

Etiquette Within and Between Large Human-Robot Teams

Christopher L. Johnson

The MITRE Corporation
7515 Colshire Drive
McLean, VA 22102-7508
cj@mitre.org

Abstract

This paper describes the Defense Advanced Research Project Agency's (DARPA) Mixed-Initiative Control of Automa-teams (MICA) program as a context for investigating human-computer etiquette. The goal of the MICA program is to enable the control of large-scale teams of semi-autonomous vehicles by a relatively small number of human operators. Program research focuses on both control-theoretic techniques for autonomous control, as well as cognitive engineering techniques for effective human interaction with the resulting control system.

Also, thoughts are presented on how shared knowledge between humans and computers could be an important component of a human-computer etiquette approach, both in managing dynamic aspects of etiquette, and in forming a basis for fundamental polite behavior.

In these working notes, I'd like to present two topics for discussion. First, I will describe a current research program in which I'm involved called *Mixed-Initiative Control of Automa-teams* (or MICA). I will describe the domain and problem the program is exploring, the basic approaches that are being taken to address human interaction, and a few of the unique issues that have arisen when combining the two. Perhaps this description can then stimulate discussion on whether an etiquette approach, as developed in this workshop, can help solve these issues, or push the research in new directions.

Second, I will present a few brief thoughts on how shared knowledge between humans and systems could contribute an important component to the successful implementation of human-computer etiquette.

Mixed-Initiative Control of Automa-teams

At the moment, the term *unmanned*, when applied in the military, generally indicates some type of teleoperated vehicle. For instance, Unmanned Aerial Vehicles (UAVs) like the Predator have been used in operations for many years. However, military planners envision future theaters

of operation rife with a variety of unmanned vehicles and robots, operating essentially under their own control.

Teleoperated vehicles continue to pose human interaction issues of their own, but this future vision of large teams of semi-autonomous vehicles requires new modes of monitoring and control—especially because these teams are not only large, but because the behavior of each individual member can itself be unique and complex. Also, the rate at which these semi-autonomous teams, sub-teams, and individual vehicles carry out their missions will typically be too fast for human operators to manage unassisted.

Even though ideally these new vehicles will often operate autonomously, we assume that the human will continue to be involved at some level, collaborating with the robots to issue goals and directives, help manage uncertainty, and inject flexibility and creativity into the system. (While this research is intended to apply to any type of semi-autonomous, physically realized agent, I will use the term *robot* for clarity.)

While a fair amount of research has begun exploring the interaction between people and individual robots, or small groups of robots in localized areas (e.g., Schultz and Parker 2002, Murphy and Rogers 2001), the problem of intensive mixed-initiative management of large teams of robots distributed over large areas has yet to be studied extensively.

The Defense Advanced Research Project Agency's (DARPA) Mixed-Initiative Control of Automa-teams (MICA) program is currently studying this very issue. The goal of MICA is to enable the control of large-scale teams of semi-autonomous vehicles by a relatively small number of human operators. By the end of the program, we want to operate in a challenge scenario deploying 10 teams, each team with a 1:30 operator to vehicle ratio.

A major focus of the program is on autonomous control of the vehicles using control-theoretic methods, developing theory, algorithms, software, modeling and simulation technologies. Working from a notional hierarchical structure, program researchers are investigating the areas

of team composition, inter-team coordination, cooperative path-planning, and uncertainty management.

However, the mixed-initiative issues involved in including humans in the control process are an equally important focus of the program. For instance, how can the system support meaningful collaboration between human decision makers and teams of autonomous entities that are operating from complex control algorithms in complex and uncertain environments? How is the performance, stability, and robustness of the system affected when an operator can take control at varying levels and at varying times? Explicitly considering human interaction issues simultaneously with technology development is historically rare in these kinds of research programs, and it has worked very well so far in the year since MICA began.

Distinguishing Features

Listed below are a few of the distinguishing features—considered as a group—of the MICA domain:

- Humans interact with physical agents
- The number of robots is large
- Each individual robot possesses its own complex behavior
- Humans and robots are distributed over a large area
- Humans can be supervisors as well as peer members of teams
- Underlying autonomous control is complex, likely difficult for human to understand on their own
- A potential exists for highly fluid team tasking and structure
- Operations occur in a military context
- Operations occur in the presence of adversarial behavior
- Situational events could occur at high speed

The challenge program for the MICA program concentrates on controlling teams of UAVs in a Suppression of Enemy Air Defense (SEAD) scenario. In other words, operators will be using semi-autonomous, armed UAVs to prepare a battlefield for air operations by incapacitating enemy radar, surface-to-air missiles (SAMs), communication links, etc. However, MICA research is intended to generalize to any domain involving autonomous team behavior, such as search and rescue, environmental exploration, etc.

As an example, an interesting issue that appears in operator interaction with autonomous control techniques is that of predictability and expectations. A standard concept in human-computer interaction design is consistency. However, in an adversarial environment, predictable behavior is often the last thing we want. Furthermore, one of the hopes for control-theoretic techniques is the discovery of new behaviors. Researchers

will need to find the right levels between consistency, unpredictability, and explanation that will still allow a human operator to understand what the control system is doing at any given time.

In such an environment, with physical agents performing combat-related tasks in real-time in the real world, etiquette concerns such as safety, trust, and efficiency of interaction are critical.

Mixed-Initiative Approaches

The MICA program was structured so that control engineers were required to form research teams with cognitive engineers, and work closely together in integrating the needs of the human operator with the developing control technology. My role on the program is as a member of the DARPA management team, advising program management on issues of supervisory control and cognitive engineering. Therefore, this section simply provides an overview of the kinds of mixed-initiative approaches the research teams are applying to the MICA problem.

The core of each approach is a traditional cognitive engineering analysis of the domain (Roth, Patterson, and Mumaw 2001), determining the kinds of information and control decisions that human operators will need to successfully interact with these large teams of robots. One interesting issue we've encountered is that we're applying these techniques to a hypothetical rather than an existing domain. Researchers can gather insight from UAV experts, current military operations experts, future military operations experts, perhaps experts in similar domains (e.g., air traffic controllers), but they have no access to experts with extensive experience in managing large teams of robots.

One group has created a baseline interface containing the kinds of information a control theorist wants to see. They then plan to work from there, putting the interface in front of more operational users and using their responses to improve the design. Another group has explored ways to visualize the operation of specific control technologies, such as policies for partially-observable Markov Decision Processes.

Others are looking at creating knowledge models with the information they gather from cognitive engineering analysis and using those models to drive components of the interaction. For instance one group is building an interaction design based around the metaphor of a playbook, where the user can control a team by selecting 'plays', adjusting the details as much or as little as needed. Another group is using a model of the tasks, roles, domain, and interactions to automatically generate displays that are appropriate to the current user and situation.

Evaluation

An obvious approach to evaluating the results of MICA research is to score the end results of a team's performance in the challenge problem. For instance, how many SAM sites were destroyed? How many team members were lost? How many collateral objects were hit? This technique is commonly used because the results are easily quantified. This is in fact one of the evaluations that will be used in the MICA program.

Because humans will be integral to the control process, we can also employ cognitive engineering techniques to analyze characteristics such as the workload of the operator, and their acceptance of the system. We also want to evaluate the effect of the operator on the control concepts of performance, stability, and robustness. Another interesting issue arises when trying to combine control engineering and cognitive engineering evaluation techniques. Control algorithms are usually tested by running large numbers of simulations—putting human subjects through thousands of simulation runs is not a viable option. Some combination of the two techniques will be need to be developed.

Besides these standard approaches to metrics and evaluation, an additional avenue we are considering is to apply evaluation ideas from the organizational/social science team performance literature (Brannick, Salas, and Prince 1997) to human-robot team evaluation. Specifically, can we evaluate the quality of the team's process in addition to the end results? This need is a common theme throughout team performance literature. Perhaps the team achieved all of its objectives, but only because of luck—perhaps the team failed, even though it performed perfectly. Developing metrics for team process is also important for finding ways to improve performance, not simply deciding whether the team succeeded or failed.

Researchers have generally compiled team process metrics into several broad categories of competencies that effective teams possess: knowledge, skills, attitudes, leadership, decision-making, etc. Because a portion of team process involves interacting with teammates, many of these competencies appear relevant to the issue of etiquette evaluation as well. For instance, what is the quality of communication within and between teams? What is the rate of correct information? What are the ratios of statement types made between members? How complex is the information communicated? In terms of etiquette design, does the etiquette improve or hinder *assertiveness*: that is, do team members confront ambiguities and uncertainties, make suggestions, maintain their position when challenged? Does the etiquette improve or hinder *feedback* and *backup* behaviors? (Feedback encourages action from another team member, backup proactively fixes the problem.)

Like etiquette, the study of team performance also concerns issues of trust and efficiency; perhaps knowledge from this field can also be applied to the study of human-computer etiquette.

Etiquette Questions

I've briefly described the MICA domain and some of the issues that we've been exploring. Below are a few additional brainstorming questions relevant to this workshop:

- Are there limitations in MICA's current approaches to addressing mixed-initiative control that an etiquette approach could improve upon?
- What would an etiquette for the MICA domain look like? Are there fundamental differences in this and etiquettes for other domains?
- Are there fundamental differences in the etiquette design process for this domain vs. other domains?
- Can we characterize the features of an etiquette and use this characterization to determine exactly how the etiquette is affecting human-computer performance? For instance, how is a MICA etiquette specifically affecting performance, stability, and robustness of the control system?
- Is it useful to consider designing etiquette into inter-agent interactions in the same way that we're thinking about human-computer etiquette? If so, how can this also benefit the larger human-team etiquette?

Contribution of Shared Knowledge to Etiquette

More generally, I am also particularly interested in what role shared knowledge can play in designing and implementing human-computer etiquette. By shared, I mean knowledge shared between the human and the computer, in forms that both can understand.

One area where shared knowledge could contribute is in managing the more dynamic aspects of human-computer etiquette. Knowing the rules of etiquette in a given situation is important, but also important is detecting changes or variations in a situation, learning new situations, dealing with breaches and exceptions, etc. If we were to implement a more dynamic etiquette, what types of situations should we represent, and at what level? How should we represent them? How can the computer detect a change to the current etiquette context? How can the computer best convey to users a change in its perceived situation, or reasons for behaving contrary to the current etiquette, and vice versa?

Working from a more naïve, informal meaning of etiquette as formalized polite behavior, another area where

shared knowledge might contribute is in providing a stronger foundation for politeness or novel etiquette. Assuming that a basis for polite behavior is anticipating the needs and goals of others you're interacting with, and assisting those needs (even occasionally at the expense of your own needs and goals), then shared knowledge between the human and computer can provide a foundation for predicting what kinds of information or direction the other will need next.

Acknowledgements

For more information about the MICA program, please see <http://www.darpa.mil/ixo/mica.asp>. The MICA Program Manager is LTC Sharon Heise. Major cognitive engineering researchers who have worked on the mixed-initiative aspects of MICA, including the techniques described above, are Aptima, BBN, Charles River Analytics, Honeywell, Iterativity, Micro Analysis and Design, and SIFT.

References

Brannick, M.T., E. Salas, & C. Prince, eds. 1997. *Team Measurement and Performance: Theory, Methods, and Applications*. Mahway, NJ: Lawrence Erlbaum Associates.

Murphy, Robin R. and Erika Rogers. 2001. Final Report for DARPA/NSF Study on Human-Robot Interaction. Available at <http://www.aic.nrl.navy.mil/hri/nsfdarpa/>.

Schultz, Alan C. and Lynne E. Parker, eds. 2002. *Multi-Robot Systems: From Swarms to Intelligent Automata—Proceedings from the 2002 NRL Workshop on Multi-Robot Systems*. Boston: Kluwer Academic Publishers.

Roth, Emilie M., Emily S. Patterson, and Randall J. Mumaw. 2001. Cognitive Engineering: Issues in User-Centered System Design. *Encyclopedia of Software Engineering*, 2nd Edition, J.J. Marciniak (Ed.). New York: Wiley-Interscience.