

Peer-to-Peer Human-Robot Interaction for Space Exploration

Terrence Fong and Illah Nourbakhsh

Computational Sciences Division

NASA Ames Research Center

Moffett Field, CA 94035 USA

{terry, illah} @email.arc.nasa.gov

Abstract

NASA has embarked on a long-term program to develop human-robot systems for sustained, affordable space exploration. To support this mission, we are working to improve human-robot interaction and performance on planetary surfaces. Rather than building robots that function as glorified tools, our focus is to enable humans and robots to work as partners and peers. In this paper, we describe our approach, which includes contextual dialogue, cognitive modeling, and metrics-based field testing.

Motivation

In January 2004, NASA established a long-term program to extend human presence across the solar system, a primary goal of which will be to establish a human presence on the moon no later than 2020, as a precursor to manned exploration of Mars (NASA 2004). In the past, NASA's human space flight programs and robotic exploration programs operated largely independently of each other. With this new program, however, significant emphasis is being placed on the development of joint human-robot systems:

NASA will send human and robotic explorers as partners, leveraging the capabilities of each where most useful. Robotic explorers will visit new worlds first, to obtain scientific data, assess risks to our astronauts, demonstrate breakthrough technologies...Human explorers will follow to conduct in-depth research, direct and upgrade advanced robotic explorers, etc.

—*NASA Vision for Space Exploration*

Moreover, a central concept of this program is that exploration activities must be sustainable over the long-term. Sustained exploration, however, will require a number of tasks to be performed daily, which will far exceed the human resources that can be sent into space. Thus, to address this problem, as well as to reduce human workload, costs, and fatigue-driven error and risk, robots will have to be an integral part of mission design.

Although robots have previously been used for scientific purposes in space (e.g., geology), in the context of this new vision, robots will also be called upon for non-scientific

work. In particular, robots will have to perform a multitude of non-science tasks. The intricate nature of these tasks will require the use of multiple levels of control and adjustable autonomy.

Considerable research has already focused on developing human and robotic systems for planetary surfaces. Scant attention, however, has been paid to joint human-robot teams. Such teams, however, are attractive for numerous applications (Cabrol 1999; Huntsberger, Rodriguez, and Schenker 2000; Jones and Rock 2002). Robots could support long-duration excursions by scouting, by surveying, and by carrying equipment. Robots could assist in site preparation, sample collection and transport, as well as provide contingency life support. Finally, robots could be used for field labor, assembly and maintenance of spacecraft and habitats.

Assuming that human cognitive, perceptual, and physical capabilities will not change significantly over the next few decades, a viable strategy is for humans to perform higher-order cognitive and perception functions while robots perform reactive, precise, and physically demanding functions (Hansen and Ford 2004). Of course, choosing the optimum distribution of responsibilities between human and robots is a difficult problem that must be addressed (Proud, Hard, and Mrozinski 2003; Rodriguez and Weisbin 2003).

However, since robots will not be 100% self-sufficient (certainly not before 2020), situations will continue to arise in which the robot fails and a human needs to intervene (e.g., via teleoperation). This is particularly true whenever the robot's autonomy is ill-suited for the task or when the robot is faced with unexpected contingencies. But, because teleoperation is not a panacea (e.g., it imposes high operator workload and requires high bandwidth communication), systems that can synergistically exploit the different strengths and capabilities of humans and robots are needed to ensure that acceptable performance is achieved.

We claim, therefore, that making human-robot interaction (HRI) effective, efficient, and natural is crucial to future space exploration. In particular, we contend that humans and robots must be able to: (1) communicate clearly about their goals, abilities, plans, and achievements (Fong 2001; Jones and Rock 2002); (2) collaborate to solve problems, especially when situations exceed autonomous capabilities (Fong 2001); and (3) interact via multiple modalities (dialogue, gestures, etc), both locally and remotely.

Objectives

Our goal is to develop HRI techniques that facilitate and support sustained, affordable human-robotic space exploration. In particular, we want to create robots that can communicate and collaborate with humans. We want to develop modular, task-oriented robots that can be interchanged, work side-by-side, and cooperatively with humans. We want robots that can interact with humans directly (i.e., in the field) or remotely (over short and long distances).

Rather than building robots that function as tools (regardless of how autonomous they might be), our work centers on *semi-autonomous robots* that operate as *peers* and that interact in a competent, natural (i.e., non-disturbing) manner. To do this, we are investigating techniques to support *peer-to-peer HRI* and methods that enable robots to collaborate with users who have little training, prior experience, or knowledge of robotics.

We agree with Cassimatis et al. (2003) that computational cognitive architectures can facilitate more natural, more productive HRI. Specifically, we believe that cognitive models can enable robots to better communicate with humans: recognizing what context and level of detail to provide, understanding when and how to interrupt the human, etc.

Consequently, we have begun addressing the following questions:

- How can models of human-based representations and human communication be used for interaction design and control (pacing, focus, etc.)?
- How can cognitive models be used to improve human-robot coordination (task allocation and joint work)? In particular, how can we choose appropriate interaction models and roles for specific work scenarios?
- How can cognitive systems be used to improve human-robot dialogue, especially in terms of user modeling, level of abstraction/context, and focus of attention?

We must note, however, that cognitively plausible models, by themselves, are insufficient for developing peer-to-peer HRI. In addition, we believe it is necessary to investigate techniques for creating robot self-awareness (understanding of capabilities, self-monitoring, fault detection, etc.), human awareness (e.g., human-oriented perception), multi-modal interaction, and social competency (engagement, adherence to norms, etc).

Approach

There are three primary components in our approach. First, we are developing a human-robot system model called *collaborative control* (Fong 2001). With this model, the human and robot engage in dialogue to work as partners. A key benefit of collaborative control is that the robot is able to ask questions of the human in order to compensate for limitations, or failures, of autonomy. As part of our current work, we are investigating techniques to determine how and when it is appropriate for the robot to interrupt humans.

Second, we intend to use computational cognitive architectures to model human behavior. Of primary interest is making the human and robot understandable and predictable

to the other. We believe that by building robots with reasoning mechanisms and representations that are similar to what humans use, we can make human-robot interaction more natural and human-like. Cognitive models also offer the potential for adaptation, i.e., allowing robots to adapt to the user as well as to predict user behavior.

Finally, we plan to conduct a regular series of evaluations (development and field tests) using relevant rover platforms and space exploration tasks, analog environments and quantitative performance metrics. We will perform detailed workflow and critical incident analysis to understand the impact of peer-to-peer human-robot interaction on tasks, to identify failure modes and to learn how to improve execution. We will place significant emphasis on assessing system performance, human performance, and robot performance.

Robot as Partner

Since 2000, we have been developing a new human-robot interaction model, collaborative control, which is based on human-robot dialogue. With collaborative control, the robot asks questions to the human in order to obtain assistance with tasks such as cognition and perception. This enables the human to function as a resource for the robot and to help compensate for the robot's limitations. Our initial results indicate that collaborative control is a highly effective paradigm for constructing and operating human-robot teams (Fong, Thorpe, and Baur 2003).

To improve collaborative control, we are developing a dialogue management system that will enable a robot: to formulate better queries than our current system (using context to disambiguate and provide detail); to understand when it is appropriate to ask questions (to avoid annoying or overwhelming the user); and to benefit from human interaction (learning and adapting to the quality of the human's responses). Our approach focuses on applying an agent-based, collaborative dialogue method to human-robot interaction.

A key aspect of this new dialogue system will be collaborative grounding, in which human and robot partners are able to establish, through multi-modal dialogue, references to the same objects in their shared, possibly remote workspace. In our research, we assume that the robot holds subjective, partial information about the environment that it exchanges with the human. We plan to use explanation and learning techniques to establish common (mutual) belief between the human and the robot. This is particularly important for helping the human acquire and maintain situational awareness.

In addition, we will investigate techniques to determine how and when it is appropriate for the robot to interrupt the human. This is needed because user interface mediated HRI lacks many of the pacing and control cues (body position, non-verbal language, etc.) that are used during face-to-face human-human interaction. Moreover, interruption is problematic because humans have cognitive limitations that restrict their ability to work during interruptions and to resume previously interrupted tasks (McFarlane 1990).

Computational Cognitive Models

In order to take the best possible advantage of the particular skills of humans and of robots in mixed-initiative teams, it

is important that robots be able to reason about how to communicate with humans for maximum effect. We believe the best way to accomplish this is to study how humans solve specific tasks, then build computational cognitive models of this ability as mechanisms for robot reasoning.

In particular, we contend that skills such as *perspective taking*, the ability of a person to take someone else's perspective, is critical to collaboration, and is therefore essential for robots collaborating with humans. Thus, we intend to use models of human perspective taking, such as in (Sofge *et al.* 2004), to improve HRI in situations where humans and robots are working side-by-side.

Moreover, we are interested in using cognitive architectures for modeling behavioral outcomes in HRI. That is, we seek to predict behavior that is part of, and which results from, interaction between humans and robots. Cognitive architectures have traditionally been of two types: process (or internal computation) models and mathematical (or behavioral) models (Sun and Ling 1998). For HRI, we believe a combination of the two approaches is needed.

For example, we are considering adapting the NASA Ames "Man Machine Integrated Design and Analysis System" (MIDAS) to robotics. MIDAS consists of an agent-based operator model (modules represent cognitive and sensor-motor functions) and models of the proximal (displays, controls, etc.) and remote environments (Smith and Tyler 1997). It has been extensively used for human performance modeling and for predicting human-system effectiveness. As such, we believe MIDAS would be useful for improving HRI design, coordination, and dialogue.

Metrics and Field Tests

To assess our progress, we are planning to conduct controlled field-experiments and usability studies. These tests will be designed to evaluate the capabilities, strengths and limitations of our tools and technologies in a structured manner. In addition, we will perform detailed workflow analysis to understand the impact peer-to-peer HRI has on these tasks, to identify failure modes and to learn how to improve our design.

Evaluating how well a human-robot team works is difficult and there is currently no universal standard by which we can determine the absolute "goodness" or capability of HRI. At the same time, however, quantitative measurements are crucial for understanding how well humans and robots are able to work together. Thus, we intend to use a variety of metrics and critical incident analysis to characterize and assess human-robot performance (Fong *et al.* 2004; Scholtz *et al.* 2004).

To evaluate system performance, one metric we will examine is *fan out*, which measures how many robots can be effectively controlled by a human (Goodrich and Olsen 2003). To assess operator performance, we will use the NASA-Task Load Index (NASA-TLX) (Hart and Staveland 1988) and will measure the level of situational awareness (Scholtz 2003). To evaluate robot performance, one metric we will apply is self-awareness (i.e., how well a robot can report on its current health, state, etc.) because it is indicative of interaction efficiency.

Our field tests are intended to assess not only human-robot team performance, but also its dependence on the roles and relationships of humans and robots. The robot role will range from near full-time teleoperation (human in manual control of robot) to near full-time autonomous (on-board task executive). The human-robot spatial relationship will range from proximal (side-by-side) interaction to long-distance, earth-lunar operations communication.

Initial Experiments

Our initial experiments (planned for 2005) will study the tradeoff in efficiency with a human-robot team conducting an in-flight inspection task and a surface assembly task while varying the roles of the two participants. We will evaluate performance under a variety of operational scenarios. The same underlying HRI system will be used for both tasks, in order to assess the flexibility of our approach.

The inspection task will be to search for and identify loose fixtures in an indoor structure analogous to the inside of the International Space Station or the planned Crew Exploration Vehicle. In the first operational scenario, the robot will serve as a videographer: sending live video to a remote human. The human will perform the perception required to find fixtures and verify that they are secure. In the second operational scenario, the robot will search for and return still images of the fixtures. The human will focus on examining the fixtures. In a third operational scenario, the robot will only return images of the fixtures that it has identified as insecure. This will allow the human to focus only on inspecting suspicious fixtures.

The surface assembly task will be for the human-robot team to transport a set of modular components to a site for assembly, placing them in approximately the correct locations for assembly. In the first operational scenario, a rover with manipulators will be teleoperated to repeatedly navigate to the component storage area, pick up an item, navigate to the construction site, and place the piece in a position, which facilitates assembly. In the second operational scenario, the human will teleoperate the manipulators during pick and place, and the rover will navigate between sites.

Open Issues

Given the potential of human-robot teams to improve space exploration, it is clear that we need to understand how such systems can best be used. A first step in this direction was taken in 1999, when NASA ARC and JSC conducted the "Astronaut-Rover Interaction Field Test" (ASRO) (Cabrol 1999). ASRO identified several operational scenarios in which EVA crewmembers and rovers can interact in a safe, productive manner. The study, however, focused only on surface exploration scenarios (i.e., geologic survey). Further research is needed to identify other space exploration tasks that can benefit from human-robot teams.

In addition, a significant issue, which needs to be addressed, is how to decide which human and robotic assets are most appropriate for a given task or mission (Rodriguez and Weisbin 2003). Often it is possible to perform tasks with robots or humans or both, but we need to understand (espe-

cially in quantitative terms) the consequences of different allocation strategies. In particular, if we “swap” a human for a robot (and vice versa), how do we distribute responsibilities and work? How do we identify appropriate roles for the human and robot? How do we decide what is “optimal”?

Finally, in order for humans and robots to work side-by-side on planetary surfaces, we need to develop reliable, robust user interfaces for field use. One approach would be to use visual gesturing (hand, arm, and body movements) for robot control. The primary challenge is to make gesture recognition sufficiently robust to function in a wide range of outdoor scenes and over a wide range of distances. Another approach would be to develop handheld or wearable devices that support human-robot communication. These interfaces would enable higher bandwidth and higher quality interaction than can be achieved through gesturing alone.

References

- Cabrol, N. 1999. Proceedings of the Astronaut-Rover interaction post Silver Lake field test debriefing: results, lessons learned, and directions for future human exploration of planetary surfaces. Technical Report NP-2004-01-334-HQ, NASA Ames Research Center, Moffett Field, CA.
- Fong, T., Kaber, D., Lewis, M., Scholtz, J., Schultz, A., and Steinfeld, A. 2004. Common metrics for human-robot interaction. in submission.
- Fong, T., Thorpe, C., and Baur, C. 2003. Multi-robot remote driving with collaborative control. *IEEE Transactions on Industrial Electronics* 50(4).
- Fong, T. 2001. *Collaborative control: a robot-centric model for vehicle teleoperation*. Ph.D. Dissertation, Carnegie Mellon University, Pittsburgh, PA. CMU-RI-TR-01-34.
- Goodrich, M., and Olsen, D. 2003. Seven principles of efficient human robot interaction. In *Proc. IEEE International Conference on Systems, Man and Cybernetics*, 3943–3948.
- Hansen, R., and Ford, K. 2004. A human centered vision of Mars exploration. Technical Report NAG9-1386, Institute for Human and Machine Cognition, University of West Florida.
- Hart, S., and Staveland, L. 1988. Development of NASA-TLX (Task Load Index): results of empirical and theoretical research. In Hancock, P., and Meshkati, N., eds., *Human Mental Workload*. North-Holland Elsevier Science.
- Huntsberger, T., Rodriguez, G., and Schenker, P. 2000. Robotics challenges for robotic and human Mars exploration. In *Proc. Robotics*.
- Jones, H., and Rock, S. 2002. Dialogue-based human-robot interaction for space construction teams. In *Proc. IEEE Aerospace Conference*.
- McFarlane, D. 1990. Coordinating the interruption of people in human-computer interaction. In *Proc. HCI Interact*.
- NASA. 2004. The vision for space exploration. Technical Report NP-2004-01-334-HQ, NASA, Washington, DC.
- Proud, R., Hard, J., and Mrozinski, R. 2003. Methods for determining the level of autonomy to design into a human spaceflight vehicle: a function specific approach. In *Proc. NIST Performance Metrics for Intelligent Systems*.
- Rodriguez, G., and Weisbin, C. 2003. A new method to evaluate human-robot system performance. *Autonomous Robots* 14(2):165–178.
- Scholtz, J., Young, J., Drury, J., and Yanco, H. 2004. Evaluation of human-robot interaction awareness in search and rescue. In *Proc. IEEE International Conference on Robotics and Automation*.
- Scholtz, J. 2003. Theory and evaluation of human robot interactions. In *Proc. Hawaii International Conference on System Science* 36.
- Smith, B., and Tyler, S. 1997. The design and application of MIDAS: a constructive simulation for human-system analysis. In *Proc. Simulation Technology and Training (SIMTECT)*.
- Sofge, D., Perzanowski, D., Skubic, M., Bugajska, M., Trafton, J., Cassimatis, N., Brock, D., Adams, W., and Schultz, A. 2004. Cognitive tools for humanoid robots in space. In *Proc. IFAC Symposium on Automatic Control in Aerospace*.
- Sun, R., and Ling, C. 1998. Computational cognitive modeling, the source of power, and other related issues. *AI Magazine* (Summer 1998).