Simulation-based Agent Support in a
Synthetic Team-based C2 Task Environment

Linda R. Elliott
Veridian Engineering
2504 Gillingham Drive, Bldg. 170
Brooks AFB, TX 78235

Scott Chaiken
AFRL Research Laboratory
2504 Gillingham Drive, Bldg. 170
Brooks AFB, TX 78235

Mathieu Dalrymple
Veridian Engineering
2504 Gillingham Drive, Bldg. 170
Brooks AFB, TX 78235

Plamen Petrov
21st Century Systems, Inc.
Omaha, NE

Alexander Stoyen
Department of Computer Science
University of Nebraska at Omaha

Abstract

In our long-term program of research in command and control (C2) teamwork and performance, we have extensively analyzed the roles, responsibilities, and interdependencies of Airborne Warning and Control System (AWACS) Weapons Director (WD) teams, using a variety of methods. The AWACS WD team serves as a vital airborne C2 node, providing airborne surveillance and command, control, and communications functions for tactical and air defense forces. They detect, identify, track, and intercept airborne threats. Our investigations seek to identify tools and techniques to facilitate performance in this complex and dynamic domain. In this paper, we describe progress toward an agent-based decision support system, the AWACS WD Intelligent Agent Assistant (IAA). The WD-IAA will facilitate decision making for decision events which are both typical and time consuming. This paper will describe approach, methodology, and potential application areas for agent-technologies in C2 training.

Introduction

Over the past several years, the overarching theme to the program of research has focused on command and control teamwork and performance, particularly that of AWACS WD teams. Our C3STARS facility (Command, Control, and Communications Simulation, Training and Research System) was configured to represent the roles, responsibilities, and task demands of AWACS WD personnel. First, we focused on conceptualization, description, and measurement of individual performance (Klinger et al., 1993). At the same time, we initiated investigations related to team performance and dynamic environments (Dalrymple, Eddy, & Schiflett, 1995; Dalrymple & Schiflett, 1997; Eddy, 1989; Eddy, et al., 1992; Elliott, Serfaty, & Schiflett, 1998; Elliott, Schiflett, & Dalrymple, 1998; Elliott, Dalrymple, & Neville, 1997; Elliott, et al., 1997; Fahey et al., in review; MacMillan, et al., 1998; Schiflett & Elliott, in review; Schiflett, et al., 1990). Finally, we are investigating various methods to model performance (Cobb et. al.., 1998; Coovert, et al., 1999; Elliott, et al., in review; Elliott, et al., 1999) and to develop tools to enhance performance (Hess, et al., 1999; Stoyen, et al., 1999). In this paper, we will describe software developed as a decision aid and training tool for AWACS WD teams, the AWACS-WD-IAA system, developed by 21st Century Systems, Inc. We begin first by describing the nature of AWACS WD team task environment.

AWACS WD Task Demands

Systematic descriptions of AWACS roles, responsibilities, requirements, interdependencies, tactics, strategies, and task demands were collected from subject matter experts, cognitive task analyses (Fahey et al., 1998; MacMillan et al., 1998; and focal-group interviews (Elliott et al., 1998; Elliott et al., in review). These data were examined to identify decision events, which were generic to performance, regardless of mission scenario, and likely to bottleneck under high tempo situations.

It is clear that AWACS WD duties exemplify core characteristics of a C2 team. They work aboard the E-3 Sentry, a very sophisticated airborne command post, equipped with state-of-the art electronics inside and topped with a 30-foot rotating radome. AWACS WDs perform in highly interdependent roles, tracking and coordinating some type of tactical action, in a manner consistent with overall strategic goals and procedures, for a defined sector of air and/or land space over a sustained period of time (Elliott et al., 1998). The AWACS team provides airborne surveillance and command, control, and communications functions for tactical and air defense forces. They detect, identify, track, and intercept airborne threats, and coordinate communications received from a number of sources, such as other WDs, the SD, air surveillance personnel, electronic combat officers (ECOs), intelligence operations, base operations, and friendly pilots. To accomplish this, they exchange, interpret and effectively weight information and optimize resource allocation decisions across team members, over time, and quite often under conditions of stress and fatigue.

Their decisions pertain to the assignment of aircraft to attack hostile and/or identify unknown entities. This also requires the maintenance of assets, through tactics that will

protect and conserve resources, such as combat air patrol and refueling. It also requires the coordination and/or sharing of resources, such as surface-to-air missile sites and various combat, reconnaissance, refueling, and search and rescue aircraft. Relevant information must be distributed to appropriate personnel and updated over time, in dynamic conditions which may require shift changes in personnel. Information is often verbal, and may be missing, degraded, passed along from unfamiliar sources, or misinterpreted by others. In addition, information is often communicated/interpreted by individuals with only partially overlapping awareness of the battlespace.

The AWACS-WD-IAA System

 The initial step in building the WD-IAA was to develop a PC-based team task environment that captures the core characteristics, roles, responsibilities, and interdependencies of AWACS WD personnel.   Many criteria were generated for the evaluation of this software (Schiflett & Elliott, in review). In this effort, a synthetic task environment (STE), based on a Java-based federation of agent technology, was developed to reflect the following features: (a) core characteristics of the AWACS WD performance roles and responsibilities, (b) team-based task, where roles may be played by "humans" or by "agents", (c) agent technology providing visual representations of recommendations, (d)   user-friendly scenario generator, and (e)   automatic / programmable performance data collection. Agent technology provides exciting opportunities for enhanced scenario generation and data collection. Also, visually-driven decision support capabilities are embedded within the AWACS STE environment.

The software is a distributed, simulated, real-time team environment comprised of air, sea, and ground assets in a combat environment (Chiara & Stoyen, 1997; Stoyen et al., 1999), primarily based on the roles and responsibilities of AWACS WD team members. The environment can contain up to two hundred physical entities (planes, ships, SAM missiles, etc.) operating with realistic yet non-classified performance characteristics in an interactive environment in which real-time decision support is available to each WD.

Visual display enhancements include 2-D representations of data and entities, time series charts and histograms, data tables, symbology, and color-coding. Much of the display formatting is end user controllable. For example, the system provides weapons directors with control over how entities are displayed onscreen. For instance, they can highlight their own resources to set them off from all the friendly resources on the battlefield. This can greatly reduce demands on attentional resources. Entities can be represented through standard symbology, as text labels, or both. Entities can also have detailed information windows temporarily created to inspect the state of an owned asset (assets that are not owned have impoverished best-guess information in their information tables).

Agent Technology

Intelligent agent technology is increasingly applied to simulation and training technology (Brenner, Zarnekow, & Wittig, 1998; Hendeler, 1996; Huhns, & Singh, 1998). "Agent" is a broadly defined term with perhaps three dominant functions in this system. The first involves the simulation of scenario entities, referred to as constructed forces. Agent

technology defines the operating characteristics and behavior of hostile and friendly entities (e.g. speed, radar range, and weapons range of aircraft). Agent specification is complete and detailed to the point where the entire scenario can be played out through a federation of numerous "agents" (i.e. a simulation with no live players).

Agent technology has also been used to simulate other WD team members. The distinction between "agent as constructive forces" and "agent as player" is largely one of degree. However, the latter sort of agent is typically far more complex, as these agents are designed to simulate another simulation player, not just a battle entity. Such agents give users the option to play with other "live" participants or participate alone, with the simulation acting out other roles in a realistic fashion. This sort of agent also defines the pedagogical goal of the simulation, in the sense that these agents can be used to implement (e.g. set policy for) optimal performance. They can also be used to demonstrate the results of flawed performance. "Player" agents are an extremely useful and yet rare capability for team task simulations. However, their development is expected to increase given the great utility of allowing individual training within a team-like context. Both entity and player agent technology are equally important to our effort. The former sense is what gives the AWACS-WD-IAA its fidelity to the real task; the latter sense provides both the model of normative behavior the user should strive for and the means (algorithms/knowledge) to effectively accomplish that behavior.

The third manner in which agent technology was utilized is for decision support. This sense of agent is not so much a simulation of a player (or simulation of an entity) but a simulation of a "coach" or "adviser". This sense of agent can be very broad, and the distinction between such agents and operational interfaces can be blurry. For instance, such agents can be imbued with the ability to seek out information over distributed networks; search through information databases; manipulate information through filtering, transforming, aggregation, and fusing of multiple, independent information steams; and to report information to the human requester. There may also be multiple agents working on several tasks at any point of time, e.g., several agents monitoring and filtering information from disparate channels, agents to aggregate and fuse relevant information, agents to select an appropriate visualization of the data to report, and so on. Some agents may be imbued with a high level of autonomy, allowing them to make critical decisions based on information found without human intervention or guidance.

For this effort, the subject controls the autonomy of the decision aide agent. The subject can ignore recommendations, accept individual recommendations, accept a group of recommendations, or allow the aide to make all decisions. If the subject tasks the aide to make all decisions, the scenario is effectively being run independently of any human intervention. This allows (a) assessment of reliability of recommendations, (b) assessment of effects of uncertainty in a dynamic environment, and (c) investigation of "what-if" scenarios, where algorithms underlying recommendations are manipulated.

Agent Implementation

Agent-modeled representations of other team members were designed from SME input for how team members would react under typical theatre operations. So for each entity or

"player" in the simulation, there is a rule base that describes its typical battle behavior and their behavioral triggers. For example, friendly aircraft will fire at a target only when it is in range of whatever ordinance it has left. Other factors are also considered, such as possible approach maneuvers. The probability of being shot down during an intercept approach will be higher given a head-on than given a rear approach. Some of this behavior can be quite specific and based on stratagems that have been useful in the past. For instance, a fighter that is "winchester" or expended all its air-to-air ordinance, can bluff a hostile bomber into diverting from its mission. In addition, High Value Airborne Assets (i.e. Rivet Joint, J-STARS, AWACS, etc.), can operate in routes as bait to lure enemy fighters within the range of friendly SAM sites.

The decision aide agent had a similar origin. That is, SMEs described the typical thinking process and heuristics that WDs use and are taught as part of their training. For example, the primary decision event targeted for assistance is that of allocating assets (friendly aircraft) to hostile or unknown entities. This is an event that can take considerable effort, when workload and task tempo increases. Information from various sources must be considered, such as threat identification, location, distance, fuel requirements, and rules of engagement. The decision aid emulates the manner in which WDs make their decisions, and provides the recommendation (or set of recommendations) immediately upon consultation. While the decision algorithm is not an optimal model, it is expected to provide reasonable options in a very short period of time.

The aide captures basic rules by which a WD makes decisions. When a WD considers which target a specific fighter should go after, he/she prioritizes the near targets first (based on the target's closure rate), the high altitude targets second, and the targets to the right of the fighter last (in case of a tie between targets on the other attributes). Exceptions to this prioritization are also considered (i.e. if it's a jammer, one may need to go after it even if there are nearer targets).

The decision aide is not based on a optimizing decision model. Instead, it is more similar to that of human decision-making, and is similar to preferential choice model such as satisficing and the lexicographic rule (Stevenson, Busemeyer, & Naylor, 1990; Simon, 1955). Satisficing strategies simplify a complex problem space by (a) simplifying thresholds on each dimension (such as categorizing speed as threatening or nonthreatening) and (b) reducing the number of dimensions necessary to make an assessment and (c) stopping the information search as soon as a satisfactory solution is generated. The lexicographic rule is similar, but prioritizes dimensions such that all options are first judged on the most important dimension, then the second, and so on. The decision aide also has simplifying rules and guidelines pertaining to assignment and allocation of assets. For instance, in scheduling the intercepts for many friendlies against many hostiles, the agent will focus in on the fighter which is furthest out (i.e. closest to the enemy) and generate an assignment for it. Once that is accomplished, it moves to the next farthest friendly fighter, and so on.

The aide is, and should be, configurable. First, it needs to be configurable in order to maintain effectiveness. The manner in which decisions are made can and will change according to the particular mission scenario and rules of engagement. Thus, certain actions, such as penetration of hostile territory, may be acceptable in one scenario and not in another. The software allows variations in its rule structure. It also allows changes with regard to the "point value" of various assets. This also enables fine-tuned research in decision process, as the agent can be tailored to be more or less risk-taking (when information is uncertain), have directional bias (more or less "aggressive", "passive" in threat assessment or rules of engagement), or bias in central tendency (decisions are always "moderate"). In addition the probabilistic nature of the environment can be manipulated by specifying the probability that the decision made will be successfully executed. When that probability is very high, the environment is deterministic and very reliable. When probabilities are lowered, scenario events will unfold in different ways, each time the scenario is run.

Configurable decision algorithms enable in-depth descriptive and prescriptive investigations of decision process. Particular heuristics, biases, and models of decision choice can be predicted and compared to algorithm function. Results from descriptive investigations can inform refinement of the decision tool. In turn the algorithm can be modified to reflect a particular decision model, and compared to other models with regard to the degree to which either model accounts for performance data. For example, threat assessment decision events have been shown to be sensitive to order effects, in that information presented first or last (depending on tempo) is given more weight, even when other information is more important. The algorithm can be adjusted to reflect this tendency, and results compared to actual data. Other facets relevant to the decision process can also be investigated with this approach, such as risk-taking, aggressiveness, and information uncertainty.

As might be expected, this approach to agent implementation has to mature over time. There are also some limitations to the agent implementation, which is to say, there are actions that the human WD can perform that are not within the capability of the agent to recommend. For instance, redirecting a fighter to a different target than the one it was originally assigned to is not currently implemented. Redirection is an action of extreme necessity, because such redirections both weaken the WD's credibility and put the redirected fighter at greater risk. Important nuances have to be weighed in order to know when a redirection is necessary. WDs must be trained as to strengths and weaknesses of the decision aide. They will sometimes need to go against recommendation in some circumstances. Scenarios will have to be developed to teach both the strengths and weaknesses of the agent.

Despite agent limitations, we expect this aide to enhance performance in complex high-tempo scenarios, where multiple decisions must be made within a short time frame. The aide has the advantage of complete reliability regardless of stress, sustained operations, or consequences of failure. It will never forget to refuel a plane because of a tense situation occurring at a different part of the scope. If supplied with the appropriate procedures, it will recommend the appropriate actions. It will also accomplish routine follow-up

actions, necessary to optimize theatre operations. For instance, once an intercept has been scheduled, it is useful to notify other friendly battle entities that the intercept is planned. Also, others should be notified if the intercept fails. These sorts of operations can be handled by the agent without bothering the human operator at all..

Validation

Validation of scenario requirements was accomplished through examination of cognitive task analysis data, training requirements, and focal group interview results. From these data, core characteristics of the task environment were identified, and task roles, responsibilities, and interdependencies clarified. This led to development of an array of scenarios varying in complexity and difficulty, populated by agent-modeled constructive forces.   SMEs also identified the types of events most likely to determine performance. These events included high-tempo stretches of enemy activity in which many simultaneous intercept decisions had to be made and keeping track of refueling needs.

Validation of decision recommendations is currently underway.   First, decision algorithms were assessed for reliability. For each scenario, data was collected through several trials in which all simulator players were agents, and all the agents were programmed to "automatically" accept all agent recommendations.   We set the probability of intercept outcomes to be deterministic (i.e. repeatable with exactly the same outcomes from trial to trial) and have verified identical performance scores at 5-minute intervals, for all roles, for almost all scenarios. (The suspect cases where this didn't occur could have been due to differences in how the simulation was run, e.g. time the simulation was cut off, how often simulation statistics were collected.)

Next, the algorithms that form the basis of the coaching agent will be reviewed by experienced AWACS WDs or their instructors. The point-values assigned to various assets will also be reviewed, as these are basic to both decision algorithms and performance outcome scores. In addition, characteristics and attributes of the scenarios and platform capabilities will be discussed with the operational personnel. Finally, an empirical test of the agent utility will be made. What follows is our basic paradigm for assessing the worth of an agent.

Experimental Design

Our experiment validates agent recommendations by comparing the performance of novice and expert groups with and without the agent. The agent was modeled to represent expert decision-making and is a heuristic, not an optimal model.   Therefore recommendations may not benefit expert WDs, except perhaps under conditions of high workload, time pressure, and stress. The agent is expected to make a greater impact on the performance of inexperienced WDs.

We created scenarios that should differentiate the performance of experienced and inexperienced WDs. In these scenarios, experienced WDs are expected to perform better than inexperienced WDs, reflected by greater loss of hostile assets and less loss of

friendly assets.   We expect that (a) inexperienced WDs will utilize the agent more frequently, (b) will accept more recommendations, (c) will perform more effectively with the agent than without, and (d) will perform as effectively as the experienced WDs when the agent is available.

 The core experimental design is a 2x2, where there are 2 levels of a within-subjects variable and 2 levels of a between-subjects variable.   All subjects will first experience indoctrination and hands-on training. They will then perform in two 30-minute scenarios, one without the agent, and one where the agent is available for consultation. The order of agent-on versus agent-off will be counterbalanced.

The between-subjects variable will be experience level, as determined by the number of flight hours in the E-3 aircraft. Those who logged more than 400 hours and who have at least 1 year Combat Mission Ready experience will be categorized as experienced. This determination reflects the current policy of the 552nd Training Squadron for AWACS WDs.

In addition, after performing in the 3 scenarios (1 training, 1 agent-on, 1 agent-off) each WD will participate in a third condition, where he/she will articulate their decisionmaking process after being instructed in think-aloud procedures.   WD decisionmaking in these scenarios will be analyzed using verbal protocol analysis.

Discussion

Here we have described the method and rationale underlying the development of a multi-purpose platform focused ultimately on enhancing the performance of AWACS WD team members. The primary emphasis was on the development of agent-based decision aiding technology. It was also developed in such a way as to increase its usefulness as a training tool.   The platform was also developed to be an effective research tool, to enable and enhance investigations of C2 team decisionmaking performance.

Preliminary feedback from AWAC WD operational personnel is very positive. Even though switch actions and other task executions differ, the WDs immediately recognized the task and quickly mastered procedural skills, moving immediately to core task goals of information retrieval, coordination, and decisionmaking. Their suggestions with regard to display and other features were noted and incorporated as iterative improvements. Their feedback on decision aiding indicated high approval of its usefulness on the job and in training.

At this time, the prototype system is ready for baseline data collection, to ascertain impact on performance per se. We will begin by collecting data from operational AWACS WD personnel at Tinker AFB, as described previously. We shall also be collecting baseline data from naïve subjects, such as university students and/or cadets at the USAF Academy, to more fully investigate and refine measures of individual and team performance.

This platform will also be made available to university researchers. We expect the platform to be a useful research tool for investigations of C2 performance, team performance, and decision-making processes. It provides complex, probabilistic, and naturalistic task demands; opportunity for experimental control and manipulation, and extensive data collection features. Scenario generation features allow manipulation of numerous task characteristics, such as task timing, tempo, information uncertainty, and unexpected events. Detailed data regarding decision events allow analysis of decision processes and investigation of decision biases/error.

Future plans include transitioning the system to trainers and researchers and pursuing opportunities for further advancement of the system. Because this STE will be used in university settings, as well as in AF training settings, it provides a fundamental platform that evolves from many sources of input. Such input reflects the iterative refinements from researchers in different disciplines and from the operational world.

While the platform is extremely useful as it stands, there are opportunities for enhancements. Certainly, the validation studies should reveal potential improvements to existing decision algorithms. Also, agent recommendations are now focused on assisting defensive counter air tactics. Other features are needed to make recommendations more effective for strike roles, such as route planning. As another example, agent recommendations are currently not capable of redirecting an aircraft once it is committed to an action. This would also be a useful capability. Feedback from the WDs will be used to identify and prioritize opportunities for augmented agent capability.

There are other opportunities for agent technology within this platform environment. One challenge has been to create simulation scenarios that are realistic, relevant, and timely (i.e. rapid scenario development), and we developed the means to create a scenario more easily, but the overall mission and flow of scenario events must still be generated beforehand. Trainers and investigators require tailored scenarios, with specific training and performance goals. To achieve this goal, one might envision agents, which would take a specific training goal and generate appropriate scenarios to train and/or test the competency of a simulation player.

As another example, agents can be utilized in a more team-level performance-monitoring role, where performance of all players are monitored and recommendations regard redistribution of workload, resource, responsibility, and/or authority. For example, this agent could identify players likely to be overloaded or run out of resources, and provide recommendations on task/resource distribution. This team-level agent can be independent of agents providing specific recommendations to individual team members. In many training and research contexts, it is often desirable to know how a simulation player does without the coach or with some other decision-aid software that is not specifically bound to the coach.

Also, we do not currently have the capability to program specific events with specific probabilities of success or failure. This will likely be an important feature to implement in our future agent-based simulations. If one can also set the outcome for specific events to be low or random, one can systematically investigate campaign "critical paths". That is

programming a particular mission to fail, allows one to see what the results of that mission failure would be relative to the entire mission (i.e. how the simulation ran after that failure). Mission components with more criticality to the outcome could thus be reinforced accordingly, a priori.

In summary, this platform represents progress towards enhanced performance of AWACS WD team members, through decision aiding, as a training tool, and as a research tool. Agent technology was utilized to provide expert-based recommendations within a naturalistic and realistic emulation of core task characteristics. It also provides the means by which WDs can practice tactics and decisionmaking strategies. It allows WDs to practice individually, yet within a team environment. It provides the rationale for each decision event, and allows replay of the scenario. We expect the platform will prove effective for enhanced decision making, training, and research. Agent-driven technology represents opportunity for an array of capabilities, and we expect further enhancements and capabilities will evolve over time.

References

Chorafas, D. N. (1997). Agent technology handbook. New York: McGraw Hill.

Cobb, M.G., Mathieu, J.E., Elliott, L. R., Dalrymple, M. A., & Ballantine, G. (April, 1998). The Impact of Environmental Complexity and Ambiguity on AWACS Crew Performance in Simulated Combat Scenarios. Presented at the Aerospace Medical Panel Symposium on Collaborative Crew Performance in Complex Operational Systems, Edinburgh, United Kingdom, April, 1998.

Coovert, M. D., Elliott, L., Foster, L. L., Craiger, J. P., & Riddle, D. (1999). Measurement in synthetic task environments for teams: A methodological typology. In H. J. Bullinger and J. Ziegler (Eds.), Human-Computer Interaction: Ergonomics and User Interfaces (Vol. 1. Pp. 686-690). Mahwah, NJ: Lawrence Erlbaum Associates.

Coovert, M. D., Riddle, D., Elliott, L., & Schiflett, S., (1999, October). Identifying the relative importance of variables related to team effectiveness. Society for Multivariate Experimental Psychology, Riverside, CA.

Coovert, M. D., Riddle, D., Foster, L. L., Gordon, T., Hoffman, K., Miles, D., King, T., Elliott, L., & Schiflett, S. (1999, August). Measurement in Synthetic Task Environments for Teams: Individual Characteristics and Team Process versus Cognitive Modeling. In M. D. McNeese (Chair), Beliefs/Emotions in adaptive interface design. HCI International '99: the 8[th] International Conference on Human-Computer Interaction, Munich, Germany

Dalrymple, M., Eddy, D., & Schiflett, S. (1995). Comparative effects of antihistamines on AWACS weapons directors performance under sustained operations (AL-TR-1995). Brooks AFB, TX: Armstrong Laboratory.

Dalrymple, M., & Schiflett, S. (1997). Measuring situational awareness of AWACS weapons directors. In A. Schopper, (Ed.), Situational awareness in the tactical air environment (SOAR 97-01). Proceedings of the Naval Air Warfare Center's First Annual Symposium at Patuxent River, MD. WPAFB, OH: CSERIAC.

Eddy, D. (1989). Selected team performance measures in a C3 environment: An annotated bibliography (USAFSAM-TR-87-25). Brooks AFB, TX: School of Aerospace Medicine.

Eddy, D., Dalrymple, M., & Schiflett, S. (1992). Comparative effects of antihistamines on aircrew mission effectiveness under sustained operations (AL-TR-1992-0018). Brooks AFB, TX: Armstrong Laboratory.

Elliott, L. R., Cardenas, R., & Schiflett, S. (1999) Development and Analysis of Team-based Rating Scales Of AWACS Performance in Distributed Mission Training. Proceedings of the Annual Research and Technology Symposium on Command and Control, Newport, RI. June, 1999.

Elliott, L. R., Serfaty, D., & Schiflett, S. G. (1998). Theory-based Development of Synthetic Task Environments for C3 Team Performance and Training Research. Paper Presented at the 1998 Command and Control Research and Technology Symposium, Naval Postgraduate School, Monterey, CA, June, 1998.

Elliott, L. R., Schiflett, S. G., & Dalrymple, M.A. (April, 1998) Modeling the Decisionmaking Process and Performance of Airborne Warning and Control System (AWACS) Weapons Directors. Presented at the Aerospace Medical Panel Symposium on Collaborative Crew Performance in Complex Operational Systems, Edinburgh, United Kingdom, April, 1998.

Elliott, L. R., Dalrymple, M.A., & Neville, K (1997). Assessing Performance Of AWACS Command And Control Teams. Paper presented at the 1997 Human Factors Conference. Alburquque, NM. (Invited to present as part of panel with GMU and ARI).

Elliott, L. R., Neville, K., Dalrymple, M. A., & Tower, S. (1997). Impact of Ambiguity on AWACS Team Communications and Decisionmaking. Paper presented at the Ninth International Symposium on Aviation Psychology, April 27-May 1, 1997, Columbus, Ohio.

Hess, S., MacMillan, J., Elliott, L., & Schiflett, S. (1999).Team-in-the-Loop, Synthetic Simulation: Bridging the Gap Between Laboratory and Field Research. Proceeding of the 43rd Annual Meeting of the Human Factors and Ergonomics Society, Houston, TX; 1999.

Elliott, L. R., Hollenbeck, J. R., Schiflett, S. G., & Dalrymple, M. (In Review). Investigation of Situation Awareness and Performance in Realistic Command and

Control Scenarios. In Review for M. McNeese, E. Salas, & M. Endsley (Eds.) "Group Situational Awareness: New Views of Complex System Dynamics" .

Fahey, R. P., Rowe, A., Dunlap, K. and DeBoom, D.(in review). *Synthetic Task Design (1): Preliminary Cognitive Task Analysis of AWACS Weapons Director Teams*. Technical Report. Brooks AFB, TX: Armstrong Laboratory.

Klinger, D.W., Andriole, S.J., Militello, L.G., Adelman, L., Klein, G. & Gomes, M.E. *Designing for performance: A cognitive systems engineering approach to modifying an AWACS human-computer interface.* AL/CF-TR-1993-0093. Wright-Patterson AFB, OH: Armstrong Laboratory, 1993.

Huhns, M.N., Singh, M.P., Gasser, L. (1998). *Readings in agents*. London: Morgan Kaufmann.

MacMillan, J, Serfaty, D., Young, P., Klinger, D., Thordsen, M., Cohen, M., & Freeman, J. (1998*). A system to enhance team decision making performance: Phase 1 Final Report.* Brooks AFB: Air Force Research Laboratory, Warfighter Training Research Division.

Schiflett, S. & Elliott, L. (in review). Development of Synthetic Team Training Environments: Application to USAF Command and Control Aircrews. To appear in H. O'Neil & D. Andrews (Eds.) *Aircrew Training: Methods, Technologies, and Assessment.*

Schiflett, S., Strome, D., Eddy, D., & Dalrymple, M. (1990). Aircrew Evaluation Sustained Operations Performance (AESOP): A triservice facility for technology transition (USAFSAM-TP-90-26). Brooks AFB, TX: USAF School of Aerospace Medicine.

Stoyen, A., Chiara, J., Schiflett, S., Ellliott, L., Walrath, L., & Dalrymple, M (June, 1999). Situational Awareness Measurement in AWACS Weapons Director Teams. Paper presented at Annual Conference on Situation Awareness, Pax River, MD.

Tecuci, G. (1998). *Building intelligent agents*. Boston, MA.: Academic Press.