

**2002 Command and Control
Research and Technology Symposium**

Topic: C2 Decision Making & Cognitive Analysis

Mixed-Initiative Control of Robotic Systems

*Jared T. Freeman, Ph.D.	Jean MacMillan, Ph.D.
Aptima, Inc.	Aptima, Inc.
Washington, DC	Woburn, MA

*Contact: Jared Freeman, Ph.D.
Aptima, Inc.
1030 15th St., NW, Suite 400
Washington, DC 20005
202-842-1548 x316 (voice)
202-842-2630 (fax)
freeman@aptima.com

Mixed-Initiative Control of Robotic Systems

Jared Freeman, Ph.D.

Aptima, Inc.
1030 15th St., Suite 400
Washington, DC 20005
freeman@aptima.com

Jean MacMillan, Ph.D.

Aptima, Inc.
12 Gill Street Suite 1400
Woburn, MA 01801
macmillj@aptima.com

Abstract

A critical topic of research concerning human interaction with robotic warriors concerns the functionality of intelligent systems to advise human operators and share control of robots with those operators. This functionality will engage human and software systems in a complex, highly interdependent exchange of information and control as human initializes systems that advise them, refine system recommendations, and trade off control of robotic forces with the system during mission execution. In research for DARPA and the U.S. Army, we have defined the Relational Knowledge Framework that defines fundamental classes of human interactions with intelligent robotics systems planning and control systems. Several cognitive issues are prominent in these interactions. They suggest that system design and training should support specific types of knowledge by operators. These concern the relations (thus, the relational knowledge framework) between (1) the current state of the battle or the system and norms, (2) system parameters and system operations, (3) system inputs and real world events, and (4) control decisions and the control interface. The framework, cognitive issues, and training and design requirements are defined.

Introduction

Robotic combat systems are a growing presence in the news, in military laboratories, and in the battlefield. Consider the news coverage concerning the shoot-down of an unmanned aerial vehicle in Afghanistan for evidence of the growing role of unmanned systems (if not fully robotic) in battle and in the public consciousness. In military laboratories, we see an increasing focus on robotics in experiments such as Future Joint Forces (FJF) at Ft. Knox and others concerning Future Combat Systems (FCS).

The cutting edge of military robotics research concerns the manner in which intelligent systems will collaborate with human operators in planning and executing battle involving both robotic and human forces. These

systems will engage human and software systems in a complex, highly interdependent exchange of information and control as human operators initialize systems that advise them, refine system recommendations, and trade off control of robotic forces with the system during mission execution. Such synchronized, human-machine command of robotic forces is called Mixed Initiative Control of Automata, or MICA, and is the focus of a large R&D program sponsored by DARPA.

A significant challenge in this program is to understand how humans and systems should interact to ensure success on the battlefield. The traditional approach has been to specify that humans need training that develops robust mental models of these systems in order to monitor and correct system performance. However, operators are generally, if not

universally incapable of building complete mental models of complex systems operating in dynamic warfare.

A more subtle view (Cohen, Parasuraman, and Freeman, 1998) is that operators must develop mental models of the system that help them to discern the contexts in which the system can and cannot be trusted to perform competently (i.e., the former are contexts of which the system is "cognizant"), and the level of accuracy to expect from the system in contexts it recognizes. This view is interesting because it suggests that the mental models operators hold can be partial. Specifically, operators need only rough models of the system's ability to discriminate different tactical situations, and fine grained models of system operation in the potentially small set of situations that the system recognizes well. This lowers the criteria for competency among operators to a more realistic level, and specifies the cognitive problem in a way that supports design and training.

In observations of FJF and FCS, and in work on DARPA's MICA program, we have developed a framework with which to further specify the requirements for operator knowledge and provide more support for designing usable systems and helpful training. We call this theory the Relational Knowledge Framework because it emphasizes the role of knowledge concerning relations between mission plan and mission state, relations between real-world entities or events and system parameters, the relative influence of various system inputs on system performance, and so forth. The framework posits several fundamental classes of human interactions with intelligent systems for planning and controlling robotic forces. We describe these below and draw specific design implications in Table 1. In the subsequent section, we present four types of relational knowledge that span many of these classes, and draw some implications for system design.

Classes of Human-System Interaction in the Relational Knowledge Framework

There are seven fundamental classes of human-system interaction with intelligent

advisory and control systems for robotic forces. We define and illustrate these classes here. In , at the end of this paper, we present implications of these classes for design.

The human must configure (or reconfigure) the system to determine which functionalities the system will apply to the mission at hand and at what levels of precision. For example, the system may employ different controllers or algorithms for different mission phases (planning vs. execution), mission type (offensive vs. defensive), or types of objectives (hard vs. soft targets). It may provide the user a choice between rapid, rough solutions and more deliberate but precise ones. These configuration interactions challenge the operator to understand the functions the system can apply to a mission, the conditions (e.g., missions with rough weather, missions with high potential for fratricide or collateral damage) under which it can competently perform, and its reliability in contexts it "understands."

The human must provide data for the system to process in the configuration specified, above. Examples include specifying current weather, targets, intelligence, and other data uniquely available to the human operator. This interaction requires the operator to possess and exercise a wide range of knowledge including the mapping of real-world events to system data requirements, and the state of those events and current system values.

The human must review system recommendations and accept, adjust, or reject them. For example, when the system generates alternative COAs, the human must review them, select among them, and potentially refine the best choice. Fundamental challenges to the operator are to think critically about complex recommendations, understand how and when to query the system for explanations of surprising recommendations, and understand how manual edits may improve or undermine plans.

The human should monitor system execution of the mission. The human must, for example, track the actual routes of robotic forces relative to planned routes to ensure that encounters with enemy forces, weather, and other dynamic obstacles to not hamper the

mission significantly. This interaction requires the human to understand the status of execution relative to the plan, understand which deviations from those plans have serious consequences, recognize events that should trigger human decisions, and understand the methods and costs of dynamically re-planning to compensate for emergent problems.

The human must monitor & refine system performance. Examples of these interactions include monitoring for sluggish system response, degraded information quality, and crashes of robots or the advisory and control system itself. Fundamental challenges to the operator are to know the norms of system performance in the given mission class, discriminate levels of degradation that significantly endanger the mission, and understand how to diagnose and work around system malfunctions.

The human must be able to take direct control of entities and functions otherwise allocated to the system. In one, current robotic force, for example, the human must manually assume control of robots that are orphaned when their C2 units are destroyed. This interaction requires the operator to understand how to transfer control of an entity from the system to the human and back again, and to possess manual control skills.

The human must balance the workload imposed by interactions with the system and by functions assigned only to human. For example, the human may be tasked both with managing the advisory system plus driving a C2 vehicle, conducting human comms, and exercising a host of other responsibilities that impose potentially large workload. To operate in this environment, the operator must monitor and even predict the workload being imposed by the human-only tasks and human-system interactions, understand the priority of tasks, and know methods of pausing selected tasks in order to conduct others.

Cognitive Issues in the Relational Knowledge Framework

The interactions, above, are associated with four cognitive issues. Among several of these

issues runs a thread concerning relational knowledge, hence the name Relational Knowledge Framework. At this, more abstract level, we can identify interesting implications for design and training. The four issues are:

- Situational awareness -- The human must understand the relationship of the battle to the plan, and the current system state relative to norms. This requires displays that emphasize departures from plans and performance norms, as well as diagnostic aids such as self-explaining system intelligence.
- Mental models of the system -- The operator must understand which inputs or parameterization actions will *significantly* influence the system. This knowledge helps the operator to invest effort in interactions that matter, inputting information that shapes system advice and control and withholding information that does not influence the system. This requires either extensive training or displays that convey the current sensitivity of the system to different inputs.
- Translation between representations -- The human must understand the real-world entities and events that correspond to system inputs and parameters. This requires sound training, but also design that simplifies this mapping. For example, map icons representing entities in the environment and relations between them (such as routes between forces and targets) should serve not only to provide situation awareness, but also as interfaces to system parameters.
- System control knowledge -- The human must have expert skills in the buttonology of controlling the system: applying filters to vast information flow, inputting and monitoring data, selecting and editing recommendations, controlling entities, etc. This demands a human-centered approach to the design of interfaces and training.

Summary

The proposed paper will present the Relational Knowledge Framework, and provide examples of each class of human interaction with an intelligent advisory and robotics control system. It will define the fundamental cognitive issues in human-system interaction in mixed-

initiative control of automata, and present concepts for supporting cognition.

Acknowledgements

Aspects of this work are funded by DARPA and the U.S. Army. The opinions expressed here are the authors' and do not necessarily reflect the views of DARPA, the U.S. Army, or the Department of Defense.

References

Cohen, Marvin S., Parasuraman, Raja, and Freeman, Jared T. (1998). Trust in Decision Aids: What is it and how can it be improved? *Proceedings of the 1998 Command and Control Research & Technology Symposium*, Monterey, CA.

Table 1: Classes of interactions, challenges, and implications for design.

Interaction Class	Challenges	Design implications
Configuration	Understand the functions the system can apply	Present well-categorized, mission-specific menus of functions
	Understand the conditions under which it can competently perform	Present reminders of mission-specific factors to which the system is insensitive, but which are of known importance to domain experts
	Understand the system's reliability in contexts it "understands."	Represent the margin of error, confidence bounds, or distribution of confidence around system estimates
Input	Map real-world events to system data requirements	Label parameters using meaningful domain labels. Provide examples. Highlight an object on all representations (e.g., a geoplot) as its parameter values are selected on another representation (e.g., a table).
	Map the state of those events to current system values.	Allow users to input categories (rather than scalar values) when they tend to categorize real-world events & entities.
Review recommendations	Think critically about complex recommendations, or simple recommendations based on complex premises or processing	Flag predictions and plans based on low certainty estimates, present the time available to resolve uncertainty, highlight information gaps and untested assumptions, and present alternative plans or assessments
	Understand how and when to query the system for explanations	Display sources of information that are relevant to each known information gap and assumption.
	Understand how manual edits may improve or undermine plans.	Provide indicators of the sensitivity of the system to various input parameters in the current context.

Interaction Class	Challenges	Design implications
Monitor mission execution	Understand status of execution relative to the plan,	Display planned route, goals (e.g., targets), and constraints (e.g., SAM sites)
	Understand which deviations from those plans have serious consequences,	Display confidence bounds with respect to route, time, and risks
	Understand recognize events that should trigger human decisions,	Make explicit the decisions the human must make. Where decisions can be scheduled, present reminders to the operator in a timely manner.
	Understand understand the methods and costs of dynamically re-planning to compensate for emergent problems	Represent the impact on mission schedule of delays due to replanning & impact on success
Monitor system performance	Know norms of system performance in the given mission class,	Represent current system performance relative to norms and thresholds given the mission type
	Discriminate levels of degradation that significantly endanger the mission,	Represent the impact of current system degradation on mission schedule & outcomes
	Understand how to diagnose and work around system malfunctions.	Build & maintain a user-extensible "tip sheet" on which users can document methods of refining system performance. Use this for reference, training, and system refinement.
Trade control		Present re-routing tools and other controls automatically in situations in which stakes and opportunities change radically.
	Understand when to transfer control	
	Understand how to transfer control	Represent who is in control (the operator or the software), the control switch, and progress towards transferring control (if the process is lengthy)
	Understand how to control entities	Implement sound UI design principles for device control and feedback.

