

Shared Control Policies for Safe Wheelchair Navigation of Elderly Adults with Cognitive and Mobility Impairments: Designing a Wizard of Oz Study

Ian M. Mitchell¹, Pooja Viswanathan², Bikram Adhikari¹, Eric Rothfels¹ and Alan K. Mackworth¹

Abstract—Older adults often lack the ability or strength to walk or to self-propel in manual wheelchairs. However, the need for a significant level of cognitive capacity to safely operate current powered wheelchair technology tends to exclude users with cognitive impairments. The resulting lack of mobility for a significant proportion of elderly residents in long term care leads to negative social, mental and physical outcomes. In this paper, we review shared (or collaborative) control policies implemented by smart wheelchair prototypes that have been tested by participants from their target populations, and highlight key usability findings and clinical insights. We discuss desirable features for control policies in new intelligent wheelchairs for cognitively impaired older adults. We then propose a method to evaluate these policies in a Wizard of Oz study conducted with cognitively impaired participants living in long term care facilities.

I. INTRODUCTION

As the population of industrialized countries ages, the number of adults living with mobility impairment is increasing. Manual and powered wheelchairs are standard approaches to address such impairments, but these technologies are unsuitable for many users [1], [2], especially among older adults, who often suffer from multiple chronic conditions. Dementia (primarily Alzheimer’s disease) is particularly prevalent: among the 1.5 million nursing home residents in the United States in 1995, 60%–80% were diagnosed with dementia [3], and the incidence of dementia is increasing. In response to the shortcomings of current powered wheelchair (PWC) technology and using techniques from autonomous robotics, researchers have proposed various intelligent or smart wheelchair systems over the past decades (see [4], [5] and the citations therein). However, fully autonomous navigation is not necessarily desirable for some of these populations; for example, wheelchair users with cognitive impairments can become confused or agitated when the chair appears to act on its own.

In earlier studies [6]–[8] we have examined intelligent wheelchair systems that provide collision avoidance (and in some cases wayfinding assistance) for the target population

This research was supported by CANWHEEL (the Canadian Institutes of Health Research (CIHR) Emerging Team in Wheeled Mobility for Older Adults Grant #AMG-100925), National Science and Engineering Council of Canada (NSERC) grants, the Canadian Foundation for Innovation (CFI) Leaders Opportunity Fund / British Columbia Knowledge Development Fund Grant #13113, the Institute for Computing, Information and Cognitive Systems (ICICS) at UBC, and TELUS.

¹Department of Computer Science, University of British Columbia, Vancouver, Canada. Email: {mitchell, bikram7, mack}@cs.ubc.ca

²Toronto Rehab Institute, University of Toronto, Toronto, Canada. Email: Pooja.Viswanathan@uhn.ca

of older adults suffering from cognitive and mobility impairments. Participants in these studies were in charge of navigating the chair, typically with a joystick, but the chair would stop if an obstacle was in the way. While participants enjoyed the autonomy offered by these systems, they often became frustrated by the simplistic collision avoidance policy.

The results of those trials encouraged us to pursue a shared control policy in which the wheelchair user and the automatic control system would collaborate to navigate safely and in the direction desired by the user to the extent possible; unfortunately, the specific policy which might be appropriate for this population is not obvious. Rather than wait to develop all the hardware and software infrastructure necessary to implement the various options, we have decided to adopt a Wizard of Oz (WoZ) experimental protocol and immediately begin an iterative design process on the shared control policy and user interface [9]. The WoZ idea is simple: A trial participant interacts with the chair as if the system were fully implemented, but it is a hidden experimenter (the “wizard”) who interprets the participant’s inputs and controls the chair’s responses.

The contributions of this paper are twofold. First, we survey smart wheelchair projects which have tested shared control systems with their target population. Second, we define a set of shared control policies—motivated by previous work in smart wheelchairs as well as by shared control for lane keeping in automobiles—suitable for implementation in a WoZ protocol. To our knowledge this is the first study of shared control using the WoZ approach, although [10] did use a WoZ experiment to examine the effectiveness of the stop-at-obstacle switched control policy in a smart wheelchair with five participants drawn from our target population.

II. RELATED WORK

We discuss related research in three areas that informs our planned trials.

A. Shared Control

By *shared control* we mean a control signal generated by combining real-time (and typically continuous) signals from multiple agents. For smart PWCs in general, the two agents are the PWC driver and a computer control system; in our WoZ trials the latter will be replaced by the teleoperator. This type of control is distinct from *supervisory control*, in which one agent (often a human operator) provides high-level and typically discrete time guidance to a second agent (often an embedded control system) by modifying modes or

parameters, or *switched control*, in which the agents take turns generating the control signal. The terms *collaborative control* or *cooperative control* might also be suitable for the shared control concept, although more often in the control literature they indicate coordinated control of multiple entities each with a single control agent, rather than multiple agents synthesizing a single control signal.

Shared control has been used in a wide variety of applications; for example, virtual fixtures in teleoperated minimally invasive surgeries (for example, [11]) or certain autopilot modes in commercial aircraft (see [12] for further citations). However, most of these applications assume a high degree of training for the human operator and/or involve control of many degrees of freedom; neither of these properties hold for our application.

By far the most relevant body of shared control research appears to be lane keeping or roadway departure avoidance systems for passenger cars: not only must these systems be amenable to novice users, but also the goal is to navigate a vehicle travelling in the plane using a control signal with one or two degrees of freedom. A number of such systems are available on selected models from well-known automobile manufacturers (see [13] and the citations therein), but here we briefly describe two examples representative of the systems which have been designed and tested in the academic literature. In both cases the goal is to keep a lookahead point (a predicted location of the vehicle a short time in the future) inside the current lane. In [14], a composite quadratic Lyapunov function is designed for a linear model of lateral control and used to synthesize a steering control signal that will maintain the lane constraints. This signal is only used if the car is at risk of departing the lane *and* the torque on the steering column from the driver is sufficiently small (a threshold test that the authors use as a proxy for diminished capability due to inattention, tiredness, illness). In [15] a model predictive control system synthesizes both steering and braking signals for a nonlinear model of vehicle motion (which includes a simple but data driven model of driver behaviour). These signals are only nonzero if the driver model is unable to maintain the lane constraints, so they are added continuously to the driver's input but only affect that input when necessary.

In the test vehicles used by both of these teams the steering assistance is effected by applying a torque to a traditional steering column, and in the experiments the drivers simulated distraction by reducing or releasing their grip on the steering wheel; consequently, there was never any conflict between the steering signals generated by the human and automation. Such conflict is far more likely in our application: WC users rarely release the joystick because it must be deflected for any movement to occur, joystick deflections are relatively larger and faster than those in a steering wheel, and our target population exhibits more frequent periods of inattention or confusion. In the domain of driving, such control signal conflicts are more likely to occur during emergency maneuvers. Three approaches to shared control road departure prevention (RDP) for a simulated emergency maneuver were considered

in [13] and compared with unassisted human control: haptic feedback (HF) in which torque was applied to the steering wheel, drive by wire (DBW) in which the steering angle was adjusted as necessary to avoid road departure, and the combination of HF and DBW. Note that in the final two cases the driver's input from the steering wheel was ignored whenever the RDP system was active. Through experiments in a vehicle simulator with 30 participants, the authors demonstrate that the HF system had essentially no effect on the path taken by the vehicle or the likelihood of road departure—the drivers simply generated more torque on the steering wheel and overrode the RDP system's input signal—whereas the DBW and DBW+HF cases dramatically and similarly reduced the likelihood of road departure.¹

Although force-feedback haptic joysticks have been used in some smart PWC systems (for example, [16], [17]), we will implement shared control policies more similar to the DBW scenario described above for a number of reasons: (a) the DBW approach worked better in [13], (b) the DBW approach has not been tested in our target population, and (c) force feedback haptic joysticks are bulky, expensive and/or lack sufficient torque. Nevertheless, we do not want to dismiss the haptic feedback channel entirely, so we will use a simple vibration actuator mounted below the joystick to provide participants with a binary signal indicating whether their joystick input is being modified or ignored. Drawing on research in the driving domain there is again support for this approach. Bandwidth feedback, in which seat vibrations were applied depending on whether the driver was inside or outside lane constraints, was found to have a positive effect on both performance and learning of a lane keeping task among novice drivers [18], with no significant difference in perceived effort or workload. Vibration alerts through the steering wheel were found to reduce reaction time and collisions in a distracted forward collision warning scenario [19], and participants reported a significantly higher perceived usefulness and overall satisfaction with the steering wheel alerts when compared with no alerts.

B. Smart Powered Wheelchairs

The field of smart PWCs is vast, so we restrict our main descriptions to projects from the last decade in which a shared control system has been tested with users who have cognitive and/or mobility impairments. Citations to older systems can be found in the comprehensive survey article [4]. Other intelligent and modified PWCs have been tested more recently on target populations (for example, [6]–[8], [20], [21]) but these systems implemented only supervisory and/or switched control policies.

1) *The Collaborative Wheelchair Assistant*: Zeng and colleagues developed the collaborative wheelchair assistant (CWA) [22], [23], for which the user prespecifies a destination and path using a graphical user interface. As the user

¹The DBW and DBW+HF cases increased the chances of hitting pylons which the drivers had been instructed to avoid but which the RDP system did not consider obstacles. In our application all obstacles will be visible to the (wizard simulated) collision avoidance system.

drives the wheelchair any motion along the path is unchanged but motion perpendicular to the path is modified by an elastic path controller: that component of the user's input is blended with a PD controller that seeks to return the chair to the original path. Parameters are chosen so that the user can deviate from the path if necessary (for example, to avoid an unexpected obstacle), but will tend to return to the path in the absence of conscious effort to the contrary.

In trials with able-bodied participants, the path guidance system resulted in fewer joystick movements and led to optimal driving performance from the first trial, whereas driving optimally in the non-assisted mode required more trials. The system was then tested by three participants with cerebral palsy and two with traumatic brain injury, all of whom were previously denied PWC use by clinicians under "conservative prescription criteria." After a few training sessions, all participants were able to drive using path guidance. Driving with CWA enabled led to decreased participant effort (when compared to driving without CWA) while still maintaining high performance. Some participants could only drive safely with the path guidance, but it was found that two participants could use the path guidance as a training tool and go on to navigate standard PWCs independently.

2) *Collaborative Wheelchair Control (Version 1)*: Urdiales and colleagues [16] implement a shared control strategy in three layers. A bottom "safeguard" layer prevents imminent collisions, while a top "deliberative" layer (activated only for users with cognitive impairments) plans a path to the goal and decomposes that path into a series of waypoints. A potential field method is used to generate "robot" control signals (as if the system were navigating autonomously). The shared control is implemented by the middle "reactive" layer, which at each instant evaluates the local efficiencies of the human and robot control signals with respect to three measures: smoothness, directness and safety. The system then blends the human and robot control signals according to weights that depend on their current and recent past average relative efficiencies.

The system was first tested with thirty hospital in-patients who had varying levels of physical and cognitive disabilities. All users successfully completed the tasks in shared control mode, and in almost every case that mode improved human performance. Somewhat surprisingly, shared control was also found to improve the robot controller's performance in many cases that prove challenging for fully autonomous systems based on potential field methods. The key finding was that shared control was able to equalize navigation performance between users, even though the participant pool had quite heterogeneous capabilities. In subsequent research [24] the authors have proposed to proactively adjust the user's control authority by comparing it to a driving database from seventy previous users and predicting user performance for situations based on the performance in similar situations from the database. A study conducted with seventeen participants with varying physical and cognitive disabilities showed that this collaborative control system compensates for skills that users do not have and improves performance of residual skills.

3) *Collaborative Wheelchair Control (Version 2)*: Carlson and Demiris [25] implement a two component shared control scheme. One component attempts to discern which of a collection of prespecified subgoals is the user's intended destination based on how well the angular direction of the joystick aligns with the direction that an autonomous controller would take to get to that subgoal. If the system is confident enough that the user is seeking a particular subgoal, then the user's input is blended with the autonomous controller's. A second component implements a standard obstacle avoidance algorithm that attempts to find a direction close to the user's input which is safe with respect to local obstacles. These two modified control signals are then blended to produce a final signal.

The system was tested with twenty-one healthy participants 17 to 47 years old and one experienced wheelchair user with mobility impairment in a realistic office environment. The system was able to increase safety for all users at a slight cost in time. It also allowed users to pay less attention to driving, as indicated by performance on a secondary task and by eye-tracking data, thus decreasing cognitive workload, visual attention and manual dexterity demands. In particular, the experienced wheelchair user benefitted from the system in high workload situations and when inattentive to the driving task.

4) *Dynamic Shared Control*: Li, Chen and Wang [26] blend the user's and an autonomous controller's direction inputs so as to increase three measures of the final shared control signal—safety, comfort and obedience—which essentially correspond to the safety, smoothness and directness measures discussed in section II-B.2. An online optimization procedure chooses the blending coefficient to maximize the minimum of these measures.

After carrying out experiments with healthy volunteers who simulated having various disabilities (physical and visual), the system was tested with older adults (75 to 84 years old) who had mobility impairment but were cognitively intact. Results showed that the system improved the smoothness of wheelchair trajectories and reduced the likelihood of collision with obstacles. More recent work [27] has proposed that the user specify intent, such as docking at a table, through the joystick and an augmented reality display; however, testing of this system by users with disabilities has not yet been reported.

C. Wizard of Oz Studies

The WoZ experimental technique [28] (also sometimes called *experimenter in the loop*) was pioneered to help design user interfaces involving natural language processing. It has since been used in many other domains, but we focus here on its applicability to human robot interaction (HRI) scenarios [29]. The WoZ approach is a powerful tool for iterative design of HRI because it allows various options to be tested before significant development effort is invested in any; however, its validity is subject to several caveats. In particular [30], the proposed system's behavior must be carefully specified, the wizard (who is subject to

human limitations) must be able to simulate the proposed system, the simulation must be convincing, and various aspects of the wizard’s behavior must be controlled (eg: how repeatably does the wizard generate the same robot action for a given participant action?). While widely adopted in the HRI community (see [29] for an extensive bibliography), WoZ approaches do not appear to have been used previously to study shared control systems.

III. A WIZARD OF OZ STUDY FOR TESTING SHARED CONTROL POLICIES

A large number of intelligent wheelchairs have been developed offering varying levels of control to the user, but only a small proportion of these systems implement shared control, and fewer yet have been tested by their target audiences (users with a combination of mobility and some other impairment that precludes them from traditional powered wheelchairs). In particular, for our target population it is not clear that study results from younger adults or healthy older adults are generalizable to older adults with dementia, who in addition to having cognitive impairment often have several other disabilities due to old age [31]. One reason for the paucity of clinical results is the enormous amount of effort required to develop, tune and test the hardware and software infrastructure before a user trial can safely occur. While the CANWHEEL team and others are working on this infrastructure, we do not want to delay the iterative design process that will be required to develop a suitable shared control interface for this target population. In this section we describe the WoZ study that we have created in order to begin this design process.

A. Study Overview

A primary method of addressing some of the validity concerns with WoZ studies is to employ scenarios that constrain the types of interaction sufficiently that the properties of those interactions may be fully specified. In our case, we choose scenarios representative of the activities of daily living for which the study participants might want to use a PWC. The *Power-Mobility Indoor Driving Assessment (PIDA)*² is an assessment instrument designed to describe an individual’s indoor mobility status in relation to their environment and roles [32]; in particular, it is conducted within a client’s own environment as opposed to on a standard obstacle course. The thirty tasks on the PIDA were chosen through an iterative consensus process involving occupational therapists and users of PWC resident in long term care facilities with the goal of providing information to therapists and caregivers for prescribing PWCs and/or designing training programs. For each task PIDA specifies a brief instruction to be read to the client and the criteria under which to score task completion on a four point scale:

- (1) Unable to complete task.
- (2) Potentially harmful collisions.
- (3) Hesitantly, requires several tries, or harmless collisions.

(4) Completely independently.

A task can also be scored “not applicable” if it is irrelevant to a client’s environment or roles.

Collecting data on all thirty PIDA tasks is impractical, so we selected scenarios based on recommendations from occupational therapists and our observations in previous studies [7], [33], [34]. These scenarios present challenges regularly faced by the vast majority of long-term care residents:

- *Parking at a Table with Limited Space*: The participant has to dock the PWC in a constrained space under a table.
- *Back-in Parking*: The participant has to back up between two chairs spaced 1m apart.
- *Left and Right turn at an intersection with obstacle and turn in place*: The participant has to slow down during approach to a 4-way intersection, check the intersection for traffic, and proceed through the intersection in a specified direction when it is clear. The participant then completes a 180° turn in place, passes back through the intersection and into the original hallway. An empty box is slid through the intersection to simulate traffic in the intersection. This scenario encompasses three PIDA tasks.
- *Maneuverability*: The participant has to drive through a short slalom course without bumping into any of the obstacles.
- *Elevator*: The participant has to drive the wheelchair into an occupied elevator, may optionally turn around, and then exit the elevator. Assistance is provided to ensure the elevator door is open. This scenario encompasses three to four PIDA tasks.

Additional details on one of these tasks is provided in section IV.

During the trials we will test three shared control policies on each of the scenarios, with the goals of determining under which policies the participants are most able to achieve high scores on the PIDA scale, which policies are most effective from the participants’ point of view, and how different policies effect various quantitative measures (such as completion times, trajectory smoothness, degree of control intervention, ...). The policies determine how the shared control signal is synthesized from the participant’s and wizard’s control signals. Control signals take the form of a joystick deflection described in polar coordinates by a magnitude and an angle. Detailed descriptions of the policies are given below.

- *Policy 1 (speed control)*: The wizard restricts the maximum magnitude control signal depending on how close the PWC is to the nearby obstacles.
- *Policy 2 (direction control)*: The wizard assumes (nearly) full control if the PWC crosses a threshold distance from an obstacle, drives the PWC to the nearest free space, and then relinquishes full control to the participant.
- *Policy 3 (autonomous driving)*: The wizard chooses a suitable path to accomplish the specified task and then drives the PWC along that path. The participant’s joy-

²<http://fhs.mcmaster.ca/powermobility/pida.htm>

stick is used as a discrete event: a significant backward deflection indicates that the chair should stop moving and a significant forward deflection indicates that the chair should start following the commands of the wizard again.

In all of these policies the wizard primarily acts as a substitute obstacle sensor, a task which a human being can perform more robustly but inaccurately and slowly than typical robot sensors. To overcome the latter two deficiencies, the control policies and wizard interface are designed to be quick to operate and the wizard will practice extensively with the obstacle safety distances to achieve repeatable performance. The inaccuracy, slow reaction time and indoor environment somewhat offset the robustness advantage of using the wizard as an obstacle sensor, but we will obviously have to take into account the degree to which real sensors' reduced robustness would effect the outcome when we evaluate our results. In the third control policy the wizard also has to synthesize a path to achieve the task, but path planning in the plane through an obstacle map is essentially a solved problem in robotics; furthermore, this policy is intended to represent an idealized smart PWC (see section III-C) and so need not be implementable.

Since we seek to design our control system around a joystick interface, the wizard does not respond to any verbal utterances by the participant, with the exception—for reasons of safety—of any obvious request to bring the wheelchair to a halt (eg: “stop”); in that case the trial would be considered abandoned.

B. Shared Control Policies

In this section we provide precise descriptions of the shared control policies. A control signal (representing a joystick deflection) is specified in polar coordinates (ρ, θ) , where $\rho \in [0, 1]$ is the (normalized) magnitude of the deflection and $\theta \in [-\pi, +\pi]$ is the angle (with $\theta = 0$ corresponding to a forward deflection and θ increasing counter-clockwise). A shared control policy takes a participant control signal (ρ_p, θ_p) and a wizard control signal (ρ_w, θ_w) and generates a shared control signal (ρ_s, θ_s) for the PWC to implement. Define a lookahead time larger than but close to the wizard's reaction time (eg: 0.5s), and let δ be the (wizard estimated) distance to the closest obstacle at that lookahead time.

The wizard's control interface is a standard wireless PS3 controller, which has two joysticks, four triggers and more than six buttons. In every policy the wizard has access to a button which immediately stops the chair and disables all motion, and another button which re-enables motion. The chair also has a number of wired emergency stop buttons in case the wireless connection fails.

Policy 1: Some small number $n \geq 1$ of obstacle distance thresholds $\{d_i\}_{i=0}^{n-1}$ and corresponding magnitude control signal constraints $\{\bar{\rho}_i\}_{i=0}^n$ are chosen such that $0 < d_0 < d_1 < \dots < d_{n-1}$ and $0 \leq \bar{\rho}_0 < \bar{\rho}_1 < \dots < \bar{\rho}_n \leq 1$. The environment around the wheelchair is partitioned into four sectors centered on the four cardinal directions $\phi_j = \{-\pi/2, 0, +\pi/2, \pm\pi\}$, and a separate δ^{ϕ_j} is estimated by the

wizard in each sector. The magnitude limit for each sector is then determined by

$$\rho_w^{\phi_j} = \begin{cases} \bar{\rho}_0, & \text{if } \delta^{\phi_j} \leq d_0; \\ \bar{\rho}_i, & \text{if } d_{i-1} < \delta^{\phi_j} \leq d_i; \\ \bar{\rho}_n, & \text{if } d_{n-1} < \delta^{\phi_j}. \end{cases} \quad (1)$$

The actual values of the parameters were determined empirically: obstacle distance thresholds and magnitude control signal constraints were chosen so that time to collision was no smaller than about two seconds.

Based on participant feedback from the pilot trial, the lowest non-zero speed for tasks which involved close approach to an obstacle (eg: parking) was designated as a *docking speed*. This speed was set very low, and the corresponding distance threshold set to essentially zero. In order to indicate to the participant that this special mode was enabled, the PWC was also brought to a complete halt for approximately one second before entering the docking speed mode.

The shared control signal is generated by first determining in which sector ϕ_k the current angle of the participant's joystick lies. Then

$$\begin{aligned} \theta_s &= \theta_p, \\ \rho_s &= \min(\rho_w^{\phi_k}, \rho_p); \end{aligned}$$

in other words, the shared magnitude is bounded by the magnitude limit for the current sector while the shared angle is unmodified. The joystick vibrator is turned on whenever the participant's magnitude control signal is being clipped significantly ($\rho_p \gtrsim 1.15\rho_s$).

Policy 2: Two obstacle distance thresholds d_0 and d_1 are defined such that $d_0 < d_1$. This policy operates on a hysteretic principle. Initially the participant begins in full control: $\rho_s = \rho_p$ and $\theta_s = \theta_p$. If the smaller threshold is crossed $\delta < d_0$, then the wizard assumes nearly full control: $\rho_s = \min(\rho_p, \rho_w)$ and $\theta_s = \theta_w$. The wizard drives the chair until it reaches a sufficiently clear area, defined by $\delta > d_1$, and then relinquishes control to the participant. The joystick vibrator is on whenever the wizard is in control. The magnitude limiting feature from policy 1 is also available to the wizard under policy 2 and was used, for example, to ensure safety during the docking phase of the parking tasks.

In order to reduce the likelihood that the wizard will have to assume control but in keeping with the intention that policy 2 is designed for users who desire a greater degree of intervention, in this policy the wizard's controller also has a small number of buttons, each of which activates a pre-recorded wayfinding audio prompt (“turn left,” “turn right,” “drive forward,” or “back up”). The wizard activates an appropriate prompt when the wheelchair's heading is sufficiently far from that of a reasonable path to the goal or if the participant does not make sufficient progress toward the goal during a period of time.

Policy 3: Before the trial begins, the wizard chooses a suitable path to accomplish the goal. Once the trial begins (indicated to the participant through an audio prompt), the wizard deflects the PS3 joystick to drive the chair along

that path. During movement $\rho_s = \rho_w$ and $\theta_s = \theta_w$. The participants have the option to stop and restart the movement by deflecting the joystick backwards or forwards respectively; otherwise their joystick signal is ignored. The joystick vibrator is never activated.

C. Discussion of Policy Choices

We had several goals in choosing the policies described above: (a) the three policies should be distinct from one another and distinct from policies that had been tested on our target population before, (b) the policies should not be too frustrating for members of our target population, (c) with suitable training the wizard should be able to safely and repeatably implement the policies, and (d) it should be possible to implement the policies robotically given appropriate sensors.

Policy 1 is based on the results of the studies conducted in [6]–[8]. These studies all examined our target population but used a simplistic control policy: stop all motion in directions in which an obstacle is detected. Based on participant feedback it seems likely that multiple and/or adaptive distance thresholds and speed settings would reduce participant frustration, and tactile feedback might improve performance. It is also in some sense a minimalistic modification of the participant’s control signal because it does not modify the directional component of that signal and merely clips the magnitude component. Finally, it can be viewed as increasing the time-to-collision so that the participant can react to nearby obstacles.

Policy 2 is a more aggressive intervention adapted from the drive-by-wire road departure prevention approach in [13]. In both cases the automation / wizard takes full control authority when necessary to avoid a collision but cedes that authority as soon as possible. The wayfinding prompts are included based on positive results in [7].

Policy 3 perhaps stands out as likely to fail the last of our criteria (eg: robotically implementable); however, we include it to represent a perfect smart PWC—the wizard will drive the chair reliably and safely to accomplish exactly the desired task—so that we can calibrate participant responses.³

The results from [7] indicate that participants vary in their acceptance of autonomous action by the wheelchair; however, these perceptions have not been explored or documented in detail. Testing the three policies above will help us to identify differences in perceptions of shared control across participants, scenarios and policies for our target population.

IV. EXAMPLE

In this section we briefly demonstrate how the two shared control policies (policies 1 and 2) might appear in one of the trial scenarios as a proof of concept that shared control policies for collision avoidance can be explored with a WoZ experimental protocol. The data for this example was

³The study described in [7] already provides us with a calibration of participant responses to a smart PWC lying at the other end of the spectrum: a completely robotic system using somewhat dated sensor and compute technology and the simplistic stop-at-obstacle control policy.




Fig. 1. The *parking at a table* task map.

generated in the lab with a healthy young adult member of the research team playing the role of the participant; consequently, we do not attempt to draw any conclusions from it about the relative effectiveness of the policies.

We choose the *parking at a table* task from section III-A as the example. The study coordinator will identify a table similar to one normally used by the participant and arrange a 1m horizontal gap between two chairs or between a chair and a table leg. The current shared control policy will be specified to the participant and the study coordinator will confirm with the participant and the wizard that the trial can begin. The participant will then be told “please park under that table.” A representative overhead view of the local environment and PWC starting position and anticipated path is shown in figure 1. Although here we show only one example run for each policy, the full experimental protocol will include multiple runs for each policy. Every effort will be made to use the same environmental setup and initial wheelchair position for each of these replications; however, it is expected that the setup will differ between participants because trials will occur in each participant’s own facility.

There are a few details to keep in mind when examining the experimental results below. First, to improve safety we have decided to limit shared control magnitude ρ_s at all times to 50% or less of peak, because all of the scenarios involve driving in confined spaces. Second, measurement of joystick angles is inaccurate when joystick magnitudes are near zero; however, in these cases the angle is irrelevant because the magnitude is near zero.

Figure 2 shows the results from running with policy 1. The top plot shows the path and heading of the chair (estimated by dead reckoning from odometry data) superimposed on an occupancy grid for the region around the desk. The occupancy grid was precomputed through a SLAM algorithm using laser rangefinder data, but is included only for visualization purposes; the PWC has no environmental sensors other than the wizard during the run. The rectangle labelled “Table” is the 1m wide target for the task, and also provides a scale for the figure. The tabletop itself does not appear in the occupancy grid because the laser rangefinder was operating in a plane at ankle height. This top plot shows that three objects were approached during the run at approximately times 35s, 53s and 69s. The other two plots show the control

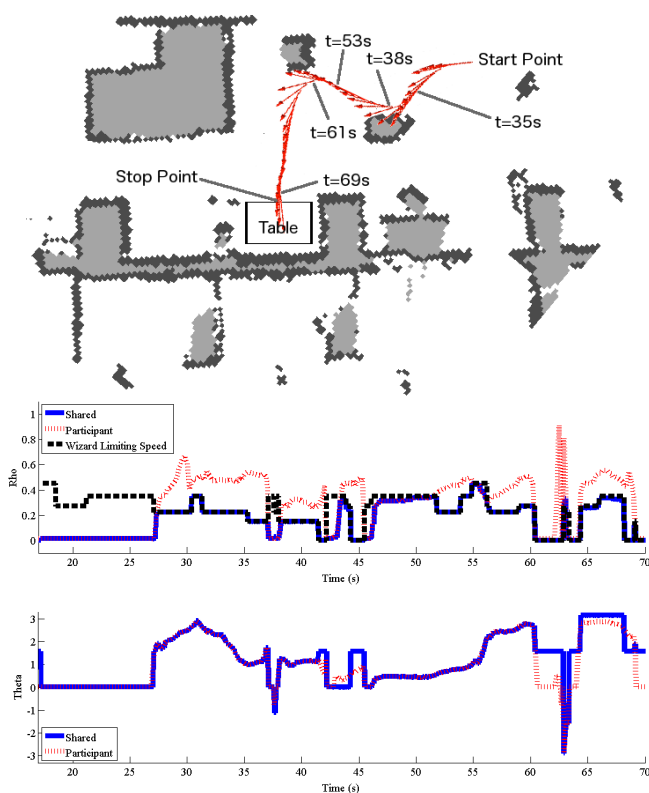


Fig. 2. Results for Policy 1.

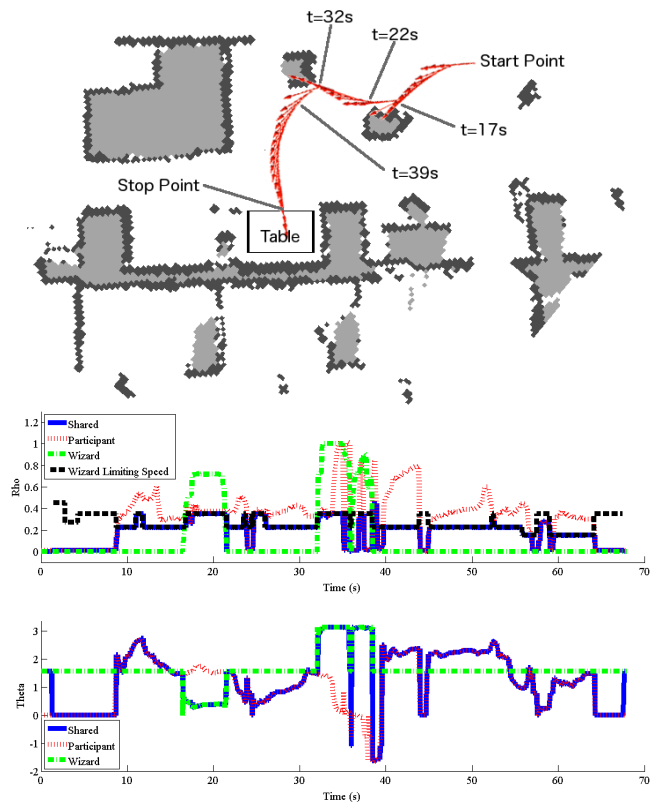


Fig. 3. Results for Policy 2.

signals from the same run. Because of the tight environment, the wizard constrains the magnitude (middle plot) throughout most of the run, especially in the first half of the run when the PWC is maneuvering through a gap that is less than 2m wide, and at the end of the run as the PWC approaches the table. The shared control angle tracks the participant's angle (bottom plot), except for times when the magnitude is near or at zero and the angle measurements are therefore suspect.

Figure 3 shows the results from running with policy 2. The PWC approaches the obstacles at times 17s and 32s, but under this policy the wizard takes full control at these times, maneuvers the chair to safety, and then relinquishes control at times 22s and 39s. In the middle plot these interventions are clearly visible as the only periods during which the wizard's magnitude is nonzero. Although the wizard can also limit magnitude in this policy, that limiting is less aggressive and never requires the full stop that was commonly encountered when running policy 1. The bottom plot shows that the shared angle tracks the participant's angle except during the intervention periods, when it tracks the wizard's angle (note that the wizard's angle is shown as $\pi/2$ when the wizard's control is inactive).

V. CONCLUSION AND FUTURE WORK

We have described shared control policies that have been tested in automobile driver assistance systems and intelligent wheelchairs. We then constructed a set of three control policies to test with our target population, which is cognitively and mobility impaired older adults living in long

term care facilities. The speed control policy represents a minimalist intervention, the direction control policy a more comprehensive intervention, and the autonomous driving policy allows us to calibrate user responses against a near perfect system. The testing and refinement of the policies will use a WoZ experimental protocol, which required us to tailor our policies to be implementable by a human wizard. We briefly described the set of five test scenarios that we will use. The WoZ approach has not previously been used to test shared control systems, but we demonstrated by running two of the policies on one of the scenarios that such testing appears feasible.

The participants for the WoZ will be drawn from cognitively impaired residents of up to six long term care facilities in the greater Vancouver area. Each participant will go through a training period, and then attempt each of the five scenarios with each of the three policies two to three times over three to five driving sessions while the study coordinator collects qualitative feedback from the participant and sensor systems on the PWC collect quantitative data about the chair's, participant's, wizard's and environment's behaviors. Quantitative data on subjective workload and user satisfaction are collected through standard questionnaires (NASA TLX and QUEST 2.0) after driving sessions. Finally, a subset of participants are invited to pre and post driving semi-structured interviews, to try to better understand their experience of interacting with the system. As of December 2013 we had completed trials with ten participants.

In addition to exploring which shared control policies are most effective, we will also collect a suite of sensor measurements during the trials—including laser rangefinder, RGB-D camera, accelerometry, participants’ joystick movements, wizard’s controller actions, odometry, and hopefully eye-tracking and physiological measures (eg: heart rate, skin conductance)—with the goal of releasing the anonymized data to provide other researchers an opportunity to see a robot’s view of this population’s behavior and environment.

ACKNOWLEDGMENT

The authors would like to thank Pouria TalebiFard for his extensive help with the wheelchair hardware, with ROS, and as a test pilot, as well as the whole CANWHEEL team for their clinical and technological feedback.

REFERENCES

- [1] L. Fehr, W. E. Langbein, and S. B. Skaar, “Adequacy of power wheelchair control interfaces for persons with severe disabilities: A clinical survey,” *Development*, vol. 37, no. 3, pp. 353–360, 2000.
- [2] R. C. Simpson, E. F. LoPresti, and R. A. Cooper, “How many people would benefit from a smart wheelchair?” *Journal of Rehabilitation Research and Development*, vol. 45, no. 1, pp. 53–71, 2008.
- [3] J. L. Payne, J.-M. E. Sheppard, M. Steinberg, A. Warren, A. Baker, C. Steele, J. Brandt, and C. G. Lyketos, “Incidence, prevalence, and outcomes of depression in residents of a long-term care facility with dementia,” *International Journal of Geriatric Psychiatry*, vol. 17, no. 3, pp. 247–253, 2002.
- [4] R. C. Simpson, “Smart wheelchairs: A literature review,” *Journal of Rehabilitation Research and Development*, vol. 42, no. 4, pp. 423–438, 2005.
- [5] P. Viswanathan, “Navigation and obstacle avoidance help (NOAH) for elderly wheelchair users with cognitive impairment in long-term care,” Ph.D. dissertation, University of British Columbia, 2012.
- [6] R. H. Wang, S. M. Gorski, P. J. Holliday, and G. R. Fernie, “Evaluation of a contact sensor skirt for an anti-collision power wheelchair for older adult nursing home residents with dementia: Safety and mobility,” *Assistive Technology*, vol. 23, no. 3, pp. 117–134, 2011.
- [7] P. Viswanathan, J. J. Little, A. K. Mackworth, and A. Mihailidis, “Navigation and obstacle avoidance help (NOAH) for older adults with cognitive impairment: a pilot study,” in *Proc. Int. ACM SIGACCESS Conf. on Computers and Accessibility (ASSETS)*, 2011, pp. 43–50.
- [8] T.-V. How, R. Wang, and A. Mihailidis, “Evaluation of an intelligent wheelchair system for older adults with cognitive impairments,” *Journal of NeuroEngineering and Rehabilitation*, vol. 10, no. 1, pp. 90–105, 2013.
- [9] P. Viswanathan, R. Wang, and A. Mihailidis, “Wizard-of-Oz and mixed-methods studies to inform intelligent wheelchair design for older adults with dementia,” in *Association for the Advancement of Assistive Technology in Europe*, 2013.
- [10] R. H. Wang, A. Mihailidis, T. Dutta, and G. R. Fernie, “Usability testing of multimodal feedback interface and simulated collision-avoidance power wheelchair for long-term-care home residents with cognitive impairments,” *Journal of Rehabilitation Research and Development*, vol. 48, no. 7, pp. 801–822, 2011.
- [11] T. Yamamoto, N. Abolhassani, S. Jung, A. M. Okamura, and T. N. Judkins, “Augmented reality and haptic interfaces for robot-assisted surgery,” *Int. Journal of Medical Robotics and Computer Assisted Surgery*, vol. 8, no. 1, pp. 45–56, 2012.
- [12] N. Matni and M. Oishi, “Reachability-based abstraction for an aircraft landing under shared control,” in *Proc. American Control Conference (ACC)*, 2008, pp. 2278–2284.
- [13] D. I. Katzourakis, J. C. F. de Winter, M. Alirezaei, M. Corno, and R. Happee, “Road-departure prevention in an emergency obstacle avoidance situation,” *IEEE Trans. Systems, Man, and Cybernetics: Systems*, vol. Early Access, 2013.
- [14] N. M. Enache, S. Mammari, M. Netto, and B. Lusetti, “Driver steering assistance for lane-departure avoidance based on hybrid automata and composite Lyapunov function,” *IEEE Trans. Intelligent Transportation Systems*, vol. 11, no. 1, pp. 28–39, 2010.
- [15] A. Gray, M. Ali, Y. Gao, J. K. Hedrick, and F. Borrelli, “A unified approach to threat assessment and control for automotive active safety,” *IEEE Trans. Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1490–1499, 2013.
- [16] C. Urdiales, J. Peula, M. Fernandez-Carmona, C. Barrué, E. Pérez, I. Sánchez-Tato, J. del Toro, F. Galluppi, U. Cortés, R. Annichiarico, C. Caltagirone, and F. Sandoval, “A new multi-criteria optimization strategy for shared control in wheelchair assisted navigation,” *Autonomous Robots*, vol. 30, no. 2, pp. 179–197, 2011.
- [17] E. B. Vander Poorten, E. Demeester, E. Reekmans, J. Philips, A. Huntemann, and J. De Schutter, “Powered wheelchair navigation assistance through kinematically correct environmental haptic feedback,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 3706–3712.
- [18] S. de Groot, J. C. F. de Winter, J. M. López García, M. Mulder, and P. A. Wieringa, “The effect of concurrent bandwidth feedback on learning the lane-keeping task in a driving simulator,” *Human Factors*, vol. 53, no. 1, pp. 50–62, 2011.
- [19] J. Chun, S. H. Han, G. Park, J. Seo, I. Lee, and S. Choi, “Evaluation of vibrotactile feedback for forward collision warning on the steering wheel and seatbelt,” *Int. Journal of Industrial Ergonomics*, vol. 42, no. 5, pp. 443–448, 2012.
- [20] R. Ceres, J. Pons, L. Calderón, A. R. Jiménez, and L. Azevedo, “A robotic vehicle for disabled children,” *IEEE Engineering in Medicine and Biology Magazine*, vol. 24, no. 6, pp. 55–63, 2005.
- [21] S. McGarry, L. Moir, and S. Girdler, “The smart wheelchair: is it an appropriate mobility training tool for children with physical disabilities?” *Disability and Rehabilitation: Assistive Technology*, vol. 7, no. 5, pp. 372–380, 2012.
- [22] Q. Zeng, C. L. Teo, B. Rebsamen, and E. Burdet, “A collaborative wheelchair system,” *IEEE Trans. Neural Systems and Rehabilitation Engineering*, vol. 16, no. 2, pp. 161–170, 2008.
- [23] Q. Zeng, E. Burdet, and C. L. Teo, “Evaluation of a collaborative wheelchair system in cerebral palsy and traumatic brain injury users,” *Neurorehabilitation and Neural Repair*, vol. 23, no. 5, pp. 494–504, 2009.
- [24] G. Peinado, C. Urdiales, J. M. Peula, M. Fernandez-Carmona, R. Annichiarico, F. Sandoval, and C. Caltagirone, “Navigation skills based profiling for collaborative wheelchair control,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 2229–2234.
- [25] T. Carlson and Y. Demiris, “Collaborative control for a robotic wheelchair: evaluation of performance, attention, and workload,” *IEEE Trans. Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 42, no. 3, pp. 876–888, 2012.
- [26] Q. Li, W. Chen, and J. Wang, “Dynamic shared control for human-wheelchair cooperation,” in *Proc. IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 4278–4283.
- [27] Z. Wei, W. Chen, and J. Wang, “3d semantic map-based shared control for smart wheelchair,” in *Intelligent Robotics and Applications*, 2012, pp. 41–51.
- [28] J. F. Kelley, “An empirical methodology for writing user-friendly natural language computer applications,” in *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI)*, 1983, pp. 193–196.
- [29] L. D. Riek, “Wizard of Oz studies in HRI: A systematic review and new reporting guidelines,” *Journal of Human-Robot Interaction*, vol. 1, no. 1, pp. 119–136, 2012.
- [30] N. M. Fraser and G. N. Gilbert, “Simulating speech systems,” *Computer Speech & Language*, vol. 5, no. 1, pp. 81–99, 1991.
- [31] J. L. Wolff, B. Starfield, and G. Anderson, “Prevalence, expenditures, and complications of multiple chronic conditions in the elderly,” *Archives of Internal Medicine*, vol. 162, no. 20, pp. 2269–2276, 2002.
- [32] D. R. Dawson, R. Chan, and E. Kaiserman, “Development of the power-mobility indoor driving assessment for residents of long term care facilities,” *Canadian Journal of Occupational Therapy*, vol. 61, no. 5, pp. 269–276, 1994.
- [33] W. B. Mortenson, W. C. Miller, J. Boily, B. Steele, L. Odell, E. M. Crawford, and G. Desharnais, “Perceptions of power mobility use and safety within residential facilities,” *Canadian Journal of Occupational Therapy*, vol. 72, no. 3, pp. 142–152, 2005.
- [34] R. H. Wang, A. Korotchenko, L. Hurd Clarke, W. B. Mortenson, and A. Mihailidis, “Power mobility with collision avoidance for older adults: User, caregiver and therapist perspectives,” *Journal of Rehabilitation Research and Development*, 2013.