Intelligent Systems
IEEE
www.computer.org/intelligent

# Autonomy and Common Ground in Human-Robot Interaction

*Kristen Stubbs, Pamela J. Hinds, and David Wettergreen*

Vol. 22, No. 2
March/April 2007

# Autonomy and Common Ground in Human-Robot Interaction: A Field Study

**Kristen Stubbs and David Wettergreen,** *Carnegie Mellon University*

**Pamela J. Hinds,** *Stanford University*

***A two-year observational study of a collaborative human-robot system suggests that the factors disrupting the creation of common ground for interactive communication change at different levels of robot autonomy.***

The use of robots, especially autonomous mobile robots, to support work is expected to increase over the next few decades.[1] However, little empirical research examines how users form mental models of robots, how they collaborate with them, and what factors contribute to the success or failure of human-robot collaboration. A few observational studies report on people and robots working together in the unstructured "real world,"[2–4] but they remain relatively rare.

Through a detailed field study, we aimed to better understand how different levels and types of autonomy affect how users make sense of the actions of remotely located robots. The context for our field observations was the Life in the Atacama project. LITA used a robot to investigate microorganisms in Chile's Atacama Desert in a way analogous to planetary exploration. The project goals were twofold: to use the Atacama Desert as a testing ground to develop technologies and methodologies relevant to Mars exploration and to generate new scientific knowledge about the Atacama Desert itself. The technology development focused on a series of mobile robots and science instrument payloads.

Our observations of users collaborating with the remote robot showed differences in how the users reached common ground with the robot in terms of an accurate, shared understanding of the robot's context, planning, and actions—a process called *grounding*. We focus on how the types and levels of robot autonomy affect grounding. We also examine the challenges a highly autonomous system presents to people's ability to maintain a shared mental model of the robot.

## Related work

Understanding how people work with robots and how to design robots to better support people is the focus of the research area known as *human-robot interaction*. Jenny Burke and Robin Murphy summarize the open HRI research questions,[2] which include the type of modeling issues we address in this study. We use behavioral theory from the fields of communication, organizational behavior, and human-computer interaction to describe how understanding the process of building common ground can inform the design of human-robot systems and HRI.

### Common-ground theory

As two individuals participate in a joint activity, they accumulate common ground—that is, "the



1541-1672/07/$25.00 © 2007 IEEE
*Published by the IEEE Computer Society*

knowledge, beliefs, and suppositions they believe they share about the activity."[5] For example, the common ground between two people playing a tennis match would include knowledge of tennis rules, who won the last match, and how to hold the racket.

Herbert Clark and Deanna Wilkes-Gibbs propose that successful collaboration requires common ground: it helps collaborators know what information their partners need, how to present information so that it's understood, and whether partners have interpreted information correctly.[6] At an interaction's start, collaborators share a certain amount of common ground. For example, if they're members of the same discipline, they likely have a common language and perspective that facilitates communication.[7] Common ground can increase over time as collaborators share common experiences,[8] but it also can be disrupted by factors such as being in and drawing information from different physical contexts.[9] This interactive process establishes the common ground between collaborators.

### Common ground and human-robot interaction

Although researchers developed the common-ground framework to understand conversation and collaboration among people, not between people and machines, recent work has extended the framework to human-computer interaction.[10,11] This research suggests that we can improve interfaces by thinking about the user's experience as a conversation in which to develop shared meaning between the user and the machine interface.

In the HRI field, Hank Jones and Pamela Hinds observed SWAT (special weapons and tactics) teams and used their findings to inform the design of robot control architectures for coordinating multiple robots.[12] Although their observations didn't include robots, their findings emphasize the importance of common ground between a robot and its user, especially when the two aren't collocated. More recently, Sara Kiesler described experiments reporting more effective communication between people and robots when common ground is greater.[13] Other researchers found that information exchange is more effective when a robot can adapt its dialogue to fit a user's knowledge.[14]

### Situation awareness

Although generally focused more on dialogue and communication, the common-ground framework overlaps with work on *situation awareness*, which Mica Endsley defines as "knowing what is going on around you."[15] Researchers recently examined SA in HRI, particularly with urban search and rescue (USAR) robots.[16–18] Empirical work indicates that USAR operators spend significantly more time trying to gain SA—assessing the state of the robot and environment—than they do navigating the robot.[16,17] 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.

From their observations in the USAR domain, Burke and Murphy propose that shared mental models contribute to SA and that communication is critical to refining these models.[2] However, they don't test this relationship directly. In the work we report here, we aimed to examine more closely how grounding occurs or is disrupted between users and a robot as the robot's autonomy increases. The science team we observed had difficulty knowing what the robot was doing and what was going on in the robot's environment, and the robot—lacking information about the science team—was unable to respond appropriately. We aim to better understand how autonomy affects the information that both the science team and the robot need to build common ground. The common-ground framework facilitates a focus on the entire "conversation" between a user and a robot rather than solely on the user's information needs as in previous SA research.

The common-ground framework facilitates a focus on the entire "conversation" between a user and a robot rather than solely on the user's information needs.

### Autonomy and the grounding process

We base our analysis of robot autonomy on the work of Raja Parasuraman, Thomas Sheridan, and Christopher Wickens,[19] who define automation as "a device or system that accomplishes (partially or fully) a function that was previously, or conceivably could be, carried out (partially or fully) by a human operator." They distinguish between autonomy types and levels, describing four basic types: information acquisition, information analysis, decision selection, and action implementation. In robotics, these autonomy types are commonly collapsed into three:[20]

- *autonomous sensing* (information acquisition and data transformation)—making observations and refining information,
- *autonomous planning* (information interpretation and decision selection)—reacting to information or deciding actions and schedule, and
- *autonomous acting* (action implementation)—executing a planned task or producing reflexive reactions.

These types decompose information analysis into data transformation during sensing and interpretation during planning.

One robotic system can have a different autonomy level for each type—that is, sensing, planning, and acting. In the LITA project, we categorized the levels according to the extent of external guidance the system required to function:

- *low autonomy*—some basic automation might be present, but both information and procedures must be provided externally;
- *moderate autonomy*—some required information will come from an external source, such as intermediate steps or proper system settings, but all procedures function independently; and
- *high autonomy*—systems can both derive needed information and proceed independently over extended periods.

Our work's most significant contribution is a better understanding of how different autonomy types and levels affect grounding between people and robots, particularly teams of people and a remote robot.

## Study methodology

The LITA project robot, Zoë, is a four-wheeled, solar-powered robot equipped with several scientific instruments, including cameras and an underbelly fluorescence imager (FI) for detecting organic molecules such as proteins and amino acids (see figure 1).

**Figure 1. Zoë, the "robotic astrobiologist" created for the Life in the Atacama project.**

**Figure 2. Science team members discuss data returned from the robot.**

For this study, we focused on a particular part of the LITA field season known as *remote science operations*. During these periods, a team of biologists, geologists, and instrument specialists (in Pittsburgh) used the robot to search for signs of life in the desert. This science team issued daily commands to the robot and received and analyzed the data products it generated (see figure 2). An engineering team of roboticists and instrument specialists (in Chile) monitored the robot, conducted troubleshooting onsite, and ensured that the science team could gather data successfully.

To collect data about both sites, we had one researcher observe the Pittsburgh science team while one to two others observed the Chile engineering team and robot. The observations involved writing detailed field notes, drawing diagrams, and taking photographs and video clips. Communication between observers across sites was limited so that each observer could focus completely on the local situation and better understand the observed group's perspective at the time. We told the scientists and engineers that our research aimed to better understand how they work with remote rovers and that the observations would continue throughout field operations. During the 2004 field season, the Pittsburgh observers logged 138 hours of observations and the Chilean observers logged 241 hours. In 2005, the observations totaled 254 hours in Pittsburgh and 239 hours in Chile.

Our data set consisted of the observers' field notes together with artifact documents, such as PowerPoint presentations, emails, and robot plans that the science team generated. An initial reading of the data revealed many communication and coordination problems between sites. Next, we identified the specific errors and miscommunications that occurred and classified them (for example, "Error in plan sent to robot" or "Miscommunication regarding interpretation of plan"). We refer to these errors and miscommunications collectively as *problems*.

We identified those problems related to common ground according to whether the science team and robot lacked mutual knowledge and, if so, what kind (for example, "Missing contextual information," "Lack of transparency into robot's behavior"). Our 2004 data coding revealed 57 separate common-ground problems during the two weeks of remote science operations; the 2005 data revealed 91 common-ground problems during 23 days of remote operations. We then used the data to trace what caused these problems, particularly those related to the robot's autonomous capabilities.

## Operational autonomy levels

Figure 3 depicts the type (sensing, planning, acting) and level (low, moderate, high) of Zoë's autonomous capabilities throughout this study. During regular operations in 2004 and 2005, the science team sent the robot plans for executing low to medium autonomous sensing or planning. In 2005, the engineering team also introduced a science autonomy system that let Zoë collect data on its own without specific commands from the science team about where to do so. This gave Zoë much greater autonomy than it had during regular operations and let us observe autonomy's impact on grounding.

### 2004 regular operations: Low autonomy

Zoë had limited autonomous capabilities in 2004. It could record data about its internal state, detect some failure conditions, and detect obstacles; but it had difficulty accurately estimating its position over the long term. (The project's planetary-exploration

goal precluded using GPS.) The robot didn't interpret any science data and performed only basic planning for scheduling science actions. As figure 3 shows, autonomy with respect to planning and acting was low, and engineers often had to drive the robot manually. They also had to command instrument operations.

***Problems.*** The problems we saw in the 2004 data related predominantly to understanding references to objects of interest—problems that emerged from lack of copresence between the science team and the robot. Herbert Clark and Susan Brennan argue that grounding becomes more difficult when people are not copresent.[8] Catherine Cramton's work on geographically distributed teams supports this argument.[9]

According to Clark and Brennan, missing contextual information jeopardizes shared understanding because "the addressee must imagine appropriate contexts for both the sender and the message."[8] We observed numerous problems with contextual information that bear on challenges users face when interacting with a remote robot. Receiving erroneous data from a robot is always a possibility. Without sufficient information about data and the context of its collection, making sound scientific judgments is difficult. In one instance, the team received a fluorescence image in which nearly half the field of view appeared to be fluorescing, signaling the possible presence of life. This caused a great deal of excitement, but it was unclear whether the data indeed represented life, the camera had malfunctioned, or some other unforeseen event had occurred. After nearly a day investigating the image, the team concluded that sunlight was responsible for the strange glow they had observed. In this case, the lack of contextual information about the data resulted in confusion and much time spent trying to deduce what could have gone wrong.

Effective reference in communication requires *perspective-taking*—that is, a speaker must take into account the listener's perspective when formulating a referring expression.[21] When two people are physically separated, gaining insight into the other's perspective is difficult. In particular, feedback is less immediate, harder to interpret, and sometimes even nonexistent. Feedback about how well the listener understands the speaker's messages is crucial to conversational grounding.[22] In 2004, the science team lacked enough information from the robot to effectively take the robot's perspective, and the robot had no means to detect or improve that situation.

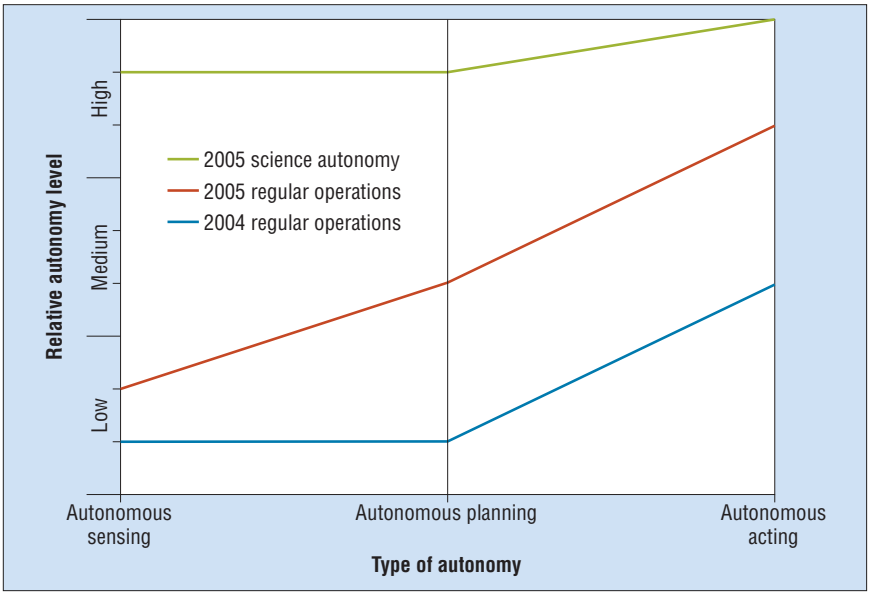


**Figure 3. Types and levels of robot autonomy in regular LITA operations and with the science autonomy system.**

***Discussion.*** During the 2004 season, the science team relied primarily on data the robot collected as well as information from engineers collocated with the science team to build common ground with the robot. At a basic level, the science team could determine what data had and hadn't been collected. However, they didn't have easy access to feedback about errors or instrument failures, so they turned to the collocated engineers, who could contact the field engineers and obtain additional contextual information about what was happening in Chile. Had these resources not been available, the grounding process would have been further impaired.

The most significant constraint on grounding at these low autonomy levels was in understanding the robot's perspective. Had the science team been able to observe the robot executing commands in the desert, they would have had enough contextual information to disambiguate problems. However, the lack of copresence combined with the lack of feedback from the robot about its actions inhibited grounding and led to frustration and errors. This observation is similar to studies of situation awareness, although we add to this work by considering the "conversation" between the science team and the robot, where the breakdowns occurred, and how the science team attempted to create common ground with the robot. In particular, we noticed that feedback from the robot was missing, as was robot awareness of and adjustment to the science team's confusion. In common-ground parlance, the conversation's *acceptance phase* was missing. The robot engaged in the *presentation phase* by providing information, but it didn't seek evidence of the science team's understanding. The conversation was therefore incomplete and led to misunderstandings.[9]

## 2005 regular operations: Moderate autonomy

In 2005, Zoë's autonomous navigation improved substantially during regular operations. Zoë could sense nearby obstacles, develop basic plans to avoid them, and act on those plans with minimal human intervention. It could drive autonomously between locations that the science team specified. In addition, as a result of problems in establishing common ground during the 2004 field season, the LITA field engineers sent a daily "robot report" to the science team as a proxy for the information that the robot should have provided autonomously. The report included which actions Zoë executed, which actions had and hadn't succeeded, instrument failures, and other contextual information.

***Problems.*** One technique the LITA science team used in both 2004 and 2005 to improve their understanding of the robot's context

**2005 regular operations (low to moderate autonomy)**
Problem: Interpreting context images

On day 3, one scientist (X) mentioned that a stereo panoramic imager image, which was supposed to include a field-of-view context for a fluorescence image, wasn't taken correctly:

*X looks at a particular SPI image and says that "this is the messed up one." X says that this was supposed to be a context image. X reads the robot report. X says that the robot moved before taking the SPI image. X says, "I'm not sure why that happened."*

Scientists X and Y spent time trying to find the FI field of view in SPI context images, but when the SPI images haven't been taken correctly, this is impossible. The science team used both the images returned from the robot as well as the robot reports to figure out what happened. On day 4, the science team talked about adjusting the commands sent to the robot to account for its backward movement and plowing (scraping away a shallow layer of soil to expose the ground beneath) 0.5 meters after an FI, before the SPI takes the context image. One scientist concluded that the robot should have moved only 1 meter, not 1.5 meters, before taking the context image:

*At 2:09 p.m., X tells Y that "we" might have to adjust the drive precise command for the FI context image. X explains that after the FI, the rover moves back 0.5 meters for the marker plow. Y says that they are imaging the marker instead of the FI. X says that they might get the FI. X says that "we" may need to adjust. X says s/he thinks that the plow is right after the FI.*

*At 4:14 p.m., X says that s/he and Y were talking. They talked about the fact that since the marker plow is done at the end of the FI, "we" need to adjust how much to move [the robot] back up. X says "we" should have asked to move 1 meter.*

After this, the science team adjusted their commands to move the robot 1 meter (days 5, 6, 7, 9, 10) and later commanded the robot to move 1.5 meters (days 9, 11, 12). On day 11, one scientist explained that the team realized they had to change back to requesting 1.5 meters instead of 1 meter:

*X says that they need the plow as a marker, so they found they did have to move up to 1.5 meters to get into the initial position.*

**2005 science autonomy system (high autonomy)**
Problem: Missing fluorescence image follow-ups

On day 1, one scientist (X) observed that the science autonomy system should have taken a full FI sequence in response to a positive chlorophyll signature (a "follow-up"), but it didn't. An engineer (Z) confirms that the system should have taken a follow-up image:

*At 10:20 p.m., X is looking at a fluorescence image on the transect associated with locale 40 and asks, "Why didn't we have a follow-up on that?" X turns to Z and asks, "Shouldn't that have initiated a follow-up?" Z replies that yes, it should have.*

On day 1, engineer Z explained that rounding errors contributed to the problem and that the system was originally designed for much longer transects than what the scientists were using:

*At 11:45 p.m., Z explains to X about some of the science on-the-fly problems that [the engineering team] had with the fluorescence imager. Z says the problem had to do with "round off" and "resource juggling." Z says that for fractional distances, the rover will always round up. X says that [the robot] went 180 meters. Z explains that the algorithm was designed for much longer distances. X explains that [the scientists] want to make the 180-meter traverse a standard procedure.*

On day 15, members of the science team and the engineer talked about other reasons why the follow-ups might not have been initiated:

*Scientist A says that s/he is going to look at the transect between 800 and 810 to try and figure out why there were three full FIs and three chlorophyll only, but it doesn't look like there was a chlorophyll follow-up. A says this has happened before. Y suggests that it could be the result of the delta in the signal between the pre and post (the difference in the signal). Z says that the algorithm uses raw signal values.*

This technical discussion continued without resolving why the robot hadn't performed follow-ups as expected.

**Figure 4. Example common-ground problems from LITA project field notes. Italics indicates exact excerpts from field notes.**

was to have its stereo panoramic imager camera take a *context image* of the area it had examined using the FI. The SPI image gave scientists additional information about the larger area in which the FI had taken its image. However, the robot didn't always take these context images correctly, and the science team had to detect these errors and determine what had happened. This problem occurred on days 3, 4, and 10. The left column of figure 4 gives details on this scenario.

Throughout this scenario, the scientists relied on the robot's data and robot report to establish common ground regarding how the robot was operating, and they used this information to adjust the commands they sent to the robot. This process mirrors conversational grounding between people in that the science team attended to the robot's feedback and adjusted their communications in hopes of being more effective. However, the adjustment was one-sided. The robot didn't learn how to better communicate with the science team; as a result, the science team wasn't always successful at deducing the robot's actions.

In a second scenario, the science team wanted to understand exactly how far the robot traveled and where it collected data products. This task was complicated because different software programs computed distances in different ways. As a result, the distances measured in the plan-creation tool differed from distances shown in the human-readable plan, and these differed from the odometric distances that Zoë reported to have traveled, the telemetric estimates of how far Zoë traveled, and the actual distances Zoë traveled.

Our data suggest that even though some science team members understood Zoë's odometry and telemetry data, it didn't help them plan paths for the robot. The team used the robot report as a definitive source for how far the robot traveled between locales. This might have been because the report was the only easily accessible source of this information.

***Discussion.*** Without the benefit of copresence, the science team used the robot reports and the data from the robot as their main sources of information about what had happened in the field. However, this feedback wasn't inadequate to establish common ground with the robot. The science team couldn't understand the robot, and the robot didn't verify the science team's understanding through an acceptance phase. Nor did the robot learn how to better refer to objects, locations, and other environmental factors so that it and the science team could expand their common ground.

In our 2004–2005 regular operations data, the major issues of copresence and inadequate feedback appear to be most associated with moderate to high levels of autonomous acting (see figure 3). The robot was acting autonomously (albeit sometimes at low levels) by driving and deploying instruments with little or no human interaction. Without contextual information or adequate feedback, the science team found it difficult to understand the autonomous actions. The robot had no means to maintain its end of the conversation by detecting the science team's difficulty in understanding the information it presented to them.

## 2005 science operations: High autonomy

The science autonomy system added to Zoë during 2005 consisted primarily of software to collect and interpret sensor, camera, and instrument data and software to plan a response, if any, to these observations. The engineering team had designed the system to let the robot collect science data as it traveled between locations of scientific interest. The science team could use the system to request autonomous collection of normal camera images and chlorophyll-only fluorescence images. If the robot detected that such a fluorescence image showed evidence of life, it would follow up by taking a full fluorescence image set. The science autonomy system gave Zoë much greater autonomy (see figure 3). It could sense, plan, and deploy instruments with little to no human intervention. The system also forced the science team to adopt a different strategy for grounding. In particular, we noticed that issues arose around why the robot made certain decisions in addition to recurring questions about objects of reference as described in regular operations.

***Problems.*** On days 1, 2, 3, 4, and 15, the science team discussed the robot's failure to perform follow-ups when it should have. The science team attempted to find out why (see figure 4, column 2, for details on this scenario). In contrast to the examples from regular operations, in this scenario the science team understands what has and hasn't been done but is baffled about why the robot made particular decisions. They attempt to reason among themselves and with an engineer about what Zoë might be "thinking," but they don't understand the robot's decision-making algorithms well enough or have enough feedback from the robot to communicate and get the data they want. The robot has no means to represent or reason about why the science team has chosen particular actions, so it can't ensure that the rationale for its actions is understood or that its decisions are consistent with the science team's larger goals.

> The lack of copresence was particularly pronounced when the science team tried to understand the robot's high-level autonomous sensing and action.

In the field, the engineering team was aware that some problems with the follow-up mechanism resulted from water on rocks or sunlight shining under the robot. Because the science team lacked this information about the context, they had to try to deduce why the robot decided not to take follow-up fluorescent images. Their grounding strategies included examining the chlorophyll fluorescent images to see how strong the signal was; they also calculated the time when the FI took the images to determine whether sunlight might have been an issue. This example shows evidence of breakdowns in shared perspectives on what was located where and why actions were or weren't taken.

***Discussion.*** With the high levels of sensing, planning, and action autonomy that the robot possessed when using the science autonomy system, the science team's problems were focused less on exactly what the robot was doing than on why the robot was making particular decisions.

The lack of copresence continued to be a constraint and was particularly pronounced when the science team tried to understand the robot's high-level autonomous sensing and action. In addition, we observed that transparency became a constraint with high-level autonomous planning. Even if the science team had been watching the robot while the science autonomy system was working, they wouldn't necessarily have had enough information to determine why the robot stopped in particular locations or failed to perform follow-up fluorescent images. The science team had to try to understand not only how the robot would react to positive or negative evidence of life but also what its analysis process was.

On the basis of the science team's strategies to understand the science autonomy system, we argue that the lack of transparency into the robot's decision-making process became the primary constraint on grounding. The robot report provided only factual information and nothing about why the robot performed measurements or follow-ups. Instead, the science team used the data to determine what might have happened and then relied on engineers to explain the algorithms behind how the robot made decisions.

Some researchers have defined transparent interactions as those in which a user can "see through" the logic behind a machine's operation. Some of this research focuses on users' understanding,[23] some on robots' explanations of their actions,[24] and some on transparency that requires no mental model.[25] Consistent with the common-ground framework, we approach transparency as a dynamic feature of the science team's interaction with the robot. Transparency therefore refers to the process of developing common ground between them. Jakob Bardram and Olav Bertelsen[26] similarly suggest that transparency can't be understood as a static feature but must reflect a deliberate formulation and refinement of understanding during the course of human-computer interaction. Although people certainly ask questions and converse about reasons for their thoughts and actions, this idea of understanding someone's logic isn't well articulated in current common-ground research. From our LITA project observations, we argue that the dynamic creation of transparency becomes a more crucial element for creating common ground as robots acquire higher levels of autonomy, particularly autonomous planning.

Figure 5 illustrates this shift from a focus

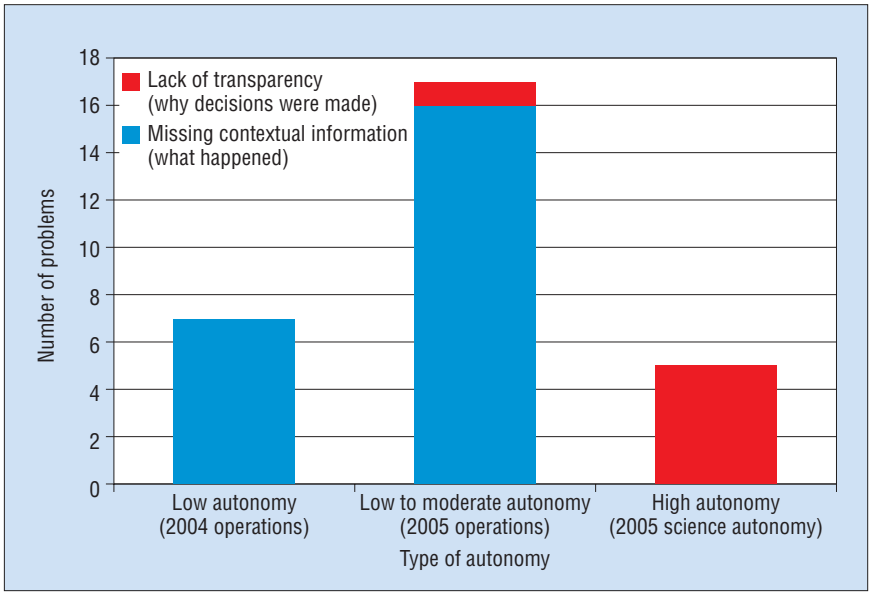


**Figure 5. Common-ground issues for problems relating to autonomy.**

on missing contextual information to a lack of transparency. From 148 total problems related to common ground that we identified from the 2004 and 2005 data, the figure shows the number for which missing contextual information or a lack of transparency was the most significant cause. As the graph indicates, the nature of the problems shifted almost entirely away from problems with missing context to issues of transparency about the robots' decisions and logic in the high-autonomy scenario. These results are even more dramatic when you consider that each problem might have occurred on multiple days and that problems related to a lack of transparency generally took more days to resolve than those related to missing contextual information.

## Implications for system design

Our results suggest that for HRI grounding to occur, particularly with remote robots, the robots must learn and adjust their behaviors on the basis of "conversations" with people. Some researchers have demonstrated a robot that can automatically adjust its dialogue in real time by exploiting its ability to create models of people or teams according to their backgrounds.[14] In addition, software systems can perfectly recall prior conversations with users, so robots might use this information to learn and adapt, just as humans do in conversational grounding. Implementing such adaptation might not be easy with current technology, but our results suggest this is a promising direction for future work and might address the recurring issue we observed with missing contextual information and confusion about objects of reference.

For remote-exploration robotics, the cost of mistakes in data collection is extremely high. Data that's not useful to the science team wastes valuable time and resources. However, delay costs are extremely low: given that the plan goes to the robot after it finishes its daily operations, the science team doesn't pay a penalty in terms of data return on the time spent revising the plan. We therefore recommend creating a software system that can participate in grounding during plan creation. This system would act as a robot proxy, providing crucial feedback to the science team and supporting transparency without consuming time or resources during plan execution. Such a system would improve the conversational grounding, which requires both parties' availability. The system could exactly recreate the robot's behavior without requiring the actual robot's participation. This helps people interacting with a remote robot understand exactly how it would respond to their requests, and it provides the immediate feedback so critical to grounding.

In addition, the system could promote transparency by actively detecting errors in a user's understanding. When a user provides inappropriate responses to questions or expresses confusion, the system could detect these grounding problems and automatically disclose its logic by providing additional information, such as the evidence it used to make a particular decision. Situation awareness research hasn't generally considered a robot with capabilities to detect and respond to grounding opportunities because SA historically hasn't focused on the conversation between users and the robot.

Software designers can use the presentation-acceptance process to drive interactions at the level of the individual actions and parameters to be sent to the robot. In particular, Herbert Clark and Deanna Wilkes-Gibbs's detailed description of the acceptance process provides specific guidance for interaction design at this low level.[6] In conversation, when a speaker presents an initial reference that isn't acceptable, either the speaker or the listener can repair, expand, or replace it (or request such a repair, expansion, or replacement). In the context of exploration robotics, we can consider an individual action and its parameters to be analogous to a reference in conversation. If need be, a scientist could repair an action by editing its parameters, expand an action by providing additional information such as a target's name, or replace an action in the plan with a different action. Presentation-acceptance for one action could then proceed as follows:

Scientist presents an action $a_i$.
Proxy system checks if $a_i$ is adequate (free of errors, consistent with other actions, and so on).
If $a_i$ is adequate:
  Proxy system accepts $a_i$.
  Proxy system provides positive evidence of acceptance.
Else:
  Proxy system presents negative evidence.
  Proxy system requests a repair (a revision, expansion, or replacement).
While scientist needs information:
  Scientist requests an expansion (further information about the inadequacy).
  Scientist presents the requested repair, $a_i'$.
Let $a_{i+1} = a_i'$. Repeat.

As Clark and Wilkes-Gibbs observe, the acceptance process is recursive.[6] From a software design point of view, it's potentially infinite; the system must be appropriately scaled to strike a balance between supporting transparency and ensuring that a plan can be completed. In addition, we're investigat-

ing how the robot can use this information to reason about higher-level science goals beyond the execution of individual actions. Beyond letting the user simulate the robot's actions,[27] this will allow the robot to build common ground with the science team regarding the relationships between different actions and the environment. If an action fails in the field, the robot can then exploit this information to repair its plan in a manner consistent with the scientists' goals.

The human-robot system that we observed wasn't a mixed-initiative system, in which the division of authority between the robot and the users could be adjusted in real time.[28] In the LITA project, only the robot could perform certain actions, and the science team couldn't exert authority in those situations. In a mixed-initiative HRI system, the grounding process would likely differ from what we observed in this study. For example, grounding between the users and robot would need to include a shared understanding of how and why authority shifts. The problem of grounding in mixed-initiative systems poses an interesting research topic.

We spent more than 800 hours observing the LITA mission and documenting the grounding process between the science team and the robot. As autonomy increased, we saw the science team's confusion about the robot's actions move away questions about the data-collection context to questions about why the robot was doing what it was doing. We also observed the grounding process become more complicated when the entire team tried to work together with the science autonomy system. Our data suggest that a team's shared mental model of an autonomous robot is more complex and variable than it is for simple devices and that it needs to be more consistent.

Higher autonomy didn't necessarily lead to better or more error-free interaction. Common-ground problems emerged whether autonomy was low or high. Our data suggests that designers must be aware of how autonomy changes the type of information needed from the robot and the type of "conversation" HRI requires. For grounding to occur with low-autonomy robots, contextual information and feedback are particularly critical; at high autonomy levels, particularly for autonomous planning, users need transparency with respect to the robot's decision making. ■

## References


1. S. Thrun, "Toward a Framework for Human-Robot Interaction," *Human-Computer Interaction*, vol. 19, nos. 1–2, 2004, pp. 9–24.
2. J.L. Burke and R.R. Murphy, *Situation Awareness, Team Communication, and Task Performance in Robot-Assisted Technical Search: Bujold Goes to Bridgeport*, tech. report CRASAR-TR2004-23, Dept. of Computer Science and Eng., Univ. of South Florida, 2004.
3. J.L. Drury, L. Riek, and N. Rackliffe, "A Decomposition of UAV-Related Situation Awareness," *Proc. 1st Ann. Conf. Human-Robot Interaction* (HRI 06), ACM Press, 2006, pp. 89–94.
4. R.M. Siino and P.J. Hinds, "Robots, Gender, and Sensemaking: Sex Segregation's Impact on Workers Making Sense of a Mobile Autonomous Robot," *Proc. 2005 IEEE Int'l Conf. Robotics and Automation* (ICRA 05), IEEE Press, 2005, pp. 2773–2778.
5. H.H. Clark, *Using Language*, Cambridge Univ. Press, 1996, pp. 38.
6. H. Clark and D. Wilkes-Gibbs, "Referring as a Collaborative Process," *Cognition*, vol. 22, no. 1, 1986, pp. 1–39.
7. S.R. Fussell and R.M. Krauss, "Coordination of Knowledge in Communication: Effects of Speakers' Assumptions about What Others Know," *J. Personality and Social Psychology*, vol. 62, no. 3, 1992, pp. 378–-391.
8. H. Clark and S. Brennan, "Grounding in Communication," *Perspectives on Socially Shared Cognition*, L.B. Resnick, R.M. Levine, and S.D. Teasley, eds., Am. Psychological Assoc., 1991, pp. 127–149.
9. C.D. Cramton, "The Mutual Knowledge Problem and Its Consequences for Dispersed Collaboration," *Organization Science*, vol. 12, no. 3, 2001, pp. 346–371.
10. S.E. Brennan and E.A. Hulteen, "Interaction and Feedback in a Spoken Language System: A Theoretical Framework," *Knowledge-Based Systems*, vol. 8, nos. 2–3, 1995, pp. 143–151.
11. T. Paek and E. Horvitz, "Uncertainty, Utility, and Misunderstanding: A Decision-Theoretic Perspective on Grounding in Conversational Systems," *Psychological Models of Communication in Collaborative Systems: Articles from the AAAI Fall Symp.*, AAAI Press, 1999, pp. 85–92.
12. H. Jones and P. Hinds, "Extreme Work Teams: Using SWAT Teams as a Model for Coordinating Distributed Robots," *Proc. 2002 ACM Conf. Computer Supported Cooperative Work* (CSCW 02), ACM Press, 2002, pp. 372–380.
13. S. Kiesler, "Fostering Common Ground In Human-Robot Interaction," *Proc. 14th IEEE Int'l Workshop Robots and Human Interactive Communication* (ROMAN 05), IEEE Press, 2005, pp. 729–734.
14. C. Torrey et al., "Effects of Adaptive Robot Dialogue on Information Exchange and Social Relations," *Proc. 1st ACM SIGCHI/SIGART Conf. Human-Robot Interaction* (HRI 06), ACM Press, 2006, pp. 126–133.
15. M.R. Endsley, "Theoretical Underpinnings of Situation Awareness: A Critical Review," *Situation Awareness: Analysis and Measurement*, M.R. Endsley and D.J. Garland, eds. Lawrence Erlbaum, 2000, pp. 1–32.
16. J.L. Burke et al., "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*, vol. 19, nos. 1–2, 2004, pp. 85–116.
17. J.L. Drury, J. Scholtz, and H.A. Yanco, "Awareness in Human-Robot Interactions," *Proc. 2003 Int'l Conf. Systems, Man and Cybernetics*, vol. 1, 2003, pp. 912–918.
18. 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*, vol. 19, nos. 1–2, 2004, pp. 117–149.
19. R. Parasuraman, T.B. Sheridan, and C.D. Wickens, "A Model for Types and Levels of Juman Interaction with Automation," *IEEE Trans. Systems, Man, and Cybernetics—Part A: Systems and Humans*, vol. 30, no. 3, 2000, pp. 286–297.
20. M.M. Veloso, "Entertainment Robotics," *Comm. ACM*, vol. 45, no. 3, 2002, pp. 59–63.
21. R.M. Krauss and S.R. Fussell, "Social Psychological Models of Interpersonal Communication," *Social Psychology: Handbook of Basic Principles*, E.T. Higgins and A. Kruglanski, eds., Guilford Press, 1996, pp. 655–701.
22. M.J. Traxler and M.A. Gernsbacher, "Improving Written Communication through Minimal Feedback," *Language and Cognitive Processes*, vol. 7, no. 1, 1992, pp. 1–22.
23. R. Sinha and K. Swearingen, "The Role of Transparency in Recommender Systems," *CHI 02: Extended Abstracts on Human Factors in Computing Systems*, ACM Press, 2002, pp. 830–831.
24. J. Herlocker, J.A. Konstan, and J. Riedl, "Explaining Collaborative Filtering Recommendations," *Proc. 2000 ACM Conf. Computer-Supported Cooperative Work* (CSCW 00), ACM Press, 2000, pp. 241–250.

## The Authors

**Kristen Stubbs** is a PhD candidate in Carnegie Mellon University's Robotics Institute. Her research interests include observational studies of human-robot systems to better understand people's mental models of them. For her thesis, she is working to design and implement software systems that will let robots participate actively in building common ground with their users. She received her MS in robotics from Carnegie Mellon. Contact her at Carnegie Mellon Univ., Robotics Inst., 5000 Forbes Ave., Pittsburgh, PA 15213; kstubbs@cmu.edu.

**Pamela J. Hinds** is an associate professor and co-director of the Center on Work, Technology, and Organization in Stanford University's Department of Management Science and Engineering. Her research focuses on the effects of technology on groups, the dynamics of globally distributed work teams, and professional service robots in the work environment. She received her PhD in social and decision sciences from Carnegie Mellon University. Contact her at Center for Work, Technology, and Organization, Dept. of Management Science and Eng., Ternan Bldg. 307, Stanford Univ., Stanford, CA 94305; phinds@stanford.edu.

**David Wettergreen** is an associate research professor in Carnegie Mellon University's Robotics Institute. His research expertise is in field robotics, where he concentrates on endurance navigation, science understanding, robust autonomy for robotic explorers, and human interaction with these explorers. He received his PhD in robotics from Carnegie Mellon. He serves as an associate editor of the *Journal of Field Robotics* and is a member of the IEEE and the International Council on Systems Engineering. Contact him at Carnegie Mellon Univ., Robotics Inst., 5000 Forbes Ave., Pittsburgh, PA 15213; dsw@ri.cmu.edu.

25. M.A. Goodrich and J.D.R. Olsen, "Seven Principles of Efficient Human-Robot Interaction," *Proc. 2003 IEEE Int'l Conf. Systems, Man, and Cybernetics*, vol. 4, IEEE Press, 2003, pp. 3942–3948.

26. J.E. Bardram and O.W. Bertelsen, "Supporting the Development of Transparent Interaction," *Human-Computer Interaction, 5th Int'l Conf., Selected Articles*, B. Blumenthal, J. Gornostaev, and C. Unger, eds., Springer, 1995, pp. 79–90.

27. B. Hine et al., "VEVI: A Virtual Environment Teleoperation Interface for Planetary Exploration," *Proc. SAE 25th Int'l Conf. Environmental Systems* (ICES 95), Soc. Automobile Engineers, 1995, pp. 615–628.

28. D. Bruemmer et al., "Mixed-Initiative Control for Remote Characterization of Hazardous Environments," *Proc. 36th Hawaii Int'l Conf. System Sciences* (Hicss 03), vol. 5, CD ROM, IEEE CS Press, 2003, p. 127a.