Human Control of Multiple Robots in the RoboFlag Simulation Environment*

Raja Parasuraman
Cognitive Science Laboratory
The Catholic University of America
Washington, DC, USA
parasuraman@cua.edu

Scott Galster
Human Effectiveness Directorate
Wright-Patterson Air Force Base
Dayton, OH, USA
scott.galster@wpafb.af.mil

Christopher Miller
Smart Information Flow Technologies
Minneapolis, MA, USA
cmiller@siftech.com

Abstract – *Human performance and supervisory control strategies were examined using the RoboFlag simulation environment. In an emulation of a multiple unmanned vehicle mission, a single operator supervised a team of six robots using automated modes or "plays" as well as manual control. A simplified form of a delegation type interface, Playbook, was used. Effects on user performance of two factors, opponent "posture" (offensive, defensive, or mixed) and environmental uncertainty (visual range of the robots: low, medium, or high), were examined in 18 participants who completed five mission trials in each of the nine combinations of these factors. Objective performance measures and subjective assessments of mental workload and situation awareness were obtained from each participant. Opponent posture had a significant effect on the percent of missions successfully completed and the duration of the games. Both opponent posture and visual range also significantly affected the use of manual and automated control strategies. Operator strategy selection and implementation are discussed with regard to performance data and the design of supervisory control interfaces that support flexible task delegation.*

Keywords: Robotics, automation, delegation, supervisory control, human-robot interface, playbook.

1 Introduction

Many military and civilian missions (e.g., reconnaissance, search and rescue) involve the use of a team of robots that have to be sent into territory that is hazardous or under the control of an adversary. Mobile robots and other unmanned vehicles play an important information-gathering role in these settings [10]. Given that completely autonomous operation is not currently technically feasible or may be undesirable in certain operational conditions, human supervision is necessary to allow for the management of unexpected events, and to ensure that mission goals are met [24]. Close attention must therefore be paid to the design of the human-robot interface, so as to allow for effective teaming and communication. The design of human-robot interfaces, which has traditionally been pursued from a purely engineering perspective, needs to be complemented by analysis and modeling of human performance [1,14].

Human supervision of robots can be seen as an extension of human control of automated agents (with variation in the level of autonomy and with the addition of perceptual and motor capabilities). Previous of human-automation coordination have revealed both benefits and costs of particular interfaces and designs [12,18]. Over-reliance, reduced situation awareness, mistrust, mode errors, loss of operator skill, and unbalanced mental workload are among the costs that have been found to be associated with particular styles of interaction between human and automation. These benefits and costs also differ depending on whether the automation supports information analysis or decision-making activities [13]. Systems that minimize human participation in higher-level decision-making processes by providing automated solutions can enhance overall system efficiency and reduce operator mental workload, but only if the automation is completely reliable, a requirement that can seldom be achieved in practice [4,7,16,17,22]. Furthermore, even if full automation of decision-making functions can be made more reliable, the required computational and engineering efforts may be considerable, with only diminishing returns in benefits obtained. This was illustrated in a recent study in which the tradeoff between the cost of additional automation and the gain in efficiency was examined using human telerobotic control of a simulated "sheepdog" herding a group of simulated "sheep" [23].

These considerations have led to the view, both in the human-automation and human-robot interaction research

* 0-7803-7952-7/03/$17.00 © 2003 IEEE.

communities, that the interface between humans and technology should be adjustable, adaptable, or adaptive [3,11,15,19]. Humans should be able to delegate tasks to automation at times of their own choosing, and receive feedback on their performance. Delegation in this sense is identical to that which occurs in successful human teams.

An interaction method that permits delegation should embody a real time approach to supervisory control [20]. Delegation architectures represent a particular type of mixed initiative interaction wherein the human sets an objective, provides more or less detailed instructions, and then "delegates" or authorizes the automation to determine the best method to proceed toward the goal within the delegation instructions. Delegation architectures should provide highly flexible methods for the human supervisor to declare goals and provide instructions and thereby choose how much or how little autonomy to delegate to automation on an instance-by-instance basis. An example of such an delegation architecture is the Playbook, which we have described elsewhere [8,9]—so named because it is based on the metaphor of a sports team's book of approved plays and the selection of these plays by the team leader (e.g., the quarterback in American football) and executed by the team members (the other players). The feasibility of this approach has been demonstrated previously in an application involving mission planning for unmanned combat air vehicles (UCAV) [8,9].

The Playbook uses a hierarchical task model to provide a common language for a human supervisor to communicate goals and intents and a Hierarchical Task Network planning system [6] to understand, reason over and either critique or complete partial plans provided by the human. A playbook permits the operator to "task" automation (such as robots or unmanned air vehicles):

- at a wide variety of functional levels of depth or granularity,
- via provision of goals, full or partial plans, and/or positive or negative constraints in any combination,
- by providing temporal, sequential or conditional constraints on task performance at varying levels of depth.

In the present study we examined the use of a highly simplified form of Playbook interface in a simulation study of human-robot teaming. We used the RoboFlag simulation environment [2,21,24]. The RoboFlag Playbook provides the ability to command simulated robots, individually or in groups, at two levels of granularity: via providing designated endpoints for robot travel or via commanding higher level behaviors (or modes or plays) such as "Patrol Border." The RoboFlag simulation was modified to emulate a typical unmanned vehicle (UV) mission involving a single operator managing a team of robots. The simulated mission goal was to send the robots from a home area into enemy territory, access and obtain a specified target, and return home as quickly as possible with minimum loss of assets.

A previous human performance study has demonstrated that performance within the RoboFlag simulation is sensitive to increased operator mental workload resulting from increased task demands from supervising greater numbers of robots [21]. This suggests that RoboFlag is an appropriate simulation environment for empirical investigations of the effects of different interface styles in supervising multiple robots. One of the postulated benefits of delegation type interfaces such as Playbook is that they allow for flexible use of automated tools in response to unexpected changes in task demands, without simultaneously increasing the operator's mental workload in using the automation [8,9]. Accordingly, we explored the use of the Playbook interface in a multiple UV simulation in which two sources of task demand were varied: (1) adversary "posture," in which the enemy engagement style was changed unpredictably between three types, offensive, defensive, or mixed; and (2) environmental uncertainty, as manipulated by variation in the effective visual range of each robotic vehicle under the control of the operator. We hypothesized that the use of the Playbook interface would allow users to respond effectively to unexpected changes in opponent posture and to increased uncertainty. We therefore assessed changes in how users tasked and supervised robots (using both manual control and the automation tools that were part of the Playbook interface) as a function of these two factors. We also examined the effects of these factors on overall mission efficiency (success rate and time to completion) and operator mental workload and situation awareness.

2 Methods

2.1 Participants

Eight males and 10 females between the ages of 18 and 33 (M = 23 years) served as paid participants.

2.2 Experimental Design

A within-subjects design was employed, with three robot Visual Range Conditions (Low, Medium, High) combined factorially with three Opponent Postures (Offensive, Defensive, Mixed), giving nine conditions in all. Visual Range was inversely related to uncertainty in the environment (e.g., low visual range = high uncertainty). Each participant completed five mission trials in each condition, for a total of 45 trials. Visual Range, represented by various ring sizes around each robot, was treated as a blocked factor with three levels, while Opponent Posture was randomized within each block. Participants filled out the NASA Task Load Index (TLX) and 3-D Situation Awareness Rating Technique (SART) questionnaires after each of the three blocks.

2.3 Apparatus and Procedure

The RoboFlag simulation ran on three separate PCs communicating under TCP/IP protocol. The human operator utilized one PC while another ran the opposing team script and the third PC displayed a central processing executive (the "Arbiter") and collected the data. The RoboFlag simulation was modified to allow a single operator (blue team) to compete against an opponent (red team) operating under scripted procedures that simulated different opponent postures. The field of engagement was divided into two halves, one for the blue team and one for the red team (see Figure 1). The operator supervised six robots and was required to send some or all of them into the opponent half of the field, capture the opponent's "flag" (located in at the center of a circular area in the opponent's half) and return safely to midfield, while simultaneously protecting their own flag (also located at the center of a circle in the home half).

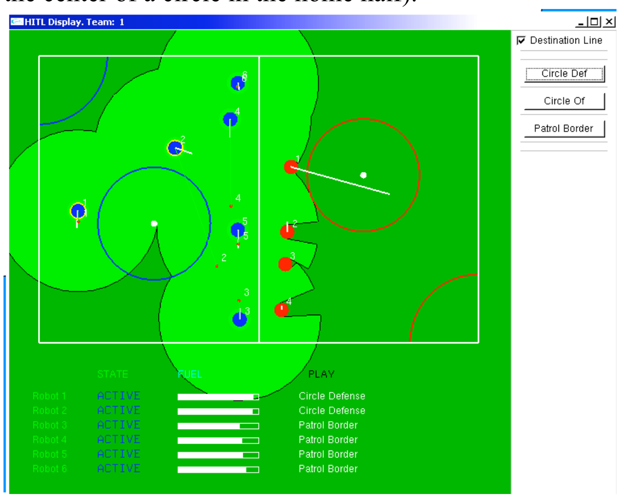

Figure 1. Human-robot interface for the Blue team.

The opponent or red team posture or stance was varied according to three available scripts: offensive, defensive, or mixed. In the offensive script (circle offense), all six red team robots attempted to capture the blue team's flag and return it to the mid-field line to win the game. In the defensive script, three red team robots defended the mid-field line (patrol border) and the other three circled around their own flag (circle defense) to prevent the blue team from penetrating the area and reaching their flag. The mixed script apportioned three red team robots to offense and three to defense (with none on patrol border). These three opponent postures were varied randomly and in an unpredictable (to the human operator of the blue team) way within each block of trials for a given visual range. The human operator had control of all six blue team robots and could assign any number of them to one of the automated plays available in the Playbook (circle offense, circle defense, patrol border) or could control them manually by giving the robot an endpoint to move to (see Figure 1).

Participants were trained by showing them the field of play, how to select and move robots, and utilize the Playbook plays. They were instructed that the only way a red team robot could be seen was if they entered into the visual range of the blue team robot, otherwise the red team robot was invisible to the blue team operator. They were shown the objective of capturing the opponent flag and crossing back into their own territory. Each participant completed a training trial for each of the nine experimental conditions and was informed in advance of the experimental condition for each trial. During data collection trials, the participant was not informed of the experimental condition for each trial. The objective of each trial was the same: cross into the opponent area with one or more robots, capture the opponent flag, and cross the mid-field line while concurrently defending their own flag from capture by the opposing team.

3 Results

3.1 Overall Performance

The performance data were submitted to a 3 (low, medium, high) visual range X 3 (offense, defense, mixed) Opponent Posture Analysis of Variance (ANOVA). The overall performance metrics included the percentage of the games that were won (mission success rate) and the time elapsed for each game (mission completion time). The results of the ANOVA indicated that there was a significant effect of Opponent Posture on the percentage of the games won, $F(2,34) = 50.89$, $p < .01$. Expectedly, the participants won 100% of the games when the opponent stance was purely defensive and no moves were made to capture the blue team flag. Excluding the defensive opponent posture condition, participants won significantly more games, $t(34) = 3.51$, $p < .05$, when they played against the red team on offense ($M = 75.56\%$, $SE = 2.6\%$) than when that team split the robots in the mixed condition ($M = 62.22\%$, $SE = 3.0\%$). Neither the effect of Visual Range nor the interaction between Visual Range and Opponent Posture were statistically significant.

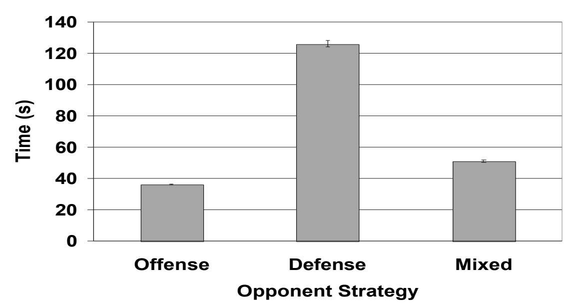

Figure 2. Averaged elapsed game times as a function of Opponent Posture.

A similar 3 X 3 ANOVA was conducted for the duration of each game (time for mission completion, whether successful or not). Game times were significantly shorter when participants played against the offensive stance ($M = 36.48s$, $SE = 0.47s$) compared to when they played against the mixed posture ($M = 51.29s$, $SE = 1.63s$), $F(2,34) = 177.92$, $p < .01$. The longest games were those where the participants played against the defensive stance ($M = 126.14s$, $SE = 4.33s$), which was also the condition when a 100% mission success rate was achieved, as described previously. These significant differences, $F(2,34) = 177.92$, $p < .01$, are illustrated in Figure 2. Thus, while participants won every game when playing against the defensive opponent posture they also took over twice as long to win the game in that condition. Conversely, participants won about 75% of the games when they played against the offensive strategy; however these games were of the shortest duration. None of the other experimental factors in the ANOVA of game duration times were statistically significant.

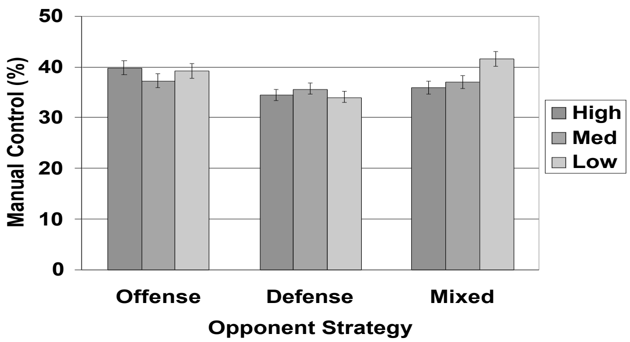

Figure 3. Percent of time robots were under manual control.

3.2 Strategy Usage

There were nine different states a robot could transition to: Inactive, Unassigned, Circle Defense, Circle Offense, Patrol Border, Manual Control, Tagged, Flagged Circle Offense, and Flagged Manual. We examined differences in the manual and automated plays used, and subsequently, what automated plays were used under which experimental conditions. The percent of time that the robots were under manual control was submitted to a 3 X 3 ANOVA. There was a significant interaction between Visual Range and the Opponent Posture, $F(4,68) = 2.96$, $p < .05$. The results, shown in Figure 3, indicate that manual control was used less frequently when the participant was playing against a defensive posture and more frequently under the offense and mixed conditions. Furthermore, Visual Range mediated the use of manual control. When the Visual Range was low, indicating a high level of uncertainty, manual control was used more often in the mixed condition. This was not the case in the offensive or defensive conditions.

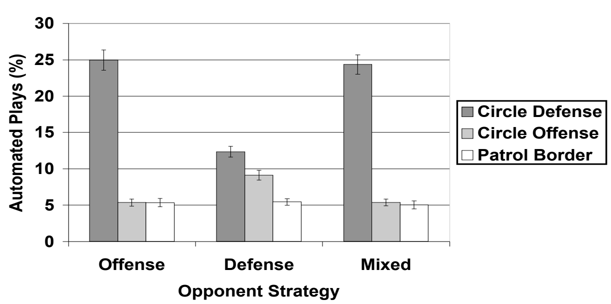

Figure 4. Percent of time robots were instructed to execute an automated play.

We also analyzed the amount of time the robots were directed to perform an automated play. All plays were grouped so that the percent of time the robots were functioning in an automated state could be determined. This percent was submitted to a 3 X 3 ANOVA. There was a significant effect of Opponent Posture, $F(2,34) = 11.34$, $p < .05$. Plays were used most often and equally when the opponent stance was offensive or mixed, but less often with a defensive opponent.

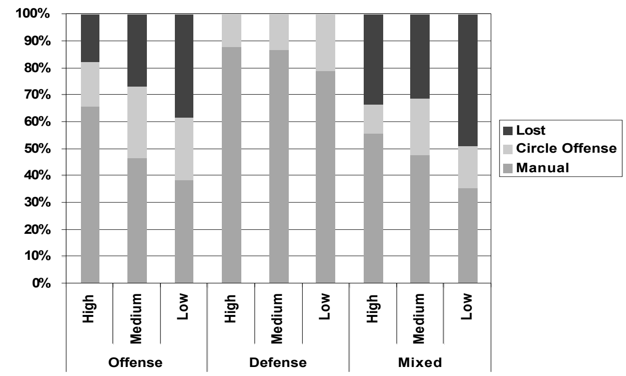

Figure 5. Games lost and strategy used by robot that won.

To better understand the type of automation used, the automated plays (Circle Defense, Circle Offense, Patrol Border) were included as an additional factor in the statistical analysis yielding a 3 X 3 X 3 ANOVA. There was a significant interaction between the Opponent Posture and the type of automation used, $F(4, 68) = 8.30$, $p < .05$. Figure 4 illustrates that the Circle Defense automated play was the most utilized play compared to the other available plays regardless of the opponent posture encountered. When the participants played against the offense or mixed condition they used the Circle Defense automated play far more often than the Circle Offense or Patrol Border automated plays. While the proportion of automated plays was similar in the Offensive and Mixed conditions, the proportion was altered in the Defensive condition. Participants utilized the Circle Defense play less often when they played against the Defense Opponent Strategy and increased the use of the Circle Offense play.

Another aspect of the use of automation or manual control concerns the capturing of the opponent's flag. Participants could either use the Circle Offense automated play or direct the robots manually to capture the flag and bring it to the mid-line to win the game. An alternative method to ending the game was for the opponent team to cross the mid-line with the participant's flag. Figure 5 shows the three possible outcomes for ending a game. There was a potential non-independence assumption violation in the factors so a statistical analysis could not be performed with these data. Rather, Figure 5 provides a profile of the percent of games won and lost as well as what strategy was employed for the robot that won each game under the nine possible experimental conditions.

3.3 Subjective Measures

The NASA TLX and 3-D SART were administered after each block (visual range) of 15 trials. Overall subjective mental workload was computed by averaging the six NASA-TLX sub-scales and submitted to a 3-way ANOVA. There was a significant difference in subjective mental workload, F(2,34) = 8.77, p < .001, across the low, medium, and high visual range conditions, with reported mental workload increasing as the visual range was reduced. There was no significant difference between the subjective assessments of overall situation awareness across the visual range conditions.

4 Discussion

Evaluations of different interfaces for human control of multiple robots can be informed by human-in-the-loop studies in which objective measures are obtained of human performance and strategy use [1,14]. Such studies are still relatively rare [5,21]. The present study using RoboFlag examined how human operators use automated tools to supervise multiple robots in response to changing task demands imposed by unpredictable changes in opponent "posture" (offensive, defensive, or mixed) and environmental uncertainty (robot visual range).

There were several results of interest. First, the results showed that the multi-level tasking of the simplified Playbook interface allowed effective user supervision of robots, as evidenced by the number of missions successfully completed (percent games won) and the time for mission execution. As expected, significantly fewer games were won when the opponent posture was mixed rather than entirely offensive. Nevertheless, users still won a moderately high proportion of games (about 62%) and in a relatively short time (about 51 seconds) in the mixed posture condition. These findings suggest, but do not prove, that the Playbook interface, as a simple example of a delegation interface, allowed users to respond effectively to unexpected changes in opponent posture by tasking robots appropriately. Further confirmation of this view requires studies in which more complex versions of the Playbook interface are evaluated.

Adaptable interfaces (such as Playbook) are also posited to allow for regulation of mental workload and maintenance of situation awareness [8,9,11]. With respect to mental workload, there was an expected effect of visual range. As the robot vision radius was reduced, environmental uncertainty increased and users reported greater overall workload. However, the experimental design we used did not allow us to examine the effects of opponent posture on workload. Presumably users found it more difficult to complete their mission when the opponent stance was mixed rather than purely offensive. It would be interesting to examine in a future study whether the Playbook interface would "balance out" the variations in user mental workload in response to changes in task demands in supervising multiple robots, as has been reported in studies with other adaptive interfaces [11].

A basic goal of this study was to examine user strategies in using either manual control or automated plays in carrying out the mission. Users had three automated plays they could use, Circle Offense, Circle Defense, and Patrol Border. Of the several results of interest, the most general was that manual control was used less frequently against a defensive opponent posture than when opponent stance was offensive or mixed. In addition, when environmental uncertainty was high due to a low robot visual range, manual control was used more often in the mixed condition. This was not the case with the offensive or defensive postures. The combination of low visual range and a mixed opponent posture represented the most challenging condition to the user. The greater use of manual control may reflect a greater perceived need for redirecting robots from previously assigned plays in this case, possibly due to reduced trust in the automation to achieve the mission goal.

Of the three plays available to the user, Circle Defense was used most often, particularly in the offensive and mixed conditions, when Circle Defense was activated far more often than the Circle Offense or Patrol Border automated plays. When the opponent stance was purely defensive, however, participants used the Circle Defense play less often and increased the use of the Circle Offense play. Both these patterns of Playbook usage seem appropriate given the task conditions. However, it will be of interest to examine whether the relative *proportion* of time a given automated play was used was appropriate or optimal (according to some specified criteria) for a given condition. We plan on conducting a Markov modeling analysis to examine this issue.

The present results provide a preliminary empirical evaluation of the use of delegation type interfaces [8,9] for human supervision of multiple robots. It should be noted however, that we used a simplified form of Playbook without the full functionality and sequential command ability that such an interface provides. Nevertheless, while our goal was not to compare Playbook to other interfaces (or to either manual control or full automation), the results nevertheless are suggestive of the effectiveness of this style of interface in allowing users flexible use of automation in response to changing task demands.

In situations where the robot to human operator ratio is high, as in the present work, low-level control of all robots becomes increasingly difficulty, mandating the use of automation at this level. At the same time, limitations in the reasoning capabilities of semi-autonomous agents and the brittleness of some automated behaviors necessitate intervention through human intelligence. These requirements indicate that the human-robot interface must support multiple levels of interaction [24]. We propose that delegation type interfaces such as Playbook allow for such multi-level interaction in a flexible manner that keeps human workload within a manageable range.

Finally, the results show that the RoboFlag simulation environment provides a viable platform for empirical evaluations of operator strategies in controlling multiple robots. We plan additional studies to extend the results obtained here and to examine other issues in human operator supervision of multiple robots.

Acknowledgements

This work was supported under the DARPA MICA Program through Contract # F30602-01-2-0577 (Sharon Heise, Program Monitor, Carl DeFranco Contract

Monitor). We thank Mark Campbell for the RoboFlag software and for valuable discussions of this research.

References

[1] J. Adams, "Critical considerations for human-robot interface development," in *Human-Robot Interaction, 2002 AAAI Fall Symposium,* AAAI Press, Menlo Park, CA, pp. 1-8, Nov. 2002.

[2] R. D'Andrea, M. Babish, "The *RoboFlag testbed,"* in *Proceedings of the American Controls Conference*, Denver, CO, June 2003.

[3] J. Crandall, M. A. Goodrich, "Principles of adjustable interactions," in *Human-Robot Interaction, 2002 AAAI Fall Symposium,* AAAI Press, Menlo Park, CA, pp. 29-38, Nov. 2002.

[4] W. M. Crocoll, B. G. Coury, "Status or recommendation: Selecting the type of information for decision aiding," in *Proceedings of the 34th Annual Meeting of the Human Factors and Ergonomics Society,* Santa Monica, CA, pp. 1524-1528, Oct. 1990.

[5] S. R. Dixon, C. D. Wickens, "Control of multiple UAVs: A workload analysis," in *Proceedings of the 12th International Symposium on Aviation Psychology. Dayton, OH*, April 2003.

[6] K. Erol, J. Hendler, D. Nau, "UMCP: A sound and complete procedure for hierarchical task network planning," in *AI Planning Systems: Proceedings of the 2nd International Conference,* K. Hammond, Ed., Los Altos, CA, pp. 249-254, 1994.

[7] S. M. Galster, R. S. Bolia, M. M. Roe, R. Parasuraman, "Effects of automated cueing on decision implementation in a visual search task," in *Proceedings of the Annual Meeting of the Human Factors and Ergonomics Society*, Minneapolis, MN, Oct. 2001.

[8] C. Miller, R. Parasuraman, *Designing for flexible human-automation interaction: Playbooks for supervisory control,* Technical Report, SIFT, MN, 2002.

[9] C. Miller, M. Pelican, R. Goldman, "'Tasking' interfaces for flexible interaction with automation: Keeping the operator in control," in *Proceedings of the Conference on Human Interaction with Complex Systems*. Urbana-Champaign, IL, May 2000.

[10] National Research Council, *Technology Development for Army Unmanned Ground Vehicles.* National Academy Press, Washington DC, 2002.

[11] R. Parasuraman, "Effects of adaptive function allocation on human performance," in *Human Factors and Advanced Aviation Technologies,* D. J. Garland, J. A. Wise, Eds., ERAU, Daytona Beach, FL, 1993.

[12] R. Parasuraman, V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, Vol. 39, pp. 230-253, 1997.

[13] R. Parasuraman, T. B. Sheridan, C. D. Wickens, "A model for types and levels of human interaction with automation," *IEEE Transactions on Systems, Man, and Cybernetics. Part A,* Vol. 30, pp. 286-297, 2000.

[14] E. Rogers, E., R. Murphy, *Human Robot Interaction*. DARPA/NSF, Washington DC, 2001.

[15] W. B. Rouse, "Adaptive aiding for human/computer control," *Human Factors*, Vol. 30, pp. 431-443, 1988.

[16] E. Rovira, K. McGarry, R. Parasuraman, "Effects of unreliable automation on decision making in command and control," in *Proceedings of the Annual Meeting of the Human Factors Society,* Baltimore, MD, Oct. 2002.

[17] N. Sarter, B. K. Schroeder, "Supporting decision-making and action selection under time pressure and uncertainty: The case of in-flight icing," *Human Factors*, Vol. 43, pp. 573-583, 2001.

[18] N. B. Sarter, D. Woods, C. Billings, "Automation surprises," in *Handbook of Human Factors and Ergonomics,* G. Salvendy, Ed., Wiley, New York, 1997.

[19] M. Scerbo, "Adaptive automation," in *International Encyclopedia of Ergonomics and Human Factors,* W. Karwowski, Ed., Taylor, New York, 2001.

[20] T. Sheridan, "Supervisory control," in *Handbook of Human Factors*. G. Salvendy, Ed., Wiley, New York, NY, pp. 1244-1268, 1987.

[21] J. Ververka, M. Campbell, "Experimental study of information load on operators in semi-autonomous systems," in *Proceedings of the AIAA Guidance, Navigation and Control Conference,* Austin, TX, 2003.

[22] C. D. Wickens, *Imperfect and unreliable automation and its implications for attention allocation, information access and situation awareness.* Technical Report ARL-00-10/NASA-00-2, University of Illinois, IL, 2000.

[23] M. Wilson, M. Neal, "Diminishing returns of engineering effort in telerobotic systems. *IEEE Transactions on Systems, Man and Cybernetics-Part A: Systems and Humans*, Vol. 31, pp. 459-465, 2001.

[24] Z. Zigoris, J. Siu, A. Hayes, "Balancing automated behavior and human control in multi-agent systems: a case study in RoboFlag," in *Proceedings of the American Control Conference*, Denver, CO, June 2003.