SmartWheeler: A robotic wheelchair test-bed for investigating new models of human-robot interaction

Joelle Pineau
for the SmartWheeler team[1]
School of Computer Science
McGill University
Montreal, QC H3A 2A7
jpineau@cs.mcgill.ca

Abstract

The goal of the SmartWheeler project is to increase the autonomy and safety of individuals with severe mobility impairments by developing a robotic wheelchair that is adapted to their needs. The platform we have built also serves as a test-bed for validating novel concepts and algorithms for automated decision-making onboard socially assistive robots.

Introduction

Many people who suffer from chronic mobility impairments, such as spinal cord injuries or multiple sclerosis, use a powered wheelchair to move around their environment. However, factors such as fatigue, degeneration of their condition, and sensory impairments, often limit their ability to use standard electric wheelchairs.

The SmartWheeler project aims at developing—in collaboration with engineers and rehabilitation clinicians—a prototype of a multi-functional intelligent wheelchair to assist individuals with mobility impairments in their daily locomotion, while minimizing physical and cognitive loads.

Many challenging issues arise in this type of application. First, there are a number of technical issues pertaining to the physical design of the wheelchair; these are only briefly mentioned below. Second there are substantial computation issues pertaining to the control of the wheelchair which require close attention. This paper outlines ongoing work targeting a number of these aspects, ranging from new approaches to path planning, to technical innovations for model learning, to the design of the human-robot control interface.

Beyond its technological components, an essential aspect of this project is a strong collaboration with clinicians, to ensure the definition of goals for the mobility functions, for the patient/wheelchair and environment/wheelchair interactions, as well as for the experimental validation of the smart wheelchair.

Our aim is to show that the robotic wheelchair reduces the physical and cognitive load required to operate the vehicle. We are therefore focusing on high-load situations, such as navigating in confined spaces (e.g. entering/exiting an elevator or a public washroom), stressful situations (e.g. exiting a building during a fire alarm), or unknown environments (e.g. transferring flights through a new airport).

Most of these tasks require basic robot navigation capabilities (mapping, localization, point-to-point motion). We make few novel contributions in this area, and rely on previous technology to implement these functionalities. Instead, the scientific objective of this project is to develop new models for the robust control of this complex interactive robot system. We are also concerned with the design of the physical interface between the robotic wheelchair and its user. A key aspect of the patient/wheelchair interface involves creating communication protocols that can ensure the quality, reduce the ambiguity, and progressively improve the effictivenes of these interactions. We investigate two such protocols: a voice-based dialogue system, and a tactile/visual interface system. Both are discussed below.

Target population

The goal of this project is to increase the autonomy and safety of individuals with severe mobility impairments by developing a robotic wheelchair that is adapted to their needs.

Through discussions with clinical collaborators at the Centre de réadaptation Constance-Lethbridge and Centre de réadaptation Lucie-Bruneau, two rehabilitation clinics in the Montreal area, we have selected a target population for this work, along with a set of challenging tasks. We have chosen to define the target population based on their *abilities*, rather than their *pathologies*. The motivation for doing so is that we can use uniform measures of performance across the target population, thereby allowing us to gauge the usefulness of the deployed robotic system.

Individuals of interest will be those who meet the reduced mobility criteria necessary to qualify for a powered wheelchair under the Régie de l'assurance maladie du Québec (the provincial public health board). There are well established guidelines for applying this criteria, and our clinical collaborators have long expertise in evaluating these.

Robot platform

SmartWheeler, shown in Figure 1, is built on top of a commercially available Sunrise Quickie Freestyle, to which we

---

[1] The SmartWheeler team includes: Amin Atrash, Jeremy Cooperstock, Robin Jaulmes, Robert Kaplow, Nan Lin, Andrew Phan, Chris Prahacs, Doina Precup and Shane Saunderson.

have added front and back laser range-finders, wheel odometers, a touch-sensitive graphical display, a voice interface, and an onboard computer. Additional devices will be integrated in the future, including stereo vision, IR sensors, and a modified joystick. All hardware and electronic design were performed in-house by staff members at McGill's Center for Intelligent Machines.

Figure 1: SmartWheeler robot platform.

The robot's basic mapping and navigation functions are provided by the Carmen robot navigation toolkit (Montemerlo et al., 2003). This toolkit can be adapted to a variety of robot platforms and has been used in the robotics community for the control of indoor mobile robots in a variety of challenging environments. It is particularly useful for validation of new algorithms because its simulator is known to be highly reliable and policies with good simulation performance can typically be ported without modification to the corresponding robot platform.

New capabilities which are developing include:

- Detection and avoidance of negative obstacles (e.g. downward staircase).
- Robust point-to-point planning and navigation.
- Shared control between autonomous controller and human user.
- Adaptive (user-specific) control strategy.
- Adapted interface for low-bandwidth communication.

The remainder of the paper discusses four ongoing areas of research pertaining to this project.

Adaptive planning in large-scale environments

As part of its task domain, the robot will be called upon to navigate robustly in very large environments. There exists a number of well known approaches for robot path planning, however they tend to roughly fall into two camps. The first group assumes deterministic effects on the part of both the robot and the environment; it can therefore scale to high-dimensional domains but is not robust to uncertainty in the motion or sensor model. An example of such algorithm is the variable resolution cell decomposition technique. The second group considers probabilistic motion effects and sensor readings, and is therefore robust to uncertainty, but generally scales poorly and can only handle small environments. An example of such algorithm is the Partially Observable Markov Decision Process (POMDP) framework.

We are developing a new approach to planning in metric environments with unifies the variable resolution cell decomposition and POMDP approaches. The variable resolution techniques allow us to select the appropriate state representation for the environment. We then infer a probabilistic motion model over this state representation by using sampling techniques. Finally, we apply POMDP solving techniques to extract an action-selection policy.

Figure 2 shows a sample map, as well as the state representation that was extracted by our approach. In this case, POMDP planning was accelerated by a factor of 400 compared to standard approaches.

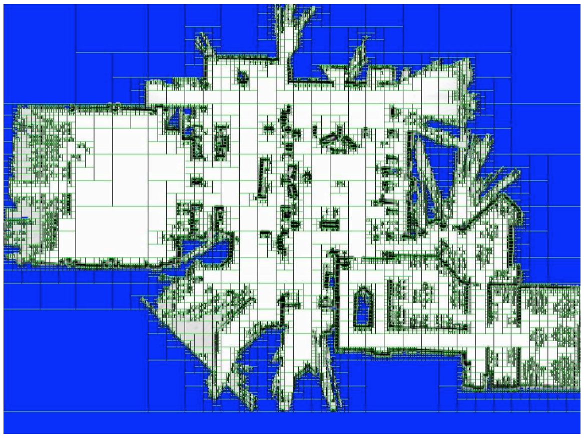

Figure 2: Variable resolution map

Learning and control under model uncertainty

The POMDP framework requires a known parametric model defining the dynamics of the problem domain. In the section above, we assume that the model is learned by acquiring motion samples and building a parametric representation. This is feasible in cases where data is inexpensive, as is the case when a simulator is available. Alternately, the parametric model could be specified by a domain expert, which again is feasible in some domains, for example in the physical world where there are natural constraints. In many domains however, for example human-robot interaction, acquiring an accurate model can be challenging.

To address this problem, we are considering ways of formulating POMDP solutions such that we take into account model uncertainty, and improve model accuracy through careful selection of queries. The framework we propose relies on a Bayesian formulation of uncertainty, which is particularly appropriate to offer a flexible trade-off between a priori knowledge engineering and data-driven parameter inference. Further information on this component is available in (Jaulmes et al., 2005a; Jaulmes et al., 2005b).

Large-scale dialogue management

A natural medium for communication between a user and an intelligent system is through voice commands. While many commercial software solutions are available for speech recognition and synthesis, there is no commercial equivalent for handling the actual dialogue (i.e. production of responses by the robot). In fact, in a spoken dialogue system, determining which action the robot (or computer) should take in a given situation is a difficult problem due to the uncertainty that characterizes human communication.

Earlier work pioneered the idea of POMDP-based dialogue managers, but were limited to small domains (e.g. 2-3 topics in a question-answer format). We are currently investigating techniques for tracking the dialogue state and efficiently selecting dialogue actions in domains with large observation spaces. In particular, we study the applicability of two well-known classes of data summarization techniques to this problem. In (Atrash and Pineau, 2005), we proposed a clustering algorithm which, through a series of EM iterations, finds a small set of summary observations. To provide a comparison against the clustering algorithm, we then presented a dimensionality reduction algorithm along the lines of Principal Component Analysis (with a few added constraints) which finds a compressed version of the observation probability model. Preliminary experiments with the SACTI dialogue corpus (Willians and Young, 2004) suggested that the simple EM-type clustering works well, even in this complex dialogue domain. This is encouraging because the clustering algorithm is simple to implement, fast to compute, and generates intuitive compressed representations. We also found that a constrained-based PCA performed on par with EM-clustering on this 450 word domain, however computation was significantly slower. We are now extending the results to topics relevant to the robot domain outlined above.

Interface design

The last component aims at designing and validating a new communication protocol for allowing users to provide high-level navigation commands to the wheelchair. Conventional control of a motorized wheelchair is typically done through a joystick device. For those unable to operate a standard joystick, alternatives include sip-and-puff devices, pressure sensors, etc. Regardless of the device used, the user input set is restricted to displacement and velocity commands. Operation of a wheelchair in this manner can result in fatigue over time, as well it is often difficult to manoeuver the wheelchair in constrained spaces (e.g. elevators, crowded rooms, etc). The prototype robotic wheelchair we are developing seeks to alleviate these challenges by allowing the users to specify high-level navigation goals (e.g. *Go to room 103.*) This requires a new communication protocol which will allow the user to input such commands.

The communication protocol we propose inputs high-level navigation goal is using EdgeWrite (Wobbrock and Myers, 2006), a unistroke text entry method for handheld devices, designed to provide high accuracy text entry for people with motor impairments. We have adapted this method for the control of a motorized wheelchair by customizing the set of strokes, input constraints, and feedback display to the task of wheelchair control.

User experiments currently under way are comparing entry of robot navigation goals using: direct map selection, menu selection, and EdgeWrite gesture entry. For each input modality, the user is shown a floor map of a building on the screen, and guided through a list of locations that must be selected quickly and accurately using the different input selection methods. We measure error rate, input time, and motion time needed to reach the target location. Early results with a control population indicate that menu selection (from a static vertical list) was twice as fast as selecting the targets directly on the map and three times as fast as entering the corresponding EdgeWrite symbol, which is as expected for this population. We are now replicating the experiment with disabled users. Since this population has significant motor constraints, we may obtain significantly different results regarding the preferred mode of input.

Discussion

This short paper highlights some of the key components of the SmartWheeler project. We are currently working on their integration and planning out a sequence of experiments with the target population. Through close collaborations with engineers and rehabilitation researchers, we hope to one day have a positive impact on the quality of life for individuals with severe mobility impairments.

It is worth noting that many of the techniques developed in this project are not specific to the mobility-impaired population, but are relevant to building service robots for a large number of applications. The SmartWheeler platform is proving to be an exciting new test-bed for exploring novel concepts and approaches in automated decision-making, human-robot interaction, and assistive robotics.

References

Atrash, A. and Pineau, J. ”Efficient Planning and Tracking in POMDPs with Large Observation Spaces”. AAAI-06 Workshop on Empirical and Statistical Approaches for Spoken Dialogue Systems. 2006.

Jaulmes, R., Pineau, J. and Precup, D. ”Active Learning in Partially Observable Markov Decision Processes”. European Conference on Machine Learning (ECML). LNAI 3729. Springer-Verlag. 601-608. 2005.

Jaulmes, R., Pineau, J. and Precup, D.. ”Probabilistic Robot Planning Under Model Uncertainty: An Active Learning Approach”. NIPS Workshop on Machine Learning Based Robotics in Unstructured Environments. 2005.

Montemerlo, M., Roy N. and Thrun, S. “Perspectives on Standardization in Mobile Robot Programming: The Carnegie Mellon Navigation (CARMEN) Toolkit”. IROS, 2003.

Williams, J.D. and Young, S. ”Characterizing task-oriented dialog using a simulated ASR channel. International Conference on Speech and Language Processing (ICSLP). 2004.

Wobbrock, J.O. and Myers, B.A. ”Trackball text entry for people with motor impairments”. ACM Conference on Human Factors in Computing Systems (CHI). 2006.