Metaphors to Drive By: Exploring New Ways to Guide Human - Robot Interaction

David J. Bruemmer[*,1], David I. Gertman[2] and Curtis W. Nielsen[1]

[1]*Robotics and Human Systems Department, Idaho National Laboratory (INL), Idaho Falls, ID, USA*

[2]*Human Factors and Advanced I&C Department, Idaho National Laboratory (INL), Idaho Falls, ID, USA*

Abstract: Autonomous behaviors created by the research and development community are not being extensively utilized within energy, defense, security, or industrial contexts. This paper provides evidence that the interaction methods used alongside these behaviors may not provide a mental model that can be easily adopted or used by operators. Although autonomy has the potential to reduce overall workload, the use of robot behaviors often increased the complexity of the underlying interaction metaphor. This paper reports our development of new metaphors that support increased robot complexity without passing the complexity of the interaction onto the operator. Furthermore, we illustrate how recognition of problems in human-robot interactions can drive the creation of new metaphors for design and how human factors lessons in usability, human performance, and our social contract with technology have the potential for enormous payoff in terms of establishing effective, user-friendly robot systems when appropriate metaphors are used.

INTRODUCTION

From a technological perspective, mobile robots have changed dramatically in the recent past and will continue to do so. Although some of these new technologies are remarkable, they have not fundamentally changed the way we, as users, interact with and perceive robots.

As an example, fielded robots are still generally directly controlled using video as a basis for supporting operator awareness of the remote environment. 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 [1-4]. Rather the problem is that the resulting interaction is often complex and the robot's actions often seem mysterious to a novice robot operator. Without the right interaction metaphor, users have a hard time anticipating how the robot will respond to their commands. This is especially true for behaviors that are adaptive or which exhibit emergent effects and/or non-determinism.

To make matters worse, robot behaviors are often accompanied by interface tools geared towards the robot's perception and internal representation of the world. It should not surprise us that these representations may not complement the operator's information processing capabilities. As a community, our well-intentioned efforts to provide increased user-bandwidth have resulted in increasingly "busy" interfaces for the user to navigate because we generally believe that information from the robot should help the operator. As the complexity of the robot's physical capabilities; perceptions; and behavior increases, we must be careful not to pass this complexity directly on to the user.

When it comes to interactions, we believe that the simpler the interaction metaphor that can still get the job done, the better. This is why, in many instances, joystick control still trumps the highly sophisticated autonomous and semi-autonomous capabilities that are now possible. Teleoperation is not easy; however, despite the technical and cognitive challenges it presents, teleoperation offers an interaction metaphor that is easily understood by even novice operators.

This paper uses several previous human-robot interaction studies to examine the problems with current metaphors for interacting with robot behaviors. It also considers the potential for new metaphors to change the way people think about and use robot behaviors. Ultimately, we need to listen to the user's "story" of what s/he would like the robot to do, i.e., how the operator and robot should act as a team and the degree to which robot information should be relayed to the operator.

Over the years, the Idaho National Laboratory (INL) has conducted a number of experiments related to robotic behavior including scalable autonomy (see [1] for a summary). During this time, it became apparent that augmenting robotic capability did not necessarily result in greater trust or enhanced operator performance. In fact, many participants preferred to use predictable robots with lower performance than robots that kept them guessing. Another insight was that adding bandwidth to the interface did not necessarily enhance performance. In one study operators with a map-based interface did equally well at the task of exploring a building as operators who got the same map, but were also given a video-based interface which required 5000 times more bandwidth [5]. In fact, operators with more data made more errors and felt more frustrated. More data is not necessarily better.

Likewise, giving the operator more control does not necessarily increase overall performance either. Specifically, one study showed that performance could be improved by restricting human input at certain times such as when the robot is navigating through a doorway [6]. Without this restriction, many operators would distract the robot by suggesting a number of minor and incorrect orientations to "help" the robot through the doorway. The basic problem

*Address correspondence to this author at the Robotics and Human Systems Department, Idaho National Laboratory (INL), Idaho Falls, ID, USA; E-mail: David.Bruemmer@inl.gov

was that operators did not understand the correct role for either themselves or the robot because they did not understand that the robot could observe more of the environment than the operator could see through the video. As designers of the robot system, we did not have a suitable metaphor to communicate the operator's role in the interaction to the operator.

User trust is a major component in operational success, and in our studies trust has been shown to be reduced by conflicts regarding the user's assumptions with respect to robot initiative [5]. Nourbkhsh [7] states "robots are real things, robots push back on the world," unless left on their own, i.e., fully autonomous, they will always be interacting in social spaces [4]. This notion of pushing back on the world means the robot will have some intelligence and make some decisions of its own accord. In our studies, the operators were not ready for the robot to push back on the world. Therefore, we need to improve the robot's ability to present its intentions to the operator such that the operator is on the "same page" as the robot and more prepared to accept the robot's initiative.

Currently, concepts for applications such as urban search and rescue assume complementary human-robotic interaction. The assumption is that by so doing we can lower the risk for victims and first responders alike. Unfortunately, effective human-robot interaction is not guaranteed and will likely only be the result of focused effort to control the level of system complexity and operator trust.

We next discuss the issues behind system complexity and operator trust and present concepts of usability and metaphor in the domain of human-robot interaction along with a potential path forward.

HUMAN FACTORS, TRUST, AND INTERFACE DESIGN

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 [8-11]. The reasons are varied.

Our Social Contract with Automation

People have an implicit social contract with automation. Things either work the way we think they should or we discard them and work around them. For example, if the brakes on our car only were to work once in a while we would tend to find another car or other form of transportation. If aircraft were likely to land at the wrong airport we would likely find other forms of transportation. Lee and See [12] maintain that "trust guides reliance when complexity and unanticipated situations make a complete understanding of the automation impractical." They also state that poor partnerships (with technology) will become increasingly costly and catastrophic ([12] p. 50). It is in our best interest to foster relationships that are built on trust and performance. Moreover, the nature of certain tasks wherein robotics are currently or envisioned to be applied such as: urban search and rescue (USAR); countermine operations; and, long duration space flight, are fraught with relatively high levels of uncertainty, and, as Lee and See note, "trust allows people to accommodate uncertainty" ([12] p. 52).

Trust and Technology Acceptance

Gertman, Novack and Marble [13] present an initial model for calculating cultural effects on performance. In that research, technology acceptance is reviewed for its contribution as a mediating variable in estimating human reliability for operator response to off-normal events in high technology systems. Factors such as culture and technology assimilation are postulated to be influencing factors and corresponding system response times, error rates, and the type of errors committed may reflect technology acceptance and trust.

Their work provides evidence that under high stress, the ability of a human and computer agent to collaboratively problem solve is dependent on the level of human trust. A natural hypothesis would be that trust might have an equally important effect on human-robot interaction. Indeed, an early study performed at the INL does indicate the importance of trust [6]. Users who did not trust the autonomy performed poorly when compared to those who did. Interestingly, these users performed better in teleoperation mode than when using the robot behaviors. This is in contrast to the users who trusted the robot behaviors and achieved the highest overall performance.

The challenge is that the same means used across the human-computer interaction community to build trust in automation may not be applicable or possible for intelligent mobile robots. Gertman, Novack and Marble [13] indicate that designers of intelligent systems usually try to increase trust by reducing operator uncertainty.

In practice this often involves providing an operator with a nearly complete understanding of the task environment. For robotic deployment in unstructured environments, this is a fool's errand. For example, imagine an urban landscape. Complete real-time understanding of the environment as the robot sees it involves a complex interplay among many noisy and incomplete sensory systems (range, radar, sonar, forward looking infrared (FLIR), motion, speed indication, etc).

What is presented to operators greatly influences their comfort with automation, however we assert that added complexity is not necessarily the way to enhance performance or acceptance.

By analogy, when we drive a car there is much information that could be collected *via* sensors and presented to the driver. For example, we could be informed of the engine revolutions per minute (RPM) prior to and during braking, the amount of brake pad remaining, readings on hydraulic pressure, the change in temperature during braking, the force by which hydraulic fluid is pulsed through the system, asymmetries present in breaking surfaces experienced by the four wheels, etc.

As drivers however, we do not want or use this information. Rather, most of us have a mental model of how a car operates and that is all we care about. When we push down on the brake pedal, we expect that the car will slow down. We don't need much more information because we use a continuous feedback loop (including vestibular sensing and visual cues) to modulate the braking behavior by pressing down more or less. This allows us to make other decisions that may be called for regarding the status of oncoming traf-

fic, pedestrians, weather conditions, and accidents that may appear up ahead.

The important point is that precious little system understanding or component data is needed to drive a car. In fact, humans generally don't even need to review the speedometer indication when accomplishing a braking task. All of the sophisticated technology (such as antilock braking) that is involved in deceleration has been reduced to a very simple metaphor that relates pedal depression to deceleration. It is hard to trust a high-dimensionality stream of multi-modal data, whereas it is easy for us to trust a pedal. The same thing is true for a steering wheel. This is one of the reasons why automotive makers have retained the steering wheel even in cars where the steering wheel is drive-by-wire and is not even physically connected to the drive train. Likewise, it should be possible to support human–robot interaction with metaphorically driven, intuitive models that abstract away complexity, focus the user's attention and reinforce notions of trust.

Components of Trust for Robots

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 pre-commitment is linked to the user's acknowledgement of value. 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. Fostering appropriate distrust may be equally vital. If the robot's map of the world begins to degrade, trust in the robot's abilities should also degrade.

To trust the robot, the human must believe that the robot will take action appropriate to the context of the situation. Although trust involves a pre-commitment, it is important to understand that trust undergoes a continual process of re-evaluation based on the user's own observations and experiences. 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. Rather, the operator develops and maintains a relationship with the robot based on an ability to accomplish a shared goal.

What Level of Understanding is Needed?

On the surface, what we argue here is in opposition to Sheridan's [14] thesis that trust (in automation) is facilitated by the operator's ability to understand the algorithms underlying automated behavior and his or her conviction that those algorithms are supportive of the operator's goals for specific situations. For many potential situations, the level of training and expertise of the user in the field may not support a deep understanding of the algorithms that underlie behaviors such as those for navigation, positioning, exploration, etc. As behaviors become more complex, it is difficult even for the developers to understand the intelligent fusion of data from many different sensors and the interplay of many indepen-

dent behaviors. 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. Is blind acceptance all that we can hope for or are there other means to build trust and facilitate understanding?

Functional Predictability

We maintain that an understanding of the behavior algorithms is not necessary to support operator trust. Note that the need for functional predictability is different than the need to understand the internal algorithmic process. The internal algorithmic process requires the user to understand how specific behavior inputs (i.e., sensing, tasking) will be mapped to specific behavior outputs (i.e., wheel or arm movement).

We define functional predictability as the ability of the user to develop a cognitive model that allows them to predict the outputs of the robot system in response to user defined input. The major difference is that functional predictability can be based on the user's own observations of the environment and task, and the robot's past behavior in similar situations whereas the algorithmic approach requires that the robot's own perceptions be known.

For the operator, we conjecture that only functional predictability is necessary to support effective tasking and trust. We propose that this cognitive mapping may be effectively communicated to the operator through a narrative or symbolic means. If we are clever, we can invoke the role of metaphor to provide an implicit understanding of what to expect for the human-robot interactions. For example, with small distributed robots that do not have absolute positioning, we may explain that the robots will spread out like a swarm of ants or bees. For an autonomous unmanned aerial vehicle (UAV) that has change detection capabilities on board, we might say that the UAV is like a hawk that tracks movement on the ground below. Whatever the story, the interface design should build upon a metaphor that supports the interaction.

Leveraging metaphors that are already understood by the user population may provide us the means to develop an intuitive interface where performance is independent of extensive training. Using effective metaphors may help us blur the line between novice and expert robot operators. To accomplish this, trust and skill need to be rapidly acquired and sustained. We particularly like the suggestions by Tognazzini [15] part of the Nielsen Norman Group who provides the following advice: "Effective interfaces do not concern the user with the inner workings of the system." We feel this to be highly applicable to robotics where designers must fight the urge to try to present everything to the operators. Tognazzini [15] admonishes us to bring metaphors alive by appealing to people's perceptions – sight, sound, touch and kinesthetic as well as triggering those memories.

EMBRACING USER CENTRIC METAPHORS

Metaphors have always played a role in robotic system employment and will continue to do so. The basic theory of operation that supports the use of a joystick is so well-accepted that we generally don't think of it as an interaction metaphor. To understand the role of this metaphor, consider that directional movement of a robot, although it may seem

like a low-level, simple functionality, is actually the effect of pulsed electricity being sent to various motors in a carefully sequenced fashion. The metaphor, simple though it is, hides the complexity of the electricity powering the motors from the user.

Furthermore, the joystick metaphor is independent of the wheel size or whether the vehicle is a skid-steer or an acumen-steer robot. The reason why the joystick metaphor has been successful is that it is simple to understand and use and applicable to a variety of vehicles and technologies. None of the mechanical, electrical, or computer engineering complexity necessary to produce reliable motion control from the robot platform is directly apparent to the user.

Although the joystick metaphor is effective in some situations, it is severely limited in its ability to scale to different modes of robot autonomy. The joystick metaphor is appropriate for an interaction model where the human directly controls robot motion. After all, direct motion control is the fundamental input and output to the joystick controller. A number of human-robot interaction studies have shown that use of a joystick is problematic when significant initiative is exhibited by the robot.

Several studies at the INL have blended human joystick control with a robotic capacity to prevent collisions and sense where it can fit. In such studies, some users come to trust and rely on robot initiative whereas others refuse to trust this capability and experience significantly reduced performance due to a "fight-for-control" between the robot autonomy and their own joystick input. In one such study, users would repeatedly try to drive the robot through an opening where it could not fit even after the robot had informed them in numerous ways (including vibrating joystick, red blockages on the interface and text messages, not to mention the inherent lack of forward progress) that it could not fit.

Communicating Intentionality

Several other studies explored the reasons for this "fight-for-control" between the human and robot initiative with the goal of increasing the user's situation awareness and trust as well as structuring the human and robot initiative to reduce opportunities for conflict. These studies indicated that the joystick itself was a confounding factor. When users move a joystick, they expect an immediate and direct response from the robot.

Although it is possible to mitigate user frustration and error through the use of interfaces that effectively communicates robot intent, the reality is that we should not use a metaphor focused on direct motion control for an interaction method that allows the robot to control its own motion. Rather, a more appropriate interface would involve communication of intentionality. Once human and robot begin to communicate at the level of goals and intentionality, new metaphors become necessary.

What might these new metaphors be? First we must consider the fundamental inputs and outputs. The joystick communicates directional motion. With higher levels of autonomy on board the robot, we want a metaphor that assumes a basic level of robot competency (e.g., navigational basics).

Ideally, the metaphor should also allow us to specify, on the fly, what level of initiative is permitted.

Perhaps a parent-child relationship is a useful metaphor? Certainly, the parent does not expect to directly control every motion and action of the child as they would a marionette doll. This metaphor may be very appropriate for robotic systems that are engineered to learn through imitation and/or reinforcement learning. On the other hand, this metaphor may not be appropriate for battlefield robots where predictability and reliability are paramount. Rather, a possible metaphor for a battlefield robot is the relationship between a bomb sniffing dog and its operator. The human sends the dog out without directly indicating exactly which path the dog should follow or how it should get around obstacles. Instead the user may indicate a direction in which the dog should start the task. Explosive ordinance disposal (EOD) personnel that use bomb sniffing dogs accept an 80 or 90% performance band and are fairly tolerant if the dog is fatigued, confused or distracted. They may infer health from the dog's performance, but don't need to access heart rate, vision, or collision avoidance information. They observe the dog's trajectory and await the outcome. This may be an effective metaphor for human-robot interactions because, in general, we expect dogs to make good use of their acute sensing abilities, know their own limitations, navigate effectively and support mixed-initiative teaming. Likewise, we do not expect dogs to exhibit high level cognitive functions such as solving complex problems or engaging in semantic dialogue. Effective metaphors must communicate limitations as well as capabilities.

PROTOCOL ANALYSIS: A USER'S TALE

One of the more powerful techniques available to aid the design process is protocol analysis. Simply stated, operators talk their way through solving a problem and the designer notes what information or feedback is required; the starting and stopping rules; the sequence of the tasks performed; what defines success; and what job performance aids are required.

For a review of protocol analysis see for example, Ericsson and Simon [16]. For our purposes, protocol analysis can be used to help develop an appropriate interaction metaphor. Protocol analysis insures that the design is mapped to the operator's mental model of what is to be accomplished, i.e., the goals for successful human-systems performance. The dialog obtained from protocol analysis is incredibly useful for the initial design phase or operational test and evaluation phase of the system design life cycle. Many experimenters routinely gain design insights during the subject debrief phase of a laboratory or field study.

What, then, are the tales that users tell? Over the last five years, we have heard a number of comments from disparate user groups that can be used to drive user-centric design. One example is that soldiers at the U.S. Army Maneuver Support Battlab in Ft. Leonard Wood, MO describe their ideal robot as one that can act as a taxi cab for a variety of sensors. After further questioning, it became apparent that the soldiers who wanted a taxi-cab literally meant that they want to be able to plug and play different sensors onto the robot (as a taxi-driver would change passengers) and then specify a destination within the map. Further questioning

indicated that they also wanted a "taxi-driver" in the form of automated software to accomplish the navigation to the destination. The soldiers want to trust the robot to take the sensor to that destination without any further input or user interaction.

The metaphor of the taxi focuses the problem on navigation and bounds the expectation of the user so that they do not expect sensor data analysis or dynamic coverage algorithms, but rather simple transportation. For more complex tasks, such as landmine detection, the robot behavior must involve a closer partnership with the soldier. The robot must be alert to danger, navigate carefully and will be depended upon to protect the human. The dog metaphor indicates to the soldier that the robot, like a dog, cannot be expected to make high-level decisions and must be kept on task. The selection of a metaphor for the end user can inform the design and implementation of levels of autonomy, interface design, and communications.

Understanding the Role of Metaphor

It is possible that part of the metaphor's power lies in the fact that metaphors activate past physical experiences that have been coded in our bodies. Thus, they are cross-modal in nature. This was noted by Lakoff and Johnson [17] and again in Erickson [18]. For a thorough review, see Saffer [19]. Hopkins and Fishwick [20] in reviewing the use of metaphor as an adjunct to modeling and implementation note that metaphors can imbue abstract ideas with concrete properties, thus making the abstraction more accessible. The key to the power of the metaphor lies in linking new ideas to well-understood objects and processes.

Work on the use of metaphors and its ties to linguistics have been championed by Lakoff [21] and Lakoff and Johnson [17]. As Hey [22] points out, we "sign on" to a computer system, we "run" a program, we use the "desktop" and the "recycle bin" or "trash can." Gibbs [23], a noted authority on figurative language and cognitive semantics, makes the point even more specifically. He maintains that there is ample empirical evidence to establish that metaphor is not just a part of speech, but is a part of people's ordinary conceptual styles, i.e., how we think. School children are taught metaphoric analysis in grade school. Psychologists often attempt to support counseling through the use of metaphors. Clients also use metaphors, and to some extent, it is the sharing of the two sets of metaphors that helps to facilitate change. One of the better characterizations of metaphorically induced change in the counseling process is found in Lyddon *et al.* [24].

In order to be effective, the metaphor must relate two objects or processes in a clear well-understood way. Metaphors are ubiquitous and are evidenced in every day life; in the choices we make, the behaviors we exhibit, and the ways in which we experience the world around us. The strongest metaphors show the least amount of variability among subjects tested.

Can the Use of Metaphor Help Us to Name, Frame, and Identify Design Solutions?

Schon [25] speaks of metaphor in design as a response to complex situations wherein the metaphor is selected or constructed on the basis of: 1) eliciting one's attention and 2) is named to frame the situation (which sounds slightly metaphorical itself). The key is to select a *limited* set of salient features and use this to orient our understanding and response.

Metaphoric thinking is difficult for the same reason that it is effective -- because it links two unfamiliar domains to reveal in each domain, elements of the other. The juxtaposition requires novel and creative thinking which occurs regularly in literature and storytelling. What really needs to happen is that robot developers need to become better storytellers.

EXPERIMENT

To understand the value of metaphor, it may be beneficial to consider a particular human-robot interaction experiment performed at the INL which evaluated the use of two new tools for navigation and exploration [26]. The new tools allow the user to "drive-by-intent." In particular, the user no longer needs to drive the robot or operate the camera payload directly. Instead, the user drops attention hotspots to which the vehicle will travel and then focus its attention. The target tool (as shown in Fig. (**1**)) allows the human to drive a target ahead of the robot instead of driving the robot itself. In this manner, the robot decides how to navigate within the local environment while the human is permitted to decide where to navigate to. The robot's autonomy is increased while the human is still engaged in the overall task.

Fig. (**1**) also shows a visual hotspot tool that allows the user to specify where in the map the camera should focus. The robot camera will automatically pan in order to focus on the indicated region regardless of how the robot platform itself is moving. Since many of our environments are relatively small, the robot is configured to always focus on the place of interest regardless of its distance from the robot. To change where the robot is looking, the operator simply moves the "look-at" icon. In this manner, the user can explore an environment while minimizing workload devoted to driving the robot or panning the camera.

The important thing in terms of understanding the power of metaphor is that the new tools do not provide new or different behaviors with respect to the robot capabilities, but rather support a new way of thinking about the robot and the task. The new metaphor engaged by the visual hotspot and target is different primarily because it is map-centric instead of robot-centric.

Moreover, the new tools can be expressed in terms that the user can easily understand. The visual hotspot tool is a way for the human to say "look over here!" Likewise, the functioning of the target can be expressed as "taking a dog for a walk on a leash." The dog doesn't always walk in a straight line, but rather uses path planning and obstacle avoidance to head towards the moving target that is the owner.

The metaphor can be extended to explain the multiple levels of robot behavior and initiative. The further the target is moved from the robot (i.e., the longer the leash), the more initiative is granted to the robot in terms of which behaviors can be used. For instance, when the target is close, the robot does not perform path planning, whereas further targets involve the robot planning a route. The user need never under-
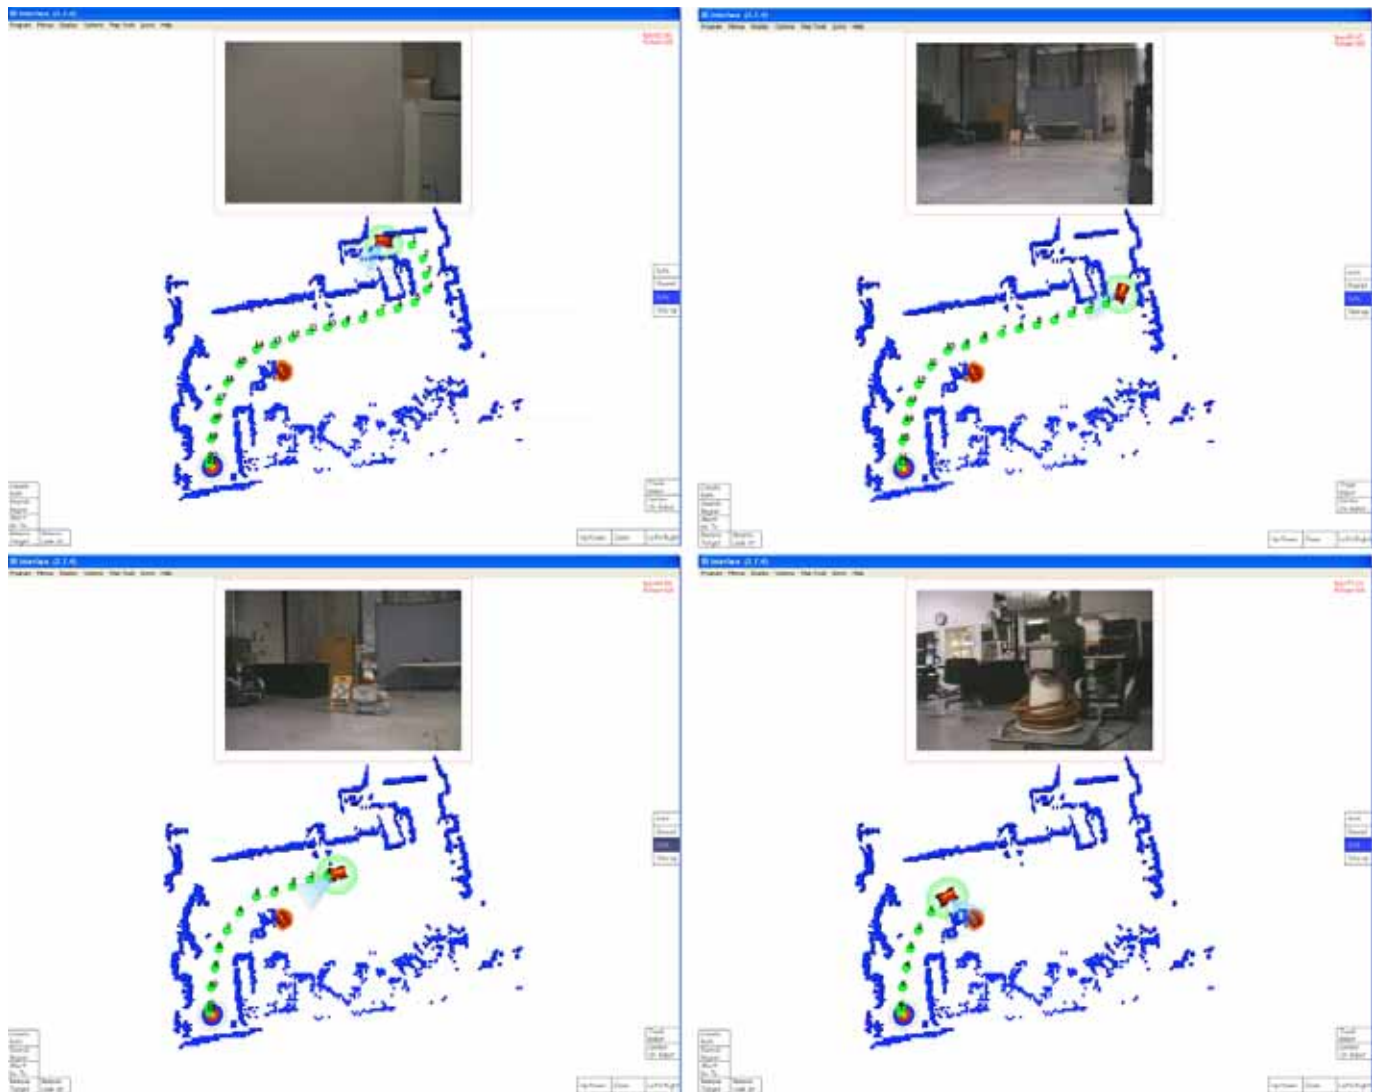
Fig. (1). Operator control unit showing the new "drive by intent" tasking tools. In the top left image, the operator has placed the target tool at the opposite side of the room and the attention hotspot at a place of interest within the room. The subsequent images illustrate the movement and orientation of the camera as the robot autonomously moves through the environment. To change where the robot is looking or the robot's destination, the operator simply moves the target icon or the attention hotspot.

stand the complexities of behavior orchestration on board the robot or how this orchestration changes depending on the length of the leash.

Experiment Design

In order to test these new tools, 153 novice participants were given the task of searching an environment for nine stuffed animals hidden throughout the environment. The environment was 1000 sq. feet with a few partitions and a maze-like feel. The stuffed animals were hidden at various corners of the environment such that the operator would not discover them by simply following the maze, but would have to pan the camera or spin the robot to see them. The starting orientation of the robot alternated between participants, so each participant explored the map differently than the previous participant. Each participant was given one of three conditions for controlling the robot.

In the first condition (joystick), operators were asked to drive the robot and operate the camera using a joystick for both sets of commands. In the second condition (target), ope-

rators used a computer mouse with the target icon to specify navigational goals. To control the movement of the camera, operators used the mouse to drag and drop the video image at the top center of the screen. The joystick was not available to participants in this condition. In the third condition (icons), operators used a computer mouse with both the target mode and the visual hotspot.

The only difference between conditions was how the participants interacted with the interface. Participants were given an a priori map of the environment because we did not want to test how well operators could discover the structure of the environment; rather, we wanted to test how well they could find items of interest. Each participant was given two minutes to find as many items as they could. Participants were not allowed to view the robot while performing the test.

Results

Analysis in this experiment consisted of measuring global task performance for a search task in terms of number of items found within the allotted time. Additionally, the

movement of the camera was measured as a means to understand the underlying behaviors that affected performance. The results indicate that operators found at least 20% more objects using the icons condition than either of the other two conditions. Furthermore, with the icons condition, operators used the camera at least 90% more than the other two conditions. The results are summarized in Table **1**. Detailed results and discussion can be found in [26].

Table 1.   Summary of Results for the Three Conditions

	Joystick	**Target**	**Icons**
Items found	4.02	4.24	5.12
Camera Use (degs/s)	2.79	5.35	11.38

Discussion

What this experiment demonstrates is that changes in design metaphors produced clear differences in human-robotic performance, namely that using the metaphors increased the operator's use of the camera which led to improved task performance. The intentionality tools that control driving and camera movement were based on simple metaphors that hide the "how" of the underlying system and support the operator's development of an adequate mental model of the interaction. Specifically, the operator is allowed to focus on commands such as "look here" and "go here" rather than "turn like this" or "move like this". By eliminating requirements for the user to remember the underlying details of the robot system we allow them instead to focus their attention on the task at hand. Understanding the user's tale of how they conceptualize the problem space and the degree to which they can accept partnership with the robot supports selection of the correct user-centric metaphor.

CONCLUSION

A Posse Ad Esse - From Possibility to Actuality

Simple design metaphors can be used to name and frame the user interface and the operator's perception of the human-robot interaction. Some of the metaphors we have considered are particularly useful in helping to transform operator thinking from the expectations of the traditional master-slave relationship into more of a partnership. We have shown that metaphors can provide new models for how to structure and interleave human and robot responsibilities and initiative that lead to improved performance over more traditional approaches. Using metaphors such as a "dog on a leash" or "human as back seat driver" may dramatically change the expectations of the user and their willingness to trust the robot and let the robot "reason" about the task, environment, and successful attainment of the goal(s).

As robot behaviors become increasingly complex, it is imperative that we find a means to hide the complexity of the system from the user while still keeping users informed and allowing them to be part of the action. By doing so, 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 using appropriate metaphors can provide the benefits of enhanced operator and system performance, decreased operator workload, and more accurate operator expectations.

In the future we plan to conduct studies where the level of workload is used as one of several factors to assess the advantage of particular metaphors over others. Experimental methods exist for determining preferred metaphors and these methods are supported through application of human factors tools such as task analysis and protocol analysis.

Of course, no one metaphor will apply to all tasks, operators, or domains. For each application, research and field testing of different metaphors is suggested. Future collaborative human-robotic design should successfully link operational goals, cognitive processes, and strategic and tactical behaviors. The solution to optimal performance and meeting mission objectives lies in developing user-centric metaphors that build appropriate trust while providing the bounds of expectation.

DISCLAIMER

This paper was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for any third party's use, or the results of such use, of any information, apparatus, product, or process disclosed in this paper, or represents that its use by such third party would not infringe privately owned rights.

REFERENCES

[1]   D. J. Bruemmer, D. A. Few, R. L. Boring, J. L. Marble, M. Walton, and C. Nielsen. "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 2005.

[2]   R. C. Arkin, *Behavior-Based Robotics,* MIT Press, Cambridge, MA. 1997.

[3]   M. Desai and H. A. Yanco, "Blending human and robot inputs for sliding scale autonomy". *Proceedings of the 14th IEEE International Workshop on Robot and Human Interactive Communication (Ro-MAN),* Nashville, TN, August 2005.

[4]   B. A. Maxwell, W. D. Smart, A. Jacoff, *et al.* "AAAI Robot Competition and Exhibition". *AI Magazine*, 25(2): 68-80, 2003.

[5]   D. J. Bruemmer, D. A. Few, M. C. Walton, *et al.* "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, 3-6, January 2005.

[6]   J. L. Marble, D. J. Bruemmer, and D. A. "Few, Lessons learned from usability tests with a collaborative cognitive workspace for human-robot teams". In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*, Washington, D.C. pp. 5-8, October 2003.

[7]   Nourbakhsh. "The wicked problem and interaction evolution", *CMU Robotics Institute*, Fall Semester, 2005.

[8]   M. A. Goodrich and E. R. Boer, "Designing human-centered automation: Tradeoffs in collision avoidance system design". *IEEE Transactions of ITS*, 1(1): 40-54, 2000.

[9]   R. Parasuraman and V. Riley, Humans and automation: Use, misuse, disuse, and abuse. *Hum. Factors*, 39(2): 230-253, 1997.

[10]   J. D. Lee and N. Moray, Trust, self-confidence, and operator's adaption to automation. *Int. J. Hum. Comput. Stud.*, 40(1): 153-184, 1994.

[11]   D. B. Kaber, and M. R. Endsley, The effects of level of automation and adaptive automation on human performance, situation awareness, and workload in a dynamic control task. *TIES*, 5(2): 113-153, 2004.

[12]   J. D. Lee and K. A. See, Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1): 50-80, 2004.

[13]   D. I. Gertman, S. Novack, J. Marble, "Culture representation in human reliability analysis". in *Proceedings opf the Inter-*

service/Industry Training, Simulation, and Education Conference (I/ITSEC), Orlando Florida, December 10-14, 2006.

[14]   T. Sheridan, *Telerobotics, automation, and human supervisory control*, MIT Press, Cambridge, MA, 1992.

[15]   B. Tognazzini. "First principles of interaction design, Nielsen Norman Group", in *Interaction Design Solutions for the Real World*. http://www.asktog.com/basics/firstPrinciples.html, retrieved February 2007.

[16]   K. A. Ericsson and H. A. Simon, *Protocol Analysis: Verbal Reports as Data* MIT Press, Cambridge, MA, 1993.

[17]   G. Lackoff and M. Johnson, *Metaphors We Live By*, University of Chicago Press, Chicago, Illinois, 1980.

[18]   T. D. Erickison. "Working with Interface Metaphors", in B. Laurel (ed.) *The Art of Human Computer Interface Design*, Addision-Wesley, 1990.

[19]   D. Saffer, The role of metaphor in interaction design. Masters Thesis submitted to the School of Design, Carnegie Mellon University, 2005.

[20]   J. F. Hopkins and P. A. Fishwick. A three dimensional human agent metaphor for modeling and simulation. *Proc. IEEE*, 89(2):131-147, 2001.

[21]   G. Lackoff, "The contemporary theory of metaphor" in A. Ortony (ed.) *Metaphor and Thought, 2nd editior,* Cambridge University Press, Cambridge, England, 1993.

[22]   J. Hey, *Metaphors we design by*. Report presented in support of Neural Theory of Thought and Language, Berkeley University, Berkeley California, 2004.

[23]   R. W. Gibbs. *"The Poetics of Mind: Figurative Thought, Language, and Understanding"*, Cambridge University Press, Cambridge, England, 1994.

[24]   W. J. Lyddon, A. L. Clay, and C. L. Sparks. Metaphor and change in counselling. *J. Couns. Dev.*, 79:269-275, 2001.

[25]   D. A. Schon. "Generative Metaphor: A perspective on problem-setting as social policy" in *Metaphor and Thought, 2nd edition*, A. Ortony (ed.), Cambridge University Press, Cambridge Mass, November 1993.

[26]   C. W. Nielsen and D. J. Bruemmer. "Hiding the system from the User: Moving from Complex mental models to elegant metaphors", in *IEEE International Symposium on Robot and Human Interactive Communications, (RO-MAN)*, Jeju Island, South Korea, August 2007.