Multi-Level Human-Autonomy Teams for Distributed Mission Management

Martin Voshell,[1] James Tittle,[1] Emilie Roth[2]
[1] Charles River Analytics, 625 Mt. Auburn St., Cambridge, MA 02138, mvoshell@cra.com; jtittle@cra.com
[2] Roth Cognitive Engineering, Stanford, CA 94305, emroth@rothsite.com

Abstract

Control of the air in envisioned large-scale battles against near-peer adversaries will require revolutionary new approaches to airborne mission management, where decision authority and platform autonomy are dynamically delegated and functional roles and combat capabilities are assigned across multiple distributed tiers of platforms and human operators. System capabilities range from traditional airborne battle managers, to manned tactical aviators, to autonomous unmanned aerial systems. Due to the overwhelming complexity, human operators will require the assistance of advanced autonomy decision aids with new mechanisms for operator supervision and management of teams of manned and unmanned systems. In this paper we describe a conceptual distributed mission management approach that employs novel human-automation teaming constructs to address the complexity of envisioned operations in highly contested environments. We then discuss a cognitive engineering approach to designing role- and task-tailored human machine interfaces between humans and the autonomous systems. We conclude with a discussion of multi-level evaluation approaches for experimentation.

Introduction

Near-peer adversaries can be expected to employ technical and tactical measures to degrade the effectiveness of traditional airborne battle management in attempt to diminish friendly forces' abilities to understand battlespace situations, communicate with peers, adapt, and respond to dynamic and changing events. Current approaches to airborne mission management require a tightly coordinated effort over reliable networks by battle managers, operators, and tactical pilots who have little or no automated decision aid support. However, with movement toward next generation air platforms, manned and unmanned multi-role aircraft, and the new complexities of operating across permissive to highly contested environments, classic approaches to airborne mission management are insufficient when facing a near-peer threat for reasons of complexity, scale, and degraded, unreliable communications. This implies the need for systems with flexible automation and well-conceived coordination protocols for granting decision-making authority to appropriate operators, pilots, or automation systems in a consistent and appropriate manner. Complexity in such envisioned environments is further heigtened by expected dramatic increases in the development and use of autonomous unmanned aerial systems across a broad range of roles, from surveillance to delivery of effects. Near future mission packages will be composed of both manned and unmanned systems in the same teams (e.g. autonomous unmanned wingmen). Successfully managing this complexity requires wholly new approaches to human systems

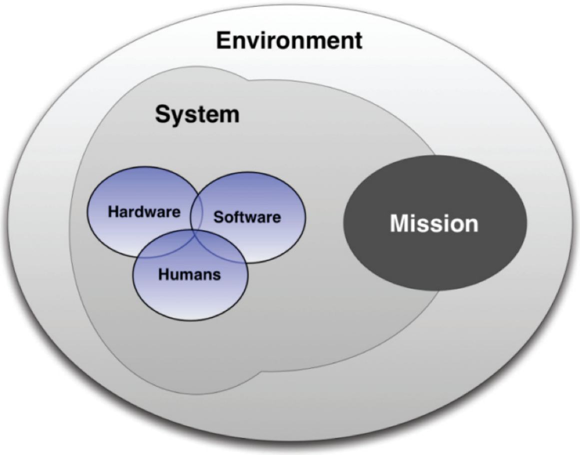

Figure 1: Adapted from USAF Report on Human Systems Integration SAB-TR-04-04; effective HSI must incorporate human factors engineering, manpower, safety, and personnel

---

integration (HSI; see Figure 1) and new architectures that enable dynamic distribution and delegation of appropriate levels of decision authority to multiple manned and unmanned systems that work collaboratively to achieve mission objectives.

Human operators will be dynamically assigned a variety of simultaneous roles and have decision spaces and timelines consistent with those roles. For example, airborne battle manager teams (ABM) have broader decision spaces and longer timelines in which to make those decisions, while manned tactical pilots with delegated decision authority have much narrower decision spaces and shorter timelines. These differences significantly impact the scope and potential roles of autonomy in supporting the human operator and the nature of human-autonomy collaboration.

An implicit system design requirement is the need for the system to be truly *coordination-centered* for the human operators and stakeholders working with advanced automation. Operators must always be "on-the-loop"- remain in positive control - with autonomy providing critical support that enables management of complex tasks. This management includes providing effective alerts for dynamic decisions, recommendations for actions and response, and decision review for inconsistencies and potential errors. This is evidenced in the Defense Science Board (DSB) Task Force's recent report on autonomy (2012), which specifically calls out the need for decision-centered Human Machine Interface (HMI) research while also identifying that insufficient research has focused on the unintended consequences of introducing such capabilities in envisioned operations. The DSB's resulting trade space of needed technology developments provides critical focus areas for new science and technology (S&T) programs, including: natural user interfaces and trusted human-system collaboration; perception and situational awareness to operate in complex battle spaces; and large-scale teaming of manned and unmanned systems.

Effective and novel approaches to human machine interaction (HMI) that go beyond mere display technology for situation awareness and classic approaches to supervisory control are required. When coupled with on-going changes in operational environments, the need for compatibility with current and future aircraft and force structures and evolving operating concepts ultimately requires detailed examination of how to effectively design, test, and integrate such new capabilities and enabling technologies into HMIs that are designed from the outset to support effective human-autonomy team coordination.

Envisioned CONOP and Architecture

Key Concepts and Requirements

System design guidance on the application of new automation technologies in complex work domains (Decker 1987; Rasmussen 1991; Woods and Dekker 2000; Woods and Hollnagel 2006) has argued that self-organizing system structures provide one of the most effective approaches for effectively meeting the dynamic requirements of new operational contexts. The allocation of the human and machine functional roles in such dynamic environments will depend not only on mission context, but will also be driven by critical considerations related to how human operators and decision-makers will "trust" the automated agents to perform complex and non-trivial tasks across their respective functional roles (Moray, Hiskes, Lee, and Muir 1995; Lee and See 2004; Kilgore and Voshell 2014).

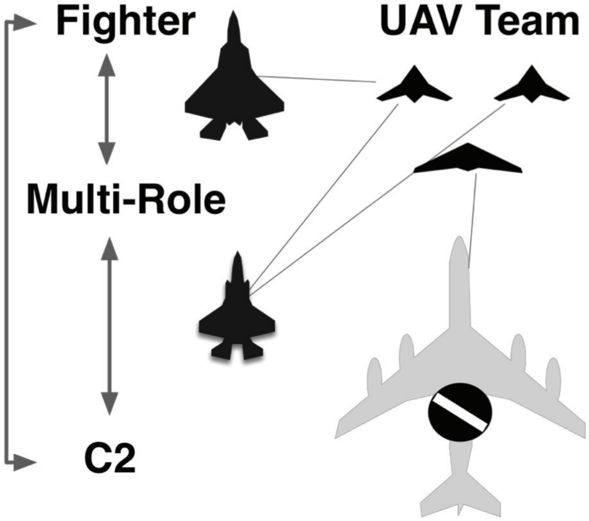

Figure 2: Operational view of multiple, mixed manned/unmanned teams.

Figure 2 provides an operational view of multiple manned and unmanned mission management nodes forming collaborative teams with varying levels of decision authority. In this envisioned concept of operations (CONOP) each aircraft in this figure, manned and unmanned, functions as a mission management node. Human operators fill multiple roles, from airborne battle managers leading teams of teams, to manned tactical pilots functioning as flight leads with both manned and unmanned wingmen.

Based on flight tactics and mission plans, some mixed manned/unmanned teams may persist, while others are dynamically formed as resource needs and assigned tasks change. Further, any mission management node can participate in multiple teams simultaneously with different roles on each team. Team leaders may be dynamically reassigned according to conditions (e.g. for reasons of degradation or attrition). Team leaders may delegate authority to team members in a manner consistent with commander's intent and their own delegated authority. For example, a manned tactical team leader may have authority to form sub-teams of unmanned vehicles to execute assigned tasks.

Functional Roles for Human Automation Teams

Automation provides many capabilities that enable this conceptual architecture. However, such envisioned operations, with pilots supervising small teams of heterogeneous unmanned platforms, transform the roles of pilots and ABMs to act more as supervisory controllers, responsible for managing a wide variety of automated subsystems and functions. Because the increased autonomy will change the traditional roles, responsibilities, and tasks of the human operators, it can also introduce opportunities for new classes of errors given new mission complexities and timelines. Because of this, the design of automation support must be sufficiently robust and tailored to key envisioned mission functions. In the current work, these functions were supported through three discrete components, Distributed Situation Understanding (DSU), Distributed Planning and Control (DPC), and Task Execution (TE) services. Since the conceptual architecture is modular and composable, the particular methods and algorithms used are less important than the functional capabilities they provide. DSU services provide capabilities for improving awareness of the battlespace to include awareness of friendly (blue) force status and capabilities and the communications network across nodes. The methods for achieving improved awareness are algorithms for distributed information integration (fusing information from multiple on-and off-board sources) and methods for maximizing the utility of limited communications bandwidth through smart distribution of information (i.e., based on mission value). DSU support services actively seek to align and synchronize the common operating picture to enable teams to collaborate.

Nodes equipped with DPC services have the ability to collaboratively and adaptively re-plan to respond to events and changing conditions, support team adaptation and role assignment, provide optional courses of action and contingencies, and other activities needed for teams to collaborate in pursuit of achieving mission goals and tasks. All nodes with asset management responsibilities are equipped with task execution services, which decompose assigned tasks to low level actions. Task execution may itself employ capabilities with varying levels of autonomy that enable a higher level tasking.

For the pilots (taking on additional supervisory control activities) and for the ABMs, integrating such automation support into these future missions will also require a significant shift away from low level parameter monitoring of the coordinating systems toward improved decision support for monitoring automation goals, handling exceptions, diagnosing anomalies, and anticipating surprise. Across the distributed systems, human machine interfaces (HMIs) must be designed to ensure positive and effective feedback and provide input-controls tailored to the human's decision support needs for both interacting with their autonomous decision support tools and with other support nodes in their teams.

HMI Design Patterns

An important consideration when designing such a conceptual distributed cognitive system is that introducing new technology fundamentally impacts the nature of the work domain (Woods and Dekker 2000). To successfully integrate advanced automation given the changing nature and impact on the work domain requires that the guiding research underlying new automation and interface designs must focus on providing operator decision support for both the current work domain and the envisioned world of work created by new distributed automation technologies.

In general, more sophisticated automated systems represent an increase in both the autonomy and authority of automated components that must be accounted for (Woods 1996). When these relationships are ignored, coordination surprises occur between humans and automation. Sarter, Woods, and Billings (1996) have identified a number of lessons learned and characterizations of brittle design patterns found in complex automation design in a variety of supervisory control systems. For example, lack of observability with complex automated agents has the potential to lead to "automation surprises", where an automated agent can become "strong and silent, and difficult to control". Another example pattern arises when goal conflict situations between distributed human/automation teams result in responsibility-authority double-binds. Double-binds manifest when a human agent is ultimately responsible for any errors or failures of the human-automation team, but the automation has the authority of action.

For advanced automation support to be both effective and trusted, system developers must design for fluent coordination between the human and automated elements of the system (Christoffersen and Woods 2002) through sufficient increases in automation *observability and directability*. To design for coordination, Hollnagel and Woods (2006) have summarized research showing that human-machine interfaces that employ better feedback about the dynamics and future behavior of automated processes increase system *observability*. As automated processes are accessed and become increasingly complex, it is critical to increase their observability. The following decision support design patterns represent high-level classes of HMI requirements necessary for effective human-automation teaming:

Observability refers to the ability of an operator to form insights into processes in the work domain or in the supporting automation. Observable feedback should provide a dynamic picture of the current situation, relevant automation activities, and information defining how these interactions may evolve in the future (Woods & Sarter 2000).

Directability refers to the use of those process insights to purposefully direct/re-direct resources, activities, and priorities as situations change (Roth, et al., in press). Effective design methods for observability create clear affordances for automation interaction. By incorporating directable interfaces into HMI designs requires identifying meaningful parameterizations of autonomous systems that enable human operators to understand supporting automation at the right level of abstraction (Myers & Morley 2003).

Directed attention helps to re-orient focus in dynamic and changing work environments (Woods 1995; Sklar & Sarter 1999; Nikolic & Sarter 2001). This includes the ability to track the focus of attention for both human and automated agents and to recognize the degree to which each may be interrupted or re-directed based on current tasking.

Joint Activity Coordination is critical for effective teaming between human and automated agents. While directed attention support represents one type of coordination between human and automated agents, supporting other types of joint activity coordination requires interfaces that actively coordinate and synchronize tasks across human and automated actors using display elements that explicitly represent automation intent and goals within the context of the human stakeholders' activities. Such displays support common ground based on the pertinent knowledge, beliefs and assumptions that are shared amongst agents involved in creating and maintaining effective joint activity (Billings 1997; Shattuck & Woods 2000; Klein et al. 2004; Nass & Brave 2005). Common ground design support is critical for coordinating with other mission elements to achieve mission objectives. Coordination may involve synchronizing activities (e.g., allocating targets or threats across different pilots; or across different members of a flight group; coordinating to minimize risk); integrating sensor assets to achieve improved situation awareness (e.g., positioning different entities in different locations so as to gain a complete joint common operational picture); and selecting and allocating weapons effectively.

Human Automation System Evaluation

Developing novel distributed mission management capabilities and effective human interfaces will require a significant amount of structured experimentation in live, virtual, and constructive environments, with multiple operator workstations fulfilling multiple roles to interact with each other and their automated teammates. Assessing the effectiveness of the human automation teams in these complex work environments requires the ability to effectively adapt evaluation approaches to emerging contexts. For this reason, standard usability evaluations are a necessary but insufficient means to fully address the wide spectrum of essential human-automation coordination mission task contexts.

We developed a work-centered multi-level system evaluation approach (based on Roth & Eggleston, 2010) that goes beyond traditional assessment of usability in a system. This approach consists of a complementary set of methods addressing three levels of evaluation that, taken together, characterize the effectiveness of the joint human automation system: *usability*, *usefulness*, and *impact*. In addition to more traditional assessments of *usability,* in early stages of design prototyping the HMI experiments provide a critical, work-centered perspective by addressing *usefulness* (assessment of whether the information content and interaction methods are sufficient to support effective human-automation decision making performance). The final level of evaluation, assessing the *impact* of the system, addresses whether this simplification has reduced the ability of the human and the automation to adequately understand and anticipate the true state of the world, and the effectiveness of the joint human automation system in meeting mission objectives. Our evaluation of *usefulness* and *impact* specifically focuses on empirical assessment of the *observability* and *directability* support different HMI design solutions provided. A critical focus for the impact assessment is related to quantitative system performance measures related to the introduction of specific complicating factors inserted as part of LVC scenarios.

Conclusions

In this paper we described a novel HSI approach to airborne mission management that distributes decision authority and mission management functions across teams of multiple manned and unmanned platforms and systems. The roles of human operators are fundamentally changed under this paradigm, and the related complexity requires the support of new automated decision aids, thus demanding new concepts for supporting human-automation teaming. The effective development of multi-level human-automation teaming will ultimately require HMI design approaches that support coordinated joint activity by making autonomous system components observable and directable as team players. Throughout this project, a Cognitive Systems Engineering research effort identified new roles, responsibilities, and tasks across the distributed mission management system in different envisioned mission contexts. The application established HMI design patterns to these newly identified human roles that have the potential to support coordinated joint activity of air battle managers, pilots, and autonomous systems. The insights gained from the study and empirical evaluation of human-automation teaming in this particular domain is applicable across many different work complex domains where increasingly advanced automation is integrated with distributed human teams in dynamic and changing environments.

Acknowledgments and Distribution

Portions of this work were sponsored by the Defense Advanced Research Projects Agency (DARPA) under the Distributed Battle Management program. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. This document has been approved for Public Release, Distribution Unlimited.

References

Alberts, D.S.., and Hayes, R.E. 2002. *Code of Best Practice for Experimentation*. Washington, DC: DoD Command and Control Research Program.

Alberts, D.S.., and Hayes, R.E. 2005. *Campaigns of Experimentation*. Washington, DC: DoD Command and Control Research Program.

Billings, C. E. 1997. *Aviation Automation: The Search for a Human-Centered Approach*. Hillsdale, NJ: Lawrence Erlbaum.

Christoffersen, K. and Woods, D. D. 2002. How to make automated systems team players. In E. Salas. ed, *Advances in Human Performance and Cognitive Engineering Research*, Volume 2. St. Louis, MO: Elsevier Science.

Decker, K.S. 1987. Distributed problem-solving techniques: A survey. *IEEE Transactions on Systems, Man, and Cybernetics*. SMC-17: 729-40.

Hollnagel, E. and Woods, D.D. 2005. *Joint Cognitive Systems: Foundations of Cognitive Systems Engineering*. Taylor & Francis.

Kilgore R. and Voshell, M. 2014. Increasing the transparency of unmanned systems: applications of ecological interface design. In: Shumaker R, Lackey S, editors. Virtual, augmented and mixed reality: applications of virtual and augmented reality. VAMR 2014. *Proceedings of the 6th International Conference, Part II*; 2014 Jun 22–27, 378–389. Heraklion, Crete, Greece. New York (NY): Springer.

Klein, G., Woods. D.D., Bradshaw, J., Hoffman, R.R., and Feltovich, P.J., 2004. Ten Challenges for Making Automation a "Team Player" in Joint Human-Agent Activity. *IEEE Intelligent Systems*, November/December: 91-95

Lee, J. and See, K. 2004. Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1): 50-80.

Moray, N.P., Hiskes, D., Lee, J., and Muir, B.M. 1995. Trust and human intervention in automated systems. In J-M Hoc, P.C. Cacciabue, and E. Hollnagel Eds. *Expertise and technology: cognition and human-computer cooperation*, 183-195. New Jersey: Lawrence Erlbaum Associates.

Myers, K., and Morley, D. 2003. Directing agents. In H. Hexmoor, C. Castelfranchi & R. Falcone (Eds.), *Agent autonomy*. Dordrecht, The Netherlands: Kluwer.

Nikolic, M. I. andd Sarter, N. B. 2001. Peripheral visual feedback: A powerful means of supporting attention allocation and human-automation coordination in highly dynamic data-rich environments. *Human Factors*: 43, 30-38.

Nass, C. I., and Brave, S. 2005. *Wired for speech: How voice activates and advances the human-computer relationship*. Cambridge: MIT press.

Rasmussen, J. 1991. Modeling Distributed Decision Making. In J. Rasmussen, B. Brehmer, and J. Leplat Eds., *Distributed decision making: cognitive models for cooperative work*, pp. 112-142. London: Taylor and Francis.

Roth, E. M., DePass, B., Scott, R., Truxler, R., Smith, S. Wampler, J. (in press). Designing collaborative planning systems: Putting Joint Cognitive Systems Principles to Practice. In P. Smith and R. R. Hoffman (Eds). Cognitive Systems Engineering: *The Future for a Changing World*. Ashgate Publishing Limited

Roth, E. M. and Eggleston, R. G. (2010). Forging new evaluation paradigms: Beyond statistical generalization. In Patterson ES, Miller J. (Eds.) *Macrocognition Metrics and Scenarios: Design and Evaluation for Real-World Teams*. Ashgate Publishing. ISBN 978-0-7546-7578-5. (pp. 203-219).

Shattuck, L. G., and Woods, D. D. 2000. Communication of intent in military command and control systems. In C. McCann & R. Pigeau (Eds.), *The human in command: Exploring the modern military experience*. New York: Plenum.

Sklar, A.E. and Sarter, N.B. 1999. "Good Vibrations": The use of tactile feedback in support of mode awareness on advanced technology aircraft. *Human Factors*, 41: 543-552.

Woods, D. D. 1995b. The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38: 2371-2393.

Woods, D. D. and Dekker, S. W. A. 2000. Anticipating the Effects of Technological Change: A New Era of Dynamics for Human Factors. *Theoretical Issues in Ergonomic Science*.

Woods, D. D. and Hollnagel, E. 2006. *Joint Cognitive Systems*. CRC Press.

Woods, D.D. and Sarter, N. 1993. Evaluating the Impact of New Technology on Human-Machine Cooperation. In J. Wise, V. D. Hopkin, & P. Stager, (Eds.), *Verification and Validation of Complex Systems: Human Factors Issues*. Berlin: Springer-Verlag.

Woods, D.D. and Sarter, N.B. 2000. Learning from automation surprises and "going sour" accidents. In N.B. Sarter and R. Amalberti (Eds.), *Cognitive Engineering in the Aviation Domain* (pp. 327-353). Hillsdale, NJ: LEA.