Chapter 12

COGNITIVELY-INSPIRED AGENTS AS TEAMMATES AND DECISION AIDS

John Yen
Pennsylvania State University, University Park, PA

Laura Strater
SA Technologies, Marietta, GA

Michael McNeese
Pennsylvania State University, University Park, PA

Xiaocong Fan
Pennsylvania State University, Erie, PA

Haydee Cuevas
SA Technologies, Marietta, GA

Sooyoung Oh
Pennsylvania State University, University Park, PA

Anthony Kim
Pennsylvania State University, University Park, PA

Dev Minotra
Pennsylvania State University, University Park, PA

Tim Hanratty
US Army Research Lab, Aberdeen Proving Ground, MD

INTRODUCTION

The advancement in sensor, communication, and information technologies have resulted in information-dense environments in which current and future warfighters must operate. In the face of dynamic, constantly changing events and conflicting reports, this vast amount of information makes it difficult for warfighters to develop and maintain a clear picture of the operational situation. This difficulty is further compounded by the challenges of asymmetric warfare where complex decisions that consider multi-dimensional factors (e.g., human, social, and cultural factors) are required not only for strategic planning, but also at the operational and tactical level, across all echelons.

To address these challenges, this research, supported by the Advanced Decision Architectures Collaborative Technology Alliance of the U.S. Army Research Lab, aims to achieve two complimentary research goals. Our first research objective is to develop a cognitively-inspired decision agent architecture that supports the delivery of relevant information to decision-makers in a timely fashion. These software agents act as both teammates and decision aids, enhance the human's capability in reasoning across multiple decision spaces, and enable dynamic team collaboration among decision makers. With its cognitive foundations, functions and components in such an architecture can be better related to the processes of human cognition. Such a relationship between computational decision aids and human decision process can contribute to the co-refinement of both the computational decision architecture and the human-agent team cognition theory.

Our second research objective is to study factors that affect human-agent interactions such that warfighters can better calibrate their automation usage and maintain global situation awareness (SA). Despite the best intentions of automation designers, the current state of technology does not produce perfect automation. On the battlefield, for example, automation, like humans, must function in the face of uncertainty and in sub-optimal environmental conditions. Warfighters, trained to operate in teams, understand that these complexities may affect the judgments of their human teammates, but sometimes fail to realize that agent teammates can be similarly affected. When warfighters fail to adequately understand the factors influencing the performance of their agent teammates, they may make poor automation usage decisions (AUDs) and may fail to appropriately trust their agent teammates. Previous research has shown that trust in team members, whether human or agent, is critical in mediating team operations, particularly with increasing levels of cognitive complexity.

The objective of our cognitively-inspired agent system (R-CAST) is to serve as decision aids and teammates of human decision makers. There are many normative, descriptive, and prescriptive decision theories that can be built within an agent system. We chose to adopt Klein's RPD model (Klein, 1997) in R-CAST for several reasons. First, the RPD model is a holistic model of human decision making processes, including activities such as seeking relevant missing information, and monitoring expected outcomes of decisions to detect anomaly. Second, operators can more easily understand a decision aid based on a naturalistic decision process (e.g., macro-cognitive process described in RPD) due to his/her familiarity and experience with the process. This encourages active human-agent collaboration along the decision process. Such a decision aid would be intelligent not only in the evaluation and choice of options, but also in social interactions with human decision makers.

With a computational RPD model as the macro-cognitive process, the R-CAST system has been used in a series of experimental studies toward a better understanding of human-agent collaboration in time-stressed decision making situations. The first experiment was conducted to help us to understand whether future warfighters can benefit from R-CAST in handling complex multi-dimensional decision tasks. Being a teammate of a human decision maker, the R-CAST system ought to be trustable and understandable. The second and third experiments examined the issue of human-agent trust and raised the interesting question of how to facilitate a suitable level of trust between warfighters and decision aids to improve the overall performance of the human-agent team. The fourth experiment was conducted to evaluate the impact of providing a mental map visualization of the agent's decision space to promote automation transparency on effective human automation-usage decisions.

The rest of the chapter is organized as follows. In section 2, we describe related research regarding human-agent team cognition and its relationship to human trust in agents. In section 3, we describe the R-CAST agent architecture, which is empowered by a computational RPD model. In section 4 we give an overview of a synthetic task designed for studying multi-dimensional decision making. In section 5, we summarize a series of human-in-the-loop experiments where the R-CAST agents served as teammates and decision aids, investigating the issue of multi-context decision making, human-agent trust, and mental map visualization. Finally, we discuss the impacts of our work and point to some directions for future studies.

HUMAN-AGENT TEAM COGNITION AND TRUST

Team cognition is traditionally involved with studying and specifying the cognitive processes that impact team performance and that vary with certain constraints. Team cognition requires that team members maintain common ground, situation awareness, interdependence, and flexibility as they pursue joint objectives. Team cognition embodies the notion that teamwork is influenced by individual differences among team members as well as the contextual environment within which a team operates. Teamwork may also be influenced by many different social-cognitive factors (McNeese 2000) or social psychological states (Wellens & McNeese, 1987) that change how teams work and perform together. For example, social psychological variables such as trust, affiliation, and status may impact the specific nature of team cognition. In particular, this chapter focuses on investigating factors that influence the interactions between human and agent team members when agents are utilized to support team cognition during a complex decision making task. In particular, we want to better understand how human operators' trust in automation may potentially influence their subsequent automation usage decisions (AUDs). Many of the factors that influence human-human team collaboration also influence human-agent collaboration.

Trust can operationally be defined not just as an emotional state or characteristic of the trustor, but rather by the behaviors and actions taken by the trustor as the result of that characteristic. This distinction is important in the current research effort as we focus on the calibration of AUDs, a behavior that reflects the user's feeling of trust. Specifically, we seek to investigate the impact of providing users with information about agent process that impacts predictability, the characteristic of the agent trustee that most influences trust in the early stages of interaction with the agent.

Human operators' beliefs about and trust in automation mediate their subsequent reliance on automated systems (such as intelligent agents), ranging from the extremes of over-reliance and complacency to under-reliance and mistrust (Cuevas, Fiore, Caldwell, & Strater, 2007; Lee & See, 2004). This can lead to at least two potentially problematic situations, misuse and disuse (Parasuraman & Riley, 1997). In misuse (over-reliance), the operator blindly follows the judgments made by the automation, thereby abdicating their role of system supervisor; while with disuse (under-reliance), the operator either ignores the automation's recommendations or delays action until system judgments can be verified, increasing decision time. While the source for these errors differs, either type of error can have critical consequences; thus, the goal of properly calibrating operators' trust in automation and their AUDs seeks to reduce both misuse and disuse of automation.

Given the broad range of factors influencing trust in automation and AUDs, it is hardly surprising that research findings have been mixed (see Beck, Dzindolet, & Pierce, 2002; Dzindolet et al. 2003; Lee & See, 2004). Despite the abundant research on the topic, little practical guidance is available to assist system designers in developing automated systems that support the operator in making appropriate AUDs, particularly in cases when optimal usage is a 'moving target,' changing with environmental conditions, workload, and other situational parameters.

Therefore, the long term goal of this research effort is to determine techniques for calibrating AUDs by enabling users to gain adequate understanding of agent functioning to allow for appropriate levels of trust in their automated teammate.

RPD-ENABLED AGENTS: THE R-CAST SYSTEM

The R-CAST agent architecture (Fan et. al., 2005a; Fan & Yen 2007) is built on top of the concept of shared mental models (Cannon-Bowers, Salas, & Converse 1990), the theory of proactive information delivery in agent teamwork (Cohen and Levesque 1991; Fan, Yen, & Volz 2005), and the recognition-primed decision framework (Klein 1997). Figure 12.1 shows the major components of R-CAST, where the RPD process is the kernel of the other functional units. Klein's RPD model describes a holistic process of human decision making, including all activities related to human decision makings. These activities include identifying missing relevant information, seeking information, interpreting information, building hypothesis, situation monitoring for detecting anomalies, and decision adaption under unexpected situations. These activities are realized in the R-CAST agent architecture such that (1) each agent can perform these functions individually, and (2) a group of agents can collaborate on these functions as a team.

Like the RPD model, R-CAST agents match the current situations with previous "experience" to determine their "similarity." This matching process is implemented in two ways in R-CAST: (1) choose the first experience with a similarity higher than a threshold, and (2) choose the experience with the highest similarity. While the first approach is a computational realization of the satisficing criteria of the RPD model, the second approach offers an alternative that is closer to rationale decision-making. This flexibility enables the developer of agent applications to choose the suitable scheme, depending on characteristics of the application.

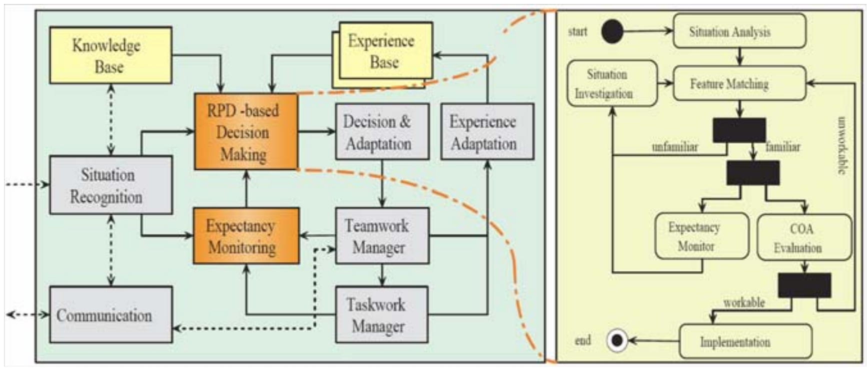

Figure 12.1
The R-CAST Agent Architecture (Fan & Yen, 2007)

The R-CAST agents realize the "seeking missing information" component in the RPD model using a backward (Prolog-style) chaining inference, starting with high-level cues of the experience being matched, using rules in the agent's knowledge base. When a piece of missing information is identified through this backward chaining inference, the agent contacts the information source associated with the missing information through agent communication protocols realized in Service Oriented Architecture (SOA). The information source can be another agent, a Web service, or a human operator. When the requested information arrives, the agent performs forward-chaining inference that fuses the newly arrived (previously missing) information with other information know to the agent. The outcomes of these fusions affects the degree a "decision experience" matches the current situation. Hence, this fusion within an agent typically corresponds to high-level information fusion (i.e., level 2 or above in the JDL data fusion model).

The R-CAST agent architecture also implements decision progress monitoring and expectancy-based decision adaptation in the RPD model. After a decision is made (either directly by an agent or indirectly by a human operator based on the agent recommendation), the agent continues to monitors the expected outcomes (called "expectancy" in the RPD model) associated with the chosen decision. If the agent detects an anomaly (e.g., some expected outcomes not fulfilled), the agent adapts the previously made decisions to deal with the anomaly.

Finally, a team of R-CAST agents realizes a collaborative RPD process (Fan, Sun, McNeese, & Yen, 2005), in which each agent anticipates information requirements of teammates and can proactively assist them by monitoring and/or seeking information relevant to their requirements.

Together, the features described above form a cognitively-inspired agent architecture that not only supports warfighters in making decisions and adapting them to changing conditions, but also in augmenting them in their capabilities to sense, fuse, and interpret information in a multi-dimensional net-centric environment .

A SYNTHETIC TASK FOR MULTI-DIMENSIONAL DECISION MAKING: THE THREE-BLOCK CHALLENGE

To conduct experiments regarding multi-dimensional decision-making of warfighters, we have designed and implemented a simulation environment called "The Three-Block Challenge." It can simulate command and control scenarios of urban operations where in close proximity (e.g., within three blocks) military operators need to quickly react to challenges related to three dimensions of mission: the humanitarian relief dimension, the peacemaking (i.e., policing) dimension, and the combat dimension (Fan et. al., 2005).

The synthetic task environment contains objects of interest such as main supply routes (MSRs) and key buildings (religious buildings, schools, and hospitals). At run time, the environment can produce three types of threats: Improvised Explosive Device (IEDs), crowds, and insurgents, which represent the targets of humanitarian, peacekeeping, and combat operations, respectively. IEDs are motionless targets, and if exploded, can cause damage to the road (e.g., MSRs) and buildings nearby. A crowd represents a group of people which may contain activists that can be friends or foes. A crowd can be of medium (M) or large (L) size, and the group size of a crowd can change over time. Two crowds can merge together if they move close enough. Another type of movable targets is insurgents, each is associated with a threat level that can be L(low), M(medium), or H(high). A target may appear, stay on, and disappear from the battle field following certain temporal and spatial patterns unknown to human operators.

There are also a limited number of friendly units—squads and Explosive Ordnance Disposal (EOD) teams—under the control of a C2 team. Each friendly unit has an associated property called "combat readiness," which is represented by a percentage value, indicating current unit ability to handle threats. The readiness value decreases after a unit is tasked to a threat, and can recover incrementally as time passes.

In our studies, a C2 team consists of an S2 suite (intelligence cell) and an S3 suite (operations cell). The roles of C2 operators have been simplified. S2 is responsible for processing incoming reports, called Spot reports; collecting relevant information from other sources; and alerting S3 of potential threats. S3 needs to process alerts from S2, and make decisions on which target to handle next and which resources (friendly units) to allocate toward that target.

Table 12.1 Credit value and resource requirements for handling targets

<table>
<tr><th colspan="3">Targets</th><th>Value</th><th>Res. Req.</th><th>Action</th></tr>
<tr><td rowspan="4">Crowd</td><td>M</td><td>w/o foe</td><td>20</td><td>1U</td><td>monitor</td></tr>
<tr><td>M</td><td>w foe</td><td>40(+10)*</td><td>2U</td><td>disperse</td></tr>
<tr><td>L</td><td>w/o foe</td><td>40(+10)*</td><td>2U</td><td>disperse</td></tr>
<tr><td>L</td><td>w foe</td><td>50(+10)*</td><td>3U</td><td>disperse</td></tr>
<tr><td colspan="3" rowspan="2">Insurgent
(3 threat levels: L, M, H)</td><td colspan="3">n=1,2,3 for L,M,H</td></tr>
<tr><td>50+50n</td><td>(n+1)U</td><td>capture</td></tr>
<tr><td colspan="3">IED</td><td>60(+20)*</td><td>1U + 1E</td><td>remove</td></tr>
</table>

‘U’ refers to “squad unit”, ‘E’ refers to EOD team.

**additional credit value when a target is near an MSR.*

Decision making in target selection and resource allocation requires the S3 suite to consider trade-offs among multiple factors: target type, threat level, the combat readiness of the available units, the unit-target distance, and how long a target has been on the field. The type and threat level of a target determine how many friendly units will be needed to handle the target. Table 12.1 lists for each type of target the credit value (the reward points a C2 team can get if a target is handled successfully), the number of resources required to handle a target, and what action S3 should take. For example, the second entry says that dispersion of a medium-sized crowd with a foe needs two squad units, and 40 points can be credited if the crowd is dispersed successfully. The last entry indicates that one squad unit and one EOD team are required to remove an IED. If an IED is removed successful, 60 points can be credited if the IED is close to buildings only or MSRs only, 80 points if it is close to both.

HUMAN-AGENT EXPERIMENTS

Toward our goal of developing theory and technologies for supporting human-agent decision-making teams using cognitively-inspired agents as teammates and decision aids, we have conducted a series of human-in-the-loop experiments to investigate issues regarding multi-context decision making, human-agent trust, and mental map visualization.

We next briefly describe an example setting for the experiments (other settings vary slightly, with or without R-CAST agents playing a role in S2 or S3 suites). For each study, we then summarize the research questions and our main findings.

A Setting with R-CAST Agents in S3 Suite

The role of S3 suite is played by an R-CAST agent (S3 agent) and a human operator. The human operator is equipped with two monitors: a map display for tracking situation development, and a tasking interface for collaborating with S3 agent to handle threats.

The map display shows all the active entities on the field, as well as MSR's, buildings, and regional boundaries. It allows a human operator to highlight the target of interest, to figure out the spatial relationships between targets, MSRs, buildings, and friendly units, and to project forward the location of moving targets. The human operator needs to determine whether there are key buildings (e.g., hospital, mosque) near a target when assigning units to the target. This is one way the social dimension is introduced into the simulated C2 environment.

The tasking interface consists of a threat table, a command panel, and a feedback display panel. The threats table shows consolidated information about the threats on the field: for each threat, it lists threat type, ID, status, crowd size, activists associated with a crowd, nearby buildings, priority, elapsed time, and the IDs of the tasks associated with prosecuting the threats.

The command panel allows the S3 human operator to physically task units to the selected target. The feedback display panel shows some statistics about tasks issued, threats cleared, and reward points earned. In the experiment, the S3 agent offers decision aids and recommendations to the S3 operator regarding units to be assigned to a target selected by the user. It is, however, the S3 operator who has the final authority for decisions on target selection and resource allocation.

Experiment 1: Supporting Multiple-Context Decision Making

In this study (Fan et. al., 2006), we focused on the challenge of switching between multiple types of decision contexts. We chose this cognitive challenge for two reasons. First, this is an important characteristic of the multi-dimensionality of the warfighter's mission. They need to deal with multiple missions, which involve different types of threats. For instance, a warfighter may need to deal with a combat mission (capturing key insurgents), police mission (controlling crowds), and humanitarian mission (making sure food, medicine, and supplies are delivered safely along logistic routes) at the same time. Second, decision-making involving different types of contexts introduces additional cognitive load, due to the need to change focus between these contexts.

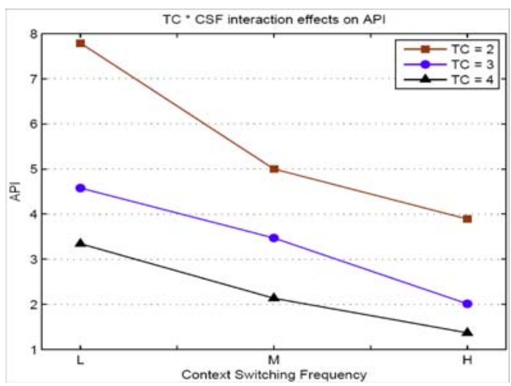

Figure 12.2

C2 team performance varies with context switching frequencies

Two sets of experiments, each with various settings of context switching frequencies and tasking complexities, were conducted. The participants were US Army Reserve Officer Training Corps (ROTC) students at Penn State. Figure 12.2 shows one of the experiment results. Overall, the study demonstrated that C2 team performance, while still limited by human cognitive capacity, could be largely improved when they were assisted by R-CAST agents capable of proactive information gathering/sharing and experience-based decision making. It also suggests that higher demand situations require more competent teammates. The experiments represent an important step forward in uncovering the nature of real-world problems that are highly relevant to the vision of the US Army.

Experiment 2: Human Trust on Cognitive Aids

As we mentioned in previous sections, human trust in agents is an important factor that affects the calibration of his/her automation usage decisions (AUDs). To improve our understanding of this phenomenon, we introduced error into the agent recommendations: the agents did not consider unit combat readiness. In this experiment, we focused on the relationship between S3 operator and S3 agent, with the S2 function simulated by an R-CAST agent without an S2 operator. Our research question is whether knowledge about the source of agent error affects human trust in the agent. Sixty command and control teams, each consisting of a human operator and two intelligent agents, allocated resources to targets in simulated urban combat scenarios. We used a mixed 2 × 2 × 2 × 3 factorial treatment design (PG × KAR × TC × RIT), where TC (task complexity) and RIT (ratio of insurgent threats) are within-subjects variables and KAR (knowledge of agent reliability/error) and PG (participant group: ROTC students vs IST students) are between-subjects variables. The experiment group was informed about the agent's source of error, whereas the control group was not informed about the source of errors. Both groups were informed about the same set of game rules and the agent's reliability measure.

The ANOVA output indicates that both the knowledge of agent reliability (KAR) and ratio of insurgent threats (RIT) had significant effects on C2 performance. Figure 12.3 shows that knowing the agent reliability helped the S3 operators rectify more inappropriate recommendations. Together with the analysis of the SAGAT and NASA-TLX measures, this study indicated that giving even a minimal basis (i.e., knowledge about the source of error from agents) for understanding conditions impacting agent reliability allowed operators to make better automation usage decisions, have better SA on the critical issues associated with automation error, and establish better trust in the intelligent agents.

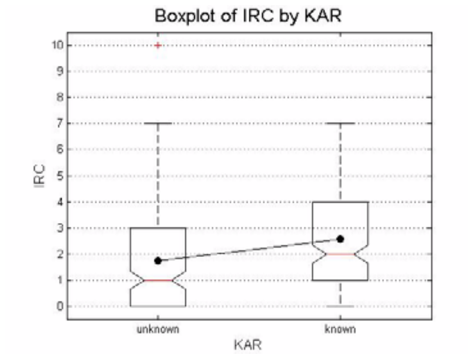

Figure 12.3
Incorrect agent recommendation corrected

This result is different from previous research regarding automation usage decisions (Dzindolet et. al. 2003) that involves randomly generated errors. This motivated us to investigate the impact of predictable versus random error on human AUDs.

Experiment 3: Agent Error Patterns and Human Trust Calibration

Our research question in this study is, are AUDs different in interactions with random agent error as compared to predictable error (that is, when error patterns are present)? To investigate this, we designed a study with both a random error condition and a systematic or predictable error condition. To ensure a fair comparison, the error rate in the two conditions were comparable, though errors were actually lower in the random error condition. The results, shown in Figure 12.4, indicates that operators teamed with agents with systematic errors that the user can make sense of make more appropriate AUDs than operators teamed with agents that make random errors.

Based on the findings of Experiment 2 and Experiment 3, we proposed a model about human-agent trust, in which the decision maker's attempt to make sense of the error patterns play an important role in his/her decision on accepting an agent recommendation or not (Strater et al, in review). Using this model, we can explain the finding described above as follows. With agent errors that the decision maker cannot make sense of (e.g., random errors), the decision maker cannot distinguish reliable recommendations from erroneous ones and thus cannot properly calibrate AUDs. For errors that the decision maker can make sense of (e.g., errors due to ignorance about combat readiness criteria), understanding error patterns allows operators to distinguish reliable recommendations from those that are not reliable and helps calibrate AUDs.

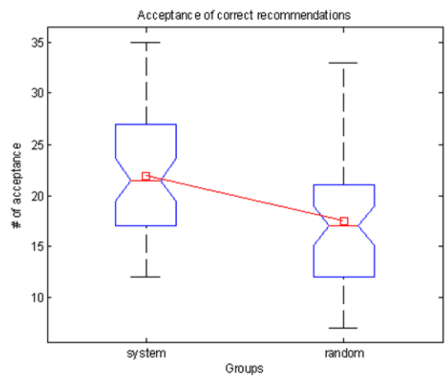

Figure 12.4

Acceptance of correct agent recommendations from agents with systematic errors vs from agents with random errors

Experiment 4: Visualizing Mental Map of RPD Agents

In order to serve as effective teammates, enhancing shared SA and stimulating the user's active participation, R-CAST agents' need the ability to expose their internal states and reasoning process to the human user. Towards this end, we implemented a visualization concept into R-CAST and conducted an experiment to examine the influence the visualization concept (information presentation) had on effective AUDs and situation awareness.

Specifically, we designed a visualization concept that helps establish a mental map of the agent's decision space while promoting transparency of the underlying agent recommendation process and higher-level SA (comprehension and projection). In this visualization, we reduce the dimension of the experience space of an RPD-agent onto a 2-D display and dynamically position the current agent-state onto the display based on shared similarity measures. Figure 12.5 shows an example of the R-CAST Visualization of the Agent Decision Space (VADS). The VADS maps a collection of past *experiences* or Common Historical Cases (CHCs) and current target icons based on their relative similarities. In addition to temporally repositioning active targets on the VADS, the conveyance of information about a target or a related CHC is augmented with the use of iconic symbols.

To assess the effects of the visualization concept, we modified the Three-Block Challenge synthetic task environment to emphasize the importance of crowd control. The study employed a 2 (visualization mode) x 2 (work load) x 2 (task complexity) mixed design. The between-group factor was the mode of visualization (experimental group utilizing mental map visualization vs. control group utilizing tradition table visualization). The two within-group factors were the scenarios' workload (5 crowds vs 10 crowds) and level of task complexity (ratio of fast-burning vs slow-burning crowd movement). The dependent variables included task performance, situation awareness, trust in automation, and subjective workload. Thirty two ROTC students from Penn State participated in the experiment, which included four 10-minute scenarios for each subject.

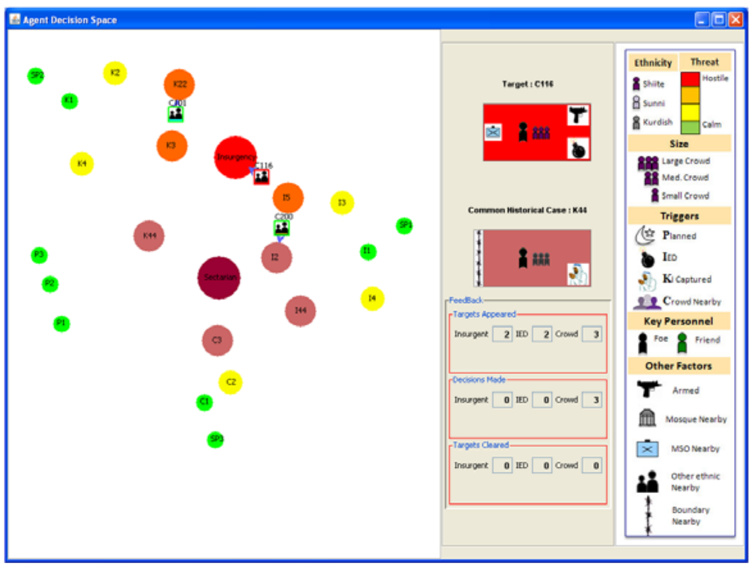

Figure 12.5
R-CAST Visualization of Agent Decision Space (VADS)

Overall, the preliminary results have revealed improvement in both task performance and SA. One of the measures calculated for task performance was the product of Score and Real-time SA. The un-normalized result for this ANOVA measure showed the experimental group scored an average 20 percent improvement over the control group. Detailed results of the experiments will be reported in a separate paper.

CONCLUSION AND IMPLICATIONS FOR THE FUTURE ARMY

The ARL ADA CTA research initiated throughout the past five years shows significant results that inform and have impact within several key areas: cognitive science, agent architectures, human-computer interaction, and automation and trust.

Looking at the area of cognitive science first, much of the results obtained are directly informative and applicable for how a command, control, and communications (C3) team addresses a very complex, evolving C3 mission to keep stability ongoing in the Iraqi warfront. The mission and scenarios incorporated into the experiments are important as it represents how team members must adaptively assign and adapt resources for Iraq multi-target, multi-faceted C3 mission. As this mission and scenarios were developed through knowledge elicitation with experts, it is both contemporary and representative of dynamic cognitive processes wherein team members must analyze changes and adapt their resources accordingly under periods of time pressure. Many studies in cognitive science utilize more toy domains with static cognitive processes, and do not involve the levels of complexity and replanning while retaining necessitated balance across ongoing multi-layered mission objectives that often have to be traded off. The experimental studies conducted also incorporate two other integrated aspects of cognitive science: 1) adaptive use of automation given understanding of trust in the agent, and 2) measuring situation awareness in addition to human performance. By manipulating elements of trust we have been able to determine how trusts develops based on what human believe the automation is capable of doing (e.g., how reliable it is). This is extremely valuable for designing and coupling humans with intelligent agent architectures in a way wherein the joint interaction can address dynamic cognitive states that arise. Because previous studies typically only use agents that are not evaluated in terms of trust with human, and utilize agents that do not employ recognition-primed decision making strategies, previous studies often only address static cognitive states. With the complexities in contemporary teamwork and advanced missions, our research is directly applicable. Because our studies look at situation awareness in addition to direct human performance measures, a deep understanding of how cognitive processes interact with automation is possible. The impact of these elements is significant because it provides a baseline for understanding how agents will need to be designed to fit human cognitive capacities, while at the same time providing a baseline on how human trust develops with differing information about the agents. Together these elements advance cognitive science and human-agent interaction in innovative ways that afford potential increases in the overall war-fighter mission.

From the perspective of agent architecture, this research has resulted in a comprehensive cognitively-inspired agent architecture that is designed to serve as both a decision aid and a teammate for warfighters. Based on Klein's Recognition-Primed Decision (RPD) model, the RPD agent provides a computational framework regarding the holistic human decision making process, which can assist warfighters in making decisions across multiple contexts/missions, collaborative sensing and seeking information relevant to the needs of warfighters, fusing information for situation understanding in a distributed net-centric environment, and detecting changes that require adaptation of decisions previously made.

From human-computer interaction perspective, one of the principal challenges in supporting close human-agent collaboration is increasing automation transparency to align the decision maker's understanding of the decision space (mental model) with that of the intelligent software agents working on their behalf. Enabling technologies that allow the conveyance of information (complex insights, experiences, and high-level concepts) and its correct application is critical. One method examined in this research to improve human-agent understanding is information visualization. In this research we examined a visualization concept designed to enhance the decision maker's perception, comprehension, and projection of the underlying knowledge space while improving shared human-agent SA. Allowing a holistic view of the agent decision space provides a transparent view not only of the agent decision process but equally important, encourages active participation from the user—allowing for the adjustment of contextual constraint, the ability to prioritize missing information and ultimately, improved decision making.

In summary, the research described in this chapter contributes to the technology for designing cognitively-inspired agents as well as our understanding about the principles regarding the macro-cognitive processes of human-agent team cognition and decision making, especially on the issue of trust in automation and the resultant automation usage decisions (AUDs). This understanding, in turn, can be adopted to enhance the design of intelligent agents. Hence, the design of cognitively-inspired agents, the findings of human-centric experiments, and the theory of human-agent team cognition form a cycle of synergistic research roadmap, which drive human/social science and computational/information technology forward toward the vision of equipping warfighters with cognitively-inspired agents as teammates and decision aids. From this, suitable levels of trust can be developed to enable effective human-agent team performance for complex multi-facet decision making across strategic, operational, and tactical missions.

ACKNOWLEDGEMENTS

We wish to thank Army Cadets and Officers at Penn State University for their participations in all of the experiments described in this chapter. This research has also benefited from many fruitful discussions with Dr. Laurel Allender at Army Research Lab, Dr. Linda Pierce at Army Research Institute, and Dr. Frank Ritter at Penn State University.

REFERENCES

Beck, H. P., Dzindolet, M. T., & Pierce, L. G. 2002. Operator's automation usage decisions and the sources of misuse and disuse. In E. Salas (Ed.), Advances in human performance and cognitive engineering research (Vol. 2) (pp. 37-78). Bingley, UK: Emerald Group.

Cannon-Bowers, J. A.; Salas, E.; and Converse, S. 1990. Cognitive psychology and team training: Training shared mental models and complex systems. *Human Factors Society Bulletin* 33:1–4.

Cohen, P. R., and Levesque, H. J. 1991. Teamwork. Nous 25(4):487–512.

Cuevas, H. M., Fiore, S. M., Caldwell, B. S., & Strater. L. 2007. Augmenting team cognition in human-automation teams performing in complex operational environments. Aviation, Space, and Environmental Medicine, 78 (5, Supp Section II), B63-70.

Deutsch, M. 1960. The effect of motivational orientation upon trust and suspicion. Human Relations 13:123–139.

Dzindolet, M. T.; Peterson, S. A.; Pomranky, R. A.; Pierce, L. G.; and Beck, H. P. 2003. The role of trust in automation reliance. *International Journal of Human-Computer Studies* 58:697–719.

Endsley, M. 1995. Toward a theory of situation awareness in dynamic systems. *Human Factors* 37:32–64.

Fan, X; Oh, S.; McNeese, M.; Yen, J.; Cuevas, H.; Strater, L.; and Endsley, M. R. 2008. The Influence of Agent Reliability on Trust in Human-Agent Collaboration. In *Proc. of the European Conference on Cognitive Ergonomics (ECCE'08)*, 2008.

Fan, X., Sun, S., McNeese, M., and Yen, J., 2005. Extending Recognition-Primed Decision Model For Human-Agent Collaboration, In Proc. AAMAS'05, pp. 945-952, The Netherlands.

Fan, X.; Sun, B.; Sun, S.; McNeese, M.; Yen, J.; Jones, R.; Hanratty, T.; and Allender, L.. 2006. RPD-Enabled Agents Teaming with Humans for Multi-Context Decision Making. In *Proc. AAMAS'06*, pages 34--41, Japan.

Fan, X., and Yen, J. 2007. R-CAST: Integrating team intelligence for human-centered teamwork. In *Proceedings of the Twenty-Second National Conference on Artificial Intelligence (AAAI'07)*, 1535–1541.

Fan, X.; Yen, J.; and Volz, R. A. 2005. A theoretical framework on proactive information exchange in agent teamwork. *Artificial Intelligence* 169:23–97.

Kaber, D. B., and Endsley, M. R. 2004. The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. Theoretical Issues in Ergonomics Science 5(2):113–153.

Klein, G. A. 1997. The recognition-primed decision (rpd) model: Looking back, looking forward. In Naturalistic decision making (Eds: C. E. Zsambok and G. Klein). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. 285–292.

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

McNeese, M. D. 2000. Socio-cognitive factors in the acquisition and transfer of knowledge. Cognition, Technology and Work, 2, 164-177.

Norling, E. 2004. Folk psychology for human modelling: Extending the BDI paradigm. In AAMAS ’04: International Conference on Autonomous Agents and Multi Agent Systems, 202–209.

Parasuraman, R., and Riley, V. 1997. Humans and automation: Use, misuse, disuse, abuse. Human Factors 39(2):230–2

Strater, L., McNeese, M., Cuevas, H. M., Fan, X., Oh, S., and Yen, J. Making Sense of Error Patterns: Toward Appropriate Human-Agent Trust, in review.

Wellens, A. R., & McNeese, M. D. 1987. A research agenda for the social psychology of intelligent machines. Proceedings of the IEEE National Aerospace and Electronics Conference (pp. 944-949). Dayton, OH: IEEE Aerospace and Electronics Systems Society.