Towards a Standardized Test for Intelligent Wheelchairs

Joelle Pineau*
Centre for Intelligent Machines
McGill University
Montreal, QC CANADA
jpineau@cs.mcgill.ca

Robert West
Centre for Intelligent Machines
McGill University
Montreal, QC CANADA
robert.west@mail.mcgill.ca

Amin Atrash
Centre for Intelligent Machines
McGill University
Montreal, QC CANADA
aatras@cs.mcgill.ca

Julien Villemure
Centre for Intelligent Machines
McGill University
Montreal, QC CANADA
julien@cs.mcgill.ca

François Routhier
Institut de réadaptation en
déficience physique de
Québec
Québec, QC CANADA
Francois.Routhier@rea.ulaval.ca

ABSTRACT

Many people who have to rely on electric wheelchairs find it hard or even impossible to fulfill daily navigation tasks with their chairs. The SmartWheeler project aims at developing an intelligent wheelchair that minimizes the physical and cognitive load required in steering it. In this paper we briefly outline the SmartWheeler project and its goals. We then argue that it is important to have a standardized test to evaluate autonomous wheelchairs in terms of performance quality, safety, and usability. No such test exists as yet for intelligent wheelchairs, but there has been an effort in the clinical community to design tests for conventional wheelchair usage. We discuss the existing Wheelchair Skills Test (WST). We then suggest a paradigm that allows us to use this test to benchmark the quality of intelligent wheelchairs, and in particular their interface, in a task context that is relevant to clinical practice in rehabilitation.

1. 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. According to a recent survey, 40% of powered wheelchair users surveyed found daily steering and maneuvering tasks to be difficult or impossible [2]; and according to the clinicians who treat them, nearly half of those patients unable to control a powered wheelchair by conventional methods would benefit from an automated navigation system [2].

*Corresponding author.

PerMIS '10, September 28-30, 2010, Baltimore, MD, USA

Such numbers make it seem likely that intelligent wheelchairs catering to those patients' needs would have a deep societal impact. One might argue that the transition to wheelchairs that cooperate with the user is at least as important as that from manual to electric wheelchairs—possibly even more important since this would mark a paradigmatic rather than merely a technological shift. However, to the best of our knowledge, there is no general method of evaluating the performance of intelligent wheelchairs yet [11]. And in particular, no formal tools exist to evaluate the safety and effectiveness of the interaction between the intelligent wheelchair and its operator.

In this paper we try to make a first step by suggesting a methodology based on work done in the clinical rehabilitation community. In particular, we investigate the use of a specific corpus of tasks, as defined by the Wheelchair Skills Test (WST) [10]. The use of such a well-defined set of tasks has many advantages for the objective evaluation of the intelligent wheelchairs. It ensures the evaluation criteria is relevant to the end-user (since the task domain was originally defined for standard powered wheelchair users), it provides a repeatable evaluation protocol between test subjects, and it admits an objective performance measure.

We first describe the SmartWheeler project and the intelligent wheelchair developed by our research team. The rationale that supports the use of a standardized test and the relevant literature are exposed. The WST is then described, followed by results pertaining to the evaluation of the human-robot interface component of our platform. Finally, future perspectives are presented.

2. THE SMARTWHEELER PROJECT

The SmartWheeler project [1, 8] 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.

Figure 1 shows a picture of the SmartWheeler platform (built on top of a commercially available Sunrise Quickie Freestyle which was extended in-house at McGill's Centre for Intelligent Machines).

Most of the software components governing the autonomous navigation are being developed by some of our collabora-

tors [3]. The main contribution of the authors is in the development and validation of the human-robot interface. Figure 2 presents an overview of the software architecture controlling the human-robot interface onboard the robot. The primary mode of interaction is a two-way speech interface. We employ a number of technologies to achieve robust interaction, including natural language processing (automatic speech recognition and grammatical parsing) and high-level dialogue management using Partially Observable Markov Decision Processes. A tactile/visual interface system is also installed, and used primarily for provided visual feedback to the human regarding the state of the dialogue system.

Figure 1: The SmartWheeler robot platform.

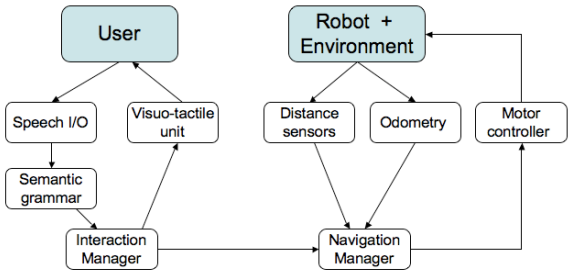

Figure 2: The SmartWheeler Interaction Architecture.

Speech provides a natural interface for human operators. Yet it is subject to significant failure rates due to the noise and ambiguity inherent in speech-based communication. Both the choice of tasks and physical environment can further affect the performance of a automated dialogue system. Thus it is imperative that we be able to carefully quantify the performance of the speech-based interface in the context of natural interactions and in a realistic environment.

3. REASONS FOR A STANDARDIZED TEST

All engineered research needs to be assessed in terms of the results it produces. Quantifying the efficacy of a robotic device designed to aid people is also a necessary step in the evaluation of its impact. A machine will only be accepted by people if it is of use to them. In this section we list the major reasons we see for adopting a standardized test for intelligent robotic wheelchairs.

In a recent review of intelligent wheelchair projects Simpson concludes that, "[while] there has been a significant amount of effort devoted to the development of smart wheelchairs, scant attention has been paid to evaluating their performance. [...] Furthermore, no smart wheelchair has been subjected to a rigorous, controlled evaluation that involves extended use in real-world settings." [11]

However, such a "rigorous, controlled evaluation" is essential particularly in the context of health-related projects like SmartWheeler. Performance and safety requirements for wheelchairs are high, since users rely heavily on the device. There must be a rigorous way of proving that these requirements are met before a wheelchair can be deployed. This is especially true for *intelligent* wheelchairs, which will eventually act at least partly in an autonomous manner. As more control is taken from the user and given to the wheelchair, it becomes more important to make guarantees about its performance. Certainly a standard evaluation scheme is also a crucial step if the use of intelligent wheelchairs is to be funded by public health services and insurance companies.

Also, one generally strives to supply a person with the wheelchair that fits them best. This is true for regular wheelchairs, and it applies equally to intelligent wheelchairs. For instance, certain features (e.g. an eye tracker) might be expensive, so one would like to dispense with them if they are not necessary. A standardized test might help figure out the best configuration for a user. Again, this will be essential for funding purposes.

Moreover, as more projects of the kind described above come into being it will be helpful to benchmark the efficacy of the technologies employed. On the one hand this can serve to assess how well the algorithms and hardware being developed within one research project work in a setting that is close to the real world, which can guide researchers towards the problems that have to be addressed next. On the other hand a standardized test facilitates the comparison of similar projects by different research teams, thus highlighting the most promising approaches.

Finally, from a practical point of view, a standardized test can be helpful during the development process because it makes work more target-driven. Keeping the test in mind can help the research team get 'boot-strapped' by providing a useful basis for thinking of possible deployment scenarios. For instance, a first English grammar for the natural language understanding component of a voice recognition system could cover the set of commands that represent the skills required in a standardized test.

4. RELEVANT LITERATURE

If a test is to be used for the reasons just listed it should be valid and reliable from a clinical point of view, i.e. it should actually measure what it is intended to measure, and do this in a reproducible way. Designing such a test can be difficult for a computer scientist or engineer lacking the necessary background in clinical rehabilitation. Fortunately the rehabilitation literature offers many possibilities.

Several wheelchair skills tests have been proposed in the literature, and Kilkens *et al.* [4] and Routhier *et al.* [9] fairly recently provided the first systematic overviews. In this section we will briefly summarize their results. Later we will describe the test we chose for evaluating the SmartWheeler

project and how we are planning to use it.

Both Kilkens *et al.* [4] and Routhier *et al.* [9] come to the conclusion that no standard test to measure wheelchair skill performance exists as yet, despite a considerable clinical and academic need for such a measure. From a clinical point of view, a standard test should allow for extrapolation of test results to assess subjects' everyday wheelchair performance, in order to guide training and facilitate the selection of a suited wheelchair. From an academic perspective, a standard test would alleviate the current difficulty in comparing study results due to the lack of a common benchmark. As a first step towards standardization, both articles give surveys about existing non-standard wheelchair tests.

Kilkens *et al.* [4] conclude that, while more research is needed to identify the skills to be included in a standard test, out of the 24 tests they reviewed only the Wheelchair Skills Test (henceforth WST) has been "adequately tested on both validity and reliability" [4] (for the results of the evaluation of the WST see [6]). Note that, although Kilkens *et al.* center their discussion on manual chairs, the WST happens to be conceived for powered wheelchairs as well.

The article by Routhier *et al.* [9] is slightly more general in that it considers tests for manual as well as powered wheelchairs and reviews not only controlled environments (as Kilkens *et al.* [4] do) but also distinguishes between three categories of test environments:

1. Real environments (observing subjects' daily wheelchair activities).
2. Controlled environments (e.g. obstacle courses).
3. Virtual environments (using a simulator).

Routhier *et al.* [9] recommend the controlled-environment para-digm. It is interesting to note that they have recently abandoned the design of their own test [10] in favor of the WST, which seems to become the 'gold standard' in the clinic and research communities, being deployed by many institutions across North America. This is due to the aforementioned reason that it is the only test that has been rigorously checked for validity and reliability in statistical terms. If there is to be a standard test for wheelchair skills in the future it seems that it will most likely be the WST.

Another reason that makes the WST particularly appropriate is that, unlike many other tests, it has not been designed for a specific target group (e.g. stroke patients) but for wheelchair users in general (manual and powered). This is important if the intelligent wheelchair shall serve as an aid to more than just a fraction of patients.

5. THE WHEELCHAIR SKILLS TEST

The WST, currently in version 4.1 [5], is being developed as part of the Wheelchair Skills Program (WSP) at Dalhousie University in Halifax, Canada, as a "standardized evaluation method that permits a set of representative wheelchair skills to be objectively, simply and inexpensively documented." [5] Extensive information about the test can be found at the WSP website (`www.wheelchairskillsprogram.ca`). The creators envision several situations in which to apply the WST:

1. In the early rehabilitation process it can serve to identify the skills that should be addressed during training.
2. It can serve as an outcome measure to compare a subject's performance before and after rehabilitation.
3. It can be used to test research hypotheses and to assist engineers in the development of new technologies.

Since the WST strives to be as general as possible, it specifies four test categories, one for each combination of wheelchair type (manual vs. powered) and test subject (wheelchair user alone vs. wheelchair user with caregiver). Some of the tasks do not apply to all of the four categories (e.g. 'Picks object from floor' is not applicable if a caregiver is present, since it is assumed that the latter rather than the wheelchair user will do this when the situation arises). For our use of the WST it is crucial to note that it is also explicitly conceived to provide a means of evaluating caregivers. Our goal is not to rate the performance of wheelchair users but that of the intelligent control system. We will do so by considering it a caregiver: Like a caregiver, the software cooperates with the wheelchair user in order to help him/her master everyday situations.

Tasks covered

The powered wheelchair version of the WST (WST-P) test covers 32 skills which are considered representative for general wheelchair performance. The assumption is that a person doing well (performance and safety) on the 32 tasks included in the WST can be considered a skilled wheelchair user because the situations he/she encounters on a daily basis will resemble those tested. In other words, the WST abstracts from a real-world setting to measurable wheelchair skills. It is based on realistic scenarios but is still standardized enough to allow for precise performance measurements. As one would expect, most tasks test navigation skills (e.g. 'Rolls forward 10 m in 30 s', 'Gets over 15-cm pot-hole'), but there are some other actions as well, e.g. those concerning the wheelchair configuration, like 'Controls recline function'. Figure 3 shows an experimenter undergoing some of the skills included in the test. One pass over all tasks takes about 30 minutes [6].

Evaluation method

The test evaluates skill performance and safety. Each skill is graded in terms of these two criteria in a binary manner: a person either passes or fails a task, and he/she does so either in a safe or in an unsafe way. The overall score consists of two numbers, which are simply percentages: one indicates the proportion of tasks that were successfully passed, the other one states how many of the tasks were carried out safely. A task is considered unsafe if injury on the patient's part seems likely or actually occurs during task completion.

The pass/fail grading method makes the evaluation simple and as objective as possible. This is reflected in the high test-retest, intra- and interrater reliabilities achieved by the WST [6][7].

The WST requires the presence of a tester (giving instructions and being in charge of conducting the test) and a spotter (ensuring safe test execution); both roles can, however, be assumed by the same person.

To summarize, the WST takes little time (around 30 minutes) and effort (no special tools required) and is easy to evaluate (just percentage scores). More important, we think that it makes most sense to adopt a test developed by the rehabilitation community and emphasize that the latter seems

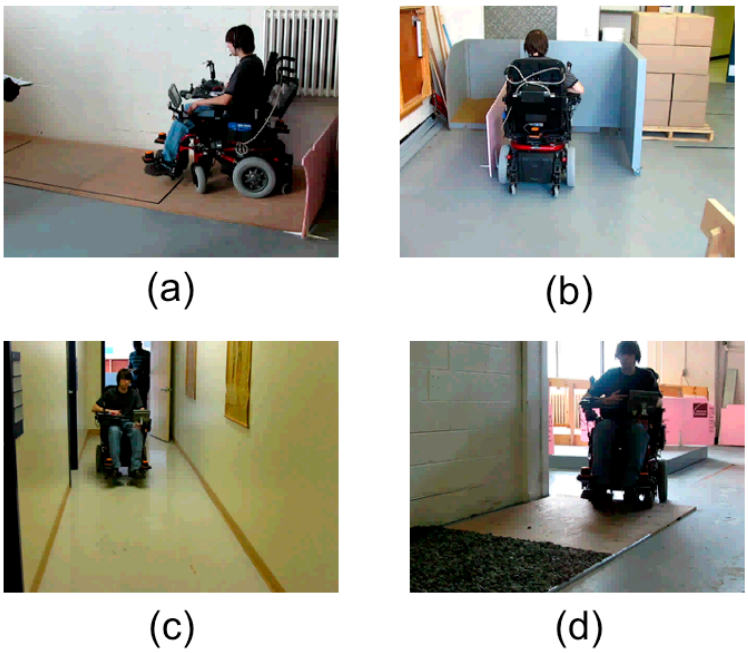

Figure 3: Various stations of the Wheelchair Skills Test, with an intelligent wheelchair. (a) The wheelchair must travel along a sloped platform. (b) The wheelchair must be aligned to the left wall. (c) The wheelchair must move forward through a door. (d) The wheelchair must travel through increased rolling resistance (in this case, gravel).

to converge on the WST as a standard. This is why we have decided to use this test in order to evaluate the SmartWheeler project and propose that it be used by similar projects, too.

6. PROPOSAL OF A TEST PARADIGM

As mentioned above, it is desirable to use a test developed by the rehabilitation research community to evaluate the performance of intelligent wheelchairs. The WST (like the other tests reviewed in [4] and [9]) was designed principally to evaluate the joint performance of the disabled person with their wheelchair, rather than evaluating specifically the person, or wheelchair, alone. This is an important aspect, one that is worth considering also in the context of evaluating intelligent wheelchairs.

The expected outcome of applying the WST consists of two numbers indicating the percentage of skills that were accomplished successfully and safely, respectively. These numbers are absolute though, and there is no obvious way of interpreting them. For instance, what does it mean if a disabled person in an intelligent wheelchair (or, to stay within our paradigm, rather the intelligent wheelchair in cooperation with a disabled person) achieved a score of 60%? Is 60% a good or a bad score? In order to attribute more meaning to the result, one should apply the test under different conditions.

A standard way of doing this in clinical practice is to use the WST with a given individual using a variety of wheelchairs. This setup measures the change in the skills exhibited by the person onboard the various wheelchair platforms, and can allow the selection of a wheelchair matched to a person's needs. In the context of intelligent wheelchairs, the WST could be applied to compare the performance and safety achieved by an individual using both a conventional powered wheelchair and an intelligent wheelchair. The difference between the two outcomes measures how helpful the intelligent software was to the wheelchair user.

The WST has also developed for assessing the efficacy of rehabilitation, by comparing the results of taking the test before and after training or modification to the wheelchair. The WST can thus be applied to evaluate the impact of incorporating different intelligent systems onboard the smart wheelchair (e.g. speech vs. tactile interface, semi-autonomous vs. fully autonomous navigation, etc.) The WST can be further used to evaluate the effectiveness of the training phase, when the human is becoming acquainted with the intelligent wheelchair.

Finally, the WST could be used to faciliate the comparison of results produced by different research teams working on similar projects.

Yet there are limitations to using such a constrained evaluation procedure The set of tasks included in the WST is very constrained, which makes it difficult to test the system for higher-level tasks such as 'Leave the house'. We will touch on this problem in the next section. Another limitation is that the presence of a qualified tester/spotter is required. However, dispensing with such personnel is possible only if an experiment does not involve actual patients. We see our methodology in between these two extremes: We are not arguing that the WST be the *only* evaluation tool used to validate intelligent wheelchairs, but rather that it serves a useful purpose to benchmark systems at a crucial point in their development, namely when the state of the project already warrants experiments with real patients, without being as advanced yet as to necessitate long-term studies in real environments.

7. PRELIMINARY EVALUATION

A preliminary evaluation was conducted to evaluate the design and implementation of the communication interface of the intelligent wheelchair. Seven healthy subjects, all of them university students without involvement in the project, were asked to go through the tasks of the WST, using appropriate vocal commands to communicate each task. The physical robot was not involved in this task; the only measures of interest were the performance of the speech recognition and the dialogue management modules through the set of WST skills. These results were reported in earlier publications [1].

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 École Polytechnique de Montréal [3]; 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 analyzing the performance of the communication interface during these experiments are presented in Figure 4. This evaluation involved the full robot capabilities, from the robust communication to autonomous navigation. However the results presented here focus primarily on the speech interface, which is the primary contribution of the authors.

As shown in Figure 4, the robot's current architecture 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.

Overall, the test subjects were satisfied by the functionality of the robot's interface and appreciated the visual feedback capabilities. While the word error was in some cases quite high, the use of probabilistic techniques allowed the system to maintain a low rate of incorrect actions, thus providing satisfactory performance overall. Some subjects felt they needed more time to get familiar with the platform to exploit it more successfully. Training time for all subjects was on the order of 30 minutes.

Based on these results, the system was judged to be sufficiently usable and robust to move forward with experiments involving the target population (disabled people). Therefore a third round of experiments involving eight subjects with mobility impairments is currently underway.

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%)

Figure 4: 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.

8. FUTURE PERSPECTIVES

The WST has been designed for evaluating skill performance and safety in a controlled environment. As just stated, we think that this paradigm is generally well-suited for the purpose of testing intelligent wheelchairs. Referring to the distinction made in [9] and summarized in section 4, we will briefly comment on the other two test categories as well:

Real environments

Observing users in their everyday setting in order to assess their performance in a standardized manner is difficult both practically and theoretically. First, from a practical point of view, it is time-consuming and thus expensive, as a clinician would have to examine the test subject's daily wheelchair performance over a sufficiently long period of time. Second, the high variance in terms of environment properties makes it conceptually hard to compare scores. Coping with this high variance is, however, one of the foremost challenges in the development of an intelligent wheelchair, so evaluating how well the device can deal with it is crucial for assessing the success of the project. Consider, for instance, a user utterance like "I'm hungry." There is no standardized way of benchmarking the wheelchair's reaction in such a situation because the best reaction depends very much on the setting: Downtown the best option might be to ask the user which restaurant he/she wants to go to, whereas at home it might be best to take him/her to the kitchen. A modified test paradigm will be necessary to rate the quality of intelligent control software in such real environments. But to rigorously assess more basic performance quality we need the more restricted and controlled type of scenario we have presented.

Virtual environments

Virtual tests involving a simulator are probably even cheaper to conduct than controlled-environment tests as proposed in this article. Routhier *et al.* [9], however, state that such tests have demonstrated a "limited applicability to assessment" mainly due to technical weaknesses of the simulators used. However, since big parts of the technology developed for intelligent wheelchairs (e.g. the interaction manager) are software rather than hardware, it might indeed make sense to evaluate these parts in a simulator. As the respective technology advances, this will clearly become more feasible

than it is today. But to assess the entire project it will be necessary to evaluate the interplay of both software and hardware. This is why we deem a controlled-environment test like the WST better suited.

9. CONCLUSION

In this paper we have suggested a methodology to quantify the performance of intelligent wheelchairs, and have applied this test to the evaluation of the speech interface of an intelligent wheelchair. Rather than designing a test from scratch we are building on work done by specialists in the field of rehabilitation. We have picked the WST, which seems to emerge as a *de facto* standard in the clinical and research communities. It is based on situations occurring in the daily lives of wheelchair users but still abstract enough to allow for precise measurements. The WST has been checked for validity and reliability by the developing team, which is crucial both principally and practically if a passing score is to be used as evidence that a wheelchair is ready to be deployed and funded by public health services and insurance companies. In this sense, a strict evaluation is a critical step towards both establishing and gauging the efficacy of intelligent wheelchairs.

Acknowledgement

Active participants in the project: The project was carried out in close collaboration with researchers at the École Polytechnique de Montréal (Paul Cohen, Sousso Kelouwani, Hai Nguyen, Patrice Boucher) who developed the intelligent wheelchair platform used for the human subject experiments. Researchers at the Universite de Montreal (Robert Forget) and at the Centre de réadaptation Lucie Bruneau (Wormser Honoré) acted as the primary clinical investigators. **Collaborators:** The project also benefitted from contributions by researchers at the Université de Montréal (Louise Demers), and also clinicians and technicians at the Centre de réadaptation Lucie Bruneau (Claude Dufour) and at the Constance-Lethbridge Rehabilitation Centre (Paula Stone, Daniel Rock, Jean-Paul Dussault). **Sponsors:** Financial support was provided through the Canadian Foundation for Innovation, the Natural Sciences and Engineering Council of Canada, and the Fonds québécois de la recherche sur la nature et les technologies. Additional funding was provided by the Fondation du Centre de réadaptation Lucie Bruneau, the Fondation Constance Lethbridge, and Robovic. Materials were donated by Sunrise Medical Canada and LiPPERT.

10. REFERENCES

[1] 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*, pages 345–356, 2009.

[2] 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.

[3] 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.

[4] O. J. E. Kilkens, M. W. M. Post, A. J. Dallmeijer, H. A. M. Seelen, and L. H. V. van der Woude. Wheelchair skills tests: A systematic review. *Clinical Rehabilitation*, 17:418–430, 2003.

[5] R. L. Kirby. *Wheelchair Skills Program (WSP), Version 4.1. Wheelchair Skills Test (WST) Manual*, 2007. `http://www .wheelchairskillsprogram.ca/eng/4.1/WST_Manual_ Version_4.1.pdf` (accessed 04/27/2008).

[6] R. L. Kirby, D. J. Dupuis, A. H. MacPhee, A. L. Coolen, C. Smith, K. L. Best, A. M. Newton, A. D. Mountain, D. A. MacLeod, and J. P. Bonaparte. The Wheelchair Skills Test (version 2.4): Measurement properties. *Arch. Phys. Med. Rehabil.*, 85:794–804, 2004.

[7] R. L. Kirby, J. Swuste, D. J. Dupuis, D. A. MacLeod, and R. Monroe. The Wheelchair Skills Test: A pilot study of a new outcome measure. *Arch. Phys. Med. Rehabil.*, 83:10–18, 2002.

[8] J. Pineau and A. Atrash. SmartWheeler: A robotic wheelchair test-bed for investigating new models of human-robot interaction. In *AAAI Spring Symposium on Multidisciplinary Collaboration for Socially Assistive Robotics*, pages 59–64, 2007.

[9] F. Routhier, C. Vincent, J. Desrosiers, and S. Nadeau. Mobility of wheelchair users: A proposed performance assessment framework. *Disability & Rehabilitation*, 25:19–34, 2003.

[10] 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.

[11] R. C. Simpson. Smart wheelchairs: A literature review. *J. of Rehabilitation Research & Development*, 42:423–436, 2005.