Changing Shape: Improving Situation Awareness for a Polymorphic Robot

Jill L. Drury
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730
+1-781-271-2034
jldrury@mitre.org

Holly A. Yanco
University of Massachusetts Lowell
Computer Science Dept.
One University Avenue
Lowell, MA 01854
+1-978-934-3642
holly@cs.uml.edu

Whitney Howell, Brian Minten, and Jennifer Casper
American Standard Robotics
719 27th Avenue N
Saint Petersburg, FL 33704
+1-727-821-7588
{whowell, bminten, jcasper}@asrobotics.com

ABSTRACT

Polymorphic, or shape-shifting, robots can normally tackle more types of tasks than non-polymorphic robots due to their flexible morphology. Their versatility adds to the challenge of designing a human interface, however. To investigate the utility of providing awareness information about the robot's physical configuration (or "pose"), we performed a within-subjects experiment with presence or absence of pose information being the independent variable. We found that participants were more likely to tip the robot or have it ride up on obstacles when they used the display that lacked pose information and also more likely to move the robot to the highest position to become oriented. There was no significant difference in the number of times that participants bumped into obstacles, however, indicating that having more awareness of the robot's state does not affect awareness of the robots' immediate surroundings. Participants thought the display with pose information was easier to use, helped their performance and was more enjoyable than having no pose information. Future research directions point toward providing recommendations to robot operators for which pose they should change to given the terrain to be traversed.

Categories and Subject Descriptors

H.5.2 [**User Interfaces**]: Evaluation/methodology, graphical user interfaces, screen design.

General Terms

Design, Experimentation, Human Factors.

Keywords

Human-robot interaction, polymorphic robots, shape-shifting robots, situation awareness, interaction design, evaluation.

1. INTRODUCTION

Studies of robot use at the World Trade Center disaster found that a disaster site presents many challenges for robot navigation [Casper 2002, Casper and Murphy 2003, Micire 2002]. Imagine that you are responsible for guiding a robot through a partially destroyed building to look for victims. The floor of the building is buckled in places and is littered with chunks of ceiling material. Exploring the damaged building requires squeezing the robot under some debris and climbing over other obstacles, all while getting an overview picture of the building's rooms as quickly as possible to speed victims' rescue. While the building's condition is still unknown, you must stay outside and rely on the robot's telemetry/sensor data to provide information about the robot's state and the environment inside of the building.

It can be difficult for the same robot to do all of these tasks, because robots that are low to the ground cannot always climb well, and cannot usually see over obstacles to get a wide view of an area. Further, even in the best of circumstances it is difficult for people to interpret information about the robot's state and surrounding environment based solely on what is presented in the robot's interface; and the rescue situation described here is hardly an ideal environment.

These two problems formed the motivation for our study. The problem of conflicting task requirements can be addressed by using shape shifting, or polymorphic, robots, since these robots can adapt themselves in real time as needed to perform different types of tasks. But the use of polymorphic robots increases the challenge of maintaining situation awareness because operators must continuously adjust their understanding of the robots' state and interpretation of what the robot is reporting as their morphology changes. In fact, operators must adjust their expectations of the robots' capabilities, as well. For example, operators need to know if they should change the robot's shape (colloquially known as changing the robot's "pose") to avoid tipping over when climbing. Also, operators may waste time in extra navigation when getting an overview of a room's condition if they do not realize that their camera is not in a high enough position.

While operators need to account for the robot's pose, providing information on pose in the interface takes up valuable resources such as screen real estate. This paper reports on a controlled experiment that explored the effect on situation awareness of providing pose information versus no pose information. We anticipate that HRI designers can use our study results to make more informed tradeoffs when deciding what information to present to operators. Further, we provide suggestions for future HRI research that go beyond the question of whether pose information should be made visible.

HRI'06, March 2–4, 2006, Salt Lake City, Utah, USA.

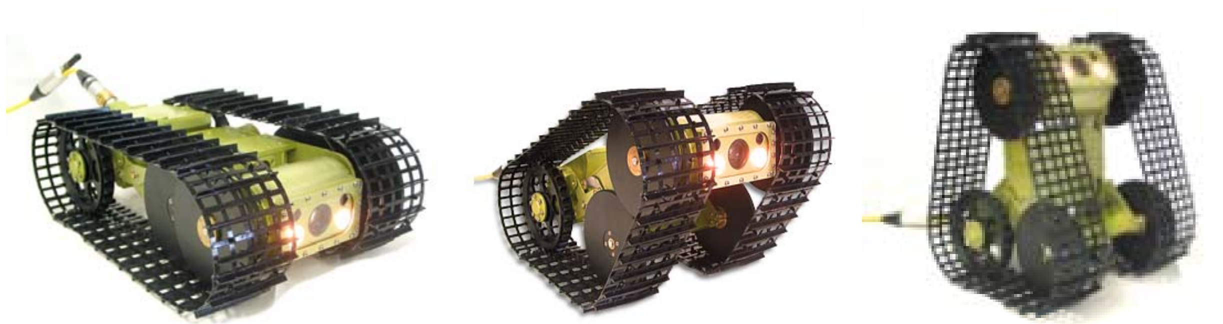

Figure 1: The VGTV-Extreme in its down position (left), partly raised position (middle) and fully raised position (right).

The next section provides additional background, followed by a description of the study methodology in section 3. Section 4 contains the hypotheses we investigated. Results and discussion may be found in section 5, and suggestions for future work complete the paper in section 6.

2. BACKGROUND AND RELATED WORK

The concepts of polymorphism and situation awareness (SA) are important to our work and merit further description.

2.1 Polymorphic Robots

We define *polymorphic robots* to be those robots that can change their shape, which changes their interaction with the world. Polymorphic robots used in remote situations require that the operator have awareness of the robot's current shape in addition to the robot's surrounding environment and other information about the robot's current state.

The iBot is a wheelchair that can drive on four wheels or rise up and balance on two of its wheels [Independence Technology 2005]. Moving to the two wheel configuration allows the robot to be taller, putting its user at eye-level with standing people. The robot can also climb or descend stairs by rotating its wheels end over end. The robot's rider is aware of the iBot's shape based upon the height of the seat. However, most polymorphic robots will not be ridden by their operator; instead, the shape information will need to be conveyed in some way.

Other examples of polymorphic robots are reconfigurable robots (e.g., [Yim et al. 2000], [Shen et al. 2002], and [Rus 2003]). Reconfigurable robots consist of a number of identical units that can be assembled into different shapes. For example, PolyBot [Yim et al. 2000] can be assembled into a four-legged walker, a fence climbing robot, a rolling robot and a two-legged tricycle rider. Reconfigurable robots can change their own shape (although quite slowly at this time). If this type of robot were to be used for human-robot interaction, it would be important to convey the current configuration (shape) of the robot to its user, at least at a high level. While the user will need to know if the modules are in a snake-like configuration, a round rolling configuration, or one of any number of other configurations, it may not be important to know the angle of each individual component with respect to its neighbors. The requirement for this type of information would depend upon the ability the robot had to carry out movement commands autonomously.

Another type of polymorphic robot is a snake robot (e.g. [Miller 2000] and [Wolf et al. 2003]). While motion control needs to occur at each of the joints, the actual shape of the robot may not be important to the robot's operator. If the robot only cares about moving the robot forward along a floor, the shape of the snake may not matter to the operator. However, if the snake must move amongst rubble or climb, the user will need to know the current configuration of the snake.

Humanoid (e.g. ASIMO [Honda 2005]) and dog-like (e.g. AIBO [Sony 2005]) robots are legged robots. As with the snake robots, it is not necessary to convey the shape of the jointed legs if the robot can be given "move forward" commands and the robot handles the control necessary to accomplish this task without falling over.

Treaded robots can also be polymorphic. The iRobot PackBot has two flippers that can be used to raise the body of the robot [iRobot 2005]. The VGTV-Extreme, used in these experiments, can raise from a flat configuration to a triangle (shown in figure 1) [Casper et al. 2004]. These two treaded robots are intended for remote operation in military and urban search and rescue environments. Due to the remote operation, the user will not be able to note the robot's shape by watching it.

2.2 Situation Awareness

The classic definition of situation awareness in the human factors literature is that of Endsley (1988): [Level 1] the perception of the elements in the environment within a volume of time and space, [Level 2] the comprehension of their meaning, and [Level 3] the projection of their status in the near future. In our earlier work [Drury et al. 2003], we determined that we needed a more detailed and tailored definition of SA to help analyze HRI. The simplest case in our definition of HRI awareness, for one human and one robot interacting together, is as follows:

> HRI awareness base case: Given one human and one robot working on a task together, HRI awareness is the understanding that the human has of the location, activities, status, and surroundings of the robot; and the knowledge that the robot has of the human's commands necessary to direct its activities and the constraints under which it must operate. [Drury et al. 2003]

Figure 2: The two testing courses, designed to have the same number and type of obstacles in different configurations.

The added challenges of shape shifting particularly affect the humans' awareness of the robots' status, location, and surroundings. Humans must know what pose the robot is in to be able to predict the robot's future behavior. In addition, they must be able to interpret the robot's sensor data based on the robot's current pose, as sensors can shift with the robot's changing shape.

There are a few studies that have explicitly examined robot operators' SA, although none of these studies has focused on polymorphic robots. Yanco and Drury [2004] found that search and rescue workers participating in their experiment spent, on average, approximately 30% of the time solely trying to gain or maintain SA, which chiefly consisted of understanding the robots' location, surroundings, and status (although this study did not break out the percentage of time spent on each of these three types of awareness). Burke et al. [2004] found that "operators spent significantly more time gathering information about the state of the robot and the state of the environment than they did navigating the robot" (p. 86). They reported that 24% of operators' communications with each other concerned the robots' state, 14% concerned the robots' location ("robot situatedness"), and 13% concerned the robots' surroundings (the "state of the environment"). Even without the challenge of polymorphism, it is clear that providing additional SA to robot operators is beneficial.

3. METHODOLOGY

Our approach was to have a set of operators navigate a polymorphic robot past surveillance checkpoints in a simulated disaster environment that required climbing over obstacles and driving through a tunnel. The course thus was most easily traversed when operators changed the robot's pose to match the terrain; the robot needed to be perfectly flat, for example, to drive through the tunnel. The experiment followed a within-subjects design, with the independent variable being interface design: pose information provided in the interface versus no pose information provided.

3.1 The Robot

Experiment participants used an American Standard Robotics VGTV-Extreme robot manufactured by Inuktun Services. It is a tracked and tethered, ground-based robot weighing approximately 14 pounds. The vehicle has six wheels, three on each side; its shape is changed by raising or lowering the front axle, which is also home to a "sensor pod" (normally containing a video camera).

3.2 Test Environment Description

Two test courses were designed to resemble office environments that had been damaged by a natural disaster such as an earthquake or a terrorist act. The NIST research on reference courses for USAR [Jacoff et al. 2000, Jacoff et al. 2001] guided us in the construction of our test courses. Our courses most closely resembled the "orange" or medium-difficulty courses because we employed rubble that the robot had to climb, ramps, and drop-offs. We were further guided in course development by the necessity of providing obstacles that the robot would need to climb over or through.

Since each participant was asked to perform two runs, we needed two courses to prevent knowledge of course layout from affecting performance when using the second interface. The courses each included five numbered tags and arrows guiding the operator in the correct direction to the next tag. We wished the participants to follow a prescribed path through each course because we wanted to force them to maneuver each of the obstacles in turn. The courses were carefully designed to be of the same difficulty and length and to require the same number of pose transitions. The courses can be seen in figure 2. The starting course was alternated between subjects.

Besides the two courses, we developed a practice area. The practice area consisted of some open space plus a board suspended between two cinder blocks approximately three inches from the ground. The purpose of this practice area was to enable the users to become familiar with the basic operations of moving the robot forward and backwards, re-aiming the camera, and changing poses to climb over obstacles.

3.3 Experiment Participants

Nineteen people participated in the experiment: 11 men and eight women. Participants' ages ranged from 25 to 55. All were computer-savvy employees of a high-technology company, required to work with computers 20 – 40 hours/week on their jobs. Nearly half, 47%, spend over 10 hours/week using a computer at home for personal use. 37% had previously used robots at least once. Of the people with

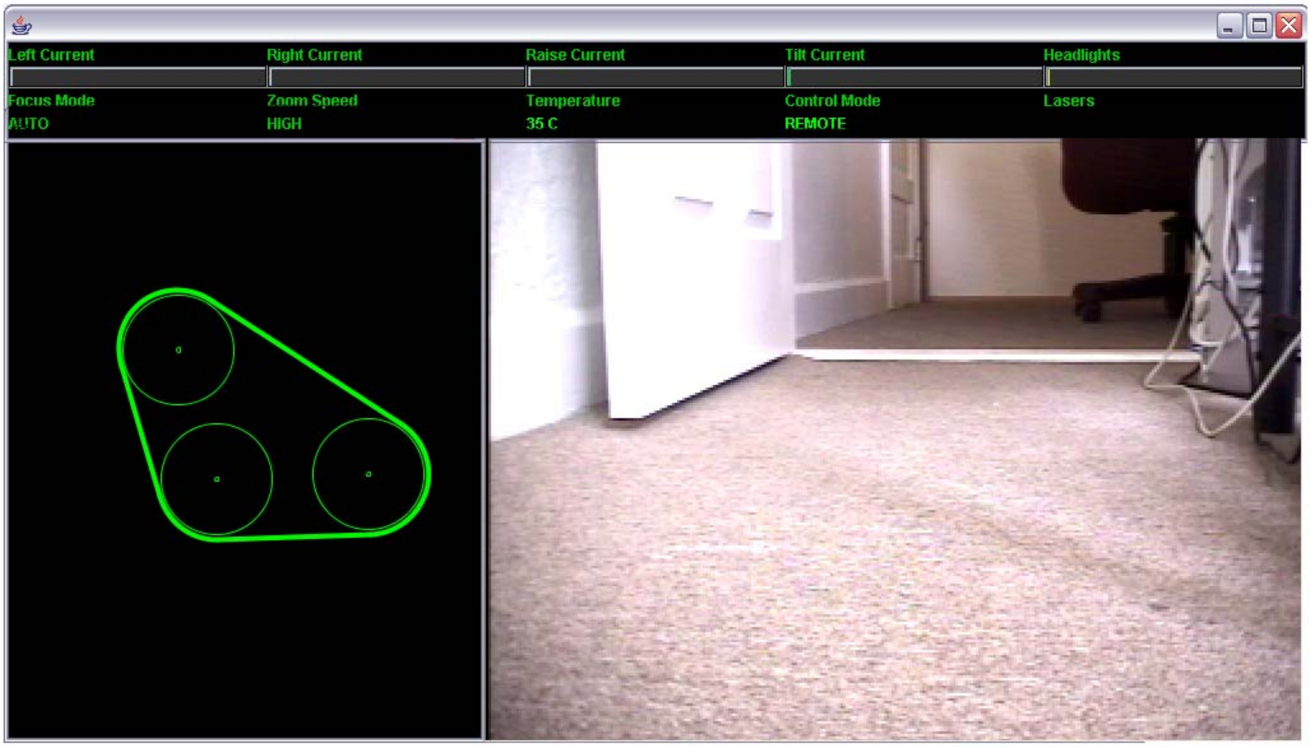

Focus Mode: AUTO
Zoom Speed: HIGH
Temperature: 35 C
Control Mode: REMOTE

Figure 3: The interface used in the experiments, showing the pose display. When testing with no pose information, the pose display was turned off and this area was empty.

previous exposure to robots, three (16%) had used industrial robots and most of the rest had used "toy" robots such as RoboSapien or household robots such as Roomba (a robotic vacuum cleaner). Over half of the participants (68%) had used remote control vehicles previously, although most rarely use them. Only one of the participants (5%) had any prior experience with search and rescue tasks; this experience consisted of some training 20 years ago with no active experience since then.

3.4 Experiment Design and Conduct

We employed two counterbalancing techniques to maintain the integrity of our within-subjects experimental design. Half of the participants began with pose information and the other half with no pose information. Within each of these two groups, half began the formal experimental runs using Course 1 and the other half began using Course 2. Thus, there were four combinations of the first interface and course. Each experiment run was conducted with a single participant. The interfaces viewed by the participants differed only in whether or not they provided pose information. The interface lacking pose information had a blank spot where the pose information was located in the other interface. (The interface is shown in figure 3.)

After signing Informed Consent forms, participants filled out a pre-experiment questionnaire to enable us to understand their previous experience with search and rescue, computers, and robots. Participants then received training on how to control the robot, including giving them a "pose reference sheet" to help them remember the various potential robot configurations and tips for which pose to use for which situation. Participants were taught how to climb over obstacles; the steps are shown in figure 4. The same researcher trained all 19 subjects using a written training document to maintain consistency. Next, we allowed participants to practice teleoperating the robot, during which time the participants often asked questions about the robot's operations that were answered by the researcher. The participants were trained using the same interface that they would subsequently use for the first run; the participants were not told in advance that they would be operating the robot with and without pose information.

Once training was completed, we moved the robot to an area of the laboratory that was screened from view. We placed the robot at the beginning of the appropriate course and asked participants to follow the arrows and numbered tags in sequential order. We told the participants that the tags represented surveillance checkpoints in a search sweep of a damaged area. The participants took between 4 and 14 minutes to complete the first run. After the first run, we moved the robot to the beginning of the other course and a researcher swapped out the interface. The total changeover took approximately one minute. The participants then performed the same task under the new conditions. Finally, the participants filled out a one-page post-test questionnaire and a researcher debriefed them, asking the participants additional questions relevant to their performance. The entire process took approximately one hour per participant.

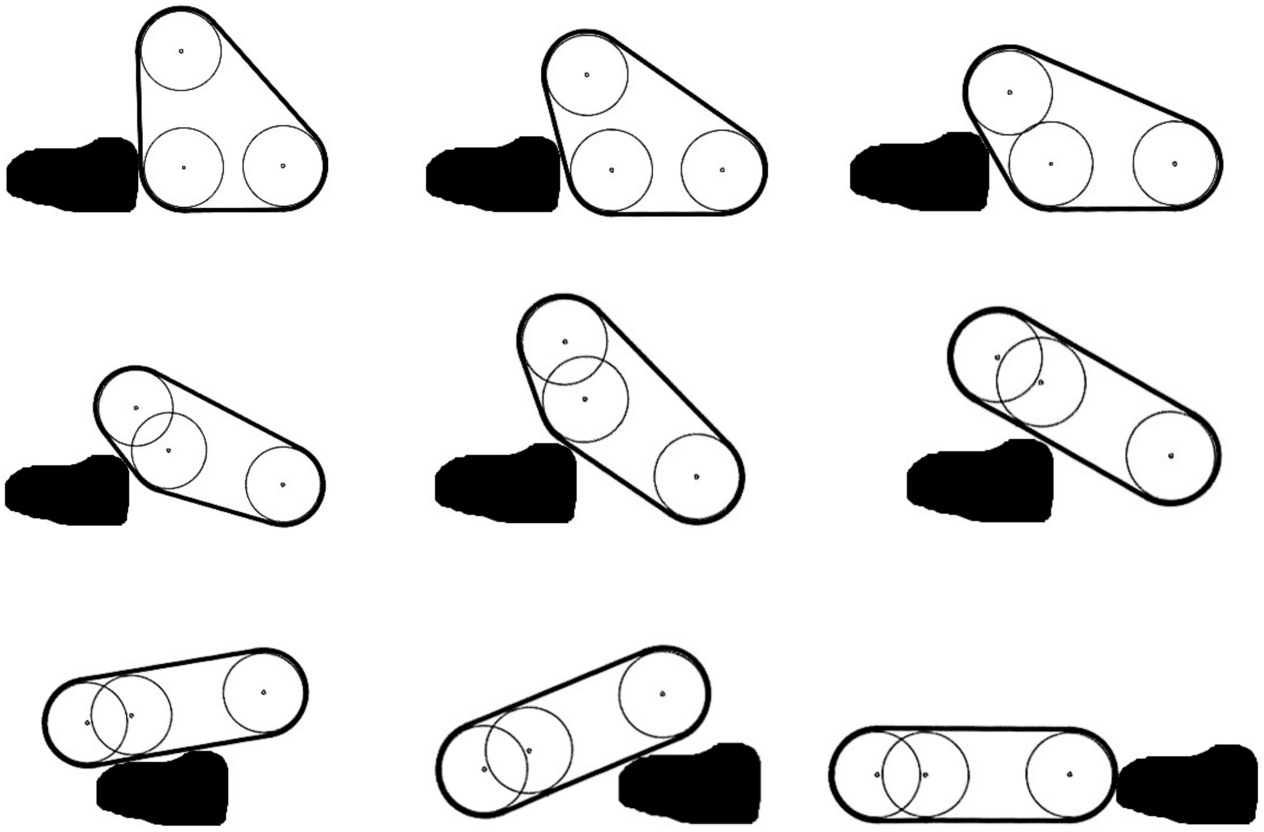

Figure 4: Climbing instructions given to participants, with the drawings above as illustrations: "When approaching an object, maintain the pose that gives the maximum vertical height of the leading track edge. To overcome the object, the robot should come to rest against it, and then be lowered so the front wheel rests against the object. Lower the robot's position until the robot begins to rise. As you continue to lower the robot to the flattened position, begin driving forward. By flattening the robot, the center of gravity will shift forward making the robot tilt forward and fall over the object. You should now be in the position to drive forward over the object. Once the object is cleared, continue on your mission choosing the appropriate pose position."

3.5 Data Collection and Measures

We collected four types of data: questionnaires, video, observer notes, and automated logs. Post-experiment questionnaires asked several Likert scale and fill-in-the-blank questions regarding the helpfulness and utility of the pose display and also solicited suggestions for interface improvements. We filmed video in two locations: pointing towards the interface from over the participants' shoulder, and in the robot area (a researcher followed the robot around).

One observer made notes on the participant during the formal runs, especially when the participant expressed confusion, frustration, or was making obvious errors. While we did not explicitly ask the participants to "think aloud" [Ericsson and Simon 1980] because we did not wish to invalidate timing data, many of the participants made comments to the observer while teleoperating the robot. (This observer was the same person who trained all the participants.)

We used the following codes for critical incidents during the runs:

- "T" when the robot tipped over or would have tipped over if one of the researchers had not pushed the robot back.
- "B" for bumps when the robot hit an object with substantial force (not simply grazing an object or interacting with the test course in a controlled fashion).
- "TU" for any incident that involved a tunnel, such as backing into it, dragging it, or bumping the top of it if the participant raised the robot's pose while still in the tunnel.
- "CL" for climbing or riding up on an object that was not intended to be climbed, such as riding up on screening that was leaning against a computer table. We noted which tread rode up on the object and how high it rode up.
- "RC" for when the robot was caught on an object such that its forward motion was hindered.

We computed Cohen's Kappa statistic to determine intercoder reliability. After chance was excluded Kappa agreement was computed to be .75 (when chance was not factored out Kappa was .81). Landis and Koch [1977] indicate that a Kappa value between .61 and .80 is "substantial agreement" (p.159), and Fleiss [1981] suggests that above .75 suggests "strong agreement" above chance.

The interfaces were well-instrumented and we had access to extensive automated log files. The log files captured every change of the joystick position and button push to raise or lower the robot's pose as well as each action taken by participants to change the camera angle.

4. HYPOTHESES

We designed the experiment to investigate six hypotheses. These hypotheses probed the components of HRI awareness that pertain to the robot operator's ability to understand the robot's status (chiefly, what pose it is in) and surroundings (for example, how high obstacles are that must be climbed).

H1: Participants will be better able to perform other tasks while changing pose when pose information is displayed.

H1 discussion: We conjecture that if a participant can see pose information, he or she may be able to focus more on other tasks while changing pose. Our rationale is that the pose information will be presented visually so the participant does not have to hold this piece of information in his or her head, freeing cognition for other uses. Norman [1988] has identified the concepts of recognizing information versus having to recall it as "knowledge in the world" versus "knowledge in the head."

H2: Participants will change the robot's pose more frequently when pose information is not displayed.

H2 discussion: We believe there could be fewer pose changes if pose is displayed in the interface. The reason for this belief is that if participants do not know what their current pose is, they will often need to bring the robot all the way to an upright position or all the way down to the flat position (that is, they change the robot to be configured in one of the "end points") to orient themselves. If participants can see the pose information in the interface, they will not have to make pose changes simply to understand their current pose.

H3: The number of times the participants tipped the robot into an unstable position will be greater when pose information is not displayed.

H3 discussion: Participants may not realize that they have left the robot in a pose that is unstable for teleoperating subsequent portions of the test course if they do not have a reminder of the current pose information in the interface.

H4: Participants will cause the robot to bump into obstacles more often when pose information is not displayed.

H4 discussion: If pose information is not presented in the interface, and if participants do not remember a robot's pose, they may bump obstacles inadvertently.

H5: Participants will cause the robot to become caught in obstacles or other parts of the environment more often when pose information is not displayed.

H5 discussion: If pose information is not presented in the interface, and if participants do not remember a robot's pose, they may cause the robot to become entangled in the environment inadvertently.

H6: Participants will cause the robot to unintentionally climb up on parts of the environment more often when pose information is not displayed.

H6 discussion: If pose information is not presented in the interface, and if participants do not remember a robot's pose, they may cause the robot to ride up on objects in the environment inadvertently.

5. RESULTS AND DISCUSSION

We hypothesized that displaying pose information would result in fewer tips (H3). Our experiments confirmed this hypothesis. Using a two-tailed paired t-test, df=18, we found a significant difference in the number of tips in the no display case vs. the pose display ($p<.03$). Eleven of the nineteen participants tipped with no pose information, while five of the nineteen tipped with pose information. The average number of tips by people who tipped without pose information is 1.74 tips (standard deviation 1.01); the average number of tips by those who tipped with pose information is 1.0 tips (standard deviation 0). Overall, the average number of tips with no pose information is 1.0 (standard deviation 1.14) and the average number of tips with pose information was 0.26 (standard deviation 0.45).

Without pose information, participants were less likely to be aware that the robot was in an improper state for climbing, which resulted in the additional tips. Additionally, the pose display showed the robot's inclination as well as its shape, which subjects could use to see that the robot was about to tip over. Measuring by the number of tips, awareness of pose was improved through the use of the display.

Another navigation problem strongly related to the pose of the robot is getting the robot caught under or on the sides of objects. For example, rather than tipping when climbing an object in an incorrect pose, it could fall over the side of the object, getting caught on that edge. We had hypothesized that there would be more incidents of the robot getting caught when pose information was not displayed (H5), which was confirmed by our experiment. There was a significant difference ($p<.03$) between the number of robot caught incidents for the two display cases (no pose: 0.37 (0.50); pose: 0.05 (0.23)). As with the number of tips, we found that fewer operators had problems getting caught when the pose information was displayed (only one of the nineteen participants) than with no pose information displayed (seven of the nineteen participants).

While SA with regard to pose was increased by displaying the pose window on the interface, SA either unrelated or only somewhat related to pose did not change. There was no significant difference in the number of bumps (no pose: 0.47 (0.61); pose: 0.58 (0.69)), tunnel incidents (no pose: 1.0 (1.0); pose: 0.89 (1.15)) and climbs (no pose: 0.26 (0.45); pose: 0.42 (0.51)). We also did not see a difference in the number of people experiencing these problems: bumps (8 of 19 without pose, 9 of 19 with pose), tunnel incidents (12 of 19 without pose, 10 of 19 with pose), and climbs (5 of 19 without pose, 8 of 19 with pose). We had hypothesized that there would be more bumps (H4) and climbs (H6) without pose information displayed because changing the robot's pose also changes its footprint and the height of the operator's view of the surroundings; we thought that misunderstanding the robot's pose may result in thinking that the robot's footprint was smaller than it actually was and thus the operator could run into (bump or ride up on) obstacles inadvertently. When we found no additional bumps or climbs, we conjecture that the participants tried to provide a wide margin between the robots and objects in the environment because they did not wish to damage anything. While they may have maneuvered the robot closer to obstacles than they thought, they had similar numbers of bumping and climbing incidents in the two conditions. It is also true that no additional awareness of surroundings was given through the interface when pose information was displayed; the robot still gave a narrow view of the environment.

We had conjectured that participants would change their pose more often when the pose information was not displayed (H2). We did not find a significant difference in the number of pose changes for the two display cases. However, we did see subjects exhibit a strategy for determining their pose when no information was displayed: subject would bring the robot all the way up to pose 5 (triangular), then would bring the robot back down to put it into pose 3 or 4. This observation is borne out by the data, which shows that there was a significant difference in the amount of time participants spent in pose 5 without pose information (p<.003).

There was no significant difference in the time required to traverse the courses under the two interface configurations (no pose: 5:08 (2:52); pose: 4:17 (1:18)). When the robot tipped, the person managing the robot's tether would put the robot back into an upright position. Had we required that the operator right the robot, we expect that we would have seen a significant difference between the two interfaces for run time, as there were significantly more tips when no pose information was displayed. However, it can be very difficult to right the robot in some situations and the additional training required for these contingencies was deemed to be overwhelming for a user just learning to use the robot.

Despite being instructed to look for all of the numbers as they were driving the robot through the course, seventeen of the nineteen participants said that they were not always aware of the number for which they were looking, indicating a high workload for the overall task. This observation of workload held true for both display cases. We believe that the workload may be one reason why we did not find evidence to support our hypothesis that operators would be able to do other tasks such as driving the robot forward while changing pose (H1). To explore this hypothesis, we measured the number of "complex moves" made by an operator. Complex moves were defined as a period of movement in which the robot pose was changed and the movement was also changed at the same time. We did not find a significant difference in the number of complex moves between the two cases.

In general, we found that participants underestimated the amount of damage they did to the arena for both display conditions, indicating that they did not have good SA throughout their runs. We looked at the participants' reports of poor SA, measured by the number of times they stated that they were confused about the robot. We found that the participants made fewer statements to this effect when the pose information was displayed (p<0.025; no pose: 0.78 (1.27); pose: 0.11 (0.30)). As discussed before, SA is improved with regard to pose information (i.e., status), but it is not with regard to understanding the current environment (i.e., surroundings). The participants did not clarify their SA comments by stating what they were confused about, but given our results, we believe that having no pose information created much greater confusion in the participants.

All participants reported that the pose display allowed them to know what pose the robot was in (all selected the highest category, 1) while without the pose display the average ranking was between "sometimes" (category 2) and "rarely" (category 3) with p<.001.

Table 1 contains a summary of the results and also states the types of SA that the hypotheses relate to most strongly, using the components of the HRI awareness definition from Drury et al. [2003]. Note that we found significant results in the case of each of the hypotheses that pertained to awareness of the robot's status (pose). These results make sense due to the fact that the only additional information that we provided to participants consisted of pose information.

Table 1. Results of Testing Hypotheses (H1 – H6) Regarding SA of Polymorphic Robots

H #	Hypothesis Summary	Related SA Component	Results
H1	Multi-tasking frequency	Surroundings, Status	Not proven. No significant difference in the number of complex moves completed.
H2	Pose change frequency	Status	No significant difference in the number of pose changes but a significant difference in the amount of time participants spent in the highest pose position.
H3	Tipping frequency	Status	Proven. Significant difference in number of tips.
H4	Bumping frequency	Surroundings	Not proven. No significant difference in the number of bumps.
H5	Entanglement frequency	Status	Proven. Significant difference in number of times robot was caught on objects in the environment.
H6	Climbing frequency	Surroundings	Not proven. No significant difference in number of unintentional climbs.

6. CONCLUSIONS AND FUTURE WORK

We found that adding the pose display did increase awareness of the robot's status. Improved pose awareness was indicated by a significant decrease in the number of tips and robot caught incidents when the pose information was displayed. However, SA measures unrelated to pose (bumping and climbing up on obstacles such as walls) were not improved by adding the pose display.

Workload was very high for the task, regardless of whether the pose was displayed or not. The subjects had to understand how to make the robot climb over the obstacles set out on the course. We believe that the operator's workload could be reduced by assisting with the shape changing task. For example, when the robot's treads came into contact with an obstacle that could be climbed (measured by sensing the height of the obstacle with a distance sensor or with the video image), the robot could raise itself into the correct pose for climbing onto the obstacle. Once on the obstacle, the robot could lower its pose. Automatically changing pose would allow the robot's operator to issue only movement commands, unless it was desirable to raise the robot up to get a better view of the environment. We believe that this type of automated assistance shows promise for reducing the operator's workload and allowing for improved task performance.

7. ACKNOWLEDGMENTS

This work was supported in part by NSF IIS-0415224 and NIST 70NANB3H1116. Thanks for Mark Micire and Mike Baker for assistance with this paper.

8. REFERENCES

[1] Burke, J.L., Murphy, R.R., Coovert, M. D., and Riddle, D.L. (2004). "Moonlight in Miami: A Field Study of Human-Robot Interaction in the Context of an Urban Search and Rescue Disaster Training Exercise." Human-Computer Interaction. 19(1-2), pp. 85 – 116.

[2] Casper, J. (2002). "Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center." MS Thesis, University of South Florida Department of Computer Science and Engineering.

[3] Casper, J.L., Micire, M.J. and Li Gang, R. (2004). "Inuktun Services Ltd. – Search and Rescue Robotics." *Proc. 3rd International Conference on Continental Earthquakes (ICCE)*, Beijing, China.

[4] Casper, J., and Murphy, R.R. (2003). "Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center." *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, June, vol. 33, pp. 367-385.

[5] Drury, J. L., Scholtz, J., and Yanco, H. A. (2003). "Awareness in human-robot interactions." *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, Washington, DC, October.

[6] Endsley, M. R. (1988). "Design and evaluation for situation awareness enhancement." *Proceedings of the Human Factors Society 32nd Annual Meetin*g, Santa Monica, CA, Human Factors Society.

[7] Fleiss, J. L. (1981). *Statistical Methods for Rates and Proportions*. New York: John Wiley & Sons. (Second Edition).

[8] Gardner, W. (1995). "On the reliability of sequential data measurement, meaning, and correction." In John M. Gottman (Ed.), *The Analysis of Change*. Mahwah, N.J.: Erlbaum.

[9] Honda (2005). http://world.honda.com/ASIMO/, accessed September 2005.

[10] Independence Technology (2005). http://www.independencenow.com/home.html, accessed September 2005.

[11] iRobot (2005). http://www.irobot.com/governmentindustrial/, accessed September 2005.

[12] Landis, J. & Koch, G. G. (1977). "The measurement of observer agreement for categorical data." *Biometrics*, (33), 159-174.

[13] Micire, M. (2002). "Analysis of the robotic-assisted search and rescue response to the World Trade Center disaster." Masters Thesis, University of South Florida, July. Also available as CRASAR-TR2002-6.

[14] Miller, G. (2000). "Snake Robots for Search and Rescue." *Proceedings of the Conference on Neurotechnology for Biomimetic Robots*, May.

[15] Norman, D. A. (1988). *The Psychology of Everyday Things*. New York: Basic Books.

[16] Rus, D. (2003). "Self-reconfiguring robots: successes and challenges." *Proceedigns of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico, August.

[17] Shen, W.-M., Salemi, B. and Will, P. (2002). "Hormone-inspired adaptive communication and distributed control for CONRO self-reconfigurable robots." *IEEE Transactions on Robotics and Automation*, 18(5), October.

[18] Sony (2005). http://www.sony.net/Products/aibo/, accessed September 2005.

[19] Wolf, A., Brown Jr., H.B., Casciola, R., Costa, A., Schwerin, M., Shammas E. and Choset, H. (2003). "A mobile hyper redundant mechanism for search and rescue tasks." *Proceedings of the 2003 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems*, Las Vegas, Nevada, USA, October, volume 3, pages 2889-2895.

[20] Yim, M., Duff, D.G. and Roufas, K.D. (2000). "PolyBot: a modular reconfigurable robot." *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, April, volume 1, pp. 514-520.

[21] Yanco, H. A. and Drury, J. L. (2004). "'Where am I?' Acquiring situation awareness using a remote robot platform." *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, The Hague, Netherlands, October.