

© [2007] IEEE. Reprinted, with permission, from [Tarek Taha, Jaime Valls Miró and Gamini Dissanayake, Wheelchair Driver Assistance and Intention Prediction using POMDPs, Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on 3-6 Dec. 2007]. This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Technology, Sydney's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it

Wheelchair Driver Assistance and Intention Prediction using POMDPs

Tarek Taha, Jaime Valls Miró and Gamini Dissanayake

ARC Centre of Excellence for Autonomous Systems

Mechatronics and Intelligent Systems Group

University of Technology Sydney

NSW2007, Australia

{t.taha, j.vallsmiro, g.dissanayake}@cas.edu.au

Abstract

Electric wheelchairs give otherwise immobile people the freedom of movement, they significantly increase independence and dramatically increase quality of life. However the physical control systems of such wheelchair can be prohibitive for some users; for example, people with severe tremors. Several assisted wheelchair platforms have been developed in the past to assist such users. Algorithms that assist specific behaviors such as door – passing, follow – corridor, or avoid – obstacles have been successful. Recent research has seen a move towards systems that predict the users intentions, based on the users input. These predictions have been typically limited to locations immediately surrounding the wheelchair. This paper presents a new assisted wheelchair driving system with large scale intelligent intention recognition based on POMDPs (Partially Observable Markov Decision Processes). The system acts as an intelligent agent/decision-maker; it relies on minimal user input; to predict the users intention and then autonomously drives the user to his destination. The prediction is constantly being updated as new user input is received allowing for true user/system integration. This shifts the users focus from fine motor-skilled control to coarse control intended to convey intention.

1. INTRODUCTION

The world's aging population and the large number of people effected by motor disabilities has motivated past researchers to develop assistive technologies in the hope of improving the quality of life of those affected. Systems such as robotic walkers [1], smart blind sticks [2] and robotic wheelchairs [3]–[8] have been developed with this goal in mind. Out of these, robotic electric wheelchairs are particularly desirable as they greatly benefit the most needy users but problems exist when they are used manually. For instance, depending on the users type of disability; safely and effectively driving the wheelchair may be difficult in general. Furthermore, such wheelchairs are also large in comparison to the passageways typical in indoors environments such as offices, nursing homes, hospitals and houses. This means that for some users tasks such as

passing through a doorway or navigating a hallway may be quite challenging.

For those who experience these considerable difficulties when driving an electric wheelchair, an intelligent system designed to assist driving would increase their mobility and increase their quality of life. This intelligent system must perform the fine motion control required in the presence of obstacles or in tight spaces. Further to this, the system should be intelligent enough to understand the users intentions and to comply with them whilst demanding the minimal input from the user. The interaction must be transparent to the user and should rely on the least possible number of communication devices. In this paper we are proposing a strategy to accomplish this.

2. RELATED WORK

If a truly user-machine integrated system is to be developed, the type of cooperation between user and machine must be comparable to the cooperation between a horse and its rider [9]; the rider navigates, while the horse avoids (small) dangerous obstacles on the ground. To achieve this level of user-machine integration the machine must grow and learn with the user so that a relationship may form (such as that between a horse and its rider) and so that the machine can predict the users intention and autonomously enact the intention with only minimal corrective input from the user.

The fundamental component of this relationship is intention recognition. The primary considerations for an intention recognition system are whether an accurate representation/model of the environment is required and whether the system is going to be reactive or deliberative. Reactive refers to systems that do not use a representation of the environment and therefore are usually weak in decision making and long term prediction as they rely on local or temporal information collected online which might not be sufficient to make correct long term plans. These reactive systems rely on directly reacting to the changes in the real world by linking the perception obtained by their sensors to their effectors; the control doesn't comply to a model but simply happens as a low level response to the perception. Systems with limited resources like processing power, memory and communication mediums often use a

reactive system but the scope of their possible applications and intelligence is limited. On the other hand, having a representation of the environment and processing the perception through an internal model reduces the ability of the robot to act quickly in new environments.

Several wheelchair platforms that help people in their daily navigation tasks have been previously developed: Rolland III [3], Bremen Autonomous Wheelchair, Sharioto [4], RobChair [5], Senario [6], VAHM [7], Wheellesley [8], and Navchair [10]. The majority of these algorithms are reactive control algorithms limited by either sets of operating modes that the user must select (manual, fully autonomous or semi-autonomous) or by the set of states they try to predict in a local scale. The drawback of such systems is their limitation to manual mode selection and the limited scope of their navigation algorithms and intention recognition.

In the last few years some wheelchair assistive reactive techniques have emerged based on capturing the users local intentions and assisting them with a limited set of motion states or tasks like avoiding obstacles, following a wall, entering a room, going to an object or a local location of interest. These algorithms are based on systems that can *act intelligently* but not *think intelligently*. In other word, they try to make the link between perception and action as direct as possible by combining decisions made by both the human and the machine [9], [10]. Even though such techniques appear to work well in predicting the user's intentions in a local scale (same room or same open space) these algorithms lack the intelligence to autonomously recognize the user's places of interest; rather, the usually visited places/locations must be manually specified prior to system operation. They also lack the larger scope of assistance that would help the user to achieve specific goal such as going to the bathroom or to the kitchen using his wheelchair. Disabled users require an intelligent system that can help achieve such tasks with as little feedback as possible whilst allowing them to remain in complete control of the system.

Unlike previous prediction methods, this paper presents a larger scale intention recognition and goal prediction assistance strategy that uses environmental knowledge to plan and interact deliberately. Typically, users requiring wheelchair assistance have a known set of target locations that they go to during their daily activities: locations like the bathroom, the kitchen or the T.V. room. Through monitoring a wheelchair user through his daily routine/activities the system presented in this paper determines the locations of interest that the use regularly frequents and builds knowledge about these locations using machine learning techniques. This knowledge is then used to predict the users intended destination and to assist the user in driving the wheelchair to this destination. The wheelchair is considered an intelligent agent with an internal representation of the environment and the intention prediction, derived via a POMDPs model of the system, is translated into a planning method in the presence of uncertainties. The intelligent controller then performs low-level navigation tasks such as local path planning, collision avoidance and actuating

the fine motion control of the wheelchair to assist the user in reaching his goal position. Whilst in motion the user performs high level planning by providing action/course correcting feedback to the system through the intention recognition algorithm. This correcting feedback can be inputted through different kinds of interfaces (e.g. a traditional joystick or a voice interface) and is not necessarily a continuous input as with shared autonomy.

3. INTRODUCTION TO POMDPs

Partially Observable Markov Decision Processes (POMDP) provides a general framework for sequential decision making in environments where states are hidden and actions are stochastic. POMDPs were recently used in assistive technologies and they proved to be successful and applicable to wide range of assistive applications [11], [12]. A POMDP model represents the dynamics of the environment, such as the probabilistic outcomes of the actions (the transition function T), the reward function R , and the probabilistic relationships between the agents observations and the states of the environment (the observation function O). In POMDP the model and the states of the environment are typically referred to as "hidden states", since they are not directly observable by the agent. The POMDP framework is a systematic approach that uses belief states to represent memory of past actions and observations. Thus, it enables an agent to compute optimal policies using the MDP framework. A POMDP model is defined by a six tuple $\langle S, A, T, R, \gamma, Z, O \rangle$:

- S : A set of states that represents the state of the system at each point in time.
- A : A set of actions that an agent can take (can depend on the current state).
- $T: A \times S \times S \rightarrow [0, 1]$: The state transition function, which maps each state action pair into a probability distribution over the state space. The next distribution over the state space depends only on the current state action pair and not on the previous state action pairs. This requirement ensures the the *Markovian property* of the process. We write $T(s, a, s')$ for the probability that an agent took action a from state s and reached state s' .
- $R: S \times A \rightarrow \mathcal{R}$: The immediate reward function which indicates the reward for doing an action in some state.
- Z : A set of observations.
- γ : A discount factor used to reduce the award given to far future steps.
- $O: A \times S \times Z \rightarrow [0, 1]$: A function that maps the action at time $t-1$ and the state at time t to a distribution over the observation set. We write $O(s', a, z)$ for the probability of making observation z given that the agent took action a and landed in state s' .

The belief is a sufficient statistic for a given history and it's update at each time step is represented by Equation 1 where $Pr(o|a, b)$ is a normalizing constant [13], [14].

$$b'(s') = \frac{O(s', a, o) \sum_{s \in S} T(s, a, s') b(s)}{Pr(o|a, b)} \quad (1)$$

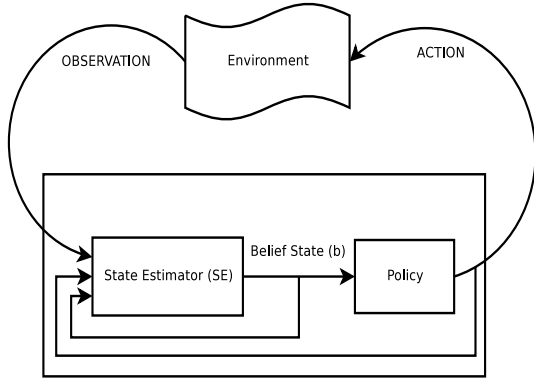


Fig. 1: A POMDP agent is made up of two main components. The state estimator module receives observations from the environment, the last action taken and the previous belief state and produces an updated belief state. The policy module maps belief state to an action.

The goal of POMDP is to find a sequence of actions a_0, \dots, a_t that maximizes the expected sum of rewards $E[\sum_t \gamma^t R(s_t, a_t)]$. Since the states are not fully observable, the goal is to maximize the expected reward for each belief [15]. The value function is calculated using the Bellman Equation 2 and the optimal policy can be obtained using Equation 3.

Considering our intention recognition problem as a planning problem, the wheelchair is transformed into an agent and a decision maker that needs to drive the wheelchair user to his destination based on a certain optimal policy (plan). The agent estimates the state of the environment as the input and generate an action based on that as depicted in Figure 1.

$$V^*(s) = \max_a [R(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V^*(s')] \quad (2)$$

$$\pi_t^* = \operatorname{argmax}_a [R(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V_{t-1}^*(s')] \quad (3)$$

4. POMDP PROBLEM SPECIFICATION

Using the POMDP framework, our prediction is transferred into a planning problem in which we are trying to find the best plan that represents the user's intention by reducing the uncertainty in the belief state, categorized by the *Destination* we are trying to reach. In our POMDP formulation, the wheelchair moves stochastically according to a fixed policy. The state space is described by the cross product of two features, the *WheelchairLocation* = $\{s_1, \dots, s_x\}$ and the *Destination* = $\{d_1, \dots, d_y\}$ resulting in a *StateSpace* = $\{s_1d_1, s_2d_1, \dots, s_xd_y\}$. The wheelchair starts from a known position and the plan finishes when the *WheelchairLocation* is the same as the *Destination*. The *Wheelchair* can have one of the following actions: $\{North, South, East, West, DoNothing\}$ indicating the global direction of travel. A reward of -1 is given for each motion step and +100 reward is given when the *Wheelchair* performs an action that leads to *Destination*. Throughout

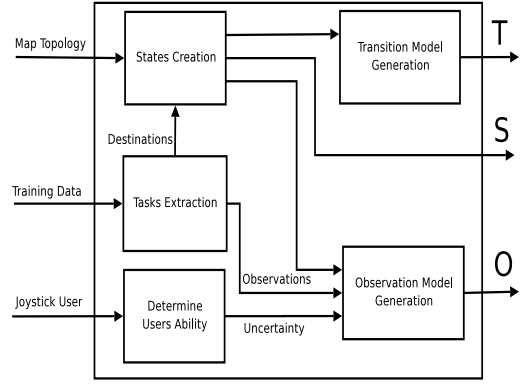


Fig. 2: The POMDP model generation architecture. The map topology together with the training data are used to determine the transition model. The training data is also used to determine the observation model of the POMDP. User's Joystick calibration determines the uncertainty in the observations.

navigation, the *WheelchairLocation* is fully observable but not the destination, and the effect of an action has a predictable deterministic effect as the example described in Equation 4:

$$\Pr(Wheelchair = S_x | Wheelchair = S_y, South) = 1 \quad (4)$$

The position of the *Destination* is completely unobservable until the wheelchair reaches its destination. At each state the joystick input is observed and is represented by a set of discrete states = $\{Up, Down, Right, Left, NoInput\}$, uncertainty in the user's input is taken into consideration while generating the observation model and will be explained in section 5-C.

5. POMDP GENERATION

To obtain a efficient POMDP system, we need to have proper *Transition*, *Observation* and *StateSpace* models. Our model generation consists of three major parts as depicted in Figure 2. These three steps will be explained in the subsections below:

A. State Space

In our assistive system, we want the user to be able to navigate in a high level topological manner. This means that the user should be focusing on driving the wheelchair from one room to another, or from one spatial location to another without having to worry about the intermediate steps that comes in between (planning wise). In order for us to do so, only significant spatial feature are considered, such as a hallway intersection, a door opening or a room.

The ability to learn tasks and represent environments [16], [17] is essential in our application as it creates the bases for the long term intention recognition and prediction. This is done by simplifying the encapsulation of spatial and activity information. For this to happen, the wheelchair should have the ability to represent the spatial information of the environment in a simplistic topological manner that can make it easier to store, extract and update information.

TABLE 1: LIST OF TASKS RECORDED FROM THE USER’S ACTIVITIES

	Start	End	Path
Task1	Lab	Office	26/D - 25/L - 24/L - 22/D - 23/N
Task2	Office	Meeting	42/U - 40/L - 43/U - 44/N
Task3	Office	Bathroom	3/D - 4/L - 5/D - 6/N

For our POMDP platform, the state space consists of two features: the *WheelchairLocation* and the intended *Destination*. The cross product of the above two feature will form the *StateSpace* = $\{s1d1, s2d1, \dots, sxdy\}$, these features are separately extracted in two different steps described below:

1) *Spatial States*: The spatial representation we are using is based on the topological graph representation of the environment, where vertices are locations in the environment and edges represent a viable path connecting two locations as a result of performing an action. In our research we are mainly targeting indoor office or home environments. For such environments there has been a lot of research done on how to build maps and extract topological representation accurately. For simplicity, we assume that the maps are already available and that the topological map representation is hand coded and programmed. It might be more convenient in the future to consider a complete system that can build maps and extract topological representations simultaneously but this is out of the scope of the current research. The map topology will be represented by a graphical tree of nodes and connections (segments), where the set of nodes *WheelchairLocation* = $\{s1, \dots, sx\}$ represents a location in the map and the connection represents a physical path that connects two locations. The hand coded spatial configuration of the domain used for planning illustrated in Figure 3.

2) *Destinations States*: Identifying places of interest is not an easy task and there is no direct method to achieve this as it is an application and environment dependent problem. For the prediction problem we are trying to solve, it’s sufficient to think about the place of interest as a spatial location in the environment where the user spends significantly most of his/her time. After observing the user’s activities we can determine the time that the user need to stay in the same place for it to be considered as a place of interest. In general staying few minutes in a certain location can nominate that location to the place of interest set. For POMDP model generation purposes we log the activities of the user over a period of time, then in that log we determine the locations of interest *Destination* = $\{d1, \dots, dy\}$ based on the above criteria.

B. Transition Model

Transition model specifies the translation from one state to another given a certain action $T(s, a, s')$. In our model specifications, the actions are global navigation commands $\{North, South, East, West, Stop\}$ and determines the spatial node that we will end up at if we are in location s and executed the action a . The transition mode is built directly from the map topology. This transition is deterministic and independent of the intention, so regardless where we want to go. The result of executing an action in the same loca-

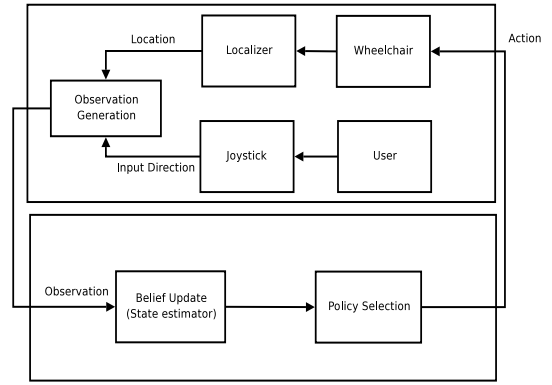


Fig. 4: The POMDP driver assistance architecture. The user’s input together with the current location generate an observation that helps in updating the belief in the destination. An appropriate action will be executed based on that belief driving the wheelchair to the next destination.

tion will be the same. For example $T(s3d1, North, s2d1) = T(s3d2, North, s2d2) = 1$ in the Figure 3.

C. Observation Model

The observation model defines the probability of observing z given that the wheelchair took an action a and landed in state s' $O(s', a, z)$. To generate a proper observation model that correctly models the user’s intention and activities, we use a training data from that particular user. In an indoor environment, the wheelchair users usually perform a repetitive set of tasks that represents navigating from one place to another. A task can be for example going from the living room to the bathroom or to the kitchen. This set of tasks can be defined by the user himself or extracted from a set of data recorded by monitoring the user’s activities. The tasks are defined by a starting location, intermediate locations, end location and the joystick inputs/observation that the user gave at each location as described in Table 1 where the path is represented by numbers corresponding to the states’ numbers and the letters corresponding to the observation in each state (L=Left, R=Right, U=Up and D=Down).

The user in many cases might be unable to give a proper joystick input due to disability or a disease causing a shaky hand for example. To best customize the POMDP model for this user, a joystick calibration is required to determine the uncertainties in the user’s inputs. This uncertainty will be a set of n probabilities describing the user’s inability to give the right joystick input, where n is the number of *JoystickInputs* = $\{Up, Down, Right, Left, NoInput\}$.

Having obtained the training data and the uncertainty, the observation model is then generated by adding the uncertainty to the frequency probability (the probability of obtaining a certain observation in a state).

6. ONLINE ASSISTANCE

Once the planning problem is formulated, we solve the POMDP to get the optimal policy π^* . While predicting online, we start with an initial belief state b_t . Since we know our currently location from our localizer, the initial belief is limited




Fig. 3: The map topology used for our intention recognition. Circles represent intersections and cannot be a destination while squares represent rooms or open spaces and can be considered as a possible destination. The numbers represent the state number and is used to build the transition model. Gray shaded rectangles represent learned destinations.

to those states in the *StateSpace* with our current location, and we will end up with a belief set size equivalent to the available destinations. For example, if our *Destination* = *Kitchen, Bathroom, T.V Room* and we know where we are, then our initial belief is distributed among these destinations and is equal to $1/3$. Based on our initial belief, we execute the optimal policy action for that belief state $\pi^*(s_t)$, calculate the reward r_t for taking that action, get an observation z_{t+1} and update our belief b_{t+1} , then repeat the procedure. This is described in Procedure 1 and illustrated in Figure 4.

Procedure 1 Online Navigation

1. Initial belief: b_t .
 2. Execute the action from the optimal policy: $\pi^*(s_t)$.
 3. Calculate the reward: r_t .
 4. Get an observation : z_{t+1} .
 5. Update the belief: b_{t+1} .
 6. Repeat until destination reached.
-

7. EVALUATION IN A SIMULATED ENVIRONMENT

To validate the proposed intention recognition architecture we simulated a training data that represents the activities of a user in the environment shown in Figure 3. The destinations are represented by the gray shaded squares and they form the set *destinations* = $\{s_{1d1}, s_{6d2}, s_{26d3}, s_{30d4}, s_{31d5}, s_{38d6}\}$. The POMDP was generated using a simulated training data with uncertainty added to the observations to represent the

user’s ability to control the joystick (in this example uncertainty on Up=10%, Down=5%, Right=15%, Left=10% and Nothing=20%). The generated POMDP problem was then solved using *zmdpSolver* [18] and the optimal policy was obtained.

The generated policy and model were tested against the tasks in the training data. For each task in the training data we start with a known location (the first state in the task) but unknown destination (equal belief among destinations) then we take observations from that task one by one, updated the belief based on the takes observation, select an action based on the optimal policy and execute that action to get to the next state. This procedure is repeated until we reach the end of the observations in the task. If the end state reached after the last observation is the same as the intended destination (the last state in the task), then the test for this particular test is considered successful, otherwise it fails. The test was successful in all of the 289 tasks in this experiment producing a 100% success rate.

An example of a single task experiment is shown in Figure 5. The wheelchair starts in spatial state s_{42} in the map shown in Figure 3 and tries to predict where the user is trying to go based on his/her inputs. Initially the user’s destination can be any of the six destinations extracted from the training data, then after the first observation, our belief is now higher in 5 of those destinations. At state s_{29} , the obtained “Right” observation increases the belief that we are going to state s_{31} or s_{38} and reduces that of going to state s_{30} . At state s_{32} , the system is very confident that the user is going to state s_{38} .

Step	1	2	3	4	5	6	7	8	9
Wheelchair Location	S42	S40	S28	S29	S31	S32	S34	S37	S38
Predicted Destination									
d1	19.94%	0.73%	0.03%	0%	0%	0%	0%	0%	0%
d2	0.28%	0.01%	0%	0%	0%	0%	0%	0%	0%
d3	19.94%	24.81%	0.46%	0.01%	0%	0%	0%	0%	0%
d4	19.94%	24.81%	33.17%	1.45%	0.08%	0%	0%	0%	0%
d5	19.94%	24.81%	33.17%	49.27%	5.26%	0.08%	0%	0%	0%
d6	19.94%	24.81%	33.17%	49.27%	94.66%	99.92%	100%	100%	100%
Action Selected	North▲	East▶	North▲	East▶	North▲	East▶	East▶	North▲	Stop
Observation	Up	Right	Up	Right	Up	Right	Right	Up	NoInput

Fig. 5: The result of a simulated experiment. The wheelchair starts in state 42 and tries to predict where the user is going to based on his joystick inputs (observations). The wheelchair in this case successfully takes the user's joystick inputs and decides on the correct actions that take the user to state 38.

8. CONCLUSION

In this paper we have presented a new method for wheelchair assistance that considers the wheelchair as a smart robotic agent, interacting with the user with the aid of a sequential decision making algorithm (POMDP). Unlike most of the currently available assistive methods that are based on semi-autonomous systems which merge wheelchair's perception and user's control with some added heuristics, our method tries to predict where the wheelchair's user is trying to go, and takes him there without any extra mode or behavioral selection. POMDP was chosen because it provides an good platform for planning and predicting under uncertainty for human-robot interaction, as we have shown in this paper. The results we obtained so far are promising and represent a realistic outcome for the smart assistance technology we are trying to develop. The methodology is currently being implemented on a real wheelchair platform, and future efforts are being devoted to design a more sophisticated "activity extrator" strategy.

ACKNOWLEDGMENTS

This work is supported by the Australian Research Council (ARC) through its Centre of Excellence programme, and by the New South Wales State Government. The ARC Centre of Excellence for Autonomous Systems (CAS) is a partnership between the University of Technology Sydney, the University of Sydney and the University of New South Wales.

REFERENCES

- [1] M. Alwan, P. J. Rajendran, A. Ledoux, C. Huang, G. Wasson, P. Sheth, "Stability Margin Monitoring in Steering-Controlled Intelligent Walkers for the Elderly", *AAAI Symposium*, November 2005.
- [2] S.J. Kang, Y. Ho, I. H. Moon, "Development of an intelligent guide-stick for the blind", *IEEE International Conference on Robotics and Automation*, 2001, Vol.4, pp. 3208-3213.
- [3] C. Mandel, K. Huebner, T. Vierhuff, "Towards an Autonomous Wheelchair: Cognitive Aspects in Service Robotics", *In Proceedings of Towards Autonomous Robotic Systems*, 2005, pp. 165-172.
- [4] D. Vanhooydonck, E. Demeester, M. Nuttin, H. Van Brussel, "Shared Control for Intelligent Wheelchairs: an Implicit Estimation of the User Intention", *Proceedings of the ASER '03 1st International Workshop on Advances in Service Robotics*, Bardolino Italy, 13-15 March 2003.
- [5] G. Pires, R. Araujo, U. Nunes, A. T. de Almeida, "RobChair-a powered wheelchair using a behaviour-based navigation", *International Workshop on Advanced Motion Control*, Coimbra., 29 June - 1 July 1998, pp. 536-541.
- [6] P. Beattie, J. M. Bishop, "Localization of the SENARIO wheelchair", *Mobile Robotics Technology for Health Care Services, Proceedings of the 1st MobiNet Symposium*, Athens Greece, pp. 287-293.
- [7] A. Pruski, M. Ennaji, Y. Morere, "VAHM: a user adapted intelligent wheelchair", *International Conference on Control Applications*, Vol.2, Iss., 2002, pp. 784-789.
- [8] A. Yanco Holly, "Development and Testing of a Robotic Wheelchair System for Outdoor Navigation", *In Proceedings of the 2001 Conference of the Rehabilitation Engineering and Assistive Technology Society of North America*, RESNA Press, 2001.
- [9] K. A. Tahboub, "Intelligent Human-Machine Interaction Based on Dynamic Bayesian Networks Probabilistic Intention Recognition", *Journal of Intelligent and Robotic Systems*, Volume 45, Number 1, pp. 31-52.
- [10] E. Demeester, A. Huntemann, D. Vanhooydonck, G. Vanacker, A. Degeest, H. V. Brussel, M. Nuttin, "Bayesian Estimation of Wheelchair Driver Intents: Modeling Intents as Geometric Paths Tracked by the Driver", *International IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing China, October 2006.
- [11] E. López, L. M. Bergasa, R. Barea y M. Escudero, "A navigation system for assistant robots using visually augmented POMDPs", *Autonomous Robots*, vol. 19, no. 1, July 2005.
- [12] J. Hoey, A. V. Bertoldi, P. Poupart, A. Mihailidis, "Assisting Persons with Dementia during Handwashing Using a Partially Observable Markov Decision Process", *In Proceedings of the International Conference on Vision Systems*, Bielefeld Germany, 2007.
- [13] L. P. Kaelbling, M. L. Littman, A. R. Cassandra, "Planning and Acting in Partially Observable Stochastic Domains", *Artificial Intelligence*, 1998, Vol. 101, pp. 99-134.
- [14] J. Pineau, G. Gordon, "POMDP Planning for Robust Robot Control". *International Symposium on Robotics Research (ISRR)*, San Francisco USA, 2005.
- [15] J. Pineau, G. Gordon, S. Thrun, "Point-based value iteration: An anytime algorithm for POMDPs", *International Joint Conference on Artificial Intelligence (IJCAI)*, August 2003, pp. 1025-1032.
- [16] T. Spexard, S. Li, B. Wrede, J. Fritsch, G. Sagerer, O. Booij, Z. Zivkovic, B. Terwijn, B. Kröse, "Biron, where are you? - enabling a robot to learn new places in a real home environment by integrating spoken dialog and visual localization", *In Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems. IEEE*, October 2006.
- [17] M. Heerink, B. J. A. Kröse, B. J. Wielinga, V. Evers, "Human-robot user studies in eldercare: Lessons learned", *In Proc. Int. Conf. on Smart Homes and Health Telematics*, Belfast Northern Ireland, June 2006.
- [18] T. Smith, "zmdpSolver", [<http://www.cs.cmu.edu/~trej/zmdp/>].