The Impact of Type and Level of Automation on Situation Awareness and Performance in Human-Robot Interaction

David Schuster, Florian Jentsch, Thomas Fincannon, and Scott Ososky

University of Central Florida, Institute for Simulation and Training 3100 Technology Parkway, Orlando, FL 32826
{dschuster,sososky}@ist.ucf.edu, florian.jentsch@ucf.edu, t_fincannon@yahoo.com

Abstract. In highly autonomous robotic systems, human operators are able to attend to their own, separate tasks, rather than directly operating the robot to accomplish their immediate task(s). At the same time, as operators attend to their own, separate tasks that do not directly involve the robotic system, they can end up lacking situation awareness (SA) when called on to recover from automation failure or from an unexpected event. In this paper, we describe the mechanisms of this problem, known as the out-of-the-loop performance problem, and describe why the problem may still exist in future robotic systems. Existing solutions to the problem, which focus on the level of automation, are reviewed. We describe our current empirical work, which aims to expand upon taxonomies of levels of automation to better understand how engineers of robotic systems may mitigate the problem.

Keywords: Human-robot interaction, robot design, situation awareness, automation.

1 Introduction

In highly autonomous robotic systems, human operators are able to attend to their own, separate tasks, rather than having to directly operate a robot to accomplish their immediate task(s). Because this is one of these systems' major benefits, robots are being developed to have high levels of autonomy, so that they can function with as little human intervention as possible [1]. Still, as robots grow in capability, they will continue to need some human supervision and occasional human intervention [2].

Human performance problems in complex and automated systems have been discussed for some time [3] [4]. Sarter [5] was one of the first to broaden the contributing factors to the problem to those beyond automation complacency [6]. For example, Sarter identified a lack of SA as a mediator between high levels of automation and poor performance. Endsley and Kiris [7] found that the out-of-the-loop performance problem affected humans working in highly automated systems. In the out-of-the-loop performance problem, operators who attend to their own, separate tasks can end up

D. Harris (Ed.): EPCE/HCII 2013, Part I, LNAI 8019, pp. 252–260, 2013.

lacking SA when they are called upon suddenly, for example, to recover from automation failures or in response to another unexpected event. The phenomenon of the out-of-the-loop performance problem has been studied in multiple domains, for example, among pilots where a loss of SA has been observed under conditions of high cockpit automation [8].

In this paper, we describe the mechanisms of the out-of-the-loop performance problem and describe why the problem may still exist in future robot systems. Existing solutions to the problem, which focus on the level of automation, are reviewed. We also describe our empirical work, which aims to expand upon taxonomies of levels of automation to better understand how engineers of robotic systems may mitigate the problem.

2 The Out-of-the-Loop Performance Problem in HRI

A common problem of capable, autonomous robots is poor SA [9]. As we introduced above, lack of SA under conditions of automation failure has come to be known as the out-of-the-loop performance problem [10] and the out-of-the-loop unfamiliarity problem [9]. Further, robustness in robot capabilities remains a challenge to robotics [11]. As new capabilities are developed, robots may be able to perform new tasks, but reliability may be limited, especially initially. For human operators to take advantage of new robot capabilities, operators must be able to recover from robot failures.

The out-of-the-loop performance problem occurs because of a shift in the importance of information within the environment that is the result of a change in the level of automation. Operators working with high-autonomy automation typically do not need the details of the automation's task to achieve their own goals; in fact, the cardinal objective of the automation is to off-load the operator of having to obtain and hold detailed knowledge and to perform continuous monitoring. Consequently, high-performing operators can be expected to ignore these details and devote their resources to their own task components. Unfortunately, failure of the automation leads to a rapid switch from these details being unimportant to suddenly being of crucial importance to task completion. The result can be very poor SA, especially when the operator's current knowledge overlaps minimally with the knowledge needed to effectively achieve task goals.

This problem is likely to be of continued concern in future robot systems. These systems are designed to offload the human through high autonomy. Additionally, a factor contributing to the out-of-the-loop performance problem, namely, imperfect reliability, is likely present, at least initially. Further, when these robots fail, they may do so in non-obvious ways. Indeed, in such systems, the need for a shift from full autonomy to operator intervention may not be marked by clearly pronounced signals, unlike in many current systems or when an obvious failure occurs in a complete subsystem. It is fairly obvious, for example, when a pilot needs to revert to an alternate method of navigation or control when an instrument in the cockpit goes completely blank or gives a value that is clearly out of the possible range. In this case, a combination of signals in the automation (the blank display or wrong value), combined with the pilot's knowledge and SA, signal that a task previously handled by automation must be performed manually.

In contrast, robot failures may include instances in which a robot provides apparently valid information that is based on incorrect sensing. Although this type of failure may occur in other systems, and robots also fail in obvious ways, a pressing problem is dealing with subtle sensing errors in a system separated from its operator by location and task assignment. That is, the operator and robot are not only in different locations, they also have different roles, making it difficult, impossible, or unnecessary to monitor each action of the robot. These types of failures are not mechanical failures (such as when the robot is stuck and cannot complete the task), but rather failures in robot sensing and intelligence (i.e., the robot completes the task but does so incorrectly). Furthermore, measuring a robot's confidence or meta-awareness of their sensors continues to be a more difficult problem than sensing itself [12]. Consequently, robot mistakes may be detectable only through cross-checking with other data, and failure at lower levels (i.e., sensing) may only have noticeable consequences at higher levels (i.e., decision-making). From the operator's perspective, a shift must take place when a robot fails; what previously did not need to be known by the operator suddenly becomes critically important. To maintain SA, interventions are thus needed to support the operator's information processing under robot unreliability. These will be important considerations for system design so that humans can effectively use the capabilities of future robots.

In summary, SA is an important construct that mediates the relationship between autonomy and performance. We assert that without consideration of other task factors, high levels of robot autonomy will lead to poor SA. Next, we discuss the potential solutions to this problem developed in other systems and how they may apply to human-robot interaction.

3 Potential Solutions

Since there are benefits to automation and operator involvement [10], a solution to the problem of "which agent does what" is needed. The solution should be human-centered and bound to the mission context [13]; it also needs to be applicable to human-robot interaction.

3.1 Selecting a Level of Automation

The degree to which a robot is involved in the task has been called the level of robot autonomy. Autonomy has origins in the Greek word autonomia, which means independence [14]. The United States Department of Defense defines an autonomous weapon system as "a weapon system that, once activated, can select and engage targets without further intervention by a human operator" [15] and an autonomous battlefield entity as one "that does not require the presence of another battlefield entity in order to conduct its own simulation in the battlefield environment" [16].

Extending these definitions to robots, their autonomy is the extent to which the behavior of a robot results from integration of its own sensing [17] and the extent to which it makes decisions not mediated by other entities, in particular, humans [18]. Robot autonomy is generally discussed as a quality integrating the robot's capabilities and authority across a series of tasks [19]. Although autonomy is an intuitive quality of robots, there are few developed quantitative models that describe robot autonomy, and fewer still that specify an operational definition.

Robot behavior can be considered to belong to the broader class of automation. Automation is any "device or system that accomplishes (partially or fully) a goal that was previously, or conceivably could be, carried out (partially or fully) by a human operator" [19]. By considering robot behavior as automation, one can describe the robot's involvement in a particular task as the level of automation.

3.2 *What* Is Automated, Not Only *How Much* Automation

The general case of automation has been widely studied [19]. An early taxonomy applicable to robots was developed by Sheridan and Verplank [20] and expanded upon by Parasuraman, Sheridan, and Wickens [19]. Importantly, their taxonomy expanded upon prior models by including what task is automated in addition to how much automation is used. Under this model, the level of automation for a robot can be described as: (a) the levels of information processing in which the robot participants (i.e., what), and (b) the conditions under which the robot participates in each process (i.e., how much).

The first two levels of this model map cleanly onto Endsley's [21] levels of SA [22]. The first two levels of the model are the two levels of diagnostic aiding [23]. Robots perform the information acquisition stage of diagnostic aiding when they gather relevant information through their sensors. Robots perform the second stage, information analysis, when they integrate multiple pieces of sensor data or when they integrate sensor data with previously stored or externally provided information. Thus, information acquisition is a precursor to information analysis, and a robot that performs both stages operates at a higher level of automation than one that only performs information acquisition. Automation that provides the higher level of diagnostic aiding leads to better decision making [24] and performance [25] in operators while lowering their workload [26].

As discussed, the what of automation has been modeled by Parasuraman et al. [19]. However, much of the applied literature has taken a how much approach to measuring and manipulating the level of automation. Although it is understandably easier to manipulate the presence or functionality of an entire system, research needs to specify the stage as well as the amount of automation. Horrey and Wickens [27] adapted this approach and found that both information acquisition (stage 1 diagnostic aiding) and information analysis (stage 2 diagnostic aiding) lead to better performance than an unaided condition on a battlefield simulation task, with the information analysis aid leading to a greater reduction in errors compared with the information acquisition aid. However, memory probe questions suggested that relevant items were processed more deeply with the lower level of diagnostic aiding [27].

Adding automation of information analysis to automation of information acquisition has been shown to have a greater effect on decisions than automation of information acquisition alone. In a study of anesthesiologists, nurses, and hospital housekeepers, for example, operating room management information was presented as either a command display, which provided recommendations, or a status display, which made decision-relevant information available [24]. When making decisions in subsequent scenarios, participants without either type of aid performed less accurately than random chance. Decision-making, both a cognitive outcome and a performance measure, was improved only by the command display (status displays did not have a significant effect on decision accuracy). Further, incorrect command displays had greater costs associated with them for trust, and users were more likely to follow erroneous recommendations that did not affect safety. From this, Dexter et al. [24] concluded that command displays were preferable, but carry additional costs when their recommendations are incorrect.

Although robot diagnostic aiding may be beneficial to SA, the literature suggests that this relationship is highly sensitive to the presence of unreliability in the robot, and that the two levels of diagnostic aiding may be differentially affected. Performance decreases as the reliability of a diagnostic aid falls [28]. While unreliable information negatively impacts performance, the effect may be stronger for information analysis automation than for information acquisition automation [29] [30].

Sarter and Scroeder [30] found that a diagnostic aid that provided recommendations (information analysis), rather than status information (information acquisition), had a greater performance cost when the automation was not reliable. Rovira and colleagues [29] found that unreliability degraded operator accuracy at three levels of increasingly automated information analysis. Unreliability did not have a significant effect on accuracy in the information acquisition condition, however.

A similar pattern of results was found in an airplane identification task by Crocoll and Coury [31]. In a study manipulating status (acquisition) and recommendation (analysis) information, the group receiving only status information was the least affected by inaccuracy in the automation. In line with this finding, Skitka, Mosier, and Burdick [32] found that introduction of imperfect automation that monitored system state lead to an increase in missed events.

Parasuraman and Wickens provided a cognitive explanation for why lower levels of diagnostic aiding may lead to better SA: "The user must continue to generate the values for the different courses of action. As a result, users may be more aware of the consequences of the choice and of the possibility that the choice may be incorrect because of a faulty automated diagnosis" [33]. This may explain the empirical findings of Horrey and Wickens [27]. When operators perform information analysis, they perform additional processing that may keep them "in the loop". Consequently, the operators' information analysis should lead to better SA during robot unreliability.

Galster, Bolia, and Parasuraman [34] found that performance on a target detection task improved when an information status cue was added, even though this cue was not perfectly reliable. When a higher level of aiding was added in the form of decision suggestion, performance was not improved unless the information status cue was also included. This suggests that under conditions of unreliability, operators may be able to recover from erroneous information provided by information acquisition automation more easily than from information analysis automation. In summary, when reliability is limited, access to lower-level data can help an operator to remain in the loop [35].

4 Discussion and On-Going Research

The out-of-the-loop performance problem affects systems that separate the human from the automation during normal operation but need human intervention if something goes wrong. Future robots will be subject to limited reliability as they deal with the complexity of their operational environments and interact with humans. Because of this, the robot will be dependent upon its human operators for successful performance, just as humans will depend upon the robot to execute mission goals. To ensure that humans do not lose SA, it is imperative that system designers consider how an operator's cognition may be affected by the autonomy and reliability of the robot.

We assert that the levels of information processing at which a robot offers assistance provide an important indicator of how the operator's SA will be affected. In other words, robots that are capable of gathering and integrating information can support higher levels of SA in the human, but only if the robot is able to perform this task reliably. When reliability is very high and unreliability has minimal impact on mission effectiveness, operators are likely to benefit from automation that integrates information and operates independently. Under the more realistic scenario of robot unreliability, it may be more beneficial for the robot to provide less integrated information to support lower levels of SA. Taking this approach could help the human mitigate the effects of robot unreliability.

Because the level of automation is both what and how much, the solution to implementing the right level of automation in a robot system will depend on the information and work needs of the human operator as well as the anticipated reliability of the robot system. It is important to consider these as two separate factors, each contributing to the autonomy of the robot. It may be neither sufficient nor desirable to lower the level of automation. Instead, high levels of automation may be possible by designing the robot to keep the human in the loop through the provision of relevant, low-level information. By strategically selecting what is automated, the robot may be able to perform its assigned functions more independently. In this way, the human will be able to maintain SA while benefiting from a robot capable of independent sensing and behavior.

This is the topic of our current empirical investigation. The purpose of our study is to determine the conditions under which diagnostic aiding contributes to operator SA given limitations of robot reliability, on the one hand, and unaided human task performance, on the other. Diagnostic aids that perform information analysis as well as information acquisition should lead to higher levels of SA than information acquisition alone. However, this relationship has been observed only under cases of perfect robot reliability. We expect the opposite under conditions of imperfect reliability. That is, information acquisition will be more beneficial form of automation than information analysis. By providing lower-level information, the human should remain in the loop. This study will confirm and extend the findings discussed while providing insight into the conditions under which the out-of-the-loop performance problem can be mitigated.

Our ultimate goal is to expand upon the existing understanding of the levels of automation by considering how the type of automation may affect levels of human SA and how that relationship varies as a function of robot unreliability. Ultimately, this may provide a path towards maximizing human SA in human-robot interaction.

Acknowledgements. The research reported in this document/presentation was performed in connection with Contract Number W911NF-10-2-0016 with the U.S. Army Research Laboratory. The views and conclusions contained in this document/presentation are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory, or the U.S. Government unless so designated by other authorized documents. Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation heron.

References

1. Cosenzo, K., Parasuraman, R., De Visser, E.: Automation Strategies for Facilitating Human Interaction with Military Unmanned Vehicles. In: Barnes, M., Jentsch, F. (eds.) Human-Robot Interactions in Future Military Operations, pp. 103–124. Ashgate, Surrey (2010)
2. Burke, J.L., Murphy, R.R., Coovert, M.D., Riddle, D.L.: Moonlight in Miami: A Field Study of Human-Robot Interaction in the Context of an Urban Search and Rescue Disaster Response Training Exercise. Human-Computer Interaction 19, 85–116 (2004)
3. Norman, D.A.: The 'Problem' with Automation: Inappropriate Feedback and Interaction, not 'Over-Automation'. Philosophical Transactions of the Royal Society of London, Series B, Biological Sciences 327(1241), 585–593 (1990)
4. Wiener, E.L.: Beyond the Sterile Cockpit. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 27(1), 75–90 (1985)
5. Sarter, N.B.: Strong, Silent, and Out-of-the-Loop: Properties of Advanced (Cockpit) Automation and their Impact on Human-Automation Interaction. Unpublished doctoral dissertation, Ohio State University, Columbus, OH (1994)
6. Parasuraman, R., Molloy, R., Singh, I.L.: Performance Consequences of Automation-induced "complacency". International Journal of Aviation Psychology 3(1), 1–23 (1993)
7. Endsley, M.R., Kiris, E.O.: The Out-of-the-Loop Performance Problem and Level of Control in Automation. Human Factors 37, 381–394 (1995)
8. Endsley, M.R.: Direct Measurement of Situation Awareness: Validity and Use of SAGAT. In: Endsley, M.R., Garland, D.J. (eds.) Situation Awareness Analysis and Measurement, pp. 147–173. Lawrence Erlbaum Associates, Mahwah (2000)
9. Wickens, C.D.: Spatial Awareness Biases. Technical report. University of Illinois at Urbana-Champaign, Savoy, Illinois (2002)

10. Endsley, M.R., Kiris, E.O.: The Out-of-the-Loop Performance Problem and Level of Control in Automation. Human Factors 37, 381–394 (1995)
11. Stancliff, S., Dolan, J.M., Trebi-Ollennu, A.: Towards a Predictive Model of Robot Reliability. Technical report, Carnegie Mellon University (2005)
12. Eski, I., Erkaya, S., Savas, S., Yildirim, S.: Fault Detection on Robot Manipulators Using Artificial Neural Networks. Robotics and Computer-Integrated Manufacturing 27, 115–123 (2010)
13. Sierhuis, M., Bradshaw, J.M., Acquisti, A., Van Hoof, R., Jeffers, R., Uszok, A.: Human-Agent Teamwork and Adjustable Autonomy in Practice. In: Seventh International Symposium on Artificial Intelligence, Robotics and Automation in Space, Nara, Japan (2003)
14. U.S. Department of Defense: Autonomy in Weapon Systems. Directive Number 3000.09. Department of Defense, Washington, DC (2012)
15. Autonomy, `http://www.etymonline.com/index.php?term=autonomy`
16. United States Department of Defense Department of Defense Modeling and Simulation glossary. Directive No. 5000.59-M. United States Department of Defense, Washington, DC (1998)
17. Franklin, S., Graesser, A.: Is It An Agent or Just a Program? A Taxonomy for Autonomous Agents. In: Jennings, N.R., Wooldridge, M.J., Müller, J.P. (eds.) ECAI-WS 1996 and ATAL 1996. LNCS, vol. 1193, pp. 21–35. Springer, Heidelberg (1997)
18. Luck, M., D'lnverno, M.: A Formal Framework for Agency and Autonomy. In: Proceedings of the First International Conference on Multiagent Systems, vol. 1, pp. 254–260. AAAI Press/MIT Press, San Francisco, CA (1995)
19. Parasuraman, R., Sheridan, T.B., Wickens, C.D.: A Model for Types and Levels of Human Interaction with Automation. IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans 30(3), 287 (2000)
20. Sheridan, T.B., Verplank, W.: Human and Computer Control of Undersea Teleoperators. Technical report, Man-Machine Systems Laboratory, Cambridge, MA (1978)
21. Endsley, M.R.: Situation Awareness Global Assessment Technique (SAGAT). IEEE 1988 National Aerospace and Electronics Conference 3, 789–795 (1988)
22. Horrey, W.J., Wickens, C.D., Strauss, R., Kirlik, A., Stewart, T.R.: Supporting Situation Asessment through Attention Guidance and Diagnostic Aiding: The Benefits and Costs of Display Enhancement on Judgment Skill. Oxford University Press, Oxford (2006)
23. Wickens, C.D., Dixon, S.R.: The Benefits of Imperfect Diagnostic Automation: A Synthesis of the Literature. Theoretical Issues in Ergonomics Science 8(3), 201–212 (2007)
24. Dexter, F., Willemsen-Dunlap, A., Lee, J.D.: Operating Room Managerial Decision-Making on the Day of Surgery With and Without Computer Recommendations and Status Displays. Anesthesia & Analgesia 105(2), 419–429 (2007)
25. Goodrich, M.A., McLain, T.W., Anderson, J.D., Sun, J., Crandall, J.W.: Managing Autonomy in Robot Teams: Observations from Four Experiments. In: ACM International Conference on Human-Robot Interaction, pp. 25–32 (2007)
26. Manzey, D., Reichenbach, J., Onnasch, L.: Human Performance Consequences of Automated Decision Aids: The Impact of Degree of Automation and System Experience. Journal of Cognitive Engineering and Decision Making 6(1), 57–87 (2012)
27. Horrey, W.J., Wickens, C.D.: Supporting Battlefield Aituation Assessment Through Attention Guidance and Diagnostic Aiding: A Cost-Benefit and Depth of Processing Analysis. Technical report. Aviation Research Lab, Savoy, IL (2001)
28. Madhavan, P., Phillips, R.R.: Effects of Computer Self-Efficacy and System Reliability on User Interaction with Decision Support Systems. Computers in Human Behavior 26(2), 199–204 (2010)

29. Rovira, E., McGarry, K., Parasuraman, R.: Effects of Imperfect Automation on Decision Making in a Simulated Command and Control Task. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 49(1), 76–87 (2007)
30. Sarter, N.B., Schroeder, B.: Supporting Decision Making and Action Selection under Time Pressure and Uncertainty: The Case of In-flight Icing. Human Factors 43(4), 573–583 (2001)
31. Crocoll, W.M., Coury, B.G.: Status or Recommendation: Selecting the Type of Information for Decision Aiding. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 34(19), 1524–1528 (1990)
32. Skitka, L.J., Mosier, K.L., Burdick, M.: Does Automation Bias Decision-Making? International Journal of Human-Computer Studies 51, 991–1006 (1999)
33. Parasuraman, R., Wickens, C.D.: Humans: Still Vital After All These Years of Automation. Human Factors: The Journal of the Human Factors and Ergonomics Society 50(3), 514 (2008)
34. Galster, S.M., Bolia, R.S., Parasuraman, R.: The Application of a Qualitative Model of Human-Interaction with Automation: Effects of Unreliable Automation on Performance. In: Tielemans, W. (chair) The Role of Humans in Intelligent and Automated Systems. North Atlantic Treaty Organization, Warsaw (2002)
35. Johnson, R.C., Saboe, K.N., Prewett, M.S., Coovert, M.D., Elliott, L.R.: Autonomy and automation reliability in human-robot interaction: A qualitative review. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 53, 1398–1402 (2009)