

Increasing Robotic Wheelchair Safety With Collaborative Control: Evidence from Secondary Task Experiments

Tom Carlson and Yiannis Demiris

Abstract—Powered wheelchairs play a vital role in bringing independence to the severely mobility-impaired. Our robotic wheelchair aims to assist users in driving safely, without undermining their capabilities or curtailing the natural development of their skills. An important research question is to determine the conditions under which shared control is most beneficial. In this paper, we describe an experiment, where a distracting secondary task caused the majority of participants to crash the wheelchair when driving without assistance. However, when they were assisted by our collaborative controller, not only did they drive safely, but they also increased their performance in the secondary task. We demonstrate that a degree of shared control is beneficial even to proficient drivers under certain circumstances, for instance when they are under a heightened workload.

I. INTRODUCTION

Smart wheelchairs are becoming a popular research platform for adaptive systems and human robot interaction. Ultimately they aim to help people who are suffering from mobility impairments (and often compound disabilities) achieve a level of independence, so that they can get on with their activities of daily living (ADLs). Many research groups have taken different approaches to tailoring the control system to the user. These range from those that offer some low level collision avoidance, to Taha *et al.* [1] who used a high level of autonomy that required relatively little user interaction or Millán *et al.* [2] who used a brain machine interface offering a very low user input resolution. There are also many hybrid systems that switch (sometimes autonomously) between different modes of operation, such as the NavChair [3]. For a more comprehensive review of smart wheelchairs, refer to Ding and Cooper [4].

We follow Nisbet's recommendations [5], whereby we keep the control user-initiated and only adapt signals where necessary, in order to perform particularly precise manoeuvres, or to avoid collisions. We have proposed an effective collaborative control methodology, which infers the user's intentions from their joystick input, along with the wheelchair's knowledge of the local environment [6]. Based on these predictions, the wheelchair alters the motor control signals to assist the user, as and when it is necessary.

When evaluating our wheelchair, we place an emphasis on measuring the performance of the user, in accordance with the recommendations of Tsui *et al.* for evaluating assistive robotic technologies [7]. A pilot study of visual



Fig. 1. Global localisation data from the camera is used alongside dynamic sensory data from the laser scanner and sonars to assist the user in performing precise manoeuvres. In this experiment, the user controls the wheelchair using the joystick in their right hand, whilst performing a secondary task on the joypad buttons with their left hand.

attention whilst driving our wheelchair yielded some surprising results; an increase in saccadic eye movements, whilst the collaborative controller was active [8]. However, it was hypothesised that the increased visual activity might be caused by the user forming an incorrect mental model of the system's behaviour. It is generally accepted that shared control is a good approach in human robot interaction, however the conditions under which it brings maximum benefit to the user are still unclear. We hypothesise that in shared tasks, the user's workload is an important factor, in addition to their capabilities and needs. We further explore the ideas of workload and participants' perceptions by using a secondary task.

Wheelchair users rarely travel without interacting with other people or the surroundings. For example, Brandt *et al.* found that 87% of the 111 people surveyed used their wheelchairs to go shopping [9]; an ADL that certainly requires divided attention. Therefore, we decided to use a secondary task to evaluate how people drive under increased workload. Parikh *et al.* [10] have used secondary tasks to compare autonomous, semi-autonomous and manual operation of a wheelchair, in terms of cognitive complexity and NASA's Task Load Index (TLX). Their approach to semi-autonomous operations differs considerably from ours; they consider the wheelchair to be autonomously following a deliberative plan, with the user intervening as and when they wish to deviate from the plan. Conversely, we consider the user to be initiating every movement of the wheelchair and

T. Carlson and Y. Demiris are with Imperial College London, SW7 2AZ, UK. tom.carlson02@imperial.ac.uk, y.demiris@imperial.ac.uk

the collaborative controller only steps in when a difficult situation is encountered (such as passing through a narrow doorway). Additionally, we did not face the problems of side-on collisions during semi-autonomous operation that were observed in [10], since our sonar sensors covered the area either side of the wheelchair. Despite the operational differences, the underlying methodology they used for evaluating their system is also well suited to our setup.

In this paper we show how our collaborative control system reduces the user's workload and improves safety when manoeuvring a powered wheelchair. Throughout these trials, we observe the participants' secondary task performance, whilst making typical manoeuvres, such as driving around cluttered offices, along corridors and passing through narrow doorways. The observations are made over one independent variable, which can take one of two states: provide adaptive assistance, or provide no assistance.

II. THE SYSTEM ARCHITECTURE

We combine our collaborative control system with a dynamic local obstacle avoidance module (DLOA) in a hierarchical manner. In essence, the user indicates their intentions via the joystick. These signals may be altered, if required, by the collaborative controller to guarantee a suitable approach trajectory to a target (e.g. doorway). The resulting motor commands are then passed to the DLOA and may be changed again, based upon the laser scanner readings. Finally, a virtual bumper is implemented, using the sonar readings; if the output motor signals from the DLOA would not cause a breach of the virtual bumper, they are fed to the wheelchair's motor control unit and the movement is executed.

A. Collaborative Control

The collaborative controller [6] is comprised of a shared controller module, which takes input from the *safe mini-trajectory* generator, along with the intention predictor, to decide exactly how to adapt the joystick signals, as shown in Fig. 2. In these experiments, the global localisation was derived from a computer vision based system that was developed to work in mapped, indoor environments (through the use of paper fiducials).

We based our intention prediction on the multiple hypothesis approach, as described in [11]. These prediction models are task based, so we defined targets of interest, such as doorways and desks, which the user may wish to drive through or approach. We also constructed a confidence function, which is monotonically increasing when moving towards a target and is based upon the Euclidean distance and the angle between the heading of the chair and the target location. For a more detailed analysis of the collaborative control architecture, please refer to [6].

B. Dynamic Local Obstacle Avoidance (DLOA)

There has been much work in the field of mobile robotics and autonomous obstacle avoidance, perhaps most famously the vector field histogram (VFH) [12], which was later

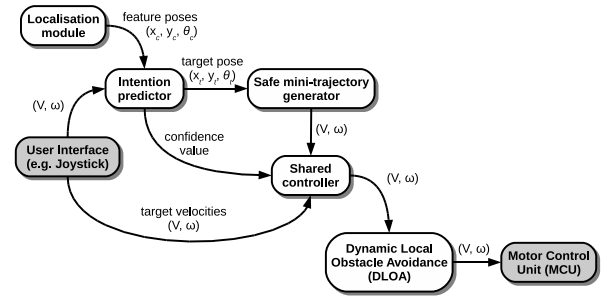


Fig. 2. The collaborative control architecture assists the user to manoeuvre the wheelchair precisely [6], whilst the new dynamic local obstacle avoidance (DLOA) module provides an additional safety mechanism. In the diagram (x_c, y_c, θ_c) and (x_t, y_t, θ_t) describe the wheelchair's current and target poses respectively. (V, ω) represent the target translational and rotational velocity tuple to be sent to the motor control unit; these are initially set according to the user input, but can be adjusted sequentially by the shared controller and the DLOA module.

adapted to be used in the context of a powered wheelchair by Levine *et al.* [3]. However, even this extensively modified version was reported to require a minimum of 18cm of clearance to pass through gaps 70% of the time, which was not flexible enough for performing our tasks. Therefore, we decided to take a fresh user-centric approach to the problem.

We based our implementation on the forward models¹ that underpin our intention prediction mechanism in the collaborative controller. A simplified mathematical model of the wheelchair's behaviour (our forward model), allows us to take the user input and estimate the wheelchair pose in the next 100ms time-step, relative to its current pose and velocity. We then defined the wheelchair's safety zone to be the boundary of the area the wheelchair would traverse in the 100ms time-step, plus a velocity-dependent error margin. This was represented in polar form, as a vector of distances (\mathbf{Z}_w) from the centre of the wheelchair, with the index (i) of each element representing the angle (θ) from the heading of the wheelchair, such that:

$$i = \left\lfloor \frac{N_L}{2} + \frac{\theta}{\delta} \right\rfloor, i \in \mathbb{Z} \quad (1)$$

where N_L is the length of the vector and δ is the angular resolution of our laser scanner.

Next, we evaluate whether or not there were any intersections with the laser range data (which was also presented as a vector of distances \mathbf{L}). An intersection would represent a collision, so we must search for the direction to travel that would not result in an intersection and is closest to the user's intended direction. To do this, we constructed Algorithm 1, which shifts \mathbf{Z}_w —yielding a rotation in Cartesian space—until it finds a suitable direction, or determines there is no safe direction. This process is illustrated in Fig. 3.

Finally, the new motor control signals are generated. If the safe direction that is computed by the DLOA is significantly different to the output from the collaborative control system, the translational velocity is reduced proportionally to this

¹A forward model estimates the next state of the system, given the current state and current inputs [11].

difference and the rotational velocity is set to achieve the newly desired direction (by using the same inverse models² that are built into the collaborative controller).

Algorithm 1 DLOA evaluates multiple forward models until it finds the direction that is both safe and closest to the user's intended direction of travel, should one exist.

Require: Z_w Wheelchair safety zone
Require: L Laser range data
Require: N_L Number of laser readings
Require: δ Angular resolution of laser readings
Require: ξ Joystick angle
Require: K_ϵ Maximum angular adjustment (we used $\frac{\pi}{4}$)

$\phi_0 := \lfloor \frac{\xi}{\delta} \rfloor$
 $\epsilon := 0$
repeat
 $\Delta := \phi_0 + \lfloor \frac{\epsilon}{2} \rfloor$
 $\text{safe} := \text{true}$
 $i := 0$
while $i < N_L$ **do**
 $j := i + \Delta$
if $j \geq 0$ **and** $j < N_L$ **then**
if $Z_w(j) \geq L(i)$ **then**
 $\text{safe} := \text{false}$
break
end if
end if
 $i := i + 1$
end while
 $\epsilon := -\text{sign}(2\epsilon + 1)(|\epsilon| + 1)$
until $\text{safe} = \text{true}$ **or** $|\epsilon| > K_\epsilon$
return $(\text{safe}, \delta\epsilon)$

III. METHODOLOGY

In this section we explain how the experiments were performed. We recruited 16 able-bodied volunteers aged between 20 and 56. Each subject took about 35 minutes to complete the experiment and fill in a brief questionnaire.

A. Primary Task

The primary task was simply to drive the wheelchair twice around the circuit shown in Fig. 4(a), without having any collisions. Each lap involved some navigation within a cluttered office environment and travelling along a short stretch of corridor, which resulted in passing through three doors of varying widths. The narrowest door was between the corridor and Room 1, being only 10cm wider than the wheelchair.

B. Secondary Task

The secondary task was chosen to be deliberately distracting and to require a certain degree of visual attention (although some studies have reported findings for memory-based secondary tasks are statistically similar to those that require visual attention [13]). This allowed us to determine how users might drive under increased workload.

A single random quadrant of the dark blue tablet PC screen was highlighted in white at random time intervals (bounded

²An inverse model estimates the control signals that are required to move a system from its current state into a desired state [11].

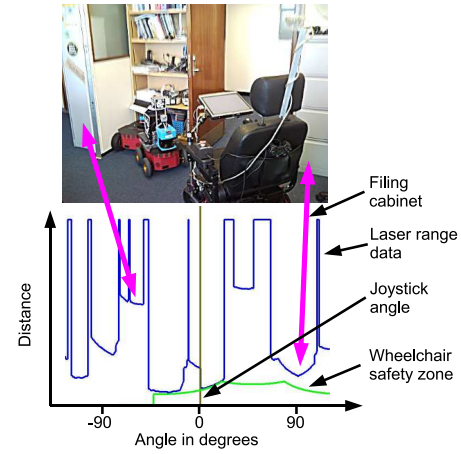


Fig. 3. As the wheelchair faces the gap between the mobile robot and the door, the joystick is set in the straight forward position. However, if the wheelchair's safety zone were centred on the joystick angle, it would intersect with the laser scan. Therefore, the dynamic local obstacle avoidance (DLOA) module shifts it approximately 45 degrees to the right, so that the wheelchair would head towards the open doorway

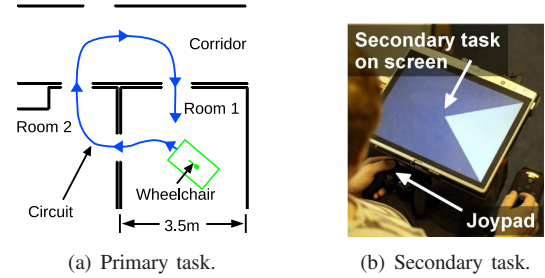


Fig. 4. The experiment: drive the route of the primary task, whilst reacting to random quadrants of the screen lighting up in the secondary task.

between 100ms and 1000ms; see Fig. 4(b)). The user had to react as quickly as possible by pressing the corresponding button on the joypad controller: i.e. the right quadrant of the screen corresponds to the east button on the joypad; the top screen quadrant corresponds to the north button etc.. If the correct button was pressed, the reaction time would be logged, the highlighted quadrant would turn momentarily green, before reverting to dark blue and the whole process would begin again. Conversely, if the incorrect button had been pressed, the quadrant of the screen that corresponded to the incorrect button would momentarily turn red and the secondary task would remain in the same state until the correct button had been pressed.

During initial trials of the secondary task, it was discovered that if the same quadrant was highlighted three times or more in a row, the participant's reaction time for the third and each successive reaction would be significantly lower than their average reaction time. They were also more likely to press the correct button, which made the secondary task relatively easy. To avoid getting stuck in this 'local minimum', we ensured that the same quadrant would only be highlighted a maximum of two times in a row and this proved to make the task more challenging.

C. Participant Feedback

Participants were asked to fill in brief a questionnaire about their experience at the end of their experiment. It was predominantly a comparative questionnaire asking them to indicate how strongly they agreed with each of the statements in Fig. 8 for each control mode, on a five point Likert scale (1 = strongly agreed, 5 = strongly disagreed).

D. Experimental Procedure

The independent variable we were testing was the wheelchair control method, which could take one of two states: provide adaptive assistance, or provide no assistance. Due to the nature of adaptive control algorithms, we decided it would be most appropriate to perform a *within subjects* experiment. This allowed us to see if providing assistance actually helped each individual, rather than testing the performance of individuals against each other. To eliminate bias from the inevitable learning effect, odd numbered participants undertook a set of trials with adaptive assistance before moving on to a set of trials without any assistance. Conversely, even numbered participants undertook the trials without any assistance, before being introduced to the adaptive assistance mode of operation.

In each set of trials, the same well-defined procedure was followed. First, the participant was given five minutes to drive the wheelchair around the office environment and along the corridor, to familiarise themselves with the active control mode. Next, whilst they were stationary, they were introduced to the secondary task (the participants were told this was a *reaction game*). They were then given a practice trial, whereby the participant was instructed to drive twice around the circuit shown in Fig. 4(a), whilst simultaneously playing the reaction game. It was reiterated that their main task was to drive safely and then to play the reaction game as best they could. Data was then recorded for the following two trials, in which the participant drove twice around the test circuit. They were then given a two minute break before undertaking the entire procedure again for the remaining wheelchair control method (either with adaptive assistance, or without assistance). The second set of trials were identical to the first, apart from the fact that the wheelchair control method was swapped and the stationary practice session of the secondary task was omitted.

For safety reasons—since we were expecting users might have minor collisions—we limited the maximum translational velocity of the wheelchair to 1 metre per second and the maximum angular velocity to 90 degrees per second.

IV. EXPERIMENTAL RESULTS

One participant had to be excluded from the trial, because they could not complete enough of the experiment to provide sufficient comparative data. They managed to undertake the secondary task and navigate safely whilst the collaborative controller was active, however, when no assistance was given, they were unable to complete a single circuit and voluntarily asked not to continue with the experiment. Although we could not use this data in our results, it demonstrates just




Fig. 5. The number of times the wheelchair collided with an obstacle.

how useful the collaborative control mechanism could be to some users.

Several collisions occurred when the collaborative controller was not assisting the user. Most of the collisions occurred on the approach to the third doorway, when driving from the corridor back into room one. This was usually because participants approached the doorway from a too shallow angle and ended up catching the driving wheels on the door-frame, which caused an emergency stop. In a few cases, the participant simply overshot the doorway and crashed into the wall with the wheelchair's footplate.

Only one participant did not crash when driving without any assistance, as can be seen in Fig. 5. On average, participants crashed 2.27 times per trial, whilst driving without any assistance. This resulted in a combined total of 34 crashes over the entire experiment. In contrast there were only two collisions when the collaborative controller was active and these could both be discounted as anomalies. Participant number seven dislodged the laser scanner with his foot, which directly resulted in a crash. Additionally, participant number ten had knocked a sonar sensor off the wheelchair in an earlier, non-assisted run; this also resulted in a collision once the collaborative controller had been activated.

Clearly, from the successful reduction in collisions, the adaptive assistance mode increases safety when driving a powered wheelchair. However, we were also interested in the impact the adaptive assistance had on the user's workload, therefore in the following subsection, we analyse the results pertaining to the secondary task, before looking at the collective feedback from participants.

A. Secondary Task

The most interesting result to come out of the secondary task was the percentage of incorrect reactions, rather than the reaction times themselves. It seems that participants were determined to react as quickly as possible under both control modes, however, when they were in complete control of the wheelchair (i.e. they were not given any assistance), they had a higher workload and were therefore more prone to make mistakes in the secondary task. Fig. 6(b) shows there was a significant improvement in the average number of incorrect reactions when the collaborative controller was assisting them, dropping from an average of 13.5% incorrect to 9.7%. The significance was verified using a paired one-tailed *t*-test with $p < 0.009$. There were 9 instances when




Fig. 6. An analysis of the secondary task in terms of percentage of incorrect reactions and average reaction times.

the percentage of incorrect reactions became worse when using collaborative control, although only participant number 6 exhibited a statistically significant increase. This may have occurred due to fatigue, since the participant was an even number, meaning the collaborative run was performed after the non-assisted run.

The mean reaction times of the participants differed negligibly between the trials where the collaborative controller was assisting them and those where they were given no help in driving, as shown in Fig. 6(c) (using the t -test, $p > 0.5$). There was slightly less variance in the results when they were given assistance, which means their reaction times seemed more predictable, however, again the t -test showed this result to be statistically insignificant, with $p > 0.3$. The mean reaction times were slightly lower whilst driving through a doorway, compared with the rest of the circuit (Fig. 6(c)). As can be seen from Fig. 7, the peak reaction times were not limited to when the chair was passing through the doorway, but often came on the immediate approach, or as the wheelchair left a doorway. This could have been because once the driver had successfully manoeuvred into a narrow doorway, in most cases, they only had to drive straight to pass through safely. Conversely, it took greater skill to align the wheelchair with the door opening in the first place, or to turn in a tight corridor when exiting a room. During these times of greater concentration, they may still have used their peripheral vision to react to the secondary task as changes occur on the screen, but due to their lack of attention, this is where the incorrect reactions appear to dominate.

B. Participant Feedback

All the participant feedback relates to driving the wheelchair whilst performing the secondary task. When doing a *between* subjects analysis of the Likert data, we did not find any statistically significant results; this is likely to be due to the subjective nature of what people find to be *easy*. Therefore, we performed a *within* subjects analysis, by comparing the Likert ranking of each statement for the case when assistance was given with that when no assistance was given, for each individual participant. This allowed us to generate the graph in Fig. 8, which shows the condition when people most strongly agreed with each statement.

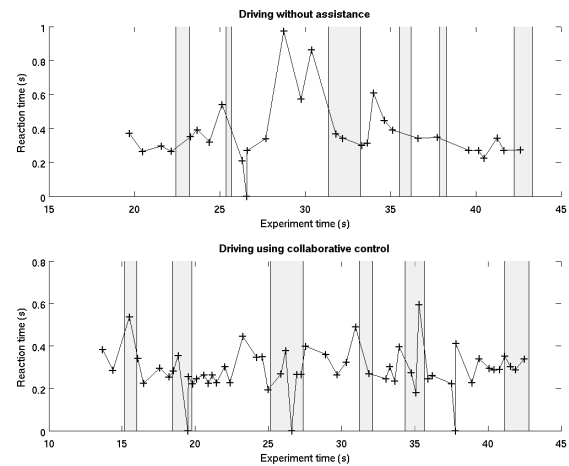


Fig. 7. A comparison of the typical reaction times for a trial. The highlighted areas represent the times that the wheelchair was in a doorway.

More than a third of the participants (six out of fifteen) reported that driving the wheelchair required greater concentration when no assistance was given, compared with when they were provided with adaptive assistance (Q3 of Fig. 8). Despite this, the majority of people considered the wheelchair easier to manoeuvre when not being given any assistance (Q1 of Fig. 8) and this is likely to be explained by the fact that most people stated that the wheelchair behaved in a less predictable manner when using the adaptive assistance mode (Q2 of Fig. 8). The result of Q4 in Fig. 8 also tends to suggest people felt more comfortable when they were not given any assistance.

Looking at Q5 of Fig. 8, we can see that for the majority of participants there was no perceived difference in the difficulty of the secondary task, despite people making fewer mistakes when the assisted mode was active as shown in Fig. 5. A marginal number of participants found the secondary task easier when the collaborative controller was active.

V. DISCUSSION

Some interesting results have been presented that suggest people's perception of how well they are performing is not necessarily aligned with their actual performance. Participants seemed less worried about colliding with furniture and

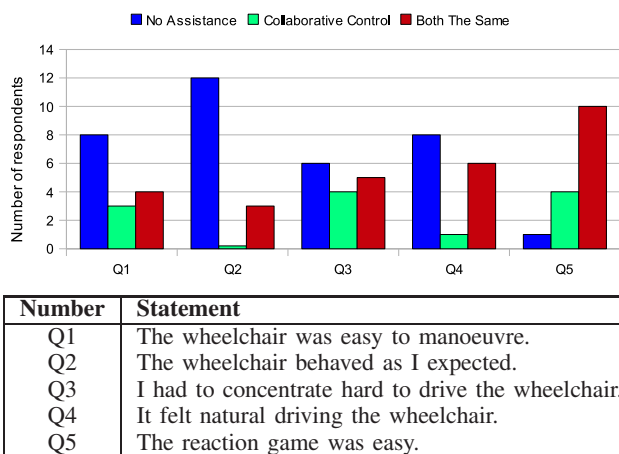


Fig. 8. Participants had to indicate their agreement with the statements (Q1-Q5), using a 5-point Likert scale. For which control method did people most strongly agree with each of the statements?

doorways than needing to feel that they were in control of the wheelchair, which echos Nisbet's message in [5]. However, it seems that the reason most people preferred not to be assisted was because they didn't always understand why the wheelchair was behaving as it did, or in other words, their mental model was not aligned with the system model, as was suggested in [8]. Some quotes from participants:

"It is easier to drive it in trial B [no assistance], probably because it brakes less."— Participant 13

Despite finding it "easier to drive", participant 13 crashed three times when not using any assistance, as can be seen in Fig. 5.

Perhaps the main problem is that it can be difficult to form a mental model of such an adaptive system, but maybe if the system gave users some appropriate feedback, they might become more comfortable. Several users made comments suggesting that the system should give them more feedback:

"There was little feedback on how the wheelchair would respond, or how to manoeuvre out of a tight section."— Participant 1

Originally, we wanted the assistance to be completely transparent to the user. However, in some situations, users are unaware as to why the wheelchair is preventing them from moving in a particular direction. For example, often when they approached the third doorway to drive back into room one, from the corridor (Fig. 4(a)), participants attempted to enter from such a shallow angle that the virtual bumper prevented them from turning right (so as not to hit the door-frame). However, in these cases, the users generally took a while to realise why the system was preventing them from turning.

VI. CONCLUSIONS

We have shown that when a user is distracted by an engaging secondary task, requiring partial visual attention, they find it difficult to drive safely around an office environment, often resulting in collisions with furniture, door frames and even people. Conversely when the user is asked

to perform the same task, but is given adaptive assistance by the collaborative controller and an additional level of safety by the dynamic local obstacle avoidance module, they do not crash and hence drive more safely. Moreover, they perform significantly better in the secondary task; their reaction times are similar, but the percentage of incorrect reactions is lower.

Although one user did not manage to perform the experiment without assistance, the majority preferred to drive without collaborative control. However, even these more proficient participants found the assistance helpful, once they were engaged in the secondary task. In naturalistic, non-laboratory environments, wheelchair users are often engaged in multiple tasks as they go about their activities of daily living, for example whilst they are shopping [9]. Our research has therefore taken an important step into disambiguating the conditions under which shared control is beneficial.

VII. ACKNOWLEDGEMENTS

The authors would like to thank all the volunteers who participated in our experiments. We would also like to thank the members of the BioART team for their insightful comments and continued support.

REFERENCES

- [1] T. Taha, J. Miro, and G. Dissanayake, "Pomdp-based long-term user intention prediction for wheelchair navigation," in *IEEE International Conference on Robotics and Automation (ICRA)*, Pasadena, LA, May 2008, pp. 3920–3925.
- [2] J. Millán, F. Renkens, J. Mouríño, and W. Gerstner, "Noninvasive brain-actuated control of a mobile robot by human eeg," *IEEE Trans Biomed Eng.*, vol. 51, no. 6, pp. 1026–1033, June 2004.
- [3] S. Levine, D. Bell, L. Jaros, R. Simpson, Y. Koren, and J. Borenstein, "The navchair assistive wheelchair navigation system," *IEEE Transactions on Rehabilitation Engineering*, vol. 7, no. 6, pp. 443–451, 1999.
- [4] D. Ding and R. A. Cooper, "Electric powered wheelchairs: A review of current technology and insight into future directions," *IEEE Control Systems Magazine*, vol. 25, no. 2, pp. 22–34, April 2005.
- [5] P. Nisbet, "Who's intelligent? Wheelchair, driver or both?" in *Proc. IEEE Intl. Conference on Control Applications*, Glasgow, Scotland, U.K., September 2002, pp. 760–765.
- [6] T. Carlson and Y. Demiris, "Human-wheelchair collaboration through prediction of intention and adaptive assistance," in *Proc. of IEEE International Conference on Robotics and Automation*, Pasadena, CA, 2008, pp. 3926–3931.
- [7] K. Tsui, H. Yanco, D. J. Feil-Seifer, and M. J. Matarić, "Survey of domain-specific performance measures in assistive robotic technology," in *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, Washington, D.C., Aug 2008, pp. 116–123.
- [8] T. Carlson and Y. Demiris, "Using visual attention to evaluate collaborative control architectures for human robot interaction," in *New Frontiers in Human Robot Interaction, a symposium at AISB 2009*, Edinburgh, UK, April 2009, pp. 38–43.
- [9] Å. Brandt, S. Iwarsson, and A. Ståhle, "Older people's use of powered wheelchairs for activity and participation," *Journal of rehabilitation medicine*, vol. 36, no. 2, pp. 70–77, 2004.
- [10] S. Parikh, V. J. Grassi, V. Kumar, and J. Jun Okamoto, "Usability study of a control framework for an intelligent wheelchair," in *IEEE International Conference on Robotics and Automation (ICRA)*, Barcelona, Spain, April 2005, pp. 4745–4750.
- [11] Y. Demiris, "Prediction of intent in robotics and multi-agent systems," *Cognitive Processing*, vol. 8, no. 3, pp. 151–158, September 2007.
- [12] J. Borenstein and Y. Koren, "The vector field histogram - fast obstacle avoidance for mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 278–288, 1991.
- [13] D. Lamble, T. Kauranen, M. Laakso, and H. Summala, "Cognitive load and detection thresholds in car following situations: safety implications for using mobile (cellular) telephones while driving," *Accident Analysis and Prevention*, vol. 31, no. 6, pp. 617–623, 1999.