"© 2009 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works."

Robotic Assistance with Attitude: a Mobility Agent for Motor Function Rehabilitation and Ambulation Support

Jaime Valls Miró, Vivien Osswald, Mitesh Patel and Gamini Dissanayake

***Abstract*— This paper presents the design of an intelligent walking aid for the frail and elderly as well as for patients who are recovering from surgical procedures, in order to enhance safer mobility for these study populations. The device augments a conventional rolling walker aid with sensing and navigational abilities to safely travel through an environment following user's perceived intentions, unless collisions or instability is imminent. The agent, embodied as a Partially Observable Markov Decision Process (POMDP), critically relies on minimal user input to seamlessly recognise user's short-term intended behaviour, constantly updating this projection to allow for inconspicuous user-robot integration. This, in turn, shifts user's focus from fine motor-skilled control to coarse indications broadly intended to convey intention. Overall, the system can afford an increase in safety for the cognitive user through preventative care - reduced number of falls or collision with surrounding objects, minimising health-care expenses as well as increasing independent living for people with gait disorders. Successful simulation and experimental results demonstrate the validity of the proposed architecture for a practical robotic rollator design.**

I. MOTIVATION

Impaired mobility is a significant problem in the frail and older population, largely as a result of the age-related decline in the musculo-skeletal and neurological systems. Physical complications of immobility include significant health issues such as bed sores, osteoporosis and deep vein thrombosis. Immobility is also associated with other functional impairments, loss of independence, and decline in quality of life. It has been demonstrated that both short- and long-term physical activities can improve health, mobility and functional abilities [1]. Even the frail and the very old can benefit from such activities. With the population ageing at an exponential rate, impaired mobility poses a serious threat to both the society and the individuals [2].

Walking aids (or "walkers" as they are commonly known) provide means whereby many frail older adults and a variety of other persons with gait disorders can maintain mobility, functional independence and social interactions. Without these assistive devices, many of them would be either chair- or bed-bound [3]. Mobility assistants have a large role to play not only in assisting the ageing population with maintaining gait stability and musculo-skeletal strength, but also in the rehabilitation process for those who have had their movement temporarily restricted, usually through surgery. Given the high incidence of falls associated with gait instability in the older population, and concomitant medical and social costs, it seems only natural that walking devices should form an important part of a successful multi-faceted prevention program [4].

Conventional walking aids are fully dependent upon the cognitive capability of the users. Their usability is dependent upon the mental ability of the user to learn, and to use the aids properly as well as safely. In the nursing home, where the prevalence of impaired mobility is the highest, cognitive impairment is equally high, and the need for individual caregiver assistance is ever present. This is only set to worsen as the world's ageing population increases and the proportion of the elderly in need for care grows. Research indicates that the number of elderly people will increase by 50% and those reaching the age of 85 and up will rise by 100% [5]. Considering the projected decline over the coming years in the number of people in active ages, those who can perform the care-taking activities amongst them, there is significant benefit in attending to these shortcomings by incorporating assistive robotic devices into the health-care sector.

II. BACKGROUND

Over the last two decades, the aforementioned circumstances have motivated researchers into developing a range of intelligent assistive robotic technologies which, one way or another, aim to improve the quality of life of those affected. Electric wheelchairs are particularly suitable for a large sample of the potential population of users given their social acceptance and ubiquity, and have seen many roboticised aid variants developed [6], [7]. Other systems such as smart blind sticks (or canes) [8] have also bee been developed with this goal in mind.

This paper focuses on walker-type support systems. Rollator walkers are best suited to ageing adults suffering from slight weakness or those who may be experiencing mild balance problems Behind the cane, they are more widely used mobility aid [9], partly due to the perception in the eyes of older adult users of not carrying the stigma of "being old", when compared with other assistive devices [10].

A reduced area of research in the development of intelligent robotic roller walkers has focused on passively-controlled prototypes where the walker can control the orientation of its front actuated steering-wheel but there is no active driving component, motion being the result of external user forces [11]. Other passive designs use servo brake controls and adaptive braking control strategies to improve manoeuvrability [12]. Passive walker are advantageous in

J. V. Miró, M. Patel and G. Dissanayake are with the Faculty of Engineering and IT, University of Technology Sydney (UTS), NSW 2007, Australia {jaime.vallsmiro, mitesh.patel, gamini.dissanayake}@eng.uts.edu.au

V. Osswald is a visiting student from the University of Technology of Belfort Montbéliard (UTBM), France vivien.osswald@utbm.fr

that they are less complex, lighter, and inherently safer as they need batteries of reduced size, or none, to operate, yet their performance is considerably dampened for these very same reasons. An increasing number of researchers in the field have paid more attention to incorporating active servo motors for power assistance, and a suite of sensors to provide functionality such as collision avoidance, navigation or adjustable motion control during walking. The "VA-PAMAID" walker allows the user explicit control of the amount of assistance provided by selecting one of three modes: manual, automatic and park [13]. A mobility assistant device developed at the CMU (predecessor of the "Nursebot" project [14]) incorporates modules for obstacle avoidance, localisation, mapping, path planning and people tracking [15]. Navigation decisions are based on a user motion model that represents a mapping of force sensor readings from a haptic device to translational and rotational commands. After obtaining these short-term trajectories representing user's directional intentions, a multi-modal shared control operates the walker. The "PAMM" developed at the MIT concentrates on the path planning aspect of a mobility assistant [16]. It is also based on a multi-modal compromise between human and robot controls: (a) the user has complete control and the PAMM provides physical support only, (b) the PAMM leads the user along a planned path at a predetermined speed, (c) similar to (b) but the user is able to control the PAMM's speed by pushing and pulling the handlebars, and (d) where the user has limited control over the path of the PAMM. The "Care-O-Bot" mobility assistant was designed as part of a large home care project for older persons at the Fraunhofer Institute [17]. It also exhibits two major modes of operation for navigation: (a) direct user control where the robot takes readings from a user intent sensor and determines the direction and speed of travel, and (b) target mode which allows users to input a destination based on a map, and the robot will guide the user to the destination in a reactive manner along the calculated route.

What these methods have in common is their limited ability to continually recognise and adapt to the situation in which the user is in, and not just during the navigational tasks, a fundamental challenge yet to be resolved for these robotic assistive agents to be effective. In addition to the usual localisation and navigation capabilities required of a mobile robot, a computer assisted walking aid should appropriately address the fact that users are not required to be aware of the intelligent agent behind the driving wheel. In practice, users should be able to use the system dependably, without specifically considering the functions of the intelligent agent, and by the same token, the cooperative agent has to have the ability to represent the uncertainty inherent in a person's behaviour. Significant steps have already been taken in this direction, such as the "Autominder" System which forms the heart of the aforementioned "Nursebot" project [14]. Here, a computerised cognitive system incorporates AI techniques and a decision theoretic approach based on a probabilistic decision making framework (a Partially Observable Markov Decision Process, or POMDP) to schedule voice reminders and navigational guidance for elderly users. The cognitive assistance to decide if and when to issue the reminders is limited to reasoning between discrepancies between supposed and observed plans, yet the same general decision theoretic model of interaction between user and cognitive agent has recently successfully incorporated user's cognitive attitudes themselves in the models. This is the case of the POMDP models proposed to assist people with more severe cognitive discapacities, such as dementia, in specific domains such as hand-washing [18].

The human-driven robotic agent hereby proposed is aimed at naturally responding to the physical interactions between user and active walking aids with a decision theoretic framework whereby user's attitudes are specifically accounted for, paying special attention to the intended behaviour or course of action, such as standing-up, strolling around, etc. Other cognitive or physical characteristics (e.g user responsiveness or health) can also be incorporated into the same framework at the expense of exactness in the solution and within the limitation of the sensing technologies.

To sum up, the proposed mechanism differs from the more traditional mixed-mode assistance provided by other active walkers in three critical aspects:

1) in assisting (inherently "noisy") users with mild cognitive or physical impairments via a decision-making mechanism that requires minimal indicative input
2) in actively addressing specific muscular strength rehabilitation via safe guided ambulation
3) in combining this task with support for other higher-level repetitive routines such as aided gait stability strolling or safe stand-up/sit-down actions

III. DESIGN OF THE WALKING ROLLING AID PLATFORM

The proposed design, based on a modified commercial rollator walking frame with four wheels, is displayed in Fig. 1. This base design has been instrumented with additional actuators and incremental encoders to the two rear wheels (front casters are passive), two infra-red (IRs) proximity sensors to detect the presence and configuration of the user (e.g. leaning forward or raising from a chair), four strain gauges (SGs), two on each of the walker's handle-bars, two contact switches in the handles, a low-level micro-controller for sensing and actuation, and a high-level control computer, as well as a laser range finder for localisation and reactive navigation.

The strain gauges employed are two Micro Measurements 250UR. The differential force measurements suministered by each pair of these sensors along the vertical axis of each handle-bar are used in conjunction with the contact switches to establish whether (and how strongly) a user is holding onto the handle-bars, in readiness to start some task such as sitting down or ambulation.

The IR subsystem sensor is made up of two Sharp GP2Y0A02YK, which are used to estimate whether the driving user is standing behind the walker (at the handle-

(a) Frontal view (b) Rear view

Fig. 1. The instrumented rollator walker platform showing the laser range finder at the front, the infra-red proximity sensors on top of the black PC controller, servo-motors and encoders.

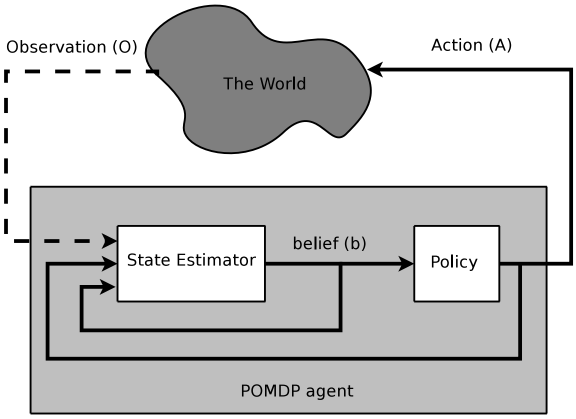

Fig. 2. The POMDP agent.

bars) and how far they are from it. Sensing range after calibration is $[20, 150]$ cm.

The motorised actuation subsystem is based around two 24VDC reversible gear-head motor with optical encoder (detailed in Fig. **??** and rotary mechanical couplings. The motors are PWM driven using a national semiconductor LMD18200 3A, 55V H-Bridge motor driver.

A compact Hokuyo URG-04LX laser range finder is also incorporated in the design for localisation and reactive navigation (it can be seen in Fig. 1a placed at the lower front of the walker platform). The URG-04LX is able to report ranges from 0.02 m to 4.0 m (0.001 m resolution) in a $240°$ arc ($0.36°$ angular resolution). Its power consumption, 500 mA @ 5V, makes it a natural choice for battery operated vehicles.

IV. HIGHER-LEVEL DECISION PLANNING

Readings from the sensorial systems in the robot are continually being observed and fused with past information to resolve for the next action to take in assisting the user with whatever task he/she is intending on doing. A decision theoretic framework to resolve for the next best action is proposed in the form of a POMDP, which allows us to take advantage of its natural fabric for sensor fusion and for the handling of the pervasive uncertainties associated in dealing with human users.

A. The POMDP Framework

POMDPs are decision theoretic models incorporating Artificial Intelligent techniques to calculate optimal control actions under uncertainty. They constitute a general framework for discrete sequential decision making in environments where there is no certainty about the actual state of the "world", i.e., states are not fully observable. Instead, as seen in Fig. 2, a policy maps computed "belief" states representing memory of past actions and observations into stochastic action policies that are expected to maximise the (discounted) sum of future rewards.

More formally, a discrete-time POMDP model is defined by $< S, A, Z, T, O, R, \gamma >$, a seven tuple which represents the dynamics of the environment as the probabilistic outcomes of the actions (the transition function T), the reward function R, and the probabilistic relationships between the agent's observations and the states of the environment (the observation function O), where :

- S: A finite set of states that represents the state of the system at each point in time.
- A: A finite set of actions that an agent can take.
- Z: A finite set of observations.
- T: $A \times S \times S \to [0, 1]$: The stochastic state transition model, 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 decision process (MDP). We denote $T(s, a, s') = Pr(s'|s, a)$ i.e. the probability that an agent took action a from state s and reached state s'.
- O: $A \times S \times Z \to [0, 1]$: 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 $O(s', a, z) = Pr(z|s', a)$ as the probability of making observation *z* given that the agent took action *a* and landed in state s'.
- R: $S \times A \to \Re$: The immediate reward function which indicates the reward for doing an action in some state.
- γ : A discount factor used to reduce the award given to future (and more uncertain) steps.

Given the POMDP model, the goal is to find the sequence of actions, or optimal policy $\pi^*(s) = \{a_0, ..., a_t\}$ which maximises the expected sum of future rewards:

$$E\left[\sum_{t=0}^{t_{max}} \gamma^t R(s_t, a_t)\right] \quad (1)$$

where t_{max} defines the time steps left in a finite horizon problem, ∞ otherwise. However, since states are not fully observable, it uses actions and noisy observations as defined by O to maintain a factored probability distribution of length $|S|$ over all possible hidden states $s_i \in S$, known as the belief

b, updated (from Bayes' Rule) at each time step according to:

$$b'(s') = Pr(s'|b, a, o) = \frac{O(s', a, o)\sum_{s\in S} T(s, a, s')b(s)}{Pr(o|a, b)} \quad (2)$$

where $Pr(o|a, b)$ is a normalising factor that ensures probabilities add up to one by summing the numerator over all possible $s' \in S$:

$$Pr(o|a, b) = \sum_{s\in S} b(s) \sum_{s'\in S} T(s, a, s')O(s', a, o) \quad (3)$$

The elements of the vector $b(i)$ indicate the conditional probability of the agent being in state $s_i \in S$, given an initial belief b_0 and the given evolution of the system so far. Since the belief is an accurate compilation of all the relevant history of the system, it can be shown that using this belief state provides just as much information as the entire action-observation history would [19], i.e., it is a sufficient statistic for selecting optimal actions $\pi^*(b')$. This is important, as it means that given the belief state, a POMDP forms a completely observable Markovian process (MDP) which only depends on the last belief, last action and current observation. Under this transformation, b' becomes the state of an observable MDP. However, what in an MDP is a discrete state space problem, becomes continuous in POMPD, and while the action space of a "belief" MDP remain the same as in a POMDP, the transition and reward functions need to be transformed accordingly. Hence, the optimal policy for any given initial belief, $\pi^*(b)$, is the one that yields the highest expected reward according to (1) for each belief state, referred to as the optimal value function $V^*(b)$ as it assigns values to belief states. It can be formulated as a solution to the following Bellman optimality equation:

$$V^*(b) = \max_{a\in A}[R(b, a) + \gamma \sum_{(b,b')\in B} T(b, a, b')V^*(b')] \quad (4)$$

which becomes a highly dimensional problem as there are an infinite number of belief states $(b, b') \in B$ - what is known as the "dimensionality curse" of POMDPs, and so far exact optimisation for larger models is still computationally intractable. There are numerous studies about finding sub-optimal policies for larger POMDPs models (e.g. [20]). The size of the proposed model detailed below in Section IV-C is beyond the reach of an exact solution. Therefore, an off-line point-based approximation has been obtained based on the SARSOP solver [21], which obtains policies by sampling over the subset of belief points reachable from the initial state under the optimal sequence of actions that iteratively converge to those beliefs.

B. Walker Agent Assistive POMDP

In the context of robot navigation, the finite states of Markovian models have traditionally represented the location of the robotic agent in a given map, either topologically [22] or as a discrete approximation of their geometric location [14]. Under these circumstances, the systematic decision-making approach of a navigational POMDP framework in based on monitoring beliefs and choosing the appropriate navigational actions. In contrast, our proposal differs significantly by transferring the planning problem into a decision-making required to find optimal policies that best match the user's cognitive attitude towards a set of assistive tasks, beyond the purely steering routines associated with some of them. In our proposition, we do not specifically incorporate a model of the user's mental state, as has recently been proposed in order to cue users with various levels of dementia to successfully complete a task such as hand-washing [**?**], [18].Instead we demonstrate how a simpler human-driven robotic agent capable of inventively exploit the physical interactions with the driving user can effectively incorporate his/hers evolving activity model into the planning process.

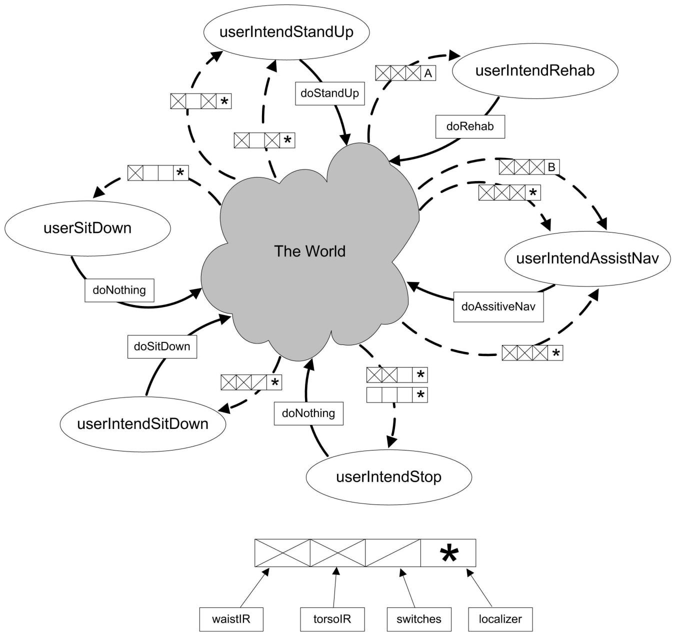

Fig. 3. Diagram of the POMDP agent dynamics with an example readings from the sensor array, indicating that the waist and torso IR have fired, one of the switched is pressed, and the location of the robot is not relevant. These dynamics are described by the stochastic parameters encoded in the transition and observation functions T and O.

C. POMDP Model Dynamics

Within this context, the proposed walker assistive agent can be best modelled by a state space S which describes the tasks `userSitDown`, `userIntendStandUp`, `userIntendSitDown`, `userIntendAssiNav`, `userIntendRehab`, `userIntendStop`. The agent evolves from an unknown initial state, and moves according to one of the following action space A: `doNothing`, `doStandUp`, `doSitDown`, `doAssistiveNav`, `doRehab` corresponding to the intended task the user is trying to perform, as perceived by the agent. A diagram of the model dynamics is depicted in Fig. 3. Actions such as standing up or down are performed by actuating the motors in a safe reactive way to aid in such endeavour. For instance, in `doStandUp` the motors slightly pull forward so as to

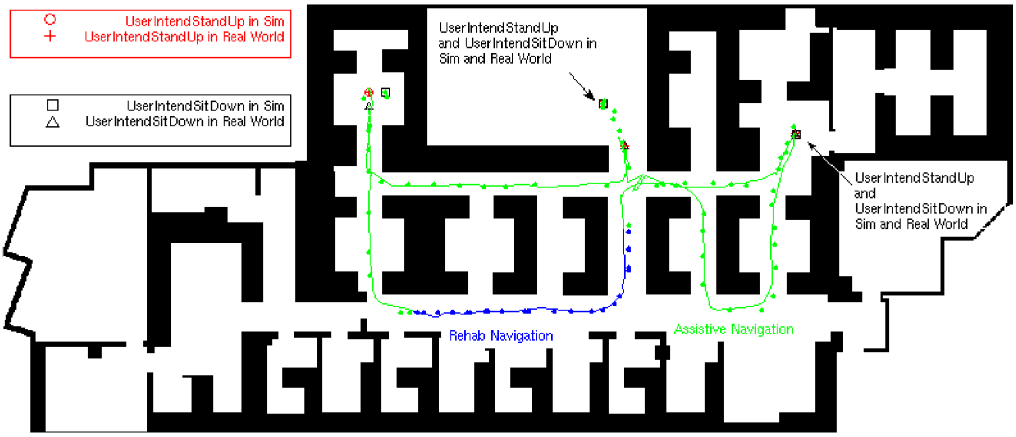

Fig. 4. A real and simulation example of the POMDP walker agent. The various tasks are depicted in different colours, and solid and dash lines refer to the navigation routines in the real and simulated experiments.

help the user to stand up. The navigational task represents an assistive ambulation function where the walker agent reacts to directional commands from the user via the strain gauges while the user strolls around. The rehabilitation task aims to substitute the labour-intensive task of an escort nurse when attending to the musculo-skeletal strength rehabilitation exercises often undertaken by elderly people in care facilities. The rationale is that when the user or the nurse take the walker to a specific location (A in Fig. 3), the walker starts an autonomous routine safely traversing between two given locations in the map, A and B for a predefined number of runs. The platform reacts safely to obstacles in the environment at all times, and is constantly monitoring new perceived indications from the user - such as the user deciding to quit and move somewhere else, or stop to talk to a friend en-route - and reacting accordingly. Further assistive tasks are planned for future work, particularly to exploit the wealth of information contained in the observations from the strain gauges.

One of the nice properties of Markov processes is their ability to formally fuse multi-sensorial information in their observation model probabilities, $Pr(o|s)$. As described in Section III, there is an array of instruments aboard the walker agent to provide diverse sensorial cues about the user, the platform itself, and its surroundings. From these, a vector of abstract discrete observation variables that the robot can make has been extracted to increase state observability and reduce the perceptual “aliasing”. For example, readings from the infra-red sensors pointing at the waist and torso of the user are fused with data from the switches to reduce ambiguity about the user trying to stand-up or sit-down while leaning on the walker robot. The increased robustness in the combination of these independent readings means a reduction in the size of the model, and consequently a better approximation to the exact solution. Our proposed model ended up with nine observation variables, fused to provide one of two discrete states (On/Off). Further, the walker agent localisation needed during the navigational tasks was assumed to be fully observable via the AMCL implementation of a Rao-Blackwell particle filter [23].

Given that the proposition is aimed at a pool of users with impaired ability to provide idealistic commands, a calibration is required to effectively learn the transition and observation functions that more closely reflect the error or uncertainties that can be expected in the user’s actions, as well as the ambiguities in subsequent observations due to sensor uncertainties and the inherent capacity of the user to properly operate the platform. In the work hereby presented, however, the training data has been restricted to a pool of able users, while awaiting ethical approval to carry out experiments on adequate subjects from an assistive living facility.

The reward function was defined such that actions that are proven to have been taken according to user’s intended activities are rewarded, whiles those direct actions that lead in the wrong direction are penalised.

V. EXPERIMENTAL RESULTS

The proposed algorithm has been evaluated within the domain of our office environment, a typical working space with desks, cubicles, people walking about, open meeting areas, corridors, etc. which can be thought of being structured similarly to a care facility for the elderly, possibly more challenging given the somewhat smaller dimensions. A 2D bird’s eye view of the map can be seen in Fig 4 overlaid with some of the results detailed next. Experiments were first simulated by manually providing the observations over typical runs, which demonstrated the viability of the decision making process itself, although they do not truly model the dynamics of the physical interactions and noises present in the world. We then asked one of the able user subjects to follow as closely as possible the same tasks in the real environment to qualitatively evaluate the effectiveness of the model under real world conditions. The results from one of these runs are depicted in Fig. 4, where coloured codes indicate the different actions taken, and solid and dash traces represent the (atemporal) trajectory followed by the walker agent (and user) during the real and simulated

TABLE I

SIMULATION OF POLICIES - MEAN REWARDS

Policy	POMDP	CE	MDP
Mean Reward	183.6	122.32	199.8

tests. The consonance in the results was also the norm in additional experiments carried out in the same context, where attempting the same observations in simulation and during the real tests (naturally nosier with real users) reproduced the same actions being taken by the proposed algorithm. Processing delays in evaluating the perceived intentions were in the order of milliseconds, hence negligible for practical purposes.

To further evaluate the performance of our proposed planning architecture more quantitatively, we compared the quality of the optimal POMDP policy against the average reward over time achieved in simulation trials by a fully observable MDP, which represents the unachievable upper-bound when there is (noise-free) perfect knowledge of the state. We have also compared against what is referred to as the certainty-equivalent (CE) policy in [18], a form of heuristic alternative which looks at the most likely state given the current belief and acts according to the policy derived for the MDP model. The results of running 1000 policy simulation trials, each for a maximum number of 100 steps, are collected in Table I. As expected, they indicate that the POMDP policy behaves more poorly than the ideal MDP upper bounds, yet the rewards that can be expected with the modelled uncertainties are close to that of the MDP. This means the policies obtained by the planner responding appropriately in the overall majority of cases within the modelled uncertainties that the system has to deal with, effectively achieving high rates of successfully "guessed" intentions. The CE policy appears far less rewarding as the belief state collapses to a single state, therefore always having to commit with little regard for sensor and action ambiguities in choosing the policy, which proves fatal is many instances.

VI. CONCLUSIONS AND FUTURE WORK

A cognitive agent capable of unobtrusively assist physically less able users with a set of mobility primitives has been presented. By making appropriate use of a range of sensorial inputs, an instrumented walker rollator platform has been proposed to assist in daily actions via a decision-making mechanism that requires minimal indicative input from the user, and potentially without the specific aid of a human caregiver. This is particularly relevant to provide gait stability support while strolling or during repetitive routines such as safe ambulation for physically strength rehabilitation, therefore increasing levels of confidence and independence. Simulation and real world experiments have shown that the decision-making agent can appropriately read and react to the user's intended behaviour, fulfilling the assistive nature of its role. Work in planned to automate the learning process and test the platform in an elderly care facility.

REFERENCES

[1] "Government action plan for health: Chronic and complex care statewide performance measures," NSW Health Department, AU, Tech. Rep., 2002.

[2] "Australian hospital statistics 2000-1," Australian Institute of Health & Welfare, AU, Tech. Rep., 2001.

[3] "Interactive national hospital morbidity data," Australian Institute of Health & Welfare, AU, Tech. Rep., 2000.

[4] N. A. R. Institute, *An Analysis of Research on Preventing Falls and Falls Injury in Older People: Community, Residential Aged Care and Acute Care Settings*. Canberra, AU: Commonwealth Department of Health and Aged Care, 2000.

[5] M. Begovich, "Arma - assisting robotic mobility aid," Umeå Institute of Design, Department of Advanced Product Design, Tech. Rep. MA2, 2005.

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

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

[8] S. Kang, Y. Ho, and I. H. Moon, "Development of an intelligent guide-stick for the blind," in *Proceedings of IEEE International Conference on Robotics and Automation*, May 2001, pp. 3208–3213.

[9] N. C. for Health Statistics, *Number of Persons Using Assistive Technology by Age and Person and Type of Device*, USA, 1994.

[10] F. Aminzadeh and N. Edwards, "Exploring seniors' views on the use of assistive devices in fall prevention." *Public Health Nursing*, vol. 15, no. 4, pp. 297–304, 1998.

[11] 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*, Arlington, VA, Nov. 2005.

[12] Y. Hirata, A. Hara, and K. Kosuge, "Motion control of passive intelligent walker using servo brakes," *IEEE Transactions on Robotics*, vol. 23, no. 5, pp. 981–990, 2007.

[13] A. J. Rentschler, R. A. Cooper, B. Blasch, and M. Boninger, "Intelligent walker for the elderly: Performance and safety testing of VA-PAMAID robotic walker," *Rehabilitation Research and Development*, vol. 40, no. 5, pp. 423–432, 2003.

[14] J. Pineau, M. Montemerlo, M. Pollack, N. Roy, and S. Thrun, "Towards robotic assistants in nursing homes: Challenges and results," *Robotics and Autonomous Systems*, vol. 42, no. 3-4, pp. 271–281, Mar. 2003.

[15] A. C. Morris, R. R. Donamukkala, A. Kapuria, A. Steinfeld, J. Matthews, J. Dunbar-Jacobs, and S. Thrun, "A robotic walker that provides guidance," in *Proceedings of the 2003 IEEE Conference on Robotics and Automation*, May 2003, pp. 25–30.

[16] S. Dubowsky, F. Genot, S. Godding, H. Kozono, A. Skwersky, H. Yu, and L. S. Yu, "PAMM - a robotic aid to the elderly for mobility assistance and monitoring," in *IEEE International Conference on Robotics and Automation*, 2000, pp. 570–576.

[17] B. Graf and M. Hgele, "Dependable interaction with an intelligent home care robot," in *In Proceedings of ICRA Workshop on Technical Challenge for Dependable Robots in Human Environments*, 2001.

[18] J. Hoey, A. V. Bertoldi, P. Poupart, and A. Mihailidis, "Assisting persons with dementia during handwashing using a partially observable markov decision process," in *Proceedings of the International Conference on Vision Systems*, Mar. 2007.

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

[20] N. Roy, G. Gordon, and S. Thrun, "Finding approximate pomdp solutions through belief compression," *Journal of Artificial Intelligence Research*, vol. 23, pp. 1–40, 2005.

[21] H. Kurniawati, D. Hsu, and W. Lee, "SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces," in *Proc. Robotics: Science and Systems*, 2008.

[22] M. E. López, , L. M. Bergasa, R. Barea, and M. S. Escudero, "A navigation system for assistant robots using visually augmented POMDPs," *Autonomous Robots*, vol. 19, no. 1, pp. 67–87, 2005.

[23] D. Fox and W. Burgard, "Markov localization for mobile robots in dynamic environments," *Journal of Artificial Intelligence Research*, vol. 11, pp. 391–427, 1999.