Human-Wheelchair Collaboration Through Prediction of Intention and Adaptive Assistance

Tom Carlson and Yiannis Demiris

Abstract— Powered wheelchair users want to be active drivers, not just passengers. However, in some situations (varying from person to person), they may require assistance; hence, research is being carried out into the development of 'smart' wheelchairs. Predominantly, this research has been derived from the field of mobile robotics, focussing on creating autonomous systems, which unfortunately tend to treat the human as little more than a precious piece of cargo. Instead, the design should be based around each individual user's abilities and desires, maximising the amount of control they are given. In this paper, we look at how collaborative control techniques can be used to achieve this, offering the user help, as and when it is required. We then evaluate the effects of this collaboration, which is built by predicting user intentions and responding to these predictions with adaptable levels of assistance.

I. INTRODUCTION

Electrically-powered wheelchairs are becoming an increasingly common solution to the lack of independence suffered by the mobility-impaired. However, a substantial number of users find it difficult to operate their chairs effectively; this can be due to a variety of physical, perceptive or cognitive impairments [17]. Ding and Cooper review the multitude of problems faced by powered wheelchair users and discuss improvements that can be made in the low-level control (velocity, traction, suspension etc.) as well as touching briefly on the higher level navigational assistance [7]. In this paper we focus on the high level control system that forms the core of our 'smart' chair.

Although many 'smart' systems are being developed, they often approach the problem from a traditional mobile robotics point of view, which means creating fully autonomous solutions that make optimal decisions based upon factors such as speed and distance travelled. In such a design, the human plays an almost insignificant role, perhaps occasionally offering a few high-level suggestions. Conversely the design approach should be to focus on the needs and abilities of the user [14], whilst considering safety to be of paramount importance. In this study, we develop an effective collaborative control system, in which the user is an integral part.

Traditionally, powered wheelchairs have been driven with a joystick, which has proven to be an intuitive solution. Unfortunately — in order to drive both efficiently and safely — this requires the user to have steady hand-control and good reactions. Some users are unable to provide this level

T. Carlson and Y. Demiris are with the Department of Electrical and Electronic Engineering, Imperial College London, SW7 2AZ, UK `tom.carlson02@imperial.ac.uk`, `y.demiris@imperial.ac.uk`
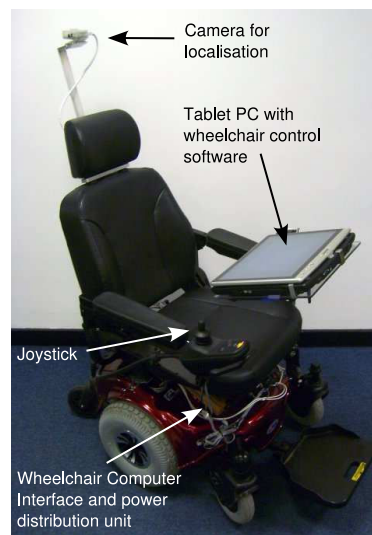
Fig. 1. The current configuration of the wheelchair. The software on the tablet PC uses the stimulus from the joystick and the camera to collaborate with the user in controlling the wheelchair motion.

of sustained control; consequently, alternative methods of interaction are being investigated. Preliminary work has been carried out in the fields of speech [16], gesture [11], [9] and gaze-direction recognition [13] for this application, as well as in more novel fields, such as brain-actuated control [15].

We believe that in many cases, a more sophisticated *intelligent* controller could compensate for the lack of steady joystick control and poor reactions, if it were not only aware of it's surroundings, but also of the user's higher-level intentions. Although we recognise that the previously mentioned multimodal input approaches can be useful in extreme cases, most of our work has been based upon human interaction with a standard wheelchair joystick.

This paper will briefly describe the work that we have undertaken in the field of collaborative control, discuss our findings and look at where our current research efforts are placed. First, we will introduce the wheelchair platform that we have developed. We will then describe the two parts of our collaborative architecture: intention prediction (or plan recognition) and adaptive assistance. Finally, after the analysis of our initial results, we will summarise our conclusions and look towards the future.

II. THE SYSTEM ARCHITECTURE

Our system is built around an EPIOC (electrically powered indoor/outdoor chair), upon which we have mounted a tablet
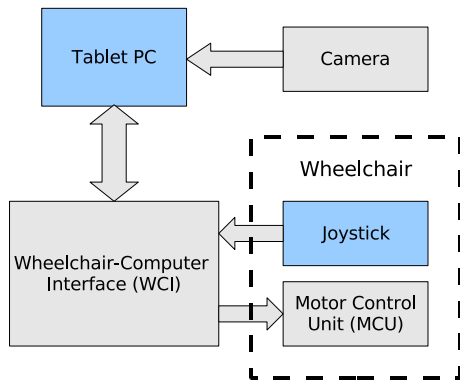
Fig. 2. This system diagram highlights the current methods of user interaction: through the joystick or the tablet PC. All the joystick commands are processed by the computer before being sent to the Motor Control Unit (MCU).
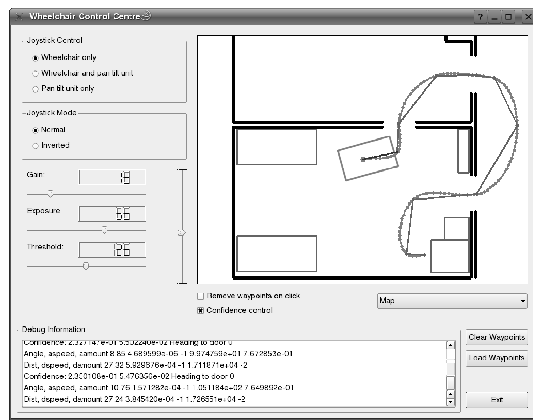

Fig. 3. The experimental GUI, displaying 9 user-instantiated waypoints, which have been interpolated with B-splines. All the features of the wheelchair control system can be easily configured by intuitively pointing and clicking with the tablet pen.

PC and interfaced it with both the joystick and motor control unit, as shown in Fig. 1. This allows us to intercept joystick signals and alter them (where necessary), before sending them to the wheelchair's motor control unit (Fig. 2). We have also developed a computer vision-based localisation system that works in mapped, indoor environments (with minimal modification of the environment).

A. Software Interface

The wheelchair control application running on the tablet PC lies at the centre of the system and is operated through a graphical user interface (GUI). The user can interactively place waypoints on the displayed map, which are automatically interpolated using B-splines, to create a smooth path. These waypoints are easily deleted or dragged around on the map at any time to amend the desired driving trajectory. The chair can then autonomously follow the given path by making use of the inverse models we have developed (discussed in more detail in Section II-C). Although we are not concerned with this type of interaction, it does form the basis of the adaptive assistance mode that will be described later, in Section III-B.

B. Localisation

In order to begin to understand what the human intends to do, the wheelchair must first be aware of its surroundings. It must also know where it is in relation to some sort of *world coordinate system*. Therefore, we will briefly discuss our current solution to the self-localisation problem.

To simplify the problem, we shall, for the moment, assume the wheelchair will be operating in a known, indoor, mapped environment. Although GPS (the Global Positioning System) would be the natural choice for an outdoor, mapped environment, it requires line-of-sight to the satellites and is therefore unsuitable for use indoors [18]. Consequently — building upon the idea of Kalkusch et al. at the Vienna University of Technology [10] — we decided to use a computer-vision based approach to determine the chair's location.

We placed fiducials (fixed 2D markers) at regular intervals on the ceiling (to prevent them from being obscured by other objects in the scene). A camera was then positioned looking directly towards the ceiling, i.e. with its z-axis perpendicular to the plane of the fiducials. To overcome the extremes of brightness caused by the lighting, an adaptive Gaussian thresholding function is applied to the images. Once a fiducial has been detected in the camera's viewport, a transformation matrix is computed — based upon the position, size and orientation of the marker — that determines the camera's position relative to that specific marker. Since the fiducial's position is known in the *global coordinate system* and the relative placement of the camera on the wheelchair is also known, we can plot the location of the chair on a map to within 5cm and 2 degrees orientation.

C. Path Following Module

If the wheelchair is going to be able to move to arbitrary points on a map, it must know how to actuate its motors to reach these positions. We use the term *inverse models*, to describe functions that generate the control commands required to reach a specified target state, given the current state of the system [6].

In our architecture, these are based on two primitive functions: a driving-forward model and a turning left/right model. The underlying mechanism of each of these models is built using a PID controller. This means the generated control signals have components which are proportional to: the error signal, the integral (or accumulation) of the error signal[1] and the derivative of the error signal[2]. In our case, the two error signals we use are the distance and angle to the target from the current location of the chair. When operating autonomously, we feed the inverse models with targets, which are successive points along the computed spline.

III. COLLABORATIVE CONTROL

A shared control system for a 'smart' wheelchair must be able to: determine the user's intention; verify the desired

---

[1] The integral affects the final spatial accuracy of the movement
[2] The derivative affects the damping, in order to prevent overshoot
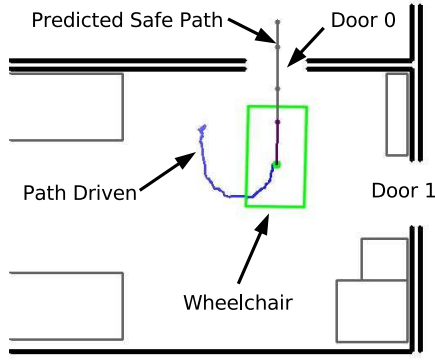
Fig. 4. The wheelchair is shown at the point where the 'Door 0' confidence crosses the threshold, as shown in Fig. 6. The path along which it has already travelled is plotted, along with four waypoints, which have been generated to form a safe passage through the doorway.

Fig. 5. A participant performing the manoeuvre shown in Fig. 4.

action is safe to perform; and, where necessary, adjust the resultant control signals to achieve the goal safely. A safe action is one that doesn't result in an impact with another object. If a crash looks likely, evasive action must be taken and many effective algorithms to implement this have been presented in the field of *route planning and collision avoidance* [12], [1], [8].

We extend the idea of *orientation correction*, where the heading of the wheelchair is constrained to fall within a certain error margin of a pre-selected goal [15], by introducing the concept of *safe mini-trajectories*. These are dynamically generated paths, which provide a safe passage from the current wheelchair position to a sub-goal (e.g. through a doorway). In addition, rather than pre-selecting a single target, we continuously update our prediction of the user's intentions, based upon the affordances of the surroundings.

In this paper we demonstrate our system using a cut-down example scenario, which will be generalised in future work. The user begins in an uncluttered office and has the option of driving around the office, or through one of two narrow doorways; *Door 0* links to the adjoining office and *Door 1* goes into the corridor (as shown in Fig. 4). The task for the wheelchair is to identify whether or not the user intends to drive through either of the doorways, and if so, guide them through safely. Therefore, we will first look at predicting the user's intentions, before deciding how to assist them in performing the desired manoeuvre. Fig. 5 shows a series of photographs of one of the trials.

A. Prediction of Intent

Many different approaches exist for intention prediction and plan recognition, as described in [5], [2], so we will explain how we came to choose our architecture. The notion of plan recognition can be split into two categories: *intended recognition* and *keyhole recognition* as defined by [4]. Essentially, *intended recognition* is when the user actively wants the system to understand their intentions, whereas the latter is when the system tries to be helpful, whilst observing the user unobtrusively. Although a wheelchair driver is actively
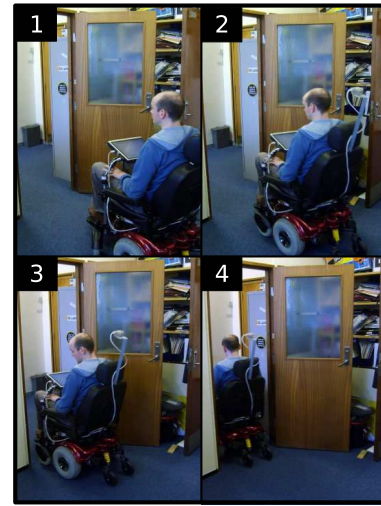
communicating with the system — in terms of moving the chair in the desired direction — they are not trying to explain their overall goal and so we should treat the plan inference as *keyhole recognition*. This way, the user can drive naturally, without the additional cognitive load of worrying whether or not the wheelchair understands their intentions; the system will try to be helpful when it believes help is required.

We perform the plan recognition using a multiple hypothesis method, following the approach we used in action recognition and imitation [6]. In this approach, all the user's known actions are represented by inverse models. Between them, they predict in parallel the required states of the system to achieve each of these tasks. By comparing the actual state of the system with these predictions, we generate a *confidence* of each task being undertaken.

In our example scenario, the driver can choose between two doorways (or neither). Therefore, we had to design a local model that represents the action *moving towards a doorway*. We achieved this by defining a confidence function $C = C_D C_\theta$, which increases when moving towards a target. This function is the product of two parts: the first (Equation 1) is computed using the Euclidean distance from the current wheelchair position (x, y) to the target (x_t, y_t), the second (Equation 3) is based upon the heading of the chair θ, compared with the angle to the target ϕ (Equation 2). The scaling factor k in Equation 3 determines the sensitivity towards the angular error and was experimentally set to 2.0.

$$C_d = exp\left\{-sqrt\{(x - x_t)^2 + (y - y_t)^2\}\right\} \qquad (1)$$

$$\phi = tan^{-1}\left(\frac{x - x_t}{y - y_t}\right) \qquad (2)$$

$$C_\theta = exp\left\{\frac{k(\pi - |\theta - \phi|)}{\pi} - k\right\} \qquad (3)$$

The choice of using exponentials as the basis for our confidence value, means that it falls off steeply as spatial or angular errors are introduced. The resultant function also
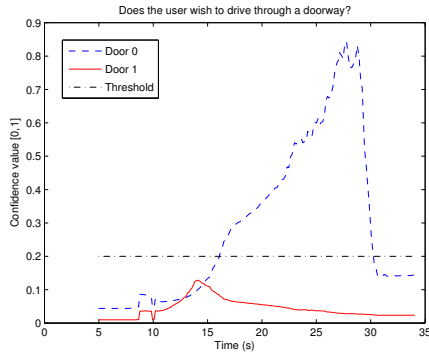
Fig. 6. The confidence functions evaluated as the user drives towards, through and away from Door 0. Note the steep drop-off in confidence due to the C_θ component, once the wheelchair has passed through the door.
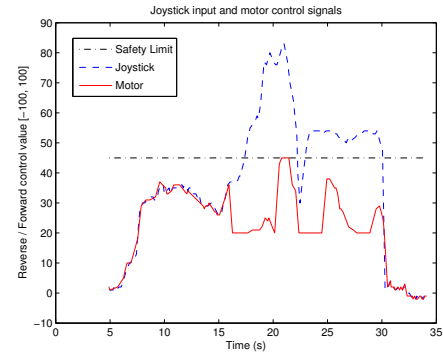

Fig. 7. The motor command signals normally follow those of the joystick. However, between 16 and 30 seconds, the assistance mode is active, so less attention is paid to the joystick data and more emphasis is placed on following the predicted path (through the waypoints shown in Fig. 4).
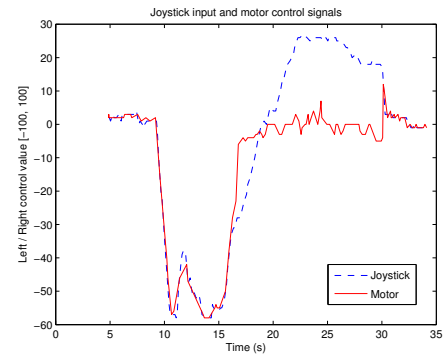

Fig. 8. Similar to Fig. 7, the steering signals are modified to prevent the wheelchair from crashing into the doorframe.

has the desirable property of scaling the output so that it falls in the interval (0, 1]. Since the confidence values of each inverse model will be competing, they can be much more effectively compared if they are known to fall on the same interval. However, we also introduce the option that the user is not performing any of the known tasks. This is achieved by introducing a confidence threshold value, below which, no assistance is given. Once this threshold has been breached, we apply winner-takes-all to determine the user's intention.

Several models can be easily generated simply by storing the coordinates of interesting targets; in our case, the two doorways. After some experimentation, we set the confidence threshold C_{thresh} to be 0.2, which allowed for a significant margin of error, preventing false positives. Fig. 6 shows how the confidence values change (and the clear separation between them) as the wheelchair performs the manoeuvre illustrated in Fig. 4.

B. Adaptive Assistance

If the system becomes very confident that a user is aiming for a specific goal, but then their input begins to deviate from the model, some assistance may be required. Alternatively they may have changed their plans; hence the need to adapt the level of assistance based upon the affordances of the situation.

Our approach is to gently guide the wheelchair towards the first waypoint of the *safe mini-trajectory*, once we are confident this is where they are headed. However, if they create large joystick signals that oppose this gentle attraction, we allow them to deviate from the target and the confidence value will naturally fall accordingly; thus allowing them to regain full control if necessary. Conversely, if they reach the first waypoint, we will prevent them from deviating from the safe path. Nonetheless, in a manner similar to that of Zeng et al. [19], the speed of the manoeuvre is still controlled by the user (it is proportional to the amplitude of the joystick *forward* value), whilst the direction is determined by the intelligent controller (such that the chair follows the *safe* path through the doorway). This continues until the corresponding

confidence value has dropped below C_{thresh}, which happens once the chair has successfully passed through the doorway. We also allow the user to reverse backwards along the safe path at any time, until the confidence value drops below C_{thresh} and they revert to *normal* control. By using this strategy, we hypothesise that the user will feel much more in control than using a rigid method which forces you to stay on a computer-controlled path at all times.

In our experiments, the *safe* path was set to be a straight line, perpendicular to and equidistant from the doorframe, that extended 60cm in each direction. Typical ammendments to the control signals are shown in Fig. 7 and 8. The driving signals sent to the motor control unit normally closely follow those of the joystick, as one would expect. However, for the period between 16 and 30 seconds — where the confidence value rises above C_{thresh} in Fig. 6 — the assisted control mode is active. This can result in significantly different motor command signals compared with the input we obtain from the joystick. It is also worth noting the safety limit we have imposed (shown in Fig. 7), this prevents the chair from accelerating rapidly and also limits it's maximum speed to 15cm/s.

3929

IV. Evaluation

In a series of short experiments, seventeen subjects (twelve male and five female, aged 20 to 46) were each asked to drive from a fixed starting point, through *Door 1* and stop when the vehicle was clear of the opening. The tablet PC time-logged a variety of important statistics, relating to the confidence values, joystick commands, motor commands, wheelchair position etc.. These were then used offline to calculate the time taken to travel through the doorway and measure of the *quality* of the trajectory. The time taken was defined as the duration for which the value corresponding to the *Door 1* confidence was greater than the confidence threshold C_{thresh}.

Each participant was required to perform a trial with the collaborative system active and a trial using only the standard joystick control. However, to eliminate biases, we changed the order in which the trials were executed, such that odd-numbered particpants started with the collaborative system active, whereas even-numbered participants began without any assistance.

Typically the performance of a control algorithm is measured in terms of speed and accuracy. Our collaborative control method exists to enable a wheelchair user to manoeuvre through a doorway, who previously would be unable to do so safely and effectively. Therefore, we place significantly more emphasis on the evaluation of accuracy compared with that of speed. However, we have included, in the interest of completeness, some results relating to the time taken for our shared control system to drive through a doorway. These are compared, in Fig. 9(a), with the time taken for a selection of able-bodied users to manoeuvre through the same doorway without the additional assistance.

When the wheelchair is driven by the assisted control mode, execution time is greater than that of an able-bodied user manoeuvring through a doorway. In fact, on average the collaborative system operates at approximately half the speed of the non-assisted mode, as can be seen in Fig. 9(a). The main reason for this is that when we designed the controller (inverse models for the wheelchair's primitive movements), we placed much greater emphasis on accuracy rather than to speed, because safety is our foremost concern. In practice, this means the chair will behave more cautiously, perhaps slowing down significantly to make safe turns, whereas a human may not decelerate to such an extent.

Next, we define a *safety deviation metric* (SDM) to measure the *quality* of the trajectory followed, whilst driving through a doorway. This is based upon $d_{min}^2[n]$, which is defined as the square of the minimum Euclidean distance between the nth point on the actual trajectory and any point on the computer-generated *safe* trajectory. Consequently, this metric places no penalty on the overall time taken to execute the manoeuvre, instead, great importance is placed on following the *safe* path as closely as possible.

$$SDM = log \left[1 + \frac{1}{N} \sum_{n=0}^{N} d_{min}^2[n] \right] \quad \forall n, \ C[n] > C_{thresh}$$
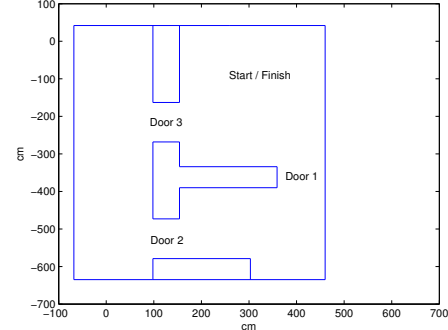
(4)

Fig. 10. The second experimental course. Participants were asked to drive from the start, through doors 1,2 and 3 (in order) to reach the finish position.

Some interesting results are presented in Fig. 9(b). The four subjects (9–12) who performed slightly better without assistance, were all male, one of whom had prior experience. However, in over 75% of the cases, the collaborative system improved the trajectories driven, giving a lower SDM compared with manual control. In more than a third of cases (1, 2, 7, 8, 13 and 17), this shift was dramatic, resulting in an improvement of over 50%. The overall improvement across all the trials is reflected by the significantly lower mean SDM achieved by the collaborative system, as show in Figure 9(c). The significance of these results was confirmed using a paired one-tailed t test ($p < 0.008$).

The large standard deviation of the SDM for the manual mode clearly shows that some users are much more adept at manoeuvring the wheelchair than others (Fig. 9(c)). This justifies the need for *adaptive* assistance, which allows them to make the most of their capabilities. Our collaborative controller provides this opportunity, resulting in a significantly smaller standard deviation of the SDM. It is also important to note that the mean variation from the safe path for the collaborative control is almost half that of the non-assisted mode. In practice, this means that on average, the collaborative controller maintains a larger safety distance from the doorframe, compared with the non-assisted mode, thereby reducing the chance of a collision. This is an encouraging result, which will enable us to move forward and test the system with representative disabled users.

We extracted data from a separate set of experiments — which investigated dexterity and shared control [3] — to again compare the SDM of the manual and assisted modes, checking the results with a paired one-tailed t test. In these trials, 20 participants, within an age range of 23 to 56 (mean 33.4, standard deviation 12.0), were asked to drive safely through three doorways (as shown in Fig. 10). This time, although the collaborative controller on average improved the trajectories driven, it was more significant for door 1 ($p < 0.011$) and door 3 ($p < 0.009$) than for door 2 ($p < 0.042$). This was most likely due to the more straight forward approach to door 2, therefore requiring less intervention from the assistance mode. Again, this highlights the importance of the adaptive controller, which provides an
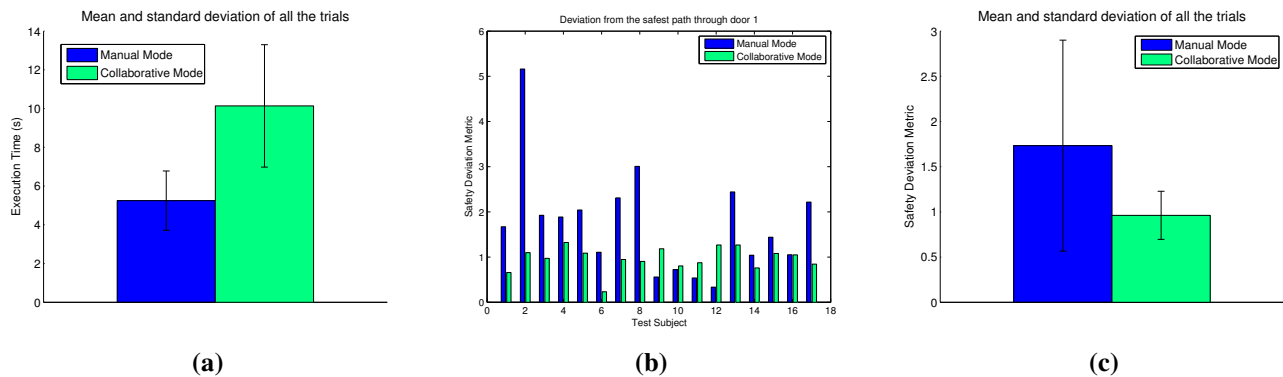
3930

Fig. 9. **(a)** The time taken to manoeuvre through a doorway using traditional joystick operation, compared with the time taken when using assistance mode (for seventeen users). **(b)** A measure of the deviation from the safest path (SDM) when driving with traditional joystick control, compared with collaborative control. **(c)** The mean and standard deviation of the SDM for seventeen users

appropriate amount of assistance as and when it is required.

V. CONCLUSIONS

This paper has presented a solid stepping stone towards creating a viable collaborative control system for use with a powered wheelchair. In order to provide useful assistance to a wheelchair driver, we aim to understand their particular needs and intentions. Our approach differs from similar works such as [15], [19], by continuously predicting the user's intentions using a multiple hypotheses method and dynamically generating *safe* trajectories. We then respond by offering adaptive assistance when a difficult task has been identified. This collaborative approach offers the user much greater control over the motion compared with traditional methods, whilst still keeping them safe.

The collaborative system has improved the quality of the trajectory driven by novice users, at a cost in terms of the time taken to perform the manoeuvre. However, an error in accuracy could be significantly more destructive than a delay in time, perhaps resulting in damage to the wheelchair, its surroundings or even in injury. Therefore, time is a small price to pay if the system empowers someone to perform activities of daily living by moving around both safely and independently.

VI. ACKNOWLEDGEMENTS

The authors would like to thank all the members of the BioART team — Simon Butler, Anthony Dearden, Matthew Johnson, Bálint Takács, Amir Vaziri, Paschalis Veskos and Kaveh Yousefi — for their continued support. We would also like to thank the participants in our experiments and José del R. Millán for his helpful comments and suggestions.

REFERENCES

[1] J. Borenstein and Y. Koren. The vector field histogram - fast obstacle avoidance for mobile robots. *IEEE Transactions on Robotics and Automation*, 7(3):278–288, 1991.
[2] S. Carberry. Techniques for plan recognition. *User Modeling and User-Adapted Interaction: The Journal of Personalization Research*, 11(1-2):31–48, 2001.
[3] T. Carlson and Y. Demiris. Collaborative control in human wheelchair interaction reduces the need for dexterity in precise manoeuvres. In *HRI'08 Workshop on Robotic Helpers*, page to appear, Amsterdam, The Netherlands, March 2008.
[4] P. R. Cohen, C. R. Perrault, and J. F. Allen. *Strategies for Natural Language Processing*, chapter Beyond Question Answering, pages 245–274. Lawrence Erlbaum Associates, 1982.
[5] Y. Demiris. Prediction of intent in robotics and multi-agent systems. *Cognitive Processing*, 8(3), September 2007.
[6] Y. Demiris and B. Khadhouri. Hierarchical attentive multiple models for execution and recognition of actions. *Robotics and Autonomous Systems*, 54:361–369, 2006.
[7] D. Ding and R. A. Cooper. Electric powered wheelchairs: A review of current technology and insight into future directions. *IEEE Control Systems Magazine*, 25(2):22–34, April 2005.
[8] S. Dubowsky, F. Genot, S. Godding, H. Kozono, A. Skwersky, H. Yu, and L. Yu. Pamm - a robotic aid to the elderly for mobility assistance and monitoring. In *IEEE International Conference on Robotics and Automation*, pages 570–576, San Francisco, 2000.
[9] P. Jia, H. H. Hu, T. Lu, and K. Yuan. Head gesture recognition for hands-free control of an intelligent wheelchair. *Industrial Robot: An International Journal*, 34(1):60–68, 2007.
[10] M. Kalkusch, T. Lidy, N. Knapp, G. Reitmayr, H. Kaufmann, and D. Schmalstieg. Structured visual markers for indoor pathfinding. In *Augmented Reality Toolkit, The First IEEE International Workshop*, 2002.
[11] S. Keates and P. Robinson. Gestures and multimodal input. *Behaviour and Information Technology*, 18(1):36–44, 1999.
[12] S. Levine, D. Bell, L. Jaros, R. Simpson, Y. Koren, and J. Borenstein. The navchair assistive wheelchair navigation system. *IEEE Transactions on Rehabilitation Engineering*, 7(6), 1999.
[13] Y. Matsumoto, T. Ino, and T. Ogasawara. Development of intelligent wheelchair system with face and gaze based interface. In *Proc. of 10th IEEE Int. Workshop on Robot and Human Communication (ROMAN 2001)*, pages 262–267, 2001.
[14] P. Nisbet. Who's intelligent? Wheelchair, driver or both? In *Proc. IEEE Intl. Conference on Control Applications*, Glasgow, Scotland, U.K., September 2002.
[15] J. Philips, J. del R. Millán, G. Vanacker, E. Lew, F. Galán, P. W. Ferrez, H. V. Brussel, and M. Nuttin. Adaptive shared control of a brain-actuated simulated wheelchair. In *Proceedings of the 2007 IEEE 10th International Conference on Rehabilitation Robotics*, pages 408–414, Noordwijk, The Netherlands, June 12-15 2007.
[16] R. Simpson and S. Levine. Adaptive shared control of a smart wheelchair operated by voice control. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems 2*, pages 622–626, Grenoble, France, September 1997.
[17] R. Simpson, D. Poirot, and F. Baxter. The hephaestus smart wheelchair system. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 10(2):118–122, June 2002.
[18] A. Smith, H. Balakrishnan, M. Goraczko, and N. B. Priyantha. Tracking moving devices with the cricket location system. In *2nd International Conference on Mobile Systems, Applications and Services (Mobisys 2004)*, Boston, MA, June 2004.
[19] Q. Zeng, E. Burdet, B. Rebsamen, and C. L. Teo. Evaluation of the collaborative wheelchair assistant system. In *IEEE Conference on Rehabilitation Robotics*, The Netherlands, June 2007.