We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,900
Open access books available

186,000
International authors and editors

200M
Downloads

154
Countries delivered to

Our authors are among the
TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Supporting Complex Robot Behaviors with Simple Interaction Tools

David J. Bruemmer[1], David I. Gertman[1], Curtis W. Nielsen[1] ,
Douglas A. Few[2] and William D. Smart[2]
[1]Idaho National Laboratory Idaho Falls, [2]Washington University in St. Louis
U. S. A.

1. Introduction

This chapter examines the potential for new mixed-initiative interaction methods and tools to change the way people think about and use robot behaviors. Although new sensors, processing algorithms, and mobility platforms continue to emerge, a remarkable observation is that for the vast majority of fielded systems, the basic inputs and outputs used to communicate between humans and robots have not changed dramatically since mobile robots were used in surprising numbers during World War II. Photographs and old books from the period show that small tracked unmanned ground vehicles were teleoperated for a variety of military missions. Today, video remains the predominant mechanism for providing situation awareness from the robot to the human. Direct control of translational and rotational velocity remains the predominant input from the human to the robot.

The problem is not that researchers have failed to create behaviors and autonomy. A preponderance of semi-autonomous capabilities are now available on laboratory robots (Arkin, 1997; Desai & Yanco, 2005; Maxwell et al., 2004). However, the resulting interaction is often complex and the robot's actions may seem mysterious to robot operators. Autonomy may actually increase the complexity of the task rather than decrease it. Without the right interaction metaphor, users do not know what to expect from their mysterious and complex robot "peers." As behavioral autonomy increases, the answer to the question: "What does it do?" may become increasingly difficult to answer. This is especially true for behaviors that are adaptive or which exhibit emergent effects and/or non-determinism. The trick then, is to make the user comfortable with these behaviors by communicating behavioral intentionality. The interface should make clear what the robot is trying to accomplish with each behavior and should provide an understanding of how the behavior will affect overall mission success. Just as the development of a windows based graphical user interface vastly augmented the impact and value of the personal computer, so likewise, an upsurge in the general utility of robots may follow closely on the heels of change in the human interface.

Until the recent past, robots did not exhibit sophisticated behaviors or take initiative to accomplish complex tasks. More often than not, the underlying theory of robot behavior was so trivial that it may have seemed unnecessary to give it much thought. An example is the basic notion that moving a joystick up and down will cause the robot to drive backwards and forwards. To someone who has already internalized this particular theory of robot

behavior, there is very little confusion. In actuality, the joystick has proved for many applications to be a valuable and effective interaction metaphor. The problem is that it is only appropriate for more direct levels of human-robot interaction. As the level of robot initiative and autonomy increases, the underlying metaphor for interaction must also change, resulting in a need for new and, at times, more sophisticated theories of robot behavior. The goal should not be to simplify the actual robot behaviors – complex environments and tasks require appropriately complex behaviors. Instead, the goal should be to simplify the user's mental model of robot behavior.

This chapter advocates that alongside new autonomous capabilities, the user must also be given interaction tools that communicate a mental model of the robot and the task, enabling the user to predict and understand robot behavior and initiative. The necessary "theory of robot behavior" should not be communicated by providing users with a developer's perspective of the system. Behavioral sciences research indicates that passing on system complexity to the user will increase stress and failure rates while decreasing task efficiency (Gertman et al., 2005). Quite to the contrary, a great deal of craft may be required to filter and fuse the system data, abstracting away from the complexity to support a higher-level, functional understanding. The hard question is how exactly to accomplish this in a principled fashion.

This chapter considers various components of interaction complexity common in human-robot interfaces. With results from several previous HRI experiments, the paper examines means employed to reduce this complexity. These include a) perceptual abstraction and data fusion, b) the development of a common visualization and tasking substrate, and c) the introduction of interaction tools that simplify the theory of robot behavior necessary for the operator to understand, trust and exploit robot behavior. The elements of operator trust are considered in detail with emphasis on how intelligently facilitating information exchange and human and robot initiative can build appropriate trust. New approaches to facilitating human-robot initiative are discussed within the context of several real-world task domains including: robotic event photographer, ground-air teaming for mine detection, and an indoor search and detection task.

2. Background

INL researchers began six years ago developing a suite of robotic behaviors intended to provide dynamic vehicle autonomy to support different levels of user intervention. Designed to be portable and reconfigurable, the Robot Intelligence Kernel (RIK) is now being used on more than a dozen different kinds of robots to accomplish a variety of missions within defense, security, energy, commercial and industrial contexts. RIK integrates algorithms and hardware for perception, world-modeling, adaptive communication, dynamic tasking, and behaviors for navigation, exploration, search, detection and plume mapping for a variety of hazards (i.e. explosive, radiological). Robots with RIK can avoid obstacles, plan paths through cluttered indoor and outdoor environments, search large areas, monitor their own health, find and follow humans, and recognize when anything larger than 10 cm has changed within the environment. Robots with RIK can also recognize when they are performing badly (i.e. experiencing multiple collisions or an inability to make progress towards a goal) and will ask for help.

A series of experiments was performed at the Idaho National Laboratory to assess the potential for autonomous behaviors to improve performance by reducing human error and increasing various measures of task efficiency. These experiments also illustrated the

opportunity for operator confusion regarding robot behavior and initiative. The first of these experiments showed that if operators were not able to predict robot behavior, a fight for control could emerge where the human tried to prevent or counter robot initiative, usually resulting in a significant performance decrement (Marble et al, 2003). Two groups emerged. One group understood and trusted the robot behaviors and achieved significant performance improvements over the baseline teleoperated system. The other group reported that they were confused by the robot taking the initiative and suffered a performance decrement when compared to their performance in the baseline teleoperation setting. This experiment and others like it showed that operator trust was a major factor in operational success and that operator trust was significantly impacted when the user made incorrect assumptions about robot behavior. The key question which emerged from the study was how the interface could be modified to correctly influence the user's assumptions.

Since these early experiments, a research team at the INL has been working not only to develop new and better behaviors, but, more importantly, to develop interaction tools and methods which convey a functional model of robot behavior to the user. In actuality, this understanding may be as important as the performance of the robot behaviors. Results from several experiments showed that augmenting robotic capability did not necessarily result in greater trust or enhanced operator performance (Marble et al., 2003; Bruemmer et al., 2005a; Bruemmer et al, 2005b). In fact, practitioners would prefer to use a low efficiency tool that they understand and trust than a high efficiency tool that they do not understand and do not fully trust. If this is indeed the case, then great care must be taken to explicitly design behaviors and interfaces which together promote an accurate and easily accessible understanding of robot behavior.

3. What Users Need to Know

One of the lessons learned from experimentally assessing the RIK with over a thousand human participants is that presenting the functionality of the RIK in the terms that a roboticist commonly uses (e.g. obstacle avoidance, path planning, laser-based change detection, visual follow) does little to answer the operators' basic question: "What does the robot do?" Rather than requesting a laundry list of technical capabilities, the user is asking for a fundamental understanding of what to expect – a mental model that can be used to guide expectations and input.

4. "What Does it Do?"

To introduce the challenges of human-robot interaction, we have chosen to briefly examine the development and use of a robotic event photographer developed in the Media and Machines Laboratory, in the Department of Computer Science and Engineering at Washington University in St Louis. The event photographer is an intelligent, autonomous robot designed to be used beyond the confines of laboratory in settings where training or education of the user set was nearly impossible. In this project, a mobile robot system acts as an event photographer at social events, wandering about the room, autonomously taking well-framed photographs of people (Byers et al., 2003; Byers et al., 2003; Smart, 2003). The system is implemented on an iRobot B21r mobile robot platform (see figure 1), a bright red cylindrical robot that stands about 4 feet tall. Mounted on top of the robot is a pair of stereo cameras, and a digital still camera (not shown in the figure), at roughly the eye-level of a

(short) human. The robot can rotate in place, and can also move forward and backward. The cameras can pan and tilt independently of the body.

Figure 1. iRobot B21 Intelligent Photographer

The system is completely autonomous, and all computation is performed on-board. This proved to be something of a problem from a human-robot interaction standpoint. Typically, when research robots are deployed in the real world, they are attended by a horde of graduate students. These students are there to make sure the deployment goes smoothly, to fix problems as they arise, and to physically extricate the robot from tricky situations when necessary. They also, however, act as translators and interpreters for members of the public watching the robot. The first question that most people have on seeing a robot operating in the real world is "What is it doing?" The attending graduate students can answer this question, often tailoring the explanation to the level of knowledge of the questioner.

However, since our system worked autonomously and rarely got into trouble, there were often no graduate students nearby. Members of the public had to interact with the system directly, and had to figure out what it was doing for themselves. This proved difficult, since the robot has no body language, none of the external cues that humans often have (such as camera bags), and was unable to answer questions about itself directly. Most of the time, it was impossible to tell if the robot was an event photographer, a security system, or simply wandering aimlessly.

The photographer was first deployed at SIGGRAPH 2002, the major computer graphics conference. Attendees at the conference generally have a technical background, and understand the basics of computer systems, cameras, and computation. Initially, we stationed a single graduate student near the robot, to answer questions about it, and hand out business cards. When asked about the robot, the student would generally start talking about navigation algorithms, automating the rules of photography, and face detection algorithms. While the listener understood each of these component technologies, they typically still did not understand what the robot was trying to accomplish. They lacked the

"big picture" view of the system, and interacted with it as if it was a demonstration of one of the components (face detection, for example). This led to significant unhappiness when, for example, the robot would move away from the human when they were trying to get it to detect their face. Some attendees actually stomped off angrily. They had been given a lecture on robotics capabilities when what they really needed to know was how to interact at a basic level and what to expect.

However, when we supplied the metaphor of "event photographer", the quality of the interaction was completely different. People immediately understood the larger context of the system, and were able to rationalize its behavior in these terms. When the robot moved before taking their picture, it was explained by "it's found someone else to take a picture of." People seemed much more willing to forgive the robot in these cases, and put it down to the fickleness of photographers. They were also much more willing to stand still while the robot lined up the shot, and often joked about the system being "a perfectionist." For the most part, people were instantly able to interact with the robot comfortably, with some sense that they were in control of the interaction. They were able to rationalize the robot's actions in terms of the metaphor ("it doesn't like the lighting here", "it feels crowded there"). Even if these rationalizations were wrong, it gave the humans the sense that they understood what was going on and, ultimately, made them more comfortable.

The use of the event photographer metaphor also allowed us to remove the attending graduate student, since passers-by could now describe the robot to each other. As new people came up to the exhibit, they would look at the robot for a while, and then ask someone else standing around what the robot was doing. In four words, "It's an event photographer", they were given all the context that they needed to understand the system, and to interact effectively with it. It is extremely unlikely that members of the audience would have remembered the exact technical details of the algorithms, let alone bothered to pass them on to the new arrivals. Having the right metaphor enabled the public to explain the robot to themselves, without the intervention of our graduate students. Not only is this metaphor succinct, it is easy to understand and to communicate to others. It lets the observers ascribe intentions to the system in a way that is meaningful to them, and to rationalize the behavior of the autonomous agent.

Although the use of an interaction metaphor allowed people to understand the system, it also entailed some additional expectations. The system, as implemented, did a good job of photographing people in a social setting. It was not programmed, however, for general social interactions. It did not speak or recognize speech, it did not look for social gestures (such as waving to attract attention), and it had no real sense of directly interacting with people. By describing the robot as an event photographer, we were implicitly describing it as being like a human even photographer. Human photographers, in addition to their photographic skills, do have a full complement of other social skills. Many people assumed that since we described the system as an event photographer, and since the robot did a competent job at taking pictures, that it was imbued with all the skills of human photographer. Many people waved at the robot, or spoke to it to attract its attention, and were visibly upset when it failed to respond to them. Several claimed that the robot was "ignoring them", and some even concocted an anthropomorphic reason, ascribing intent that simply wasn't there. These people invariably left the exhibit feeling dissatisfied with the experience.

Another problem with the use of a common interaction metaphor is the lack of physical cues associated with that metaphor. Human photographers raise and lower their cameras, and

have body language that indicates when a shot has been taken. The robot, of course, has none of these external signs. This led to considerable confusion among the public, since they typically assumed that the robot was taking no pictures. When asked why they thought this, they often said it was because the camera did not move, or did not move differently before and after a shot. Again, this expectation was introduced by our choice of metaphor. We solved the problem by adding a flash, which actually fired slightly after the picture was taken. This proved to be enough context to make everyone happy.

5. Developing a "Theory of Robot Behavior"

The need for human and robot to predict and understand one another's actions presents a daunting challenge. If the human has acquired a sufficient theory of robot behavior, s/he will be able to quickly and accurately predict: 1) Actions the robot will take in response to stimuli from the environment and other team members; 2) The outcome of the cumulative set of actions. The human may acquire this theory of behavior through simulated or real world training with the robot. Most likely, this theory of behavior (TORB) will be unstable at first, but become more entrenched with time. Further work with human participants is necessary to better understand the TORB development process and its effect on the task performance and user perception.

It may be helpful to consider another example where it is necessary for the human to build an understanding of an autonomous teammate. In order to work with a dog, the policeman and his canine companion must go through extensive training to build a level of expectation and trust on both sides. Police dog training begins when the dog is between 12 and 18 months old. This training initially takes more than four months, but critically, reinforcement training is continuous throughout the dog's life (Royal Canadian Police, 2002). This training is not for just the dog's benefit, but serves to educate the dog handlers to recognize and interpret the dog's movements which increase the handler's success rate in conducting task.

In our research we are not concerned with developing a formal model of robot cognition, just as a police man need not understand the mechanisms of cognition in the dog. The human must understand and predict the emergent actions of the robot, with or without an accurate notion of how intelligent processing gives rise to the resulting behavior. Many applications require the human to quickly develop an adequate TORB. One way to make this possible is to leverage the knowledge humans already possess about human behavior and other animate objects, such as pets or even video games, within our daily sphere of influence. For example, projects with humanoids and robot dogs have explored the ways in which modeling emotion in various ways can help (or hinder) the ability of a human to effectively formulate a TORB (Brooks et al. 1998). Regardless of how it is formed, an effective TORB allows humans to recognize and complement the initiative taken by robots as they operate under different levels of autonomy. It is this ability to predict and exploit the robot's initiative that will build operator proficiency and trust.

7. Components of Trust

There is no dearth of information experimental or otherwise suggesting that lack of trust in automation can lead to hesitation, poor decision making, and interference with task performance (Goodrich & Boer, 2000; Parasuraman & Riley, 1997; Lee & Moray, 1994; Kaber & Endsley, 2004). Since trust is important, how should we define it? For our purpose, trust

can be defined as a pre-commitment on the part of the operator to sanction and use a robot capability. In general, this precommittment is linked to the user's understanding of the system, acknowledgement of its value and confidence in its reliability. In other words, the user must believe that the robot has sufficient utility and reliability to warrant its use. In terms of robot behavior, trust can be measured as the user's willingness to allow the robot to accomplish tasks and address challenges using its own view of the world and understanding of the task. The behavior may be simple or complex, but always involves both input and output. To trust input, the human must believe that the robot has an appropriate understanding of the task and the environment. One method to build trust in the behavior input is to diagnose and report on robot sensor functionality. Another is to present an intelligible formatting of the robot's internal representation as in the instance of a map. Fostering appropriate distrust may be equally vital. For instance, if the robot's map of the world begins to degrade, trust in the robot's path planning should also degrade.

To trust the output, the human must believe that the robot will take action appropriate to the context of the situation. One example of how to bolster trust in this regard is to continually diagnose the robot's physical capacity through such means as monitoring battery voltage or force torque sensors on the wheels or manipulators. Although trust involves a pre-commitment, it is important to understand that trust undergoes a continual process of reevaluation based on the user's own observations and experiences. In addition to the value of self-diagnostic capabilities, the user will assess task and environment conditions, and monitor the occurrence of type I and type II errors, i.e., "false alarms" and "misses", associated with the robot's decision making.

Note that none of these components of trust require that the operator knows or has access to a full understanding of the robot system. The user does not need to understand every robot sensor to monitor behavior input; nor every actuator to trust the output. Neither does the human need to know volumes about the environment or all of the decision heuristics that the robot may be using. As behaviors become more complex, it is difficult even for developers to understand the fusion of data from many different sensors and the interplay of many independent behaviors. If an algorithmic understanding is necessary, even something as simple as why the robot turned left instead of right may require the developer to trace through a preponderance of debugging data. Put simply, the user may not have the luxury of trust through algorithmic understanding. Rather, the operator develops and maintains a relationship with the robot based on an ability to accomplish a shared goal.

8. Reducing Interaction Complexity

In his 1993 book, *Introduction to the Bootstrap*, Hans Hofmann, states that "The ability to simplify means eliminating the unnecessary so that the necessary may speak." One of the criticisms that the INL research team's early efforts received from colleagues, domain experts, practitioners and novice users was that the interface used to initiate and orchestrate behaviors was simply too complex. There were too many options, too many disparate perspectives and too many separate perceptual streams. This original interface included a video module, a map module, a camera pan – tilt – zoom module, a vehicle status window, a sensor status window and a obstruction module, to name a few.

As developers, it seemed beneficial to provide options, believing that flexibility was the key to supporting disparate users and enabling multiple missions. As new capabilities and behaviors were added, the interface options also multiplied. The interface expanded to

multiple robots and then to unmanned aerial vehicles (UAVs) and unattended ground sensors. Various application payloads were supported including chemical, radiological and explosive hazard detection capabilities. Now these all had to be represented and supported within the interface. In terms of perspectives, the interface needed to include occupancy grids, chemical and radiological plumes, explosive hazards detection, 3D range data, terrain data, building schematics, satellite imagery, real-time aerial imagery from UAVs and 3D representation of arm movement to support mobile manipulation. The critical question was how to support an ever increasing number of perceptions, actions and behaviors without increasing the complexity of the interface. To be useful in critical and hazardous environment such as countermine operations, defeat of improvised explosive devices and response to chemical, biological, radiological or nuclear hazards, the complexity of the task, particularly the underlying algorithms and continuous data streams must somehow not be passed on directly to the human.
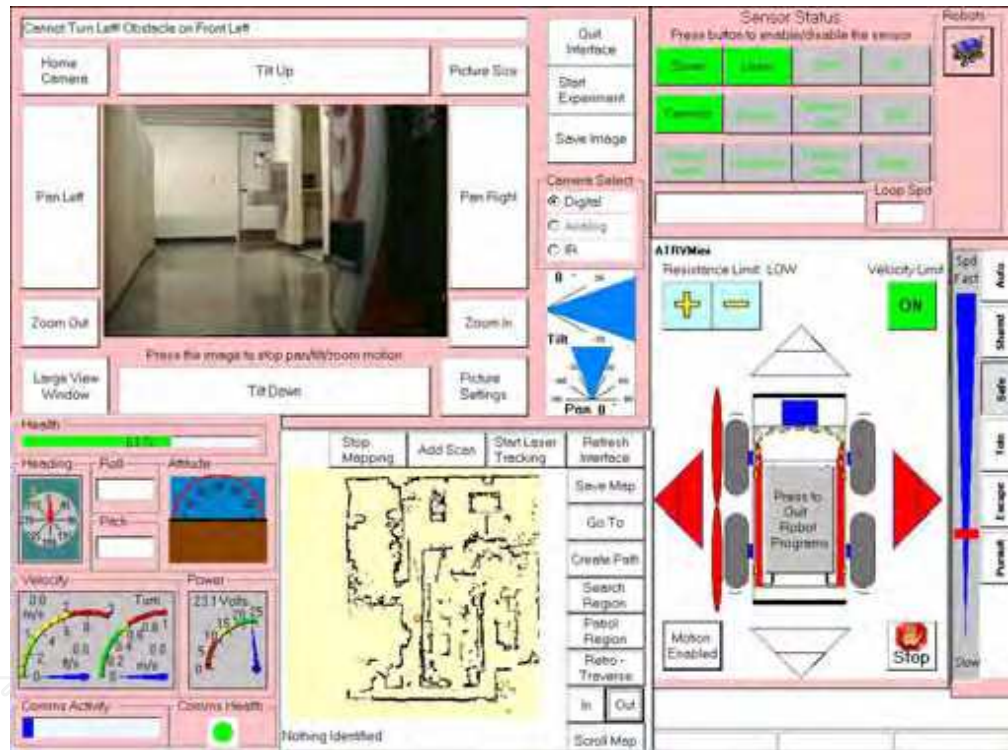

Figure 2: The Original RIK Interface

9. Data Abstraction and Perceptual Fusion

At a low level, robots must process multiple channels of chaotic, multi-dimensional sensor data that stream in from many different modalities. The user should not have to sift through this raw data or expend significant cognitive workload to correlate it. The first step to facilitating efficient human-robot interaction is to provide an efficient method to fuse and filter this data into basic abstractions. RIK provides a layer of abstraction that underlies all

robot behavior and communication. These abstractions provide elemental constructs for building intelligent behavior. One example is the ego-centric range abstraction which represents all range data in terms of an assortment of regions around the robot. Another is the directional movement abstraction which uses a variety of sensor data including attitude, resistance to motion, range data and bump sensing to decide in which directions the robot can physically move. Figure 3 shows how sensor data and robot state is abstracted into the building blocks of behavior and the fundamental outputs to the human.
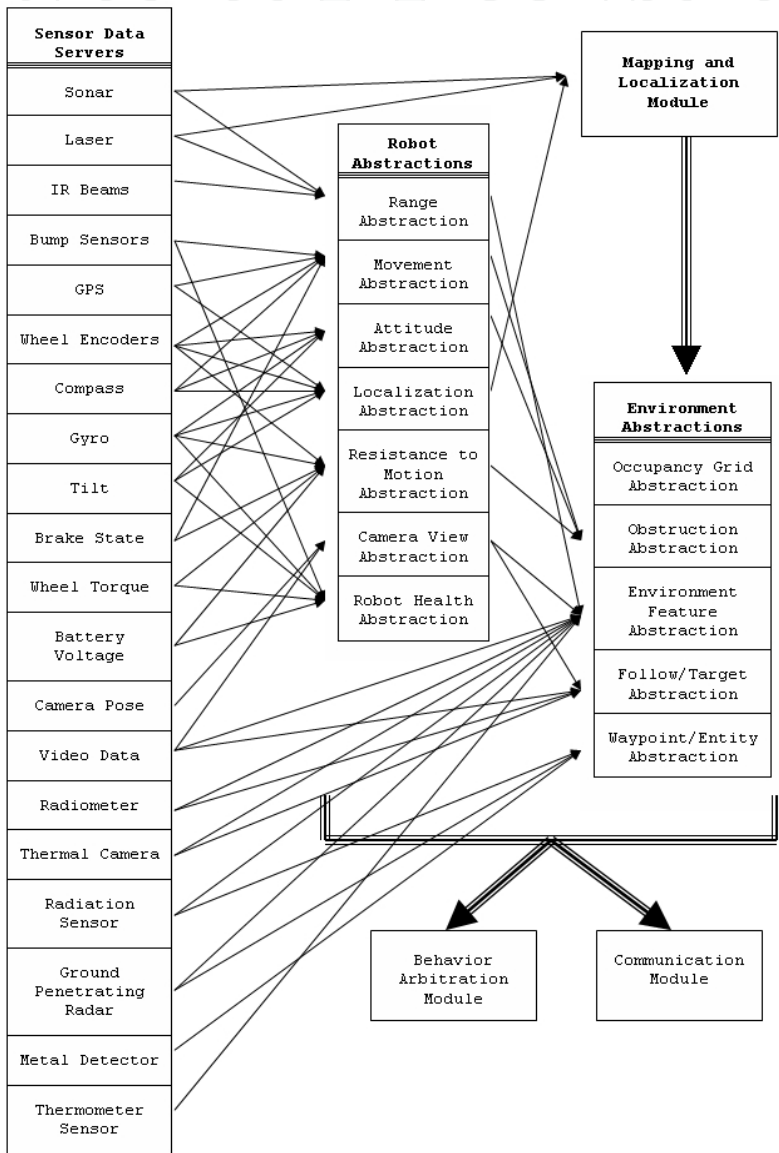
Figure 3. Robot and environment abstractions

The challenges of robot positioning provide an example of how this data fusion takes place as well as show how the abstractions are modulated by the robot to communicate information rather than data. To maintain an accurate pose, it is necessary to probabilistically fuse global positioning, simultaneous mapping and localization, inertial sensors and then correlate this with other data that might be available such as aerial imagery, a priori maps and terrain data. This data fusion and correlation should not be the burden of the human operator. Rather, the robot behaviors and interface intelligence should work hand in hand to accomplish this in a way that is transparent to the user. Figure 4 below shows how the Robot Intelligence Kernel – the suite of behaviors running on the robot fuses position information towards a consistent pose estimate.
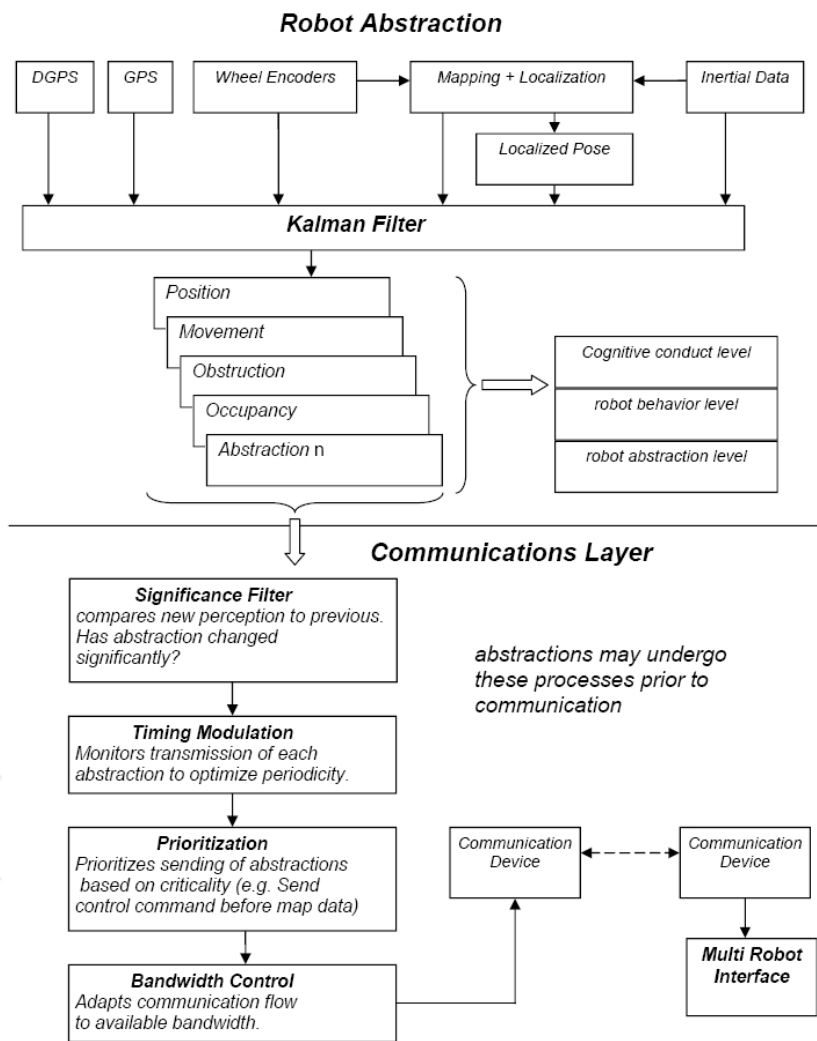

Figure 4. The use of perceptual and robot abstractions within the INL Robot Intelligence Kernel

These abstractions are not only the building blocks for robot behavior, but also serve as the fundamental atoms of communication back to the user. Rather than be bombarded with separate streams of raw data, this process of abstraction and filtering presents the user with only the useful end product. The abstractions are specifically designed to support the needs of situation awareness.
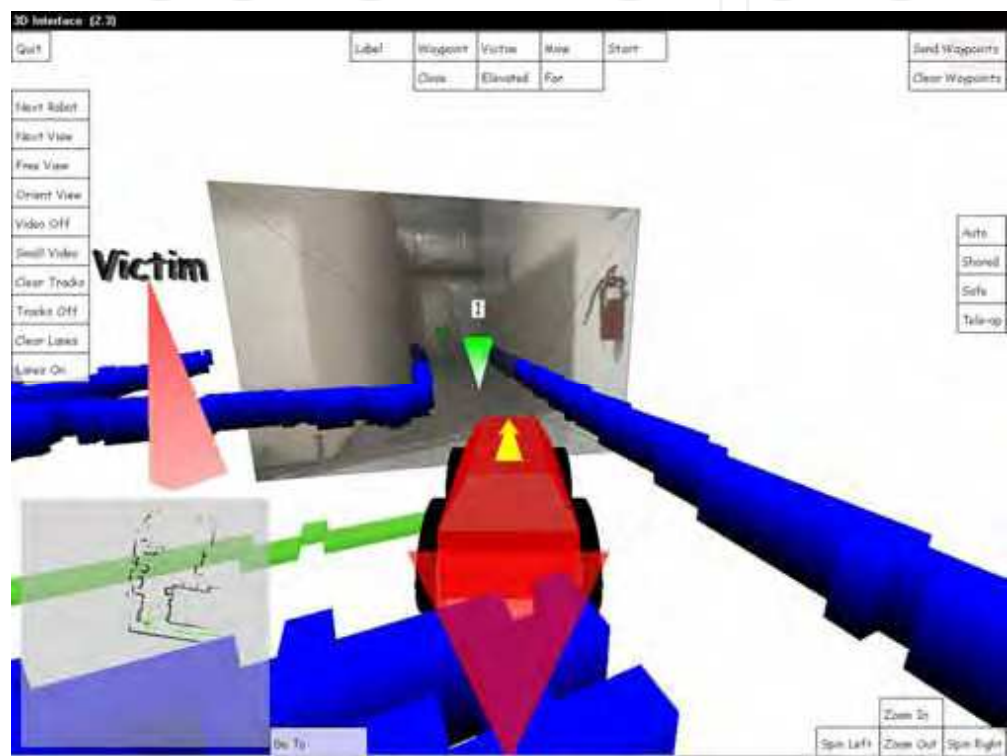
10. Fusing Disparate Perspectives

Figure 5. Interface showing fusion of video and map representation

Combining different perspectives into a common reference is another way to reduce complexity. On the interface side, efforts have been undertaken to visually render video, map data and terrain data into a seamless, scalable representation that can be zoomed in or out to support varying levels of operator involvement and the number of robots being tasked. Figure 5 below shows how video can be superimposed over the map data built up by the robot as it navigates. The Interface used with the Robot Intelligence Kernel and shown numerous times throughout this experiment was developed jointly with Brigham Young University and the Idaho National Laboratory (Nielsen & Goodrich 2006). Unlike traditional interfaces that require transmission of live video images from the ground robot to the operator, the representation used for this experiment uses a 3D, computer-game-style representation of the real world constructed on-the-fly. The digital representation is made possible by the robot implementing a map-building algorithm and transmitting the map information to the interface. To localize within this map, the RIK utilizes Consistent Pose Estimation (CPE)

developed by the Stanford Research Institute International (Konolige 1997). This method uses probabilistic reasoning to pinpoint the robot's location in the real world while incorporating new range sensor information into a high-quality occupancy grid map. When features exist in the environment to support localization, this method has been shown to provide approximately +/- 10 cm positioning accuracy even when GPS is unavailable.

Figure 6 shows how the same interface can be used to correlate aerial imagery and provide a contextual backdrop for mobile robot tasking. Note that the same interface is used in both instances, but that Figure 5 is using an endocentric view where the operator has focused the perspective on a particular area whereas Figure 6 shows an exocentric perspective where the operator is given a vantage point over a much larger area. Figure 6 is a snap shot of the interface taken during an experiment where the functionality and performance benefit of the RIK was assessed within the context of a military countermine mission to detect and mark buried landmines.
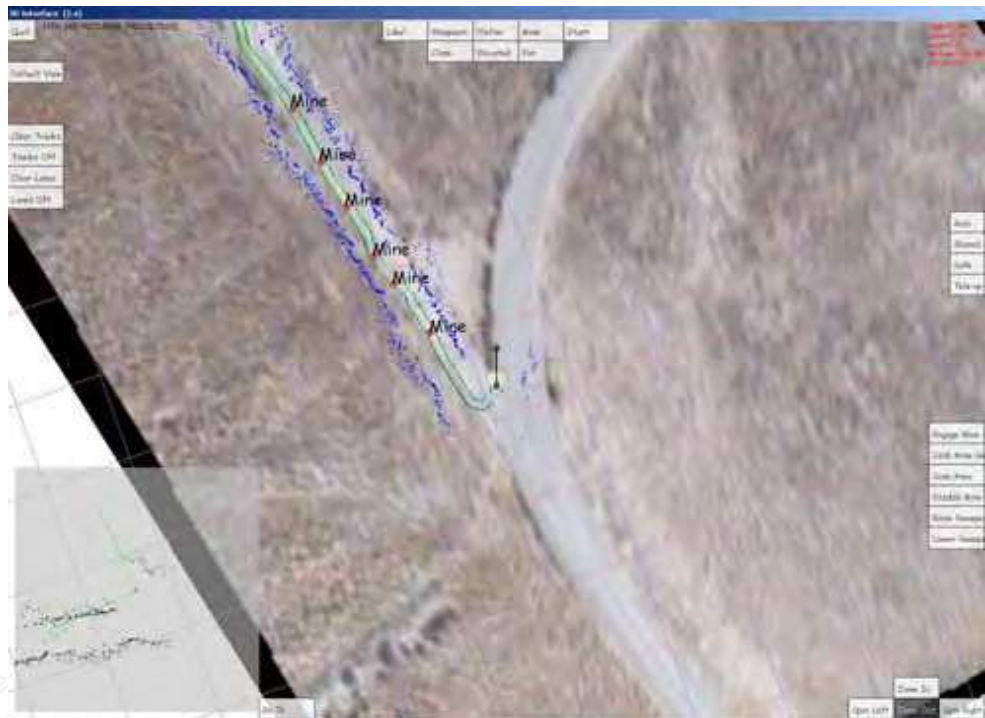

Figure 6. Interface showing fused robot map and mosiaced real-time aerial imagery during a UAV-UGV mine detection task

This fusion of data from air and ground vehicles is more than just a situation awareness tool. In fact, the display is merely the visualization of a collaborative positioning framework that exists between air vehicle, ground vehicle and human operator. Each team member contributes to the shared representation and has the ability to make sense of it in terms of its own, unique internal state. In fact, one lesson learned from the countermine work is that it may not be possible to perfectly fuse the representations especially when error such as positioning inaccuracy and camera skew play a role. Instead, it may be possible to support

collaboration by sharing information about landmarks and key environmental features that can be identified from multiple perspectives such as the corners of buildings or the intersection between two roads. The benefits of this strategy for supporting collaboration will be discussed further in Case Study One.

Simplifying the interface by correlating and fusing information about the world makes good sense. However, sensor fusion is not sufficient to actually change the interaction itself – the fundamental inputs and outputs between the human and the robotic system. To reduce true interaction complexity, there must be some way not only to abstract the robot physical and perceptual capabilities, but to somehow abstract away from the various behaviors and behavior combinations necessary to accomplish a sophisticated operation.

11. Understanding Modes of Autonomy

Another source of complexity within the original interface was the number of autonomy levels available to the user. When multiple levels of autonomy are available the operator has the responsibility of choosing the appropriate level of autonomy. Within the original interface, when the user wished to change the level of initiative that the robot is permitted to take, the operator would select between five different discrete modes. In teleoperation mode the user is in complete control and the robot takes no initiative. In safe mode the robot takes initiative only to protect itself or the environment but the user retains responsibility for all motion and behavior. In shared mode the robot does the driving and selects its own route whereas the operator serves as a backseat driver, providing directional cues throughout the task. In collaborative tasking mode, the human provides only high level intentions by placing icons that request information or task-level behavior (i.e. provide visual imagery for this target location; search this region for landmines; find the radiological source in this region). Full autonomy is a configuration rarely used whereby the system is configured to accept no human input and to accomplish a well-defined task from beginning to end. Table 1 below shows the operator and robot responsibilities for each autonomy mode.

Autonomy Mode	Defines Task Goals	Supervises Direction	Motovates Motion	Prevents Collision
Teleoperation Mode	Operator	Operator	Operator	Operator
Safe Mode	Operator	Operator	Operator	Robot
Shared mode	Operator	Operator	Robot	Robot
Collaborative Tasking Mode	Operator	Robot	Robot	Robot
Autonomous Mode	Robot	Robot	Robot	Robot

Table 1: Responsibility for operator and robot within each of five autonomy modes

Figure 7 below shows how the various behaviors that are used to support tasking in different autonomy modes.

The challenge with this approach is that operators often do not realize when they are in a situation where the autonomy on the robot should be changed and are unable to predict how a change in autonomy levels will actually affect overall performance. As new behaviors and intelligence are added to the robot, this traditional approach of providing wholly separate

modes of autonomy requires the operator to maintain appropriate mental models of how the robot will behave in each mode and how and when each mode should be used. This may not be difficult for simple tasks, but once a variety of application payloads are employed the ability to maintain a functional understanding of how these modes will impact the task becomes more difficult. New responses to this challenge will be discussed in Case Study Two.
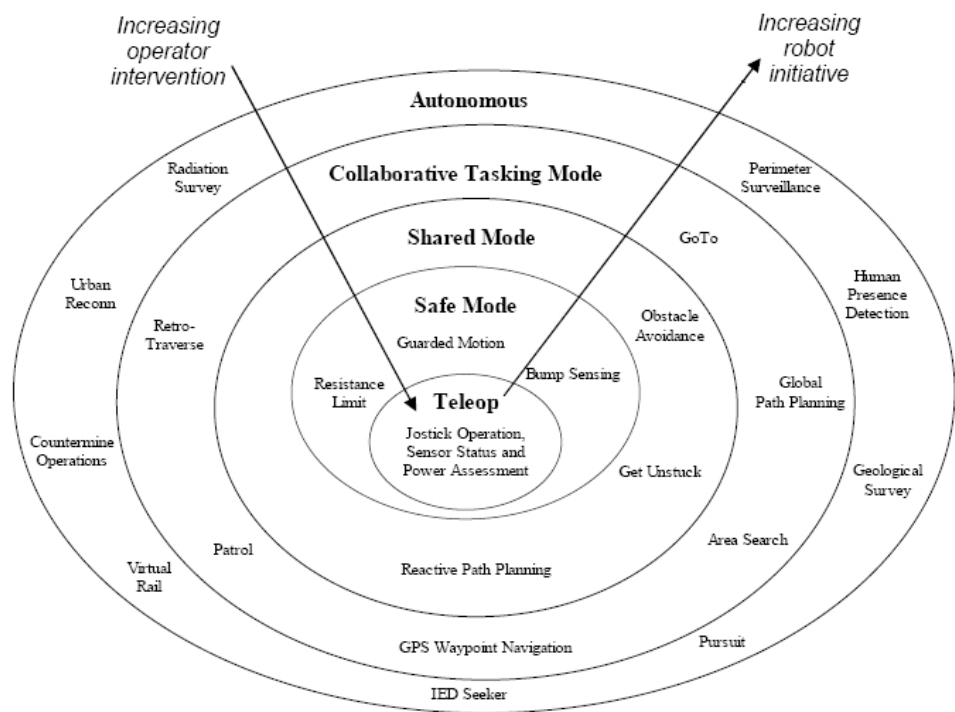

Figure 7. Behaviors associated with each of the five modes of autonomy

12. Case Study One: Robotic Demining

Landmines are a constant danger to soldiers during conflict and to civilians long after conflicts cease, causing thousands of deaths and tens of thousands of injuries every year. More than 100 million landmines are emplaced around the world and, despite humanitarian efforts to address the problem, more landmines are being emplaced each day than removed. Many research papers describe the challenges and requirements of humanitarian demining along with suggesting possible solutions (Nicoud & Habib, 1995; Antonic et. al., 2001). Human mine sweeping to find and remove mines is a dangerous and tedious job. Moreover, human performance tends to vary drastically and is dependent on factors such as fatigue, training and environmental conditions. Clearly, this is an arena where robot behaviors could someday play an important role. In terms of human-robot interaction, the need to locate and mark buried landmines presents a unique opportunity to investigate the value of shared representation for supporting mixed-initiative collaboration. A collaborative representation

is one of the primary means by which it is possible to provide the user with insight into the behavior of the unmanned systems.

13. Technical Challenges

It has long been thought that landmine detection is an appropriate application for robotics because it is dull, dirty and dangerous. However, the reality has been that the critical nature of the task demands a reliability and performance that neither teleoperated nor autonomous robots have been able to provide. The inherent complexity of the countermine mission presents a significant challenge for both the operator and for the behaviors that might reside on the robot. Efforts to develop teleoperated strategies to accomplish the military demining task have resulted in remarkable workload such that U.S. Army Combat Engineers report that a minimum of three operators are necessary to utilize teleoperated systems of this kind. Woods et al. describe the process of using video to navigate a robot as attempting to drive while looking through a 'soda straw' because of the limited angular view associated with the camera (Woods et al., 2004). If teleoperation is problematic for simple navigation tasks, the complexity of trying to use video remotely to keep track of where the robot has been over time as well as precisely gauge where its sensor has covered. Conversely, autonomous solutions have exhibited a low utility because the uncertainty in positioning and the complexity of the task rendered the behaviors less than effective. Given these challenges, it seemed prudent to explore the middle ground between teleoperation and full autonomy.
The requirement handed down from the US Army Maneuver Support Battlelab in Ft. Leonard-Wood was to physically and digitally mark the boundaries of a 1 meter wide dismounted path to a target point, while digitally and physically marking all mines found within that lane. Previous studies had shown that real-world missions would involve limited bandwidth communication, inaccurate terrain data, sporadic availability of GPS and minimal workload availability from the human operator. These mission constraints precluded conventional approaches to communication and tasking. Although dividing control between the human and robot offers the potential for a highly efficient and adaptive system, it also demands that the human and robot be able to synchronize their view of the world in order to support tasking and situation awareness. Specifically, the lack of accurate absolute positioning not only affects mine marking, but also human tasking and cooperation between vehicles.

14. Mixed-Initiative Approach

Many scientists have pointed out the potential for benefits to be gained if robots and humans work together as partners (Fong et. al., 2001; Kidd 1992; Scholtz & Bahrami, 2003; Sheridan 1992). For the countermine mission, this benefit cannot be achieved without some way to merge perspectives from human operator, air vehicle and ground robot. On the other hand, no means existed to support a perfect fusion of these perspectives. Even with geo-referenced imagery, real world trials showed that the GPS based correlation technique does not reliably provide the accuracy needed to support the countermine mission. In most cases, it was obvious to the user how the aerial imagery could be nudged or rotated to provide a more appropriate fusion between the ground robot's digital map and the air vehicle's image. To alleviate dependence on global positioning, collaborative tasking tools were developed that use common reference points in the environment to correlate disparate internal representations (e.g. aerial imagery and ground-based occupancy grids).

As a result, correlation tools were developed that allow the user to select common reference points within both representations. Examples of these common reference points include the corners of buildings, fence posts, or vegetation marking the boundary of roads and intersections. In terms of the need to balance human and robot input, it was clear that this approach required very little effort from the human (a total of 4 mouse clicks) and yet provided a much more reliable and accurate correlation than an autonomous solution. This was a task allocation that provided significant benefit to all team members without requiring significant time or workload.

The mission scenario which emerged included the following task elements.

a) Deploy a UAV to survey terrain surrounding an airstrip.
b) Analyze mosaiced real-time imagery to identify possible minefields.
c) Use common landmarks to correlate UAV imagery & umanned ground vehicle (UGV) occupancy map
d) UGV navigates autonomously to possible minefield
e) UGV searches for and mark mines.
f) UGV marks dismounted lane through minefield.

The behavior decomposition allows each team member to act independently while communicating environmental features and task intent at a high level.

To facilitate initiative throughout the task, the interface must not only merge the perspectives of robotic team members, but also communicate the intent of the agents. For this reason, the tools used in High Level Tasking were developed which allow the human to specify coverage areas, lanes or target locations. Once a task is designed by the operator, the robot generates an ordered waypoint list or path plan in the form of virtual colored cones that are superimposed onto the visual imagery and map data. The placement and order of these cones updates in real time to support the operator's ability to predict and understand the robot's intent. Using a suite of click and drag tools to modify these cones the human can influence the robot's navigation and coverage behavior without directly controlling the robot motion.

15. Robot Design

Figure 8: The Arcturus T-15 airframe and launcher

The air vehicle of choice was the Arcturus T-15 (see Figure 8), a fixed wing aircraft that can maintain long duration flights and carry the necessary video and communication modules. For the countermine mission, the Arcturus was equipped to fly two hour reconnaissance missions at elevations between 200 and 500ft. A spiral development process was undertaken to provide the air vehicle with autonomous launch and recovery capabilities as well as path planning, waypoint navigation and autonomous visual mosaicing. The resulting mosaic can be geo-referenced if compared to a priori imagery, but even then does not provide the positioning accuracy necessary to meet the 10cm accuracy requirements for the mission. On the other hand, the internal consistency of the mosaic is very high since the image processing software can reliably stitch the images together.

Carnegie Mellon University developed two ground robots (see Figure 9) for this effort which were modified humanitarian demining systems equipped with inertial systems, compass, laser range finders and a low-bandwidth, long range communication payload. A MineLab F1A4 detector which is standard issue mine detector for the U. S. Army, was mounted on both vehicles together with an actuation mechanism that can raise and lower the sensor as well as scan it from side to side at various speeds. A force torque sensor was used to calibrate sensor height based on sensing pressure exerted on the sensor when it touches the ground. The mine sensor actuation system was designed to scan at different speeds to varying angle amplitudes throughout the operation. Also, the Space and Naval Warfare Systems Center in San Diego developed a compact marking system that dispenses two different colors of agricultural dye. Green dye was used to mark the lane boundaries and indicate proved areas while red dye was used to mark the mine locations. The marking system consists of two dye tanks, a larger one for marking the cleared lane and a smaller one for marking the mine location.

Figure 9: Countermine robot platform

16. Experiment

The resulting system was rigorously evaluated by the Army Test and Evaluation Command (TECO) and found to meet the Army's threshold requirement for the robotic countermine mission. A test lane was prepared on a 50 meter section of an unimproved dirt road leading off of an airstrip. Six inert A-15 anti tank (AT) landmines were buried on the road at varying depths. Sixteen runs were conducted with no obstacles on the lane and 10 runs had various obstacles scattered on the lane. These obstacles included boxes and crates as well as

sagebrush and tumble weeds. The ARCS was successful in all runs in autonomously negotiating the 50 meter course and marking a proofed 1-meter lane. The 26 runs had an average completion time of 5.75 minutes with a 99% confidence interval of +/- 0.31 minutes. The maximum time taken was 6.367 minutes.

Figure 10: Proofed Lane and Mine Marking

The robot was able to detect and accurately mark, both physically and digitally, 130 out of 135 buried mines. Throughout the experiment there was one false detection. The robot also marked the proved mine-free lanes using green dye. The robot was able to navigate cluttered obstacles while performing various user-defined tasks such as area searches and the de-mining of roads and dismounted lanes.

Figure 11. Interface shows the operator the position of mines detected along a road

17. Discussion

When compared to the current military baseline, the mixed-initiative system produced a fourfold decrease in task time to completion and a significant increase in detection accuracy.This is particularly interesting since previous attempts to create robotic demining systems had failed to match human performance. The difference between past strategies and the one employed in this study is not that the robot or sensor was more capable; rather, the most striking difference was the use of mixed-initiative control to balance the capabilities and limitations of each team member. Without the air vehicle providing the tasking backdrop and the human correlating it with the UGV map, it would not have been possible to specify the lane for the robot to search. The research reported here indicates that operational success was possible only through the use of a mixed-initiative approach that allowed the human, air vehicle and ground vehicle to support one another throughout the mission. These findings indicate that by providing an appropriate means to interleave human and robotic intent, mixed initiative behaviors can address complex and critical missions where neither teleoperated nor autonomous strategies have succeeded.

Another interesting facet of this study is to consider how the unmanned team compares to a trained human attempting the same task. When comparing the robot to current military operations, the MANSCEN at Ft. Leonard Wood reports that it would take approximately 25 minutes for a trained soldier to complete the same task accomplished by the robot, which gives about a four-fold decrease in cycle time without putting a human in harm's way. Furthermore, a trained soldier performing a counter-mine task can expect to discover 80% of the mines. The robotic solution raises this competency to 96% mine detection. Another interesting finding pertained to human input is that the average level of human input throughout the countermine exercises, namely the time to set up and initiate the mission, was less than 2% when calculated based on time. The TECO of the U.S. Army indicated that the robotic system achieved "very high levels of collaborative tactical behaviors."

One of the most interesting HRI issues illustrated by this study is the fact that neither video nor joystick control was used. In fact, due to the low bandwidth required, the operator could easily have been hundreds of miles away from the robot, communicating over a cell phone modem. Using the system as it was actually configured during the experiment, the operator had the ability to initiate the mission from several miles away. Once the aerial and ground perspectives are fused within the interface, the only interaction which the soldier would have with the unmanned vehicles would be initiating the system and selecting the target location within the map.

18. Case Study Two: Urban Search and Rescue

This case study evaluates collaborative tasking tools that promote dynamic sharing of responsibilities between robot and operator throughout a search and detection task. The purpose of the experiment was to assess tools created to strategically limit the kind and level of initiative taken by both human and robot. The hope was that by modulating the initiative on both sides, it would be possible to reduce the deleterious effects referred to earlier in the chapter as a "fight for control" between the human and robot. Would operators notice that initiative was being taken from them? Would they have higher or lower levels of workload? How would overall performance be affected in terms of time and quality of data achieved?

19. Technical Challenge

While intelligent behavior has the potential to make the user's life easier, experiments have also demonstrated the potential for collaborative control to result in a struggle for control or a suboptimal task allocation between human and robot (Marble et al., 2003; Marble et al., 2004; Bruemmer et al., 2005). In fact, the need for effective task allocation remains one of the most important challenges facing the field of human-robot interaction (Burke et al., 2004). Even if the autonomous behaviors on-board the robot far exceed the human operators ability, they will do no good if the human declines to use them or interferes with them. The fundamental difficulty is that human operators are by no means objective when assessing their own abilities (Kruger & Dunning, 1999; Fischhoff et al., 1977). The goal is to gain an optimal task allocation such that the user can provide input at different levels without interfering with the robot's ability to navigate, avoid obstacles and plan global paths.

20. Mixed-Initiative Approach

In shared mode (see table 1), overall team performance may benefit from the robot's understanding of the environment, but can suffer because the robot does not have insight into the task or the user's intentions. For instances, absent of user input, if robot is presented with multiple routes through an area shared mode will typically take the widest path through the environment. As a result, if the task goal requires or the human intends the exploration of a navigable but restricted path, the human must override the robot's selection and manually point the robot towards the desired corridor before returning system control to the shared autonomy algorithms. This seizing and relinquishing of control by the user reduces mission efficiency, increases human workload and may also increase user distrust or confusion. Instead, the CTM interface tools were created to provide the human with a means to communicate information about the task goals (e.g. path plan to a specified point, follow a user defined path, patrol a region, search an area, etc) without directly controlling the robot. Although CTM does support high level tasking, the benefit of the collaborative tasking tools is not merely increased autonomy, but rather the fact that they permit the human and robot to mesh their understanding of the environment and task. The CTM toolset is supported by interface features that illustrate robot intent and allow the user to easily modify the robot's plan. A simple example is that the robot's current path plan or search matrix is communicated in an iconographic format and can be easily modified by dragging and dropping vertices and waypoints. An important feature of CTM in terms of mixed-initiative control is that joystick control is not enabled until the CTM task is completed. The user must provide input in the form of intentionality rather than direct control. However, once a task element is completed (i.e target is achieved or area searched), then the user may again take direct control. Based on this combined understanding of the environment and task, CTM is able to arbitrate responsibility and authority.

21. Robot Design

The experiments discussed in this paper utilized the iRobot "ATRV mini" shown on the left in Figure 12. The robot utilizes a variety of sensor information including compass, wheel encoders, laser, computer camera, tilt sensors, and ultrasonic sensors. In response to laser and sonar range sensing of nearby obstacles, the robot scales down its speed using an event

horizon calculation, which measures the maximum speed the robot can safely travel in order to come to a stop approximately two inches from the obstacle. By scaling down the speed by many small increments, it is possible to insure that regardless of the commanded translational or rotational velocity, guarded motion will stop the robot at the same distance from an obstacle. This approach provides predictability and ensures minimal interference with the operator's control of the vehicle. If the robot is being driven near an obstacle rather than directly towards it, guarded motion will not stop the robot, but may slow its speed according to the event horizon calculation. The robot also uses a mapping, localization system developed by Konolige et al.

Figure 12

22. Experiment

A real-world search and detection experiment was used to compare shared mode where the robot drives, but the human can override the robot at any time, to a Collaborative Tasking Mode (CTM), where the system dynamically constrains user and robot initiative based on the task element. The task was structured as a remote deployment such that the operator control station was located several stories above the search arena so that the operator could not see the robot or the operational environment. Plywood dividers were interspersed with a variety of objects such as artificial rocks and trees to create a 50ft x 50ft environment with over 2000 square feet of navigable space. Each participant was told to direct the robot around the environment and identify items (e.g. dinosaurs, a skull, brass lamp, or building blocks) located at the numbers represented on an a priori map. In addition to identifying items, the participants were instructed to navigate the robot back to the Start/Finish to complete the loop around the remote area. This task was selected because it forced the participants to navigate the robot as well as use the camera controls to identify items at particular points along the path. The items were purposely located in a logical succession in an effort to minimize the affect of differences in the participants' route planning skills.

In addition to the primary task of navigating and identifying objects the participants were asked to simultaneously conduct a secondary task which consisted of answering a series of basic two-digit addition problems on an adjacent computer screen. The participants were instructed to answer the questions to the best of their ability but told that they could skip a problem by hitting the <enter> key if they realized a problem appeared but felt they were

too engaged in robot control to answer. Each problem remained present until it was responded to, or the primary task ended. Thirty seconds after a participant's response, a new addition problem would be triggered. The secondary task application recorded time to respond, in seconds, as well as the accuracy of the response and whether the question was skipped or ignored.

During each trial, the interface stored a variety of useful information about the participant's interactions with the interface. For instance, the interface recorded the time to complete the task to be used as a metric of the efficiency between the methods of control. For the CTM participants, the interface also recorded the portion of time the robot was available for direct control. The interface recorded the number of joystick vibrations caused by the participant instructing the robot to move in a direction in which it was not physically possible to move. The number of joystick vibrations represent instances of human navigational error and, in a more general sense, confusion due to a loss of situation awareness (See Figure 13). The overall joystick bandwidth was also logged to quantify the amount of joystick usage. Immediately after completing a trial, each participant was asked to rank on a scale of 1 to 10 how "in control" they felt during the operation, where 1 signified "The robot did nothing that I wanted it to do" and 10 signified, "The robot did everything I wanted it to do."
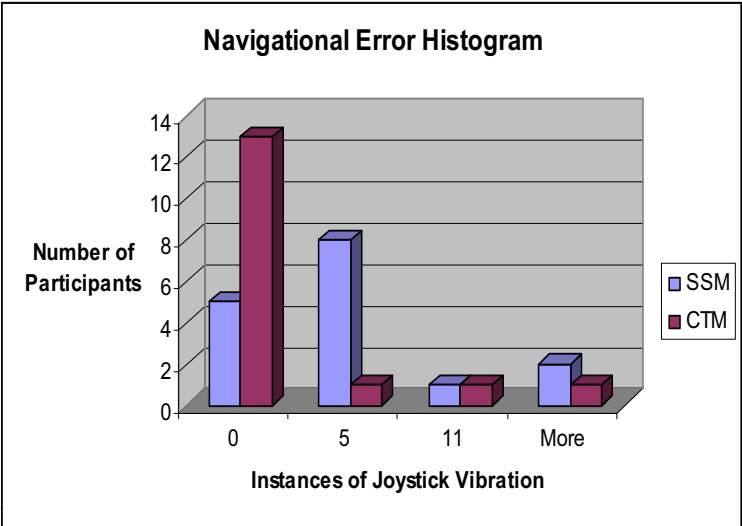

Figure 13. Navigational Error

All participants completed the assigned task. Analysis of the time to complete the task showed no statistically significant difference between the shared mode and CTM groups. An analysis of human navigational error showed that 81% of participants using CTM experienced no instances of operator confusion as compared to 33% for the shared mode participants (see Figure 13). Overall, shared mode participants logged a total of 59 instances of operator confusion as compared with only 27 for the CTM group.

The CTM participants collectively answered 102 math questions, while the shared mode participants answered only 58. Of questions answered, CTM participants answered 89.2% correctly as compared to 72.4% answered correctly by participants using shared mode. To further assess the ability of shared mode and CTM participants to answer secondary task

questions an analysis was performed on the average response time for each group. CTM participants had a statistically significant average response time of 25.1 seconds as compared to 49.2 seconds for those using shared mode. Together these results indicate that the participants using the collaborative tasking tools experienced a substantial decrease in the required workload to complete the task. In addition, CTM participants enjoyed a higher overall feeling of control as compared to shared mode participants (see Figure 14).
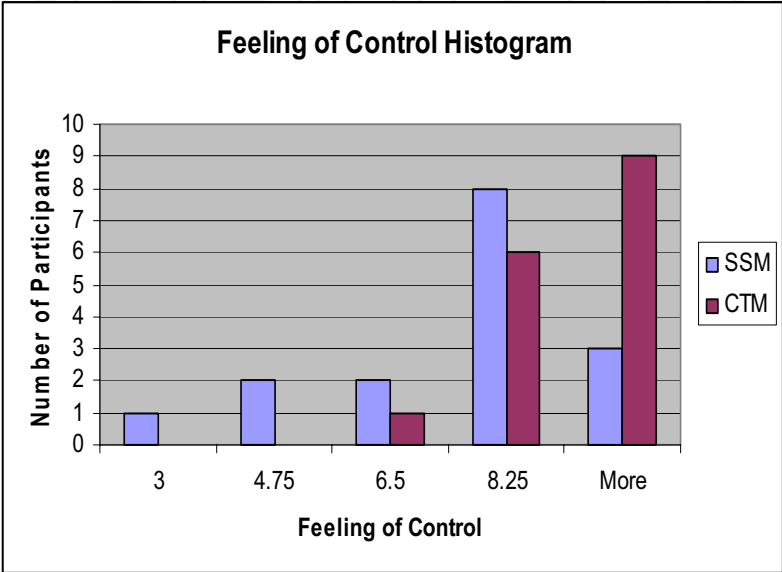
Figure 14. Feeling of Control

23. Discussion

This experiment provides validation of the collaborative tasking tools that have been implemented as part of the RIK. The experiment showed that from an engineering perspective, the blending of guarded motion, reactive obstacle avoidance and global path planning behaviors on board the robot can be used effectively to accomplish a search and detection task. Of greater significance to the Human-Robot Interaction (HRI) community is the fact that this experiment represents a definitive step away from the supervisory control paradigm where the human may accept or decline robot initiative, while remaining at all times in the leadership role for all task elements. Instead, the collaborative tasking tools presented here arbitrate leadership in a facilitative manner to optimize overall team performance. By constraining operator initiative at the right times, CTM reduces human confusion and frustration. Data from this study suggests that the CTM serves in a suprising fashion to increases users' feeling of control by taking control away from them. Something we had not initially predicted. Although the HRI community has long used the phrase "mixed initiative" to describe the goal of team members blending their input together,the findings of this paper imply that rather than "mixing" initiative, human-robot teaming may benefit when initiative is "facilitated" to avoid conflict and optimize task allocation.

24. Conclusion

All too often, increased robot capability has been handed down to the user in the form of increased interaction complexity. As we have seen, autonomous robotic behaviors do not necessarily provide a performance benefit and may lead to operator confusion and distrust if the system does not support a mental model that can be easily adopted or used by operators. As robot behaviors become increasingly complex, it is imperative that we find a means to hide complexity while still keeping users in the know and allowing them to be part of the action. By so doing, the operator can be freed up to successfully oversee the performance of greater numbers of robots while maintaining a greater sense of his or her own situation awareness. We believe that the mixed-initiative tools discussed within this chapter provide evidence that if we are willing to move beyond the existing tools and metaphors, it is possible to craft mixed-initiative interaction methods that can enhance operator and system performance, and decrease operator workload.

Ideally, the interaction metaphor that underlies these tools should be functionally linked to the application such that the human need no longer consider individual actions, perceptions or behaviors, but rather can begin to reason about the task elements that are native to the application domain. Thus, one might say that an effective interaction metaphor provides a readily apparent abstraction from robot behaviors into mission tasking. The more elegant the interaction metaphor, the less the user will have to think about the robot and the more s/he can begin to think about the environment and task in their own terms. It is because of this that an elegant interaction metaphor will, by its very nature, simplify the theory of robot behavior necessary for the operator to efficiently interact with the system. By simplifying the necessary theory of robot behavior it may be possible to reduce uncertainty, bring users' expectations into line, reduce the dimensionality of human-robot communication and narrow the possible outcomes. Ultimately, the overall complexity of the system does not change. Instead of the human performing the cognitive mapping between intent and robot behaviors, the interface must now play a role in accomplishing this mapping. Intelligent interface tools can be used to orchestrate robot behaviors according to the operative interaction metaphor.

24. References

Antonic, D., Ban, Z. & Zagar, M. (2001) Demining robots - requirements and constraints. *Automatika*, 42(3-4)

Arkin, R. C. (1997) *Behavior-Based Robotics*. MIT Press. Cambridge, MA

Brooks, R. ; Brazeal, C., Marjanovic, M., Scassellati, B., & Williamson, M. (1998) The Cog Project: Building a Humanoid Robot. *In Computation for Metaphors, Analogy, and Agents*, Springer Lecture Notes in Computer Science, 1562, C. Nehaniv, Ed., Springer-Verlag, Berlin

Bruemmer, D. J. ; Few, D. A., Walton, M. C., Boring, R. L., Marble, J. L., Nielsen, C. W., & Garner, J. (2005) . Turn off the television!: Real-world robotic exploration experiments with a virtual 3D display . *In Proceedings of the Hawaii International Conference on System Sciences (HICSS)*. Waikoloa Village, Hawaii. January 3-6

Bruemmer, D.J. ; Few, D.A., Boring, R. L., Marble, J. L., Walton, M. C. & Nielsen C. W. (2005) Shared Understanding for Collaborative Control. IEEE *Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans.* Vol. 35, no.4, pp. 505-512. July

Bruemmer, D J; Boring, R. L,. Few. D.A., & Walton M. C. (2004). I Call Shotgun: An Evaluation of a Mixed-Initiative Control for Novice Users of a Search and Rescue Robot *In proceedings of 2004 IEEE International Conference on Systems, Man & Cybernetics,* The Hague, ND October 10–13

Burke, J. L. ; Murphy, R. R., Coovert, M. D. & Riddle D. L. (2004) Moonlight in Miami: A field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise, *Human-Computer Interaction,* 19:85–116

Byers, Z. ; Dixon, M., Goodier, K., Grimm, C. M., & Smart, W. D. (2003) An Autonomous Robot Photographer, In *Proceedings of the IEEE/RSJ International Conference on Robots and Systems (IROS 2003,* Vol 3, pages 2636-2641, Las Vegas, Nevada.

Byers, Z., Dixon, M., Goodier, K., Smart, W. D. & Grimm, C.M. (2003) Say Cheese!: Experiences with a Robot *In Proceedings for the Fifteenth Innovative Applications of Artificial Intelligence Conference (IAAI 2003),* John Riedl and Randy Hill (editors), pages 65-70, Acapulco, Mexico

Conley, D. T., & Goldman, P. (1994) Facilitative Leadership: How Principals Lead Without Dominating. Eugene, Oregon: Oregon School Study Council, August

Desai. M. & Yanco H. A. Blending Human and Robot Inputs for Sliding Scale Autonomy.(2005) *Proceedings of the 14th IEEE International Workshop on Robot and Human Interactive Communication (Ro-MAN).* Nashville, TN, August

Fischhoff, B. ; Slovic, P., & Lichtenstein S, (1977) Knowing with Certainty: The appropriateness of extreme confidence *Journal of Experimental Psychology: Human Perception and Performance,* 3:552-564

Fong, T. ; Thorpe, C. & Baur, C. (2001) Collaboration, dialogue,and human robot interaction. *In 10th International Symposium of Robotics Research,* Lourne, Victoria, Australia, November

Gertman, D. I., Blackman, H. S., Marble, J. L., Byers, J., & Smith, C. (2005) The SPAR-H Human Reliability Analysis Method, NUREG/CR-6883, US Nuclear Regulatory Commission, Washington, DC

Kidd, P.T. (1992) Design of human-centered robotic systems. In Mansour Rahimi and Waldemar Karwowski, editors, *Human Robot Interaction,* pages 225–241. Taylor and Francis, London, England

Konolige, K. (2004) Large-scale map-making. *In Proceedings of the National Conference on AI (AAAI),* San Jose, CA

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology,* 77, 1121-1134

Maxwell B. A. ; Smart W. D., Jacoff A., Casper, J., Weiss, B., Scholtz, J., Yanco, H., Micire, M., Stroupe, A., Stormont, D., Lauwers, T., (2004) The 2003 AAAI Robot Competition and Exhibition. *AI Magazine.* Vol. 25, no. 2, pp. 68-80

Marble, J. L. ; Bruemmer, D. J., and Few, D. A. (2003) Lessons Learned from Usability Tests with a Collaborative Cognitive Workspace for Human-Robot Teams. In *Proceedings of the IEEE International Conférence on Systems, Man and Cybernetics,* Washington, D.C., October 5-8

Marble, J. L. ; Bruemmer D. J., Few, D. A. and Dudenhoeffer D. D. (2004) Evaluation of supervisory vs. peer-peer interaction for human-robot teams *In Proceedings of the 37th Annual Hawaii International Conference on Systems Sciences*, Big island, Hawaii, January 4-8

Nicoud, J. D. &. Habib, M. K (1955)The pemex-b autonomous demining robot: perception and navigation strategies. *In Proceedings of intelligent robots and systems*, pages 419–424, Pittsburgh, PA

Nielsen Curtis W. and Michael A. Goodrich. Testing the usefulness of a pan-tilt-zoom (PTZ) camera in humanrobot interactions. In Proceedings *of the Human Factors and Ergonomics Society Meeting*, San Francisco, CA, 2006.

Pacis, E.B., Everett, H. R., Farrington,N., . Bruemmer D. J, (2004) Enhancing Functionality and Autonomy in Man-Portable Robots, In *Proceedings of the SPIE Defense and Security Symposium 2004.* 13 -15 April

Royal Canadian Mounted Police: Police Dog Services, K-9 Recruit Successfully Completes Training, 2002. [Online]. Available: http://www.rcmp-grc.gc.ca/html/dogs.htm

Scholtz, J. & Bahrami, S. (2003) Human-robot interaction: development of an evaluation methodology for the bystander role of interaction. *In Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, volume 4, pages 3212–3217

Sellner, B. ; Heger, F. W., Hiatt, L. M, Simmons, R. & Singh, S. (2006) Coordinated multiagent teams and sliding autonomy for large-scale assembly. *Proceedings of the IEEE*, 94(7):1425–1444

Sheridan,T. B. (1992) Telerobotics, automation, and human supervisory control. MIT Press, Cambridge, MA

Smart, W. D. ; Grimm, C. M., Dixon, M. and Byers, Z. (2003) (Not) Interacting with a Robot Photographer,. *In Human Interaction with Autonomous Systems in Complex Environments*: Papers from the 2003 AAAI Spring Symposium, Stanford University, March 24-26,David Kortenkamp and Michael Freed (editors), pages 181-186, Stanford California

Woods, D. D., Tittle, J., Feil, M., F.& Roesler, A. (2004) Envisioning human-robot coordination in future operations. *IEEE Transactions on Systems, Man, and Cybernetics*, Part C, 34(2):210–218, May.

Human Robot Interaction

Edited by Nilanjan Sarkar

Human-robot interaction research is diverse and covers a wide range of topics. All aspects of human factors and robotics are within the purview of HRI research so far as they provide insight into how to improve our understanding in developing effective tools, protocols, and systems to enhance HRI. For example, a significant research effort is being devoted to designing human-robot interface that makes it easier for the people to interact with robots. HRI is an extremely active research field where new and important work is being published at a fast pace. It is neither possible nor is it our intention to cover every important work in this important research field in one volume. However, we believe that HRI as a research field has matured enough to merit a compilation of the outstanding work in the field in the form of a book. This book, which presents outstanding work from the leading HRI researchers covering a wide spectrum of topics, is an effort to capture and present some of the important contributions in HRI in one volume. We hope that this book will benefit both experts and novice and provide a thorough understanding of the exciting field of HRI.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

David J. Bruemmer, David I. Gertman, Curtis W. Nielsen, Douglas A. Few and William D. Smart (2007). Supporting Complex Robot Behaviors with Simple Interaction Tools, Human Robot Interaction, Nilanjan Sarkar (Ed.), ISBN: 978-3-902613-13-4, InTech, Available from: http://www.intechopen.com/books/human_robot_interaction/supporting_complex_robot_behaviors_with_simple_interaction_tools

INTECH
open science | open minds