

Using a Robot Proxy to Create Common Ground in Exploration Tasks

Kristen Stubbs, David Wettergreen, and Illah Nourbakhsh
Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue Pittsburgh, PA 15213
{kstubbs, dsw, illah}@cmu.edu

ABSTRACT

In this paper, we present a user study of a new collaborative communication method between a user and remotely-located robot performing an exploration task. In the studied scenario, our user possesses scientific expertise but not necessarily detailed knowledge of the robot’s capabilities, resulting in very little common ground between the user and robot. Because the robot is not available during mission planning, we introduce a robot proxy to build common ground with the user. Our robot proxy has the ability to provide feedback to the user about the user’s plans before the plans are executed. Our study demonstrated that the use of the robot proxy resulted in improved performance and efficiency on an exploration task, more accurate mental models of the robot’s capabilities, a stronger perception of effectiveness at the task, and stronger feelings of collaboration with the robotic system.

Categories and Subject Descriptors

H.1.2 [Models and Principles]: User/Machine Systems—*Human factors*; H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Evaluation/methodology*; I.2.9 [Computing Methodologies]: Artificial Intelligence—*Robotics*

General Terms

Experimentation, Human Factors

Keywords

human-robot interaction, exploration robotics, common ground, robot proxy

1. INTRODUCTION

In this work, we focus on improving human-robot interaction within the domain of exploration robotics. We define robotic exploration tasks broadly as those in which a robot

co-investigates an unknown environment with a remote human partner. Exploration is an important domain of study because of its applicability to a wide variety of problems, which range from searching for signs of life on other planets to investigating debris after a building collapse (e.g. [18, 2]). In particular, we are interested in exploration which involves the deployment of autonomous robots that work in complex, real-world settings. In these situations, our users are not likely to be experts in robotics, and they may possess inaccurate mental models of robotic technologies. At the same time, these users often possess sophisticated domain knowledge which the robot does not. In order to facilitate successful exploration, we are interested in promoting shared understanding between users and robots. That is, we wish to increase users’ understanding of robots and foster accurate mental models, and, at the same time, enhance robots’ understandings of users and their goals in order to drive robots’ decision-making processes.

1.1 Collaboration Models for Exploration Robotics

Current collaboration model. Current exploration robotics systems follow a communication model similar to the encoder-decoder model of information processing [14]. The user possesses goals which he/she would like to accomplish and uses an interface to encode those goals into machine-readable actions. These actions are sent to the robot, which uses a planner to decode and schedule the necessary low-level commands and an executive process to direct the execution of the commands. After execution has completed, the robot returns the resulting data to the user. Particularly in the case of domains featuring remote, asynchronous communication, such as planetary robotics, the interaction is essentially “open-loop”: the user cannot communicate with the robot while developing plans, and the robot cannot communicate with the user during execution.

Common ground collaboration model. As an alternative to the current collaboration model, the focus of the proposed research is to build common ground between users and robots explicitly. As defined by Herbert Clark and colleagues, common ground between two participants in a joint activity is “the knowledge, beliefs, and suppositions they believe they share about the activity” [5, p. 38]. Common ground is required in order for individuals to communicate and collaborate successfully [4, 3]. The process by which common ground is established between individuals is referred to as *grounding*. While we have witnessed the application of common ground theory to collocated, social

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HRI’08, March 12–15, 2008, Amsterdam, The Netherlands.
Copyright 2008 ACM 978-1-60558-017-3/08/03 ...\$5.00.

human-robot interaction (Section 1.2), here we examine how to promote common ground during a remote collaboration between a human and robot in which both agents possess a large amount of domain-specific knowledge which does not necessarily intersect. In this paper, we present the results of a user study of a system based on the use of a robot proxy to improve common ground between a user and robot collaborating on an exploration task.

1.2 Common Ground and HRI

Although the common ground framework was developed to understand conversation and collaboration among people, not between people and machines, recent work has extended the framework into the field of human-computer interaction [1, 16]. This research suggests that interfaces can be improved by thinking about the user’s experience as a conversation in which shared meaning between the user and the interface must be developed.

In the field of HRI, Jones and Hinds [12] observed SWAT teams and used their findings to inform the design of robot control architectures to coordinate multiple robots. Although their observations did not include robots, their findings emphasize the importance of common ground between a robot and its user, especially when the two are not collocated. Kiesler and colleagues [13] describe experiments reporting more effective communication between people and robots when common ground is greater. When a robot adapts its dialog to fit the knowledge of the user, more effective information exchange results [20]. Common ground theory has been applied as a means to drive conversational interactions between humans and collocated, social robots [17, 15]; our work focuses on using common ground theory to improve interactions with a remotely-located partner that does not engage in natural language conversation.

Situation Awareness. Although generally focused more on dialog and communication, the common ground framework overlaps with work on situation awareness (SA). Endsley [9] defines SA as “knowing what is going on around you.” SA been examined in the HRI domain, particularly with urban search and rescue (USAR) robots [7, 2, 8, 21]. Empirical work indicates that USAR operators spend significantly more time trying to gain SA—assessing the state of the robot and the environment—than they do navigating the robot [2, 8]. This work tends to focus on “real time” interaction (with teleoperated robots), so its applicability is less clear for HRI with autonomous robots that are remotely and asynchronously commanded.

Prior Work. Our common ground-based approach to exploration robotics has been validated by two years of observations of a remote exploration task, the NASA-funded Life in the Atacama (LITA) project [18, 19]. We focused on remote science operations: a science team composed of biologists, geologists, and instrument specialists (located in Pittsburgh) used a robot called Zoë to explore the desert and search for signs of life. The science team produced day-long command sequences for the robot, and they received and analyzed data products generated the previous day. An engineering team, composed of roboticists and instrument specialists (located in Chile), monitored the robot, conducted troubleshooting on-site, and ensured that the science team was able to gather data successfully.

Our application of common ground theory to exploration robotics results from a detailed analysis of the problems




Figure 1: Our robot proxy-based exploration robotics mission planning process.

which we observed during the LITA mission. These problems included errors and miscommunications that often stemmed from a lack of mutual knowledge between the science team and the robot, such as information about the robot’s capabilities and the science team’s underlying goals [18]. Common ground theory offers us a model for human-robot interaction based on substantial empirical evidence of how people collaborate with each other (e.g., [3, 6]).

1.3 The Robot Proxy

In contrast with previous work utilizing common ground theory to improve social, one-on-one interactions, we consider how to design the interaction between a user and robot who are not physically co-located and who do not communicate using natural language dialog. Because of the high cost of executing plans on a remotely-located robot (as robot runtime is very expensive), in this paper we introduce the use of a robot proxy, which allows the user to participate in the grounding process during plan creation **before the plan is sent to the robot** (Figure 1).

As the user inputs a set of actions for the robot, the robot proxy monitors these actions, simulates their execution, and asks targeted questions designed to improve common ground with respect to these actions. For example, if the robot proxy is simulating the execution of two sequential actions in the plan and detects that, as a result, two instruments will target close but non-intersecting areas, the robot proxy would ask the user whether the instruments were intended to target the same area and help adjust the plan if need be. Conventional, conversation-based grounding requires both people to participate at the same time. A robot proxy-based system does not require the remotely-located robot to be available for real-time interactions with the user; instead, the robot proxy provides crucial feedback to the user and supports transparency without consuming time or resources during plan execution. By engaging in a “conversation” with the robot proxy, the user has the opportunity to learn about the robot’s capabilities and provide additional information to the robot about the user’s goals without the need to communicate with the physical robot itself. When the user is satisfied with the plan, the robot proxy delivers the plan to the robot for execution (Figure 1).

At this time, our robot proxy has the ability to provide feedback to the user about plans before they are executed; a more complex proxy which is able to support a richer conversation with the user is currently under development. This proxy will actively work to infer the user’s goals and ask questions about those goals in order to obtain information above and beyond a simple list of commands (as opposed to passively illustrating what happens when a list of commands is executed, as in a normal simulator).

The results of the study presented here serve to demonstrate how the use of a robot proxy can improve task performance and efficiency, help foster more accurate mental

models of the robot’s capabilities, and create a stronger perception of effectiveness at the task as well as stronger feelings of collaboration with the robotic system.

2. STUDY DESIGN AND METHOD

In this experiment, we compared a robot proxy-based interface which could provide feedback to users about their plans for the robot before execution with an interface that could only pass plans from the user to the robot. We used a between-subjects design: each participant was randomly assigned to one of two conditions, the Robot Proxy condition or the Control condition. No physical robot was used in the study and all data were simulated.

The goals of the study were to understand the impact of a robot proxy-based interface on three particular areas relevant to common ground and exploration robotics tasks:

- *Task performance.* Which group is more efficient at completing the task successfully? How many correct and incorrect plans does each group send to the robot?
- *Mental model development.* After completing the task, which group knows more about the robot’s capabilities and can make accurate predictions about the robot’s behavior in novel situations?
- *Self-evaluation of performance.* How does the robot proxy-based interface affect participants’ perceptions of their own performance and their feelings of collaboration with the system?

It is important to note that within the exploration robotics domain, particularly planetary exploration, communication with the robot is often infrequent and costly. In the case of the Life in the Atacama project, the science team in Pittsburgh could only communicate with the robot twice a day: once to receive data from the robot and once to transmit a new plan to the robot. Because of this asynchrony, the amount of time required by the planning process is a less significant concern than in other human-robot interactions. To be consistent with this aspect of the exploration robotics domain, our investigation of task efficiency focuses primarily on how many communication cycles are required to complete the task rather than the amount of time spent by the user to create plans for the robot.

2.1 Participants

Thirty-six participants were recruited from Carnegie Mellon University; eighteen were assigned to each of the two conditions. All participants were graduate students or staff selected for their background in computer science (e.g., members of the School of Computer Science). Participants were compensated for their time upon completion of the study; they received either refreshments or US\$10 cash.

2.2 Procedure

After arriving at the lab, each participant was seated at a desktop computer. The experimenter explained that the computer would provide a description of the task and guide the participant through the task. The computer displayed the following scenario to each participant: “In this game, you will work with a Personal Exploration Rover (PER) which is located at an archeology site. Scattered around the site are fragments of a stone tablet covered in dirt. Each

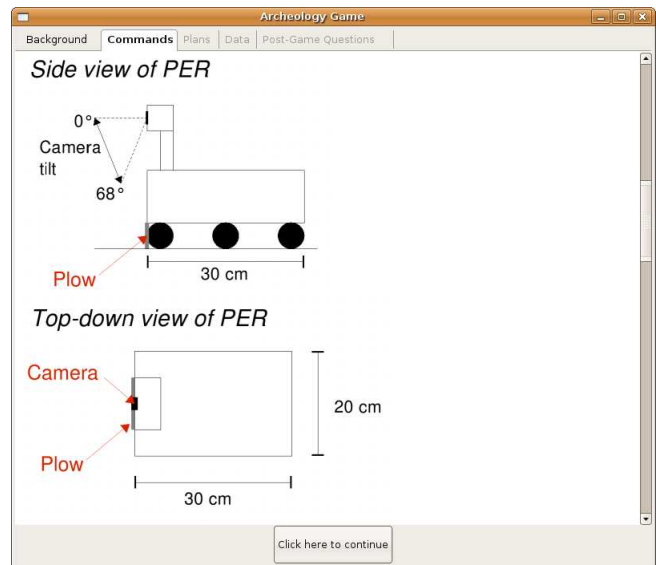



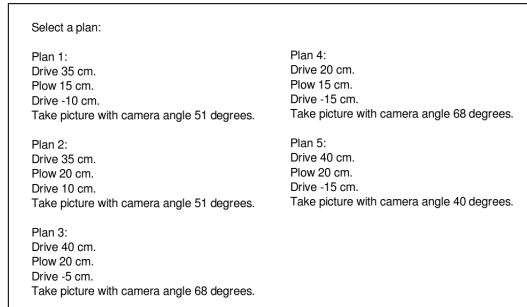
Figure 2: The side and top-down views of the robot as presented to study participants.

piece of the tablet contains words which can be combined to form a message. You must use the robot’s plowing abilities to scrape the dirt off of the tablet fragments and reconstruct the message. Once you have examined all of the fragments, you will be asked what you think the complete message is.” The participant was then shown a summary of the entire procedure for playing the game (instructions only relevant to the Robot Proxy group are shown here in boldface text; these instructions were not shown to the Control group):

1. The computer will present you with a set of plans for the fragment.
2. You choose the plan that you want the PER to execute. It’s multiple choice: you choose one plan from a set of five plans.
3. Once you have selected a plan, you choose **whether you would like feedback on the plan or whether you are ready for the robot to execute the plan.**
 - **If you choose to receive feedback, the system will analyze your plan and provide you with additional information about it. You can then tell the robot to execute the plan, or you can select a different plan. You may request feedback no more than two times before telling the robot to execute a plan.**
4. When you command the robot to execute the plan, it will execute the plan and return a picture. After it has executed the plan, the PER automatically “resets” (goes back to the location where it was before it executed the plan).
5. You decide whether to:
 - Send another plan to study the current fragment.
 - Go back to a previously visited fragment. The PER can automatically navigate to any fragment



(a) An example overhead map showing the location of the robot and the fragment.



(b) An example list of plans.

Figure 3: For each fragment, participants were given (a) an overhead map and (b) a set of five possible plans.

you have already seen so you can study it again. However, you may study the same fragment no more than three times all together.

- Go on to a new fragment. The PER can autonomously navigate to a new fragment so that you can send a plan to study it.

The participant was also provided with a specific list of robot commands available to use as well as diagrams depicting the robot's shape and size, which were available to the participant throughout the game (Figure 2).

The archeology site contained three fragments. For each fragment, the participant was given a map indicating the location of the robot and the nearest fragment as well as a set of five possible plans the robot could execute (Figure 3). The participant was asked to choose one of these five plans for the robot to execute given that only one plan was correct (only the correct plan would result in a complete picture of the fragment). Additionally, participants in the Robot Proxy group had the option of requesting feedback about a possible plan up to two times per fragment. This feedback consisted of an image containing a scale drawing of the robot, the location of the targeted fragment, and the field of view of one of the robot's instruments (Figure 4). The feedback was designed to provide both contextual information about the robot's surroundings (the location of the fragment) as well as to encourage an accurate mental model of the robot's capabilities (the field of view of the instrument).

After the participant selected a plan to be executed, he/she was shown the resulting image and given the opportunity to review this data. Figure 5 shows the image returned from




Figure 4: An example of the type of feedback shown to participants in the Robot Proxy group.



Figure 5: A screenshot containing an image of the first fragment after it has been completely cleaned.

the correct plan for the first fragment. The participant could send up to three plans for each fragment; however, each time a participant chose to re-send a plan for a fragment, the participant was given a new set of five plans from which to choose. This prevented participants from using a process of elimination to find correct plans.

Thus, during the activity, each participant completed two major activities multiple times: selecting a plan for the robot (planning) and examining the image that was returned from a selected plan (data review). We refer to a *cycle* as one planning session followed by one data review session. Each participant examined three different fragments for a total of three trials. After completing all three trials, the participant was asked a set of questions about his/her experiences. These questions included self-evaluation questions and questions intended to evaluate the participant's mental model of the robot (Section 2.4). The entire process lasted approximately thirty to forty-five minutes per participant.

2.3 Simulation

Because the task involved a non-collocated robot, software could be used to simulate a physical robot without sacrificing the fidelity of the human-robot interaction. Software was

Table 1: Dependent Variables

Variable	Measure
<i>Task Performance</i>	
Accuracy	Did the participant successfully identify the secret message?
# Cycles	How many cycles were required in order for the participant to reveal the entire secret message?
Review-Data Ratio	What proportion of the participant’s time spent on the task was used to review data from the robot?
<i>Mental Model Development</i>	
Quiz Score	What percentage of questions about the robot’s capabilities did the participant answer correctly? (14 questions, included both true-false and multiple-choice questions)
<i>Self-Evaluation of Performance</i>	
Effectiveness	To what extent did the participant agree or disagree that he/she was efficient at performing the task and felt confident during the task? (4 questions)
Fun	To what extent did the participant agree or disagree with the statement, “I had fun playing this game.”?
Collaboration	To what extent did the participant agree or disagree with the statement, “When developing plans, I felt I was collaborating with the system.”?

used to simulate the robot’s actions and the data returned from the robot. At the end of the experiment, the participant was informed that he/she had been using a simulated robot. It is important to note that the simulated robot’s actions were not stochastic: the robot always executed plans consistently and perfectly; poor-quality images of fragments were solely the result of the incorrect plans chosen by participants.

2.4 Dependent Variables

Table 1 illustrates the dependent variables measured in our study. The performance and mental model variables were derived from the requirements of the task itself. Mental model questions included questions about the robot’s physical properties as well as questions about how the robot would perform in situations similar to (yet slightly different from) those seen during the task. For all six self-evaluation questions, participants were given a Likert scale from 1 to 5 and asked how strongly they agreed with a particular statement (1 = “Strongly disagree”, 5 = “Strongly agree”). We used factor analysis to confirm that one question on perceived efficiency and three questions on confidence could be combined into a coherent factor “Effectiveness”; the Cronbach’s alpha of this factor was calculated to be 0.74, which suggests that the factor is internally consistent. The question about the participants’ feelings about collaborating with the robot was motivated by work by Hinds et al. on human-robot collaborative tasks [11].

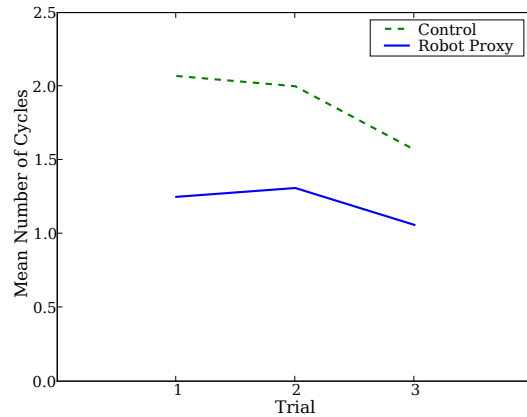


Figure 6: Mean number of cycles per trial for participants who successfully completed the task.




Figure 7: Mean number of incorrect plans sent to the robot per trial.

3. RESULTS

Our data analysis focused primarily on understanding differences between the Robot Proxy and Control participants and how participants’ performance changed over the three trials for the dependent variables given in Table 1. The multivariate correlations between dependent variables are shown in Table 2.

3.1 Task Performance

Overall, 30 of 36 participants successfully completed the task by revealing the entire secret message; a plot of the mean number of cycles used per trial is shown in Figure 6. To better understand participants’ performance, we conducted a two-way repeated measures analysis of variance (ANOVA) on the data from the 30 participants who were successful with condition as a between-subjects variable and trial number as a within-subjects factor. There was a main effect for condition ($F[1, 28] = 37.52, p < .001$), indicating that participants in the Robot Proxy group needed significantly fewer cycles than those in the Control group. The main effect of trials was also significant ($F[2, 56] = 4.44, p <$

Table 2: Multivariate Correlation

	# Cycles Trial 1	# Cycles Trial 2	# Cycles Trial 3	Review-Data Ratio	Quiz Score	Effectiveness	Fun	Collaboration
# Cycles Trial 1	1.000	0.450	0.319	0.352	-0.168	-0.272	0.127	-0.060
# Cycles Trial 2		1.000	0.481	0.499	-0.290	-0.224	0.045	-0.143
# Cycles Trial 3			1.000	0.613	-0.581	-0.274	0.190	-0.270
Review-Data Ratio				1.000	-0.252	-0.259	0.222	-0.204
Quiz Score					1.000	0.355	-0.235	-0.134
Effectiveness						1.000	0.005	0.283
Fun							1.000	0.230
Collaboration								1.000

Statistically significant ($p < .05$)

.05). This shows that participants required significantly fewer cycles during the later trials, which indicates that learning occurred over the trials. There was no significant interaction effect between condition and trials.

We also ran a two-way repeated measures ANOVA on the number of correct plans sent to the robot with condition as a between-subjects variable and trial number as a within-subjects factor. There was a main effect for condition ($F[1, 34] = 11.17, p < .01$): participants in the Robot Proxy group sent significantly more correct plans to the robot. 91% of trials in the Robot Proxy condition resulted in a correct plan being sent to the robot as opposed to 59% of trials in the Control condition. There was no significant main effect of trials; this was the expected result because each trial ended as soon as one correct plan was sent to the robot. There was also no significant interaction effect between condition and trials.

The ANOVA on the number of incorrect plans also showed a main effect for condition ($F[1, 34] = 38.4, p < .001$), meaning that participants in the Robot Proxy group sent significantly fewer incorrect plans to the robot than participants in the Control group (Figure 7). We also observed a significant main effect for trials ($F[2, 68] = 3.46, p < .05$), which indicates that participants sent significantly fewer erroneous plans to the robot during the later trials, which provides further evidence of learning. There was no significant interaction effect between condition and trials.

In addition, we conducted a two-way repeated measures ANOVA on the review-data ratio (the proportion of time spent on data review) with condition as a between-subjects variable and trial number as a within-subjects factor. The main effect of condition was highly significant ($F[1, 34] = 36.4, p < .001$), meaning that participants in the Robot Proxy group used much less of their time reviewing data from the robot than did participants in the Control group (Figure 8). There was no significant main effect of trials nor a significant interaction effect between condition and trials. One possible explanation for the main effect of condition is that participants in the Control group may need more time to review the data because they must both interpret the data and use it to improve their mental models. By contrast, robot proxy users may update their mental models based on the feedback they receive during the planning process and so do not need to spend as much time reviewing the data. This finding is also supported by our correlation analysis, which indicates that the proportion of time spent reviewing data was significantly positively correlated with the number of cycles used in each trial (Table 2).




Figure 8: Mean proportion of time spent reviewing data from the robot per trial.

3.2 Mental Model Development

After completing the activity, participants were asked fifteen questions related to the robot’s physical structure and capabilities. One question of the fifteen was not answered correctly by any participant and was therefore dropped from the analysis. A plot of the least squares mean score of each question by condition is shown in Figure 9. We conducted a multiple analysis of variance on participants’ scores for each quiz question with condition as a between-subjects variable and question number as a within-subjects factor. The results indicated that there was no main effect for condition ($F[1, 34] = 2.55, p > .1$). We found a significant main effect for question number ($F[12, 23] = 15.26, p < 0.001$) as well as a small interaction effect between question score and condition ($F[12, 23] = 1.86, p < 0.1$). This indicates that, while there was no significant difference in average total quiz score between the groups ($M_{RP} = 54\%, SD_{RP} = 15\%$, $M_{Control} = 45\%, SD_{Control} = 18\%$), whether or not a participant answered a particular question correctly was related to his/her group membership. This result is also reflected in Figure 9: while the average difficulty of questions varied, members of the Robot Proxy group tended to score higher on most questions.

We also observed a negative correlation between total quiz score and the number of cycles required for each trial; the magnitude of the correlation increased over time (Table 2). This suggests that participants who scored higher on the




Figure 9: Least squares mean of score on each question by condition.

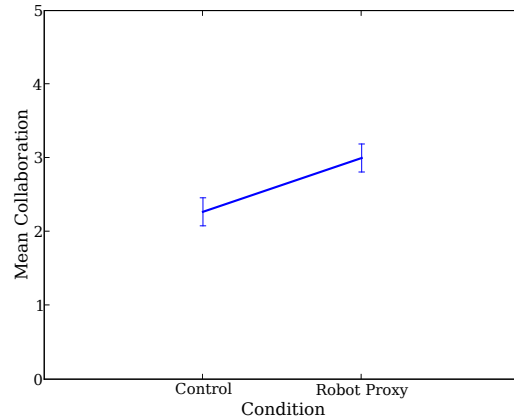


Figure 11: Least squares means of self-evaluation of collaboration with the system.




Figure 10: Least squares means of self-evaluation of effectiveness.

quiz needed fewer cycles to complete the task. This provides evidence that a higher score on the mental model quiz was associated with better performance on the task.

3.3 Self-Evaluation of Performance

Regression analysis was used to ascertain if the presence of the robot proxy could explain the differences in participants' self-evaluation of their performance on the task. The regression of efficiency on condition was significant ($M_{RP} = 3.82$, $M_{Control} = 3.08$, $r^2 = 0.19$, $p < .01$) (Figure 10). The regression of collaboration on condition was also significant ($M_{RP} = 3.0$, $M_{Control} = 2.28$, $r^2 = 0.11$, $p < .05$) (Figure 11). There was no significant difference with respect to participants' ratings of the task as fun. This shows that participants' ratings of their own effectiveness and feelings of collaboration with the system were strongly impacted by their interaction with a robot proxy.

4. DISCUSSION

Our participants were representative of the highly-trained scientists who currently participate in robotics exploration

missions due to participants' strong mathematics backgrounds and lack of direct experience remotely controlling robots for scientific exploration. The task was fairly short and straightforward: there were only three trials, and most participants successfully determined the secret message. While our results indicated that participants did learn over the course of the three trials, we found that the Robot Proxy group performed significantly better than the Control group on the first trial, a strong example of one-trial learning. With a longer, more complex task, we would expect to see more differences between the two groups in terms of learning, performance, and mental model development; this could be verified in a future experiment. We would also expect more significant differences between the groups based on the type of the feedback provided; different forms of feedback (i.e. three-dimensional images, video, etc.) could be compared in future studies. However, the basic feedback used in this study was still sufficient to highlight the benefits of a robot proxy, and we plan to conduct further studies when our full proxy implementation is complete.

Through this study, we have demonstrated proof-of-concept for using a robot proxy to increase common ground between a user and a remotely located robot as they complete an exploration task. Users who could request feedback about their plans before those plans were sent to the robot were more accurate (sent many fewer incorrect plans) and more efficient (required fewer cycles in order to successfully complete the task). Robot proxy users were also able to develop a better mental model of the robot, which was correlated with improved efficiency. The use of the robot proxy also explains participants' stronger feelings of effectiveness at the task and collaboration with the system, the benefits of which have been shown in [11]. In addition, we found that that individuals who completed the task in fewer cycles also spent less time reviewing data from the robot relative to their total amount of time spent on the task. Experiments involving a greater number of trials are needed, but we hypothesize that this review-data ratio could be used as a real-time, quantitative estimate of the common ground between a robot and user.

Future challenges include reformulating the task to be longer and more challenging. We are also interested in com-

paring the results of the study when a real robot is used (as opposed to a simulated robot). In previous work, we have documented the significance of problems stemming from a lack of common ground due to errors or failures when a real robot executes a plan [18, 19]. Based on these results, we are interested in creating a system on-board the robot which uses the information gained by the robot proxy to help the robot make better decisions in the event of failures at execution time. We are also currently analyzing how robot-proxy grounding could apply to other domains based on Clark and Brennan’s analysis of grounding constraints [6]. We expect that the use of a real robot and more challenging task would reveal further significant differences in task performance and mental model accuracy between participants who use a robot proxy during planning and those who do not.

5. ACKNOWLEDGMENTS

This research was supported by the NASA ASTEP program under grant NAG5-12890. The authors would like to thank Sara Kiesler, whose assistance with data analysis was invaluable. We would also like to thank all colleagues who contributed to this study, in particular Debra Bernstein, Sonia Chernova, Rachel Gockley, Colin McMillen, and Cristen Torrey.

6. REFERENCES

- [1] S. E. Brennan and E. A. Hulstien. Interaction and feedback in a spoken language system: A theoretical framework. *Knowledge-Based Systems*, 8:143–151, 1995.
- [2] J. L. Burke, R. L. Murphy, M. D. Covert, 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-Computer Interaction*, 19(1–2):85–116, 2004.
- [3] H. Clark and C. Marshall. Definite reference and mutual knowledge. In A. K. Joshi, B. L. Webber, and I. A. Sag, editors, *Elements of discourse understanding*, pages 10–63. Cambridge University Press, 1981.
- [4] H. Clark and D. Wilkes-Gibbs. Referring as a collaborative process. *Cognition*, 22(1):1–39, 1986.
- [5] H. H. Clark. *Using Language*. Cambridge University Press, 1996.
- [6] H. H. Clark and S. E. Brennan. Grounding in communication. In L. B. Resnick, R. M. Levine, and S. D. Teasley, editors, *Perspectives on socially shared cognition*, pages 127–149. APA, 1991.
- [7] J. L. Drury, L. Riek, and N. Rackliffe. A decomposition of UAV-related situation awareness. In *Proceedings of the First Annual Conference on Human-Robot Interaction*, pages 89–94, March 2006.
- [8] J. L. Drury, J. Scholtz, and H. A. Yanco. Awareness in human-robot interactions. In *International Conference on Systems, Man and Cybernetics 2003*, volume 1, pages 912–918, October 2003.
- [9] M. R. Endsley. Theoretical underpinnings of situation awareness: A critical review. In M. R. Endsley and D. J. Garland, editors, *Situation Awareness: Analysis and Measurement*, chapter 1, pages 1–32. Lawrence Erlbaum, 2000.
- [10] E. Falcone, R. Gockley, E. Porter, and I. Nourbakhsh. The Personal Rover Project: The comprehensive design of a domestic personal robot. *Robotics and Autonomous Systems*, 42:245–258, 2003.
- [11] P. J. Hinds, T. L. Roberts, and H. Jones. Whose job is it anyway? A study of human-robot interaction in a collaborative task. *Human-Computer Interaction*, 4(1–2):151–181, 2004.
- [12] H. Jones and P. Hinds. Extreme work teams: Using SWAT teams as a model for coordinating distributed robots. In *Proceedings of Computer Supported Cooperative Work 2002, New Orleans, Louisiana*, pages 372–380, November 2002.
- [13] S. Kiesler. Fostering common ground in human-robot interaction. In *Proceedings of the IEEE International Workshop on Robots and Human Interactive Communication (RO-MAN)*, pages 729–734, 2005.
- [14] R. M. Krauss and S. R. Fussell. Social psychological models of interpersonal communication. In E. T. Higgins and A. Kruglanski, editors, *Social psychology: Handbook of basic principles*, pages 655–701. Guilford Press, 1996.
- [15] S. Li, B. Wrede, and G. Sagerer. A computational model of multi-modal grounding. In *Proceedings of the ACL SIGdial Workshop on Discourse and Dialog, in conjunction with COLING/ACL 2006*, pages 153–160. ACL Press, 2006.
- [16] T. Paek and E. Horvitz. Uncertainty, utility, and misunderstanding: A decision-theoretic perspective on grounding in conversational systems. In *Psychological models of communication in collaborative systems: Papers from the AAAI Fall Symposium, November 5–7, North Falmouth, Massachusetts*, pages 85–92, 1999.
- [17] K. Severinson-Eklundh, H. Huttenrauch, and A. Green. Social and collaborative aspects of interaction with a service robot. *Robotics and Autonomous Systems, Special Issue on Socially Interactive Robots*, 42(3–4), 2003.
- [18] K. Stubbs, P. Hinds, and D. Wettergreen. Challenges to grounding in human-robot interaction: Sources of errors and miscommunications in remote exploration robotics. In *Proceedings of the First International Conference on Human-Robot Interaction*. ACM, 2006.
- [19] K. Stubbs, P. Hinds, and D. Wettergreen. Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems, Special Issue on Interacting with Autonomy*, 22(2):42–50, March–April 2007.
- [20] C. Torrey, A. Powers, M. Marge, S. R. Fussell, and S. Kiesler. Effects of adaptive robot dialogue on information exchange and social relations. In *Proceedings of the First Annual Conference on Human-Robot Interaction*, pages 126–133. ACM, March 2006.
- [21] H. A. Yanco, J. L. Drury, and J. Scholtz. Beyond usability evaluation: Analysis of human-robot interaction at a major robotics competition. *Human-Computer Interaction*, 19(1–2):117–149, 2004.