On the design and validation of an intelligent powered wheelchair: Lessons from the SmartWheeler project

Joelle Pineau, Amin Atrash, Robert Kaplow and Julien Villemure

Abstract New-generation, intelligent, powered wheelchairs promise to increase the mobility and freedom of individuals with serious chronic mobility impairments. And while rapid progress continues to be made in terms of the engineering capabilities of robotic wheelchairs, many projects fall short of the target in terms of ease of use, conviviality, and robustness. This paper describes the SmartWheeler, a multi-functional intelligent wheelchair, which leverages state-of-the-art probabilistic techniques for both autonomous navigation and user interaction modeling, to provide a novel robust solution to the problem of assistive mobility. We also discuss the use of standardized evaluation in the development and testing of such technology.

1 Introduction

Many people suffering from chronic mobility impairments use a powered wheelchair to help them move around their environment. However, there are many factors which may make the use of such wheelchairs difficult or impractical, including fatigue, muscle spasms, or sensory impairments, to name just a few. It has been reported that up to 40% of patients found daily steering and maneuvering tasks to be difficult or impossible [1]. Considering that there are 4.3 million users of power mobility in the US alone [2], the potential for intelligent solutions to assistive mobility is immense.

Several intelligent wheelchair platforms have been developed over the years. A thorough review of research in this area is available for interested readers [3]. Technical innovation on these projects usually targets one or more of three key aspects: autonomous navigation, safe interaction with the environment, and interaction between the user and the robotic wheelchair. In general, researchers have been most successful at advancing technology in the first two of these aspects, arguably be-

Joelle Pineau
McGill University, Montreal QC CANADA, e-mail: jpineau@cs.mcgill.ca

cause physical environments are easier to handle than humans. Yet one of the greatest challenges for the smooth integration of assistive mobility technology is in providing a usable interface between the technology and the person.

This paper provides an overview of the SmartWheeler platform, a multi-functional intelligent wheelchair designed to assist individuals with mobility impairments in their daily locomotion. The robotic components are integrated onboard a commercially available powered wheelchair. The main innovation of this project is in incorporating state-of-the-art artificial intelligence techniques to tackle both autonomous navigation of the wheelchair, and robust human-robot interaction.

While the long-term objective of our work is to increase the autonomy and safety of individuals in the context of their daily activities, the development of the platform to date has focused on satisfying a specific corpus of tasks, as defined by the Wheelchair Skills Test [4]. The use of such a well-defined set of tasks has many advantages for the objective evaluation of the robot platform; in particular, it allows a direct comparison with non-intelligent powered wheelchairs using an objective performance measure.

This paper provides an overview of the physical robot platform. We briefly describe the autonomous navigation system, and focus at more length on the human-robot interaction manager. Finally we describe the evaluation protocol and preliminary results. Formal experiments with the target population are still underway, and will be reported in later publications.

2 Wheelchair platform design

The current SmartWheeler platform, developed between 2006 and 2009 at the Center for Intelligent Machines (McGill University), is shown in Figure 1. The robotic components are integrated onboard a commercially available Sunrise Quickie Freestyle wheelchair. The robot senses obstacles in its environment through two (one forward-facing, one rear-facing) SICK laser range-finders. Robot positioning information is acquired through custom-made wheel odometers. Both of these components are essential for the autonomous navigation of the robot. Different sensors could be used (e.g. sensors, infrared), but the precision and reliability of the information is of the utmost importance for achieving robust navigation.

Communication between the intelligent wheelchair and the user occurs through a Lilliput 8” touch-sensitive LCD and a two-way voice interface. The voice input (microphone) modality is used primarily for acquiring user commands, whereas the display is used primarily to provide feedback to the user about the wheelchair’s understanding of the dialogue.

The robot platform also includes an onboard computer, which interfaces with the wheelchair’s motor control board to provide navigational commands. Wireless communication is also available onboard, though not required for the main functionalities of the platform ,as described below.

The robotic components listed above could easily be ported to standard commercial platforms other than the Sunrise Quickie Freestyle. In practice, the primary challenge of such an integration is in interfacing the onboard computer with the proprietary wheelchair controller.

Fig. 1 SmartWheeler robot platform.

Figure 2 displays an overview of the software architecture operating onboard the robot. The primary modules of interest are the Navigation Manager and the Interaction Manager. Sections 3 and 4, respectively, provide a detailed discussion of the main models and algorithms used in these two components.

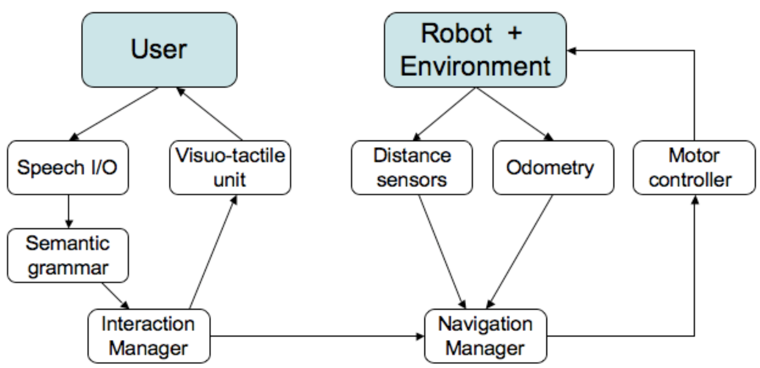

Fig. 2 SmartWheeler Architecture

3 Autonomous navigation

Core components of the autonomous robot navigation system include automatic map-building, localization within the environment, avoidance of obstacles (both stationary and moving), and automatic path planning. In recent years, statistical techniques have been used extensively to provide robust mathematical frameworks for these operations [5]. The SmartWheeler leverages most recent algorithmic techniques for performing these basic operations. The overall software architecture is written within the publicly available Player application [6]. Specific robot navigation routines, including mapping, localization, and low-level path planning and localization are implemented using the Carmen robot navigation toolkit [7].

A number of sophisticated statistical techniques–such as particle filtering—are included within the Carmen toolkit to perform state estimation based on the information acquired through the laser range-finders and odometry measurements. This allows the robot to build robust estimates of its environment, and its location within that space. The path planning methods included, however, are somewhat less robust, in that they assume the pose of the robot is known and do not take state uncertainty into account when making control decisions.

As part of the SmartWheeler project, we developed and implemented probabilistic models of decision-making, which are more robust for navigating in uncertain domains. The primary mathematical framework for this component is the Partially Observable Markov Decision Process (POMDP), which provides a stochastic model for sequential decision-making under uncertainty [8, 9]. One of the advantages of the POMDP paradigm is its ability to optimize plans contingent on partial state observability. In recent years, the development of efficient algorithms for POMDPs has allowed the framework to handle increasingly large domains [10]. Yet applicability for robot navigation remains elusive, due to the curse of dimensionality.

We recently developed a method for automatically generating a POMDP representation of an environment using variable resolution decomposition techniques. We can then apply state-of-the-art POMDP solution methods to optimize the action-selection policy of this model, such as to provide a near-optimal control strategy [11]. Using this method, the spatial discretization of the planning problem is generated automatically from any given metric map of the environment (sensor-built or hand-coded). As shown in Figure 3, this is done by taking advantage of the regular structure of indoor environments, such as identifying open spaces and abstracting them into a small number of states, while preserving significant resolution in areas near objects or walls. This results in a compact, yet accurate model of the environment. By applying POMDP planning methods using this representation, we can achieve robust planning under position uncertainty. Mapping uncertainty may also be included using more recently developed methods [12].

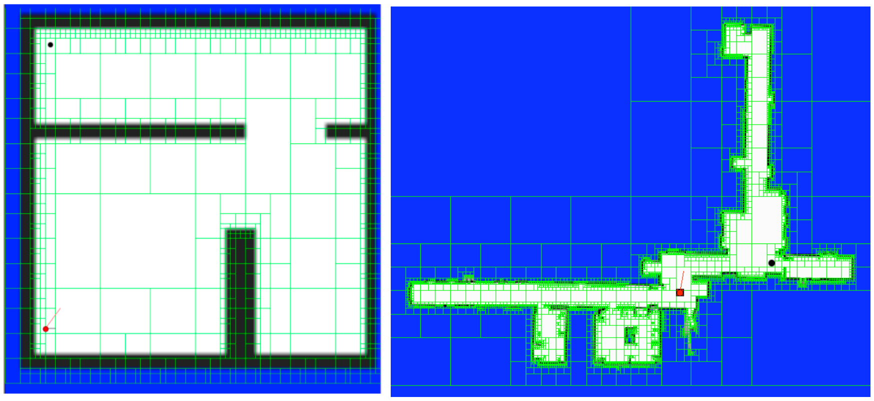

Fig. 3 Variable resolution decomposition of a simple robot navigation domain (left) and the first floor of the McConnell Engineering building (right). Light green lines indicate the state space representation considered for POMDP-based path planning.

4 User interaction

The Interaction Manager provides the mechanism for communication between the user and the wheelchair. Various communication modalities have been explored on-board smart wheelchairs, to supplement the traditional joystick input. Voice recognition is a popular choice as it requires little user training, as well as offering high information bandwidth at a low cost. Yet achieving reliable performance with a voice-driven wheelchair can be difficult due to high recognition error rates.

The goal of our system is to provide robustness to noise in the recognition system, as well as allow the user to speak naturally to the wheelchair, without needing to memorize a specific set of commands. To handle this, the SmartWheeler architecture includes a combination of intelligent systems techniques.

The speech recognition component is handled by Dragon Naturally Speaking, one of the more reliable commercially available software packages for speech to text transcription. We then process the output of the speech recognizer using a semantic grammar [13], which removes unnecessary lexical and syntactic information from the transcribed text. The final, and most important step, is the actual dialogue management. This module is responsible for the production of appropriate responses by the robot. There is no available software package for this step, and in fact the development of such systems is a highly active area of research, motivated largely by complex automated calling systems now in place in many organizations

In the context of our intelligent wheelchair project, it was imperative to have a dialogue management system which offers both robustness and flexibility, such as to maximize ease-of-use. To achieve robustness, we turned once again to the POMDP paradigm. In this context, the POMDP model is used by the Interaction Manager to select appropriate responses subject to the observations acquired through the speech interface. Here again, the ability for the POMDP model to reason about uncertainty in the state (in this case the user's spoken word and intention) is crucial to the robustness of the system. In particular, the POMDP model can suggest clarification questions whenever the input received is incomplete or ambiguous. Figure 4 shows a short dialogue between the intelligent wheelchair and a test subject.

User command	Speech recognition output	Robot action
"Turn right."	`turn right`	Turns right 90 deg.
"Stop."	`this`	*Please repeat your command.*
"Turn left a little."	`turn left little`	Turn left 22.5 deg.
"Drive forward."	`drive forward`	Drive forward (max. 100m).
"Follow right wall."	`right wall`	Drive along the right wall (distance to wall = 0.65m).
"Turn around."	`turn around`	Turn to the left 180 deg.
"Follow left wall."	`left all`	*Please repeat your command.*
"Follow left wall."	`fall left wall`	Drive along the left wall (distance to wall = 0.65m).
"Stop."	`stop`	Stop

Fig. 4 Sample dialogue between a test subject and the intelligent wheelchair. The first column shows the user's actual spoken words. The second column reports the output of the automated speech recognition module. The third column shows the action chosen by the POMDP-based Interaction Manager (entries in *italic* represent clarification queries; all other action choices are transmitted to the Navigation Manager.)

While the POMDP offers the appropriate mathematical tools for achieving robust interaction, it can be quite inflexible. In particular, standard POMDP planning algorithms require a (mathematically) accurate predictive model of the dynamics of the conversation. In general, it is challenging to define such a model *a priori*, as human behavior is difficult to quantify, and varies substantially between individuals. To address this issue, we leveraged machine learning techniques to build the model directly from observed data [14]. Figure 5 shows the effect of learning on the performance of the Interaction Manager; these results were compiled using a simulated user .

5 Validation methodology

As outlined by Simpson in his review paper, there have been substantial efforts directed towards the development of smart wheelchairs, yet much less attention has been devoted to the evaluation of these prototypes [3]. The evaluation methodology for the SmartWheeler was developed in close consultation with rehabilitation researchers and clinicians, with the goal of having an evaluation protocol that met a number of criteria:

Fig. 5 Improvement in the quality of actions selected by the Interaction Manager as a function of the observed training data.

- it could be undertaken by users with disabilities,
- it did not require extensive user training or customization,
- it allowed comparison between standard and intelligent powered wheelchairs,
- it was sufficiently repeatable to allow aggregation of user results,
- it admitted an objective performance measure.

We settled on using the Wheelchair Skills Test, developed by researchers at Dalhousie University [4]. This testing procedure consists of a set corpus of motion tasks, which the user must accomplish safely and to completion. It was originally designed for objectively evaluating manual wheelchair skills, then extended to powered wheelchairs. It has not previously been used to evaluate intelligent wheelchairs, yet applies readily without modification. In the case of the SmartWheeler, the main difference is that the user controls the chair through vocal (rather than manual) commands, and the internal controller is responsible for the low-level navigation.

The test includes tasks with a wide range of difficulty. We focus on the subset of skills which are relevant for robot-assisted mobility. Figure 6 shows an experimenter undergoing testing for some of the skills included in the test. Based on the selected set of tasks, we specified the vocabulary and grammar necessary for the speech interface. Neither of these are so restrictive as to require the user to use specific phrases. The initial vocabulary contained approximately 60 keywords which triggered appropriate grammatical rules and served as the input set for the Interaction Manager. The vocabulary itself, and the probability distributions describing the expected frequency of specific words for certain tasks were extended in the course of the user experiments using simple statistical techniques. Individual customization was not applied, given that each user only ran the test a single time. Rather, the learning was used to improve the model for the full population.

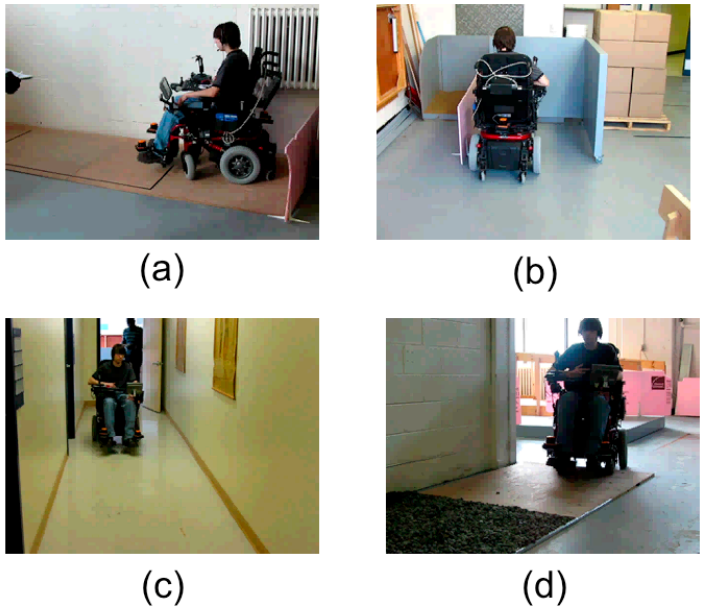

Fig. 6 Various stations of the Wheelchair Skills Test. The wheelchair must (a) travel along a sloped platform; (b) be aligned to the left wall; (c) move forward through a door; (d) travel through increased rolling resistance (in this case, gravel).

6 Results and discussion

A preliminary evaluation of the Interaction Manager involving seven healthy subjects, all of them university students without involvement in the project, was performed early on [15]. These results provided an initial set of data for the model customization. A second round of experiments involving eight healthy subjects, all of them clinicians in local rehabilitation centers but without involvement in the project, was performed more recently. These experiments were performed on a different robotic platform developed at Ecole Polytechnique de Montréal [16]; this second platform features substantial differences from the SmartWheeler in terms of hardware and autonomous navigation software, however the user interaction modules are the same. Results of these experiments are recorded in Figure 7. A third round of experiments involving eight subjects with mobility impairments is currently underway (also using the second robot).

As shown in Figure 7, the Interaction Manager provides a robust architecture for handling communication with the user. Users were able to complete the test using between 114 and 219 commands. The word error rate for some subjects (subjects 4 and 8) was quite high. However, the appropriate use of queries allowed the system to reach a performance level comparable to that of other users, as shown by the low incidence of incorrect actions.

Subject id	Number of commands	Word error rate	Number of queries	Number of correct actions	Number of incorrect actions
1	136	8.8%	10	121	5 (3.7%)
2	159	13.8%	18	136	5 (3.1%)
3	165	13.5%	11	152	2 (1.2%)
4	201	23.6%	37	155	9 (4.5%)
5	114	6.2%	13	97	4 (3.5%)
6	219	2.3%	10	208	1 (0.5%)
7	210	13.1%	25	175	10 (4.8%)
8	141	19.3%	26	111	4 (2.8%)

Fig. 7 Performance of the Interaction Manager for the Wheelchair Skills Test. The second column shows the number of vocal commands issued by the user throughout the test. The third column reports the raw speech recognition error rate. The fourth column shows the number of clarification queries issued by the robot in cases where the command was misunderstood or ambiguous. The fifth column presents the number of correct actions carried by the robot, as identified by human labeling of video sequences. Finally, the last column reports the number of times the robot selected an incorrect actions; users were instructed to recover from such situations by issuing a *Stop* command, or starting a new command.

Overall, the test subjects were satisfied by the functionality of the interface and appreciated the visual feedback capabilities. Some subjects felt they needed more time to become familiar with the system to exploit it more successfully. Training time for all subjects was on the order of 30 minutes. The system was judged to be sufficiently usable and robust to move forward with experiments involving the target population.

The current experimental protocol is quite constrained, both in time and type of tasks evaluated. This is useful to allow formal testing with a number of subjects. But it has important limitations in terms of evaluating the long-term usability of the system. However, a number of challenges remain before the intelligent wheelchair is ready for deployment in natural living environments. From a practical engineering perspective, it is necessary to extend the intelligent wheelchair's knowledge of the world. This means acquiring maps of the environment such as to allow navigation over a much larger footprint. It also means extending the vocabulary and grammar to accommodate a larger variety of commands. We foresee a number of important technological challenges in carrying out these extensions. First, there is the issue of computational scalability of the models and algorithms. Some of the inference techniques used by both the Navigation and Interaction component will require more efficient approximation to handle larger dimensional domains. However the most important challenge will likely be to develop models and algorithms which allow life-long learning, so that an intelligent wheelchair can automatically adapt to new environments, new habits, and new activities, along with its passenger.

Acknowledgements The project involved close collaboration with researchers at the Ecole Polytechnique de Montreal (Paul Cohen, Sousso Kelouwani, Hai Nguyen, Patrice Boucher), at the Universite de Montreal (Robert Forget, Wormser Honore, Louise Demers), at the Centre de readaptation Lucie Bruneau (Claude Dufour) and at the Constance-Lethbridge Rehabilitation Centre (Paula Stone, Daniel Rock, Jean-Paul Dussault). Financial support was provided through the Canadian Foundation for Innovation, the Natural Sciences and Engineering Council of Canada, and the Fonds quebecois de la recherche sur la nature et les technologies. Additional funding was provided by the Fondation du Centre de readaptation Lucie Bruneau, the Fondation Constance Lethbridge, and Robovic. Materials were donated by Sunrise Medical Canada and LiPPERT.

References

1. L. Fehr, E. Langbein, and S. B. Skaar. Adequacy of power wheelchair control interfaces for persons with severe disabilities: A clinical survey. *J. of Rehabilitation Research & Development*, 37:353–360, 2000.
2. R. C. Simpson, E. F. LoPresti, and R. A. Cooper. How many people would benefit from a smart wheelchair? *J. of Rehabilitation Research & Development*, 45:53–72, 2008.
3. R. C. Simpson. Smart wheelchairs: A literature review. *J. of Rehabilitation Research & Development*, 42:423–436, 2005.
4. F. Routhier, C. Vincent, J. Desrosiers, S. Nadeau, and C. Guerette. Development of an obstacle course assessment of wheelchair user performance (OCAWUP). *Technology & Disability*, 16:19–31, 2004.
5. S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
6. T. H. J. Collett, B. A. MacDonald, and B. P. Gerkey. Player 2.0: Toward a practical robot programming framework. In *Proceedings of the Australasian Conference on Robotics and Automation*, 2005.
7. M. Montemerlo, N. Roy, and S. Thrun. Perspectives on standarization in mobile robot programming: The Carnegie Mellon Navigation (CARMEN) Toolkit. In *International Conference on Robotics and Systems*, pages 2436–2441, 2003.
8. L.P. Kaelbling, M.L. Littman, and A.R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.
9. E.J. Sondik. *The optimal control of partially observable Markov processes*. PhD thesis, Stanford University, 1971.
10. J. Pineau, G. Gordon, and S. Thrun. Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research*, 27:335–380, 2006.
11. R. Kaplow, A. Atrash, and J. Pineau. Variable resolution decomposition for robotic navigation under a POMDP framework. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.
12. A. Guez and J. Pineau. Multi-tasking SLAM. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2010.
13. M. White. Open ccg: The open nlp ccg library. Software available at http://openccg.sourceforge.net, 2001.
14. A. Atrash and J. Pineau. A bayesian reinforcement learning approach for customizing human-robot interfaces. In *International Conference on Intelligent User Interfaces (IUI)*, 2009.
15. A. Atrash, R. Kaplow, J. Villemure, R. West, H. Yamani, and J. Pineau. Towards the deployment of an intelligent wheelchair in a standardized rehabilitation environment. *Interaction Studies*, 2009.
16. W. Honoré, A. Atrash, P. Boucher, R. Kaplow, S. Kalouwani, H. Nguyen, J. Villemure, R. West, F. Routhier, P. Stone, C. Dufour, J.-P. Dussault, D. Rock, P. Cohen, L. Demers, R. Forget, and J. Pineau. Human-oriented design and initial validation of an intelligent powered wheelchair. In *RESNA Annual Conference*, 2010.