Intelligent assistance for human supervisory control of multiple robots

Robin R. Murphy† and Erika Rogers‡

†Dept. of Mathematical and Computer Sciences
Colorado School of Mines
Golden, CO 80401

‡Dept of Computer Science
California Polytechnic State University
San Luis Obispo, CA 93407

ABSTRACT

This paper reviews on-going collaborative efforts between the Colorado School of Mines and Clark-Atlanta University in cooperative assistance for coordination and control of multiple vehicles. It reports on progress in developing an intelligent assistance agent (IAA) for aiding a human operator in diagnosing problems and generating recovery strategies in remote ground robots. The current work has focused on the identification and incorporation of categories of additional information from multiple robots and other agents. These categories are: *mission-related sources*, such as peer robots working nearby; *facility-related sources*, such as security cameras; and *opportunistically available agents*, such as overhead satellites or humans working in the area as part of another mission. The incorporation of additional sources of information requires enhancements to the previously developed teleVIA architecture. In particular, the teleVIA IAA must provide more strategic management support (reminding operator of resources, impact of reallocation of resources on mission, etc.), sophisticated viewpoint and data management and presentation, and simplified control of the additional agents for diagnosis and recovery activities (move to robot and examine, corroborate sensor X, etc.) These enhancements are encapsulated in software agents within the IAA.

Keywords: mobile robots, multi-sensor integration, intelligent assistance, intelligent agents, multi-agents

1. INTRODUCTION

Multiple ground robots are needed for a variety of emerging domains: Urban Search and Rescue (USAR), tactical reconnaissance, planetary exploration, and treaty verification. These applications will require robots to operate in a partially known environment with high likelihood of surprises, often under time pressures, and with limited communication. The robots must be robust, both individually and as a team.

Due to the underconstrained environment and need for robustness, we consider a semi-autonomous control paradigm as an acceptable evolutionary intermediate step to full autonomy. Under the semi-autonomous paradigm, each robot is given a set of tasks to perform autonomously, while a human is responsible for planning, task verification, and most importantly, troubleshooting. If a robot runs into trouble, the human can quickly assess the situation and make the appropriate corrections. An example of this occurred during the Mars Pathfinder mission, when mission controllers used the lander camera to determine that Sojourner's obstacle avoidance sensors were malfunctioning, preventing Sojourner from moving ahead to its next scientific target, the rock "Shark." (Sojourner was essentially hallucinating that an unsurmountable obstacle was directly in its way, and began going in circles in an attempt to avoid it.) The mission controllers overcame the problem by temporarily turning off the obstacle avoidance system. Sojourner then climbed over the rock "Wedge" and made its way successfully to Shark.

However, relying on a human to supplement on-board diagnostics is not necessarily a satisfactory solution. Diagnosis and recovery is hard for a human to perform because of the limited sensing modalities and unnatural viewpoints. Another problem is cognitive fatigue. These problems become exacerbated by the presence of additional sources of information (teams of peer robots, surveillance cameras, unmanned aerial vehicles, etc.) which may have different sensing modalities and different viewpoints. Furthermore, ground robots may be involved in time-critical

Email: rmurphy@mines.edu (after Aug. 3, 1998: murphy@csee.usf.edu), erogers@csc.calpoly.edu

missions, where a human must immediately assess the situation and generate a recovery method, without impacting the overall mission or requiring extensive communication. Our current efforts are concentrating on the development of an *intelligent assistance agent* which can aid the human operator in identifying and managing supplemental resources without endangering the overall mission objectives.

This paper first describes the types of additional information sources available to an unmanned ground mobile robot, and the potential impact of that information on the operator. Next, it presents an extension of our work with *teleVIA*, an intelligent assistance architecture for mobile robots, which attempts to mitigate the negative impacts. Although *teleVIA* is novel, we discuss how it is based on concepts from HCI, semi-autonomous control and web-based agents. Finally, the paper reports on the current status of our research.

2. IMPACT OF ADDITIONAL INFORMATION SOURCES

Our experiences with teams of heterogeneous multiple robots has suggested that robots can serve as additional sources of information. We have also identified other resources which the operator can exploit and control. However, as we illustrate below, more information does not always make the human operator's task easier.

2.1. Sources of Additional Information

The human operator may be able to draw upon one or more of three distinct types of information resources: mission related information sources, facility related sources, and opportunistic sources. Each category of source has different implications for human control.

Mission-related sources of information are agents directly involved in accomplishing the mission objectives. These agents are working with the ground robot and may be redeployed to acquire additional information, or effect a cooperative recovery solution. An example of mission related sources is a team of tactical ground robots performing building clearing. Each robot is a potential source of information about a failed peer robot.

Mission-related sources can be either loosely or tightly coupled to the failed robot. A *loosely coupled team* is one where the agents are not necessarily working in the same proximity, such as teams of tactical ground robots or UGV/UAV teams. In this case, the human operator will have to divert a peer robot from its current task, navigate it to the failed robot, and observe the situation. This is shown in Fig. 1a. An outdoor mobile robot, C2, has been damaged by brushing up against a tree. The tree limb has knocked off the camera from the panning mast, but the camera is still transmitting images. From the images, the human operator cannot tell if the camera is damaged and sideways or if the entire vehicle is overturned. The operator temporarily suspends the activities of another outdoor robot, Silver Bullet, and sends it to investigate. With the additional information source of Silver Bullet, it is straightforward for the operator to determine the vehicle is intact, but the camera is dismounted. Note that the human operator must consciously identify that there is an outside source of information, and decide that it is worth reallocating that resource.

A *tightly coupled team* is one where the agents are working together already. An example is shown in Fig. 1b, where Silver Bullet and Bujold are working cooperatively for a simulated Urban Search and Rescue task. Bujold is an industrial inspection robot with great mobility over short ranges of rough terrain. She is powered and controlled by Silver Bullet, another autonomous mobile robot, which literally carries Bujold to an interesting site, then deploys her. Bujold can be tracked and controlled by Silver Bullet or by a human operator teleoperating Bujold through Silver Bullet. Silver Bullet is an explicit source of information about the state of Bujold as they work together. However, Fig. 2 shows that Bujold, in turn, can provide information about Silver Bullet. Because of the tight coupling, the issues for the operator in identifying the additional source of information and deploying it are less problematic than for loosely coupled teams.

One or more robots may be operating in a facility which has implicit sensing resources, such as overhead cameras. These resources are referred to as *facility-related sources*. They usually are available directly to the operator, not to the robots themselves. Examples of this category of resource are nuclear power plants which often have overhead cameras, and security cameras in buildings. Fig. 3 shows the output of a black and white security camera looking down at a mobile robot and associated gear.

One of the hardest sources of additional information for a human operator to exploit is *opportunistic sources*. These are resources which may not be explicitly attached to the mission or to the facility, but could provide important information if the human knew of its availability. Examples of opportunistic sources are satellite observers or humans near by ("I can see that from here, it looks like the robot is stuck in some mud.")

Figure 1. a) Loosely coupled, and b) tightly coupled mission-related robots providing additional information.

Figure 2. a) Small robot redeployed to check out suspect sensor. b) Robot's eye view of damage.

Figure 3. View from a security camera of robot lab.

Figure 4. One robot trying to corroborate what the failed robot is “seeing.”

2.2. Disadvantages of Additional Information for Operator

Additional information may be difficult to exploit by a human operator for several reasons. First, the incoming observations may not be necessarily helpful or be easy to extract relevant information from. The observations may be similar to the existing data, providing more of the same information, but nothing new. The opposite extreme is that the data may be too different, either in terms of viewpoint or sensing modality. The operator may have problems reconciling different viewpoints or frames of references from observers. Likewise, sensing modalities such as SAR often require expert training to interpret correctly. The operator may not be qualified to exploit the new information, and, worse yet, the availability of the data may distract him from considering more meaningful options in a timely fashion. Finally, there is not guarantee that the quality of the additional data is sufficient for the diagnostic task. Fig. 3 illustrates the visual challenges of scene interpretation with poor resolution and focus, plus an unnatural viewing angle.

Second, acquiring additional information is an active task, not passive. The operator must identify sources of additional information, decide if they are suitable and can be deployed without risking the overall mission objectives, and then deploy those resources. Clearly, the operator must have a situational awareness in order to make the correct decision. The operator must also be aware of how those resources are controlled (e.g., navigate the peer robot over to the failed robot) and how to exploit them. Fig. 4 shows a peer robot looking in the same direction as a failed robot thinks it is looking at. Since the failed robot has a camera pan/tilt fault, it is not looking at the wheelbarrow; the operator can compare images from the robots to ascertain that something has gone wrong with the camera positioning system. But this type of corroboration scenario requires the operator to know the similarities between robots. This can be very difficult to infer with heterogeneous robots, where corroboration may be indirect. Also, the process of aligning the two robots for corroboration may be challenging for teleoperation; ideally the operator could give a simple command to deploy the peer robot for corroboration.

A third challenge for an operator is identifying novel recovery options. Not only can additional agents be deployed to aid in diagnosis, they may be part of the solution. Fig. 5a shows a scenario where the robot on the right has had a camera failure. The robot cannot complete its mission without cameras and cannot safely navigate. However, it can use sonars to follow a peer robot, which can navigate, back to a safe place. Fig. 5b shows a more fanciful scenario, where one robot might physically rescue a failed agent.

2.3. Operator Assistance

In addition to the explicit challenges associated with acquiring, exploiting and managing additional information, the operator may have to overcome these challenges under unfavorable conditions. These unfavorable conditions include time pressure and the pressure associated with the mission context (importance, penalty of mission failure, risk to humans and/or hardware). Also, the robot may display a unique failure, and the human operator may have no clear sense of how to handle it. Likewise, the human operator may be a novice, or unfamiliar with the mission.

Therefore, we believe that an intelligent system should assist the operator by providing:

Figure 5. Rescue as a recovery method: a) visually impaired robot using sonars to follow a rescue robot to safety. b) rescue robot acting as a tow truck.

- Contextual and domain knowledge about the mission
- Insight into robot's "mind:" what is going wrong? what has it tried for recovery?
- Meta-knowledge about the robot's behaviors and capabilities
- Cognitive support for the human's perceptual and problem solving activities: viewpoint/data management, "reminders"
- Operationalization of diagnosis and recovery commands, and subsequent autonomous execution

3. TELEVIEW: INTELLIGENT ASSISTANCE AGENT

Our current work dedicated to developing and implementing an intelligent assistance agent (IAA) to satisfy the needs of an operator in exploiting and managing additional sources of information. In particular, the IAA is intended to support human supervisory control of multiple robots and information sources for 1) information gathering and 2) recovery, including rescue of disabled robot.

In general, the IAA can be viewed as performing two roles. In one role it will serve as an expert consultant as to the availability, applicability, and translation/display of all sources of information (failed robot plus additional sources) for the human. In the other role it will act as the manager of the additional sources of information and translate high level human commands for actively deploying those resources. the IAA is responsible for operationalizing requests from the human for information about robot tasks: for example, the human as part of a diagnostic routine may wish to test two sensors to see if they return the same value. This can be a complex activity, where sensors and their effectors are deallocated from the current task, and servo-ed to a corroborating position. The human implicitly assumes that the robot will default to a fail-safe set of behaviors while the diagnostic process continues; the IAA must take this implicit assumption and make it explicit to the robot. The IAA must also provide an feedback or confirmation to the human that the task is being carried out. The IAA may also gather information from any directly connected sensors, e.g., surveillance cameras, observing robot, etc. This requires the IAA to instantiate the behaviors needed for active perception.

The IAA architecture is novel in that it treats the remote robots and human as computational agents possessing unique knowledge and intelligence. It relies on a third computational agent called the intelligent assistant to act as an intermediary between the human and the robot. This agent resides on the local system; it doesn't move and it doesn't perceive. Rather, it supports the perception and problem-solving capabilities of the human and the robot by selectively filtering and enhancing perceptual data obtained from the robot, as well as generating hypotheses about execution failures which cannot be solved by the remote.

Figure 6. Overview of IAA architecture.

3.1. Overview of IAA Architecture

As shown in Fig. 6, the IAA is a middle agent. The human operator (*cognitive agent*) interfaces with the situated agents (robots, autonomous information sources) through the IAA. The IAA is not strictly a software agent: it may have direct command over facility-related resources such as security cameras.

The IAA itself consists of software agents. Interface agents are responsible for translating human and robot commands and data into a common format, and *visa versa*: they essentially handle I/O. The majority of software agents act as experts. The *external information experts* are responsible for accessing the status and suitability of any possible sensing resources. One external information agent might be dedicated to interfacing with the sensing resources directly under the command of the IAA (ex., security camera). Another might be focused on maintaining a list of resources that could be used opportunistically. The *problem solving experts* attempt to provide problem solving guidance and hypothesis maintenance based on the context. Likewise, the *perceptual display experts* attempt to filter and present the most useful data based on the mission context and the status of problem solving. The *activity history monitors* serve as a situational memory buffer: in some cases, problem solving is facilitated by understanding the overall pattern of degradation of performance, rather than just receiving data about the current, failed state of the robot. However, we have found that this type of data must also be carefully filtered and displayed according to context, or else the “data dump” becomes an impediment to the human operator.

3.2. Information Coordination and Control

Clearly, effective, context-based information coordination and control is a key component to the IAA. The IAA uses a blackboard scheme to observe and manage the information posted independently by the robotic and human intelligences. In our application of the blackboard, Each software agent has internal routines called *knowledge sources* which read and post information to a global, asynchronous data structure called the *blackboard*. The knowledge sources of the robot post their information about the status of the robot. The operator reads the status and can use the knowledge stored by the intelligent assistant about previous or related cases to generate new directives such as task plans, sensor configurations, specification of parameters, response to anomalous situations, etc. The operator, by definition a knowledge source, communicates with the intelligent assistant and the robot via a graphical interface managed by the assistant.

The blackboard is the heart of the cooperative intelligent assistant. It is where the evolutionary results of the problem-solving effort are captured. The effort is divided into three main components, or global panels, the *context*

Figure 7. Tele-VIA Blackboard.

panel provides a global representation of the current situation, the state of the interactive exception handling process (what has failed, what is the status of problem solving, what resources are available), and any reconfiguration or reallocation of those resources. The *hypothesis panel* represents what the robot believes has gone wrong (locally generated hypotheses) and any results of local testing, in addition to what the human operator believes and is pursuing. The *attention panel* is responsible for posting the displays most suitable for diagnosing the situation. It also posts the attention directives, which offer suggestions to the human operator as to the steps in verifying a particular hypothesis, the utility of deploying additional information sources, etc. This logical partitioning is based on components of a cognitive model of visual interaction described in,¹ and is designed to facilitate transfer of information between human perception and problem-solving during a visual reasoning task.

The system's user interface presents two logical user views: monitor mode, which is used solely for monitoring the robot's normal behavior, and failure (or problem-solving) mode, which appears when a request for assistance is sent by the robot. All panels, when visible, are presented as interactive direct manipulation windows, with a variety of icons, text and buttons used to display the underlying frame-based information. Color-coding is used to indicate correspondences between related data distributed across several panels. A sample screen dump is shown in Fig. 8

4. RELATED WORK

Our approach to human-agent interactions in complex environments is similar in scope, but different in emphasis, to work conducted by Webber et al in translating natural language (NL) instructions into intentions and expectations,² and the use of discourse in collaboration.³ Our work is concentrating on first establishing the reasonable intentions and expectations of the human and software agents by cognitive modeling, rather than on language issues.

Our characterization of intelligent agents is similar to the Belief, Desire and Intention (BDI) framework,^{4,2} especially in acknowledging that defining the informational, motivational, and cognitive state of the agent is key in both designing a software agent in general and in designing a software agent to supplement these states in humans and other agents. However, we believe that it is premature to attempt to represent these states through formal logic, even with simplifying assumptions as per,⁴ without a fundamental understanding of the cognitive process.

Our architectural organization is most similar to Decker et al,⁵ who have constructed a web-agent architecture consisting of interface agents, task agents, and information agents. Our IAA can be viewed as a middle-agent,⁶ containing these agents, where the IAA uses visual problem solving to control the interface to the human, task knowledge (intentions and expectations) is triggered by the current context, and information agents provide broad problem solving resources. The IAA is broader, where interface and information agents exist for facilitating interactions with the robotic and other software agents as well as with the human agent. The IAA serves as a middle-agent because it is not strictly a provider or a requestor, but acts as a go-between the human and other agents in the system. As with,⁶ the IAA exploits preferences and capabilities.

6. SUMMARY AND FUTURE WORK

This article has presented work in progress on developing an intelligent assistance agent (IAA) for aiding a human operator in diagnosing problems and generating recovery strategies in remote ground robots. The current work has focused on the incorporation of three types of additional information: mission-related sources, such as peer robots working nearby; facility-related sources, such as security cameras; and opportunistically available agents, such as overhead satellites or humans working in the area as part of another mission.

The incorporation of additional sources of information requires enhancements to the previously developed teleVIA architecture. In particular, the teleVIA IAA must provide more strategic management support (reminding operator of resources, impact of reallocation of resources on mission, etc.), sophisticated viewpoint and data management and presentation, and simplified control of the additional agents for diagnosis and recovery activities (move to robot and examine, corroborate sensor X, etc.) These enhancements are encapsulated in software agents within the IAA.

Current work is concentrating on cognitive modeling of the use of visually presented data for human-agent interactive problem solving, and on autonomous corroboration and recovery behaviors for remote robots. Further work is needed to represent the current state of the failed robot, its local knowledge, and to extract and quantify the informational utility of additional strategic resources for a given context.

ACKNOWLEDGMENTS

This work has been supported in part by the National Science Foundation under grant CDA-9617309, and the CRA Distributed Mentoring Program. The authors would like to thank Pavani Reddy for her contribution to interface facility-related resources to teleVIA. Ann Brigante, Damian Diaz, and Charlie Ozinga helped with the robot scenarios.

REFERENCES

1. E. Rogers, "A cognitive theory of visual interaction," in *Diagrammatic Reasoning: Computational and Cognitive Perspectives*, B. Chandrasekaran, J. Glasgow, and N. H. Narayanan, eds., pp. 481-500, AAAI/MIT Press, (Menlo Park, CA), 1995.
2. B. Webber, N. Badler, B. DiEugenio, C. Geib, L. Levison, and M. Moore, "Instructions, intentions and expectations," in *Computational Theories of Interaction and Agency*, P. Agre and S. Rosenschein, eds., pp. 619-635, MIT Press, 1996.
3. C. Rich and C. Sidner, "Collagen: When agents collaborate with people," in *Proc. of First Int. Conference on Autonomous Agents*, pp. 284-291, ACM Press, 1997.
4. A. Rao and M. Georgeff, "Bdi agents: From theory to practice," in *Proc. of First Int. Conference on Multi-Agent Systems*, V. Lesser, ed., pp. 312-319, AAAI Press, 1995.
5. K. Decker, A. Pannu, K. Sycara, and M. Williamson, "Designing behaviors for information agents," in *Proc. of First Int. Conference on Autonomous Agents*, pp. 404-412, ACM Press, 1997.
6. K. Decker, K. Sycara, and M. Williamson, "Middle-agents for the internet," in *Proc. of Int. Joint Conference on Artificial Intelligence (IJCAI)*, pp. 578-583, 1997.
7. G. R. Edwards, R. H. Burnard, W. L. Bewley, and B. L. Bullock, "The ground vehicle manager's associate," Tech. Rep. AIAA-94-1249-CP, AIAA, 1994.
8. G. K. H. Pang and H. C. Shen, "Intelligent control of an autonomous mobile robot in a hazardous material spill accident - a blackboard structure approach," *Robotics and Autonomous Systems* 6, pp. 351-365, 1990.
9. E. Rogers, R. R. Murphy, and B. Ericson, "Agent-based expert assistance for visual problem solving," in *Proc. of 1st International Conference on Autonomous Agents*, pp. 156-163, ACM Press, 1997.