



The effect of robotic wheelchair control paradigm and interface on user performance, effort and preference: An experimental assessment



Ahmetcan Erdogan^{a,d}, Brenna D. Argall^{a,b,c,d,*}

^a Physical Medicine and Rehabilitation, Northwestern University, 710 North Lake Shore Drive Chicago, IL 60611, United States

^b Electrical Engineering and Computer Science, Northwestern University, 2145 Sheridan Road Evanston, IL 60208, United States

^c Mechanical Engineering, Northwestern University, 2145 Sheridan Road Evanston, IL 60208, United States

^d Rehabilitation Institute of Chicago, 345 East Superior Street Chicago, IL 60611, United States

HIGHLIGHTS

- Empirical comparison of multiple shared-control paradigms for a robotic wheelchair.
- Multi-session evaluation on real hardware, with multiple interfaces.
- Study participants with and without spinal cord injury.
- Differences in paradigm preference and performance, influenced by interface.
- Suggestion that should offer multiple control sharing options to end-user.

ARTICLE INFO

Article history:

Available online 12 May 2017

Keywords:

Control sharing
Smart wheelchair
Robotics autonomy
Control interface

ABSTRACT

The exact manner in which control is shared between a human and an autonomous system is a crucial factor for assistive robots that provide physical support to people with severe motor impairments. There has however been little comparative study between different control-sharing paradigms within the field of assistive robotics. We present a control architecture for a robotic “smart” wheelchair that allows for the seamless interchange and evaluation of any number of control-sharing paradigms, and so facilitates comparative study between them. We present an implementation of four control sharing paradigms, and results from a study that compares all four to each other and teleoperation, and moreover using multiple control interfaces and across multiple sessions. Experimental results suggest that (i) task performance metrics differ with each control interface, (ii) performance increases with increasing autonomy assistance however it is not statistically significantly different between higher levels of autonomy, (iii) metrics related to user effort show a decrease with increasing autonomy, which is more emphasized with more limited control interface and (iv) how much the autonomy is utilized differs greatly between control paradigms, but not control interfaces. Moreover, (v) for almost all performance metrics, there is a consistent performance increase in Session 2 compared to Session 1, and for both control interfaces. Lastly, subjective questionnaires (control paradigm preference and perceived utility) reveal both similarities and differences between SCI and uninjured subjects. No single control paradigm is the clear winner in performance or preference, suggesting that it will be important to offer end-users multiple control options to accommodate their individual needs and preferences.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

The potential for robotics technologies to transform the lives of those with severe motor impairments is impressive. For many people with severe paralysis, the very assistive machines meant to

* Corresponding author at: Electrical Engineering and Computer Science, Northwestern University, 2145 Sheridan Road Evanston, IL 60208, United States.

E-mail addresses: ahmetcan.erdogan@northwestern.edu (A. Erdogan), brenna.argall@northwestern.edu (B.D. Argall).

improve their quality of life are in practice a burden to operate, or entirely inaccessible—because of the limited control interfaces that are available to them [1]. Robotics autonomy is able to offload some of this control burden, by sharing control with the human operator. However, each operator is unique in their motor abilities, personal preferences and required or desired levels of assistance. From the myriad of ways to implement shared-control, the method which is most *optimal* according to typical robotics metrics may not be the most *accepted* by the user. This choice may be due not only to users preferring to retain as much control as they are able,

but also because the performance of control sharing paradigms crucially depends on the source of the control signals—that is, on both the artificial intelligence control paradigm that generates the autonomy signal as well as the signal from human operator, which depends on their motor abilities and choice of control interface.

Within the literature, the most common studies of physically assistive shared-control robots validate whether the proposed assistive controller is superior to direct teleoperation [2], or whether full autonomy is superior to shared control [3]. There has been limited comparative study between different control-sharing paradigms within the field of smart wheelchairs specifically, and assistive robotics in general. It is, however, reasonable to expect that different control sharing paradigms will perform differently depending on the control interface that is available to the end-user and the motor impairment of the patient.

We present in this paper a control architecture for a robotic “smart” wheelchair that allows for the seamless interchange and evaluation of any number of control-sharing paradigms. Our autonomy system is designed to be strap-on to any commercial power wheelchair system. We provide an overview of our hardware architecture, and its integration with the electronics of a commercial wheelchair.

Our architecture not only allows for the interchange of multiple control-sharing paradigms, but furthermore works immediately with various commercially-available control interfaces—including those accessible to people with severe paralysis. Our control architecture thus makes possible *comparative study* between different control formulations and using different control interfaces. Performing such studies is the only way to systematically assess and evaluate which are the superior performing and preferred ways to share control between physically assistive robots and humans, or if such a superiority even exists. *Especially* when the support provided is physical—that the robot is attached to, or supporting, the human’s body—getting the control sharing right is essential.

Our experimental work evaluates five control paradigms that represent different levels of autonomy in a two-session experiment with seven subjects with Spinal Cord Injury (SCI) (at the C4-5-6, C5-6, C4-5, C5, C6, C6, and C7 levels) and seven subjects without injury. These control paradigms are chosen from common paradigms published within the literature that assist a user in driving a robotic wheelchair. Our aim is to provide systematic comparisons of task performance metrics, user effort and user preference as the (i) navigation assistance paradigm, (ii) injury condition and (iii) control interface are modulated. Learning effects over multiple sessions also are evaluated.

The rest of the paper is organized as follows. Section 2 provides a review of related literature in smart wheelchair research. Section 3 details the control and hardware architectures of our NURIC SMART WHEELCHAIR. The four implemented autonomous assistance paradigms are described in Section 4. The experimental protocol and results are provided in Sections 5 and 6 respectively. Section 7 concludes the paper.

2. Background

The potential for robotics autonomy to aid in the operation of powered wheelchairs has been recognized for decades [4], and a survey of epidemiological data estimates that between 1.4 and 2.1 million individuals would benefit from the use of a smart wheelchair [5]. Recent years in particular have been marked by the development of control paradigms that, rather than being fully autonomous, share control between the human and the robot—in part because keeping the human in the loop can lead to more robust operation, and also in acknowledgment of the reality that most end-users wish to retain as much control as is possible [6,7]. We focus our review of the related literature on smart wheelchairs

to the question of control, and specifically on how control is allocated between the human and the autonomy.

Fully autonomous smart wheelchair systems are effective in achieving high-level goals for outdoor navigation [8,9] and vision-based navigation [10–12]. For fully autonomous executions, it generally is necessary for the human to provide the navigation goal. Mechanisms include brain-machine interfaces [13], face-pose recognition [14], deictical controllers [15], touch screens [16] and graphical user interfaces [17]. In all of these examples, once the goal is provided by the user, the robotic system afterwards operates entirely autonomously—the human provides no further input. It is however possible to override the autonomy by switching to teleoperation in some implementations [14,18]. In our comparative study, we include an approach that operates as a discrete switch between full autonomy and teleoperation, depending on the user’s signal to relinquish the control and the autonomy’s confidence in the user’s intent (Section 4.4).

Rather than relying on the user only at the onset, other approaches have the operator provide input throughout the entire smart wheelchair trajectory. (Note, however, that the feasibility of providing continual control input depends on the user’s motor impairment.) When the robot and the user provide signals in the same control space, then it is necessary to somehow reason between the two signals—whether by fusion or arbitration—and create a single control output.

One approach is to *continually blend* the user and autonomy signals in a weighted sum, where the blending ratio may be constant or changed with respect to metrics like comfort, transparency or safety [19–21] and compared with *a priori* generated baseline commands [22,23]. In our comparative study, we evaluate a similar blending approach, where the user and autonomy signals are summed with changing weights depending on the safety of the blended signal. Two such formulations are assessed, where one infers immediate user intent from only the teleoperation commands (Section 4.2), and the other infers higher-level intent by additionally processing sensor information about perception goals (Section 4.3).

A fundamentally different approach is to *partition the control space* between the human and the autonomy. On a non-holonomic system such as a wheelchair, the steering angle can be handled by one operator (i.e. human or autonomy) while the other controls linear speed. When the autonomy is in control of the steering angle, generally the aim is to find a heading which is as close to the user input as possible but also is free of collisions [24–27]. Having the autonomy instead control linear speed often is used for safety, where speed is decreased in the face of an imminent collision [28]. The idea of constraining the movement of the robot for safety also is used outside of control-partitioning schemes, for example reducing speed when there are large discrepancies between the user input and safe paths [29,30], or stopping the wheelchair before a collision [31,32] and turning away from obstacles [33–35]. In our comparative study, we evaluate a similar signal filtering approach, where the signal in both control dimensions is filtered for safety (Section 4.1).

The analysis of *multiple assistance paradigms* in a comparative study however is given limited attention in the smart wheelchair literature. One Wizard-of-Oz study compares two filtering methods and full autonomy in simulated domain [36]. However, the “autonomous” signals do not actually involve any robotics autonomy; rather a human experimenter provides these control signals. Moreover, no quantitative data is reported. While using human supervisor input to simulate the autonomy is employed elsewhere (e.g. in studies that learn when to provide assistance [37–39]), studies that involve real hardware and actual autonomy are essential for capturing and evaluating the complete human–robot interaction.

The effectiveness of control sharing with *different control interfaces* likewise is analyzed in a limited number of works. An exception are studies with (non-robotic) wheelchairs driven by Brain Computer Interfaces (BCI), where the performance of the assistance is compared to, for example, head gesture control [40] or intrusive and non-intrusive vision-based interfaces [41]. Also, differences in joystick control between novice and experts users are studied in [42]. We know of no study however that compares different commercially available control interfaces within a robotic wheelchair framework. These control interfaces differ in continuity, bandwidth and ease of use; for example a 2-axis joystick with a continuous 2-D signal versus a sip-and-puff with a discrete 1-D signal. These interfaces are far from perfect, but they are ubiquitous, extensively user-validated and covered by insurance. In our comparative study, therefore, the performance of and preference between control paradigms is evaluated with both a 2-D continuous interface and 1-D discrete interface.

In summary, smart wheelchair research is a rich field in which multiple ways to allocate control between the human and autonomy are proposed, implemented and evaluated. However, we know of *few studies* within the literature that explicitly compare multiple smart wheelchair control paradigms within an end-user study, of which *none* include multiple control interfaces and real hardware. We suspect that an optimal control-sharing paradigm might even be *user-specific*, and dependent on characteristics like impairment level (which could render some control paradigms inaccessible), chosen control interface (which dictates the control signal dimensionality and continuity), personal preference (for more or less independence or assistance), therapy requirements (which might encourage active participation by the user to facilitate recovery) and fluctuating user ability (which might decrease or increase over time with illness or rehabilitation).

To address these concerns, our approach allows for the control-sharing paradigm to be *customized* to each individual user. This customization is facilitated by a *modular software architecture*, that easily allows for the various control-sharing paradigms to be turned on or off—and where each paradigm furthermore is parameterized by tunable knobs. Moreover, we leverage this software modularity to *implement and assess multiple control-sharing paradigms* to investigate quantitative differences between control-sharing paradigms and under the control of two input interfaces. In order to begin to assess the effects of learning artifacts, the experiments are repeated in a second session, by both uninjured and SCI volunteers.

3. The NURIC SMART WHEELCHAIR

Here we present the control and hardware architecture of our NURIC SMART WHEELCHAIR.

3.1. Control architecture

The control framework [43] consists of a modular system of software components which implement a set of autonomous behaviors—broadly generalized as high-level and low-level behaviors—and a set of functions that arbitrate between these behaviors (Fig. 1). The user is able to select a custom set of high-level behaviors, low-level behaviors and arbitration functions. Each moreover is parameterized in a way to be customizable to each user.

In particular, there exists a set \mathcal{F}_{hi} of high-level behaviors $f_{hi}(\cdot)$, that each output a goal g

$$g \leftarrow f_{hi}(\mathbf{x}) \quad (1)$$

based on the current state, \mathbf{x} , of the robot and the observable environment.

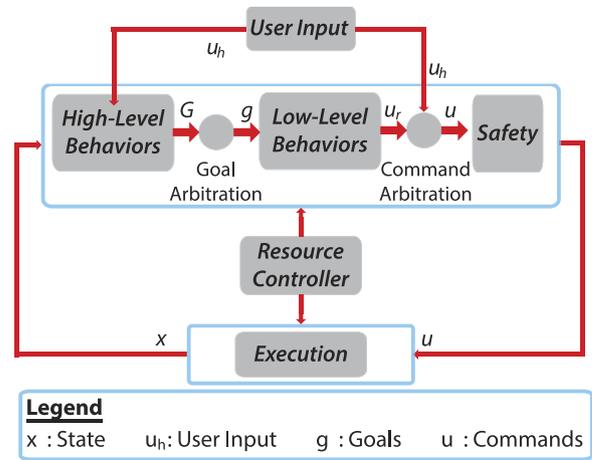


Fig. 1. Control architecture of the NURIC SMART WHEELCHAIR. High-level behaviors feed goals to low-level behaviors, whose output goes through a safety check. Goal arbitration reasons between multiple candidate goals, while command arbitration reasons between the autonomy and human commands.

At each execution cycle goals are generated—multiple, one or no goals, depending on which behaviors f_{hi} are active and whether any autonomously perceivable goals exist within the environment. The set of generated goals is \mathcal{G} (which might be empty). Perception algorithms autonomously perceive goals for tasks identified by end-users as difficult [1] to negotiate due to tight spatial constraints, such as traversing doorways [44], docking at tables and desks [45] and driving up ramps [46]. The *goal arbitration* module reasons between the goals $g \in \mathcal{G}$ —checking for conflicts, feasibility, perception confidence and, most importantly, agreement with the human-generated commands \mathbf{u}_h . A confidence c_g is computed for each element $g \in \mathcal{G}$. From this, the highest-confidence goal $g^* \in \mathcal{G}$ is selected. (Or no goal, if none of the c_g are above threshold.)

There also exists a set \mathcal{F}_{lo} comprised of low-level behaviors $f_{lo}(\cdot)$ that output a control command \mathbf{u}_r

$$\mathbf{u}_r \leftarrow f_{lo}(\mathbf{x}, g) \quad (2)$$

based on a goal g and state \mathbf{x} .

At each execution cycle, a single command \mathbf{u}_r is generated from each active behavior f_{lo} operating on the most confident goal g^* —for example, f_{lo} might consist of the generation of a trajectory able to achieve g^* while also avoiding obstacles, and then a command \mathbf{u}_r able to drive that trajectory. The *command arbitration* module reasons between the autonomy-generated command \mathbf{u}_r and the human-generated command \mathbf{u}_h . Specifically, arbitration function $\beta(\cdot)$

$$\mathbf{u} \leftarrow \beta(\mathbf{u}_h, \mathbf{u}_r) \quad (3)$$

generates command \mathbf{u} that is executed by the robot system. All of the control commands $(\mathbf{u}, \mathbf{u}_h, \mathbf{u}_r)$ consist of two components: translational speed v and rotational speed ω .

Lastly, there exists a *resource controller* to resolve conflicts between competing behaviors. The resource controller registers the data and control signal needs of each behavior, as well as what data that behavior provides to the system as a whole.

The entire architecture is implemented within the Robot Operating System (ROS) [47]. Each behavior is a ROS node, and each arbitration paradigm is specified as a parameter managed by a dynamic parameter server (ROS dynamic reconfigure). Therefore, the ability to create a custom set of behaviors is accomplished simply by specifying the node names in the launch file, and the customization of which type of arbitration is set simply through

a flag in the parameter server. All behaviors register with the resource controller as they are launched, and errors are thrown in the case of any irreconcilable conflicts (e.g. two behaviors claiming the same interface buttons for different uses). The ability to generate a custom set of behaviors and arbitration functions thus is seamless and immediate.

3.2. Control interfaces

It is extremely important (and often overlooked) to consider the different forms that the user-generated signal u_h might take. Depending on a person's type and level of motor impairment, different control interfaces are available to them. The most common is the 2-axis joystick, which provides a continuous-valued 2-D signal. However, for those with severe motor impairments and/or paralysis, the interfaces available to them are much more limited—for example, the sip-and-puff which is operated by issuing discrete 1-D signals via respiration (as hard and soft inhalations and exhalations). Not only are these control signals lower dimensional and discrete, they also are issued at a much slower rate.

Characteristics of the user's control signal necessarily impact how they interact with the autonomous control system. Thus, the performance of the human-robot system as a whole is impacted by this control interface. It might even be the case that which control-sharing paradigm performs best is specific to the type of control interface used by the human. Moreover, it would not be surprising to see a correlation between the effort and difficulty in generating control signals (i.e. through more limited interfaces) and the amount of desired assistance—where the more difficult it is for the human to operate the interface, the more assistance from the robotics autonomy is welcome.

While there are numerous commercial interfaces for operating powered wheelchairs, we can categorize the form of the signals they produce largely into the following groups:

- Continuous, covers the full (2-D) control space, high bandwidth. (e.g. hand- and head-operated joystick).
- Discrete, covers part (1-D) of the control space, low bandwidth. (e.g. sip-and-puff, switch-based headrest).

Our system is able to handle both types of human input. Each input group again is implemented in our system as a distinct ROS node, and how they publish information takes into consideration constraints on the signal (e.g. repeatedly publishing a discrete low-bandwidth signal). The behaviors thus are able to be largely agnostic to the form of the control interface. There are exceptions—for instance, reasoning about agreement between u_p and u_r should only happen within dimensions where the human is actively issuing control commands. To handle these exceptions, the type of control interface also is a flag in the parameter server, and so interface constraint information is readily available throughout the system.

3.3. Hardware

The hardware design adopted in this work prioritizes customization, modularity, low-cost and the use of commercial hardware to facilitate practical adoption by users.

Modifying commercially-available powered wheelchairs is common within the smart wheelchair literature [12,48–51]. Commercial systems have the advantage of being extensively validated by thousands of users, and the fact that expensive parts are covered by insurance provides financial feasibility. Our NURIC SMART WHEELCHAIR is built on a commercially available powered wheelchair, a Permobil C300 (Timra, Sweden), which we then outfit with additional components including a computer, electronics and sensors (Fig. 2).



Fig. 2. NURIC SMART WHEELCHAIR hardware configuration. The base system consists of an RGB-D sensor, a mini-PC, converter boards and wheel encoders. Additional hardware can be added based on a user's needs and preferences.

To interface with the wheelchair electronics, we leverage the fact that all mid- and top-range wheelchair lines from all of the major manufacturing companies offer an electronics box (*expandable input*) add-on that communicates directly with the wheelchair's firmware. While these expandable interfaces were developed to allow for the use of third-party control interfaces to drive the chair, we can exploit this technology to pass our control command u directly to the proprietary wheelchair motion controller. Specifically, our computer-generated control commands can be modulated to *mimic* a regular inductive joystick and passed directly to the wheelchair control. Therefore, our system is easily applied to any mid- to top-range wheelchair from all of the major manufacturing companies. The expandable input device on our NURIC SMART WHEELCHAIR is the OMNI interface from R-Net Electronics (Christchurch, UK).

Hardware add-ons to the base system include on-board computing (mini-PC), electronics boards, override buttons, wheel encoders and a top-mounted RGB-D sensor (Asus Xtion). The on-board computing system is directly powered by the wheelchair batteries.

Control interfaces, through which the human provides input to drive the wheelchair, are directly connected to the mini-PC computer via USB. An Arduino board is used to read external button interactions, while an analog board is used to mimic the inductive interface voltage when passing control signals to the wheelchair controller. The only custom piece of hardware (besides mounting pieces) on our system are the wheel encoders.¹ Wheel encoders typically are not available on commercial wheelchairs—however most autonomous controllers for mobile robots rely on knowing how far the robot has moved. Our custom encoder is low-cost and consists of auxiliary wheels that are in contact with the inner rim of the wheelchair wheels, mounting parts and small encoders that measure the orientation of these auxiliary wheels. The Arduino board is used to count the encoder ticks that provide quadrature outputs. All the additional (e.g. mounting) parts are printed by 3D printers.

Additional sensing elements can be added—for additional cost, but also additional functionality and robustness. From a software standpoint, each sensor simply is a ROS node which registers the information it offers with the resource controller, and becomes available to behaviors and environmental constructs (e.g. navigation costmaps) which utilize that type of information. A schematic of the hardware architecture is provided in Fig. 3.

¹ In order to not violate ADA standards, we did not mount encoders on the outside of the wheels (which would increase the width of the wheelchair). All commercially-available (and reasonably priced) encoders were just barely too big for mounting on the interior of the wheel. We do however expect a commercial solution in the near future.

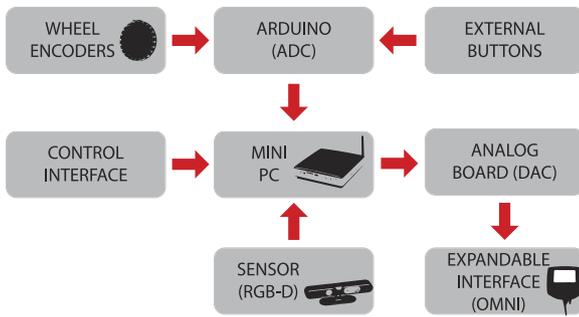


Fig. 3. Hardware schematic. The mini-PC receives sensor information and human-generated commands directly, and encoder and button-press information through an Arduino board. Control commands are calculated and sent to the expandable (OMNI) interface via a digital analog converter.

3.4. Safety

We wrap a safety check around our control loop that constantly monitors the state of the world, specifically for collisions, before any commands are sent through the OMNI interface. Control commands (passed through the OMNI interface) then are executed by the proprietary wheelchair electronics. This eliminates the need for a stability check on the low-level controller, since any control commands are executed via the electronics of the commercial wheelchair which includes multiple safety checks. (The alternative is to send control signals directly to the motors of the wheelchair [51], which does require stability checks but as a trade-off is able to mitigate the noise sensitivity of a feedback controller, for example due to external disturbances from caster wheel forces.)

As a final safety mechanism, a manual override button on the armrest disengages the autonomy and turns the system to a regular manually-driven powered wheelchair.

4. Navigation assistance paradigms

We focus in this work on how control is shared with the human. It is unlikely that a single control-sharing paradigm will be best-suited for all assistance scenarios, not to mention for all users and their unique abilities and preferences. We anticipate that flexibility and variability in how control is shared with the human will be a fundamental threshold for the adoption of smart wheelchair technologies within larger society.

We describe here four paradigms for control sharing implemented on our NURIC SMART WHEELCHAIR. The purpose of each is to compute a control input $\mathbf{u} = [v, \omega]$ for the system to execute, that takes into consideration user signal $\mathbf{u}_h = [v_h, \omega_h]$, autonomy signal $\mathbf{u}_r = [v_r, \omega_r]$ and environment information. A schematic representation of the four control-sharing paradigms is given in Fig. 4.

4.1. Signal filtering with immediate goals

The operation of this assistance paradigm is to filter the user signal based on the output of an autonomous planner, simply to avoid collisions. The autonomy signal (\mathbf{u}_r) is derived from a local controller executing a collision-free path planned to the *immediate goal*. The immediate goal is calculated (via forward projection for time Δt) from the user's current command (\mathbf{u}_h), and evaluated for safety—by checking whether the immediate goal lies within an obstacle. (In our experimental work, $\Delta t = 0.5$.²) If found to be

unsafe, the user signal (speed v_h , heading ω_h) is *filtered* to be no greater than the autonomy command [29,30]. Thus,

$$\beta(\mathbf{u}_h, \mathbf{u}_r) \triangleq \min(\mathbf{u}_h, \mathbf{u}_r). \quad (4)$$

Pseudo-code for this paradigm is given in Algorithm 1.

Algorithm 1: Filtering with Immediate Goals

Autonomy
 $g \leftarrow \text{ForwardProjection}(\mathbf{u}_h, \Delta t)$
 $\mathbf{u}_r \leftarrow \text{Planner}(g)$
Arbitration
if $\text{CheckCollision}(\mathbf{u}_h) \wedge \mathbf{u}_r < \mathbf{u}_h$ **then**
 $\mathbf{u} \leftarrow \mathbf{u}_r$
else
 $\mathbf{u} \leftarrow \mathbf{u}_h$

4.2. Signal blending with immediate goals

The operation of this assistance paradigm is to blend the user signal with the output of an autonomous planner, again simply to avoid collisions. The autonomy signal (\mathbf{u}_r) is derived from a local controller executing a collision-free path planned to the *immediate goal*. The immediate goal is calculated as in Section 4.1. If found to be unsafe, the user signal (speed v_h , heading ω_h) is linearly *blended* with the autonomy command [19–21] in an iterative manner, that steps control away from the user based on *safety* constraints. Thus,

$$\beta(\mathbf{u}_h, \mathbf{u}_r) \triangleq \alpha \cdot \mathbf{u}_h + (1 - \alpha) \cdot \mathbf{u}_r \quad (5)$$

where blending parameter α is iteratively decremented (by $\Delta\alpha$) until the forward projection of blended command \mathbf{u} is safe, with safety being calculated as in Section 4.1. (In this study $\Delta\alpha = 0.25$, an empirically-determined compromise between computational efficiency and allowing the user to remain as in control as possible.) Pseudo-code for this paradigm is given in Algorithm 2.

Algorithm 2: Blending with Immediate Goals

Autonomy
 $g \leftarrow \text{ForwardProjection}(\mathbf{u}_h, \Delta t)$
 $\mathbf{u}_r \leftarrow \text{Planner}(g)$
Arbitration
 $\alpha \leftarrow 1$
 $\mathbf{u} \leftarrow \mathbf{u}_h$
while $\text{CheckCollision}(\mathbf{u})$ **do**
 $\mathbf{u} \leftarrow \alpha \cdot \mathbf{u}_h + (1 - \alpha) \cdot \mathbf{u}_r$
 $\alpha \leftarrow \alpha - \Delta\alpha$

4.3. Signal blending with a high-level goal

The operation of this assistance paradigm is to blend the user signal with the output of an autonomous planner, in order to achieve a high-level goal (and maintain safety). The autonomy signal (\mathbf{u}_r) is derived from a local controller executing a collision-free path planned to a *high-level goal*. The high-level goal is detected through perception algorithms that process (RGB-D) sensor information, and is inferred to be the user's goal through a *confidence measure* based on agreement with the user's current command (\mathbf{u}_h). The confidence c_g associated with an observed goal g is calculated based on the distance d and heading ϕ to the current goal g .

$$c_g = c_p \cdot \left(\gamma \cdot \left(\frac{2}{1 + e^\phi} \right) + (1 - \gamma) \cdot \left(\frac{2}{1 + e^d} \right) \right) \quad (6)$$

where c_p is the perception confidence observing the same doorway goal multiple times. The parameter γ dictates that when the robot

² With a maximum linear velocity of 0.352 m/s, a 0.5 s look-ahead provides sufficient distance for the controllers to safely respond to obstacles.

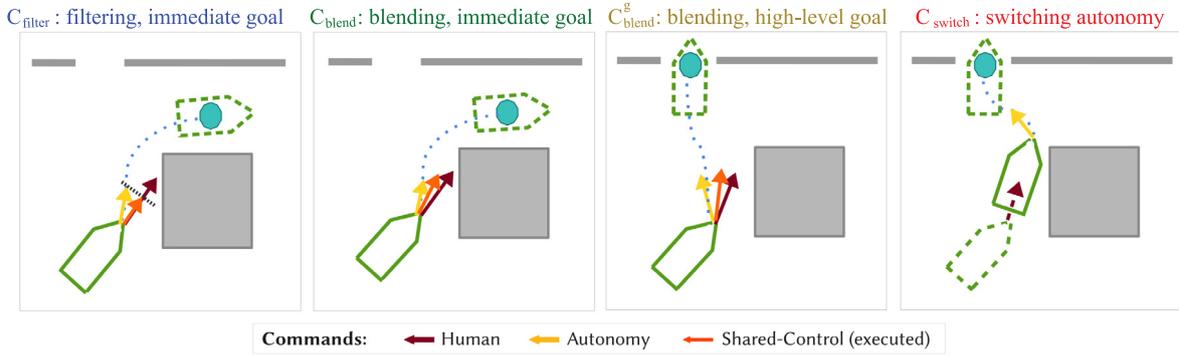


Fig. 4. The four control-sharing paradigms of our study which differ according to how the autonomy commands are generated (immediate or perception goals) and how the control is shared between user and autonomy (filtering, blending, or switching). In particular, navigation goals (blue circle) for the autonomy are inferred simply from a brief (0.5 s) forward projection of the human’s current control command (as in C_{tele} and C_{blend}) or from a higher-level perception goal (doorway) detected from sensor data (as in C_{blend}^g and C_{switch}). In either case, the same planner is used to generate a path (dashed blue line) to the goal, and the same controller is used to drive that path. Doorway shown as a gap in the top gray line, robot footprint shown as a green outline and obstacle as a gray box.

is far from g ($d > 1$ m), robot’s heading with respect to the goal has higher importance, while distance matters most when the robot is near the goal. In our implementation, a goal is above threshold when its confidence is at least 80% and also at least 25% higher than any other goals (empirically determined). Readers are referred to [43] and [52] for further details.

If the confidence in the goal is above threshold, and the user is issuing commands, the user signal (speed v_h , heading ω_h) is linearly blended with the autonomy command in an iterative manner, that steps control away from the user based on (Euclidean) distance to the goal. Thus,

$$\beta(\mathbf{u}_h, \mathbf{u}_r) \triangleq \alpha \cdot \mathbf{u}_h + (1 - \alpha) \cdot \mathbf{u}_r \quad (7)$$

where blending parameter α is continuously decremented as the wheelchair gets closer to the goal g . Pseudo-code for this paradigm is given in Algorithm 3. In our implementation $\text{DistToGoal}(g)$ function is: $\alpha = 1/(1 + e^{-\tau_c \cdot (d - d_c)})$, where time constant $\tau_c = 3$ and distance coefficient $d_c = 1.5$ m are empirically determined.

If there is no high-level goal ($\neg \text{Confident}(g)$)—due to either no candidate goal being perceived or the confidence in the perceived goal being low—the blending paradigm based on safety, as described in Algorithm 2, is performed. The key differences between Algorithms 2 and 3 are the use of a higher-level perception goal and how blending happens near that goal. Many smart wheelchair systems aim to provide assistance in achieving higher-level goals [8,18], and the introduction of a perception system for that purpose [10,11] also introduces additional complexity and uncertainty which might affect control-sharing performance and preference.

Algorithm 3: Blending with a High-Level Goal

Autonomy
 $g \leftarrow \text{Perception}(\text{sensors})$
 $\mathbf{u}_r \leftarrow \text{Planner}(g)$
Arbitration
if $\text{Confident}(g) \wedge \mathbf{u}_h \neq \emptyset$ **then**
 $\alpha \leftarrow \text{DistToGoal}(g)$
 $\mathbf{u} \leftarrow \alpha \cdot \mathbf{u}_h + (1 - \alpha) \cdot \mathbf{u}_r$
else
 Algorithm 2

4.4. Autonomy switching with a high-level goal

The operation of this assistance paradigm is to filter the user signal entirely, in order to achieve a high-level goal (and maintain safety). The autonomy signal (\mathbf{u}_r) is derived from a local controller

executing a collision-free path planned to high-level goal, detected as in Section 4.3. If the confidence in the goal is above threshold, and the user stops issuing commands, a control switch occurs and the autonomy command is executed in full [10,11,18]. Thus,

$$\beta(\mathbf{u}_h = \emptyset, \mathbf{u}_r) \triangleq \mathbf{u}_r. \quad (8)$$

Pseudo-code for this paradigm is given in Algorithm 4.

If there is no high-level goal ($\neg \text{Confident}(g)$) as in Section 4.3, the blending paradigm based on safety, as described in Algorithm 2, is performed.

Algorithm 4: Autonomy Switching

Autonomy
 $g \leftarrow \text{Perception}(\text{sensors})$
 $\mathbf{u}_r \leftarrow \text{Planner}(g)$
Arbitration
if $\text{Confident}(g) \wedge \mathbf{u}_h = \emptyset$ **then**
 $\mathbf{u} \leftarrow \mathbf{u}_r$
else
 Algorithm 2

5. Control sharing experiment

In this experiment, subjects were asked to perform a closed loop in our laboratory which included four doorway traversals. Several of the previously proposed smart wheelchair systems in the literature include doorway traversal as an evaluation task [17,21,25,53–57]. Additionally, it is one of the tasks identified in the Powered Wheelchair Skills Test [58] used by clinicians to establish whether an individual has the requisite skills to safely pilot a powered wheelchair. Traversing a doorway is challenging because the clearance on either side of the wheelchair is only approximately 10 cm according to ADA standards, which means a high level of steering accuracy, and by extension visual acuity and manual dexterity, are required. The tight tolerances are challenging for autonomous robotic navigation as well.

5.1. Hypotheses

Our experiment aimed to evaluate the following hypotheses:

- H1: User preference and performance metrics do not provide experimental evidence for the superiority of any single navigation assistance paradigm.
- H2: Performance and preference change based on the control interface.



Fig. 5. Experiment control interfaces. Left: Joystick and manual override buttons. Right: Switch-based headrest array.

- H3: Performance differences between control paradigms are less pronounced when operating a continuous, high-bandwidth interface (i.e. the joystick), as compared to a discrete, low-bandwidth interface (i.e. the head array).
- H4: Participants prefer paradigms with more assistance when using a more limited control interface.
- H5: Uninjured subject performance is significantly different than SCI performance.
- H6: Performance improves between sessions, for both subject populations, because of human learning.
- H7: Reliance on the autonomy increases between sessions, for both subject populations, because of familiarization.

5.2. Participants

Experiment participants included seven SCI subjects (36–68 years old) and seven uninjured subjects (23–37 years old). On average, it had been 23.6 ± 11 years since injury for our SCI subjects and they used a powered wheelchair for 21 ± 11.4 years. Most of this experience was with a 2-axis joystick, though one SCI subject did drive the wheelchair with a sip-n-puff for his first five years post-injury (and after with a joystick for 20 years). The uninjured subjects had varying experience with robotic systems but were mostly naive to powered wheelchair driving. All subjects indicated a high comfort level with technology (7-point Likert scale, SCI: 6.52 ± 0.60 and Uninjured: 6.29 ± 0.80).

The experiment was approved by the Northwestern University Institutional Review Board (STU00201743) and informed consent was obtained from all subjects. An occupational therapist was present for all SCI sessions, who oversaw subject transfer into our NURIC SMART WHEELCHAIR and constantly monitored the subject's condition.

Each participant performed the task in two sessions on separate days. The experiment was expected to take two hours, which is beyond the sitting tolerance of some SCI subjects. For this reason, upon the recommendation of the therapist, one SCI subject split both sessions into two (for a total of four sessions), and a second SCI subject split Session 1 into two (for a total of three sessions).

5.3. Control paradigms and interfaces

The five control paradigms were: teleoperation (C_{tele}), signal filtering (C_{filter}), signal blending with immediate goals (C_{blend}), signal blending with a high-level goal (C_{blend}^g) and autonomy switching (C_{switch}). Note that teleoperation had no assistance from any autonomy paradigm (100% human control).

The two control interfaces were: a hand-operated 2-axis joystick and a switch-based headrest array (Fig. 5). The joystick is a proportional interface, where speed scales with the amount of

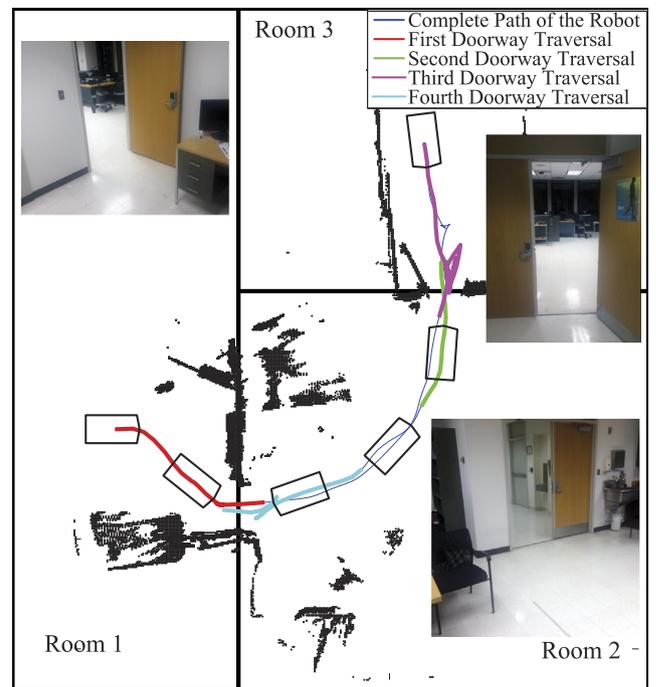


Fig. 6. Sample experimental run under direct teleoperation (C_{tele}). Phases of a single run shown as differently colored lines. Sensor data projected to the ground plane shown as black dots.

joystick deflection. The headrest array is not a proportional interface, and all movements happen at constant speeds. The headrest interface is operated by touching (with the head) the back panel to drive forwards and backwards (toggled by a button mounted on the armrest), and touching the right and left panels to rotate clockwise and counterclockwise, respectively.

The linear speed of the wheelchair was limited to 0.35 m/s, and the angular speed to 0.4 rad/s. Discrete values for the headrest interface were set to 95% of these values.

We also augmented our base hardware configuration with a Hokuyo laser rangefinder, both to validate the RGB-D sensor and increase robustness and safety during the experiment.

5.4. Experimental protocol

The overall experiment was divided into four phases for each door, each of which started two meters away from the door and ended when the user safely traversed the doorway. A sample run is given in Fig. 6 where each phase is color coded.

For each experimental run, the subject began in Room 1, passed through Door 1 into Room 2, and then drove freely to then pass through Door 2 into Room 3. The supervisor then took control to turn the wheelchair around and position it for the second starting position. The progression was then reversed, from Room 3 into Room 2 (via Door 3) and from Room 2 into Room 1 (via Door 4). (Doors 1 and 4 are the same doorway traversed in opposite directions; and likewise for Doors 2 and 3.) Of note is that Room 1 was cluttered, and so obstacles needed to be navigated around when approaching Door 1 and passing through Door 4.

Participants were instructed to drive the wheelchair as they felt comfortable. They could drive backwards, however there was no assistance from the autonomy backwards.³ At the start of the each run, the current control sharing paradigm was introduced verbally.

³ The ratio of backward commands during doorway traversal was less than 1.5% (over all subjects, assistance paradigms and sessions).

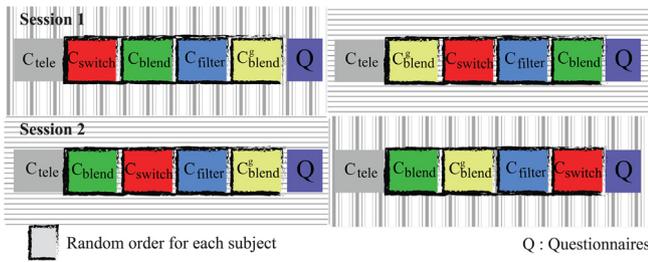


Fig. 7. Schematic representation of the experimental protocol (example ordering for a single subject). A randomized presentation order for each assistance paradigm (colored boxes) was defined for each subject prior to the experiment.

At any point during the experiment, participants could switch (via a button press) to fully manual operation if they did not feel comfortable with the autonomy. Also, a secondary control interface operated by the experimenter could be used to take over control at any time for safety.

Each control paradigm was utilized twice with both control interfaces, for a total of 20 trials per participant in a session. All sessions began with the teleoperation base condition, C_{tele} , and the presentation order of the remaining navigation assistance paradigms followed a predefined randomization. The allocation of which interface was used first also was randomized across subjects. Upon completing all sessions with the initial control interface, subject filled out the questionnaires. The second half of the session used the remaining control interface, again starting with C_{tele} and using a different set of randomized control paradigms (Fig. 7).

The second day trials were based on the same protocol with a different randomization of presentation order for the navigation assistance paradigms and interchanging the order of the control interface. Thus, a subject that started with the 2-axis joystick in Session 1 would do the headrest array trials first in Session 2.

5.5. Performance metrics

To assess the utility and utilization of the various control paradigms and test the hypotheses, the following metrics were computed:

- **Task Completion Time:** $T = t_N - t_0$, provides a measure of system performance.
- **Contribution from Autonomy:** R is computed differently for each paradigm, and provides insight into the relative utilization of autonomy compared to the user inputs:

$$R = \begin{cases} \frac{1}{N} \sum_{t_0}^{t_N} \left\| \frac{\bar{\mathbf{u}}_h^t - \bar{\mathbf{u}}^t}{\bar{\mathbf{u}}_h^t} \right\|, & \text{if } C_{filter} \\ \frac{1}{N} \sum_{t_0}^{t_N} \left\| \frac{\bar{\mathbf{u}}^t}{\bar{\mathbf{u}}_h^t} \right\|, & \text{if } C_{blend} \\ \begin{cases} R \text{ as in } C_{blend}, & \text{if } |\mathbf{u}_h^t| > 0 \\ 1, & \text{otherwise,} \end{cases} & \text{if } C_{switch} \end{cases}$$

- **Similarity of Commands:** $S = \frac{1}{N} \sum_{t_0}^{t_N} \|\bar{\mathbf{u}}_h^t - \bar{\mathbf{u}}^t\|$, provides insight into the agreement between the autonomous commands and the user inputs.
- **Fluency of User Commands:** $F = \frac{1}{N} \sum_{t_0}^{t_N} (1 - |\bar{\mathbf{u}}_h^t - \bar{\mathbf{u}}_h^{t-1}|)$, shows the continuity of the user signals, and provides insight into how much effort the user issued.
- **Number of Interface Interactions:** $E = \sum_{t_0}^{t_N} (\bar{\mathbf{u}}_h^t > \zeta)$, counts the number of time the user interacts with the control interface.

- **Distance to Obstacles:** $D = \frac{1}{N} \sum_{t_0}^{t_N} \min(\epsilon, 0.4 m)$, averages the minimum distance to obstacles ϵ (capped by 0.4 meter) during doorway traversal.

Here t_0 and t_N are respectively the starting and ending times for each door traversal, $\bar{\mathbf{u}}^t$ is control command \mathbf{u}^t normalized by the maximum speed limits, $\bar{\mathbf{u}}_h^t$ is the similarly normalized user command \mathbf{u}_{ht} , $\|\cdot\|$ represents L_2 -norm and $|\cdot|$ the absolute value, and N is the number of samples. Lastly, ζ represents the dead-zone of the joystick interface covering 25% of the joystick range. Each instance of exiting the dead-zone counted as a single joystick interaction. Interactions with the head array interface were already discrete and therefore counted directly.

All manual-override interactions (by the participant, or the experimenter) also were recorded, as well as any collisions with the door frames or objects in the environment.

5.6. Subjective metrics

After the successful completion of all trials with a given control interface, participants completed a questionnaire to (i) indicate their most and least preferred control paradigms and (ii) evaluate each control paradigm.

In the first questionnaire, participants were asked to indicate their most/least preferred and the most/least useful paradigms. Ties between paradigms were allowed. Subjects furthermore responded to the following questions on a 7-point Likert scale:

- Q1: The assistance from the robot was useful and I can achieve the task more easily with the robot's assistance.
 Q2: The assistance from the robot complemented my abilities and contributed to the success of the task.
 Q3: I am confident in the robot's ability to help me and I trust the robot to do the right thing at the right time.

5.7. Data analysis

All experimental data was collected via the ROS pipeline and the majority was sampled at 25 Hz (with the exception of computationally expensive topics such as the 2-D costmap, which was sampled at 7 Hz). MATLAB was used to segment out the doorway time intervals and to pre-process the data for statistical analysis in IBM SPSS.

For each performance metric in each session (e.g. Task Completion Time in Session 1), one factor repeated measure ANOVA (Analysis of Variance) was performed to determine significant differences ($p < 0.05$) between the five navigation assistance paradigms. Once the significance of the dataset was established, multiple post-hoc pairwise comparisons were performed by using Bonferroni Confidence interval adjustments. In order to analyze differences between sessions, an additional pairwise comparison for each navigation assistance paradigm was performed within both sessions (e.g. Task Completion Times with direct teleoperation in Session 1 versus Session 2).

6. Experimental results and discussion

We evaluate the performance of the four control paradigms and direct teleoperation across two experimental sessions. Experimental results suggest that performance metrics change with the type of autonomy and depending on the control interface. Performance metrics that relate to user effort generally decrease over sessions, with both interfaces. The extent to which users relied on the autonomy did not significantly change between control interfaces, but did between control paradigms. Table 1 furthermore

Table 1
Experimental evidence supporting each hypothesis.

Hypothesis	Supported	Not Supported
H1	✓ Figs. 8–11, 14, 15	–
H2	✓ Figs. 8, 10–15	✓ Fig. 9
H3	✓ Fig. 9	✓ Figs. 8, 10–15
H4	✓ Figs. 14, 15	–
H5	✓ Figs. 10, 12	✓ Figs. 8, 9, 11, 13
H6	✓ Fig. 10	✓ Figs. 8, 9, 11–13
H7	–	✓ Figs. 12, 13

summarizes the hypotheses (Section 5.1) that are either supported or not supported by the experimental evidence.

In the following analysis, we loosely order the control sharing paradigms in terms of *increasing autonomy* as $C_{tele} \rightarrow C_{filter} \rightarrow C_{blend} \rightarrow C_{blend}^g \rightarrow C_{switch}$. The reasoning is that for C_{filter} the user's signal is not augmented but only capped, for C_{blend} the user's signal is augmented but no higher-level inference of their goal is performed as in C_{blend}^g , and for C_{switch} it is possible for the executed signal to have been generated 100% by the autonomy.

For all figures, the notation * implies $p < 0.05$ and the notation ** implies $p < 0.01$. The details of the statistical significance analysis for each performance metric are given in Appendix A.

6.1. User performance and effort under different assistance paradigms and control interfaces

On the topic of user performance, we hypothesized that (i) none of the tested assistance paradigms would be perform significantly superior to other methods (H1), (ii) performance would change based on the control interface even under same assistance paradigms (H2), and (iii) users would benefit from autonomy more when the wheelchair is driven with limited control interfaces (H3).

In general, we find performance metrics to improve with increasing contribution from the autonomy. However, pairwise comparisons between three of the control paradigms (C_{blend} , C_{blend}^g and C_{switch}) generally are not statistically significant. More aggressive controllers do perform slightly better than other assistance paradigms, and this improvement is more pronounced in the headrest array. However, these assistance paradigms also have a tendency to result in higher dissimilarity between the user and autonomy commands, which in practice could affect the effort of the user and maybe their preference. These paradigms further result in a higher variance in the utilization of autonomy and similarity to the user commands.

The following subsections discuss these results in greater detail, and also their impact on our hypotheses.

6.1.1. Task completion time

The Task Completion Time for each control paradigm in each session, averaged over subjects, trials and doors, is provided in Fig. 8. We do not observe a significant difference between the autonomy paradigms, with the exception of C_{filter} , for both interfaces. The similarity in the completion time between these assistance paradigms support hypothesis H1. Unexpectedly, we also do not observe major performance difference between control interfaces in those cases (\neg H3). There is a visibly-higher change with autonomy (top plot, left→right) with head array, however, with no significance due to the extreme variability across subjects.

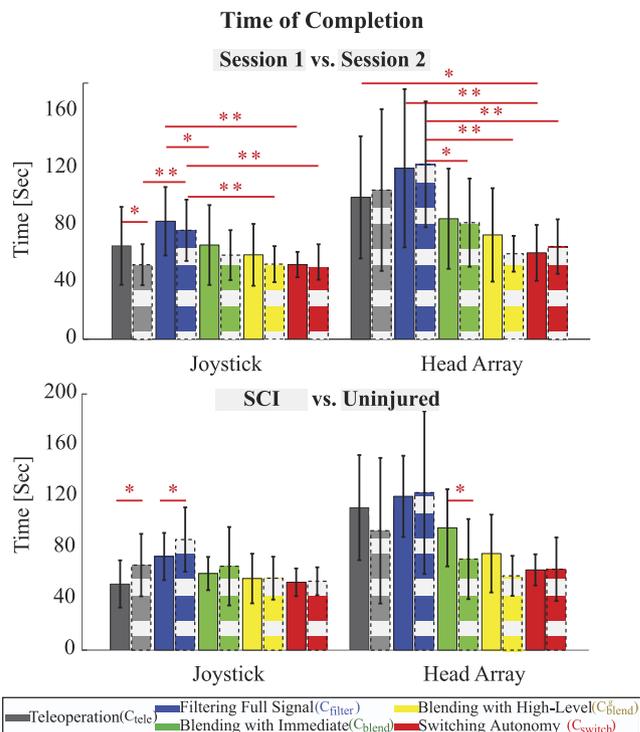


Fig. 8. Task Completion Time under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

6.1.2. Interface interactions

The number of interactions with a control interface is shown to differentiate between expert and novice drivers of a powered wheelchair, with expert drivers shown to have fewer interactions [42]. In a similar manner, this metric contains information with respect to the performance of the control sharing—one could argue that the most capable control sharing algorithm would also require fewer inputs from the user to help them accomplish the task. Moreover, it is a powerful metric to identify the effort of the user—a higher number of interactions requires more participation and attention from the human.

The Number of Interface Interactions metric for each control paradigm in each session, averaged over subjects, trials and doors, is provided in Fig. 9. We see that with the head array in particular significantly fewer interface interactions are required with increasing autonomy (top plots, left→right, $C_{tele} \rightarrow C_{switch}$). This trend is observed roughly equally in each subject group, with few statistically significant differences between SCI and uninjured subjects (bottom plot).

6.1.3. Human command fluency

Similar to Number of Interface Interactions, the Fluency of User Commands metric contains information about user effort and task performance. Somewhat unexpectedly, the head array trials are more fluent than the joystick (Fig. 10). The most probable reason for this is the higher activation cost of each action with the head array interface, and the discrete nature of the interface—when the command is issued (i.e. headrest is touched) for multiple consecutive time steps, the fluency between them by definition is 1 (since $u_h^t = u_h^{t-1}$). These results support hypothesis H2. The more pronounced fluency differences between control paradigms under the joystick however refute hypothesis H3.

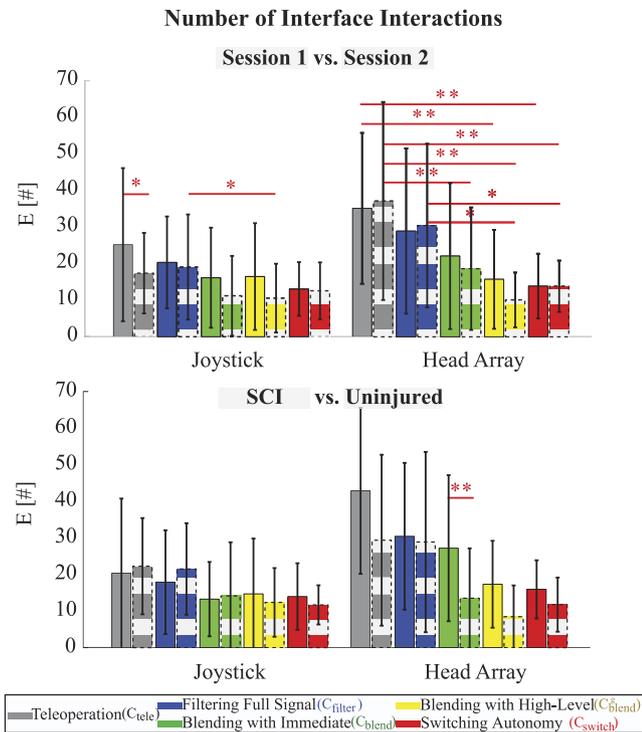


Fig. 9. Number of Interface Interactions under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

6.1.4. Distance to obstacles

Number of Interactions and Fluency of the User Command metrics do not necessarily correlate to the safety of the trials. Fig. 11 reports the Distance to Obstacles for each control paradigm and subject group. Autonomy significantly increases the average distance to the obstacles compared to teleoperation—despite requiring fewer user commands. We do not however observe difference between the various autonomy paradigms or between the different control interfaces.

6.1.5. Contributions from the autonomy

We additionally report in Fig. 13 the Ratio of Autonomy metric for each control paradigm and subject group. In both control interfaces, there is an unsurprising (and sometime significant) trend of increasing contribution from the autonomy across the increasing autonomy ordering of the navigation assistance paradigms—with the exception of the ordering of C_{blend}^g and C_{switch} . In fact, our increasing autonomy ordering of $C_{blend}^g < C_{switch}$ often is reversed, since users often opt to not relinquish control. It would be interesting to observe whether these trends hold in a richer task environment, or with greater familiarity with the autonomy—that is, whether users begin to relinquish more control, or continue to prefer to retain more control.

6.1.6. Summary

In summary, we see performance differences between paradigms (H2) but not that consistently favor one over another. User preference (see Section 6.4 and Fig. 14) shows a general trend towards higher autonomy paradigms, but this was by no means definitive. Overall, from the quantitative data, it is not obvious that a single navigation assistance paradigm is most preferred or performs better with statistical significance (H1). We do observe marked performance differences between the two interfaces, supporting

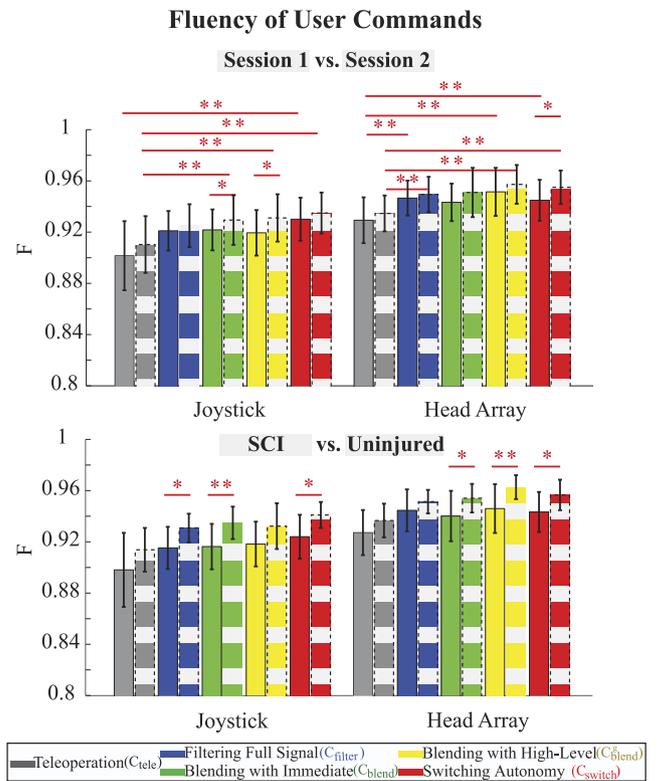


Fig. 10. Fluency of User Commands under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

hypothesis H2. In particular, most task execution metrics (with the exception of the Number of Interface Interactions) significantly differ with control interfaces on either the majority of or all assistance paradigms.

Comparing the relative effect of autonomy under different control interfaces shows that joystick trials benefit less from the autonomy compared to the head array trials (H3). Specifically, the Number of Interface Interactions metric follow steeper improvement trends with the head array. In fact, performance similar to that of the joystick interface is achieved only when using more aggressive autonomy paradigms such as C_{blend} , C_{blend}^g and C_{switch} . We do not however observe this significant behavior with other performance metrics. It would be interesting to see whether these performance metrics might further improve if the autonomy were to explicitly take into account the constraints of the control interface [59].

6.2. Learning under different assistance paradigms and control interfaces

Based on the potential for learning both with and without autonomy, between Session 1 and Session 2, we hypothesized that (i) a significant performance increase would be observed between sessions (H6) and (ii) reliance on the autonomy would increase because of increased familiarity (H7).

Between Sessions 1 and 2, a consistent (sometimes significant) performance increase across all assistance paradigms is observed (H6). These increases are not dependent on the control interface (Figs. 8–10). Such a learning artifact with a robotic assistive device might be due to the shared-control autonomy paradigms requiring time to fully understand and utilize. However, it also

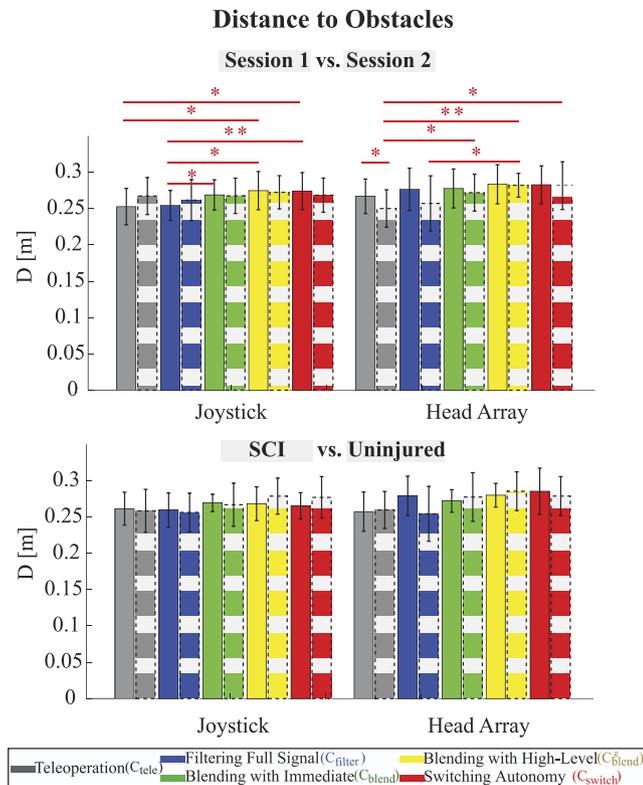


Fig. 11. Distance to Obstacles under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

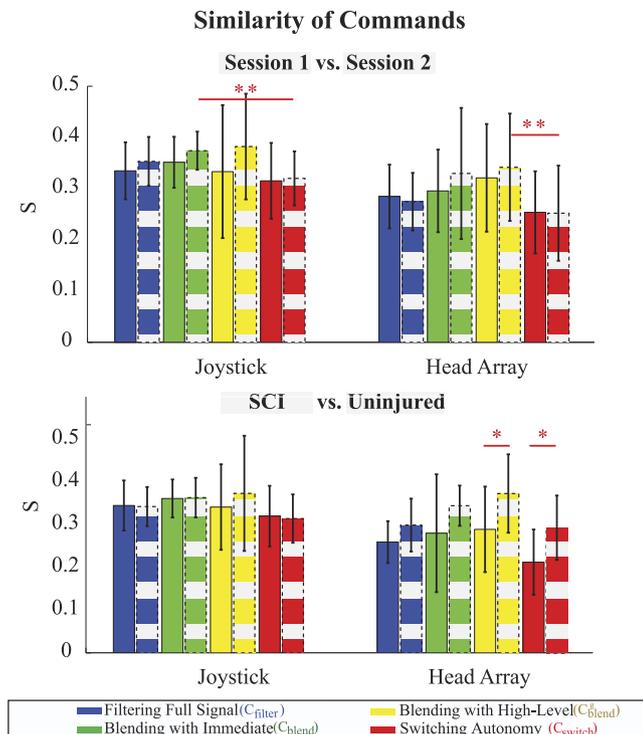


Fig. 12. Similarity of Commands under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

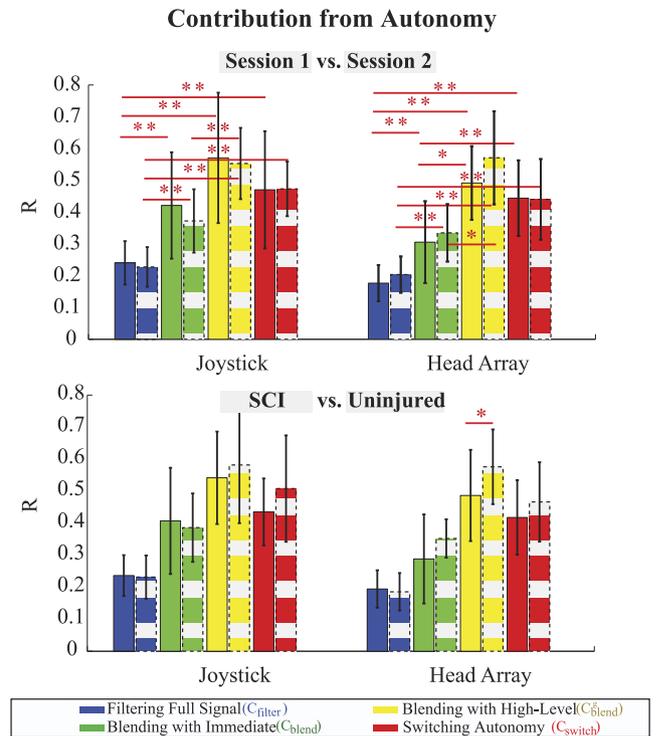


Fig. 13. Contribution from Autonomy under different control paradigms. Top: Session 1 (solid) vs. Session 2 (striped). Statistical comparisons are made between navigation paradigms within a given session and between sessions. Bottom: SCI (solid) vs. uninjured (striped) subjects. Statistical comparisons are made (for each control paradigm) between SCI and uninjured subjects.

means that performance may be increased with practice. It would be interesting to observe whether a more extensive longitudinal study (>2 sessions) would identify a steady-state response in user performance.

With respect to reliance on the autonomy, the most striking difference we see is between C_{filter} and each of the three other autonomy paradigms. These differences are significant in both sessions, and for both interfaces. Between sessions however, we do not observe any differences in reliance, for any paradigm ($-H7$).

6.3. User performance and effort within subject groups

On the topic of differences between subject groups, we hypothesized that a significant performance difference would be observed between SCI subjects and uninjured volunteers ($H5$).

We observe however few differences between subject groups. The one big exception is in command fluency (Fig. 10), for which we see a significant difference between the subject groups for the majority of autonomy paradigms ($H5$)—though, surprisingly, not for teleoperation. The decrease in fluency is likely due to the impairment in the SCI subjects' motor control. In all other metrics, we see significant differences between subject groups for at most 2 of the 5 paradigms, and often for no paradigms ($-H5$).

6.4. User preference under different assistance paradigms and control interfaces

Fig. 14 (top) shows that the SCI subjects' most preferred paradigms when operating a joystick to be similarly distributed to uninjured subjects' for many of the paradigms, with notable differences arising however with teleoperation (favored by SCI subjects) and C_{blend}^g (favored by uninjured subjects). There are no

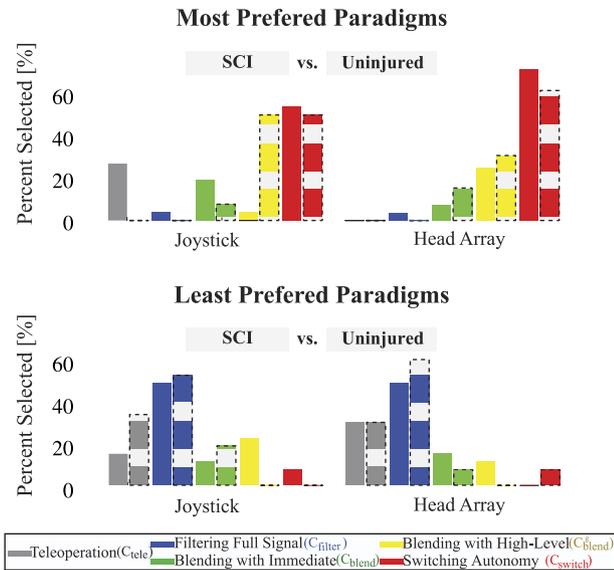


Fig. 14. Most and least preferred navigation paradigms with joystick and headrest array. SCI subjects (solid) vs. uninjured subjects (striped).

such notable preference differences between subject groups when operating the head array. Moreover, for the head array, neither group selected teleoperation as their preference, and more of the distribution weight lies with autonomy paradigms that provide more assistance—which (weakly) supports hypothesis H4.

These general observations hold for the selection of least preferred paradigm as well. That is, a notable difference in preference is observed between the subject groups again for teleoperation (least preferred by uninjured subjects) and C_{blend}^g (least preferred by SCI subjects) when operating a joystick. The distributions are otherwise similar between subject groups.

For both subject groups, when operating either interface, there is a general trend of preferring paradigms with stronger autonomy and not preferring paradigms with weaker or no autonomy. Support for H4 thus is inconclusive: while subjects do prefer paradigms with greater autonomy when using a more limited control interface, they show this same preference when operating a richer interface as well.

Fig. 15 presents results for the questions about autonomy utility, contribution and trust. We see a general trend of higher scoring with increasing autonomy, for both subject groups and interfaces, which again inconclusively supports hypothesis H4.

A low opinion of C_{filter} is observed in both subjects groups, regardless of interface and across all questions. One notable difference between subjects groups is seen in their opinion of the paradigms with stronger autonomy—with SCI subjects rating these paradigms lower in regards to utility and contribution than uninjured subjects. Interestingly, SCI subjects prefer one of these paradigms (C_{switch}) just as often as uninjured subjects do (Fig. 14), in spite of finding this paradigm to be less capable in comparison to uninjured subjects (and similarly capable to other paradigms which they prefer less). Further interesting, these differences in opinion between subject groups are present only when operating a joystick, and disappear when operating a head array. A possible explanation could be interface familiarity, since the SCI subjects all are regular operators of a joystick interface.

Between Sessions 1 and 2, we find SCI subjects to change their selection of most preferred paradigm 58% of the time, and for both interfaces. Interestingly, for uninjured subjects, preference changes only 17% of the time when operating with the joystick but,

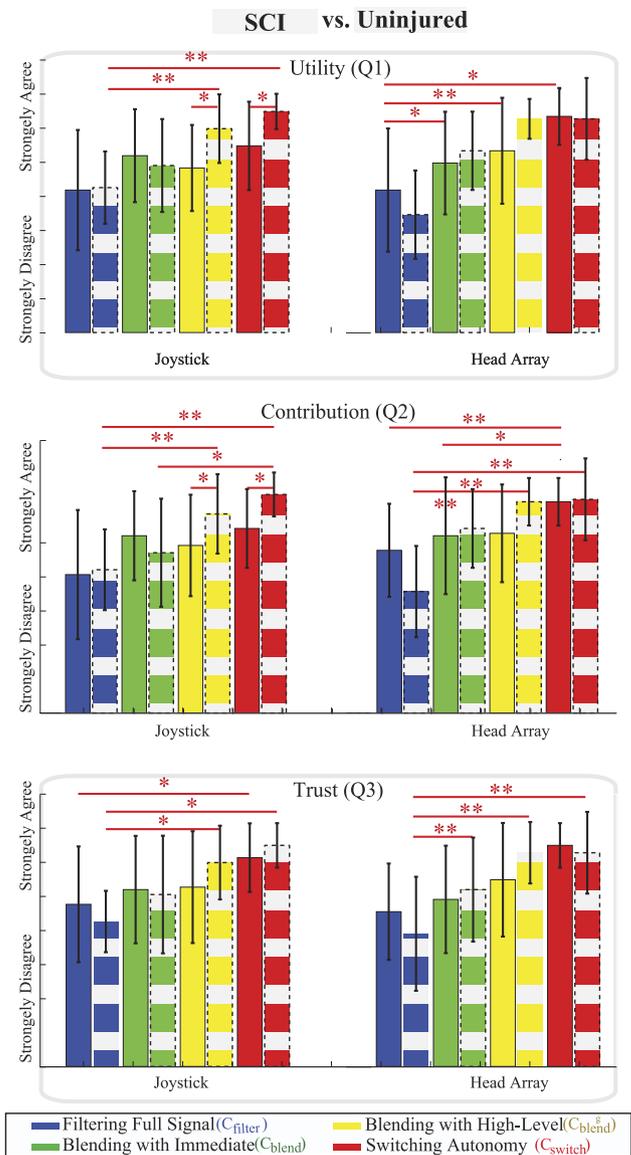


Fig. 15. Subjective evaluation of all navigation assistance paradigms on a 1-7 Likert Scale. SCI subjects (solid) vs. and uninjured subjects (striped).

like SCI subjects, changes 58% of the time with the head array—which is a more difficult interface. Thus, for a nontrivial number of participants, preference over control-sharing paradigms does appear to change over time.

We also asked the participants to indicate why they choose a particular paradigm as their most/least preferred. The most common answer for most preferred was a confidence in the robot’s ability to help them through the task. The most common reason for not preferring a navigation assistance paradigm was a lack of smoothness and/or hesitation. Two SCI subjects additionally reported preferring assistance paradigms that helped them to avoid walls they were “intentionally aiming” for to test the autonomy.

A comparison between the preferred paradigm answers with the questionnaire answers reveals no significant correlation between paradigm preference and the utility, contribution and trust subjective results.

6.5. Interventions and collisions

The performance metric that matters most for safety is the number of collisions and interventions either from participant or

experimenter. For trials involving any amount of autonomy, there were no head-on collisions and the experimenter did not need to intervene to prevent collisions. This was not the case for teleoperation which did see emergency stops by the experimenter to prevent collisions (1 joystick, 31 head array) and head-on collisions that were not prevented (3 joystick, 19 head array).

From a total of 896 trials with autonomy using the joystick (14 subjects driving a path that consisted of four doorway traversals, with four different navigation assistance paradigms, twice in a single session) the total number of side swipes was 51. For the head array trials, this number was 87. We do not observe consistent difference in the number of swipes across assistance paradigms. Considering that during 224 teleoperation trials, a total of 29 side swipes occurred with the joystick and 88 with the head array, autonomy effectively decreased the number of collisions.

There were zero instances of participants manually overriding the autonomy.

7. Conclusion

We have presented a control architecture for a robotic “smart” wheelchair that allows for the seamless interchange and evaluation of control-sharing paradigms. Its software and hardware components are modular and customizable—allowing for various control formulations and sensors to be swapped in or out. The interchangeable design facilitates comparative studies between different control sharing paradigms. The exact manner in which control is shared between a human and robotics autonomy is a crucial factor for assistive robots, and so the ability to easily assess and compare different control sharing paradigms is important.

We also have presented an implementation of four control sharing paradigms, and experimental results from a two-session study that compares all four to each other and teleoperation, and moreover using multiple control interfaces. We have shown that autonomy does increase safety, as measured by fewer collisions and higher distance to obstacles. Analysis of the quantitative results showed that performance metrics in general improve with increasing contribution from the autonomy, however, pairwise comparisons between three most dominant control paradigms were not significant. There also were marked differences between the amount of effort put forth by the human for different control paradigms, measured by a decreasing number of interactions with the interface and an increase in control signal frequency. We furthermore observed statistically significant differences in performance depending on the control interface, however, we did not observe strong evidence towards subjects benefiting from autonomy with more limited control interface.

Comparing the results of SCI and uninjured subject groups, we observed significant differences in regards to command fluency, but otherwise found few differences in performance metrics. Some differences were observed in regards to paradigm preference and opinion, but otherwise the distributions of most and least preferred paradigms, and opinions of the paradigms’ capabilities, were strikingly similar between the subject groups. While a general trend of preferring paradigms with stronger autonomy was observed, each paradigm was selected as most and least preferred at least once.

These results suggest that there is high variability in the most effective and accepted assistance paradigms over subjects. They further suggest that it would be desirable to provide options in control sharing to the end-users of robotic wheelchairs.

Our future work will extend this evaluation to more complex tasks, that will require intent inference, and a larger end-user population. Moreover, a longitudinal study would allow us to further investigate the adaptation of users to the autonomy and its effects on user’s performance and preference. We additionally will explore the customization of control-sharing paradigms by the end-users themselves.

Acknowledgment

This work was supported by grant from U.S. Office of Naval Research under the Award Number N00014-16-1-2247, which we gratefully acknowledge. The authors would also like to thank Jessica Presperin Pedersen OTR/L, ATP/SMS for helping with patient recruitment and experimental supervision.

Appendix A. Statistical significance analysis

This section outlines the details of the statistical significance analysis performed for all performance metrics. In particular, performed analysis can be summarized as (i) one factor repeated measure ANOVA (Analysis of Variance) for each session, (ii) if ANOVA is significant ($p < 0.05$), post-hoc pairwise comparisons using Bonferroni Confidence interval adjustments for each assistance paradigm, (iii) pairwise comparisons using Bonferroni Confidence interval adjustments between sessions, (iv) pairwise comparisons using Bonferroni Confidence interval adjustments between control interfaces, and (v) pairwise comparisons using Bonferroni Confidence interval adjustments between subject groups.

A.1. Time of completion

Fig. 8 demonstrates the Time of Completion metric over control paradigms, sessions, control interfaces and subject groups. Joystick trials show that within navigation paradigms during Session 1 there is a statistically significant difference, as observed in the ANOVA results [$F(4, 52) = 7.726, p < 0.01$]. Pairwise comparisons show an increase in time of completion with C_{filter} compared to C_{blend} ($p < 0.05$) and C_{switch} ($p < 0.01$).

Performance in Session 2 follows a similar pattern: within navigation paradigms there exists a statistically significant difference [$F(2.535, 32.957) = 11.981, p < 0.01$]. Pairwise comparisons show an increase in time of completion with C_{filter} compared to C_{tele} , C_{blend}^g , and C_{switch} ($p < 0.01$).

Interestingly, completion times largely do not change—for better, or for worse—in Session 2 in comparison to Session 1. The exception however is teleoperation with joystick ($p < 0.05$), indicating that participants do perform the task significantly faster on their own in Session 2.

Operating the wheelchair with the headrest interface also presents a statistically significant difference in Task Completion Time within navigation paradigms during Session 1 [$F(4, 52) = 8.102, p < 0.01$] and Session 2 [$F(4, 52) = 12.097, p < 0.01$]. Pairwise post-hoc analysis demonstrates the statistical significant increase in Time of Completion metric (faster task completion) in C_{switch} compared to C_{tele} ($p < 0.05$) and C_{filter} ($p < 0.01$) during Session 1. Statistical differences between teleoperation and switching autonomy disappears during Session 2. C_{filter} performance however is still significantly worse than other autonomous assistance paradigms— C_{blend} ($p < 0.05$), C_{blend}^g ($p < 0.01$) and C_{switch} ($p < 0.01$).

Between the control interfaces, pairwise comparisons reveal statistically significant differences during teleoperation trials and C_{filter} (Sessions 1 and 2) and also for C_{blend} and C_{switch} in Session 2 ($p < 0.05$).⁴ Overall, performance changes with the introduction of autonomy are more visible with the head array trials, especially compared to the poor performance during teleoperation.

We additionally report in Fig. 8 (bottom plot) Task Completion Times for each control paradigm and subject group, averaged over sessions, trials and doors. Comparisons between SCI and uninjured

⁴ Pairwise comparisons between interfaces are not displayed visually in Fig. 8–13 to minimize visual cluster.

subjects reveal limited statistically significant differences in completion time with the four assistance paradigms (C_{filter} in Session 1 and C_{blend} in Session 2 ($p < 0.05$)). Interestingly, when teleoperating with the joystick SCI subjects have significantly faster completion times than uninjured subjects—presumably because of prior experience operating the interface. This difference however does not present with the head array.

A.2. Number of interface interactions

Fig. 9 demonstrates the Number of Interface Interactions metric over control paradigms, sessions, control interfaces and subject groups. For the joystick trials, there exists statistical significance between navigation paradigms during Session 1 [$F(4, 52) = 3.498, p < 0.05$] and Session 2 [$F(4, 52) = 3.498, p < 0.05$]. There is an observable (and only significant between C_{filter} and C_{blend}^g ($p < 0.05$)) trend of a slight decrease in the Number of Interface Interactions with increasing autonomy, in both Sessions. Users providing fewer inputs with higher autonomy is not surprising, as the assistance paradigms are offloading more effort from the user.

For the headrest array trials, we do also observe a statistically significant difference between navigation paradigms during Session 1 [$F(4, 52) = 9.229, p < 0.01$], and during Session 2 [$F(2.207, 28.693) = 11.501, p < 0.01$]. Pairwise comparisons within paradigms show a statistically significant difference between teleoperation and the autonomous assistance paradigms C_{blend}^g and C_{switch} ($p < 0.01$) in Session 1, and, C_{blend} , C_{blend}^g , and C_{switch} in Session 2 ($p < 0.01$). Moreover, we do observe a statistical significant difference between C_{filter} and the two paradigms C_{blend}^g and C_{switch} ($p < 0.05$) in Session 2.

Between the control interfaces, pairwise comparisons for each navigation assistance paradigm reveal statistically significant difference in C_{tele} ($p < 0.01$) during Session 2 with joystick requiring fewer interactions, as expected.

We additionally report in Fig. 9 (bottom plot) Number of Interface Interactions for each control paradigm and subject group, averaged over sessions, trials and doors. Statistically significant difference between subject groups is observed in C_{blend} with the head array ($p < 0.01$). Specifically, uninjured volunteers provided fewer interactions compared to the SCI subjects.

A.3. Fluency of user commands

Fig. 10 demonstrates the Fluency of User Commands metric for control paradigms, sessions, control interfaces and subject groups. There is statistical significance between the joystick driven navigation paradigms during Session 1 [$F(2.199, 28.583) = 8.425, p < 0.01$] and Session 2 [$F(4, 52) = 10.598, p < 0.01$]. Post-hoc pairwise comparisons show that there is a significant difference in fluency of user commands between the two extremums— C_{tele} and C_{switch} ($p < 0.01$). During Session 2, users provide significantly smoother signals (compared to the teleoperation runs) in all assistance paradigms except the filtering approach ($C_{blend} p < 0.01, C_{blend}^g p < 0.01$ and $C_{switch} p < 0.01$). In particular, the Fluency of User Commands increases with autonomy which indicates that users provide smoother inputs.

For the headrest array trials, we also observe statistical significance between navigation paradigms in Session 1 [$F(4, 52) = 10.556, p < 0.01$] and Session 2 [$F(4, 52) = 11.379, p < 0.01$]. Post-hoc pairwise comparisons do show similar results to joystick trials, with teleoperation being less fluent than C_{filter} ($p < 0.01$), C_{blend}^g ($p < 0.01$) and C_{switch} ($p < 0.01$) during Session 1 and Session 2. Moreover, fluency increases on day two for C_{blend} and C_{blend}^g with the joystick ($p < 0.05$) and C_{switch} with the head array ($p < 0.05$).

Pairwise comparisons between the control interfaces reveal statistically significant differences ($p < 0.01$) with all assistance

paradigms. In all cases, use of the head array is more fluent than the joystick.

We additionally report in Fig. 10 (bottom plot) the Fluency of User Commands metric for each control paradigm and subject group, averaged over sessions, trials and doors. Comparisons between SCI and uninjured subjects reveal the same statistically significant differences in fluency (pairwise comparisons). Moreover there is an observable (and sometimes significant) trend in which the SCI fluency is less than that of the uninjured subjects, for all control paradigms. This trend is present for both interfaces.

A.4. Distance to obstacles

Fig. 11 demonstrates the Distance to Obstacles metric for control paradigms, sessions, control interfaces and subject groups. There is statistical significance between the joystick driven navigation paradigms during Session 1 [$F(4, 52) = 3.826, p < 0.01$], but not in Session 2 [$F(4, 52) = 0.665, p > 0.1$]. During Session 1, distance to obstacles is statistically significantly between the first two paradigms (C_{tele} and C_{filter}) and the rest of the autonomy assistance (C_{blend} , C_{blend}^g and C_{switch}) ($p < 0.05$)—with the exception of C_{tele} and C_{blend} combination.

For headrest trials, there exists no statistically significant difference within navigation assistance paradigms in Session 1 [$F(4, 52) = 0.515, p = 1.181$] but Session 2 control paradigms are significantly different [$F(2.319, 30.145) = 4.823, p < 0.01$]. Post-hoc pairwise comparison shows that trials during C_{tele} is significantly closer to obstacles compared to the three paradigms C_{blend} ($p < 0.05$), C_{blend}^g ($p < 0.01$) and C_{switch} ($p < 0.05$). There also exists a statistically significant difference between C_{filter} and C_{blend}^g ($p < 0.05$) during Session 2.

Pairwise comparisons between the interfaces for each control paradigm reveal statistically significant differences for C_{filter} in Session 1 ($p < 0.05$), and C_{tele} in Session 2 ($p < 0.05$).

We additionally report in Fig. 11 (bottom plot) the Distance to Obstacles metric for each control paradigm and subject group, averaged over sessions, trials and doors. Comparisons between SCI and uninjured subjects reveal no statistical significant difference (pairwise comparisons).

A.5. Similarity of commands

Fig. 12 demonstrates the Similarity of Commands metric for control paradigms, sessions, control interfaces and subject groups. We do not see any statistical differences through the ANOVA results analysis of joystick use with all navigation assistance paradigms in Session 1 [$F(3, 29) = 0.711, p > 0.1$] but in Session 2 [$F(3, 29) = 4.286, p = 0.01$]. Post-hoc pairwise comparisons of the navigation assistance paradigms show that only C_{switch} is significantly different than C_{filter} ($p < 0.01$). There is, although not significant, a noticeable trend of increasing similarity from Session 1 to 2.

For headrest trials, there exists a statistically significant difference within navigation assistance paradigms in Session 1 [$F(3, 29) = 3.295, p < 0.05$], and in Session 2 [$F(3, 29) = 4.238, p < 0.01$]. Post-hoc pairwise comparisons of the navigation assistance paradigms show that only C_{switch} is significantly different from C_{blend}^g ($p < 0.01$) during Session 2.

Pairwise comparisons between the interfaces for each control paradigm reveal statistically significant differences for all paradigms in Session 1 ($p < 0.01$), and C_{filter} and C_{switch} in Session 2 ($p < 0.05$). In all cases the joystick trials produced higher similarity.

We additionally report in Fig. 12 (bottom plot) Similarity of Commands for each control paradigm and subject group, averaged over sessions, trials and doors. Statistically significant difference between subject groups is only observed in C_{blend}^g and C_{switch} with the head array ($p < 0.05$). In both cases, SCI commands are less similar to autonomous commands compared to uninjured subjects.

A.6. Contribution from autonomy

Fig. 13 demonstrates the Contribution from Autonomy metric for control paradigms, sessions, control interfaces and subject groups. For the joystick trials, we do see statistically significant changes of the contribution of the autonomy to the executed control signal \mathbf{u} within the navigation assistance paradigms for Sessions 1 [$F(3, 39) = 15.297, p < 0.01$] and 2 [$F(3, 39) = 36.598, p < 0.01$]. Unsurprisingly, it contributes the greatest amount during the blending with high-level goals and autonomy switching runs. Specifically, C_{filter} is significantly lower than all other navigation assistance paradigms, in both Sessions ($p < 0.01$). Moreover, C_{blend}^g contributes significantly more compared to C_{blend} ($p < 0.01$) during Session 2.

For the head array trials, there also is a statistically significant difference within the navigation assistance paradigms during Session 1 [$F(3, 29) = 31.791, p < 0.01$] and Session 2 [$F(3, 39) = 29.430, p < 0.01$]. Post-hoc pairwise comparisons in between navigation assistance paradigms show similar results to the joystick: C_{filter} contributes significantly less than all other navigation assistance paradigms, in both sessions. Moreover, C_{blend}^g contributes significantly more compared to C_{blend} ($p < 0.01$). With the head array we additionally see some instances of a significantly greater contribution from the autonomy on day one when C_{blend} is compared with C_{blend}^g ($p < 0.01$) and C_{switch} ($p < 0.05$).

Between the control interfaces, pairwise comparisons reveal that C_{filter} and C_{blend} contribute significantly less with head array than they are doing with joystick, during Session 1.

We additionally report in Fig. 13 (bottom plot) Contribution from Autonomy for each control paradigm and subject group, averaged over sessions, trials and doors. Comparisons between SCI and uninjured subjects reveal the statistically significant difference between subject groups in C_{blend}^g with the head array ($p < 0.05$). Specifically for this case, uninjured subjects utilized autonomy more effectively compared to SCI subjects.

Appendix B. Supplementary data

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.robot.2017.04.013>.

References

- [1] L. Fehr, W.E. Langbein, S.B. Skaar, Adequacy of power wheelchair control interfaces for persons with severe disabilities: A clinical survey, *J. Rehabil. Res. Dev.* 37 (3) (2000) 353–360.
- [2] F. Leishman, O. Horn, G. Bourhis, Smart wheelchair control through a deictic approach, *Robot. Auton. Syst.* 58 (10) (2010) 1149–1158.
- [3] S.P. Parikh, V. Grassi, Integrating human input with autonomous behaviors on an intelligent wheelchair platform, *IEEE J. Intell. Syst.* 22 (2) (2007) 33–41.
- [4] R. Simpson, Smart wheelchairs: A literature review, *J. Rehabil. Res. Dev.* 42 (4) (2005) 423–438.
- [5] R. Simpson, E. LoPresti, R. Cooper, How many people would benefit from a smart wheelchair? *J. Rehabil. Res. Dev.* 45 (1) (2008) 53–72.
- [6] E.A. Biddiss, T.T. Chau, Upper limb prosthesis use and abandonment: A survey of the last 25 years, *Prosthet. Orthot. Int.* 31 (3) (2007) 236–257.
- [7] A. Lankenau, T. Röfer, A versatile and safe mobility assistant, *IEEE Robot. Autom. Mag.* 8 (1) (2001) 29–37.
- [8] C. Gao, M. Sands, J.R. Spletzer, Towards autonomous wheelchair systems in urban environments, in: *Field and Service Robotics*, in: Springer Tracts in Advanced Robotics, vol. 62, 2010.
- [9] C.T. Lin, C. Euler, P.J. Wang, A. Mekhtarian, Indoor and outdoor mobility for an intelligent autonomous wheelchair, in: *Proceedings of the International Conference on Computers Helping People with Special Needs - Volume Part II, ICCHP*, 2012.
- [10] J.-A. Fernández-Madriral, C. Galindo, J. González, Assistive navigation of a robotic wheelchair using a multihierarchical model of the environment, *Integr. Comput.-Aided Eng.* 11 (4) (2004) 309–322.
- [11] T. Gomi, A. Griffith, Developing intelligent wheelchairs for the handicapped, in: *Assistive Technology and Artificial Intelligence*, in: Lecture Notes in Computer Science, vol. 1458, Springer, Berlin, Heidelberg, 1998.
- [12] F. Pasteau, V.K. Narayanan, M. Babel, F. Chaumette, A visual servoing approach for autonomous corridor following and doorway passing in a wheelchair, *Robot. Auton. Syst.* 75 (2016) 28–40.
- [13] I. Iturrate, J. Antelis, A. Kubler, J. Minguez, A noninvasive brain-actuated wheelchair based on a P300 neurophysiological protocol and automated navigation, *IEEE Trans. Robot.* 25 (3) (2009) 614–627.
- [14] A. Escobedo, A. Spalanzani, C. Laugier, Multimodal control of a robotic wheelchair: Using contextual information for usability improvement, in: *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, 2013.
- [15] J.D. Crisman, M.E. Cleary, Progress on the deictically controlled wheelchair, in: *Assistive Technology and Artificial Intelligence, Applications in Robotics, User Interfaces and Natural Language Processing*, 1998.
- [16] A. Argyros, P. Georgiadis, P. Trahanias, D. Tsakiris, Semi-autonomous navigation of a robotic wheelchair, *J. Intell. Robot. Syst.* 34 (3) (2002) 315–329.
- [17] R.A.M. Braga, M. Petry, L.P. Reis, A.P. Moreira, IntellWheels: Modular development platform for intelligent wheelchairs, *J. Rehabil. Res. Dev.* 48 (9) (2011) 1061–1076.
- [18] A.K. Makhtar, H. Yussof, H. Al-Assadi, L.C. Yee, M. R.M. Tomari, Y. Kobayashi, Y. Kuno, Development of smart wheelchair system for a user with severe motor impairment, in: *Procedia Engineering International Symposium on Robotics and Intelligent Sensors, IRIS*, 2012, pp. 538–546.
- [19] M. Fernandez-Carmona, B. Fernandez-Espejo, J. Peula, C. Urdiales, F. Sandoval, Efficiency based collaborative control modulated by biometrics for wheelchair assisted navigation, in: *Proceedings of the IEEE International Conference on Rehabilitation Robotics, ICORR*, 2009.
- [20] Q. Li, W. Chen, J. Wang, Dynamic shared control for human-wheelchair cooperation, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, 2011.
- [21] C. Urdiales, J. Peula, M. Fernandez-Carmona, R. Annicchiarico, F. Sandoval, C. Caltagirone, Adaptive collaborative assistance for wheelchair driving via CBR learning, in: *Proceedings of the IEEE International Conference on Rehabilitation Robotics, ICORR*, 2009.
- [22] C. Urdiales, E.J. Pérez, G. Peinado, M. Fdez-Carmona, J.M. Peula, R. Annicchiarico, F. Sandoval, C. Caltagirone, On the construction of a skill-based wheelchair navigation profile, *Proc. IEEE Trans. Neural Syst. Rehabil. Eng.* 21 (6) (2013) 917–927.
- [23] M. Fdez-Carmona, C. Urdiales, F. Sandoval, Reactive adapted assistance for wheelchair navigation based on a standard skill profile, in: *Proceedings of the IEEE International Conference on Robotics and Biomimetics, ROBIO*, 2014.
- [24] D. Miller, M. Slack, Design and testing of a low-cost robotic wheelchair prototype, *Auton. Robot. Syst.* 2 (1) (1995) 77–88.
- [25] R.C. Simpson, S.P. Levine, D.A. Bell, L.A. Jaros, Y. Koren, J. Borenstein, NavChair: An assistive wheelchair navigation system with automatic adaptation, in: *Assistive Technology and Artificial Intelligence*, in: Lecture Notes in Computer Science, vol. 1458, Springer, Berlin, Heidelberg, 1998.
- [26] J. Philips, J. del R. Millan, G. Vanacker, E. Lew, F. Galan, P.W. Ferrez, H.V. Brussel, M. Nuttin, Adaptive shared control of a brain-actuated simulated wheelchair, in: *Proceedings of the IEEE International Conference on Rehabilitation Robotics, ICORR*, 2007.
- [27] A.V. Nguyen, L.B. Nguyen, S. Su, H.T. Nguyen, The advancement of an obstacle avoidance bayesian neural network for an intelligent wheelchair, in: *Proceedings of the IEEE International Conference on Engineering in Medicine and Biology Society, EMBC*, 2013.
- [28] T. Röfer, A. Lankenau, Architecture and applications of the Bremen autonomous wheelchair, *Inf. Sci. Inf. Comput. Sci.* 126 (1–4) (2000).
- [29] T. Carlson, Y. Demiris, Collaborative control for a robotic wheelchair: Evaluation of performance, attention, and workload, *IEEE Trans. Syst. Man Cybern. B* 42 (3) (2012) 876–888.
- [30] T. Carlson, Y. Demiris, Human-wheelchair collaboration through prediction of intention and adaptive assistance, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, 2008.
- [31] Y. Satoh, K. Sakaue, An omnidirectional stereo vision-based smart wheelchair, *EURASIP J. Image Video Process.* 2007 (1) (2007) 1–11.
- [32] P. Viswanathan, J.J. Little, A.K. Mackworth, A. Mihailidis, Navigation and Obstacle Avoidance Help (NOAH) for older adults with cognitive impairment: A pilot study, in: *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS*, 2011.
- [33] S.P. Parikh, V. Grassi, V. Kumar, J. Okamoto, Incorporating user inputs in motion planning for a smart wheelchair, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, 2004.
- [34] S.P. Parikh, V. Grassi, V. Kumar, J.J. Okamoto, Usability study of a control framework for an intelligent wheelchair, in: *Proceedings of the IEEE International Conference on Robotics and Automation, ICRA*, 2005.
- [35] R. Simpson, E. LoPresti, S. Hayashi, I. Nourbakhsh, D. Miller, The smart wheelchair component system, *J. Rehabil. Res. Dev.* 41 (3B) (2004) 428–442.
- [36] I.M. Mitchell, P. Viswanathan, B. Adhikari, E. Rothfels, A.K. Mackworth, Shared control policies for safe wheelchair navigation of elderly adults with cognitive and mobility impairments: Designing a wizard of oz study, in: *Proceedings of the American Control Conference, ACC*, 2014.

- [37] H. Soh, Y. Demiris, When and how to help: An iterative probabilistic model for learning assistance by demonstration, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2013.
- [38] A. Kucukyilmaz, Y. Demiris, One-shot assistance estimation from expert demonstrations for a shared control wheelchair system, in: Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN, 2015.
- [39] D. Vanhooydonck, E. Demeester, A. Hüntemann, J. Philips, G. Vanacker, H.V. Brussel, M. Nuttin, Adaptable navigational assistance for intelligent wheelchairs by means of an implicit personalized user model, *Robot. Auton. Syst.* 58 (8) (2010) 963–977.
- [40] J.S. Nguyen, S.W. Su, H.T. Nguyen, Experimental study on a smart wheelchair system using a combination of stereoscopic and spherical vision, in: Proceedings of the IEEE International Conference on Engineering in Medicine and Biology Society, EMBC, 2013.
- [41] J.S. Ju, Y. Shin, E.Y. Kim, Vision based interface system for hands free control of an intelligent wheelchair, *J. NeuroEng. Rehabil.* 6 (1) (2009) 1–17.
- [42] G.U. Sorrento, P.S. Archambault, F. Routhier, D. Dessureault, P. Boissy, Assessment of Joystick control during the performance of powered wheelchair driving tasks, *J. NeuroEn. Rehabil.* 8 (1) (2011) 1–11.
- [43] B.D. Argall, Modular and adaptive wheelchair automation, in: *Experimental Robotics*, in: Springer Tracts in Advanced Robotics, vol. 109, Springer International Publishing, 2016.
- [44] M. Derry, B. Argall, Automated doorway detection for assistive shared-control wheelchairs, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA, 2013.
- [45] S. Jain, B.D. Argall, Automated perception of safe and oriented docking locations with alignment information for assistive wheelchairs, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2014.
- [46] M. Nejati, B. Argall, Automated incline and drop-off detection for assistive powered wheelchairs, in: Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN, 2016.
- [47] M. Quigley, K. Conley, B.P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A.Y. Ng, ROS: an open-source robot operating system, in: ICRA Workshop on Open Source Software, 2009.
- [48] E. Demeester, A. Hüntemann, D. Vanhooydonck, G. Vanacker, H. Van Brussel, M. Nuttin, User-adapted plan recognition and user-adapted shared control: A Bayesian approach to semi-autonomous wheelchair driving, *Auton. Robots* 24 (2) (2008) 193–211.
- [49] E. Prassler, J. Scholz, P. Fiorini, A robotics wheelchair for crowded public environment, *IEEE Robot. Autom. Mag.* 8 (1) (2001) 38–45.
- [50] P. Boucher, A. Atrash, S. Kelouwani, W. Honoré, H. Nguyen, J. Villemure, F. Routhier, P. Cohen, L. Demers, R. Forget, J. Pineau, Design and validation of an intelligent wheelchair towards a clinically-functional outcome, *Auton. Robots* 10 (58) (2013).
- [51] D. Sinyukov, R. Desmond, M. Dickerman, J. Fleming, J. Schaufeld, T. Padir, Multi-modal control framework for a semi-autonomous wheelchair using modular sensor designs, *Intell. Serv. Robot.* 7 (3) (2014) 145–155.
- [52] A. Gail, M. Derry, B.D. Argall, Using machine learning to blend human and robot controls for assisted wheelchair navigation, in: Proceedings of IEEE International Conference on Rehabilitation Robotics, ICORR, 2013, pp. 1–6.
- [53] L. Chen, S. Wang, H. Hu, K. McDonald-Maier, Bezier curve based trajectory planning for an intelligent wheelchair to pass a doorway, in: Proceedings of the UKACC International Conference on Control, 2012.
- [54] L. Montesano, J. Minguez, L. Montano, Modeling dynamic scenarios for local sensor-based motion planning, *Auton. Robots* 25 (3) (2008) 231–251.
- [55] E. Demeester, M. Nuttin, D. Vanhooydonck, H.V. Brussel, A model-based, probabilistic framework for plan recognition in shared wheelchair control: experiments and evaluation, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2003.
- [56] P. TalebiFard, J. Sattar, I.M. Mitchell, A risk assessment infrastructure for powered wheelchair motion commands without full sensor coverage, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2014, pp. 3592–3597.
- [57] S. Gulati, B. Kuipers, High performance control for graceful motion of an intelligent wheelchair, in: Proceedings of IEEE International Conference on Robotics and Automation, ICRA, 2008, pp. 3932–3938.
- [58] L.R. Kirby, L.A. Worobey, R. Cowan, J.P. Pedersen, A.W. Heinemann, T.A. Dyson-Hudson, M. Shea, C. Smith, P.W. Rushton, M.L. Boninger, Wheelchair skills capacity and performance of manual wheelchair users with spinal cord injury, *Arch. Phys. Med. Rehabil.* 2016 (2016) 30242–30248.
- [59] A. Broad, B. Argall, Path planning under interface-based constraints for assistive robotics, in: Proceedings of the International Conference on Automated Planning and Scheduling, ICAPS, 2016.



Ahmetcan Erdogan is a Post-Doctoral Fellow in the Department of Physical Medicine & Rehabilitation at the Northwestern University. He has received his B.Sc. degree in Electrical Engineering from Istanbul Technical University, Turkey in 2006 and the M.Sc. degree in Electrical Engineering specialized in control of dynamical systems from Högskolan Dalarna, Sweden, in 2008. He received his Ph.D. degree in Mechatronics Engineering from Sabanci University, Turkey in 2014. His research interests include the optimal design of physical assistive devices, the implementation of assist-as-needed control methodologies, and human motor control system analysis and haptics.



Brenna Argall is the June and Donald Brewer Junior Professor of Electrical Engineering & Computer Science at Northwestern University, and also an assistant professor in the Departments of Mechanical Engineering and Physical Medicine & Rehabilitation. Her research lies at the intersection of robotics, machine learning and human rehabilitation. She is director of the assistive & rehabilitation robotics laboratory (**argallab**) at the Rehabilitation Institute of Chicago (RIC), the premier rehabilitation hospital in the United States. The mission of the **argallab** is to advance human ability by leveraging robotics autonomy. Argall is a 2016 recipient of the NSF CAREER award, and her Ph.D. in Robotics (2009) was received from the Robotics Institute at Carnegie Mellon University, as well as her M.S. in Robotics (2006) and B.S. in Mathematics (2002). Prior to joining Northwestern and RIC, she was a postdoctoral fellow (2009–2011) at the École Polytechnique Fédérale de Lausanne (EPFL), and prior to graduate school she held a Computational Biology position at the National Institutes of Health (NIH).