Intention Driven Assistive Wheelchair Navigation

Tarek Taha
t.taha@cas.edu.au

Jaime Valls Miró
j.vallsmiro@cas.edu.au

Gamini Dissanayake
g.dissanayake@cas.edu.au

ARC Centre of Excellence for Autonomous Systems (CAS)
Faculty of Engineering, University of Technology Sydney (UTS)
NSW2007, Australia

ABSTRACT

This paper presents an intelligent decision-making agent to assist wheelchair users in their daily navigation activities. The system has the ability to predict the users' intended destination at a larger scale, that of a typical office or home arena. This system relies on minimal user input - obtained from a standard wheelchair joystick - in conjunction with a learned Partially Observable Markov Decision Process (POMDP), to estimate and subsequently aid in driving the user to the destination. The projection is constantly being updated, allowing for true user-platform integration. This shifts users' focus from fine motor-skilled control to coarse guidance, broadly intended to convey intention. Successful simulation and experimental results on a real automated wheelchair platform demonstrate the validity of the approach.

1. INTRODUCTION

The world's aging population and the large number of people affected by motor disabilities has motivated past researchers to develop assistive technologies to give otherwise immobile people freedom of movement, dramatically increasing independence and 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 given their social acceptance and ubiquity. Yet depending on the users' type of disability, safely and effectively driving the wheelchair may be difficult in general, for example people with severe tremors. Furthermore, wheelchairs are also large in comparison to the passageways typical of indoors environments such as offices, nursing homes, hospitals and the home environment. This means that for some users, apparently trivial tasks such as passing through a doorway or navigating a hallway may be quite challenging. An intelligent driving system is needed to assist these users by understanding and complying with their intentions. The interaction should take place transparently to the user, demanding a minimal input from them to automatically perform the required fine motion control for the given situation.

In the robotic-assistive scenario being considered here, user and wheelchair should be regarded as two inter-dependant intelligent agents. Each of them has it's own knowledge of their surroundings, preception and level of autonomy/mobility. The user should always be the dominant partner in this relationship but 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.

This paper presents an intention recognition and goal prediction assistance strategy that uses environmental knowledge to plan and interact with the user in a deliberate manner, but at a larger scale. Typically, disabled users requiring wheelchair assistance in a constrained environment such as the home have a known set of target locations that they go to during their daily activities, such as bathroom, kitchen or T.V. room. Through monitoring a wheelchair user through his usual routines, the technique hereby proposed first determines the locations of interest that the user regularly frequents and builds knowledge about these locations using machine learning techniques. This knowledge is then used to predict the users intended destination from sensed inputs, as derived from a POMDP model. In this scheme of things, the wheelchair is considered an intelligent agent with an internal representation of the environment which effectively translates these target destinations into a plan of actions to reach them in the presence of uncertainties. An intelligent controller subsequently performs the lower level navigational tasks such as local path planning, collision avoidance and actuating motion control. It is important to emphasize that whilst in motion the user remains in complete control of the system, providing continuous (or discrete) action/course correcting feedback to the system through the intention recognition algorithm.

2. THE POMDP FRAMEWORK

Partially Observable Markov Decision Processes (POMDP) provide a general framework for sequential decision making in environments where states are hidden (not fully observable) and actions are stochastic [10]. 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). POMDPs can be regarded as a systematic approach that uses belief states to represent memory of past actions and observations.

A POMDP model is defined by $< S, A, T, R, Z, \gamma, O >$, a seven tuple where:

- 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

Table 1: POMDP Model Variables

	Values
Actions (A)	$North, South, East, West, DoNothing$
Observations (Z)	$Up, Down, Right, Left, NoInput$
States (S)	$s_1d_1, s_2d_1, s_3d_1...s_xd_y$

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 $Markovian\, property$ of the process. We define $T(s, a, s')$ as the probability that an agent took action a from state s and reached state s'.

- R: $S \times A \rightarrow \Re$: 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 future (and more uncertain) 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 define $O(s', a, z)$ as 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 is updated at each time step according to (1), where $Pr(o|a, b)$ is a normalizing constant [10], [11].

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

Given a POMDP model, the goal is to find a sequence of actions, or an optimal policy $\{a0, ..., at\}$ 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 [12]. The function $V^*(s)$ that solves the Bellman Equation (2) is called the value function, and its associated optimal policy can be formulated using Equation (3).

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

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

Using the POMDP framework, our intention recognition problem is transferred into a planning problem where the wheelchair is transformed into a decision maker agent required to find the best plan (optimal policy) that represents the user's intention by reducing the uncertainty in the belief state, categorized by the destination the user is trying to reach.

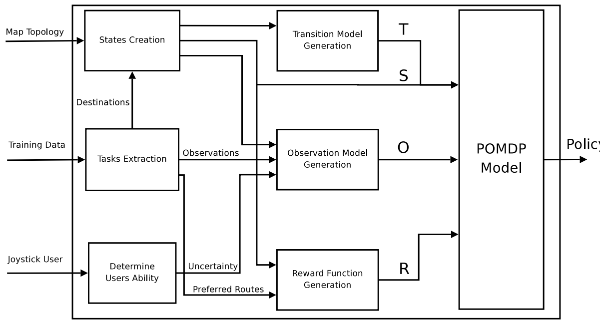

Figure 1: The POMDP model generation architecture. The map topology together with the training data are used to determine the transition model. The training data is also used to determine the observation model of the POMDP. User's joystick calibration determines the uncertainty in the observations.

2.1 An Overview of the Proposed POMDP

For the wheelchair navigation problem driven by user intention presented here, the state space can be best described by the cross product of two features, wheelchair locations $s_i = \{s_1, ..., s_x\}$ and destinations $d_j = \{d_1, ..., d_y\}$ resulting in the state space $S = \{s_1d_1, s_2d_1, ..., s_xd_y\}$. The wheelchair starts from a known postion and the plan finishes when the wheelchair location is the same as one of the destinations. The wheelchair can move according to one of the following actions: $A = \{North, South, East, West, DoNothing\}$ indicating the global direction of travel. It is assumed the wheelchair location is fully observable via a localizer, but the destination is completely unobservable until the wheelchair reaches its destination. Also, the expected effect of an action has a predictable deterministic effect. For instance, in the example described by (4) when the wheelchair is at location s_2 and takes the action $North$, it always ends up at location s_1:

$$Pr(\mathit{Wheelchair} = s_1 \,|\, \mathit{Wheelchair} = s_2, North) = 1 \quad (4)$$

At each state the joystick input is observed and is represented by a set of discrete states $Z = \{Up, Down, Right, Left, NoInput\}$, while the uncertainty in the user's input that represents their capacity to operate the platform is taken into consideration when generating the observation model O. A summary of the POMDP variables is listed in Table 1. In the next section more insight is provided in into how the proposed model is actually generated.

3. MODEL GENERATION

To obtain a efficient POMDP model, we need to properly define the state space (S), transition model (T), observation model (O) and the reward function (R). These are the four major inputs to the POMDP model, as depicted in Fig. 1 within the bigger picture of the overall POMDP model generator.

3.1 State Space (S)

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 come 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.

Table 2: 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

The ability to learn tasks and represent environments [13, 14] 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.

For our POMDP platform, the state space consists of two features: the wheelchair location s_i and the intended destination d_j. The cross product of the two feature will form $S = \{s_1d_1, s_2d_1, ..., s_xd_y\}$. These features are separately extracted in the two different steps described below:

3.1.1 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 $\{s_1, ..., s_x\}$ 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 in the work presented here is illustrated in Fig. 3.

3.1.2 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 needs to stay in the same place for it to be regarded as a place of interest. 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 $\{d_1, ..., d_y\}$ based on this time criteria.

3.2 Transition Model (T)

Transition model specifies the translation from one state to another given a certain action $T(s, a, s')$. In our model specifications, actions represent global navigation commands *{North, South, East, West, Stop}* and determine the spatial nodes s' that we will end up at if we are at location s and executed action a. The

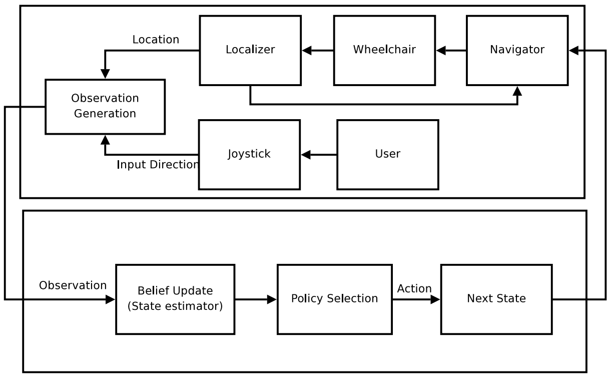

Figure 2: The POMDP driver assistance architecture. The user's input together with the current location generate an observation that helps in updating the belief in the destination. The appropriate action will be selected based on that belief, and the next state will then be determined and given to the navigator to drive the wheelchair to the next state.

transition model 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 location will be the same. For example, in reference to the map presented in Fig. 3, $T(s_3d_1, North, s_2d_1)$ = $T(s_3d_2, North, s_2d_2)$ = 1.

3.3 Observation Model (O)

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 training data from that particular user. In an indoor environment, 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 provided at each of these locations as described in Table 2, where the path is represented by numbers corresponding to the node location of the wheelchair s_i and letters are just a shorthand for the observation at each location $(U = Up, D = Down, R = Right, L = Left, N = NoInput)$.

Given that the proposed solution is aimed at a pool of users unable to provide proper joystick inputs due to some form of disability, a joystick calibration is required to best customize the POMDP model, effectively determining the uncertainties that can be expected for each user's inputs. This uncertainty will be a set of n probabilities describing the user's ability (or inability) to give the right joystick input, where n is the discrete number of possible joystick inputs, or observations, which in this work is restricted to $Z = \{Up, Down, Right, Left, NoInput\}$, i.e. $n = 5$.

It's worth mentioning that in the current stage of our research we decided to use the wheelchair's joystick as the main interface to capture the user's input. This decision is purely based on the simplicity of using such an interface as it's the natural and standard interface that most wheelchair users are accustomed to. Our intention recognition framework can be easily adjusted to use

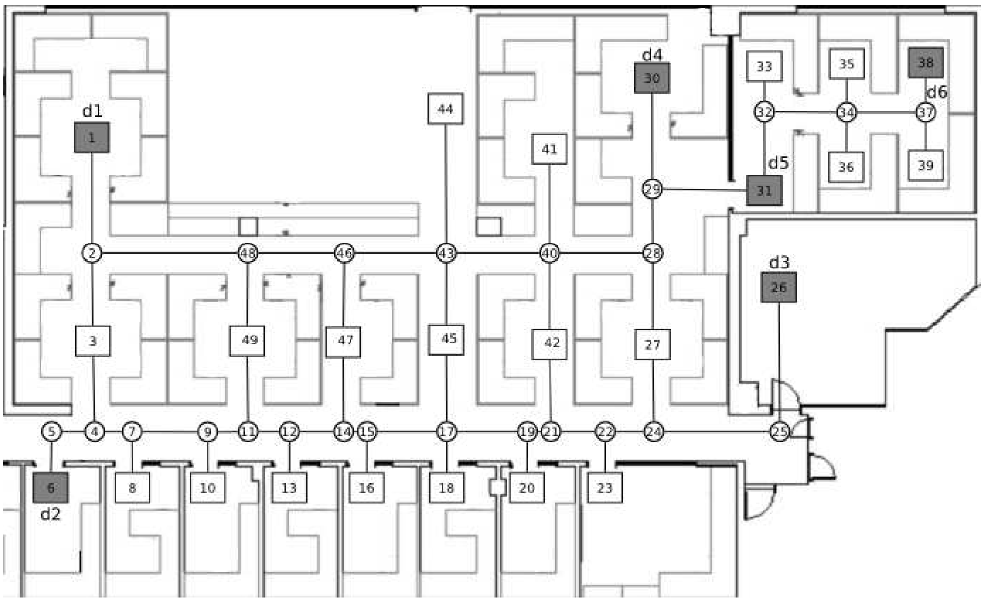

Figure 3: The map topology used for our intention recognition. Circles represent intersections and cannot be a destination while squares represent rooms or open spaces and can be considered as a possible destination. The numbers inside the circles represent the possible wheelchair locations and are used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

other interface (e.g. a head-movement interface [15]).

3.4 Reward Function (R)

The reward function is defined such that all actions that take the user through the preferred routes are rewarded by +10 and those direct actions that lead the wheelchair to a destination are rewarded +100 whilst all remaining actions are rewarded -1 as shown by 5

$$R(s,a,s',o) = \begin{cases} +100 & \text{if } s' \in d_j \\ +10 \times \frac{n_s}{n_t} & \text{if } T(s,a,s') \in Tasks \\ -1 & \text{otherwise} \end{cases} \quad (5)$$

where n_s is how many times this route segment ss' has been traversed in the training data while going to destination d, and n_t is how many times in total a route segment starting from s has been traversed while going to the same destination. This will give more rewards to routes preferred by the user while trying to reach the destination.

4. ONLINE ASSISTANCE

Once the planning problem is formulated and the model is generated as shown in Fig. 1, we solve the POMDP to get the optimal policy π^*. Given the policy, the level of online interaction assistance acceptable by a user will most likely be dependent on the specific class the target user might fall under, more specifically the level of disability such as users with severe or light hand tremors. It is likely this will determine the level of supervised control each users will prefer to have when circulating with the wheelchair. With that in mind, two assistive interaction levels have been implemented and tested within our intention recognition framework:

4.1 Continuous intermediate user input

In this level of interaction, the user is required to give an input at each intermediate step. The wheelchair expects an input from the user as it's approaching the next node, but if by the time it is reached no input is given, it will stop and wait until an input is received. Once the input is received, the belief will be updated (therefore the most likely destination) before making any further navigation decisions. While predicting online, we start with an initial belief state b_t, but since we know our current location from the localizer, the initial belief is limited to those states in the state space S with our current location, ending up with a belief set size equivalent to the available destinations. For example, if our destination set is $Kitchen, Bathroom, T.V Room$ and we know where we are (from localizer), 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 as illustrated in Procedure 1.

Procedure 1 Online Navigation - Continuous Input

1. Initial belief (Uniformly distributed): b_t
2. Execute the action from the optimal policy: $\pi^*(s_t)$
3. Calculate the reward: r_t
4. Get an observation (pause if not given and wait for it) : z_{t+1}
5. Update the belief: b_{t+1}
6. Repeat until destination reached

4.2 Minimal user input

In this interaction level, the wheelchair has more control over the navigation process representing a higher level of trust in the intelligent agent (wheelchair). This level of interaction is convenient for users that have difficulty using the joystick and would prefer a higher level of guidance. The wheelchair starts by predicting from the user's current location the most likely destination. It expects a user's input only if there is ambiguity in the action/route selection (more than one possible destination exist

Step	1	2	3	4	5	6
Wheelchair location	S22	S24	S27	S28	S29	S30
Predicted Destination						
d1	0.725%	0.957%	0.971%	0.027%	0%	0%
d2	0.725%	0.027%	0%	0%	0%	0%
d3	24.63%	13.56%	0.019%	0%	0%	0%
d4	24.63%	32.55%	33%	33.32%	92.3%	100%
d5	24.63%	32.55%	33%	33.32%	3.85%	0%
d6	24.63%	32.55%	33%	33.32%	3.85%	0%
Observation	Right	Up	Up	Up	Up	NoInput
Action Selected	East►	North▲	North▲	North▲	North▲	Stop■
Translate to location	S24	S27	S28	S29	S30	S30

Table 3: The result of an experiment on a real platform. The wheelchair starts in location 22 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 location 30.

with high probability). The user can still contribute his input at any stage during the navigation to override or enhance the wheelchair's decisions. The initial belief in this case is extracted from the tasks history and it represents the probability of going to any of the destinations knowing that we started in location s_i. The online navigation procedure is described in Procedure 2.

Procedure 2 Online Navigation - Minimal Input

1. Extract initial belief from recorded tasks list: b_t
2. Execute the action from the optimal policy: $\pi^*(s_t)$
3. Calculate the reward: r_t
4. Get an observation or $NoInput$ if no input from user: z_{t+1}
5. Update the belief: b_{t+1}
6. Repeat until destination reached

5. EXPERIMENTAL RESULTS

To validate the proposed intention recognition architecture we simulated a training data that represents the activities of a user in the environment shown in Fig. 3. The total number of states is 294 (49 possible locations 6 of which are also destinations). The destinations are represented by the gray shaded squares forming the set $\{s_1d_1, s_6d_2, s_{26}d_3, s_{30}d_4, s_{31}d_5, s_{38}d_6\}$. 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 model was solved off-line using zmdpSolver [16] and the optimal policy was obtained after around 10 hours of running on a linux machine (2.93 GHz, 4 MB Cache, X6800 Core 2 with 4 GB DDR2-RAM).

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 location in the task) but unknown destination (equal belief among destinations) then we take observations from that task one by one, update the belief based on the obtained 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 location 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 navigation task on a real wheelchair platform using continuous interaction is shown in Table 3. The wheelchair used was the one described in [17] and it measures 1.2×0.7m. The wheelchair's size is considered large compared to the environment and driving in such a constrained environment can be a challenging task. In this example, the user was giving observations at each state to indicate where he wants to go. Initially, the user can be going to any of the pre-determined destinations, therefore the belief is uniformly distributed among them. With the first observation, the belief is updated and the next state is determined based on the appropriate selected action and the wheelchair navigates to that state autonomously. This is repeated until the user reaches his destination.

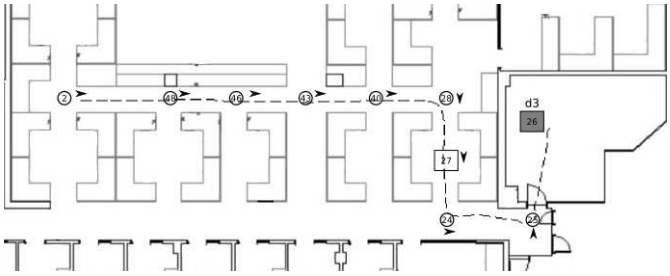

Figure 4: The result of a real Wheelchair experiment showing the path (dashed line) and observations (arrows). The wheelchair starts in location 2 and drives the user successfully to location 26 by updating the belief at each step from the obtained observation.

In another longer example with continuous intermediate inputs, the user was able to navigate from locations 2 to 26 (destination d_3) successfully and the route followed during the navigation is shown in Fig. 4.

In a third example, a minimal input interaction level was used to drive the user also from locations 2 to 26, similar to that shown in Fig. 4. The user contributed only 2 inputs during this navigation demonstrating a high level of trust in the wheelchair. At the beginning of the experiment a belief of where the user might be going is extracted from the tasks history of the user. This initial belief is used to decide if an action can be taken directly or a user input is required to update the belief and reduce the ambiguity. The results are summarized in Table 4, where we can clearly see that at location s_{28} the belief was high in 4 destinations requiring different actions. In this case a user input is required to strengthen the belief in one of those destinations making it easier to select the correct action.

All the above mentioned experiments were performed on only one user. We hope that soon we will be able to conduct experiments on a larger number of users in a nursing home, but this requires Research Ethics Committee approval that we are in the process of attaining.

6. CONCLUSION

In this paper we have presented a novel methodology for wheelchair assistance that regards the wheelchair as a smart robotic agent, capable of interacting with the user with the aid of a sequential decision making algorithm (POMDP). Unlike most

Step	1	2	3	4	5	6	7	8	9	10
Wheelchair location	S2	S48	S46	S43	S40	S28	S27	S24	S25	S26
Predicted Destination										
d1	6.45%	0.22%	0%	0%	0%	0%	0%	0%	0%	0%
d2	6.45%	0.11%	0%	0%	0%	0%	0%	0%	0%	0%
d3	12.9%	14.76%	14.81%	14.81%	14.81%	14.81%	86.85%	99.60%	99.99%	100%
d4	19.35%	22.14%	22.21%	22.21%	22.21%	22.21%	3.43%	0.10%	0%	0%
d5	21.87%	25.03%	25.11%	25.11%	25.11%	25.11%	3.88%	0.12%	0%	0%
d6	32.98%	37.74%	37.86%	37.86%	37.86%	37.86%	5.84%	0.18%	0.001%	0%
Observation	Right	NoInput	NoInput	NoInput	NoInput	Down	NoInput	NoInput	NoInput	NoInput
Action Selected	East ▶	East ▶	East ▶	East ▶	East ▶	South ▼	South ▼	East ▶	North ▲	Stop ■
Translate to location	S48	S46	S43	S40	S28	S27	S24	S25	S26	S26

Table 4: The results of a navigation task with minimal interaction, the user gave only two inputs during navigation.

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 modal or behavioral selection. We also proposed in this paper two levels of interaction between the user and the wheelchair, one of which requires the user to continuously give an input to the system, while the other expects him to give inputs only if it's not clear what the user's intended destination is. We feel that the latter form of interaction assistance better represents a realistic strategy for intention driven wheelchair navigation. Having said that, we are also aware that this level of interaction will not suit all users, as there is realistically no one-size-fits-all type of solution for the scenario presented here. Hence, in the future more efforts are to be directed towards properly studying the satisfaction of the user in the proposed interaction levels and how they can be further improved to occur in the most natural and acceptable manner. Despite the limitation of experiments having been conducted with a reduced set of users so far, we feel they still demonstrate the promising capabilities of the system. We are also currently dedicating considerable efforts to explore recent advances in active policy generation whilst navigating [18], and automated activity monitoring and tasks extraction.

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.

7. REFERENCES

[1] M. Alwan, P. J. Rajendran, A. Ledoux, C. Huang, G. Wasson, and P. Sheth. Stability margin monitoring in steering-controlled intelligent walkers for the elderly. In *Proceedings of AAAI Symposium*, 2005.

[2] S. Kang, Y. Ho, and I. Hyuk Moon. Development of an intelligent guide-stick for the blind. *Proceedings of IEEE International Conference on Robotics and Automation*, 4:3208–3213, 2001.

[3] D. Vanhooydonck, E. Demeester, M. Nuttin, and H. V. Brussel. Shared control for intelligent wheelchairs: an implicit estimation of the user intention. In *Proceedings of the International Workshop on Advances in Service Robotics*, 2003.

[4] C. Mandel, K. Huebner, and T. Vierhuff. Towards an autonomous wheelchair: Cognitive aspects in service robotics. In *Proceedings of Towards Autonomous Robotic Systems*, pages 165–172, 2005.

[5] G. Pires, R. Araujo, U. Nunes, and A. T. de Almeida. Robchair-a powered wheelchair using a behaviour-based navigation. *Proceedings of International Workshop on Advanced Motion Control*, pages 536–541, 1998.

[6] P. Beattie and J. M. Bishop. Localization of the senario wheelchair. In *Proceedings of MobiNet Symposium Mobile Robotics Technology for Health Care Services*, pages 287–293, 1997.

[7] A. Pruski, M. Ennaji, and Y. Morere. Vahm: a user adapted intelligent wheelchair. *Proceedings of the International Conference on Control Applications*, 2:784–789, 2002.

[8] H. A. Yanco. Development and testing of a robotic wheelchair system for outdoor navigation. 21:145–147, 2001.

[9] K. A. Tahboub. Intelligent human-machine interaction based on dynamic bayesian networks probabilistic intention recognition. *Journal of Intelligent and Robotic Systems*, 45(1):31–52, 2006.

[10] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.

[11] J. Pineau and G. Gordon. Pomdp planning for robust robot control. In *Proceedings of International Symposium on Robotics Research*, 2005.

[12] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for pomdps. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 1025–1032, August 2003.

[13] T. Spexard, S. Li, B. Wrede, J. Fritsch, G. Sagerer, O. Booij, Z. Zivkovic, B. Terwijn, and B. Krose. Biron, where are you? enabling a robot to learn new places in a real home environment by integrating spoken dialog and visual localization. *International IEEE Conference on Intelligent Robots and Systems*, pages 934–940, Oct. 2006.

[14] M. Heerink, B. J. A. Krose, B. J. Wielinga, and V. Evers. Human-robot user studies in eldercare: Lessons learned. June 2006.

[15] P.B. Taylor and H.T. Nguyen. Performance of a head-movement interface for wheelchair control. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2:1590–1593, 2003.

[16] T. Smith. zmdpsolver. [http://www.cs.cmu.edu/˜trey/zmdp/].

[17] T. Taha, J. V. Miro, and D. Lui. An efficient path planner for large mobile platforms in cluttered environments. In *Proceedings of IEEE Conference on Robotics, Automation and Mechatronics*, pages 225–230, 2006.

[18] S. Paquet, L. Tobin, and B. Chaib-draa. An online pomdp algorithm for complex multiagent environments. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 970–977, 2005.