# Shared Understanding for Collaborative Control

David J. Bruemmer, Douglas A. Few, Ronald L. Boring, Julie L. Marble, Miles C. Walton, and Curtis W. Nielsen

***Abstract*—This paper presents results from three experiments in which human operators were teamed with a mixed-initiative robot control system to accomplish various indoor search and exploration tasks. By assessing human workload and error together with overall performance, these experiments provide an objective means to contrast different modes of robot autonomy and to evaluate both the usability of the interface and the effectiveness of autonomous robot behavior. The first experiment compares the performance achieved when the robot takes initiative to support human driving with the opposite case when the human takes initiative to support autonomous robot driving. The utility of robot autonomy is shown through achievement of better performance when the robot is in the driver's seat. The second experiment introduces a virtual three-dimensional (3-D) map representation that supports collaborative understanding of the task and environment. When used in place of video, the 3-D map reduced operator workload and navigational error. By lowering bandwidth requirements, use of the virtual 3-D interface enables long-range, nonline-of-sight communication. Results from the third experiment extend the findings of experiment 1 by showing that collaborative control can increase performance and reduce error even when the complexity of the environment is increased and workload is distributed amongst multiple operators.**

***Index Terms*—Dynamic autonomy, human–robot interaction (HRI), mixed initiative, shared control.**

## I. Introduction

REMOTE deployment of mobile robots offers a compelling opportunity to merge human intelligence with machine proficiency. However, most mobile robots are currently controlled in a teleoperated fashion, such that the robot is used as a passive tool. Despite the high operator workload, communication constraints, and poor visibility, video remains the primary means of providing situation awareness, and the cognitive burdens for all decisions are placed on the operators [1]. Unless the utility of sharing control with the robot can be demonstrated in terms of performance, the potential for humans and robots to collaborate as team members will not be realized in operational settings. Human–robot interaction (HRI) evaluations can help to address these challenges.

Yanco *et al.* [2] have identified major shortcomings in current HRI evaluations. The first is that the designers of the system are often enlisted as test users. The second is that HRI evaluations are commonly informal and rarely provide controlled, objective assessment. Another shortcoming to previous HRI studies has been a lack in the number of participants. In response, the present experiments do not use system designers as test participants, involve a large participant pool, and use objective performance measures such as completion time, items found, joystick usage, communication bandwidth, and human navigational error to empirically assess the performance of a human–robot team.

Previous HRI experiments with the Idaho National Laboratory (INL) control architecture have focused on particular applications including search and rescue and remote characterization of hazardous environments [3], [4]. Such studies focus on usability and, by involving target users, make it possible to tailor the robotic behaviors and interface for a particular application. The present experiments are not intended to apply technology to a particular application domain (e.g., urban search and rescue) or investigate the use of robots by one particular group of users. In fact, for many of the application domains where intelligent robots may be used in the future, target users of such systems have not yet been identified [1]. Instead, the goal of this research is to investigate the fundamental challenges of sharing control and promoting collaborative understanding between humans and robots. Accordingly, the experiments reported here do not use experts, but rather draw on a varied pool of novice users. As such, the experiments are not to be considered as usability tests, but rather as an exploration of basic principles surrounding cooperation between humans and robots. The present experiments with novice users are intended to directly complement application specific research endeavors that involve target users. In [5], Yanco and Drury provide an excellent example of how this can be done by performing a usability study that places the technology used in the present study into the hands of Federal Emergency Management Agency (FEMA) domain area experts.

## II. Previous Research

Attempts to design behaviors for mobile robots typically fall into three categories: deliberative, reactive, and hybrid control approaches. With deliberative approaches, the robot is given or builds a world model which it then uses to plan and act, as typified by Simmons' work [6]. In these systems, it is assumed that environments are static and predictable [7]. Performance degrades once positioning is corrupted or the robot becomes otherwise unable to relate its internal model to the real world. Brooks fundamentally disagreed with the notion of static environments [8] and countered with a reaction-based architecture designed to handle dynamic environments [9]. The advantage of the reaction- or behavior-based approach [10] is that it allows the robot to respond directly to the environment rather than to a flawed internal representation. The INL robot architecture is a hybrid control architecture [11], [12] designed to capitalize on the planning aspects of the deliberative paradigm and the

Manuscript received August 13, 2004; revised February 15, 2005 and March 14, 2005. This work was supported in part by the U.S. Department of Energy under a Laboratory Directed Research and Development grant to D. J. Bruemmer. This paper was recommended by the Guest Editors.

D. J. Bruemmer, D. A. Few, R. L. Boring, J. L. Marble, and M. C. Walton are with the Idaho National Laboratory, Idaho Falls, ID 83415 USA (e-mail: david.bruemmer@inl.gov).

C. W. Nielsen is with the Computer Science Department, Brigham Young University, Provo, UT 84602 USA (e-mail: curtisn@cs.byu.edu).

Digital Object Identifier 10.1109/TSMCA.2005.850599

responsiveness of the reactive paradigm. In terms of HRI, the benefit of this hybrid approach is that the reactive elements support robust guarded motion and autonomous behavior, while the map-building and planning capacity allows the robot to provide situation awareness and communicate intentionality to the user.

An increasing number of researchers from the fields of human factors, cognitive science, and robotics are working to develop new HRI methods for remote operation of mobile vehicles (see [13] for an overview). Casper and Murphy present a post-hoc analysis of the rescue efforts at the World Trade Center in September 2001 where robots were used for the first time to assist in real, un-staged search and rescue operations [14]. Burke *et al.* present a field study on HRI in an urban search and rescue training task [15]. Yanco *et al.* [2] present an analysis of the 2002 American Association for Artificial Intelligence (AAAI) Robot Rescue Competition where robot systems were used to compete in a mock search and rescue operation. In each study, the authors noted that it was difficult for operators to navigate due to an inability to understand the robot's position and/or perspective within the remote environment. Collectively, these findings suggest that in order to move beyond the limitations discussed in the literature, there must be interface methods that more effectively promote a shared understanding of the environment and task. The present study investigates the hypothesis that in order to realize the potential benefits of collaborative control, the human operator must be able to understand robot intentionality and predict robot behavior.

An ability to understand and predict robot behavior is especially important when the robot has the ability to function in multiple modes of autonomy. There have been numerous taxonomies developed to discuss the various levels of autonomy a robot might have when interacting with a human [16]–[20]. Safeguarded teleoperation [17], [21], [22], which allows the robot to protect itself from collisions, can be considered a form of mixed-initiative interaction [23], [24], where both the human and robot take initiative to accomplish a task objective. However, although it allows the robot to take a role in navigation, safeguarded teleoperation does not permit a dynamic sharing of roles and responsibilities customarily found in an effective team [25]. In fact, almost all robotic systems, including those with autonomous capabilities, are currently used within a human supervisory control schema such that the only responsibility of the robot is to enact human commands (with more or less autonomy depending on the complexity of the commanded task) [18]. It is sometimes assumed that autonomy (i.e., full independence) is the ultimate goal for remote robotic systems [26]. One purpose of the research reported here is to suggest that effective teamwork, where the robot is a peer, is an equally profitable aim. Collaborative control as discussed by Fong *et al.* [27], [28] is a notable example of research that has attempted to treat the robot as a peer. Note that sharing control with the robot as a peer is more than allowing the robot to take initiative. It means that the robot may take a leadership role over certain task elements and that communication and support flow in both directions. Within Fong's work, sharing of control is mediated by explicit semantic dialogue between robot and human. Building on the concept that the robot can be viewed as a trusted team member, the research presented here investigates the opportunity for human–robot teaming that can support a dynamic sharing of roles and responsibilities. Whereas Fong's work allows for textual dialogue, the goal of the present study is to provide not only textual dialogue, but a collaborative, cognitive workspace—a virtual three-dimensional (3-D) representation that supports a shared understanding of the task and environment.

## III. System Design

Through the Office of the Secretary of Defense (OSD) Joint Robotics Program, the Space and Naval Warfare Systems Center San Diego and the INL have worked together to develop, mature, and integrate promising robotics technologies from throughout the robotics community including components for perception, communication, behavior and world modeling. Currently, the INL Robot Intelligence Architecture is being used to unite the selected components into a behavior kernel that can be transferred to fieldable unmanned ground vehicle systems. Permeating this intelligence architecture, multiple levels of autonomy provide the user with an ability to coordinate a variety of reactive and deliberative robot behaviors. The experiments reported here are based on this intelligence architecture.

### *A. Levels of Autonomy*

Four robot modes of control are available from the interface [3], [4], affording the robot different types of behavior and levels of autonomy.

1) *Tele mode* is a fully manual mode of operation, in which the operator must manually control all robot movement.
2) *Safe mode* is similar to Tele Mode, in that robot movement is dependent on manual control. However, in Safe Mode, the robot is equipped with a level of initiative that prevents the operator from colliding with obstacles.
3) In *shared mode*, the robot can relieve the operator from the burden of direct control, using reactive navigation to find a path based on perception of the environment. Shared Mode provides for a dynamic allocation of roles and responsibilities. The robot accepts varying levels of operator intervention and supports dialogue through the use of a finite number of scripted suggestions (e.g., "Path blocked! Continue left or right?") and other text messages that appear in a text box within the graphical interface.
4) *Autonomous mode* consists of series of high-level tasks such as patrol, search region or follow path. In Autonomous Mode, the only user intervention occurs on the tasking level; the robot itself manages all decision-making and navigation.

To investigate the challenges of sharing control, the experiments reported here focus on the middle ground (i.e., levels two and three) that falls between teleoperation and full robotic autonomy. Although the experiments restricted each participant to only one level of control, normal operation permits the user to switch between all four modes of autonomy as the task constraints, human needs and robot capabilities change. For instance, the tele mode can be useful to push open a door or shift a chair out of the way, whereas the autonomous mode is especially useful if human workload intensifies or in an area where communications to and from the robot are sporadic.

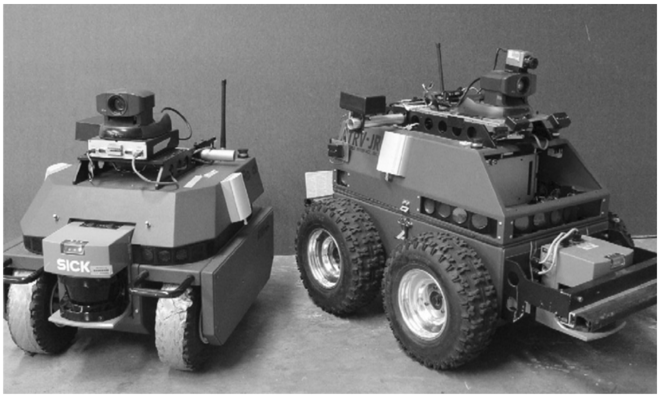

Fig. 1. Robots used for experiments 1–3.

### B. Robot Implementation

The control architecture discussed in this paper is the product of a spiral or iterative development cycle where behaviors have been evaluated in the hands of users, modified, and tested again. In fact, many of the strategies and interface components that originally seemed elegant from a conceptual standpoint, proved to be frustrating for users. For example, during a preliminary experiment [3], it was noted that although most participants felt a high level of control, some participants indicated on subjective questionnaires that they were confused by the robot initiative while in the shared mode. A review of video footage indicated that automatic initiation of the get-unstuck behavior could lead to confusion and a fight for control between robot and human. Consequently, the functionality of the shared mode was changed so that the robot prompts the user for permission before initiating the get-unstuck behavior.

Since no single platform is appropriate for all tasks, the INL has developed a behavior architecture that can port to a variety of robot geometries and sensor suites and which is being used as a standard by several HRI research teams throughout the community. Experiments discussed in this paper utilized the iRobot "ATRV mini" and the "ATRV Jr" shown in Fig. 1. On each robot, the behavior architecture utilizes a variety of sensor information including inertial sensors, compass, wheel encoders, laser, computer vision, thermal camera, infrared break beams, tilt sensors, bump sensors, sonar, and ultrasonic sensors.

Using a technique described in [29], a *guarded motion* behavior permits the robot to take initiative to avoid collisions. In response to laser and sonar range sensing of nearby obstacles, the robot scales down its speed using an event-horizon calculation, which measures the maximum speed the robot can safely travel in order to come to a stop approximately two inches from the obstacle. By scaling down the speed by many small increments, it is possible to insure that regardless of the commanded translational or rotational velocity, guarded motion will stop the robot at the same distance from an obstacle. This approach provides predictability and ensures minimal interference with the operator's control of the vehicle. If the robot is being driven near an obstacle rather than directly toward it, guarded motion will not stop the robot, but may slow its speed according to the event horizon calculation.

### C. Interface Design

The default configuration of the interface consists of a single touch screen display containing five sizeable windows (see Fig. 2). The upper left-hand window on the screen contains a video feed from the robot as well as controls for pan, tilt, and zoom. Frame size, frame rate, and compression settings can be accessed from a subwindow, but were held constant throughout the experiments reported here. The upper right-hand window contains status indicators and controls that allow the operator to monitor and configure the robot's sensor suite as needed. The lower right-hand window pertains to movement within the local environment and provides indications of direction and speed of robot motion, obstructions, resistance to motion, and feedback from contact sensors. The interface indicates blockages that impede motion in a given direction as red ovals next to the iconographic representation of the robot wheels (lower left of Fig. 2) indicating that movement right and left is not possible because of an object close to the left side wheels. These indicators allow the operator to understand why the robot has overridden a movement command. Since the visual indications can sometimes be overlooked, a force feedback joystick was also implemented to resist movement in the blocked direction. The joystick vibrates if the user continues to command movement in a direction already indicated as blocked.

At the far right, the user can select between different levels of robot autonomy. The lower central window provides an emerging map of the environment and allows the user to initiate a number of waypoint-based autonomous behaviors such as *search region* and *follow path*. Participants were permitted to use the controls on this window to zoom the map in and out. The lower left-hand window contains information about the robot's operational status such as communication activity, power and feedback regarding the robot's pitch and roll. When driving the robot directly, operators give directional commands using the joystick. For each of the three experiments, participants were explicitly instructed on how to use the onscreen controls and the joystick.

### D. Virtual 3-D Display

The goal of the 3-D display (see Fig. 3) is to provide a workspace for collaborative understanding between the human and robot. The virtual 3-D component has been developed by melding technologies from the INL [30], Brigham Young University (BYU) [30], and Stanford Research Institute (SRI) International [31], [32]. The 3-D virtual display is not based on true 3-D range sensing, but rather by extruding a two-dimensional (2-D) map to provide the user with a malleable perspective. To build the map, the INL control system uses a technique developed at SRI called Consistent Pose Estimation (CPE) that allows for efficient incorporation of new laser scan information into a growing map and addresses the problem of loop closure: how to register new laser information when the robot returns to a previously explored area.

The map produces the basis for the 3-D representation that includes obstacles and other semantic entities that are of significance to the operator such as start location, labels, and waypoints. These items can be inserted by the robot to

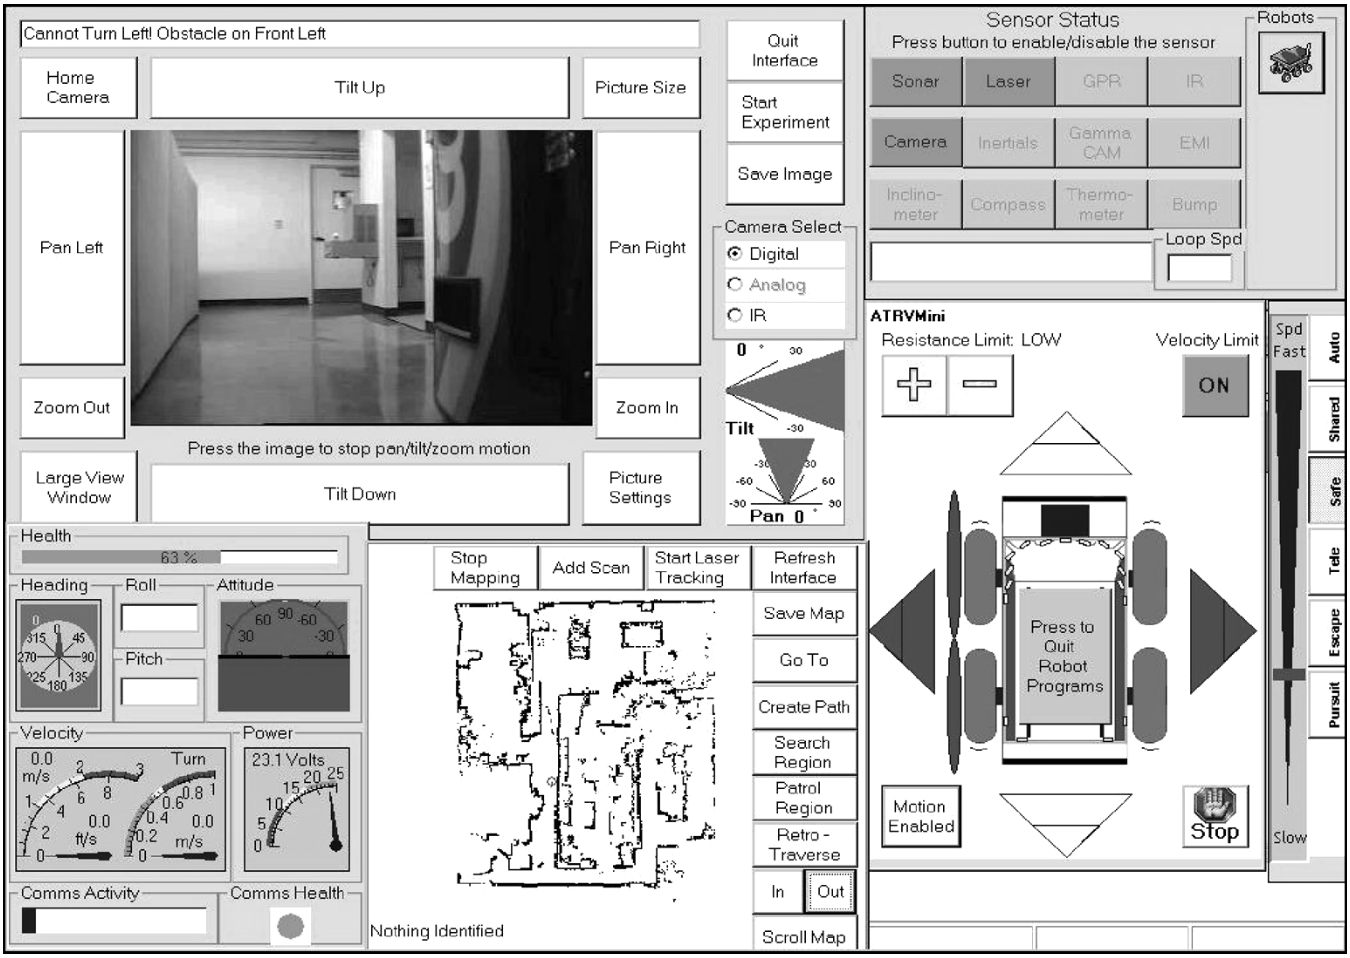


Fig. 2. Standard interface configuration.

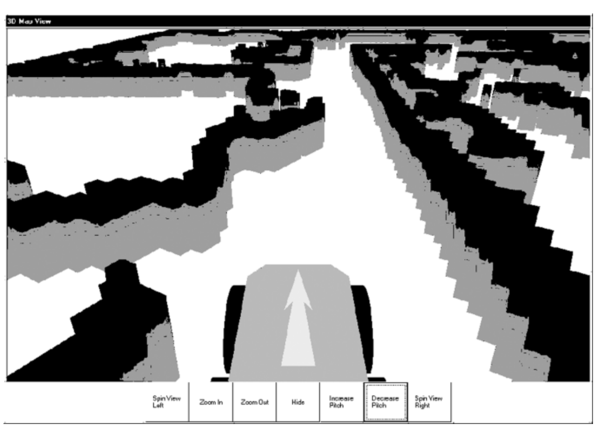

Fig. 3. Virtual 3-D display.

indicate percepts or intentions; likewise, the operator may insert entities from a drop-down menu. The operator may also insert translucent still images, excerpted from the robot video, which are overlaid onto the corresponding area of the 3-D map display, providing a means to fuse real and virtual elements. By changing the zoom, pitch, and yaw of the field of view, it is possible to move from an egocentric perspective (i.e., looking out from the robot), to a fully exocentric view, where the entire environment can be seen at once.

## IV. Experiment 1

Experiment 1 was intended to show that the behaviors on board the robot (e.g., guarded motion and autonomous navigation) were useful and could be supported by the graphical interface. Furthermore, Experiment 1 was designed to provide a means to compare autonomous driving with direct joystick control.

### A. Participants

The first study included 107 participants drawn as volunteers from attendees of the INL annual science and engineering exposition. The participants consisted of 46 females and 61 males, ranging in age from 3 to 78 years old, with a mean age of 14. Participants were asked basic questions including their age and gender, and whether they had experience in remote systems operation. It was determined by self-report that none of the participants had experience remotely controlling robots, or had knowledge of or access to the remote environment. Furthermore, none had prior experience with or knowledge of the interface or robot control system. It was determined that all of the participants could be regarded as novice users. Participants were assigned to either the shared or safe control modes alternately based on their sequence in participation.

### B. Procedure

A $20' \times 30'$ robot search environment was created for this test. The participants controlled the robot from a remote station, thereby ensuring they had no visual cues from the environment. To facilitate realistic maneuvering through an urban environment, the robot's search arena featured several obstacles. The central area was divided using conventional office dividers, while four cylindrical pylons were also placed strategically to force participants to maneuver effectively. Five objects were placed throughout the arena in locations that remained fixed across participants. These consisted of two mannequins, a stuffed dog, a disabled robot, and a small, simulated explosive device. The placement of these items also made the actual driving task more challenging, since operators were told not to drive into or over the objects. Also, certain objects remained hidden except from certain vantage points. Participants were given 60 s to locate as many of the five items in the search area as possible. Participants were assigned to alternating conditions so as to ensure approximately equal numbers of participants in each condition. No participant was allowed to operate the robot in more than one trial. Each participant was instructed on the use of the joystick for controlling the robot but was given no opportunity to practice driving prior to the trial. Additionally, each participant was instructed on the robot's camera controls (e.g., pan, tilt, and zoom). Runs alternated between use of safe mode, in which the robot takes initiative only to protect itself from collisions and shared mode, in which the robot drives autonomously but accepts periodic intervention from the operator. For participants using the shared mode, it was explained that they should let the robot do the driving, but that if they wanted to redirect the robot, the robot would temporarily yield control to their joystick commands.

### C. Results

The effects of participant age, gender, and operational mode were compared against the total number of objects that were located and identified. The results were analyzed by age in five-year intervals up to 20 years old; thereafter they were grouped in ten-year intervals. This made it possible to determine if there were differences in performance that might be possible with age, especially due to developmental differences in pre-adults. There was no significant difference in the number of objects found across participants of different ages, $F(8, 96) = 1.64, p = 0.12$. Although the 15–20 year-old age group had the highest overall average, analysis of the data showed that the fluctuations in the number of objects found by different age groups were not statistically different. There was no difference in the number of objects found due to gender. Females statistically found the same number of objects as did males, $M = 2.54$ and $M = 2.68$, respectively, $F(1, 103) = 0.31$, $p = 0.58$. There were no significant two- or three-way interactions between gender, age, and operational mode. Analysis of age and gender showed that the data would permit combination of the sample in order to assess the effects of operational mode. There was a significant difference due to operational mode, $F(1, 103) = 4.83$, $p < 0.05$. Participants who used the shared mode found an average of 2.87 objects, while those who used the safe mode found an average of 2.35 objects.

### D. Discussion

Although this experiment is not intended to support a careful comparison of age and gender groupings, it does support the claim that the interface allowed a wide variety of participants to find objects successful. As with all the experiments reported here, the use of volunteers does present the possibility that not all participants were equally motivated. Fortunately, with a duration of only 60 s, the task offered little opportunity for effects such as boredom that may be linked to task duration. Moreover, use of alternating conditions can be expected to reduce effects due to differing levels of motivation or interest.

Participants were able to find objects successfully in both the safe and shared modes, indicating that both the guarded motion used in safe mode and the autonomous navigation behaviors used in the shared mode were usable by participants. Across all age and gender groupings, performance was significantly better in the shared mode than in the safe mode, providing evidence that the robot's ability to navigate the environment can actually exceed the ability of a human operator. The performance benefit experienced by allowing the robot to navigate suggests the potential to use robot initiative and autonomy not only as a last resort (i.e., when communication fails or operator workload increases), but as the basis for sustained collaborative interaction.

Taken on its own, this first study demonstrates the utility of robot autonomy, but leaves many questions to be answered by further experiments. The first experiment did not look beyond overall performance (e.g., items found) to discern the reasons for the observed difference in performance between the safe and shared modes. In response to this limitation, it was determined that the next experiments should empirically measure differences in operator workload, operator error, and operator confusion in order to provide deeper insight. Also, the first experiment utilized a relatively small search environment. Areas of the environment required careful maneuvering, but the task was not designed to reward path planning or strategy. The question was raised of whether, in a more complex environment that required cognitive intelligence, the robot's ability to make decisions and navigate autonomously would fall short of the human's ability to maintain situation awareness. The second and third experiments were designed to answer each of these questions.

Experiment 1 also raised the question of how useful the streaming video provided by the interface actually was to users. Especially in tight spaces (where situation awareness is important to prevent collisions) participants often found the entire visual field filled by an immediate obstacle; conversely, the visual field could fail to show an obstacle if it was outside of the current visual field. One hypothesis was that in such instances video promoted a false sense of situation awareness and led to operator confusion. Consider the common scenario of a robot approaching a doorway in the safe mode. The door frame disappears from the video feed long before the robot has reached the doorway. However, the operator, already viewing video feed from the next room, may believe that the robot is already through the door. To prevent a collision with the doorframe, the robot may stop and refuse to move forward. Although the robot communicates that it is blocked in front, the user may be confused by the lack of obstacles in the visual feed.

Put simply, the default interface used in Experiment 1 did not provide the operator with a window into the "mind" of the robot or an accurate concept of the spatial environment. Experiment 2 was designed to explore the use of a new interface component intended to do just that.

## V. EXPERIMENT 2

Observations from Experiment 1 suggested that video may not provide an adequate means for the operator to predict robot behavior or understand its intentions. On the other hand, humans are visually centric and prefer pictures and diagrams when attempting to understand or communicate [33]. In order to address the HRI limitations observed in Experiment 1, some means was required to support collaborative understanding and yet take full advantage of the functional utility associated with visual representation. In addition to these human factors, there were also significant engineering reasons for assessing alternatives to video presentation of the remote environment. Video demands high-bandwidth and continuous communication and is therefore ill-suited for many of the very environments where robots are most needed. Except for short ranges, transmission of high-bandwidth video is only possible when line of sight can be maintained either with a satellite or another radio antenna [34], [35]. For instance, high-bandwidth video cannot be transmitted through layers of concrete and rebar, making it inappropriate for urban terrain. Likewise, forest and jungle canopy precludes reliable transmission of video.

In response to these human and engineering factors, work began to develop a new interface component that could provide a collaborative representation of the environment and help the operator to understand the perspective and intentions of the robot. A virtual 3-D display (see Fig. 3) was developed as a means to give users insight into the reason for robot initiative and diminish the possibility of disorientation. The second experiment was designed to assess the effectiveness of the virtual 3-D display. In particular, the goal was to contrast the 3-D map representation used within the new interface with the presentation of streaming video from the old interface. Unlike the first experiment that examined the difference between Shared and Safe Modes, the second experiment used only Safe Mode in order to ensure that control mode did not complicate the analysis of performance with the different information displays.

### A. Participants

The experiment was performed over a seven-day period within the St. Louis Science Center and utilized 64 visitors who volunteered to take part in the experiment. The majority of participants were high school students from schools in the St. Louis area. These students were not preselected, but rather volunteered to take part in the study while visiting the Science Center. Age was not recorded due to the fact that most participants were of a similar age.

As before, the experiment was set up as a remote deployment such that the operator control station was located several stories above the robot arena so that the operator could not see the robot or the operational environment. The arena was built by the production staff of the Science Center and utilized artificial rocks, artificial trees, mannequins, and plywood dividers to create a maze environment (see Fig. 4). Due to the distance and physical occlusions separating the control station from the actual robot environment, analog video was not possible. Instead, state-of-the-art video compression was used to digitize the analog video into an MJPEG format and wirelessly transmit from the robot to a nearby access point connected to the building's network. The building's wired network was then used to transfer the video data two stories up to the operator. Exploiting the wired infrastructure in place throughout the building made it possible to provide continuous, reliable video at a high frame rate. The presentation speed and resolution of this video exceeded that possible through an entirely wireless data link. This configuration ensured that the comparison between video and the 3-D map display was not merely a function of current communication bandwidth constraints, but rather an investigation of the fundamental differences between an interface based primarily on viewing raw video and one which presented the environment perceived by the robot.

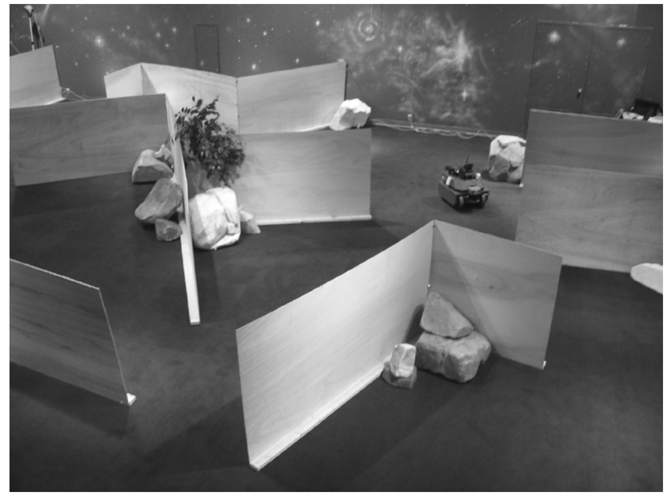

Fig. 4. Partial view of the arena built at the St. Louis Science Center.

### B. Procedure

Each participant was given basic instructions on how to use the interface, and, as with the previous experiment, no participant was permitted to drive the robot until the start of the trial run. Participants were assigned to alternating conditions so as to ensure equal numbers of participants in each condition. No participant was allowed to operate the robot more in more than one trial. Each trial run was exactly 3 min. This time limit helped to insure that the measured performance was a function of the interface rather than a function of operator interest or time spent on task.

At the beginning of each run, the robot programs were restarted so that the map built by the previous participant was erased. Each participant was told to direct the robot around the environment in order to build as large a map as possible. This task was selected to assess differences between the presentation modes because it involved spatial reasoning. The task required operators to perceive the frontiers of the map and direct the robot to explore as many new areas as possible in a limited time. All participants were given access to the same 2-D map component (see Fig. 5) within which the robot presents the

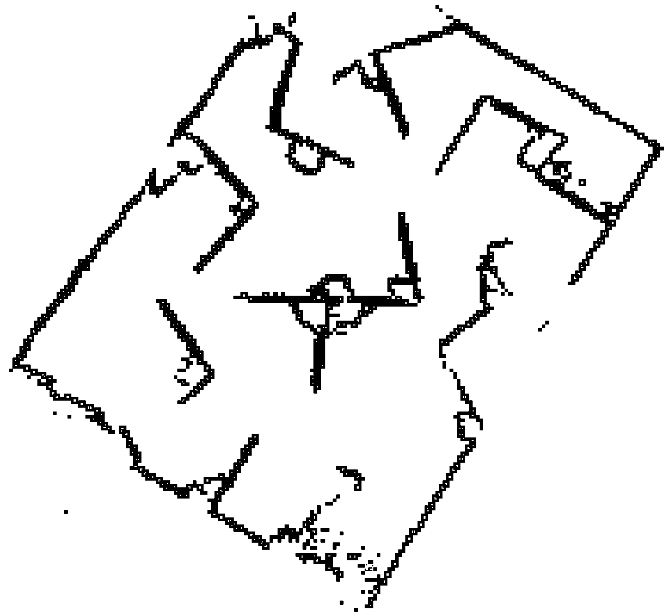

Fig. 5. Near-complete map built by one of the participants.

map that it builds as it explores new territory. Exactly half of the participants used the same interface as in Experiment 1 (see Fig. 2). These participants were able to use both the 2-D map and the video module. For the other half, the same interface was used except that the new, virtual 3-D interface module entirely occluded the video module.

During each trial, the interface stored a variety of useful information about the participant's interactions with the interface. Joystick bandwidth was recorded as the number of messages sent from the joystick indicating a change of more than 10% in the position of the stick. This information was used as an indirect measure of workload [36], [37]. The interface also recorded the number of joystick vibrations caused by human navigational error. For each trial, the map produced by the robot was saved. In order to assess performance, a software algorithm was implemented to calculate the percentage of the full map that was present in each of these saved maps. This approach provided a reasonable assessment of the operator's ability to explore the environment in the time available. Immediately after completing a trial, each participant was asked to rank on a scale of 1 to 10 how "in control" they felt during the operation, where 1 signified "The robot did nothing that I wanted it to do" and 10 signified, "The robot did everything I wanted it to do."

### C. Results

In the 3 min provided, 80% of participants explored over half of the total environment. One person, a 3-D display participant, was able to build the entire map in the allotted 3 min. As described above, task performance was calculated by comparing the map generated during the exploration task with the complete map of the task environment. This comparison showed no significant statistical difference between the use of the video interface module $(M = 0.71)$ and the virtual 3-D map module $(M = 0.61)$, $F(1,31) = 0.558$, $p = 0.070$. Using joystick bandwidth as an indication of human workload (see Fig. 6) and joystick vibration as a metric for human navigational error (see Fig. 7), analysis shows that operators using the virtual 3-D display worked less and demonstrated fewer instances of human navigational error. On average, the joystick bandwidth for participants using the virtual 3-D display was 1057.50 messages from the interface to the robot, compared to 1229.07 for operators using video feed, $F(1,31) = 2.024$, $p < 0.05$. Human navigational error for participants using the virtual 3-D display averaged 11.00 instances, compared to an average of 14.29 for the video participants, $F(1,31) = 0.399$, $p < 0.05$.

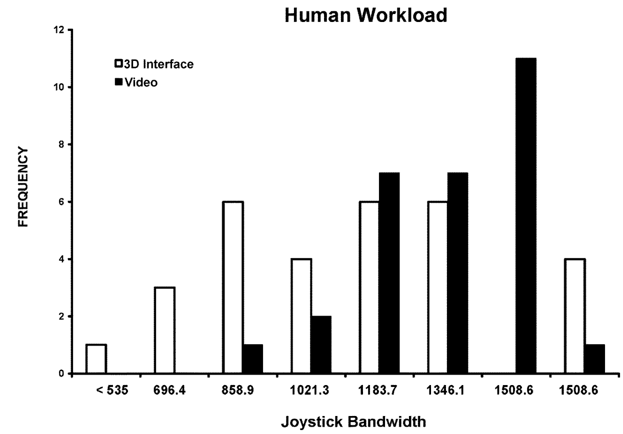


Fig. 6. Human workload as measured by joystick bandwidth.

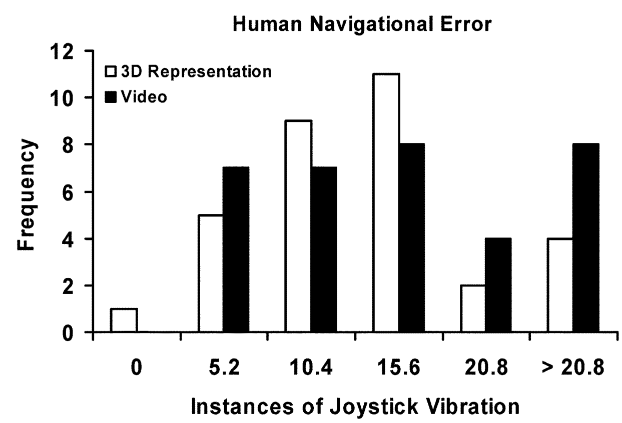


Fig. 7. Human navigational error as measured by joystick vibration.

In addition to reduced workload and fewer navigational errors, use of the virtual 3-D display slightly increased the operator's subjective "feeling of control" while operating the robot. The average feeling of control for the 3-D display was 7.219 compared with an average of 7.059 for the video, $F(1.31) = 0.497$, $p < 0.05$.

### D. Discussion

The second experiment provided initial validation for the effectiveness of the 3-D map representation. Results show that use of the virtual 3-D display resulted in no significant performance decrement and provided reduced workload, fewer navigational errors, and a heightened sense of control. A motivation for the development of the virtual 3-D display had been to promote a shared understanding of the task and environment. To assess the effectiveness of the virtual 3-D display in this regard, it is useful to consider that a decrease in joystick vibrations not only represents a reduction in operator navigational error, but also a reduction in the instances where the operator failed to understand the reason for robot initiative. Recall that the joystick vibrates only if the operator commands movement in a direction in which the robot has already taken initiative to prevent a collision. These results indicate progress toward the goal of providing a representation that supports situation awareness and a shared understanding of the environment. More broadly, these results provide evidence that it may be possible to support situation awareness needs of human operators without using video. This finding provides an important counterpoint to opinion within the field of HRI that reliable, continuous video is essential for remote navigation [38].

One limitation of the study is that the data do not provide a means to discuss the usability of the 3-D display based on age or gender. Although these effects were not deemed central to the invented hypothesis, it is recognized that future usability testing, focused on age and gender effects would be necessary to generalize about the overall usability of the 3-D display.

From an engineering perspective, this experiment shows that it is possible to build a map on-the-fly and communicate it back to a remote user fast enough to support real-time robot navigation. The significance of this result to the area of remote systems can be seen most clearly when one considers the reduction in communication bandwidth made possible by using the 3-D map display. Whereas the video alone required 3 000 000 bits per second (bps), the total interface bandwidth while using the virtual 3-D interface was only 64 000 bps. This bandwidth savings allows control to extend into new domains using data transmission methods that can be used in underground bunkers, caves, nuclear reactors, and urban search and rescue sites where it is often impossible to maintain a video feed.

However, the fact that the human–robot team can function effectively without video is no reason to disregard the potential benefits of video in those instances when video is available. Experience with operators and subject area experts from Energy, Defense and Emergency Management contexts indicate that operators expect and can exploit video in remarkable ways [3], [5], [14]. Many applications require the human to play a role in visual search and detection. Although the second experiment suggested that video could be replaced with the 3-D representation, the optimal interface will likely provide a dynamic balance between the video and virtual displays.

## VI. Experiment 3

The third experiment was designed to revisit the comparison between modes of control undertaken in Experiment 1 and explore possible interactions with the use of the virtual 3-D map display discussed in Experiment 2. A previous usability study by Marble *et al.* had shown that the shared mode offered the greatest potential for operator confusion and frustration [3]. Consequently, it was hypothesized that the shared mode might provide the greatest potential for the virtual 3-D display to reduce human navigational error and workload. Therefore, this experiment was designed to explore how the new 3-D display would affect the comparison between control modes.

Experiment 3 was also intended to show that the benefits of sharing control between the human and operator demonstrated in the first experiment were not merely due to the high cognitive workload placed on the operator. The typical assumption found in the literature is that robot autonomy trails behind human performance, but may be useful as operator workload increases or communications fail [17], [39], [40]. It was hoped that Experiment 3 could provide conclusive evidence that robot initiative can significantly improve performance even when data link connectivity is maintained and human workload is minimal. In order to minimize individual human workload, the control task was separated into specific operator functions: navigation, which depends on an exocentric display; driving, which uses an egocentric display; and operation of an application payload, which can be controlled independently from the robot. In addition to minimizing individual human workload, an added benefit of assigning different roles was that it afforded an opportunity to observe the exchange of information between team members in different roles. As Scholtz points out, the roles of human operators do not remain static and interfaces should be able to adapt accordingly.

### A. Participants

This experiment included 120 volunteers grouped into teams of six members. The participating teams consisted of one team of teachers, three teams of eighth grade students, and the remainder of the teams being drawn from local high schools. Unlike the volunteers who comprised the participant pool for the last experiment, Experiment 3 participants signed up in advance to take part in the study. Participants were recruited from a solicitation of local schools through the museum's outreach program. Participants knew and selected the other people in their team prior to participation in the experiment. Age and gender were not recorded due to the fact that most participants were of similar age and the fact that gender was mixed for each team.

### B. Procedure

The experiment was run over seven days at the St. Louis Science Center. Teams of participants were assigned to alternating conditions so as to ensure equal numbers of teams in each condition. No participant was allowed to operate the robot in more than one trial. As in the previous experiment, the robot was located in the lower level of the Science Center, while the control center was located on the top level. Experiment 3 used the same environment as was used in Experiment 2 with the same lighting and placement of obstacles. Three mannequins were placed in locations designed to force participants to coordinate in order to discover aspects regarding each particular mannequin's location. The mannequins remained in place throughout the entire experiment. The starting point of the robot alternated between two different locations such that an equal number of shared mode and safe mode runs were begun from each starting point. For this experiment, the control interface components were divided across three separate stations, each with its own monitor and input devices. No interface component was visible at more than one control station. Two participants manned each station resulting in a total of six people dedicated to robotic system control. The stations were arranged in an arc such that the participants at each station could communicate easily with the others, but could not see the other displays.

The first control station was dedicated to the application payload, which in this case was a pan, tilt, and zoom camera. Using a joystick that allowed operation of the various camera controls,

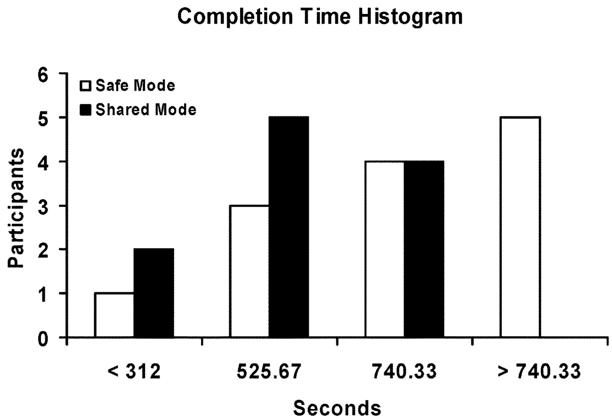


Fig. 8. Time to complete the task in safe and shared modes.

the application payload participants used the visual feedback from the robot to seek out the three mannequins and to provide navigational advice. The second control station was dedicated to driving the robot. Participants were permitted to see the virtual 3-D window, the local environment window, the sensor status window, and the robot state window (see Fig. 2). Primarily, the operators at the driving station used the 3-D virtual display, but were constrained to an egocentric perspective which precluded a global view of the environment. The final station was the navigation station. The navigators had access to the 2-D map being built as the robot traveled through its environment. This gave them a bird's eye view of the environment and the robot's position in it. In addition, the participants at the navigation station were given a hardcopy of an *a priori* map that showed the locations of the three mannequins. Having two participants at each station was not necessary, but ensured that workload was minimal. Task completion required the three groups to self-organize in order to arrive at and gain a visual lock on all three of the mannequins as quickly as possible. As in the previous experiment, joystick usage was measured as an indication of operator workload and instances of joystick vibration were recorded as an indication of operator error and confusion.

### C. Results

To assess whether there were differences in performance between the teams or modes due to start position, a $t$-test was performed on the completion time data by start position, $t(19) = 0.72$, $p > 0.05$, and on joystick bandwidth, $t(19) = 0.5$, $p > 0.05$. Because no differences were found based on start position, the data was merged across this analysis. On average, less time was required for the participants using the shared mode (see Fig. 8). The mean completion time for shared mode participants was 466.8 s compared to a mean completion time of 641.1 s for the safe mode participants, $F(1,9) = 3.64$, $p < 0.05$. Safe mode participants demonstrated a greater workload, as assessed by joystick movements, than that of their shared mode counterparts, $M = 2743.8$ and $M = 1725.6$, respectively, $F(1,9) = 0.296$, $p < 0.05$. Using joystick vibration as a metric for human navigational error shows that safe mode participants made 25.1 errors compared to 16.8 errors for the shared mode participants, $F(1,9) = 5.92$, $p < 0.05$. An analysis of performance in the safe mode indicates that joystick vibration and joystick bandwidth were correlated to task duration $r(7) = 0.761$, $p < 0.05$ and $r(7) = 0.729$, $p < 0.05$, respectively. Similar analysis of the shared mode performance indicated only a correlation between joystick vibration and task duration, $r(9) = 0.659$, $p < 0.05$ but no correlation between task duration and joystick bandwidth $r(9) = 0.311$, $p > 0.05$.

### D. Discussion

As with the first experiment, shared mode participants experienced increased performance efficiency when compared to their safe mode counterparts. The results of Experiment 3 strengthen the case for improved human–robot team performance when the robot is able to take initiative to respond to the environment and assist the operator. The results from Experiment 3 imply that reducing the workload placed on the human driver and increasing the importance of strategy and intelligence does not diminish the performance benefits of sharing control between human and robot team members. Previous research has shown that effective teams utilize a shared mental model of the task and current situation [25], [42]. Likewise, the findings from Experiment 3 suggest that in order to fully realize the benefits of sharing control between human and robot team members, it is advantageous to provide a shared understanding of the environment. The results also suggest that a virtual 3-D map may be as useful to the operator as real-time video. Unlike most interfaces for remotely controlling a mobile robot, the virtual 3-D display presents the user with the same information used by the robot to make decisions, which may make it easier for the human to predict robot behavior and understand occasions of robot initiative.

In many operational scenarios, it is not only possible but probable that the roles of driving, navigating, and operating the application payload will be spread amongst multiple human operators. Several researchers have pointed out the high cognitive burden associated with remote deployment of mobile robots and have argued that effective control requires multiple human operators [13]–[15]. Although detailed analysis of these different roles (i.e., driver, navigator, payload operator) was beyond the scope of this experiment, anecdotal observations (recorded during the experiment and also during debriefing conversations with operators after the experiments) suggest interesting areas for further investigation. One observation was that just as performance can be degraded by a fight for control between the driver and robot, there is the potential for similar conflicts between human operators. In fact, in more than one instance, the operators responsible for driving the robot said they came to trust robot suggestions and robot initiative over that of their human team members. For example, when told to turn by other human team members, the driver would sometimes refuse, stating that it was better to let the robot choose the path. On the other hand, human drivers sometimes chose to override robot initiative, just as they sometimes chose to ignore advice from the navigators or payload operators. Further experimentation will be necessary to characterize the reasons for these choices and quantify their effect on team performance. One explanation found in the literature is that team success depends on the ability of each team member to understand the perspective of other members [25]. If this is true, the most effective human–robot teams will be those that utilize a collaborative model of the environment and task. Such research questions provide a fertile ground for further experimentation into the challenges of sharing control, not only between human and robot, but also between humans.

## VII. Conclusion and Further Work

In [43], Woods *et al.* argue that far from simplifying the control problem, the introduction of autonomy and robot initiative places new demands on the operator and often increases the need for a sophisticated understanding of the robot and interface. To address these challenges, the experiments discussed in this study offer a form of collaborative control and representation that may reduce instances of human navigational error, decrease human workload, improve the operator's "feeling of control," and increase overall performance. The purpose of this study has not been simply to argue that autonomous capabilities and better interfaces can improve remote navigation, but rather to provide evidence for a form of collaborative control where robots are regarded as peers and effectively used as trusted team members.

Collectively, the experiments suggest the value of a collaborative representation that allows the human and robot to predict behavior and communicate intent. Philosopher of the mind, D. Dennett introduced the concepts of physical stance, design stance, and intentional stance when referring to intelligent systems. Most operators can interpret how joystick movements will be translated into robot movements (i.e., design stance) and can predict whether a robot can ascend stairs by seeing its physical construction (i.e., physical stance). In fact, an understanding of physical and design considerations may be sufficient to maintain control under a human supervisory control paradigm. However, once the robot becomes viewed as a team member, something more is needed. For intelligent systems, Dennett believes the human must be able to interpret and predict robot behavior based on an intentional stance—an understanding of the robot as a rational agent whose behavior is governed by intentional states [44]. For example, in terms of the present research, an intentional stance is necessary to predict the robot's path as it attempts to autonomously complete a search and detection task. According to Dennett's reasoning, allowing the operator to see the environment through the robot's eyes facilitates intentional stance predictions. The research presented has attempted to provide this common perspective through the collaborative representation presented in the virtual 3-D display. The resulting methods of collaborative representation and control provide a means to move toward a notion of true teaming where responsibilities and roles shift dynamically to permit mutual support. Together the experiments indicate that sharing control between humans and robots has the potential to provide a compelling control alternative across a broad range of tasks and applications including many of those discussed in the DARPA/NSF Human–Robot Interaction roadmap [5].

## References


[1] J. L. Burke, R. R. Murphy, E. Rogers, V. J. Lumelsky, and J. Scholtz, "Final report for the DARPA/NSF interdisciplinary study on human–robot interaction," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 34, no. 2, pp. 103–112, May 2004.

[2] H. A. Yanco, J. L. Drury, and J. Scholtz, "Beyond usability evaluation: Analysis of human–robot interaction at a major robotics competition," *J. Human–Comput. Interaction*, vol. 19, pp. 117–149, 2004.

[3] J. L. Marble, D. J. Bruemmer, and D. A. Few, "Lessons learned from usability tests with a collaborative cognitive workspace for human–robot teams," in *Proc. IEEE Conf. Syst., Man, Cybern.*, Waikoloa, HI, 2003.

[4] J. L. Marble, D. J. Bruemmer, D. A. Few, and D. D. Dudenhoeffer, "Evaluation of supervisory vs. peer–peer interaction for human–robot teams," in *Proc. 37th Annu. Hawaii Int. Conf. Syst. Sci.*, Waikoloa, HI, 2004.

[5] H. A. Yanco and J. Drury, "'Where am I?' Acquiring situation awareness using a remote robot platform," in *Proc. IEEE Conf. Syst., Man Cybern.*, 2004, pp. 2835–2840.

[6] R. Simmons, "Concurrent planning and execution for autonomous robots," *IEEE Control Syst. Mag.*, vol. 12, no. 1, pp. 46–50, Feb. 1992.

[7] A. D. Mali, "On the behavior-based architectures of autonomous agency," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 32, no. 3, pp. 231–242, Aug. 2002.

[8] R. Brooks, "Planning is just a way of avoiding what to do next," Mass. Inst. Technol., Cambridge, Working Paper 303, 1987.

[9] R. A. Brooks, "A robust layered control system for a mobile robot," *IEEE Trans. Robot. Autom.*, vol. 2, no. 1, pp. 14–23, Mar. 1986.

[10] R. C. Arkin, *Behavior-Based Robotics*. Cambridge, MA: MIT Press, 1997.

[11] D. J. Bruemmer, J. L. Marble, and D. D. Dudenhoeffer, "Mutual initiative in human–machine teams," in *Proc. IEEE Conf. Human Factors Power Plants*, 2002, pp. 7.22–7.30.

[12] R. Simmons, R. Goodwin, K. Z. Haigh, S. Koenig, and J. O'Sullivan, "A layered architecture for office delivery robots," in *Proc. Int. Conf. Auton. Agents*, 1997, pp. 245–252.

[13] R. R. Murphy, "Human–robot interaction in rescue robotics," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 34, no. 2, pp. 138–153, May 2004.

[14] J. Casper and R. R. Murphy, "Human–robot interactions during the robot-assisted urban search and rescue response at the world trade center," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 33, no. 3, pp. 367–385, Jun. 2003.

[15] J. L. Burke, R. R. Murphy, M. D. Coovert, and D. L. Riddle, "Moonlight in Miami: A field study of human–robot interaction in the context of an urban search and rescue disaster response training exercise," *Human–Comput. Interaction*, vol. 19, pp. 85–116, 2004.

[16] B. Trouvain, H. L. Wolf, and F. E. Schneider, "Impact of autonomy in multi-robot systems on teleoperation performance," in *Proc. 2003 Workshop Multi-Robot Syst.*, 2004, pp. 253–264.

[17] M. A. Goodrich, D. R. O. Jr, J. W. Crandall, and T. J. Plamer, "Experiments in adjustable autonomy," in *Proc. IJCAI Workshop Auton., Delegation, Control: Interacting Auton. Agents*, Seattle, WA, 2001, pp. 1624–1629.

[18] R. Parasuraman, T. B. Sheridan, and C. Wickens, "A model for types and levels of human interaction with automation," *IEEE Trans. Syst., Man, Cybern, A, Syst. Humans*, vol. 30, no. 3, pp. 286–297, May 2000.

[19] V. Riley, "A general model of mixed-initiative human–machine systems," in *Proc. 33rd Annu. Human Factors Soc. Conf.*, 1989, pp. 124–128.

[20] T. B. Sheridan and W. L. Verplank, "Human and computer control of undersea teleoperators," MIT Man–Mach. Syst. Lab., Cambridge, MA, 1978.

[21] T. Fong, C. Thorpe, and C. Baur, "A safeguarded teleoperation controller," in *Proc. IEEE Int. Conf. Adv. Robot.*, Budapest, Hungary, 2001.

[22] E. Krotkov, R. Simmons, F. Cozman, and S. Koenig, "Safeguarded teleoperation for lunar rovers: From human factors to field trials," in *Proc. IEEE Planetary Rover Technol. Syst. Workshop*, 1996, pp. 1587–1608.

[23] D. Kortenkamp, R. P. Bonasso, D. Ryan, and D. Schreckenghost, "Traded control with autonomous robots as mixed initiative interaction," in *Proc. AAAI Spring Symp. Mixed Initiative Interaction*, Stanford, CA, 1997.

[24] K. Myers and D. Morley, "Human directability of agents," in *Proc. K-CAP*, 2001.

[25] J. Yen, J. Yin, T. Ioerger, M. Miller, D. Xu, and R. Volz, "CAST: Collaborative agents for simulating teamwork," in *Proc. 17th Int. Joint Conf. Artif. Intell.*, 2001, pp. 135–142.

[26] DARPA Grand Challenge. [Online]. Available: http://www.darpa.mil/grandchallenge

[27] T. Fong, C. Thorpe, and C. Baur, "Robot, asker of questions," *Robot. Auton. Syst.*, vol. 42, no. 3–4, pp. 235–243, 2003.

[28] ——, Collaboration, dialogue, and human–robot interaction. presented at Proc. 10th Int. Symp. Robot. Res. [Online]. Available: http://www.ri.cmu.edu/pub_files/pub3/fong_terrence_w_2001_1/fong_terrence_w_2001_1.pdf

[29] E. B. Pacis, H. R. Everett, N. Farrington, and D. J. Bruemmer. Enhancing functionality and autonomy in man-portable robots. presented at Proc. SPIE Defense Secur. Symp. [Online]. Available: http://www.spawar.navy.mil/robots/pubs/spie5422-49.pdf

[30] D. J. Bruemmer, D. A. Few, R. Boring, M. Walton, J. L. Marble, C. Nielsen, and J. Garner, "'Turn off the television!:' Robotic exploration experiments with a 3-D abstracted map interface," in *Proc. 38th Hawaii Int. Conf. Syst. Sci.*, Waikoloa, HI, 2005.

[31] K. Konolige, "Large-scale map-making," in *Proc. AAAI*, San Jose, CA, 2004, pp. 457–463.

[32] J. S. Gutman and K. Konolige, "Incremental mapping of large cyclic environments," in *Proc. CIRCA*, Monterey, CA, 1999, pp. 318–325.

[33] H. Pashler, "Coordinate frame for symetry detection and object recognition," *J. Exper. Psych.: Human Perception Perform.*, vol. 16, pp. 150–163.

[34] BreezNet 802.11 Data Sheet. [Online]. Available: http://www.alvarionusa.com/RunTime/Materials/PDFFiles/version_4.204_IEEE.PDF

[35] Freewave FGR Series Data Sheet. [Online]. Available: http://www.freewave.com/fgr900oem.html

[36] G. J. Khoury and G. V. Kondraske, "Measurement and Continuous Monitoring of Human Workload Associated With Manual Control Devices," NASA JSC, Houston, TX, Tech. Rep. 91-011R, 1991.

[37] D. Clarke, S. Yen, G. V. Kondraske, G. J. Khoury, and K. J. Maxwell, "Telerobotic Network Workstation for System Performance and Operator Workload Monitoring," NASA JSC, Houston, TX, Tech. Rep. 91-013R, 1991.

[38] M. Baker, R. Casey, B. Keyes, and H. A. Yanco, "Improved interfaces for human–robot interaction in urban search and rescue," in *Proc. IEEE Conf. Syst., Man, Cybern.*, 2004, pp. 2960–2965.

[39] C. Nielsen, M. Goodrich, and J. Crandall, "Experiments in human–robot teams," in *Proc. Int. Workshop Multi-Robot Syst.*, 2003, pp. 241–252.

[40] B. Trouvain, H. L. Wolf, and F. E. Schneider, "Impact of autonomy in multi-robot systems on teleoperation performance," in *Proc. Workshop Multi-Robot Syst.*, 2003, pp. 253–264.

[41] J. Scholtz, "Human–robot interactions: Creating synergistic cyber forces," in *Proc. Workshop Multi-Robot Syst.*, 2002, pp. 177–184.

[42] N. J. Cooke, E. Salas, J. A. Cannon-Bowers, and R. Stout, "Measuring team knowledge," *Human Factors*, vol. 42, pp. 151–173, 2000.

[43] D. Woods, J. Tittle, M. Feil, and A. Roesler, "Envisioning human–robot coordination in future operations," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 34, no. 2, pp. 210–218, May.

[44] D. C. Dennett, *True Believers: The Intentional Strategy and Why It Works*, A. F. Heath, Ed. Oxford, U.K.: Oxford Univ. Press, 1981.

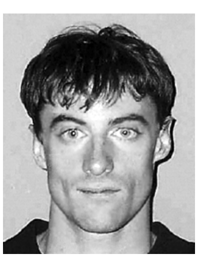

**David J. Bruemmer** received the B.A. degrees in computer science and religion from Swarthmore College, Swarthmore, PA, in 1998.

He is currently a Principal Research Scientist at the Idaho National Laboratory (INL), Idaho Falls, where he serves as Technical Director for Unmanned Ground Vehicle Systems. As a Consultant to the Defense Advanced Research Projects Agency, he worked to coordinate development of autonomous robotics technologies through several offices and programs. Since arriving at the INL, he has led development of a robot intelligence architecture that has been ported to a variety of robots for applications including remote characterization, mine sweeping operations, military reconnaissance, and search and rescue operations. His interests include autonomous robot behavior, swarm robotics, mobile manipulation, and teaming between unmanned air and ground vehicles.

Mr. Bruemmer is Program Chair for the 2006 Conference on Human–Robot Interaction and was the Guest Editor of *Intelligent Systems Magazine*.

**Douglas A. Few** received the B.S. degree in computer science from Keene State College, Keene, NH, in 2003.

He is currently a Principal Research Scientist with the Robotic and Human Systems Group, Idaho National Laboratory (INL), Idaho Falls. Prior to joining INL, he supported robotic research efforts around the world in as a Software Support Engineer for iRobot Corporation's Research Robot Division. Additionally, he served as a Project Manager for CHI Engineering Services, Portsmouth, NH, a small engineering firm catering to the needs of the natural gas industry. His interests include human–robot interactions, mixed autonomy robotic systems, generic robot control architectures, and system reliability algorithms.

**Ronald L. Boring** received the B.A. degree in psychology and German from the University of Montana, Missoula, the M.A. degree in human factors and experimental psychology from New Mexico State University, Las Cruces, and the Ph.D. degree in cognitive science from Carleton University, Ottawa, ON, Canada.

He is an Adjunct Professor of psychology at Brigham Young University, Provo, UT, and is also with the Idaho National Laboratory (INL), Idaho Falls. Prior to joining the INL, he was a Usability Engineer for the Microsoft Corporation and the Expedia Corporation, and also was a Guest Researcher in human–computer interaction (HRI) at the National Research Council of Canada. He has published in a wide variety of human factors and human–computer interaction forums. In addition to HRI, his primary research emphasis since joining the INL has been in human reliability analysis on projects for NASA, the US Nuclear Regulatory Commission, and the Department of Homeland Security.

Prof. Boring has served on the conference committees for the Association for Computing Machinery's annual conferences on Computer–Human Interaction and for the annual meetings of the Human Factors and Ergonomics Society.

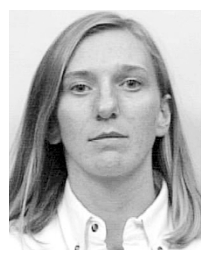

**Julie L. Marble** received the Ph.D. degree in cognitive psychology and human factors from Purdue University, West Lafayette, IN, in 2001.

She joined the Idaho National Laboratory (INL), Idaho Falls, as a Principal Scientist in February 2001. In addition to her work at the INL, she is an Adjunct Professor with the University of Idaho, Moscow, where she teaches an undergraduate and graduate course, Cognitive Engineering of Interactive Software Design, which covers information visualization techniques and extraction of user requirements for software systems. Her research interests include information presentation, display design, human–robot interaction, and human reliability analysis. She has performed research in the design of displays for autonomous and teleoperated robots under normal and failure conditions, and display interaction with controls and control systems.

**Miles C. Walton** received the B.S. degree in computer science from Brigham Young University, Provo, UT, in 1980.

He is a Software Engineer in the Robotics and Human Systems Department, Idaho National Laboratory, Idaho Falls. He is currently working with autonomous ground vehicles (AGVs) and operator control systems. His latest work on AGVs is in the development of operator interfaces that allow an effective mix of operator-directed controls and autonomous vehicle action. He is also working on a team developing a system for automated radiation surveys of large outdoor areas using AGVs.

**Curtis W. Nielsen** received the B.S. and M.S. in computer science in 1999 and 2003, respectively, from Brigham Young University, Provo, UT, where he is currently pursuing the Ph.D. degree in computer science.

His research interests include computer graphics, robotics, human–robot interaction, and search and rescue.