



April 2007

Integrating Human Inputs with Autonomous Behaviors on an Intelligent Wheelchair Platform

Saranghi P. Parikh
U.S. Naval Academy

Valdir Grassi Jr.
University of Pennsylvania

R. Vijay Kumar
University of Pennsylvania, kumar@grasp.upenn.edu

Jun Okamoto Jr.
University of Sao Paulo

Follow this and additional works at: https://repository.upenn.edu/meam_papers

Recommended Citation

Parikh, Saranghi P.; Grassi, Valdir Jr.; Kumar, R. Vijay; and Okamoto, Jun Jr., "Integrating Human Inputs with Autonomous Behaviors on an Intelligent Wheelchair Platform" (2007). *Departmental Papers (MEAM)*. 82.
https://repository.upenn.edu/meam_papers/82

Copyright 2007 IEEE. Reprinted from *IEEE Intelligent Systems*, Volume 22, Issue 2, April 2007, pages 33-41.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/meam_papers/82
For more information, please contact repository@pobox.upenn.edu.

Integrating Human Inputs with Autonomous Behaviors on an Intelligent Wheelchair Platform

Abstract

Researchers have developed and assessed a computer-controlled wheelchair called the Smart Chair. A shared control framework has different levels of autonomy, allowing the human operator complete control of the chair at each level while ensuring the user's safety. The semiautonomous system incorporates deliberative motion plans or controllers, reactive behaviors, and human user inputs. At every instant in time, control inputs from three sources are integrated continuously to provide a safe trajectory to the destination. Experiments with 50 participants demonstrate quantitatively and qualitatively the benefits of human-robot augmentation in three modes of operation: manual, autonomous, and semiautonomous. This article is part of a special issue on Interacting with Autonomy.

Keywords

experimental tests, shared motion control, smart wheelchairs, user evaluations, wheelchair navigation

Comments

Copyright 2007 IEEE. Reprinted from *IEEE Intelligent Systems*, Volume 22, Issue 2, April 2007, pages 33-41.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

Integrating Human Inputs with Autonomous Behaviors on an Intelligent Wheelchair Platform

Sarangi P. Parikh, *US Naval Academy and York College of Pennsylvania*

Valdir Grassi Jr., *University of Pennsylvania and University of São Paulo*

Vijay Kumar, *University of Pennsylvania*

Jun Okamoto Jr., *University of São Paulo*

A computer-controlled wheelchair's shared control framework allows users complete control of the chair while ensuring their safety.

Nearly five million individuals in the US have limited arm and hand movement, making it difficult or impossible for them to use computers and products with embedded computers, such as wheelchairs, household appliances, office electronic equipment, and robotic aids.¹ Although some current wheelchair systems have embedded computers,

they have very little computer control and require precise, low-level control inputs from the user; interfaces are similar to those found in passenger cars. The rider must continuously specify the chair's direction and, in some cases, velocity using a joystick-like device. Unfortunately, many users who could benefit from powered wheelchairs lack these fine motor skills. For instance, those with cerebral palsy might not be able to guide a chair through a narrow opening, such as a doorway, without repeatedly colliding into the sides. These types of physically challenging environments can be frustrating and require a lot of user effort.

At the University of Pennsylvania's General Robotics, Automation, Sensing, and Perception Lab, we developed the Smart Chair, a smart wheelchair with intelligent controllers that lets people with phys-

ical disabilities overcome these difficulties. By outfitting the wheelchair with cameras, a laser range finder, and onboard processing, we give the user an adaptable, intelligent control system.

It's easy to imagine situations where the chair must respond to different types of inputs. For example, consider the chair navigating to a user-specified museum exhibit using an automatically generated deliberative plan, while also avoiding museum visitors via reactive controllers and being diverted by the user for a stop at the water cooler. Scenarios like this motivated our research (for information on other research in this area, see the "Related Work in Human-Robot Shared Control" sidebar). More generally, we address a human-in-the-loop motion planning and control framework for human-robot augmentation in an assistive technology. We systematically bring together three

Related Work in Human-Robot Shared Control

Over the past decade, more and more industrial products have used embedded computers. However, product development has emphasized designing new products for niche markets, ignoring the more difficult problem of creating products for universal access. This is especially true regarding assistive technology, where selection interfaces are often inadequate for people with physical disabilities who might not be able to perform coordinated movements efficiently. The direct inputs that users provide to intelligent machines are particularly insufficient because the performance of human-in-the-loop systems, such as wheelchairs, is affected by the user's ability to interact with the embedded computer and sensors.¹

Human-robot interaction has become an increasingly popular research topic. Although it's not a new area, many aspects of HRI haven't been explored fully. Integrating new technology with humans has been at the level of supervisory control, where the user manages the robotic system while it behaves autonomously.^{2,3} However, the more complex task of a human user and a robot sharing control to accomplish a mutual goal has received less attention.^{4,5} In mixed-initiative systems, the human user and robot share control by actively participating.⁶ By adjusting the system's autonomy, the user can collaborate with the robot at different autonomy levels. Similar to two individuals working together, as the person is able to predict the robot's behavior, trust is established. Gaining system trust means having less anxiety and frustration.

Extensive research on computer-controlled chairs has examined how the chairs use sensors and intelligent control algorithms to minimize human intervention.⁷ Numerous research groups have developed novel robotic wheelchairs. Wheelchair researchers have taken different approaches to incorporate human inputs into the control loop. One strategy is to let the user give directions to the chair in the form of commands and use the autonomous system to ensure safety by avoiding obstacles.⁸ Another is to have the wheelchair perform specified behaviors, such as following a person or tracking a line.^{9,10} At an even higher level, it's beneficial to automatically navigate to locations on a map.¹¹ At this level, landmarks or known targets are used to navigate to the desired location.¹²

References

1. S. Dubowsky et al., "PAMM: A Robotic Aid to the Elderly for Mobil-

ity Assistance and Monitoring: A 'Helping-Hand' for the Elderly," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 00)*, vol. 1, IEEE Press, 2000, pp. 570–576.

2. R.R. Murphy and E. Rogers, "Cooperative Assistance for Remote Robot Supervision," *Presence*, vol. 5, no. 2, 1996, pp. 224–240.
3. S. Thrun et al., "MINERVA: A Second-Generation Museum Tour-Guide Robot," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 99)*, vol. 3, IEEE Press, 1999, pp. 1999–2005.
4. H.A. Yanco and J. Drury, "Classifying Human-Robot Interaction: An Updated Taxonomy," *Proc. IEEE Int'l Conf. Systems, Man, and Cybernetics*, vol. 3, IEEE Press, 2004, pp. 2841–2846.
5. J. Scholtz, B. Antonishek, and J. Young, "Evaluation of a Human-Robot Interface: Development of a Situational Awareness Methodology," *Proc. 37th Ann. Hawaii Int'l Conf. System Sciences (HICSS 04)*, IEEE CS Press, 2004, p. 50130.3.
6. J.L. Marble et al., "Evaluation of Supervisory vs. Peer-Peer Interaction with Human-Robot Teams," *Proc. 37th Ann. Hawaii Int'l Conf. System Sciences (HICSS 04)*, IEEE CS Press, 2004, p. 50130.2.
7. J. Wagner et al., "Provar Assistive Robot Interface," *Proc. 6th Int'l Conf. Rehabilitation Robotics (ICORR 99)*, 1999; www.rehabrobotics.org/icorr1999/attendees/papers/wagner.pdf.
8. D. Miller and M. Slack, "Design and Testing of a Low-Cost Robotic Wheelchair Prototype," *Autonomous Robots*, vol. 2, no. 1, 1995, pp. 77–88.
9. R.C. Simpson and S.P. Levine, "Automatic Adaptation in the Navchair Assistive Wheelchair Navigation System," *IEEE Trans. Rehabilitation Eng.*, vol. 7, no. 4, 1999, pp. 452–463.
10. S. Patel et al., "Sensor Based Doorway Navigation for a Nonholonomic Vehicle," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 02)*, vol. 3, IEEE Press, 2002, pp. 3081–3086.
11. G. Bourhis and Y. Agostini, "Man-Machine Cooperation for the Control of an Intelligent Powered Wheelchair," *J. Intelligent and Robotic Systems*, vol. 22, nos. 3–4, 1998, pp. 269–287.
12. T. Gomi and A. Griffith, "Developing Intelligent Wheelchairs for the Handicapped," *Assistive Technology and Artificial Intelligence, Applications in Robotics, User Interfaces and Natural Language Processing*, LNCS 1458, Springer, 1998, pp. 150–178.

diverse and at times contradictory goals in motion planning: deliberative, reactive, and user-initiated. Our experimental results show that we're able to plan deliberative paths, use reactive controllers for safety, and integrate human inputs into our smart wheelchair system. We demonstrate the ease with which a human user can interact with the Smart Chair, intervening in real time while the chair executes an autonomous task. This flexibility lets the human user and the autonomous system truly share control of the motion. Through extensive studies of 50 individuals,

we show that it's possible to evaluate an assistive technology both quantitatively and qualitatively. We designed a single, repeated measure experiment and used performance metrics for the quantitative analysis that are related to everyday wheelchair activities. Our results indicate that user frustration levels and user preferences are inversely proportional to performance results. This is a valuable finding that other researchers in the human-robot systems field might want to consider and follow up on when developing similar motion control systems.

Background

Our research provides an integrated solution to motion planning and control with inputs from three sources. At the highest level, the human operator can specify goal destinations using a simple visual interface by pointing to locations or features on a display.² This input lets the chair automatically generate a deliberative plan incorporating prior world knowledge. At the intermediate level, the human operator must use reactive controllers to avoid obstacles and features that the sensors detect. At the lowest level,

the human operator can directly provide velocity commands using a joystick.

As a side note, although we've used a joystick as our manual input device for these experiments, this might not be the most convenient device for all users. So, we've also implemented other types of input devices. We've studied and successfully demonstrated voice command, a finger mouse, and a vision-based hand recognition system to replace the joystick and let the user manually control the wheelchair.

Before we conducted any experiments, the University of Pennsylvania's Institutional Review Board for human subjects research approved the study. As part of the protocol, we obtained informed consent from all the participants. In the experimental design's preliminary stages, we brought our wheelchair system to a rehabilitation facility. Although disabled individuals couldn't test drive the system at the rehab center, we demonstrated various functionalities in real time. The intended user group and rehabilitation specialists gave us useful feedback. Our intelligent wheelchair was well received, and the overall reaction was extremely positive.

We conducted our feasibility studies with several individuals. Unfortunately, several factors, such as the chair's physical size and transporting issues, prevented us from testing the system with users with disabilities. Although it would have been interesting to have users with disabilities drive the system, we still accomplished our goal of testing our shared control framework on a human-robot cooperative system. The experiments and analysis we present here can apply to other human-robot interaction systems, not just smart wheelchairs.

The Smart Chair

Our motorized wheelchair is equipped with onboard processing and a suite of sensors (see figure 1). The omnidirectional camera, mounted over the user's head, lets the user view 360 degrees around the wheelchair. The projector system displays images onto the lap tray and lets the user send commands to the wheelchair through a visual interface. The projector and camera systems act in concert, forming a feedback system where occluding various parts of the projected image affects user interaction.

Along with the vision system, a laser scanner is mounted in the front of the wheelchair, under the user's feet. The laser measures distances at every half degree through a 180-




Figure 1. The Smart Chair is a motorized wheelchair equipped with onboard processing and multiple sensors. The system can function as a completely autonomous mobile robot platform.

degree scan. Similarly, IR proximity sensors on the back of the chair detect any obstacles located behind the wheelchair. Lastly, encoders on the motors provide a dead reckoning system for the wheelchair. This is augmented, when necessary, by vision-based localization using an omnidirectional camera and landmarks on the ceiling. A simple triangulation scheme is implemented with the recognized landmarks to determine the wheelchair's position and orientation. Although this is convenient for an indoor application, when the system is deployed outdoors, we'll need to incorporate a global positioning unit. We've performed preliminary experiments using a GPS unit and believe it's a feasible upgrade to the system. Other work discusses the wheelchair platform in greater detail.²

Shared control framework

In most assistive devices, the operator and

the robot must share control of the system. While humans bring experience and global knowledge, a robot can help increase precision and reduce fatigue when assisting in tasks. We've developed a computer-mediated motion planning and control framework that lets the human user share control with the robotic assistive technology on various levels.

In the semiautonomous mode of operation, the user initially selects the goal and the chair generates a deliberative plan using an approximate navigation function. The user lets the system begin behaving autonomously, but the user can interrupt or change the system's behavior at any time the task is executing. We designed a control framework that lets us smoothly integrate different inputs. Particularly, we're mainly concerned with the deliberative plan, reactive behaviors (such as avoiding obstacles), and any user-initiated commands.

To reach the desired goal, the robot must follow a motion plan that satisfies the constraint $\phi(q) < 0$, where $\phi(q)$ is the scalar potential field or navigation function. This constraint is satisfied when the robot moves along the negative gradient of the potential, $-\nabla\phi(q)$. If the user's input is consistent with the goal and any necessary obstacle avoidance, we let the user maintain complete control of the wheelchair. If the user's input isn't consistent with the goal, we have two options: modify the user's input to conform to the goal or temporarily abandon the goal. Modifying the user's input gives the user limited control over the system and keeps the robot motion consistent with $\phi(q) < 0$, meaning that the goal hasn't been abandoned. This also prevents us from having to globally replan, making the system more efficient for the user. In our experiment, we modified the user's input when necessary to reach the final destination. You can find further technical information and details on the specific controllers in other work.³

We can also vary the amount of autonomy a human user or machine has for each person. Those individuals without fine motor control might want the system to be less sensitive to sudden changes they accidentally apply to the system. In the experiment, we didn't allow this variable and instead represented the semiautonomous mode as a single condition. If the user intervened with the autonomous behavior, we considered it semiautonomous.

Experimental method

We aimed to investigate the viability of the developed shared control framework and the algorithms, with the goal of evaluating the different levels of human-robot cooperation.

Participants

We recruited 50 individuals for the study. For purposes of consent, participants had to be at least 18 years old. We didn't exclude individuals on the basis of age, gender, economic status, or race, and we didn't limit the study to those with computer experience because the interface was so simple.

Each participant spent approximately one hour running experiments and answering questionnaires. We obtained 43 complete data sets, which we used in the evaluations. We discarded the first five data sets because we clarified the instructions after those trials. Two other data sets had incomplete questionnaires. We analyzed demographic data (that is, gender, age, and education completed) across all groups by performing an

analysis of variance (ANOVA). We encountered no significant statistical differences at a level of significance $\alpha = 0.05$.

Independent variables

We were interested in three operation modes that let us navigate a mobile system. Each mode corresponds to a different level of robot autonomy.

Autonomous mode. In each environment, the system constructed a deliberative plan using a potential field placed on the map. By following the potential function's negative gradient, the system drove the wheelchair from the initial position to the user-specified destination. However, the deliberative controller

We aimed to investigate the viability of the developed shared control framework and the algorithms, with the goal of evaluating the different levels of human-robot cooperation.

is locally modified to accommodate obstacles via reactive controllers.

Manual mode. We instructed the users to drive the chair using the joystick. Even in the manual mode, the obstacle avoider is always turned on, and the user might sometimes have limited control. For safety reasons and according to the human subjects board, the collision avoidance system had to be active at all times. This is different from the semiautonomous mode. In the manual mode, the system will only prevent the user from colliding into obstacles and won't perform any tasks at a higher level. In a real-world situation, the user would ultimately decide whether the obstacle avoider was activated, and either the user or the chair's provider could set this option.

Semiautonomous mode. Here, individuals had the opportunity to share control with the wheelchair. When human inputs are incorporated into the system, different scenarios

can occur (see the previous section). So, the semiautonomous mode incorporates all the behaviors we've already described: a deliberative plan with reactive behaviors and user-initiated inputs. For example, if the operator wanted to approach a water fountain, he or she would need to (temporarily) manually override the autonomous behavior. However, once the user was ready to continue to the original destination, he or she would let go of the joystick and the chair would switch to the semiautonomous mode.

To evaluate an intelligent robotic wheelchair's usefulness and each mode's efficacy, we conducted a study where participants drove the chair using the different modes.

Experimental design

We randomly assigned the order of the experiments to each subject to avoid biases related to the sequence in which we tested the modes. When evaluating the data, we combined the groups only if there was no significant statistical difference (according to ANOVA or a t-test) between them. Participants drove the wheelchair in each mode in two rooms. These two environments let the user spend more time in each mode for a more accurate assessment. We performed all the experiments in two rooms with similar environments on a single floor, and we conducted the training exercises in hallways.

During each experiment, we gave the user a math test to take as a secondary task. As other researchers have suggested,⁴⁻⁶ we used the secondary-task performance as a measure in our evaluations. The math test consisted of simple third-grade arithmetic operations. We placed the test on the wheelchair's lap tray so the individuals could read the problems. We asked users to vocally answer the problems so that we could record the answers, which freed up both of the user's hands for navigation. Such a secondary task is similar to the types of interruptions that might occur during normal operation—the user might be looking at a map or having a conversation with another individual while driving to a destination. We emphasized to all subjects that the math task wasn't as important as the driving test.

Another secondary task that we asked all users to do was to pick up a specified object near the chair's path. This task is also similar to the types of tasks a user might perform while moving to a destination—for example, the user might want to pick up a ringing telephone or stop by the refrigerator to grab an

apple. Again, we told the participants to pick up the small item if they could reach it—and if they couldn't, to disregard it. We used this secondary task to assess how the individual's workload varied throughout the three levels of robot autonomy. It also let the user assess each mode's flexibility.

Procedure

We used a well-defined protocol for the experimental study. After the subjects signed a consent form, they filled out a pre-experimental questionnaire, which we used to learn about the individual's wheelchair and robotic experience. Wheelchair experience includes any time spent in a wheelchair (for instance, due to an accident) and controlling it using a standard joystick. Robotic experience refers to any experience with computer-controlled platforms that can behave autonomously. We then followed the subsequent steps for each participant.

Step 1. We brought subjects to a training environment where they could practice driving the wheelchair in each mode. All subjects spent approximately 10 to 15 minutes becoming familiar with the wheelchair.

Step 2. We brought the individual into an environment and read the instructions out loud. We asked the user to navigate from an initial position to a marked final destination in the room. The user had to avoid tables, chairs, and unmodeled obstacles in the environment. While the user was performing this primary navigation task, he or she was also completing the secondary tasks (solving simple arithmetic problems and picking up an object). To avoid bias, some participants began in one environment while others began in the other. Similarly, some individuals began using the autonomous mode while others started with the manual mode. Once the first experiment was done, we brought the user into the other room and asked him or her to run the same operation mode again.

Step 3. After the participant had run the mode in both environments, he or she filled out a post-task questionnaire and the NASA Task Load Index (TLX) form to evaluate the mode of operation he or she experienced.

We repeated the three steps for the second and third modes of operation. After we ran all three modes of operation, we asked the user to complete a post-experimental questionnaire.

Dependent measures

Throughout the experiments, we collected both quantitative and qualitative data. Either the onboard computer or the test observer gathered the quantitative data directly. We compiled the qualitative data from the questionnaires and any verbal or written user comments. In addition, observers noted comments that might not have been obvious in the collected data.

Quantitative assessments. We measured

- Time to complete a given task. This might be influenced by other factors such as the presence of obstacles, which can affect the user's driving speed.

Although we measured some methods' effectiveness purely by quantitative data, having a human interactive system means we're also interested in the user's assessments.

- Number of human-robot interactions. These occur when the user makes decisions such as mode selection. We simply use the number of times the user interacts with the joystick to count the human-robot interactions. Thus, we consider multipart movements (that is, going straight and then turning right) to be multiple interactions.
- Secondary-task effectiveness (that is, the simple math test and picking up an object).
- Distance traveled. Along with time to completion, this lets us determine the most efficient mode.
- Number of collisions, including situations where the user hit an unmodeled obstacle while driving the chair.

Ground truth. Ground truth is necessary to determine the wheelchair's position and orientation during the experiments. We represent ground truth using odometry from the wheelchair. Simple tests conducted to observe factors that could contribute to irregu-

larities in odometry, such as belt or wheel slippage, reveal that such slip is minimal or nonexistent in our test environment.¹¹ Quantitatively, in the transverse direction the average error is 2.5 percent. Because our tests show only slight deviations from an ideal performance, we accept odometry as ground truth in this article. We also note that the odometry data corroborated with our indoor vision-based localization.

Qualitative assessments. Although we measured some methods' effectiveness purely by quantitative data, having a human interactive system means we're also interested in learning about the user's assessments. Most of the questions on the questionnaires let the user rank items on a scale of zero to 10. By comparing the ratings of the three modes of operation, we were able to draw conclusions regarding workload, efficiency, and frustration. In particular, we sought to answer several questions:

- Does the user feel safe and in control of the wheelchair?
- In terms of system capabilities, does the user feel the system is rigid or flexible? How does each mode compare in terms of rigidity?
- How do the three operational modes compare to each other in terms of frustration and effort?
- What is the user's overall reaction to the intelligent wheelchair system?

Experimental results

We discussed both quantitative and qualitative measures when assessing the wheelchair platform. The user's overall satisfaction and the system's success include careful assessment of both the measurable quantities as well as the end user's impressions.

Experimental trajectories

Because each subject had the opportunity to drive the wheelchair in three available control modes, we first discuss some sample trajectories.

Autonomous mode. Figure 2 shows the wheelchair's path when using the autonomous mode. The chair drives completely autonomously using the deliberative motion plan in tandem with the system's sensors. The dashed trajectory is the path the chair would have taken if the environment was completely known and free of unmodeled

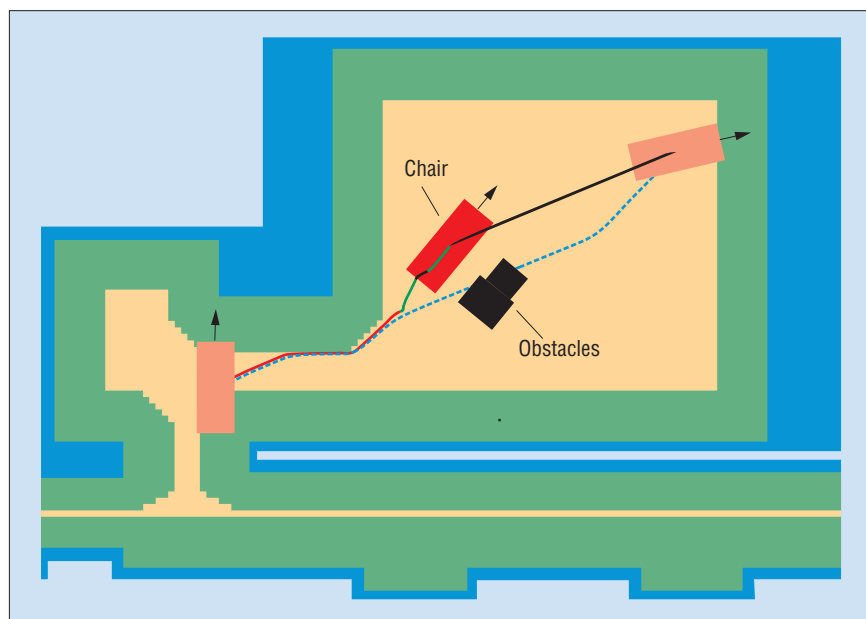


Figure 2. Sample trajectory taken by the wheelchair in room 1 while in the autonomous mode. The trajectory taken is a result of the deliberative controller (potential field controller) and local, reactive behaviors. The solid line shows when the deliberative controller is invoked. The green line segments represent locations where the activated obstacle avoidance behavior is composed with the deliberative behavior. The blue line is the trajectory that the chair takes if there are no obstacles.

obstacles. All trajectories taken by different users in the autonomous mode are similar to this sample trajectory. After at least 50 runs, we concluded that the autonomous system efficiently and predictably drives the wheelchair in the same manner each time.

Manual mode. Figure 3 illustrates a manual path a user took in room 1. Manual mode trajectories taken by all the participants are similar to this sample trajectory.

Semiautonomous mode. Figure 4 shows a trajectory that one of the users took. In the figure, we point out the different controllers that are used. Hence, during autonomous behavior (figure 4a), the system will follow a path that moves toward the selected goal. In figure 4b, the user's input is consistent with the goal. Here, the chair motion is modified to accommodate the user while the chair continues to move toward the goal. The part of the trajectory labeled figure 4c is a result of reactive behavior.

Quantitative results

We used several measurable quantities to compare the various modes of operation.

Time to completion. We recorded the total

time it took to complete the navigation task. To compare the three modes, we computed the average time to completion for each mode. We treated the collected data as a repeated measure across the three modes. In the experiments, we noticed that the semi-autonomous mode required the greatest amount of time (the mean was 78.1 seconds), while the autonomous was the fastest (68.5 seconds). Because all the users showed exactly the same trend, we concluded that in terms of time to completion, the modes' order isn't significant. Using ANOVA, we determined that the difference across the three modes was statistically significant, $F(2, 40) = 66.45$, $p < 0.0001$. We ran post hoc tests to determine which modes varied from the others. We concluded, for a 95-percent confidence interval, that there's a significant difference between the autonomous and semiautonomous modes and between the manual (70.5 seconds) and semiautonomous modes.

Although we expected the semiautonomous mode to have a similar completion time as the manual and autonomous modes, it was different. We attribute this to our observation of individuals driving slower than usual when switching from the autonomous to the manual mode. We noticed that individuals took manual control of the system

once they were closer to obstacles. At this point, they drove much slower than they normally did in the manual-only mode. The autonomous system had brought them closer to the obstacles than they would have gotten if they were in the manual-only mode. Being closer to the obstacles also meant that the users activated the obstacle avoidance algorithm more often than when manually driving. According to the post hoc tests, the autonomous and manual modes didn't differ to a statistically significant degree. As our observations showed, when the chair was closer to the obstacles, the manual driving was slower than the autonomous behavior. Thus, the autonomous mode (shorter distance, but slower near the obstacles) and the manual-only mode (further from obstacles, but moving at a faster speed) have similar mean completion times.

Human-robot interactions. One way to measure a user-operated system's complexity is to count the number of times the user must interact with the system to accomplish the desired task. On our smart wheelchair platform, the user can drive the chair in three ways. Each mode requires the user to input some command to the chair. In the autonomous mode, the subject selects the destination and the chair automatically drives to the goal while avoiding collisions. In this mode, we count just one user interaction because the user doesn't input any other information into the chair. Although we can count an infinite number of human interactions in the manual mode, we're able to count the approximate number of times the user moves the joystick. This lets us quantify the human-robot interactions in the manual mode. Similarly, in the semiautonomous mode, we can count the times the person uses the joystick.

Again, we used ANOVA to determine if the modes were statistically different from each other. This time, the number of human-robot interactions per minute was the repeated measure across the modes. We calculated the F ratio, which led us to conclude that changing the wheelchair from a manual to semiautonomous system significantly reduces the number of human-robot interactions necessary for successful functionality. From the data we collected, we recognized that, on average, twice as many human-robot interactions occur in the manual mode compared to the semiautonomous mode. Specifically, approximately 14 interactions occur per

minute in the semiautonomous mode and an average of 32 interactions occur per minute in the manual mode. We can confidently conclude that manual driving requires more effort than driving in the semiautonomous mode, in terms of necessary user interactions. The calculated value of $F(1, 40) = 51.61$ is much greater than $F_{critical}$ for $\alpha = 0.05$. So, we can conclude that a significant statistical difference exists between the two modes, so they can't be grouped together. Our experimental data further supports Holly Yanco's doctoral dissertation,⁸ which found that users issued fewer commands in the semiautonomous mode.

Cognitive complexity. Another measure of complexity can be made by observing an individual's cognitive abilities. In our experiments, we tested this by having the users perform a secondary task, which we used to study the mode's complexity. Each participant took three similar math tests, one in each mode of operation. We used the number of math questions answered correctly per minute to determine each mode's cognitive complexity. Using the math scores as a repeated measure, we performed a statistical analysis and compared the three modes' efficiency.

We observed that individuals solve more problems correctly per minute in the autonomous mode. However, when comparing the manual and semiautonomous modes, the results were inconclusive. We also observed that individuals who didn't perform efficiently in the autonomous mode appeared to be distracted by the secondary task of picking up the object. We should note that the path the wheelchair takes in the autonomous mode makes picking up the object very difficult. Although we told individuals to pick up the object only if possible, some were distracted because they couldn't reach the object, and, as a result, they stopped solving the math problems. On the other hand, a few individuals devoted more cognitive resources to answering the math problems and consequently didn't attempt to pick up the object. Anecdotal evidence suggests that as an individual's experience with the intelligent wheelchair system increases, the cognitive workload decreases in the automated and semiautomated modes. Our quantitative data supports this conclusion when comparing the manual and semiautonomous modes with the fully automated system. We found that 70 percent of the total subjects were cognitively most efficient in the autonomous mode. However, no

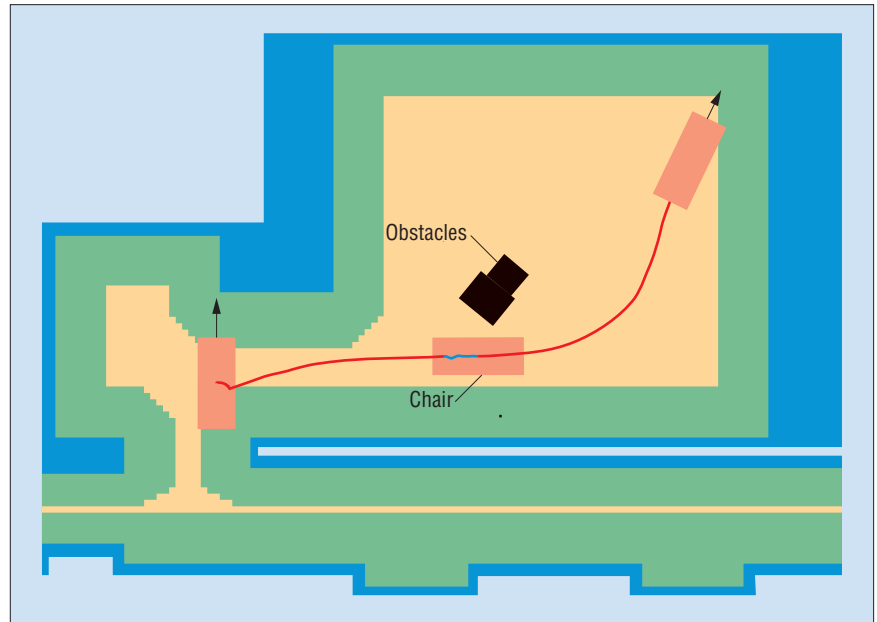


Figure 3. Sample trajectory (red) taken by the user in room 1 during the manual mode. The blue line segments represent the activation of obstacle avoidance and modification of the human input.

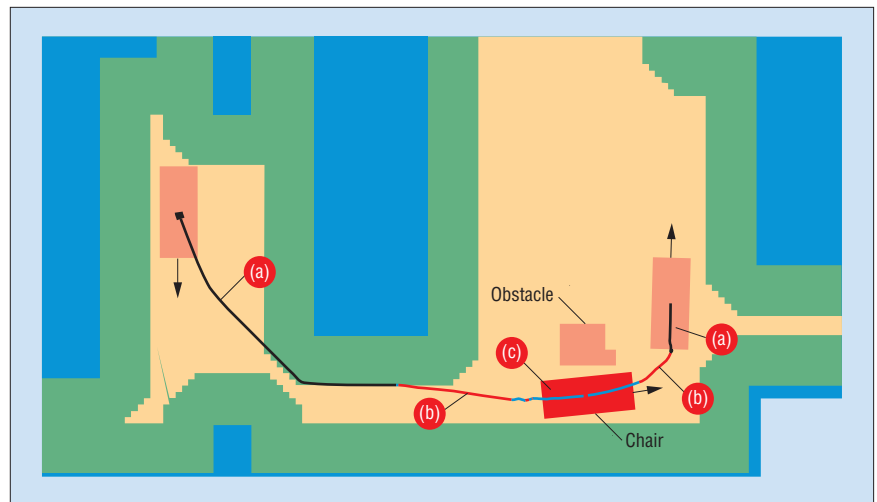


Figure 4. A sample wheelchair trajectory using a deliberative plan, combined with user input (semiautonomous mode): (a) the deliberative path taken, (b) the user's input, which is consistent with the deliberative plan, and (c) when the obstacle avoidance behavior is activated.

further quantitative data supports this conclusion when comparing the manual and semiautonomous modes of operation.

Distance traveled. When analyzing navigation tasks, another method of measuring efficiency and effectiveness is the distance traveled. Although the distances should be similar, the modes differed noticeably. The

autonomous mode always took the path that required the least amount of travel. There wasn't a significant difference between the manual and semiautonomous mode because in both of these modes, users purposely took the longer path so that they could pick up the object.

Path-length analysis indicated a significant difference between modes. Post hoc tests indi-

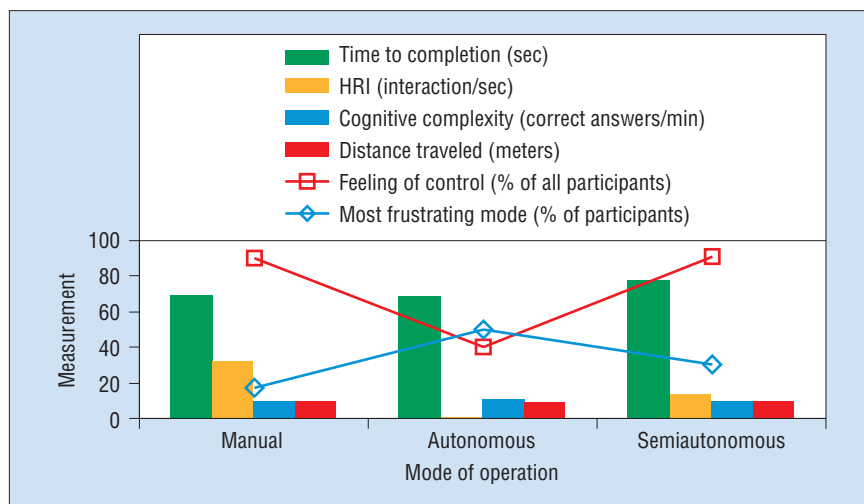


Figure 5. Quantitative and qualitative analysis.

cated that the autonomous system takes a significantly shorter path than the other modes.

Along with the distance traveled, we're interested in the system's speed. When we compared the modes' speeds, we assumed that the autonomous system would be the fastest. However, this wasn't always the case due to the system's safety boundary condition. This boundary condition required the system to drive slower than the allowed maximum speed near obstacles. Although the chair's speed slowed down in all modes when it was near obstacles, in the manual mode, individuals have more control over how far they wish to drive from obstacles. So, many individuals took paths that were further away from obstacles and therefore weren't forced to slow down. Regarding speed, statistically, there was no difference between any of the operating modes. This could be because we didn't have enough subjects—or, those who drove very fast offset those who drove slower. Thus, the manual and semiautonomous modes' average speeds were similar to the autonomous mode's average speed.

Observations and qualitative results

Throughout the experiments, the navigation task remained the same. This let us focus on the different modes' effectiveness. An overwhelming majority of the participants (greater than 90 percent) were satisfied with the system and found it easy to use. Users also judged their performance in each mode. We asked users to rate the mode on a scale from zero (inadequate) to 10 (adequate) in various categories. In terms of safety, more than 93 per-

cent of the users rated all three modes at 5 or higher. On the other hand, when asked whether they felt in control of the system, only 40 percent rated the autonomous system higher than 5. However, more than 90 percent of participants felt in control when driving in the semiautonomous and manual modes. When asked about system capabilities, 88 percent rated the overall system as very flexible.

The level of frustration appears to be related to how much control the user felt in each mode. Users were most frustrated in the autonomous mode, where they felt that they had the least control. Participants found the manual system the least frustrating and also felt that it gave them the most control. A little more than half of the participants found the semiautonomous mode more frustrating than the manual mode, but 91 percent of the total subjects felt in control in the shared control mode. Individuals might also have been frustrated by the semiautonomous mode when the obstacle avoider behavior became activated; they might have thought it was unnecessary or an unintelligent response by the system. Again, we left the obstacle avoidance algorithm on during all trials for safety. However, it's an optional feature that users can turn off (altogether or in certain scenarios) if they decide it's unnecessary.

A majority of users (90 percent) thought that the manual mode required more effort than the autonomous mode, and 60 percent felt that the manual system required more effort than the semiautonomous mode. A full 90 percent felt that the autonomous system required the least amount of effort.

A more detailed look at the semiauto-

nomous behavior helped us determine what percentage of the total time users spent in the manual mode versus the autonomous mode. While operating the system in the semiautonomous mode, individuals spent 40 percent of their time manually driving the chair and let the autonomous system control the chair the remaining 60 percent of the time.

Figure 5 shows results for both the quantitative and qualitative analysis. Using quantitative measures, the autonomous system is clearly the most efficient. However, our qualitative analysis shows the inverse relationship between frustration and control in each mode.

As we expected, the autonomous mode requires the least amount of effort from the user and is the most efficient in terms of time to completion, number of human-robot interactions, cognitive efficiency, and distance traveled. On the other hand, it's the most rigid mode, and it doesn't let the user make minor changes to the planned path to accommodate last-minute user decisions. For instance, during our experiments, users couldn't modify the path to get within range of the object that they needed to pick up. This lack of interaction and control over the chair caused frustration in some situations.

The manual mode requires the most effort from the user and is overall the least efficient. Despite this, users considered this mode the least frustrating. This is mainly because the manual mode gives the user complete motion control, letting the user drive the chair to any desired position.

The semiautonomous mode requires less effort from the user than the manual mode. Also, significantly fewer human-robot interactions occurred in the semiautonomous mode than in the manual mode. Along with providing a lighter workload, the semiautonomous mode let users control the wheelchair whenever they wanted, letting them make small changes in the path. Despite these advantages, there's an inverse relationship between "feeling in control" and "feeling frustrated." Individuals have less control in the semiautonomous mode than the manual mode, which they found frustrating.

Our shared control framework is generally applicable to a wide range of systems in which commands come from both humans and machines. This work is unique in that it reconciles different, possibly conflicting

inputs, at the continuous level. However, as we've shown, the experimental studies' results don't always follow intuition. The most significant finding was that although the overwhelming majority of users preferred semiautonomous control over the manual or autonomous modes, the autonomous mode performed consistently better in most quantitative comparisons. This suggests that integrating disparate control commands at the continuous level might not be feasible, and it points to the need for a hierarchical, discrete, and continuous framework for control. Our study also highlights the necessity of conducting experiments with human users when constructing a human-robot system. Although our semiautonomous system might not have quantitatively performed better than the autonomous system, the human operator measured a successful system differently. ■

Acknowledgments

We thank the many individuals who participated in this study.

References

1. V. Kumar, T. Rahman, and V. Krovi, "Assistive Devices for People with Motor Disabilities," *Wiley Encyclopedia for Electrical and Electronic Engineers*, J. Webster, ed., John Wiley & Sons, 1998.
2. R. Rao et al., "Human Robot Interaction: Application to Smart Wheelchairs," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 02)*, vol. 4, IEEE Press, 2002, pp. 3583–3588.
3. S.P. Parikh et al., "Incorporating User Inputs in Motion Planning for a Smart Wheelchair," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 04)*, vol. 2, IEEE Press, 2004, pp. 2043–2048.
4. J.L. Marble et al., "Evaluation of Supervisory vs. Peer-Peer Interaction with Human-Robot Teams," *Proc. 37th Ann. Hawaii Int'l Conf. System Sciences (Hicss 04)*, vol. 5, IEEE CS Press, 2004, p. 50130.2.
5. O. Nakayama et al., "Development of a Steering Entropy Method for Evaluating Driver Workload," *Int'l Congress and Exposition*, SAE tech. report 1999-01-0892, SAE Int'l, 1999.

The Authors



Sarangi P. Parikh is a faculty member in the Mechanical Engineering Group at York College of Pennsylvania. Her research interests include human-robot interaction and shared control of intelligent systems (specifically, computer-mediated motion planning and control with applications to assistive technology). She received her PhD in mechanical engineering and applied mechanics from the University of Pennsylvania. She's a member of the American Society of Mechanical Engineers and the American Society for Engineering Education. Contact her at Mechanical Eng., Dept. of Physical Sciences, York College of Pennsylvania, 890 Grantley Rd., McKay Hall 104, York, PA 17405; sarangi@alumni.upenn.edu.



Valdir Grassi Jr. is a postdoctoral researcher at the Department of Mechatronics and Mechanical Systems Engineering at the Escola Politécnica of the University of São Paulo. His research interests include mobile robot architecture and navigation, motion planning, and human-robot interaction. He received his PhD in mechanical engineering from the University of São Paulo. He's a member of the IEEE. Contact him at the Dept. of Mechatronics and Mechanical Systems Eng., Escola Politécnica da USP, Av. Prof. Mello Moraes, 2231, 05508-030 São Paulo, SP, Brazil; vgrassi@usp.br.



Vijay Kumar is on the faculty of the University of Pennsylvania's Department of Mechanical Engineering and Applied Mechanics, with a secondary appointment in the Department of Computer and Information Science. He's also the UPS Foundation Professor and the Chairman of Mechanical Engineering and Applied Mechanics at the University of Pennsylvania. His research interests include robotics and networked multiagent systems. He received his PhD in mechanical engineering from Ohio State University. He's a fellow of the American Society of Mechanical Engineers and the IEEE and a distinguished lecturer in the IEEE Robotics and Automation Society.

He served on the editorial board of *IEEE Transactions on Robotics and Automation*, *Journal of Franklin Institute*, and *ASME Journal of Mechanical Design*. Contact him at the Dept. of Mechanical Eng. & Applied Mechanics, Univ. of Pennsylvania, 3330 Walnut St., Levine Hall, GRW 470, Philadelphia, PA 19104; kumar@grasp.upenn.edu.



Jun Okamoto Jr. is an associate professor at the Department of Mechatronics and Mechanical Systems Engineering at the Escola Politécnica of the University of São Paulo. His research interests include mobile autonomous robots and the use of sensor information in the decision-making process, particularly omnidirectional vision. He received his PhD in mechanical engineering from the University of São Paulo. He's a member of the IEEE. Contact him at the Dept. of Mechatronics and Mechanical Systems Eng., Escola Politécnica da USP, Av. Prof. Mello Moraes, 2231, 05508-030 São Paulo, SP, Brazil; jokamoto@usp.br.

6. J.W. Crandall and M.A. Goodrich, "Characterizing Efficiency of Human Robot Interaction: A Case Study of Shared-Control Teleoperation," *Proc. 2002 IEEE/RSJ Int'l Conf. Intelligent Robots and Systems*, IEEE Press, 2002, pp. 1290–1295.
7. S. Patel et al., "Sensor Based Doorway Navigation for a Nonholonomic Vehicle," *Proc. IEEE Int'l Conf. Robotics and Automation (ICRA 02)*, vol. 3, IEEE Press, 2002, pp. 3081–3086.
8. H.A. Yanco, "Shared User-Computer Control of a Robotic Wheelchair System," doctoral dissertation, Dept. of Electrical Eng. and

Computer Science, Massachusetts Inst. of Technology, 2000.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.