

THE INTERNATIONAL
C2 JOURNAL

VOLUME 1, NUMBER 2, 2007

SPECIAL ISSUE

*Decision Support for Network-Centric
Command and Control*

GUEST EDITOR

Sandeep Mulgund
The MITRE Corporation

Adaptive Automation for Human-Robot Teaming
in Future Command and Control Systems

Raja Parasuraman
Michael Barnes
Keryl Cosenzo



THE INTERNATIONAL C2 JOURNAL

David S. Alberts, Chairman of the Editorial Board, *OASD-III, CCRP*
Joseph R. Lewis, Managing Editor

The Editorial Board

Berndt Brehmer (SWE), *Swedish National Defence College*
Reiner Huber (GER), *Universitaet der Bundeswehr Muenchen*
Viggo Lemche (DEN), *Danish Defence Acquisition and Logistics Organization*
James Moffat (UK), *Defence Science and Technology Laboratory (DSTL)*
Mark Nissen (USA), *Naval Postgraduate School*
Ross Pigeau (CAN), *Defence Research and Development Canada (DRDC)*
Mink Spaans (NED), *TNO Defence, Security and Safety*

Associate Editors

Gerard Christman, *U.S. OSD Technical Services - Femme Comp Inc.*
R. Scott Cost, *Johns Hopkins University Applied Physics Laboratory*
Raymond J. Curts, *Strategic Consulting, Inc*
Paul K. Davis, *RAND Corporation*
Petra M. Eggenhofer, *Munich Bundeswehr University*
Elliot Entin, *Aptima*
Stuart Grant, *Defence Research and Development Canada (DRDC)*
Juergen Grosche, *FGAN-FKIE, Germany*
Paul Labbé, *Defence Research and Development Canada*
Michael Malm, *Swedish Defence Research Agency*
Sandeep Mulgund, *The MITRE Corporation*
Philip W. Pratt, *Northrop Grumman*
Jens Roemer, *Fachschule der Bundeswehr für Informatiktechnik*
Pamela Savage-Knepshield, *U.S. Army Research Laboratory, Human Research & Engineering Directorate*
Keith Stewart, *Defence Research and Development Canada (DRDC)*
Andreas Tolk, *Old Dominion University*

About the Journal

The International C2 Journal was created in 2006 at the urging of an international group of command and control professionals including individuals from academia, industry, government, and the military. The Command and Control Research Program (CCRP, of the U.S. Office of the Assistant Secretary of Defense for Networks and Information Integration, or OASD-NII) responded to this need by bringing together interested professionals to shape the purpose and guide the execution of such a journal. Today, the Journal is overseen by an Editorial Board comprising representatives from many nations.

Opinions, conclusions, and recommendations expressed or implied within are solely those of the authors. They do not necessarily represent the views of the Department of Defense, or any other U.S. Government agency.

Rights and Permissions: All articles published in the International C2 Journal remain the intellectual property of the authors and may not be distributed or sold without the express written consent of the authors.

For more information

Visit us online at: www.dodccrp.org

Contact our staff at: publications@dodccrp.org



Adaptive Automation for Human-Robot Teaming in Future Command and Control Systems

Raja Parasuraman (George Mason University)

Michael Barnes (Army Research Laboratory)

Keryl Cosenzo (Army Research Laboratory)

Abstract

Advanced command and control (C2) systems such as the U.S. Army's Future Combat Systems (FCS) will increasingly use more flexible, reconfigurable components including numerous robotic (uninhabited) air and ground vehicles. Human operators will be involved in supervisory control of uninhabited vehicles (UVs) with the need for occasional manual intervention. This paper discusses the design of automation support in C2 systems with multiple UVs. Following a model of effective human-automation interaction design (Parasuraman et al. 2000), we propose that operators can best be supported by high-level automation of information acquisition and analysis functions. Automation of decisionmaking functions, on the other hand, should be set at a moderate level, unless 100 percent reliability can be assured. The use of adaptive automation support technologies is also discussed. We present a framework for adaptive and adaptable processes as methods that can enhance human-system performance while avoiding some of the common pitfalls of "static" automation such as over-reliance, skill degradation, and reduced situation awareness. Adaptive automation invocation processes are based on critical mission events, operator modeling, and real-time operator performance and physiological assessment, or hybrid combinations of these methods. We describe

the results of human-in-the-loop experiments involving human operator supervision of multiple UVs under multi-task conditions in simulations of reconnaissance missions. The results support the use of adaptive automation to enhance human-system performance in supervision of multiple UVs, balance operator workload, and enhance situation awareness. Implications for the design and fielding of adaptive automation architectures for C2 systems involving UVs are discussed.

Introduction

Unmanned air and ground vehicles are an integral part of advanced command and control (C2) systems. In the U.S. Army's Future Combat Systems (FCS), for example, uninhabited vehicles (UVs) will be an essential part of the future force because they can extend manned capabilities, act as force multipliers, and most importantly, save lives (Barnes et al. 2006). The human operators of these systems will be involved in supervisory control of semi-autonomous UVs with the need for occasional manual intervention. In the extreme case, soldiers will control UVs while on the move and while under enemy fire.

All levels of the command structure will use robotic assets such as UVs and the information they provide. Control of these assets will no longer be the responsibility of a few specially trained soldiers but the responsibility of many. As a result, the addition of UVs can be considered a burden on the soldier if not integrated appropriately into the system. Workload and stress will be variable and unpredictable, changing rapidly as a function of the military environment. Because of the likely increase in the cognitive workload demands on the soldier, automation will be needed to support timely decision-making. For example, sensor fusion systems and automated decision aids may allow tactical decisions to be made more rapidly, thereby shortening the "sensor-to-shooter" loop (Adams 2001; Rovira et al. 2007). Automation support will also be mandated because of the high cognitive workload involved in supervising *multiple* unmanned

air combat vehicles (for an example involving tactical Tomahawk missiles, see Cummings and Guerlain 2007).

The automation of basic control functions such as avionics, collision avoidance, and path planning has been extensively studied and as a result has not posed major design challenges (although there is still room for improvement). How information-gathering and decision-making functions should be automated is less well understood. However, the knowledge gap has narrowed in recent years as more research has been conducted on human-automation interaction (Billings 1997; Parasuraman and Riley 1997; Sarter et al. 1997; Wiener and Curry 1980). In particular, Parasuraman et al. (2000) proposed a model for the design of automated systems in which automation is differentially applied at different levels (from low, or fully manual operation, to high, or fully autonomous machine operation) to different types or stages of information-processing functions. In this paper we apply this model to identify automation types best suited to support operators in C2 systems involved in interacting with multiple UVs and other assets during multi-tasking missions under time pressure and stress. We describe the results of two experiments. First, we provide a validation of the automation model in a study of simulated C2 involving a battlefield engagement task. This simulation study did not involve UVs, but multi-UV supervision is examined in a second study.

We propose that automation of early-stage functions—information acquisition and analysis—can, if necessary, be pursued to a very high level and provide effective support of operators in C2 systems. On the other hand, automation of decisionmaking functions should be set at a moderate level unless very high-reliability decision algorithms can be assured, which is rarely the case. Decision aids that are not perfectly reliable or sufficiently robust under different operational contexts are referred to as *imperfect* automation. The effects on human-system performance of automation imperfection—such as incorrect recommendations, missed alerts, or false alarms—must be considered in evaluating what level of automation to implement (Parasuraman and Riley 1997; Wickens and Dixon 2007). We also

propose that the level and type of automation can be varied during system operations—so-called *adaptive* or *adaptable* automation. A framework for adaptive and adaptable processes based on different automation invocation methods is presented. We describe how adaptive/adaptable automation can enhance human-system performance while avoiding some of the common pitfalls of “static” automation such as operator over-reliance, skill degradation, and reduced situation awareness. Studies of human operators supervising multiple UVs are described to provide empirical support for the efficacy of adaptive/adaptable automation. We conclude with a discussion of adaptive automation architectures that might be incorporated into future C2 systems such as FCS. While our analysis and empirical evidence is focused on systems involving multiple UVs, the conclusions have implications for C2 systems in general.

Automation of Information Acquisition and Analysis

Many forms of UVs are being introduced into future C2 systems in an effort to transform the modern battle space. One goal is to have component robotic assets be as autonomous as possible. This requires both rapid response capabilities and intelligence built into the system. However, ultimate responsibility for system outcomes always resides with the human; and in practice, even highly automated systems usually have some degree of human supervisory control. Particularly in combat, some oversight and the capability to override and control lethal systems will always be a human responsibility for the reasons of system safety, changes in the commander’s goals, and avoidance of fratricide, as well as to cope with unanticipated events that cannot be handled by automation. This necessarily means that the highest level of automation (Sheridan and Verplank 1978) can rarely be achieved except for simple control functions.

The critical design issue thus becomes: What should the level and type of automation be for effective support of the operator in such systems (Parasuraman et al. 2000)? Unfortunately, automated aids have not always enhanced system performance, primarily due to

problems in their use by human operators or to unanticipated interactions with other sub-systems. Problems in human-automation interaction have included unbalanced mental workload, reduced situation awareness, decision biases, mistrust, over-reliance, and complacency (Billings 1997; Parasuraman and Riley 1997; Sarter et al. 1997; Sheridan 2002; Wiener 1988).

Parasuraman et al. (2000) proposed that these unwanted costs might be minimized by careful consideration of different information-processing functions that can be automated. Their model for effective human-automation interaction design identifies four stages of human information processing that may be supported by automation: (stage 1) information acquisition; (stage 2) information analysis; (stage 3) decision and action selection; and (stage 4) action implementation. Each of these stages may be supported by automation to varying degrees between the extremes of manual performance and full automation (Sheridan and Verplank 1978). Because they deal with distinct aspects of information processing, the first two stages (information acquisition and analysis) and the last two stages (decision selection and action implementation) are sometimes grouped together and referred to as *information* and *decision* automation, respectively (see also Billings 1997). The Theater High Altitude Area Defense (THAAD) system used for intercepting ballistic missiles (Department of the Army 2007) is an example of a fielded system in which automation is applied to different stages and at different levels. THAAD has relatively high levels of information acquisition, information analysis, and decision selection; however, action implementation automation is low, giving the human control over the execution of a specific action.

Parasuraman et al. (2000) suggested that automation could be applied at very high levels without any significant performance costs to early-stage functions, i.e. information acquisition and analysis, particularly if the automation algorithms used were highly reliable. However, they suggested that for high-risk decisions involving considerations of lethality or safety, decision automation should be set at a moderate level such that the human operator is not only

involved in the component processes involved in the decisionmaking cycle, but also makes the final decision (see also Wickens et al. 1998). As noted previously, decision aids are often imperfect. Several empirical studies have compared the effects on operator performance of automation imperfection at different information-processing stages. Given that perfectly (100 percent) reliable automation cannot be assured, particularly for decisionmaking functions, it is important to examine the potential costs of such imperfection. Crocoll and Coury (1990) conducted one such study. They had participants perform an air-defense targeting (identification and engagement) task with automation support that was not perfectly reliable. The automation provided either status information about a target (information automation) or a recommendation concerning its identification (decision automation). Crocoll and Coury found that the cost of imperfect advice by the automation was greater when participants were given a recommendation to implement (decision automation) than when given only status information, which they had to use to make their own decision (information automation).

A more extensive study by Rovira et al. (2007) using a more realistic simulation confirmed and extended these findings. They examined performance in a sensor-to-shooter simulation involving a battlefield engagement task with automation support. The simulation consisted of three components shown in separate windows of a display: a terrain view, a task window, and a communications module (see Figure 1). The right portion of the display showed a two-dimensional terrain view of a simulated battlefield, with identified enemy units (red), friendly battalion units (yellow), artillery units (green), and one blue friendly headquarter unit. The participants made enemy-friendly engagement selections with the aid of automation in the task window. They were required to identify the most dangerous enemy target and to select a corresponding friendly unit to engage in combat with the target. Like Crocoll and Coury (1990), Rovira et al. (2007) also compared information and decision automation, but examined three different forms of decision automation to examine the generality of the effect. The information automation gave a list

of possible enemy-friendly unit engagement combinations with information (such as the distance between units, distance from headquarters, etc.), but left the decision to the operator. The three types of decision automation gave the operator: (1) a prioritized list of all possible enemy-friendly engagement choices; (2) the top three options for engagement; or (3) the best choice (these were termed low, medium, and high decision automation, respectively). Also, importantly, Rovira et al. (2007) varied the level of automation unreliability, which was fixed in previous studies. The decision aid provided recommendations that were correct 80 percent or 60 percent of the time at the high and low levels of reliability. This manipulation of automation reliability allowed for an assessment of the range of automation imperfection effects on decisionmaking, as well as a comparison to previous work on complacency in which automation reliability was varied and operator reliance on automated aiding was evaluated (Parasuraman et al. 1993).

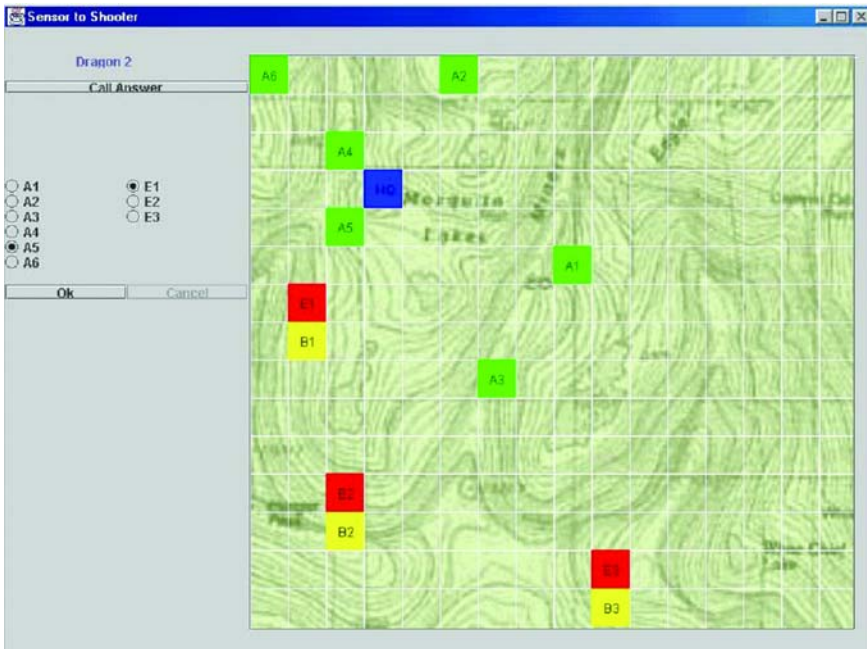


Figure 1. Sensor to shooter simulation of battlefield engagement task, with automated support (Rovira et al. 2007).

Rovira et al. (2007) found that when reliable, automation enhanced the accuracy of battlefield engagement decisions. This benefit was greater for decision than for information automation, with the benefit being greater as the level of decision automation (advisories) increased from low (prioritized list of all possible options) to medium (top three choices) to high (best choice) (see Figure 2). More importantly, reliable decision automation significantly reduced the time for target engagement decisions. This is important for shortening the overall sensor-to-shooter time in tactical C2 operations. However, when the automation provided incorrect assessments on occasion, the accuracy of target engagement decisions at those times declined drastically for all three different levels of decision automation (low, medium, and high), but not for information automation. The accuracy of engagement decisions was significantly reduced (see Figure 2) when incorrect advisories were provided by the automation, but this cost was obtained only for decision and not for information automation.

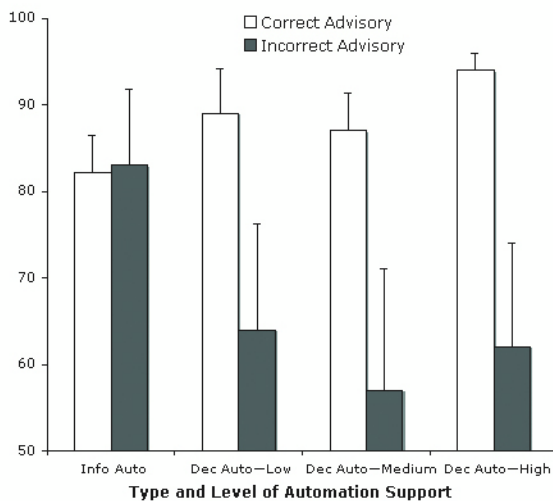


Figure 2. Decision accuracy under different forms of automation support for correct and incorrect advisories (Rovira et al. 2007).

These results confirm the previous study of Crocoll and Coury (1990) and point to the generality of the effect for levels of decision automation in C2 systems. Furthermore, this effect was particularly prominent at a high (80 percent) overall level of automation reliability and less so at a lower reliability (60 percent). This finding reflects one of the paradoxes of decision automation: for imperfect (less than 100 percent) automation, the greater the reliability, the greater the chance of operator over-reliance (because of the rarity of incorrect automation advisories), with a consequent cost of the operator uncritically following unreliable advice. These results also demonstrate that the cost of unreliable decision automation occurs across multiple levels of decision automation, from low to high.

These findings may be interpreted within the framework of the human-automation interaction model proposed by Parasuraman et al. (2000). Information automation gives the operator status information, integrates different sources of data, and may also recommend possible courses of action, but not commit to any one (Jones et al. 2000). However, this form of automation typically does not give values to the possible courses of action, which decision automation does, and so is not only in some sense “neutral” with respect to the “best” decision, but also does not hide the “raw data” from the operator. Thus for automation that is highly reliable yet imperfect, performance is better with an information support tool because the user continues to generate the values for the different courses of action and, hence, is not as detrimentally influenced by inaccurate information. Additionally, a user of decision automation may no longer create or explore novel alternatives apart from those provided by the automation, thus leading to a greater performance cost when the automation is unreliable.

These findings suggest that in high-risk environments, such as air traffic control and battlefield engagement, decision automation should be set at a moderate level that allows room for operator input into the decisionmaking process. Otherwise there is a risk of operators uncritically following incorrect advisories provided by highly but not perfectly reliable decision support systems. At the

same time, information automation can assist the user in inference and provide information relevant to possible courses of action. As Jones et al. (2000) have demonstrated, very high levels of automation can be effective so long as automated help is confined to information analysis and support of inference generation rather than prescription of best decision outcomes.

Adaptive Automation for Human-Robot Teaming

Designing automation so that different levels of automation are applied appropriately to different stages of information processing represents one method to ensure effective human-system performance. A related approach is to vary the level and type of automation *during* system operations, adaptively, depending on context or operator needs. This defines adaptive or adaptable automation (Opperman 1994): Information or decision support that is not fixed at the design stage but varies appropriately with context in the operational environment. Adaptive/adaptable automation has been proposed as a solution to the problems associated with inflexible or “static” automation (Inagaki 2003; Parasuraman and Miller 2006; Scerbo 2001), although the concept has also been criticized as potentially increasing system unpredictability (Billings and Woods 1994; but see Miller and Parasuraman 2007).

Adaptive systems were proposed over 20 years ago (Hancock et al. 1985; Parasuraman et al. 1992; Rouse 1988), but empirical evaluations of their efficacy have only recently been conducted in such domains as aviation (Parasuraman et al. 1999), air traffic management (Hilburn et al. 1997; Kaber and Endsley 2004), and industrial process control (Moray et al. 2000). Practical examples of adaptive automation systems include the Rotorcraft Pilot’s Associate. This system, which aids Army helicopter pilots in an adaptive manner depending on mission context, has successfully passed both simulator and rigorous in-flight tests (Dornheim 1999). In the context of C2 systems, adaptable automation allows the operator to define conditions for automation decisions during mission planning and

execution (Parasuraman and Miller 2006), while adaptive automation is initiated by the system (without explicit operator input) on the basis of critical mission events, performance, or physiological state (Barnes et al. 2006).

The efficacy of adaptive automation must not be assumed but must be demonstrated in empirical evaluations of human-robot interaction performance. The method of invocation in adaptive systems is a key issue (Barnes et al. 2006). Parasuraman et al. (1992) reviewed the major invocation techniques and divided them into five main categories: (1) critical events; (2) operator performance measurement; (3) operator physiological assessment; (4) operator modeling; and (5) hybrid methods that combine one or more of the previous four methods. For example, in an aircraft air defense system, adaptive automation based on critical events would invoke automation only when certain tactical environmental events occur, such as the beginning of a “pop-up” weapon delivery sequence, which would lead to activation of all defensive measures of the aircraft (Barnes and Grossman 1985). If the critical events do not occur, the automation is not invoked. Hence this method is inherently flexible and adaptive because it can be tied to current tactics and doctrine during mission planning. This flexibility is limited by the fact that the contingencies and critical events are anticipated. A disadvantage of the method is its insensitivity to actual system and human operator performance. One potential way to overcome this limitation is to measure operator performance and/or physiological activity in real time. For example, operator mental workload may be inferred from performance, physiological, or other measures (Byrne and Parasuraman 1996; Kramer and Parasuraman in press; Parasuraman 2003; Wilson and Russell 2003). The measures can provide inputs to adaptive logic (which could be rule or neural network-based), the output of which invokes automation that supports or advises the operator appropriately, with the goal of balancing workload at some optimum, moderate level (Parasuraman et al. 1999; Wilson and Russell 2003).

Wiener (1988) first noted that many forms of automation are designed in a “clumsy” manner, increasing the operator’s mental workload when his or her task load is already high, or providing aid when it is not needed. An example of the former is when an aircraft flight management system must be reprogrammed, necessitating considerable physical and mental work during the last minutes of final descent, because of a change in runway mandated by air traffic control. Thus, one of the costs of static automation can be *unbalanced* operator workload. There is good evidence that adaptive automation can offset this and balance operator mental workload, thereby protecting system performance from the potential catastrophic effects of overload or underload. For example, Hilburn et al. (1997) examined the effects of adaptive automation on the performance of military air traffic controllers who were provided with a decision aid for determining optimal descent trajectories of aircraft—a Descent Advisor (DA) under varying levels of traffic load. The DA was either present throughout, irrespective of traffic load (static automation), or came on only when the traffic density exceeded a threshold (adaptive automation). Hilburn et al. found significant benefits for controller workload (as assessed using pupillometric and heart rate variability measures) when the DA was provided adaptively during high traffic loads, compared to when it was available throughout (static automation) or only at low traffic loads.

These findings clearly support the use of adaptive automation to balance operator workload in complex human-machine systems. This workload-leveling effect of adaptive automation has also been demonstrated in studies using non-physiological measures. Kaber and Riley (1999), for example, used a secondary-task measurement technique to assess operator workload in a target acquisition task. They found that adaptive computer aiding based on the secondary-task measure enhanced performance on the primary task. The results of these and other studies indicate that adaptive automation can serve to reduce the problem of unbalanced workload, without the attendant high peaks and troughs that static automation often induces.

In addition to triggering adaptive automation on the basis of behavioral or physiological measures of workload, real-time assessment of situation awareness (SA) might also be useful in adaptive systems (Kaber and Endsley 2004). Reduced SA has been identified as a major contributor to poor performance in search and rescue missions with autonomous robots (Murphy 2000). In particular, probing the operator's awareness of changes in the environment might capture transient or dynamic changes in SA. One such measure is *change detection* performance. People often fail to notice changes in visual displays when they occur at the same time as various forms of visual transients or when their attention is focused elsewhere (Simons and Ambinder 2005). Real-time assessment of change detection performance might therefore provide a surrogate measure of SA that could be used to trigger adaptive automated support of the operator and enhance performance. Parasuraman et al. (2007) examined this possibility in a multi-task scenario involving supervision of multiple UAVs in a simulated reconnaissance mission.

Parasuraman et al. (2007) used a simulation designed to isolate some of the cognitive requirements associated with a single operator controlling robotic assets within a larger military environment (Barnes et al. 2006). This simulation, known as the Robotic NCO simulation, required the participant to complete four military-relevant tasks: (1) a UAV target identification task; (2) a UGV (unmanned ground vehicle) route planning task; (3) a communications task with an embedded verbal SA probe task; and (4) an ancillary task designed to assess SA using a probe detection method, a change detection task embedded within a situation map (see Figure 3). The UAV task simulated the arrival of electronic intelligence hits ("elints") from possible targets in the terrain over which the UAV flew. When a target was detected, it was displayed in the UAV view as a white square. Participants were told that when a target was presented, they were to zoom in on that location, which opened up a window of the UAV view. They were required to identify the target from a list of possible types (for which they had received prior training). Once identified, the target icon was then displayed on the situation map. At the same time as the UAV continued on its flight path

over the terrain, the UGV moved through the area following a series of pre-planned waypoints indicating areas of interest (AOI). During the mission, the UGV would stop at obstacles at various times and request help from the operator. A blocking obstacle required the operator to re-plan and re-route the UGV, whereas a traversable obstacle required that the participant resume the UGV along its pre-planned path. Participants received messages intermittently while performing the UAV and UGV tasks. The communications task involved messages presented both visually in a separate communications window (see Figure 3) and acoustically over a speaker. The messages requested updates on the UGV/UAV status and the location of particular targets (to assess SA, as described further below). Participants also had to monitor the communications for their own call sign (which they had been given previously). The communications task also included embedded messages designed to assess the participants' SA. The SA questions, which were divided into two types according to Endsley's (1995) taxonomy of SA—perception and comprehension—were presented during the mission and required a Yes or No response from the participant.

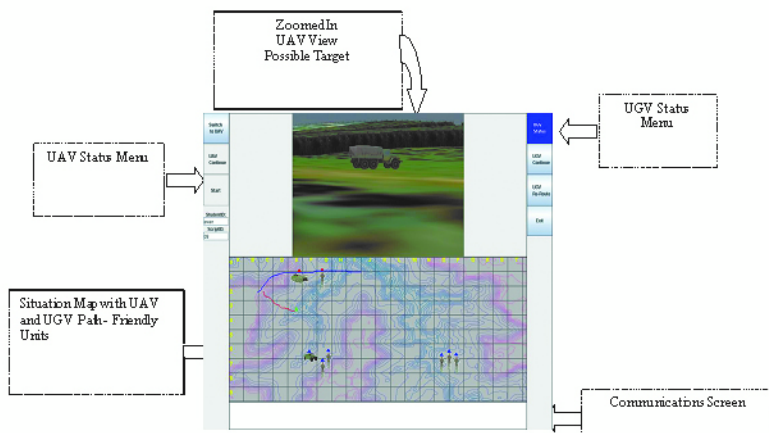


Figure 3. Robotic NCO simulation.

Parasuraman et al. (2007) examined the potential efficacy of adaptive automation based on real-time assessment of operator change

detection performance, on performance, situation awareness, and workload in supervising multiple UVs under two levels of communications task load. They used an adaptive automation invocation method first developed by Parasuraman et al. (1996) known as “performance-based” adaptation. In this method, individual operator performance (i.e., change detection performance) is assessed in real time and used as a basis to invoke automation. In contrast, in static or “model-based” automation, automation is invoked at a particular point in time during the mission based on the model prediction that operator performance is likely to be poor at that time (Parasuraman et al. 1992). This method is by definition not sensitive to individual differences in performance since it assumes that *all* operators are characterized by the model predictions. In performance-based adaptive automation, on the other hand, automation is invoked *if and only if* the performance of an individual operator is below a specified threshold at a particular point in time during the mission. If a particular operator does not meet the threshold, automation is invoked. However, if the threshold is exceeded in another operator, or in the same operator at a different point in the mission, the automation is not invoked. Thus, performance-based adaptation is by definition context-sensitive to an extent, whereas model-based automation is not.

Parasuraman et al. (2007) examined performance under three conditions: (1) manual; (2) static automation, in which participants were supported in the UAV task with an automated target recognition (ATR) system, thereby off-loading them of the responsibility of identifying targets; and (3) adaptive automation, in which the ATR automation was invoked if change detection accuracy was below a threshold (50 percent), but not otherwise. Each of these conditions was combined factorially with two levels of task load, as manipulated by variation in the difficulty of the communications task. The results were broken down by performance in the pre-automation invocation and post-invocation phases. In the pre-invocation phase before automation was implemented, there were no significant differences in change detection accuracy between conditions. However, there was a significant effect of automation condition in the

post-invocation phase. Participants detected significantly more icon changes in the situation map in the static and adaptive automation conditions compared to manual performance. Furthermore, change detection performance was better in the adaptive automation condition than in the static automation condition. These findings were mirrored in the analysis of the verbal SA questions that probed the operators' awareness of the relative threat imposed by enemy units. Finally, overall operator mental workload, as assessed subjectively (on a scale from 0–100) using a simplified version of the NASA-TLX scale (Hart and Staveland 1988), was lower in the adaptive and static automation conditions compared to the manual condition. These benefits of adaptive automation for change detection, SA, and workload are summarized in Figure 4.




Figure 4. Effects of static and adaptive automation on change detection accuracy, SA, and mental workload. From Parasuraman et al. (2007).

Adaptive or Adaptable Systems?

The results of the study by Parasuraman et al. (2007) indicate that both overall SA and mental workload were positively enhanced by adaptive automation support keyed to real-time assessment of change detection performance. Coupled with the previous findings using adaptive automation triggered by behavioral and physiological measures of workload, the empirical evidence to date supports the use of adaptive automation to enhance human-system performance in systems involving supervision of multiple UVs. However, before adaptive systems can be implemented using these methods, a critical issue must be addressed: Who is “in charge” of adaptation, the human or the system?

In adaptive systems, the decision to invoke automation or to return an automated task to the human operator is made by the system, using any of the previously described invocation methods. This immediately raises the issue of user acceptance. Many human operators, especially those who value their manual control expertise, may be unwilling to accede to the “authority” of a computer system that mandates when and what type of automation is or is not to be used. Apart from user acceptance, the issue of system unpredictability and its consequences for operator performance may also be a problem. It is possible that the automated systems that were designed to reduce workload may actually increase it. Billings and Woods (1994) cautioned that truly adaptive systems may be problematic because the system’s behavior may not be predictable to the user. To the extent that automation can hinder the operator’s situation awareness by taking him or her out of the loop, unpredictably invoked automation by an adaptive system may further impair the user’s situation awareness. (There is evidence to the contrary from several simulation studies, but whether this would also hold in practice is not clear.)

If the user explicitly invoked automation, then presumably system unpredictability will be lessened. But involving the human operator in making decisions about when and what to automate can increase

workload. Thus, there is a tradeoff between increased unpredictability versus increased workload in systems in which automation is invoked by the system or by the user, respectively (Miller and Parasuraman 2007). These alternatives can be characterized as “adaptive” and “adaptable” approaches to system design (Opperman 1984; see also Scerbo 2001). In either case, the human and machine systems adapt to various contexts, but in adaptive systems, automation determines and executes the necessary adaptations, whereas in adaptable systems, the operator is in charge of the desired adaptations. Although we have provided evidence for the efficacy of adaptive automation, it is important to keep in mind that adaptable automation may provide an alternative approach with its own benefits (see Miller and Parasuraman 2007). In adaptable systems, the human operator is involved in the decision of what to automate, similar to the role of a supervisor of a human team who delegates tasks to team members, but in this case, tasks are delegated to automation. The challenge for developing such adaptable automation systems is that the operator should be able to make decisions regarding the use of automation in a way that does not create such high workload that any potential benefits of delegation are lost. There is some preliminary evidence from studies of human supervision of multiple UVs that adaptable automation via operator delegation can yield system benefits (Parasuraman et al. 2005). However, much more work needs to be done to determine whether such benefits would still hold when the operator is faced with the additional demands of high workload and stress.

Architectures for Adaptive Automation in Command and Control

The question remains whether adaptive or adaptable automation would be useful in a specific environment representative of battlefield C2 systems. The research reported in this paper is part of a broader science and technology program aimed at understanding the performance requirements for human robot interactions (HRI) in future battlefields (Barnes et al. 2006). Some of the initial

results have shown that the primary tasks that soldiers are required to perform (and will continue to perform in the future) place severe limits on their ability to monitor and supervise even a *single* UV, let alone multiple UAVs and UGVs. The results of a number of modeling and simulation studies indicate that the soldier's primary tasks (e.g., radio communications, local security), robotic tasks (e.g., mission planning, interventional teleoperation), and crew safety could be compromised during high workload mission segments involving HRI (Chen and Joyner 2006; Mitchell and Henthorn 2005). Some variant of an adaptive system would therefore be particularly well suited to these situations because of the uneven workload and the requirement to maintain SA for the primary as well as the robotic tasks.

What adaptive automation architectures should be considered? Based on the empirical results discussed previously, we suggest that:

1. Information displays should adapt to the changing military environment. For example, information presentation format (e.g., text vs. graphics) can change depending on whether a soldier is seated in a C2 vehicle or is dismounted and using a tablet controller.
2. Software should be developed that allows the operator to allocate automation under specified conditions *before* the mission (as in the Rotorcraft Pilot's Associate).
3. At least initially, adaptive systems should be evaluated that do *not* take decision authority away from the operator. This can be accomplished in two ways:
 - (a) an advisory asking permission to invoke automation, or
 - (b) an advisory that alerts the operator that the automation will be invoked unless overridden.
4. For safety or crew protection situations, specific tactical or safety responses can be invoked without crew permission.

However, it will be important to evaluate options in a realistic environment before final designs are considered. For example, performance-based adaptive logic has proven to be effective in laboratory

simulations but may not be practical during actual missions because of the temporal sluggishness of performance measures (tens of seconds to many minutes). There is also the danger of adding yet another secondary task to the soldier's list of requirements (although it might be possible to assess performance on "embedded" secondary tasks that are already part of the soldier's task repertoire). Physiological measures could in principle be used to invoke adaptive automation rapidly (Byrne and Parasuraman 1996; Parasuraman 2003) because they have a higher bandwidth, but their suitability for rugged field operations has yet to be demonstrated reliably. Also, the type of tasks to automate will depend on engineering considerations as well as performance and operational considerations. Finally, there is the system engineering challenge of integrating different levels of automation into the overall software environment. Fortunately, there are simulation platforms that capture a more realistic level of fidelity for both the crew tasking environment and the software architecture making the design problem more tractable.

The future promises to be even more complex, with multiple, highly autonomous robots of varying capabilities being supervised by a much smaller team of human operators who will only be alerted when the situation becomes critical (Lewis et al. 2006). However, all the performance problems discussed—over-reliance, lack of situation awareness, skill degradation, and workload increases during crises—will, if anything, worsen. Automation architectures that give operators the ability to adapt to the current situation while engaging them in more than passive monitoring will become an increasingly important part of the HRI design process.

Given the complexity of the soldier's environment, the multitude of tasks that can be potentially automated, and the potential human-performance costs associated with some forms of automation, how do we decide not only how to implement the automation but also what tasks to automate? One approach we are pursuing is to develop an "automation matrix" that prioritizes proposed tasks for automation for UV operations (Gacy 2006). The matrix provides the framework for determining what is automated and the impact

of such automations. Soldier tasks define the first level of the matrix. The next level of the matrix addresses the potential automation approaches. The final component of the matrix provides weightings for factors such as task importance, task connectedness, and expected workload. A simple algorithm then combines the weights into a single number that represents the overall priority for development of that automation. From this prioritized list, various automation strategies (i.e., adaptive or adaptable) can be evaluated. We plan further analytical and empirical studies to investigate the utility of this approach to human-automation interaction design for UV operations in C2 systems.

References

- Adams, T.K. 2001. Future warfare and the decline of human decision-making. *Parameters* 2001–2002, Winter: 57–71.
- Barnes, M., and J. Grossman. 1985. *The intelligent assistant for electronic warfare systems*. (NWC TP 5885) China Lake: U.S. Naval Weapons Center.
- Barnes, M., R. Parasuraman, and K. Cosenzo. 2006. Adaptive automation for military robotic systems. In NATO Technical Report *RTO-TR-HFM-078 uninhabited military vehicles: Human factors issues in augmenting the force*. 420–440.
- Billings, C. 1997. *Aviation automation: The search for a human-centered approach*. Mahwah: Erlbaum.
- Billings, C.E., and D.D. Woods. 1994. Concerns about adaptive automation in aviation systems. In *Human performance in automated systems: Current research and trends*, ed. M. Mouloua and R. Parasuraman. 24–29. Hillsdale: Lawrence Erlbaum.
- Byrne, E.A., and R. Parasuraman. 1996. Psychophysiology and adaptive automation. *Biological Psychology* 42: 249–268.

- Chen, J.Y.C., and C.T. Joyner. 2006. *Concurrent performance of gunner's and robotic operator's tasks in a simulated mounted combat system environment* (ARL Technical Report ARL-TR-3815). Aberdeen Proving Ground: U.S. Army Research Laboratory.
<http://www.arl.army.mil/arlreports/2006/ARL-TR-3815.pdf>
- Crocoll, W.M., and B.G. Coury. 1990. Status or recommendation: Selecting the type of information for decision aiding. In *Proceedings of the Human Factors Society, 34th Annual Meeting*, Santa Monica, CA. 1524–1528.
- Cummings, M.L., and S. Guerlain. 2007. Developing operator capacity estimates for supervisory control of autonomous vehicles. *Human Factors* 49: 1–15.
- Department of the Army. 2007. *THAAD theatre high altitude area defense missile system. The website for defense industries - army*.
www.army-technology.com/projects/thaad/ (accessed Jan 26, 2007).
- Dornheim, M. 1999. Apache tests power of new cockpit tool. *Aviation Week and Space Technology* October 18: 46–49.
- Endsley, M.R. 1995. Measurement of situation awareness in dynamic systems. *Human Factors* 37: 65–84.
- Gacy, M. 2006. *HRI ATO automation matrix: Results and description*. (Technical Report). Boulder: Micro-Analysis and Design, Inc.
- Hancock, P., M. Chignell, and A. Lowenthal. 1985. An adaptive human-machine system. *Proceedings of the IEEE Conference on Systems, Man and Cybernetics* 15: 627–629.
- Hart S.G., and L.E. Staveland. 1988. Development of the NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human mental workload*, ed. P.A. Hancock and N. Meshkati. 139–183. Amsterdam: Elsevier Science/North Holland.

- Hilburn, B., P.G. Jorna, E.A. Byrne, and R. Parasuraman. 1997. The effect of adaptive air traffic control (ATC) decision aiding on controller mental workload. In *Human-automation interaction: Research and practice*, ed. M. Mouloua and J. Koonce. 84–91. Mahwah: Erlbaum Associates.
- Inagaki, T. 2003. Adaptive automation: Sharing and trading of control. In *Handbook of cognitive task design*, ed. E. Hollnagel. 46–89. Mahwah: Lawrence Erlbaum Publishers.
- Jones, P.M., D.C. Wilkins, R. Bargar, J. Sniezek, P. Asaro, N. Danner et al. 2000. CoRAVEN: Knowledge-based support for intelligent analysis. *Federated Laboratory 4th Annual Symposium: Advanced Displays & Interactive Displays Consortium*. 89–94. University Park: U.S. Army Research Laboratory.
- Kaber, D., and J. Riley. 1999. Adaptive automation of a dynamic control task based on workload assessment through a secondary monitoring task. In *Automation technology and human performance: Current research and trends*, ed. M. Scerbo and M. Mouloua. 129–133. Mahwah: Lawrence Erlbaum.
- Kaber, D.B., and M. Endsley. 2004. The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science* 5: 113–153.
- Kramer, A., and R. Parasuraman. In press. Neuroergonomics—application of neuroscience to human factors. In *Handbook of psychophysiology*, 2nd ed, ed. J. Caccioppo, L. Tassinary, and G. Berntson. New York: Cambridge University Press.
- Lewis, M., J. Wang, and P. Scerri. 2006. Teamwork coordination for realistically complex multi robot systems. In *Proceedings of the HFM-135 NATO RTO conference on uninhabited military vehicles: Augmenting the force*. 17.1–17.17. Biarritz, France.
- Miller, C., and R. Parasuraman. 2007. Designing for flexible interaction between humans and automation: Delegation interfaces for supervisory control. *Human Factors* 49: 57–75.

- Mitchell, D.K., and T. Henthorn. 2005. *Soldier workload analysis of the Mounted Combat System (MCS) platoon's use of unmanned assets*. (ARL-TR-3476), Army Research Laboratory, Aberdeen Proving Ground, MD.
- Moray, N., T. Inagaki, and M. Itoh. 2000. Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of Experimental Psychology: Applied* 6: 44–58.
- Murphy, R.R. 2000. *Introduction to AI robotics*. Cambridge: MIT Press.
- Opperman, R. 1994. *Adaptive user support*. Hillsdale: Erlbaum.
- Parasuraman, R. 2003. Neuroergonomics: Research and practice. *Theoretical Issues in Ergonomics Science* 4: 5–20.
- Parasuraman, R., and C. Miller. 2006. Delegation interfaces for human supervision of multiple unmanned vehicles: Theory, experiments, and practical applications. In *Human factors of remotely operated vehicles: Advances in human performance and cognitive engineering*, ed. N. Cooke, H. L. Pringle, H.K. Pedersen, and O. Connor. Volume 7. 251–266. Oxford: Elsevier.
- Parasuraman, R., and V.A. Riley. 1997. Humans and automation: Use, misuse, disuse, abuse. *Human Factors* 39: 230–253.
- Parasuraman, R., E. de Visser, and K. Cosenzo. 2007. *Adaptive automation for human supervision of multiple uninhabited vehicles: Effects on change detection, situation awareness, and workload*. Technical report.
- Parasuraman, R., M. Mouloua, and B. Hilburn. 1999. Adaptive aiding and adaptive task allocation enhance human-machine interaction. In *Automation technology and human performance: Current research and trends*, ed. M.W. Scerbo and M. Mouloua. 119–123. Mahwah: Erlbaum.
- Parasuraman, R., M. Mouloua, and R. Molloy. 1996. Effects of adaptive task allocation on monitoring of automated systems. *Human Factors* 38: 665–679.

- Parasuraman, R., R. Molloy, and I.L. Singh. 1993. Performance consequences of automation-induced “complacency.” *The International Journal of Aviation Psychology* 3: 1–23.
- Parasuraman, R., S. Galster, P. Squire, H. Furukawa, and C.A. Miller. 2005. A flexible delegation interface enhances system performance in human supervision of multiple autonomous robots: Empirical studies with RoboFlag. *IEEE Transactions on Systems, Man & Cybernetics – Part A: Systems and Humans* 35: 481–493.
- Parasuraman, R., T. Bahri, J. Deaton, J. Morrison, and M. Barnes. 1992. *Theory and design of adaptive automation in aviation systems*. (Progress Report No. NAWCADWAR-92033-60). Warminster: Naval Air Warfare Center.
- Parasuraman, R., T.B. Sheridan, and C.D. Wickens. 2000. A model of types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans* 30: 286–297.
- Rouse, W. 1988. Adaptive interfaces for human/computer control. *Human Factors* 30: 431–488.
- Rovira, E., K. McGarry, and R. Parasuraman. 2007. Effects of imperfect automation on decision making in a simulated command and control task. *Human Factors* 49: 76–87.
- Sarter, N., D. Woods, and C.E. Billings. 1997. Automation surprises. In *Handbook of human factors and ergonomics*, ed. G. Salvendy. 1926–1943. New York: Wiley.
- Scerbo, M. 2001. Adaptive automation. In *International encyclopedia of ergonomics and human factors*, ed. W. Karwowski. 1077–1079. London: Taylor & Francis.
- Sheridan, T.B. 2002. *Humans and automation: Systems design and research issues*. New York: Wiley.

- Sheridan, T.B., and W.L. Verplank. 1978. *Human and computer control of undersea teleoperators*. Cambridge: MIT, Man Machine Systems Laboratory.
- Simons, D.J., and M.S. Ambinder. 2005. Change blindness: Theory and consequences. *Current Directions in Psychological Science* 14(1): 44–48.
- Wickens, C.D., A. Mavor, R. Parasuraman, and J. McGee. 1998. *The future of air traffic control: Human operators and automation*. Washington: National Academy Press.
- Wickens, C.D., and S.R. Dixon. 2007. The benefits of imperfect diagnostic automation: A synthesis of the literature. *Theoretical Issues in Ergonomics Science* 8: 201–212.
- Wiener, E.L. 1988. Cockpit automation. In *Human factors in aviation*, ed. E.L. Wiener and D.C. Nagel. 433–461. San Diego: Academic Press.
- Wiener, E.L., and R.E. Curry. 1980. Flight deck automation: Promises and problems. *Ergonomics* 23: 995–1011.
- Wilson, G., and C.A. Russell. 2003. Real-time assessment of mental workload using psychophysiological measures and artificial neural networks. *Human Factors* 45: 635–644.