

Learning to Identify Users and Predict Their Destination in a Robotic Guidance Application

Xavier Perrin, Francis Colas, Cédric Pradalier, and Roland Siegwart

Abstract. User guidance systems are relevant to various applications of the service robotics field, among which: smart GPS navigator, robotic guides for museum or shopping malls or robotic wheel chairs for disabled persons. Such a system aims at helping its user to reach its destination in a fairly complex environment. If we assume the system is used in a fixed environment by multiple users for multiple navigation task over the course of days or weeks, then it is possible to take advantage of the user routine: from the initial navigational choice, users can be identified and their goal can be predicted. As a result of these prediction, the guidance system can bring its user to its destination while requiring less interaction. This property is particularly relevant for assisting disabled person for whom interaction is a long and complex task. In this paper, we implement a user guidance system using a dynamic Bayesian model and a topological representation of the environment. This model is evaluated with respect to the quality of its action prediction in a scenario involving 4 human users, and it is shown that in addition to the user identity, the goals and actions of the user are accurately predicted.

1 Introduction

Robots are more and more present in the daily life, not only in the industry but also at home as toys or as service robots such as vacuum cleaners. There is also a growing demand in the health-care domain for smart assistive device such as intelligent wheelchairs. The present paper is focused on an intelligent system designed to help the elderly or disabled people in their daily activities. For these people, moving in their houses or passing through doorways may represent challenging tasks. We developed a semi-autonomous robot for improving user mobility while minimizing the required input, i.e. having an interaction process adapted to low throughput devices such as single switches, sip and puff systems, brain machine interfaces, or simple voice recognition. More precisely, at each crossing, the robot proposes a direction of

Xavier Perrin, Francis Colas, Cédric Pradalier, and Roland Siegwart
ETHZ, Zurich, Switzerland
e-mail: xavier.perrin@mavt.ethz.ch

travel to the user who will then either agree or disagree. The better the propositions are, the faster and easier the human-robot interaction is.

In this work, a dynamic Bayesian network (DBN) is used in order to learn the habits of multiple human users of a robotic helper. By *habits*, we mean the *succession of navigational tasks* one executes in a known environment: in a retirement home for instance, one resident wakes up in his bedroom, goes to the bathroom, then to the common room where the breakfast is served, and then goes to other rooms for his daily activities. Another resident will have other preferred locations for his own activities. All these accumulated additional information allow the robotic assistant to help its current human user from the first movements of a new travel until the destination. From the system's point of view, the actual identity of its user is not a relevant information. A user is merely defined by his activity pattern.

In the next section, we describe related works in recognition techniques for goal, user, or activity. In section 3, we detail the developed DBN, represented as a graphical model. Section 4 presents experiments in simulation as well as with a real robot and their results. The final discussion appears in section 5.

2 Related Works

Our aim is to ease the navigation of users, therefore we will consider activity recognition only from the point of view of navigation. In this case, inferring the user's intention requires techniques for plan recognition, which are used in a broad variety of domains, such as motion prediction, speech understanding, video surveillance, and so on. Uncertainty is inherent in plan inference, as the robot does not know in advance the intended destination of a user. Furthermore, many ways can lead to the same destination while one way can lead to several places. Probabilistic reasoning techniques are used in almost all works (review in [1]), as they help to express and maintain the beliefs in the possible goal destinations.

In the particular domain of *intention recognition* for navigation, two aspects were studied recently: local intention recognition (immediate action or location in the vicinity of the wheelchair from uncertain input) and global intention recognition (goal destination from local decision). Our work focuses on this latter issue, assuming that the recognition of the immediate intention is solved. In our test, this will be achieved by an interaction device with low uncertainty, such as a joystick or a reliable speech recognition. Inferring a global intention requires the ability to localize the robot on an available map of the environment. In a discrete environment represented by adjacent cells, Verma and Rao [12] described a Dynamic Bayesian Network (DBN) in the form of a graphical model composed of a Partially Observable Markov Decision Process (POMDP) enhanced with the notion of goal locations. Relying on reinforcement learning techniques, the agent explores how it should behave in order to reach three possible goals. Based on this acquired knowledge, the system could infer the most probable goal from a set of possibly noisy observations. Taha et al. [10] achieved similar results by fusing the robot position and the possible goals as the new state definition in a common POMDP. Vasquez [11] described a growing

hidden Markov model algorithm dedicated to continuous learning, clustering, and making inference about car motions in a parking lot or people motions in a hallway.

Topological maps were successfully used in the works of Taha [10]. This environment representation is compact and matches the human's natural description of a path better than metric representations [7] (e.g. "go on the left at the second crossing" instead of "go straight for 100 meters, then turn left"). Many techniques exist for the map construction, e.g. the generalized Voronoi diagram, and applications based on imprecise human drawings have also been reported [9].

User models are often used in computer applications (e.g. intelligent help) or online search or e-commerce sites (e.g. book recommendation based on other customer's choices) [8]. Based on databases containing information from numerous users and some observations from the current user activity, reasoning techniques infer the next user action and try to help him. With the aging population, some research domains are focused on the *activity recognition* in so-called smart environments, i.e. environments where sensors have been installed in order to monitor the activity of human beings. The accumulated data are used to train algorithms which later serves to determine human activities, supervise the user's condition and medication, or detect anomalies [2, 4, 6, 13]. Some researches try to further determine the attributes (coffee drinker, smoker) of multiple users based on their location [3, 5].

As a summary, the studies on activity recognition perform well in capturing the user habits, but they can only monitor ongoing activities, not help the user to perform a task. On the other hand, the studies on intention recognition for facilitating the user's motion do not try to learn the typical user's daily habits. In this article, we propose a system that learns the habits of multiple human users controlling a robot and exploit this knowledge for the navigational control of the robot. In case of an unknown user, the same system first infers the identity of the user from the first movement in the environment before being able to exploit this knowledge for helping its user specifically. The robotic system is given a topological representation of the environment a priori. While incrementally learning the global intentions of a known user, it anticipates the user's destination and proposes better actions, as regular patterns of actions are performed day after day by the same user.

3 Model Description

The graphical model shown in figure 1 represents the dynamic Bayesian network which composes the core of the system. Our DBN is an extension of a Markov Decision Process (MDP), with the agent's *State* S_t and the *Action* A_t (lower part). The state S_t is the observed pose of the robot, in terms of position in a given topological map (a node) and orientation.¹ The action A_t is among the repertoire: forward, left, right, u-turn, and stop. An action always connects to nodes of the map. The MDP is enhanced by the notion of *Goal state* G_t , indicating the current global goal, and the

¹ In relatively small and not too dynamic environments as the ones considered for this work, localization can be considered as a solved issue.




Fig. 1 Our dynamic Bayesian model for learning the habits of multiple users.

variable *Goal reached* R_t , a Boolean value indicating whether the agent's current state S_t is a goal state or not ($P(R_t|G_t S_t) = 1$ iff $G_t = S_t$).

The term $P(A_t|S_t G_t R_t)$ is specified separately according to whether or not the goal is reached. First, when $R_t = 0$, $P(A_t|S_t G_t)$ is the action model leading to each goal G_t from any state S_t as deduced from the topological map (i.e no reinforcement learning technique used). Then, when $R_t = 1$, the probability of staying still $P(A_t = \text{stop})$ is higher than the other actions, so that the user can easily stop when it reaches a goal location. The last term of the MDP, the transition model $P(S_t|S_{t-1} A_{t-1})$, is also deduced from the topological map.

With the goal of learning the user habits for multiple users, we added two more variables. First, the user is symbolized with the variable U_t . Second, the variable G_{last} , representing the last visited goal, makes possible the learning of the succession of place visits, as will be shown later. Both the goal model $P(G_t|G_{t-1} R_{t-1} U_{t-1} G_{last})$ and the user model $P(U_t|U_{t-1} G_{t-1} G_{last})$ result from submodels, where predefined terms related to the persistence of a goal (keeping the same goal or switching to another one), resp. the persistence of a user, are combined with terms learned online representing the knowledge acquired by the system. These submodels are described in details in the appendix. When reaching a goal during an online supervised learning phase, the known user and the particular G_t and G_{last} are used for updating the histograms of the probability distributions of the goal and user models.

From the model shown in figure 1, we compute at each time step several probability distributions. After the learning phase, the user is not known any more. Nevertheless, given the known variables S_{t-1} , A_{t-1} , and G_{last} , we can infer his identity by computing $P(U_t|S_{t-1} A_{t-1} G_{last})$ using Bayes' rule and the law of total probability:

$$\begin{aligned}
 & P(U_t|S_{t-1} A_{t-1} G_{last}) \\
 &= \frac{1}{Z} \sum_{G_{t-1}} \left\{ \begin{aligned} & P(G_{t-1}) \sum_{U_{t-1}} [P(U_{t-1}) P(U_t|U_{t-1} G_{t-1} G_{last})] \\ & \times \sum_{R_{t-1}} [P(R_{t-1}|G_{t-1} S_{t-1}) P(A_{t-1}|G_{t-1} S_{t-1} R_{t-1})] \end{aligned} \right\} \quad (1)
 \end{aligned}$$

where the term $\frac{1}{Z}$ corresponds to a normalization factor. For a guidance robot, we also want to infer the goal the robot should be aiming to or the action it should propose to the unknown user. These two distributions can be computed as follows:

$$\begin{aligned}
& P(G_t | S_{t-1} A_{t-1} G_{last}) \\
&= \frac{1}{Z} \sum_{G_{t-1}} \left\{ P(G_{t-1}) \sum_{R_{t-1}} \left[\frac{P(R_{t-1} | G_{t-1} S_{t-1}) P(A_{t-1} | G_{t-1} S_{t-1} R_{t-1})}{\sum_{U_{t-1}} \{P(U_{t-1}) P(G_t | G_{t-1} R_{t-1} U_{t-1} G_{last})\}} \right] \right\}
\end{aligned} \tag{2}$$

$$\begin{aligned}
& P(A_t | S_{t-1} A_{t-1} G_{last}) \\
&= \frac{1}{Z} \sum_{S_t} \left\{ \sum_{G_{t-1}} \left[\frac{P(S_t | A_{t-1} S_{t-1}) P(G_{t-1})}{\sum_{R_{t-1}} \left\{ \frac{P(R_{t-1} | G_{t-1} S_{t-1}) P(A_{t-1} | G_{t-1} S_{t-1} R_{t-1})}{\sum_{U_{t-1}} \left[\frac{P(U_{t-1}) P(G_t | U_{t-1} G_{t-1} R_{t-1} G_{last})}{\sum_{G_t} \left\{ \frac{P(G_t | U_{t-1} G_{t-1} R_{t-1} G_{last})}{\sum_{R_t} \left[\frac{P(R_t | G_t S_t)}{P(A_t | G_t S_t R_t)} \right]} \right\}} \right]} \right\}} \right] \right\}
\end{aligned} \tag{3}$$

If the user is known by the robot (e.g. through manual or visual identification), we can use a Dirac distribution for $P(U_{t-1})$ and recompute all the above equations with this additional information, the system giving back the learned user habits.

4 Experiments

We run experiments of our multi-user guidance robot in our laboratory environment, using a differential-drive robot (fig. 2a). Figure 3 shows that the topological decomposition of the environment is made of goal nodes and connecting nodes. The former ones represent either people's desks (D1-D6) or common rooms like the cafeteria (C), the printer room (P), the robot lab (RL), or the bathroom (B). Four users share the robot, each with a particular desk and typical sequence, all starting from the entrance E. User 1 executes the sequence D3-D4-P-C-D3-D2-D3, user 2 D4-RL-C-D3-D4-RL-D4, user 3 D2-C-D2-D3-D2-P-D2, and user 4 D6-RL-C-D4-D6-RL-D6 (fig. 4a-d). Additionally each user goes to the bathroom at a random point in his sequence. These sequences are repeated 20 times during the learning phase. In order to speed up this process, we used a simulation. However, for the tests on the 21th




Fig. 2 (a) The differential drive robot used in our experiments. (b) Example of action proposed to the user, who then either agrees or disagrees orally (use of a speech recognition software).




Fig. 3 Topological map of our laboratory used for the experiments with an entrance (E), desks (D1-D6), a printer room (P), a cafeteria (C), a robot lab (RL), and toilets (T).

day, a real robot was used.² In a second phase, the real robot was used by each of the user, the robot having no preliminary knowledge about the current user. At each time step, the system computes the probability distribution over the user, the goal, and the action with, respectively, equations 1, 2, and 3. This last distribution is used to propose an action to the user (see fig. 2b), who can confirm or not. In this precise experiment, this was achieved through a speech recognition system.³

Results of the user recognition are shown in figure 5, which present the evolution of the user probability distribution during a trial of the test phase. Starting from a uniform distribution, the actual user is correctly inferred as soon as a revealing action, mostly leading towards his desk, is executed. In the scenarios, the different user habits shared some common goals, or even some common sequences of goals. These similarities can be seen in figure 5, where the probability in the most probable user decreases and the one(s) for other user(s) increases (e.g. fig 5b, the sequence between goals D3 & D4 is shared by users 1 & 2). We can also notice that when a user goes to the bathroom, his inferred probability decreases slightly. As every user can go to the bathroom at any time, this is not a discriminant observation, the models tending thus slightly to a uniform due to the transition in the user submodel.

For quantifying our model, we compare the model prediction of the user, the goal, and the action with the real values. Based on their distributions computed at each time step, we can check if its maximum is the actual user, intended goal and proper action proposition. We thus introduce a measure μ_{max} , being the mean of the number of times the max of a distribution matches the real value. For example, the formula for the goal is $\mu_{max} = \frac{1}{N} \sum_{t=1}^N \delta(\hat{G}_t, \text{argmax}_{G_t} P(G_t | S_{t-1} A_{t-1} G_{last}))$, with N the number of steps, \hat{G}_t the intended goal, and $\delta(a, b)$ the Kronecker function.

The results are displayed in table 1. As can be seen when comparing with figure 5, the user are properly recognized, the differences being explained by the amount of steps until a revealing action. Concerning the goals and the actions, their probabilities are related to the learned habits of the users and are also influenced by the geographical location of the visited goals. As an example, user 3 visited three

² In simulation, the robot moves instantaneously to the next node, the inputs to our DBN S_{t-1} , A_{t-1} , and G_{last} being identical to when using the real robot.

³ The software performed well enough to assume a perfect recognition of the user agreement (“Yes”) or disagreement (“No”).

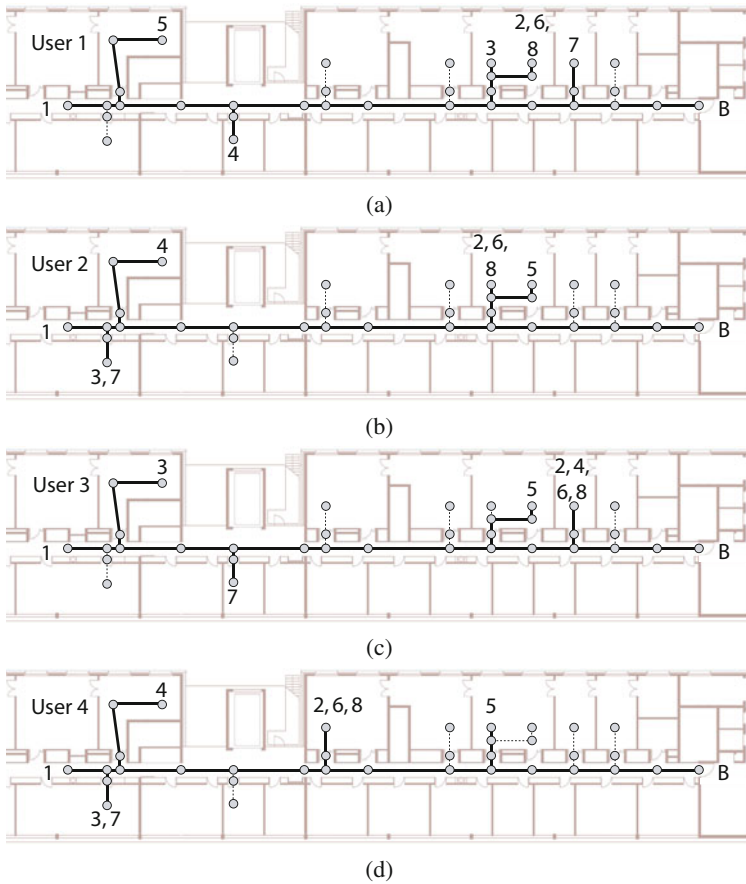


Fig. 4 Habit description of the four users. The succession of visited places is marked with increasing numbers. Each user visits the bathroom (B) at random among the sequence.

goals on the right when leaving his office, while the fourth one was on the left. Accordingly, the probability of going right, resp. of the three right goals were higher than going left to the bathroom. Furthermore, we can notice that the action inference is much better than the goal inference. Indeed, most of the time, the goal confusion is between two nearby goals that share a part of the sequence. As a consequence, the goal could be erroneously inferred during the first steps of the travel, while the action inference will be wrong only at a branching node.

Finally, we want to assess the robustness of our system. To this end, we run scenarios where two users took over the robot when user 2 reaches the cafeteria (C). First, the results with user 3 are displayed in figure 6a. As can be seen, while leaving C, the robot thought user 2 was going to D3, but also increased the probability in user 1, 3, and 4 as they all go in that direction after being in C. At each revealing action, the probability in specific users changes. But when reaching D2, user 3 starts

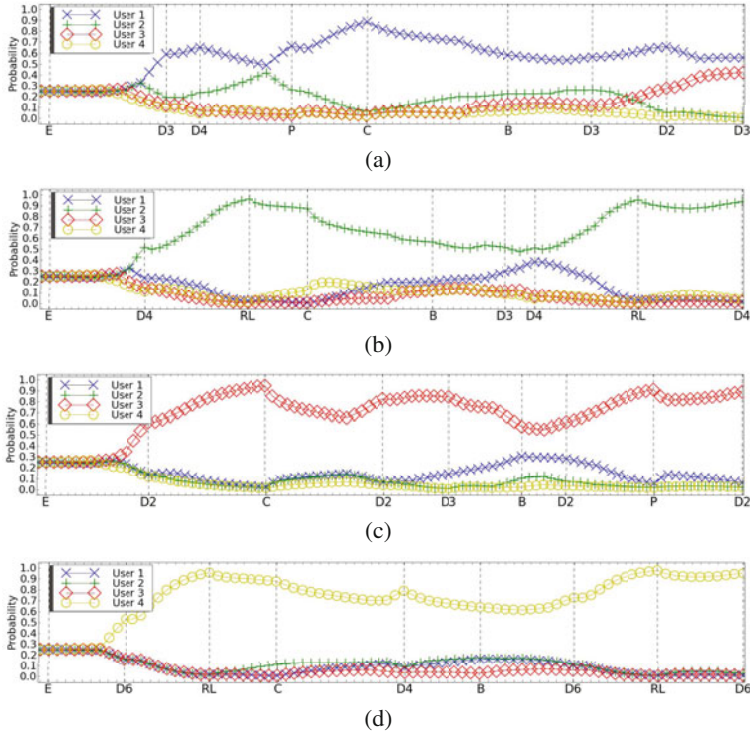


Fig. 5 User probability when (a) user 1, (b) user 2, (c) user 3, (d) user 4 was steering the robot. The dotted lines indicate the steps where a user stopped at a goal.

Table 1 Model performance for the user, goal, and action predictions of each of the four single user scenarios (left columns) and the two taking-over scenarios (right columns).

μ_{max}	Standard usage				Robot taken from user 2		Random
	User 1	User 2	User 3	User 4	User 3	User 4	
User	0.89	0.91	0.94	0.96	0.76	0.56	0.25
Goal	0.64	0.63	0.57	0.84	0.54	0.52	0.09
Action	0.94	0.93	0.94	0.96	0.89	0.92	0.20

being the most probable one. The increasing probability of user 1 is explained by the shared common goals with user 3, and starts to decline as soon as user 3 comes back to D2. In figure 6b, the results when user 4 takes the robot from user 2 are displayed. As user 4 first goes to the desk of user 2, this sequence is not revealing enough because only the former action differs from the habit of user 2. Then, going to the toilets does not bring any information about the user’s identity. Finally, when user 4 reaches his desk, the remaining sequences are specific enough for the system to infer its user correctly. Overall, despite a performance loss in both goal and user recognition, table 1 shows that the actions proposed to the user are still relevant.




Fig. 6 Scenarios where (a) user 3 and (b) user 4 take the robot from user 2 in the cafeteria. The user probability is displayed. The dotted lines indicate the steps where a user stops at a goal. The black region indicates when a new user is steering the robot.

5 Conclusion

In this paper, we introduced a probabilistic model able to learn the daily habits of different users of a guidance robot. When the robot accumulated enough knowledge about a particular user, it is able to actively help its user to reach a goal destination. As shown in the experiments starting with an unknown user, the same model is able to infer the identity of the user after just a few revealing actions have been made, or specific goals have been reached. Having then a high confidence in a particular user, the system is again able to actively help its user. Furthermore, experiments where a different user takes over the guidance robot from an initial user have shown the ability of the system to recognize the switch between the users.

We tested our system with a reasonable environment size and user number, but it would be interesting to test its robustness to an increased user number or to complex user patterns. If leading to pattern overlaps, the system should still be able to help the users as they would go to the same location. However, the extreme case would be if there is a uniform distribution on the goals when leaving a particular one (e.g. the bathroom). For these cases, discretizing the user sequences over different periods of the day or increasing the history length should help.

Acknowledgements. This work has been supported by the Swiss National Science Foundation NCCR ‘IM2’ and by the EC-contract number BACS FP6-IST-027140. This paper only reflects the authors’ view and funding agencies are not liable for any use that may be made of the information contained herein.

Appendix: Submodels for the Goal and User Models

In this appendix, we first describe the submodel used for building the goal model in the main system, then the submodel for the user model. In the goal model $P(G_t|G_{t-1} R_{t-1} U_{t-1} G_{last})$, the goal at time t is dependent from a lot of variables issued from the previous time step. Some relations, e.g. between G_t and G_{last} , are user-dependent and some others, e.g. between G_t and R_t , can be predefined. A submodel is a Bayesian model on his own, with specific relation among the variables, which is then used in order to infer the probability distribution over some variables given some other known variables. In the case of the goal model, our submodel is defined as follows:

$$P(G_t G_{t-1} R_{t-1} U_{t-1} G_{last}) = P(G_{t-1} R_{t-1}) P(G_t|G_{t-1} R_{t-1}) P(U_{t-1} G_{last}|G_t)$$

We have $P(G_{t-1} R_{t-1})$ which is a uniform distribution. Then, $P(G_t|G_{t-1} R_{t-1})$ defines the persistence of a goal: if a goal is not reached, the probability of having the same goal in mind is much higher than changing, whereas when the user reaches a goal, the other places become equally probable with a preference of staying in the current place. This term is predefined. Finally, the $P(U_{t-1} G_{last}|G_t)$ distribution is derived from a histogram learned by the system when a user stops at a goal location. From this model, we compute the probability distribution $P(G_t|G_{t-1} R_{t-1} U_{t-1} G_{last})$, which is then used in the main model as the goal model, as follows:

$$P(G_t|G_{t-1} R_{t-1} U_{t-1} G_{last}) \propto P(G_t|G_{t-1} R_{t-1})P(U_{t-1} G_{last}|G_t)$$

The user model is build in a similar manner than the goal model and is the following:

$$P(U_t U_{t-1} G_{t-1} G_{last}) = P(U_{t-1}) P(U_t|U_{t-1}) P(G_{t-1} G_{last}|U_t)$$

Again, $P(U_{t-1})$ is a uniform distribution. $P(U_t|U_{t-1})$ describes the persistence of the human user, i.e. keeping the same user is more probable than switching between users. $P(G_{t-1} G_{last}|U_t)$ is also derived from a histogram, which is also learned when a user reaches a goal location. The probability distribution computed in order to build the user model is $P(U_t|U_{t-1} G_{t-1} G_{last})$:

$$P(U_t|U_{t-1} G_{t-1} G_{last}) \propto P(U_t|U_{t-1})P(G_{t-1} G_{last}|U_t)$$

References

1. Carberry, S.: Techniques for Plan Recognition. User Modeling and User-Adapted Interaction 11, 31–48 (2001)
2. Hristova, A., Bernardos, A.M., Casar, J.R.: Context-aware services for ambient assisted living: A case-study. In: ISABEL2008, vol. 1, pp. 1–5 (2008)
3. Liao, L., Fox, D., Kautz, H.: Location-Based Activity Recognition using Relational Markov Networks. In: IJCAI 2005, pp. 773–778 (2005)

4. Lu, C.-H., Ho, Y.-C., Fu, L.-C.: Creating Robust Activity Maps Using Wireless Sensor Network in a Smart Home. In: Conference on Automation Science and Engineering, pp. 741–746 (2007)
5. Matsuo, Y., Okazaki, N., Izumi, K., Nakamura, Y., Nishimura, T., Hasida, K., Nakashima, H.: Inferring Long-term User Property based on Users' Location History. In: IJCAI 2007, pp. 2159–2165 (2007)
6. Mori, T., Urushibata, R., Shimosaka, M., Noguchi, H., Sato, T.: Anomaly detection algorithm based on life pattern extraction from accumulated pyroelectric sensor data. In: Proceedings of the IEEE IROS, pp. 2545–2552 (2008)
7. Rawlinson, D., Jarvis, R.: Ways to Tell Robots Where to Go - Directing Autonomous Robots Using Topological Instructions. *IEEE Robotics & Automation Magazine* 15(2), 27–36 (2008)
8. Schickel-Zuber, V., Faltings, B.: Inferring User's Preferences using Ontologies. In: AAAI 2006, vol. 1, pp. 1413–1418 (2006)
9. Setalaphruk, V., Ueno, A., Kume, I., Kono, Y., Kidode, M.: Robot navigation in corridor environments using a sketch floor map. In: International Symposium on Computational Intelligence in Robotics and Automation, vol. 2, pp. 552–557 (2003)
10. Taha, T., Miró, J.V., Dissanayake, G.: POMDP-based Long-term User Intention Prediction for Wheelchair Navigation. In: Proceedings of the IEEE ICRA, pp. 3920–3925 (2008)
11. Vasquez Govea, A.D.: Incremental Learning for Motion Prediction of Pedestrians and Vehicles. PhD thesis, Institut National Polytechnique de Grenoble (2007)
12. Verma, D., Rao, R.: Goal-Based Imitation as Probabilistic Inference over Graphical Models. In: Weiss, Y., Schölkopf, B., Platt, J. (eds.) *Advances in NIPS 18*, pp. 1393–1400. MIT Press, Cambridge (2006)
13. Zheng, H., Wang, H., Black, N.: Human Activity Detection in Smart Home Environment with Self-Adaptive Neural Networks. In: International Conference on Networking, Sensing and Control (ICNSC 2008), pp. 1505–1510 (2008)