"Where Am I?" Acquiring Situation Awareness Using a Remote Robot Platform *

Holly A. Yanco
University of Massachusetts Lowell
Computer Science Dept., One University Ave.
Lowell, MA 01854 USA
holly@cs.uml.edu

Jill Drury
The MITRE Corporation
202 Burlington Road
Bedford, MA 01730 USA
jldrury@mitre.org

Abstract - *Human-robot interaction with urban search and rescue (USAR) robots needs to provide operators with a means of maintaining situation awareness (SA), especially since the USAR operators usually cannot see the robots that they are directing. We used a technique from human-computer interaction known as usability testing, plus implicit and explicit SA measurement techniques, to investigate USAR operators' levels of SA and strategies for maintaining SA. We found that operators developed different SA strategies, spent an average of 30% of their time solely in SA activities, had less SA of the space behind the robot than in front or on the sides, did not use automatically-generated maps to gain SA., and had difficulty maintaining SA when in the autonomous mode.*

Keywords: Situation awareness, human-robot interaction, HRI, urban search and rescue, human-computer interaction, HCI.

1 Introduction

We have been working to understand how human-robot interaction (HRI) can better support users in safety-critical situations: situations in which an error or failure could result in death, injury, loss of property, or environmental harm [Leveson 1986]. Safety-critical situations constitute a serious challenge for robot designers due to the vital importance that robots perform exactly as intended and support humans in efficient and error-free operations. We have been studying HRI in the urban search and rescue (USAR) domain as a prime example of a safety-critical application.

Prior to this year, we analyzed HRI chiefly during USAR robotic competitions. At competitions, the robots were most often controlled by the people who developed them. These developers were not typical users: they often did not have the same backgrounds or degree of experience with computerized systems (meaning, we expect that the developers were more used to working with technology than USAR domain experts on average, and less experienced in performing authentic USAR tasks). According to the principles of human-computer interaction (HCI), it is not possible to get a true assessment of the suitability and usability of an interface without testing it with representative users. Thus, we designed an investigation, described below, to determine whether two different interfaces for USAR robots adequately support USAR domain experts known as first responders.

We had previously noted [Drury, Yanco, and Scholtz 2003] that most problems encountered when navigating robots have resulted from the humans' lack of awareness of the robot's location, surroundings or status. While we called it "human-robot awareness," first responders speak of "maintaining situation awareness." Situation awareness (SA) is defined by Endsley [1988] as "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future." We can modify this definition for HRI, giving a definition of situation awareness as *the perception of the robots' location, surroundings, and status; the comprehension of their meaning; and the projection of how the robot will behave in the near future*.

Since it was clear to us that HRI needs to support users in attaining and maintaining SA, we designed an investigation specifically to probe SA acquisition and maintenance as supported by the robots' interfaces.

The rest of this paper contains a brief overview of the SA literature, our investigation methodology, results, and conclusions.

---

2 Situation Awareness

Because the literature surrounding SA is very large, we confine ourselves to SA measurement, plus literature specific to SA of robotic systems.

2.1 SA Measurement

Hjelmfelt and Pokrant [1998] state that experimental methods for measuring SA fall into three categories:

Subjective: subjects rate their own SA

Implicit performance: experimenters measure task performance, assuming that a subject's performance correlates with SA, and that improved SA will lead to improved performance

Explicit performance: experimenters directly probe the subjects' SA by asking questions during short suspensions of the task.

Two of the most frequently used measurement methods fall into different categories: Taylor's [1990] Situational Awareness Assessment Technique (SART) is a subjective measure, while Endsley's [1988] Situation Awareness Global Assessment Technique (SAGAT) falls under the explicit performance category. These and other examples of often-cited measurement methods are described in Table 1.

Table 1. SA Measurement Techniques

Category	Technique	Reference	Description
Subjective	Situational Awareness Assessment Technique (SART)	[Taylor 1990]	SART was developed through analysis with pilots. Operators rate on a 10-dimensional bipolar scale the degree to which they perceive (1) a demand on operator attention, (2) supply of attention and (3) understanding of the situation. The 10 components of the scale are combined to provide an overall score.
Subjective	Crew Awareness Rating Scale (CARS)	[McGuinness 1999]	A generic 8-part questionnaire addressing both the mental content and mental processing of SA with respect to 4 separate functions: • Perception – assimilation of new information • Comprehension – understanding of information in context • Projection – anticipation of possible future developments • Integration – synthesis of above with one's courses of action For each, subjects were asked to rate (from 1 – 4 with 1 being best and 4 being worst): • Content of that aspect – is it reliable and accurate? • Processing of that aspect – is it easy to maintain?
Implicit performance	"Mini sitreps"	[McGuinness 2002]	Short situation reports (mini sitreps) are used to provide an objective measure of the match between the subject's understanding of the situation and the actual situation at that point in time. The mini sitreps in McGuinness 2002 contained: • Current enemy positions • Assessment of enemy intent • Assessment of current operation • Assessment of future developments/outcomes • Any deviations from original plans
Explicit performance	Situation Awareness Global Assessment Technique (SAGAT)	[Endsley 1988]	Employs periodic, randomly-timed freezes in a simulation scenario during which all of the operator's displays are temporarily blanked. A series of questions is asked of the operator with the intent of probing their understanding of the situation.
Explicit performance	Real-time probes	[Endsley et al. 1998]	Uses SAGAT-type questions, but doesn't involve blanking the screen. Measures response time as well as accuracy.

2.2 SA of Robotic Systems

A number of studies have stated the importance of designing the human-robot interaction so that the operator can maintain SA. Some recent examples are Dudenhoeffer et al. [2001], Johnson et al. [2002], Green and Oh [2003], Lewis et al. [2003], and Hughes and Lewis [2004], as summarized below.

- Dudenhoeffer et al. [2001] hypothesize that modeling and simulation will help to develop interfaces that provide SA.
- Johnson et al. [2002] evaluate how well operators can perform tasks with an original vs. improved interface.
- Green and Oh [2003] discuss using unmanned aerial vehicles (UAVs; airborne robots) in enclosed spaces.
- Lewis et al. [2003] compare users' performance with a gravity-referenced display versus a standard fixed camera with separate attitude indicator.
- Hughes and Lewis [2004] investigate how multiple cameras with independent control can potentially improve HRI performance (where performance is defined by numbers of targets located and identified).

All of these studies used implicit performance measures for SA to a greater or lesser extent (e.g., subjects in the experiment described by Hughes and Lewis [2004] reported targets found under several camera configuration conditions), although none of the papers made a strong connection between their dependent variables and the implications for SA. In particular, the dependent variables chosen by Johnson et al. [2002] are very pertinent examples of implicit measures of operators' SA levels. In addition, Scholtz [2003] discusses roles that humans take on when working with robots and how their SA can be measured explicitly or implicitly.

3 Methodology

We performed a "usability test" as practiced by HCI experts. Usability tests involve observing typical users (often only 3 to 5 users) performing representative tasks under realistic conditions, usually while "thinking aloud" (voicing their thoughts) [Ericsson and Simon 1980]. If all users perform the same tasks and each has difficulty at certain points, then the interface elements that are in use during those points are likely candidates for redesign.

We tested four first responders who had no prior experience working with robots (although some had experience with remote controlled cars or airplanes). For each user, we conducted a half-day of testing on two different experimental robotic systems, which included training for each system as well as practice and testing runs. In all cases, the human directing the robot (the "operator") could not see the robot while it was in operation. This is a very important distinction: it is much easier to direct a robot that a human can see rather than to perform so-called "remote" robot operations. The operators were asked to find victims (represented by numbered tags) in the NIST USAR test arena. This arena replicates a partially destroyed building [Jacoff et al. 2000]. The operators were not allowed to see the arena prior to their runs, further increasing task complexity.

For each run, we videotaped the robot's progress in the arena and recorded which parts of the arena it covered, captured the operator's manipulation of the interface and also the operator's voice, and conducted post-run interviews with the operator. We measured SA implicitly using such measures as amount of time spent panning the camera and the number of times the robot bumped elements in the environment. Further, we analyzed the tapes of the robot and the operator's screen/voice to determine the operator's SA acquisition strategies and explicit self-assessment of SA (e.g., "I have no idea where the robot is right now.")

While the subjects tested two robot systems, the results discussed below are for only one of the two systems. The second system had many hardware failures during the runs, preventing us from obtaining usable data.

4 Results and Discussion

Our quantitative results are summarized in Table 2. We found that individual users have different strategies for acquiring SA. We had hypothesized that the strategies would mostly be influenced by the design of the robot system, since the design of the system directly influences the interactions that a user has with it. However, in our tests, we found that four subjects had different strategies for acquiring SA on the same robot system.

In our tests, we found that 12 – 63% of each run was spent acquiring SA to the exclusion of all other activities.[1] An average of 30% of each run was spent acquiring SA while no other task was being done.[2] Despite this time spent trying to acquire SA, users often expressed confusion about where their robots were located relative to various landmarks and whether their robots were near obstacles. For example, users backed robots into walls, asked "have I been here before?", and stated, "I have no idea where I am."

---

[1] In Subject 4's Run 3, no time was spent by the operator to acquire SA because the robot was in autonomous mode. This number is not included in our range.

[2] This figure includes Subject 4's Run 3 that used no time to acquire SA while not moving the robot. Removing this zero results in an average of 33% of run time spent only acquiring SA.

Table 2. Data from the experiments

Subject, Run	% Run acquiring SA	% SA camera panning	% SA robot panning	% Run with camera off-center	Number of obstacle encounters
1, 1	63	81	19	20	0
1, 2	56	43	57	1	7
1, 3	43	54	46	0	0
2, 1	28	26	74	9	2
2, 2	12	73	27	32	4
2, 3	28	23	77	0	7
3, 1	25	70	30	48	0
3, 3[1]	24	13	87	2	1
4, 1	13	64	36	0	2
4, 2	36	66	34	0	1
4, 3	0[2]	0	0	0	5
Average	29.8	46.6	44.3	10.2	2.6

We observed two primary methods for acquiring SA on the robot platform that was tested. The first involved moving the camera; operators could pan, tilt and zoom the camera to look around the environment. The second involved moving the robot back and forth to look around, not for any navigation purpose. Most of the subjects developed this strategy when they found that the robot could be turned more quickly than the camera could be turned. We found that two of the subjects spent more time panning the camera than the robot (Subjects 1 and 4), while the other two (Subject 2 and 3) spent more time panning the robot than the camera.

One problem observed with camera panning is that the camera was sometimes left off-center when the operator resumed driving, which can negatively impact an operator's SA. In a prior study [Yanco et al. 2004], we found that an operator drove with his camera off center for over half of his run, causing him to hit more obstacles than usual during that run. The camera was pointed to the left, so the operator would see a clear area, but there would be an obstacle in front of the robot.

4.1 SA and Mapping

Another method for acquiring SA would be to use the map generated by the robot system and displayed on the system's interface. However, the subjects noted that the map was not very useful. One subject stated, "I have to keep a mental image of the map because this [pointing to the map on the interface] right now is useless." Another subject said, "The mapping down here, for me personally, isn't really helping a whole lot. It's hard to tell where the robot is with respect to this [pointing to the map]."

We believe that this problem occurred due to the fact that the depiction of the robot on the map was a small gray dot, while walls were also marked using gray pixels. If the robot was moving, you could see the gray dot representing the robot move around the screen. However, since the map is presented diagonally to the right, below the video screen, the operator could not watch the map updating while concentrating on the video screen to drive the robot.

The mapping system could be improved by using a clearer icon showing where the robot is on the map. Additionally, allowing the users to mark landmarks on the map could help. Subjects were often surprised to find that they had come back to a known location such as a previously identified tag or the entrance to the arena. If landmarks such as these were marked on the map, the operator could see the robot's progress through the space more clearly.

4.2 SA During Autonomous Driving

Subject 4 chose to put the robot in an autonomous mode during the majority of the third run. In the autonomous mode, the robot drives around the arena,

[1] Data from Subject 3's Run 2 was not analyzed due to problems with the robot system during the run.

[2] During Subject 4's Run 3, the autonomous mode was primarily used. Since the robot did most of the driving, the subject did not pan the camera or robot to acquire SA.

avoiding obstacles (or, rather, avoiding as many obstacles as the autonomy algorithm can manage). During the run, the subject did not need to acquire SA in order to be able to navigate the robot. The subject noted, "This [autonomous driving mode] is really helpful because now you can really kind of use the mapping since I don't have to worry about where I'm going."

During this run, the subject was able to pay more attention to the interface, noting details such as the health status of the robot. The subject remarked, "There's a lot of information on the screen. It'd take a lot of practice to take it all in and use it."

In the middle of the run, the map created by the robot grew too large to be properly displayed on the interface, leading the subject to ask if the robot had left the arena since it had gone off the edge of the map.

After approximately eight minutes of driving autonomously, the subject noted, "Looks like it keeps going over the same area, so I'm going to have to… after it backs itself out of here, I'm going to intervene maybe and get it out. I don't think I'm going to have much luck." The subject then switched back to a driving mode where a joystick was used to control the robot's movement. The switch was made because the robot had moved into a dark area and didn't seem to be making any progress towards emerging from the area. The subject described being confused about where the robot was and how to get out of the area.

While an autonomous robot can explore areas that a human operator may not be able to go (in this case, the robot ventured into a covered area that the operator had not gone into in either of the prior two runs), the operator may lose SA during the period of autonomy. Since the operator is not driving, it does not seem necessary to keep track of every turn the robot is making. However, the operator needed to take control back from the robot when it had problems, putting the operator into a situation where SA had not been acquired and the video image was dark.

4.3 Directional SA

During the experiments, we observed the robot bump obstacles in the environment an average of 2.6 times per run. Of the 29 hits during all of the subjects' runs, 12 or 41% of the hits were on the rear of the robot. If SA were the same everywhere around the robot, we'd expect to see an even distribution of hits.

It is easiest to obtain awareness of the robot's surroundings around the front of the robot. The front of the robot has a color video camera with pan, tilt and zoom capabilities, a laser rangefinder, and several sonar sensors. This observation appears to be supported by the fact that just 4, or 14%, of bumps happened on the front of the robot.

It is a bit harder to gain awareness on the side of the robot, with only two sonar sensors pointing out from each side. However, the primary direction of motion is usually forward or backward, leading to less opportunity for hitting the side of the robot. Additionally, the obstacles to the sides of the robot just passed from view as the robot moved forward.

The rear of the robot only has two sonar sensors, one on the left and one on the right. We believe that the gap in sensing is causing many of the rear hits. In fact, while in safe mode (a mode that is supposed to prevent the robot from hitting objects in the environment), we observed the robot back up into obstacles several times. This SA problem could be remedied through the addition of more sonar sensors on the rear of the robot. Adding a rear view camera might also help to correct the problem of lack of SA on the rear of the robot.

4.4 Losing and Regaining SA

All subjects expressed confusion as to where they were during some portion of their runs. In one run, a subject said, "I have no idea of where I am." Three minutes later, he added, "I'm all disoriented now." In another run, another subject asked, "Where am I at?" One subject started her second run by stating, "I'm going to look where I start so I know where I'm going this time."

We noted that SA is most often regained when a subject finds an object that has already been seen before, whether it's a numbered tag or a particular location. One subject noted, "I've lost my bearings of where I am." Three minutes later, the subject said, "I'm back to where I came in." While the subjects can reacquire SA in this manner, they don't know how they got back to that point. In a rescue situation, an inability to describe how to get to a victim would be a significant problem.

5 Conclusions and Future Work

We expected that remote robot operators would spend a measurable amount of time acquiring or regaining SA but were surprised that operators spent such a large fraction of their time maintaining SA to the exclusion of all other activities (30% on average). This large time commitment devoted to SA, however, presents clear opportunities for targeting improvements to the human-robot interface. We have mentioned several possible improvements here: a clearer depiction of the robots' position on the map plus the ability for the operator to place landmarks on the map, and information presented to the operator from new sensors that point backwards. Since operators often watched video to the exclusion of the other sensor readouts (and the occasional detriment to their level of SA), we are exploring a design approach that overlays sensor data on the video display [Baker et al. 2004, Hestand and Yanco 2004]. By doing so, we hypothesize that the operator will be able to pay attention to a greater variety of sensor input simultaneously and enhance their SA.

6 Acknowledgements

This work was supported by NIST 70NANB3H1116 and NSF IIS-0308186. Michael Baker and Brenden Keyes of UMass Lowell assisted with the data collection. Thanks to Brian Antonishek, Will Becker, Elena Messina, Jean Scholtz, Brian Weiss, and Jeff Young, all of NIST, for their assistance with data collection in the NIST test arenas. Thanks also to our test subjects and the developers of the robot systems who participated in the tests.

7 References

1. Baker, M., Casey, R., Keyes, B., and Yanco, H. A. (2004). "Improved interfaces for human-robot interaction in an urban search and rescue task." *Proc. IEEE Conference on Systems, Man and Cybernetics*, October 2004, this volume.
2. Dudenhoeffer, D. D., Bruemmer, D. J., Davis, M. L. (2001). "Modeling and simulation for exploring human-robot team interaction requirements." *Winter Simulation Conference 200*1: 730-739
3. Drury, J. L., Scholtz, J., and Yanco, H. A. (2003). "Awareness in human-robot interactions." In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, Washington, DC, October 2003.
4. Endsley, M. R. (1988). "Design and evaluation for situation awareness enhancement." *Proceedings of the Human Factors Society 32nd Annual Meeting*, Santa Monica, CA, Human Factors Society.
5. Endsley, M. R., Selcon, S. J., Hardiman, T. D., and Croft, D. G. (1998). "A comparative analysis of SAGAT and SART for evaluations of situation awareness." *Proceedings of the 42nd annual meeting of the Human Factors and Ergonomics Society*, Chicago, October 1998.
6. Ericsson, K. A. and Simon, H. A. (1980). "Verbal reports as data." *Psychological Review***,** 87: 215 - 251
7. Green, W. E., and Oh, P. Y. (2003). "An aerial robot prototype for situational awareness in closed quarters." *IEEE/RSJ International Conference of Intelligent Robots and Systems (IROS)*, pp. 61-66, Las Vegas, NV, Oct. 2003.
8. Hestand, D. and Yanco, H. A. (2004). "Layered sensor modalities for improved human-robot interaction." *Proc. IEEE Conference on Systems, Man and Cybernetics*, October 2004, this volume.
9. Hjelmfelt, A. T. and Pokrant, M. A. (1998). "Coherent tactical picture," CNA RM 97-129, Alexandria, Virginia: Center for Naval Analyses.
10. Hughes, S., and Lewis, M. (2004). "Robotic camera control for remote exploration." *Proc. CHI 2004 Conference on Human Factors in Computing Systems*, April 2004, Vienna, Austria, 511- 517.
11. Jacoff, A., Messina, E., and Evans, J. (2000). "A standard test course for urban search and rescue robots." *Proceedings of the Performance Metrics for Intelligent Systems Workshop*, August 2000.
12. Johnson, C. A., Koku, B., Kawamura, K., and Peters II, R. A. (2002). "Enhancing a human-robot interface using sensory egosphere." *Proc. IEEE International Conference on Robotics and Automation*, Washington D.C., May 2002. IEEE Press: pp. 4132 – 4137.
13. Leveson, N. G. (1986). "Software safety: why, what and how." *ACM Computing Surveys* **18**(2): 125 – 162, June 1986.
14. Lewis, M., Wang, J., Hughes, S., and Liu, X. (2003). "Experiments with attitude: attitude displays for teleoperation." *Proc. 2003 IEEE International Conference on Systems, Man, and Cybernetics*, Washington, DC, October 5-8, pp. 1345-1349.
15. McGuinness, B. (1999). "Situational awareness and the CREW awareness rating scale (CARS)." *Proc. 1999 Avionics Conference*, Heathrow, 17 – 18 November 1999. ERA Technology Report 99-0815 (paper 4.3)
16. McGuinness, B. and Ebbage, L. (2002). "Assessing human ractors in command and control: workload and situational awareness metrics." *Proceedings of the 2002 Command and Control Research and Technology Symposium*, Monterey, CA.
17. Scholtz, J. (2003). Theory and evaluation of human-robot interactions. In *Proceedings of the Hawaii International Conference on Systems Sciences*, January 2003.
18. Taylor, R. M. (1990). "Situational awareness rating technique (SART): The development of a tool for aircrew systems design." *Situational Awarenss in Aerospace Operations (AGARD-CP-478)*, pp. 3/1 – 3/17. Neuilly Sure Seine, France: NATO-AGARD.
19. Yanco, H. A., Drury, J. L., and Scholtz, J. (2004). "Beyond usability evaluation: analysis of human-robot interaction at a major robotics competition." *Human-Computer Interaction*, Vol. 19, No. 1 & 2, pp. 117 - 149.