] or fall detection [8]. To preserve a maximum of privacy and comfort while at the same time collecting data necessary for the application, many approaches include the use of ambient sensors such as motion sensors and light barriers. Since the data collected from these sensors does not carry identifying information, use of any such application in settings where more than one person – the patient – moves or resides becomes difficult.

Complex sensors, such as cameras and microphones are rarely accepted in living spaces. Body-worn sensors are often forgotten or ignored due to discomfort. Binary sensors such as light barriers and motion sensors are easy to retrofit, have relatively little power consumption and can be installed unobtrusively. A no-requirements sensor model also enables us to install more complex sensors (such as laser scanners or depth-finding cameras) as required. The necessary information can be extracted from their data by partitioning the sensors' range and converting activity in each partition to a binary signal.

To separate data from multiple persons moving in a space monitored by binary sensors, we present a multi-target tracking algorithm using Bayesian estimation and multi-hypothesis tracking. This algorithm makes no assumptions on the selection and placement of sensors or sensing technology. Tracking takes place on a graph of the sensors and their spatial relation. It is thus not helpful in determining the precise location of a present person, but at (or below) room-level accuracy. This algorithm performs particularly well on low-resolution data, such as when only few binary sensors are used. We test the algorithm across various sets of sensors, varying by placement and number. A decreasing number of sensors will likely have an impact on the tracking accuracy, but is important in regard to energy consumption, costs and user acceptance. We show that data from two residents in an apartment can be separated with high (>90%) accuracy, and that the selection and placement of sensors can play a significant role in tracking accuracy.

The remainder of this article is structured as follows: Section II summarizes related works on multi-target tracking and activity monitoring in the home using binary sensors. Section III describes the theoretical principles surrounding data association and multi-hypothesis tracking for single- and multi-target tracking. Section IV explains how the approach was evaluated, including data preparation, the evaluation function

 ISBN: 978-1-61208-421-3

as well as the sensor placement concept. The results of the evaluation are presented in section V. Section VI concludes the article.

II. Related Work

Prior work has shown that data collected from sensor networks allow for the deduction of information used in activity monitoring, care assessments and behavior modeling. Target tracking, in particular multi-target tracking, is a task often applied to visual data such as video feeds and images. The practical application of multi-target tracking in binary or low-resolution home sensor networks has been to little research.

A. Target Tracking in Home Sensor Networks

Wilson and Atkeson [9] describe an algorithm for tracking of multiple persons and their activity status in a binary sensor network. In this work, the authors use a transition matrix representing transition probabilities between sensors. By keeping track of the targets' identities, personal motion models emerge. The data association is achieved using a particle filter. During a five-day experiment in a house instrumented with 49 sensors (contact switches, motion sensors), data during two-person scenarios was correctly assigned 82.1% of the time.

Krüger et al. [10] use a particle filter and *action plans* to assign sensor events from motion sensors and light switches to tracks and simultaneously identify the target. Action plans describe action sequences in terms of sensor data. These plans can be synthesized or learned from historic data. For the evaluation, an office corridor was equipped with six light switches and six motion sensors. The mean squared error across time and all targets is reported as approximately 0.26 for two-person scenarios. The work shows how – similar to trained motion models – previous knowledge of a person's plans can help tracking individuals in binary sensor networks.

Oh and Sastry [11] perform tracking on data of binary sensor networks and passage connectivity graphs. The graphs are calculated from transition probability matrices. A tracking algorithm, derived from the Viterbi algorithm, pruning strategies and multiple target tracking extensions are presented. No evaluation on real world data is conducted.

Marinakis et al. [12] derive the topology of a sensor network in terms of transition times and probabilities from data of unspecified sensors. The authors use Monte Carlo Expectation Maximization to assign activity to agents (people present) in order to build a graph of the sensor network. 95% of the topology of simulated node graphs is recovered correctly. The results for a trial using a network of cameras and photocell-based sensors are not reported.

B. Activity Monitoring in Home Sensor Networks

Numerous studies show that data collected from sensor networks in living spaces allow for the deduction of information relevant in applications of activity monitoring, care assessments and behavior modeling.

Logan et al. showed that ambient motion-based sensors provide the most useful information for detection and classification of daily in-home activities in a study compared to RFID, on-body and on-object sensors. In their study, infrared motion sensors yielded the best results overall, although classification performance on this data was better on activities that are strongly correlated with locations in the home, such as "watching TV" and "meal preparation" [13].

Data from binary sensors can also be used to calculate average room residence time and frequency: Assessment tests are partly realizable by using recordings from light barriers and reed contacts alone [14]. The authors argue that light barriers alone do not constitute sufficient evidence of a person entering a room, because people may change directions between rooms. It is suggested to combine light barriers with sensors covering larger areas. Room residence times are calculated by manually labeling the sensors constituting a room using a floor plan and knowledge of the sensors placements. In a similar study, the authors model user behavior of a resident from the probability of location at a certain time of day and the frequency of presence in a location in a defined period of time [15]. Models are created for rooms individually (bathroom, bedroom, living room, kitchen). Based on the number of anomalous behavior detected, the authors conclude that the models' performance varies by room: Presence in the bathroom is best modeled duration-based, while the timeslot-based model yielded better results for the other rooms.

Frenken et al. [6] use ambient sensors in an attempt to automate measurement of mobility and gait velocity, as required in the Timed Up and Go assessment [16]. For this, five flats are equipped with home automation sensors and one with an additional laser range scanner. It is shown that the data is suitable to compute gait velocity at home. While data from the laser range scanner is proven to be more precise than home automation sensor data, no statistical post-processing or filtering was performed on the latter.

III. Approach

We define a sensor graph of sensors $s_1, \cdots, s_N$ as a weighted, directed graph $G = (V, L)$, where $V = \{1, \cdots, N\}$ is the set of nodes in the graph representing the sensors, and L is the set of all edges (u, v) for which there is a direct passage from the sensing region of sensor u to the sensing region of v which does not intersect any other sensing regions. Informally, two sensors u, v are connected if it is possible for a person to traverse from the sensing region of u to the sensing region of v without activating any other sensor.

Each resident in the target space is represented by a discrete Bayesian filter on an unweighted, undirected graph consisting of sensors as nodes and edges representing their spatial adjacency. For our evaluation data, this graph was published by Crandall et al. [17] (Figure 1). If the adjacency relations are not known, they can be approximated by a path planning algorithm [18] using a floor plan, if available, or generated from historic data [19].

A. Tracking of individuals

Bayesian filters estimate the state of a dynamic system from noisy data. We choose a probability distribution to represent the location of each individual, because it helps estimating a more precise location later on, especially when sensor regions overlap. More importantly still, it helps the tracker to recover more quickly when a noisy measurement is assigned to the individual's track. Lastly, we aim to replace the manually constructed, unweighted graph with a weighted graph that is automatically constructed from in-situ recorded data and transition probabilities between sensors as weights (cf. [12]).

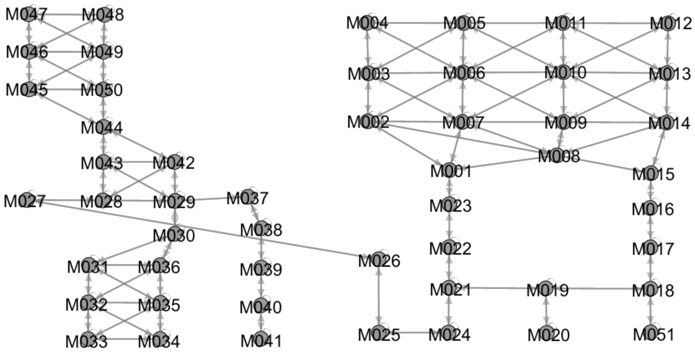

Figure 1. Graph of sensors (with their internal IDs and their spatial relations used in the evaluation (adapted from [17]).

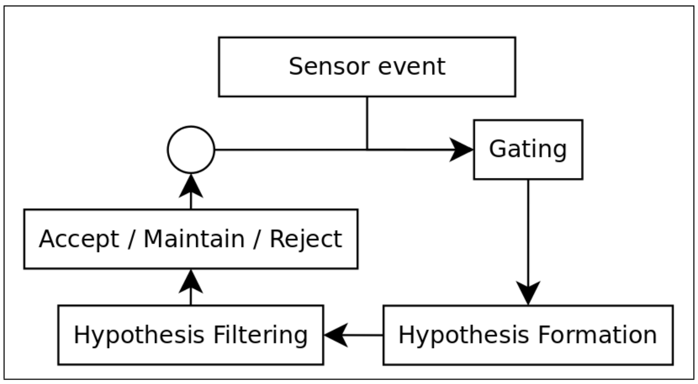

Figure 2. Hypothesis formation overview.

All sensors are subject to noise, but many motion and home automation sensors come with an additional source for noise: measurement delay. Many sensors do not measure or report measurements for a specified amount of time after triggering. This period can last from a few seconds to several minutes due to radio communication regulations. This results in sensors missing the presence or movement, thus breaking the continuity of measurements of a motion track.

At which point and how many filters are created – that is to say, how many individuals are assumed to be present – depends on the performance (*belief*) of the previously existing track: When new measurements cause the current data to be more likely when assigned to more or fewer tracks (= individuals) than before, a new filter is spawned or an existing filter is discarded. The data could be bundled into larger updates within reasonable time frames (cf. [20]), but in our case an update occurs for each new sensor event.

B. Multi-target tracking

When new sensor data arrives, hypotheses are created by considering all possible assignments of the data to existing and new tracks (”hypotheses”) until the filter’s *window size* is reached. This is particularly useful in a low-resolution setting like ours, where individuals may occlude each other in sensor readings for any period of time.

The window size in multi-hypothesis tracking (MHT) describes the maximum number of events (or time steps) that are considered before choosing a likely hypothesis. Windowing is necessary to limit the number of possible hypotheses and to limit the information loss in case no acceptable hypothesis remains and the data is discarded. The influence of the window size on tracking accuracy has been shown previously [21]. For our evaluation, we use a window size of 10 events.

The idea of multi-hypothesis tracking dates back to 1979, when Donald B. Reid published ”An algorithm for tracking multiple targets” [22]. Reid’s algorithm was developed to work on data from a continuous scale sensor (e.g., radar). Therefore, Reid speaks of associating measurements to clusters. In the work presented here, the target space is discrete (nodes on a graph), and targets and their locations are stored as a probability distribution over the space using Bayesian filters.

There are several significant differences between Reid’s original work and the approach described here. In accordance with Reid’s *type 2* sensor, our sensor model expects *positive reports* only, meaning that we consider only sensor data reporting activity. However, tracks are updated per hypothesis, rather than generated and filtered individually (*hypothesis-oriented MHT*). This means that hypotheses are not constructed from *compatible* tracks, but all possible combinations of updates of existing hypotheses. Furthermore, the tracker is updated every time a sensor reports activity. Because of this, and the fact that our state space is discrete, computational complexity is reduced. For a more detailed description of track- and hypothesis-oriented MHT, see Blackman [20].

For each triggered sensor, a new hypothesis based on all previously existing hypotheses is created, in which the triggered sensor is

- considered noise and discarded,
- used to update one of the existing filters, or
- assigned to a new filter.

Due to the exponential growth of the number of possible hypotheses ($> 4.74 \times 10^{13}$ for 20 events), we must employ a number of filters to optimize computation efficiency.

All hypotheses must pass a gating function before they are considered for evaluation (see Figure 2). In our case, this gating function is a simple comparison of the prior probability of each filter to a threshold value. Afterwards, hypotheses are filtered based on confidence, noise ratio and similarity. This procedure is performed until a single hypothesis remains or the window size is reached. In the former case, the hypothesis is accepted, the underlying Bayesian filters updated, and the window size reset. In the latter case, all hypotheses are evaluated. If no single, dominating hypothesis can be found, all hypotheses are discarded and the underlying filters reset.

The size of the window strongly influences the performance of the algorithm. A larger window size will result in a larger number of correct associations, but also in a larger number of discarded sensor events [21].

Figure 2 depicts the general multi-hypothesis tracking logic. For a more in-depth description of multi-hypothesis tracking, see Blackman [20] or Reid [22].

IV. EVALUATION

A. Data Preparation

The data used for this evaluation was recorded at the Center for Advanced Studies in Adaptive Systems (CASAS) at the University of Washington [23]. It shows activity of two residents of a smart home environment, residing in a 4-room, 2-story apartment for approximately 8 months. For our evaluation, we use subsets of the data recorded by the 50 motion sensors mounted to the ceiling. The smart home is also equipped with contact sensors on doors and cabinets, temperature, water and electricity sensors. For our purposes, however, motion sensors offer the most precise and least noisy data.

We use data for which at least both residents are present and active. We choose time frames

- that last at least 20 minutes or contain at least 300 sensor events,
- in which both residents change rooms at least once, and
- in which neither resident is inactive for more than 20% of the time.

The result are twenty time frames, with 330 to 910 sensor events with durations between 24 and 530 minutes. After selection, each of the 13321 sensor events was labelled as originating from Resident 1, Resident 2 or a third person using the manually labelled events and the laboratory's floor plan.

B. Data Association

The algorithm can track any number of targets. However, our intended area of application – small households – allows us to use an evaluation function that is tailored towards few targets (1-3). For this evaluation, the algorithm was optimized to track two targets by using an evaluation function that favors one- and two-track hypotheses. Equation (1) describes the evaluation function, where h is the hypothesis in question, $conf(p_n)$ is the belief of the Bayesian filter at the most recent event location n, $\|p\|$ is the number of paths (= targets) in h, and m is the expected number of targets in the sensor space.

$$eval(h) = \frac{\sum_{i=1}^{n} conf(p_n)}{\frac{\|p\|^2+m}{m+1}} \quad (1)$$

C. Sensor Placement

To get a better understanding of how the number of sensors affects tracking accuracy, we also run the algorithm on subsets of the original set of sensors in decreasing size (40, 30 and 20 sensors). Instead of choosing the sensors randomly, we chose characteristics of sensors we deemed possibly influential on tracking performance:

1) Number of neighboring sensors: Based on the assumption that sensors in doorways, which usually have few neighboring sensors, are critical in tracking room transitions, we remove those in larger areas with many neighboring sensors. The number of neighboring sensors can be calculated from the sensor graph.

2) Duration of stay: Given that tracking stationary targets is much simpler than moving targets, we consider subsets of sensors that cover areas in which the average duration of stay is short. The duration of stay can be calculated from the duration between consecutive sensor events in recorded data.

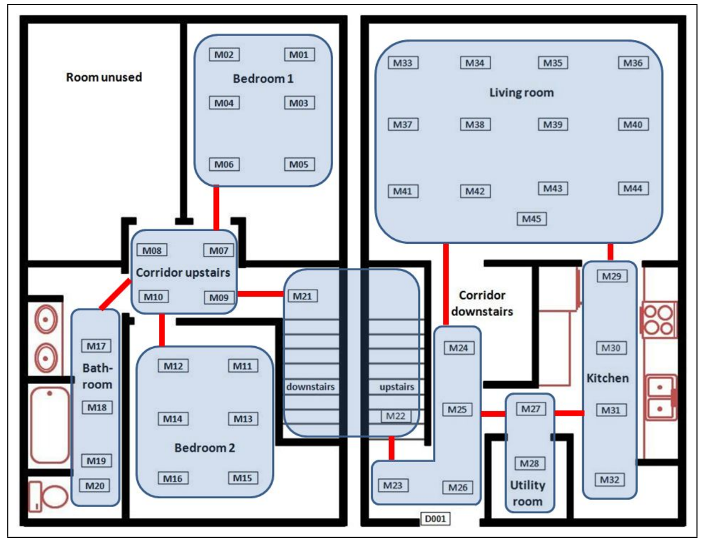

Figure 3. Sensors clustered to represent rooms.

3) Activity: Considering the application of in-home activity monitoring, it is imperative that the placement of sensors for tracking accuracy improvement does not interfere with the necessity of covering those areas in which the majority of activity is taking place. Thus, we select and filter sensors based on the amount of activity covered. The amount of activity covered by a sensor is simply calculated by the number of times it is triggered.

These criteria were used to create subsets of data of varying size, selected by increasing, as well as decreasing order of the respective criterium (cf. Figure 4).

D. Sensor Clustering

The procedure of selecting subsets of sensors for tracking performance evaluation was also conducted for sets of 10 sensors. However, due to the selection criteria, most of the sets had removed whole rooms, and in one case all data from one individual. Thus, in order to evaluate tracking performance on 10 sensors, we cluster the sensors by rooms and spatial adjacency (see Figure 3), and treat the resulting clusters as individual sensors. This also results in a more realistic scenario, in which motion sensors often cover different size areas up to whole rooms.

For this evaluation, we use data from all sensors, but we replace the sensor IDs with IDs for their corresponding cluster. This way, we make use of all sensor events but decrease their spatial resolution.

V. RESULTS

Tracking accuracy using all sensors is 90.3%. This is the percentage of the 13321 sensor events across all time frames that are correctly associated to any of the targets. The accuracy of individual time frames ranges from 62.1% to 99.5%, with a median of 93.1%. The error is composed of false associations (events that are falsely associated to another target, median 5.88%), no associations (events that could not be associated with any target, 0.59%) and noise (events that are falsely discarded as noise, 0.44%).

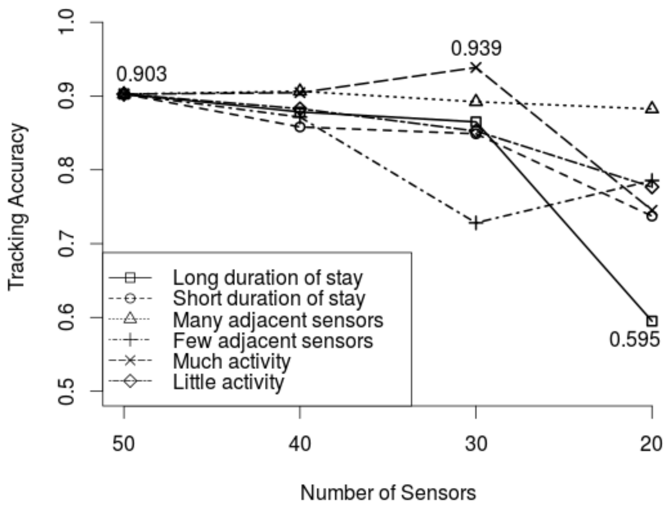

Figure 4. Tracking performance across sensor subsets by size.

Figure 4 shows tracking accuracy across sensor subsets. The sets vary by the number of included sensors (x-axis) and their selection criterion. As can be seen, tracking accuracy generally decreases with reduced sensor count. This is to be expected as the resolution of the tracking space decreases and situations with overlapping motion increases. Down to 30 sensors, tracking accuracy decreases only slightly for all but one sensor set. For the set of sensors with much overall activity, accuracy even increases slightly. The graph also shows that performance variation increases with the number of sensors. While tracking accuracy varies between 85.8% and 90.6% with 40 sensors, with 20 sensors accuracy ranges between 59.5% and 88.2%.

Tracking accuracy on the clustered data set is 77.2%.

VI. Conclusion

The article at hand describes an algorithm for tracking of multiple targets in a space monitored by binary sensors. It enables the separation of sensor data generated by multiple persons in smart home environments without the need for identifying sensors. The algorithm makes use of a graph consisting of sensors as nodes and their spatial relations as vertices. Compared to other related works, the algorithm works particularly well in low-resolution settings (i.e., with few binary sensors). It was shown that tracking accuracy can be improved by placing sensors based on activity characteristics. For example, sensors with many neighboring sensors provide a consistently higher accuracy than those with few, and sensors in places where the duration of stay is long on average prove to be less beneficial than those where duration of stay is short.

The data suggests that the decrease of tracking accuracy resulting from smaller sets of sensors (i.e., decreased target space resolution) can be largely absorbed by selective placement of sensors. It was shown that tracking two targets in a network of 20 or more can be achieved for over 90% of the time. The algorithm tracked correctly on ten clusters of motion sensors 77.1% of the time.

It must be noted that differences in tracking performance may not only be due to advantageous sensor placement, but also due to favorable data: While tracking in space with many adjacent sensors works well, it neglects in part space where tracking might be particularly difficult but useful, such as in narrow hallways. The share of total events covered by the different subsets of sensors range from 11 to 98%.

The experiment presented here gives insight into the importance of sensor placement for multi-target tracking using binary sensors. The next step will be to find the ideal sensor setup for the data used in this evaluation, which may be a mixture of the sensor subsets and criteria examined here. Furthermore, the algorithm's performance with more than two targets must be evaluated.

It is further planned to include identifying information in the algorithm so as to not only associate the data to tracks, but to identify the target. This way, the sensor graph can be replaced by an individual motion model, further improving tracking accuracy.

Acknowledgment

The authors would like to thank the team of the Center for Advanced Studies in Adaptive Systems (CASAS) at the Washington State University for making their data publicly available.

This work was funded by the German Ministry for Education and Research (BMBF) within the research project *Cicely* (grant 16SV5896).

References

[1] J. Vaupel and H. Lundstrom, “The future of mortality at older ages in developed countries.” 1994.

[2] Eisenmenger, M. and Pötzsch, O. and Sommer, B., “Germany's population by 2050, Results of the 11th coordinated population projection,” Federal Statistical Office, Tech. Rep., 2006.

[3] H. Rothgang, R. Müller, and R. Unger, Themenreport ”Pflege 2030”: was ist zu erwarten - was ist zu tun? Bertelsmann Stiftung, 2012.

[4] W. H. Organization et al., “The european health report 2009: Health and health systems,” 2009.

[5] A. H. Association et al., The AHA annual survey database. American Hospital Association, 2005.

[6] T. Frenken, B. Vester, M. Brell, and A. Hein, “atug: Fully-automated timed up and go assessment using ambient sensor technologies,” in Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2011 5th International Conference on. IEEE, 2011, pp. 55–62.

[7] A. Fleury, M. Vacher, and N. Noury, “Svm-based multimodal classification of activities of daily living in health smart homes: sensors, algorithms, and first experimental results,” Information Technology in Biomedicine, IEEE Transactions on, vol. 14, no. 2, 2010, pp. 274–283.

[8] A. Sixsmith, N. Johnson, and R. Whatmore, “Pyroelectric ir sensor arrays for fall detection in the older population,” in Journal de Physique IV (Proceedings), vol. 128. EDP sciences, 2005, pp. 153–160.

[9] D. Wilson and C. Atkeson, “Simultaneous tracking and activity recognition (star) using many anonymous, binary sensors,” in Pervasive computing. Springer, 2005, pp. 62–79.

[10] F. Krüger, M. Kasparick, T. Mundt, and T. Kirste, “Where are my colleagues and why? tracking multiple persons in indoor environments,” in Intelligent Environments, International Conference on. IEEE, 2014, pp. 190–197.

[11] S. Oh and S. Sastry, “Tracking on a graph,” in Proceedings of the 4th international symposium on Information processing in sensor networks. IEEE Press, 2005, p. 26.

[12] D. Marinakis, G. Dudek, and D. Fleet, “Learning sensor network topology through monte carlo expectation maximization,” in Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on. IEEE, 2005, pp. 4581–4587.

[13] B. Logan, J. Healey, M. Philipose, E. Tapia, and S. Intille, “A long-term evaluation of sensing modalities for activity recognition,” in UbiComp 2007: Ubiquitous Computing. Springer, 2007, pp. 483–500.

[14] E.-E. Steen, T. Frenken, M. Frenken, and A. Hein, “Functional assessment in elderlies homes: Early results from a field trial,” in Ambient Assisted Living. Springer, 2014, pp. 3–17.

[15] E.-E. Steen, T. Frenken, M. Eichelberg, M. Frenken, and A. Hein, “Modeling individual healthy behavior using home automation sensor data: Results from a field trial,” Journal of Ambient Intelligence and Smart Environments, vol. 5, no. 5, 2013, pp. 503–523.

[16] D. Podsiadlo and S. Richardson, “The timed” up & go”: a test of basic functional mobility for frail elderly persons.” Journal of the American geriatrics Society, vol. 39, no. 2, 1991, pp. 142–148.

[17] A. Crandall and D. Cook, Human Behavior Recognition Technologies, H. W. Guesgen and S. Marsland, Eds. IGI Global, Mar. 2013.

[18] T. Frenken, E.-E. Steen, M. Brell, W. Nebel, and A. Hein, “Motion pattern generation and recognition for mobility assessments in domestic environments.” in Proceedings of the 1st Int. Living Usability Lab Workshop on AAL Latest Solutions, 2011, pp. 3–12.

[19] S. Müller, A. Helmer, E.-E. Steen, M. Frenken, and A. Hein, “Automated clustering of home sensor networks to functional regions for the deduction of presence information for medical applications,” Wohnen–Pflege–Teilhabe–Besser leben durch Technik , 2014.

[20] S. Blackman, “Multiple hypothesis tracking for multiple target tracking,” Aerospace and Electronic Systems Magazine, IEEE, vol. 19, no. 1, 2004, pp. 5–18.

[21] S. Müller and A. Hein, “Multi-target tracking in home sensor networks,” in Proc. of the Workshop on AI Problems and Approaches for Intelligent Environments. Prague, Czech Republic: CEUR-WS.org, Aug. 2014, to appear.

[22] D. Reid, “An algorithm for tracking multiple targets,” Automatic Control, IEEE Transactions on, vol. 24, no. 6, 1979, pp. 843–854.

[23] D. Cook and M. Schmitter-Edgecombe, “Assessing the quality of activities in a smart environment,” Methods of information in medicine, vol. 48, no. 5, 2009, p. 480.

 ISBN: 978-1-61208-421-3