INL/CON-07-13035
PREPRINT

Hiding the System From the User: Moving From Complex Mental Models to Elegant Metaphors

Ro-man 2007

Curtis W. Nielsen
David J. Bruemmer

August 2007

Hiding the System from the User: Moving from Complex Mental Models to Elegant Metaphors

Curtis W. Nielsen and David J. Bruemmer

Robotics and Human Systems, Idaho National Laboratory, Idaho Falls, ID, USA
e-mail: {curtis.nielsen, david.bruemmer}@inl.gov

***Abstract*— In previous work, increased complexity of robot behaviors and the accompanying interface design often led to operator confusion and/or a fight for control between the robot and operator. We believe the reason for the conflict was that the design of the interface and interactions presented too much of the underlying robot design model to the operator. Since the design model includes the implementation of sensors, behaviors, and sophisticated algorithms, the result was that the operator's cognitive efforts were focused on understanding the design of the robot system as opposed to focusing on the task at hand. This paper illustrates how this very problem emerged at the INL and how the implementation of new metaphors for interaction has allowed us to hide the design model from the user and allow the user to focus more on the task at hand. Supporting the user's focus on the task rather than on the design model allows increased use of the system and significant performance improvement in a search task with novice users.**

I. Introduction

Six years ago, researchers at the Idaho National Laboratory (INL) began developing a suite of behaviors that was intended to provide dynamic vehicle autonomy for robots in order to support human operators. A series of early experiments showed the potential for these behaviors to improve performance by reducing human error and increasing various measures of task efficiency [2]. In addition to these benefits, these experiments also illustrated the opportunity for operator confusion regarding robot behavior and robot initiative. Early 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 counteract robot initiative, which usually resulted in a significant decrease in performance.

Participants in the experiments could often be divided into two groups. One group demonstrated that they understood and trusted the robot behaviors and robot initiative which resulted in significant performance improvements over a baseline system. The other group of participants demonstrated that they were confused by the robot initiative and actually suffered a performance decrease when compared with their performance on a baseline system. This experiment and subsequent ones like it illustrate that operator trust is a major factor in operational success and that this trust is significantly impacted when the user makes incorrect assumptions about robot behaviors and initiative.

A user's assumptions about a system, as described by Norman [9], are based on the user's previous experiences and what the user perceives they can do with the system. Norman's description very closely coincides with Endsley's three levels of situation awareness: namely, *perception* of what the system offers, *comprehension* of how the system can be used, and *projection* of how the interaction will affect the system [4]. Taking Norman and Endsley's descriptions in conjunction would suggest that fundamental elements of situation awareness are related to the accuracy of the user's assumptions about how the system works. Throughout this discussion, we refer to the operator's *mental model* as the operator's assumptions and expectations of how commands are issued, acted upon, and visualized by the system.

More specific to the mobile robot domain, Murphy has said that "More sophisticated mobility and navigation algorithms without an accompanying improvement in situation awareness support can reduce the time spent on a mission by no more than 25 percent [10]. Another way to view this statement in light of the current discussion is that even if robots are improved or made more intelligent we should not expect a significant increase in overall task performance unless there is also a correlating improvement in the user's mental model of how to use the robot.

The problem is that as the robot is endowed with more sophisticated behaviors, algorithms, and sensors, the inherent complexity of the human-robot system increases which in turn increases the difficulty for the user to develop a correct mental model of the robot system. For example, Dennett points out that when interacting with intelligent systems the user needs to have a better understanding about the intentional stance of the system [3]. With intelligent robots, this means that the operator must understand when the robot will perform actions of its own initiative, what those actions will be, and how those actions will affect the system, task, and environment. Even in systems that increase functionality and options without necessarily increasing "intelligence", the operator is still required to understand how the various modes and settings of the system will affect the interaction.

The challenge when developing human-robot systems is to increase the capabilities, behaviors, and initiative of the robot while supporting the development of a sufficient mental model [6, 9]. The development of sufficient mental models is facilitated when the robot design model is hidden from the operator and replaced with metaphors that the user readily understands. These metaphors must support the complexity of the system design but only reveal necessary choices and information to the operator. This paper describes the solutions that have been developed by the INL in collaboration with Brigham Young University (BYU) to allow the increase in design model complexity for more capable robots while supporting the operator's development of correct assumptions and expectations for the robot system.

II. INCREMENTAL COMPLEXITY

To begin, it is beneficial to illustrate how the incremental development of mobile robots can lead to more complex interfaces and interactions, thereby minimizing the advantages of sophisticated systems by complicating the user's interaction with the system. In this section, we review the iterative development of the INL robot intelligence kernel [2], which we imagine is similar to many robot development cycles. The discussion illustrates changes made to the robot and how those changes were reflected on the interface. As the robot and interface changed so did the operator's cognitive workload and mental model of the system.

The first phase of development allowed the movement of a robot in response to joystick commands. The interface to the robot consisted of a joystick whose two main axes mapped to the forward and turn velocities on the robot. Information from the robot was shared with the operator from a video camera attached to the front of the robot. The operator's mental model was that the robot would move in the direction that the joystick was pushed. Video provided some information about how the robot moved and where the robot was within the environment however, the operator was required to remember previously visited areas of the environment.

Next, a pan-tilt unit was added to the camera on the robot. The interface was augmented to allow the operator to view and control the pan and tilt angles of the camera. The pan-tilt angles of the camera were illustrated using iconic representations and the control of the pan-tilt camera was performed by buttons on the interface, a second joystick, or a more sophisticated joystick that could support both robot navigation and camera movement. The operator's mental model now demanded support for how the camera's movements will affect the representation of the interface particularly when the robot is moved while the camera is off-center. The cognitive workload added to the user includes remembering which commands control which aspects of the robot and camera and remembering the camera's orientation relative to the robot.

To support the operator in more difficult tasks, the robot was further equipped with a variety of sensors that were designed to help the operator understand the robot's situation within the environment. Some of the sensors were to provide information about the robot itself (e.g. heading; attitude; pitch; roll; forward velocity; turn velocity; battery power, communications activity, GPS, proximity sensors, IR, and overall health of the robot) and some were to provide information about the structure of the environment (e.g. laser range finder, sonar). Upon adding these sensors to the robot, the interface was also augmented to show the operator the measurements and status of all the sensors. The operator's cognitive workload increased as the operator was required to interpret the meaning of the observed information. The operator's mental model also required an understanding of how the information could be used to effectively control the robot.

In an effort to better support the performance of navigation tasks, the algorithms and robot capabilities were increased to support behavioral intelligence. To utilize the behavioral modes of the robot, the interface provided a set of buttons that the user could select which would change either the level of initiative of the robot or the manner in which the robot was tasked by the operator. Some of the modes included: teleoperation; safe; shared; autonomous; pursuit; escape; track; search region; patrol region; waypoints; and go-to. With the different behaviors, the operator's mental model demanded an understanding of how each behavior or operational mode affected the movement of the robot, specifically, the idiosyncrasies between how the different modes respond to commands and the environment.

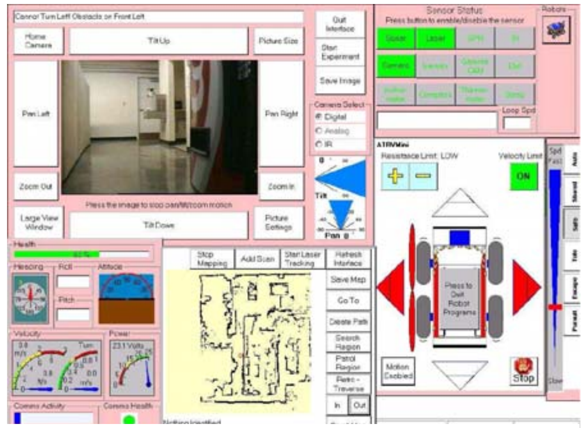

Figure 1. The original INL interface.

The corresponding interface that was developed for this level of robot functionality included video, map, camera, vehicle status, sensor status, and obstruction modules. The interface is shown in Figure 1.

As developers of the robot system design it seemed beneficial to illustrate the new robot capabilities directly to the user, believing that flexibility was the key to supporting disparate users and enabling a variety of missions. Furthermore, from the developer's perspective, this type of interface is helpful to verify the functionality of the low-level robot systems and to debug and solve technical difficulties with the system. Why should the development be done any other way?

The problem was that we were developing the system from the designer's perspective, not from the user's perspective. In fact, despite the capabilities of the robot, early criticism that we received from colleagues, domain experts, practitioners, and novice users was that the interface was simply too complex. There were too many options, too many disparate perspectives and too many separate perceptual streams. Even after describing the system to intelligent, competent individuals, they would often settle on a single mode of interaction that they understood well, regardless of the actual efficiency of the interaction for the specified task.

While some may say that the interface could still be manageable with training, we realized that in order to meet our research goals we required even more system complexity. The interface needed to be expanded to support multiple robots, which could require a replication of the complete display for each robot. Moreover, as we expand the domains and tasks, we would require the interface to support communication with and control of multiple UAVs and unattended ground sensors. In terms of information sets, the operator may also require information relating to occupancy grids, chemical and radiological plumes, explosive hazards

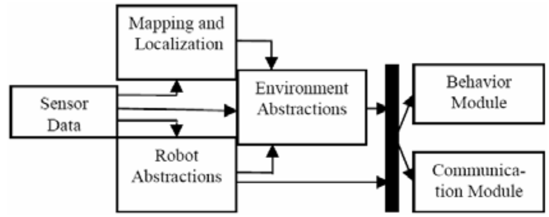

Figure 2. Information flow through the INL robot intelligence kernel.

detection, 3D range data, terrain data, building schematics, satellite imagery, real-time aerial imagery, and representations of movable mechanical arms. Clearly, even with training, the system would overwhelm users if we were to simply implement the new technology on the robot and naively show all the available information to the operator.

The critical question that we have worked to solve is how to support an ever increasing number of perceptions, actions, and behaviors while hiding from the operator the details of the design of the system that are not necessary for a particular task. Such a solution should minimize the cognitive workload of the operator and support the operator's development of an appropriate mental model of the robot capabilities by reducing the complexity of the interface and the human-robot interaction. In order to hide the robot design complexity from the user, we focus on three areas: using data abstractions on the robot, fusing data sets with the interface, and providing seamless autonomy interactions.

III. Using Data Abstractions

Most mobile robots have a variety of sensors that can be used to relay information about the status of the robot or measurements about the environment. The user should not have to sift through this raw data or expend significant cognitive workload to comprehend or correlate the information. The first step to facilitating efficient human-robot interaction is to provide an efficient method for the robot itself to combine and filter the raw sensor data into basic meaningful abstractions before the information is ever shared with the operator. The combination and filtering that takes place can happen at multiple levels. For example, information about the status of the robot could be combined into a numerical value representing the health of the robot. Furthermore, instead of conveying raw range data from various sensors, a map could be built and shared with the operator. The system used at the INL provides layers of abstraction that underlie all robot behavior and communication with the operator. The flow of information as it is received by sensors, abstracted, and sent to the operator is illustrated in Figure 2.

To give an idea of the complexity that is being abstracted by the INL robot system before any information is transmitted to the user, consider the list of previously mentioned sensors with inputs available to the robot. Information used to determine the robot pose could come from GPS, SLAM algorithms, inertial sensors, and wheel encoders. Range data (sonar, laser, IR) could be used to create an ego-centric robot abstraction that describes the proximity of obstacles around the robot from the robot's perspective and can further be used to build an exocentric map and facilitate localization of the robot within the map. In these examples, it is much more beneficial to have the robot perform the data combination and correlation than the operator.

IV. Fusing Data Sets with the Interface

Even though much of the raw sensor information is abstracted on the robot, there are still relationships that exist between the different sets of data that could be exploited with the user interface. Using the interface to combine different sets of information into a common reference frame is another way to reduce the operator's mental model of the system.

In collaboration with BYU, a pseudo-3D interface was developed that visually rendered video, map data and camera orientation into a seamless, scalable representation that could be zoomed in or out to support varying levels of operator involvement [11]. As shown in Figure 3, a traditional robot interface might have at least four different frames of reference that the operator must cognitively correlate in order to comprehend the relationships between the different sets of information. In contrast, the 3D interface (Figure 4) has a single reference frame from which all the information is viewed by the operator. Instead of illustrating to the operator the details of the system (that there are multiple sets of information that may or may not be related), the integration of video, camera pose, robot pose, and map tell the user how the information is related. The 3D interface has been shown to significantly increase the operator's ability to navigate the robot through planar maze-like environments [7] while also supporting the use of a pan-tilt-zoom camera [8].

In addition to providing a single reference frame for the

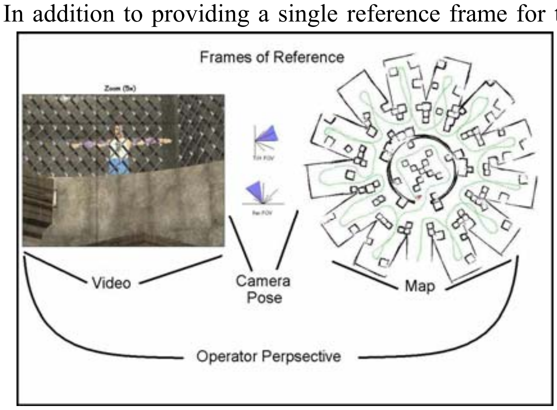

Figure 3. Frames of reference with a 2D interface.

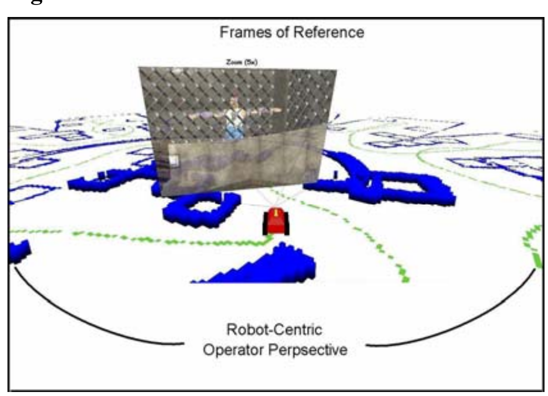

Figure 4. Frames of reference with a 3D interface.

operator to view the information, we also believe that the 3D interface is successful in navigation tasks because it supports the development of an accurate mental model of the interaction better than the traditional 2D interface. In particular, the 3D interface supports what we term the "correlation of action and response" which is similar to Endsley's 3rd level of situation awareness (projection) [4]. In order for a novice user to develop an appropriate mental model about how the robot responds to commands, it is beneficial for the system to reinforce the user's assumptions by behaving the same each time the operator issues the same command.

What actually happens with the robot is likely to be the same irrespective of how the information is displayed on the interface, however, the manner in which the interface illustrates the movement of the robot can affect the operator's trust and confidence with the system. For example, with the 3D interface, when the user moves the robot forward it shows the robot moving "into" the display because the user's perspective is tethered to the robot. In cases where the video is off-center, the robot still moves "into" the display (not "into" the video). In contrast, consider the 2D interface. When the video is in front of the robot, and the robot is moved forward, the robot appears to move "into" the video. However, if the video is to the side of the robot, movement forward appears to "slide" the video. In both cases the same action is performed however, the interface changes differently.

Maps on traditional interfaces present similar challenges. When the robot is facing up with respect to the map, then forward motion moves the robot icon "up" with respect to the display. However, if the robot is facing any other direction, the illustration of the robot moves in the direction the robot is facing, not the direction the joystick was moved. This means that the operator is required to remember the details of the system, including the camera pose, robot pose, and map orientation in order to understand why the interface is not changing consistently each time. Trust in the system can erode when some of this information is forgotten and the interface appears to behave unexpectedly.

Although the 3D interface was beneficial for teleoperation tasks, when the interface and robot system were tested by search and rescue personnel in a search task, the movement of the video around the robot icon on the interface was hard for them to use and they requested that the video be kept stationary [12]. The recommendation was taken and a hybrid 2D-3D interface was developed (See Figure 5). To maintain as much support for the operator's understanding of the orientation of the camera with respect to the robot, we continue to render the field of view of the camera in the 3D portion of the display, but the video is placed in a static 2D perspective above the 3D map.

Simplifying the interface by correlating and fusing related information is a step in the right direction, however sensor fusion is not sufficient to actually change the nature of the interaction itself or the fundamental inputs and outputs between the human and the robotic system. To reduce the interaction complexity, we must be able to not only abstract the robot physical and perceptual capabilities, but also the elemental behaviors and behavior combinations necessary to accomplish a sophisticated operation.

Figure 5. Illustration of the "look here" and "go here" metaphors.

V. PROVIDING SEAMLESS AUTONOMY

Earlier we mentioned that adding intelligent behaviors to the robot was challenging because on the one hand it seemed to improve performance on some tasks [5], but on the other hand it also made the system difficult to use because the operator had the responsibility of choosing the appropriate level of autonomy for the given situation. The problem with leaving the operator with the responsibility to change autonomy levels is that the operator may not recognize the need to switch modes of autonomy or may not understand how the autonomy level will relate to the task. In general, requiring the operator to have an understanding of the appropriate context for each autonomy mode is cognitively challenging and prone to error [1]. Even when the robot is given the ability to automatically switch between the discrete autonomy modes based on its own assessment of the situation, the operator often feels frustrated and confused about what the robot is doing and why it usurped control, which leads to confusion and, often, a fight for control between the operator and the robot. This fight for control generally emanates from poor communications between the robot and the human as to why the decision was made and how the decision is being actuated.

Our solution to the fight for control is to support seamless levels of interaction based on the operator indirectly controlling the intent of the robot as opposed to directly controlling the movement of the robot. The indirect control is performed by allowing the operator to define the desired intent from a human-centered reference frame. Instead of selecting between "discrete levels of autonomy" the operator uses locally defined intentions to exert more direct control over the robot while more globally defined intentions augment robot initiative and reduce the operator's navigational requirements. The advantage of indirect control through specifying intent is that the "fight-for-control" problem with direct control is eliminated because issues with communications latency and bandwidth requirements are mitigated. Furthermore, the interaction methodologies do not change for higher or lower levels of robot-initiative.

The metaphors we used are to allow the user to point to the desired area of interest and allow the operator to tell the robot to "go here" or "look here". From the interface, this metaphor is actuated by using a mouse and placing either a target icon (for "go here") or a look-at icon (for "look here") at the place of interest. Figure 5 illustrates the interface as the robot starts a complex task of: "look-at this place while moving to that place". What these metaphors really do for the user is to hide the "how" of the underlying system and replace it with a powerful and easily understood metaphor that supports the operator's development of a correct mental model of the interaction.

VI. Evaluating the New Metaphors

To determine whether or not the new metaphors actually improve performance on a complex task a user-study was performed using visitors to the 2006 INL Science Expo.

A. Experiment Design

The purpose of the experiment was to test different modes of human-robot interaction in a search and identify task. Participants used an ATRV-mini robot equipped with a pan-tilt camera, laser range finder, and the INL robot intelligence kernel to search an environment for nine hidden stuffed animals. The environment was a simple maze as shown in Figure 6. The animals were hidden at the places indicated by circles and the robot start location is represented by the rectangle on the map. The starting orientation of the robot alternated between participants, so every other participant explored the map differently than the previous participant. The animals were hidden on the ground such that they generally could not be stumbled upon by simply driving through the maze; rather the camera or the robot needed to be rotated from side to side for the operator to see the animals. It was not necessary to use the tilt of the camera to see the animals.

In accordance with a between-subjects experimental design, each participant was given one of three conditions for the experiment. In the first condition (which we term the "joystick" condition), operators were asked to drive the robot and operate the camera with a joystick. By moving the joystick the operator controlled the navigation of the robot. The pan-tilt camera was controlled by moving the hat on top of the joystick. This condition amounted to controlling the camera and the robot through traditional teleoperation.

In the second condition (the "icon+video" condition), operators used the mouse to drag 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 this condition the camera controls were teleoperated, like the joystick condition, however the navigation of the robot was performed with the new "move here" metaphor.

In the third condition (the "icons" condition), operators controlled the movement of the robot by using the computer mouse and dragging the target icon to the desired robot destination. The pan of the camera was controlled similarly by dragging the look-at icon. In this condition the tilt of the camera was disabled. The joystick was not used by participants in this condition. In this condition both the camera and navigation of the robot utilized the new "look here" and "move here" metaphors.

In each of the conditions, the participants had the same perspective of the robot's environment, the same interface, and the same software and algorithms running on the robot. The only difference between conditions was how the participants interacted with the interface according to the aforementioned conditions. 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 the hidden items. Each participant was told how to control the robot with the condition they would be using and that they had two minutes to find as many of the stuffed animals as possible. Participants did not control the robot prior to the experiment.

B. Analysis Methods

Analysis in this experiment consisted primarily of task performance. Specifically, how many items each participant found as indicated by their verbal acknowledgment that they observed the item. Secondarily, we measured average velocity, distance traveled, and how much the camera was used as a means to understand how the operator utilized the system to achieve their performance. The significance of results are based on an unpaired t-test with a sample size of n = 51 unless otherwise specified.

C. Results

There were 153 students that participated in the user-study. Most of the participants were between the 5th and 7th grades. Results indicate that operators with the "icons" condition found, on average, 27% more items than operators with the "joystick" condition (X_{icons} =5.12, $X_{joystick}$ =4.02, $p< 0.01$) and 20% more items than operators with the "icon+video" condition (X_{icons} =5.12, $x_{icon+video}$ = 4.24, $p< 0.05$). There was not a significant difference in the number of items found between the joystick and icon+video conditions (p =0.622). Figure 6 shows the comparison between the numbers of items found with the different conditions.

There was not a significant difference between the average velocities or the distance traveled with any of the three conditions ($p> 0.20$). However, there was a significant difference in how much the camera was moved. Specifically, the camera was moved 308% more with the icons condition

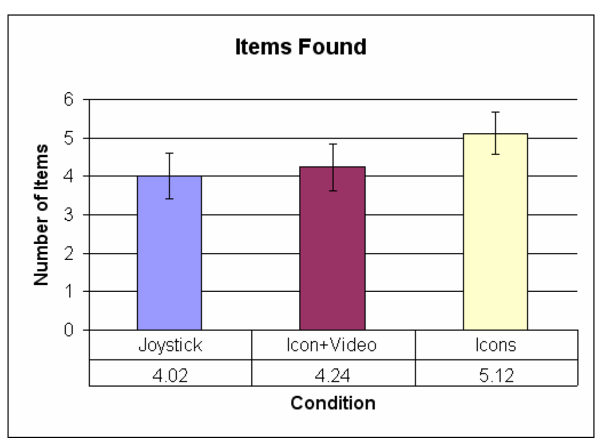

Figure 6. The average number of items found for the experiment (95% confidence).

than with the joystick condition (X_{icons} = 11.38 degrees/s, $X_{icon+video}$ =2.79 degrees/s, $p< 0.001$), and 112% more with the icons condition than the video+icon condition (X_{icons} = 11.38 degrees/s, $X_{icon+video}$ =5.35 degrees/s, $p< 0.001$). Furthermore, the camera was used 92% more with the video+icons condition than the joystick condition ($X_{icon+video}$ =5.35 degrees/s, $X_{joystick}$ =2.79 degrees/s, $p< 0.005$). Figure 7 shows the comparisons of the camera use with the different conditions.

D. Discussion

The joystick condition for this experiment required the most complex mental model to operate the robot. The user had to remember how the robot and camera responded to individual joystick commands which lead to the poorest performance and the least utilization of the camera. When the design of the movement of the robot was abstracted and hidden from the user via the "move here" metaphor, the user was able to spend nearly double the time operating the camera as shown by the icon+video results. Although this approach improved use of the camera, it did not significantly improve the task performance. In the icons condition however, the design of the movement of the camera was also abstracted which allowed for significant improvements in the use of the camera over the other two conditions which also supported improved task performance. By abstracting and simplifying the camera and robot controls, the operator was better able to utilize the system for the search task.

This experiment was performed with novice participants and it should not be assumed that the observed results will immediately transfer to more domain-specific applications. Future work will use domain specific end-users to explore how similar metaphors could be applied to domains such as urban search and rescue or hazardous material detection.

VII. CONCLUSIONS

When designing human-robot systems, it is important to support the user's development of a sufficient mental model of how the robot system should be used. When too much information about the underlying robot design is presented to the operator it can be overwhelming and lead to frustration when the operator does not remember all the details and idiosyncrasies of the system design. We have found that a better approach, in terms of task performance, is to hide the design complexity and instead allow the user to focus on the requirements of the task by providing more powerful metaphors for the interaction. Such a solution will help engender user's trust and confidence in robots as the interaction requirements are simplified and come more in line with the goals and demands of the user and less prominently associated with the goals of the system developer.

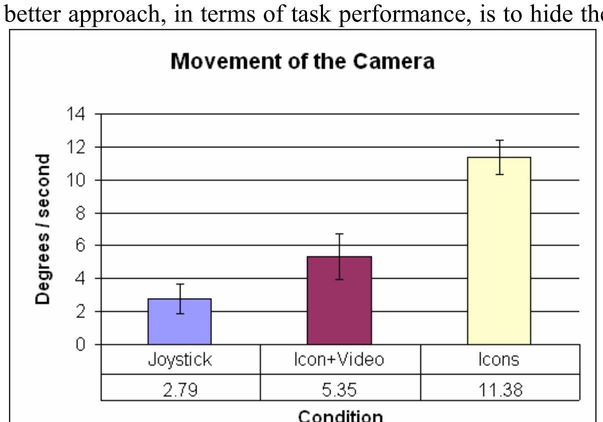

Figure 7. The average movement of the camera for the experiment (95% confidence).

We have shown that simplifying the interaction by using new metaphors can lead to significant improvement in performance in a search task as compared to more traditional interactions. The results achieved to date illustrate viable solutions for single robots and fairly simple tasks. Future work will focus on supporting more complex single robot tasks such as mobile manipulation and will address the operator's ability to command, control, and visualize information from multiple unmanned ground and air vehicles.

VIII. REFERENCES

[1] D. J. Bruemmer, J. L. Marble, D. D. Dudenhoeffer, M. O. Anderson, and M. D. McKay. Mixed-initiative control for remote characterization of hazardous environments. *Hawaii International Conference on System Sciences (HICSS) 2003*, Waikoloa Village, Hawaii, January 2003.

[2] 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, 5(4):505-512, Jul. 2005.

[3] D.C. Dennett. True believers: The intentional strategy and why it works. In A.F. Heath, editor, *Scientific Explanation*. Oxford University Press, Oxford, England, 1981.

[4] Mica R. Endsley. Design and evaluation for situation awareness enhancement. *In Proceedings of the Human Factors Society 32nd Annual Meeting*, pages 97–101, Santa Monica, CA, 1988.

[5] D. A. Few, D. J. Bruemmer, M. C. Walton. Improved Human-Robot Teaming through Facilitated Initiative*. In Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (ROMAN).* Hatfield, United Kingdom, September 6-8 2006.

[6] Michael A. Goodrich. Using models of cognition in HRI evaluation and design. In Proceedings of the AAAI 2004 Fall Symposium Series, Arlington,VA, October 21-24, 2004.

[7] Curtis W. Nielsen and Michael A. Goodrich. Comparing the usefulness of video and map information in navigation tasks. *In Proceedings of the 2006 Human-Robot Interaction Conference*, Salt Lake City, UT, 2006.

[8] Curtis W. Nielsen and Michael A. Goodrich. Testing the usefulness of a Pan-tilt-zoom (PTZ) camera in human-robot interactions. *In Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting*, San Francisco, CA, 2006.

[9] Donald A. Norman. *The design of everyday things*. Doubleday, New York, NY, 1988. Previously published as The psychology of everyday things.

[10] Robin R. Murphy. Humans, robots, rubble, and research. *Interactions*, 12(2):37–39, March + April 2005.

[11] Bob W. Ricks, Curtis W. Nielsen, and Michael A. Goodrich. Ecological displays for robot interaction: A new perspective*. In International Conference on Intelligent Robots and Systems IEEE/RSJ,* Sendai, Japan, 2004.

[12] Holly A. Yanco, Michael Baker, Robert Casey, Brenden Keyes, Philip Thoren, Jill L. Drury, Douglas A. Few, Curtis W. Nielsen, and David J. Bruemmer. Analysis of human-robot interaction for urban search and rescue. *In Proceedings of IEEE International Workshop on Safety, Security, and Rescue Robotics*, Gaithersburg, MD, 2006.