

Using Cognitive Architectures to Study Issues in Team Cognition in a Complex Task Environment

Paul R. Smart^a, Katia Sycara^b and Yuqing Tang^b

^aElectronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK;

^bRobotics Institute, School of Computer Science, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213-3890, US

ABSTRACT

Cognitive social simulation is a computer simulation technique that aims to improve our understanding of the dynamics of socially-situated and socially-distributed cognition. This makes cognitive social simulation techniques particularly appealing as a means to undertake experiments into team cognition. The current paper reports on the results of an ongoing effort to develop a cognitive social simulation capability that can be used to undertake studies into team cognition using the ACT-R cognitive architecture. This capability is intended to support simulation experiments using a team-based problem solving task, which has been used to explore the effect of different organizational environments on collective problem solving performance. The functionality of the ACT-R-based cognitive social simulation capability is presented and a number of areas of future development work are outlined. The paper also describes the motivation for adopting cognitive architectures in the context of social simulation experiments and presents a number of research areas where cognitive social simulation may be useful in developing a better understanding of the dynamics of team cognition. These include the use of cognitive social simulation to study the role of cognitive processes in determining aspects of communicative behavior, as well as the impact of communicative behavior on the shaping of task-relevant cognitive processes (e.g., the social shaping of individual and collective memory as a result of communicative exchanges). We suggest that the ability to perform cognitive social simulation experiments in these areas will help to elucidate some of the complex interactions that exist between cognitive, social, technological and informational factors in the context of team-based problem-solving activities.

Keywords: collective cognition, social information processing, distributed cognition, team cognition, social simulation, group performance, cognitive modeling, cognitive architecture, ACT-R

1. INTRODUCTION

The advent of global information and communication networks, such as the World Wide Web, has led to a recent explosion of interest concerning the socially-distributed nature of human cognition. As it becomes increasingly possible to harness the cognitive processing capabilities of multiple individuals within a distributed, networked environment, so attention has focused on the factors that enable groups to function as the processors of information.¹ This burgeoning research interest is particularly apparent in the context of the nascent discipline of Web Science,² with notions such as collective intelligence,³ augmented social cognition,⁴ social machines,⁵ the global brain,⁶ human computation,⁷ social computation⁸ and technology-mediated social participation systems⁹ serving as the just some of the conceptual anchors for ongoing research efforts.

Aside from the research into Web-based systems, research concerning socially-distributed cognition in more restricted organizational contexts is also of increasing interest. A particular focus of attention is what has been called ‘team cognition’,^{10,11} which refers to the “cognitive processes or activities that occur at a team level” [11, p. 256]. Team cognition recognizes that many cases of real-world problem solving and decision making rely on the coordinated interaction of multiple individuals, often working within a distributed, network-enabled environment that provides access to a range of shared information resources and collaborative technology solutions. Given that teams often play a central role in enabling organizations to make sense of their environment and act accordingly,

Further author information: (Send correspondence to Paul R. Smart)

Paul R. Smart: E-mail: ps02v@ecs.soton.ac.uk

it is imperative that we develop a better understanding of the factors that influence the dynamics of team-level cognitive processing.

The advent of modern sensor and networking technologies has clearly provided new opportunities to harness the collective cognitive potential of groups of individuals; however, it is not always clear how we should engineer technological and social environments to support socially-distributed information processing. When it comes to collective or team sensemaking,^{12,13} for example, should the individual team members be allowed to engage in frequent communication with one another in order to coordinate their interpretations of an object, event or situation, or should more restrictive communication policies be enforced? While it may be natural to assume that full connectivity and frequent communication are desirable features of a networked environment – and this is certainly consistent with the technological trend towards communication networks of ever greater reliability and bandwidth – research within the social psychological community suggests that precipitant forms of information sharing can sometimes lead to deficits in group performance.¹⁴ Cognitive biases are also a major concern when it comes to team cognition, especially since cognitive biases are sometimes more extreme at the collective, as opposed to the individual, level.¹

One approach to the study of team cognition is to rely on multi-agent simulation techniques (sometimes referred to as social simulations). By using such techniques, the profile of inter-agent communication and the dynamics of information exchange can be systematically manipulated in order to observe their effects on collective cognitive outcomes. While such techniques have proven useful in investigating a number of social psychological phenomena, most notably social influence,¹⁵ they have sometimes been criticized in terms of their cognitive sophistication and fidelity. Sun,¹⁶ for example, argues that we should move towards the use of what are called cognitive architectures in performing ‘cognitive social simulations’, i.e., social simulations in which the individual agents are implemented as computational cognitive models that draw on the representational and computational resources of a particular cognitive architecture. Cognitive architectures, in this case, are frameworks that make particular commitments about the kinds of mental representations and computational procedures that are sufficient to explain important aspects of human cognition, such as problem solving, memory and learning.¹⁷ A cognitive architecture serves as a framework for the development of sophisticated computational cognitive models that can then be incorporated into multi-agent social simulations. This supports investigations into the complex interactions that exist between social and cognitive factors in the context of team-based tasks, and it provides insights into the sorts of human factors interventions that might be required to improve team performance.

The current paper describes an ongoing effort to develop a cognitive social simulation capability based around a particular cognitive architecture, called Adaptive Control of Thought-Rational (ACT-R).^{18,19} The specific aim of our research is to develop a generic capability for running cognitive social simulation experiments and then apply this capability to a specific task. The capability we aim to develop is hereafter referred to as the ACT-R Cognitive Social Simulation Capability, or ACT-R/CSSC, and the task that is the focus of our current modeling and simulation efforts is a task that has previously been used to investigate socially-distributed problem solving as part of what is called the ELICIT experimentation framework. ELICIT, in this case, is an acronym that stands for the Experimental Laboratory for Investigating Collaboration, Information Sharing and Trust. It represents a sustained effort to advance our understanding of the factors that affect collective performance in task contexts that incorporate both cognitive and social elements. The task that is commonly used as part of the ELICIT experimentation framework is called the ELICIT task. This is a task that requires individuals to reason over bodies of task-relevant information and to also share information with others in the context of a distributed, network-enabled task environment. The ELICIT task is thus one that draws on factors that are spread across the cognitive, social, informational and technological domains. These features make the task suitable for testing the ability of the ACT-R/CSSC to support experimental studies of team-based problem solving behavior.

The main aim of the current paper is to outline the motivation for adopting cognitive architectures in the context of social simulation experiments. We also describe the approach adopted with respect to the development of the ACT-R/CSSC and outline some focus areas for future work with the simulation capability.

2. COGNITIVE ARCHITECTURES AND COGNITIVE SOCIAL SIMULATION

Sun¹⁶ defines a cognitive architecture as “a domain-generic computational cognitive model that captures essential structures and processes of the individual mind for the purpose of a broad (multiple domain) analysis of cognition and behaviour” (p. 33). A cognitive architecture is thus a framework that captures some of the relatively invariant features of the human cognitive system – those features that are deemed to be more-or-less constant across domains, tasks and individuals. One example here is the mechanisms that support the storage and retrieval of information from long-term memory. Although a number of features of the task environment may affect the ability of subjects to recall information (see Section 3.3), the mechanisms that actually realize the recall process are unlikely to change from one task to another*.

A number of cognitive architectures have been developed within the cognitive scientific community to support the computational modeling of cognitive behavior. These include ACT-R,^{18,19} State, Operator and Result (SOAR)^{23–25} and Connectionist Learning with Adaptive Rule Induction ON-line (CLARION).^{16,26} Among these, ACT-R is probably the one that has received the most research attention. The ACT-R architecture has been used to model human cognitive performance in a wide variety of experimental contexts[†], and it has generated findings of predictive and explanatory relevance to hundreds of phenomena encountered in the cognitive psychology and human factors literature. This has earned it a reputation as the cognitive architecture that is probably the “best grounded in the experimental research literature” [28, p. 24].

As mentioned in Section 1, the use of cognitive architectures to study issues in team cognition requires the development of computational cognitive models that draw on a cognitive architecture to impose psychologically-realistic constraints on agent behavior. For example, the agents in cognitive social simulation experiments may possess learning and mnemonic capabilities similar to those of their human counterparts. The integration of such cognitively-sophisticated agents into social simulation experiments results in a form of social simulation that Sun¹⁶ refers to as ‘cognitive social simulation’. As discussed by Sun,^{16,26} there are a number of factors that motivate the use of cognitive social simulation experiments to better understand aspects of collective behavior. These include the following:

- By providing a framework that captures the task-invariant features of human cognition, cognitive architectures provide a realistic basis for modeling individual agents. This tends to improve the predictive validity of simulation experiments that seek to shed light on the dynamics of team cognition.
- Cognitive architectures can also serve to provide a deeper level of explanation regarding specific simulation results. In particular, they highlight the role that specific cognitive processes play in terms of the emergence of phenomena at the collective or social level, and they thus support the explanation of macro-level (social) phenomena in terms of micro-level (cognitive) processes. Such explanatory capabilities are important when it comes to understanding the role that specific cognitive properties play in shaping collective behavior and performance.
- Aside from their predictive and explanatory virtues, the cognitive plausibility and sophistication of cognitive architectures makes them suitable for experimental manipulations that would be impractical (and sometimes impossible) to perform with human subjects[‡]. Although many studies using cognitive architec-

*Actually this assumption only holds in certain situations. In some task contexts, the mechanisms that are used to retrieve information from memory may be many and varied. For example, human subjects may exploit features of the physical and social environment in order to prompt the recall of specific information from either internal (i.e., bio-memory) or external memory stores.^{20–22}

[†]The ACT-R website²⁷ provides access to a broad range of academic publications covering areas such as problem solving, learning, language processing, decision making, and perceptual processing.

[‡]While cognitive social simulations should not be seen as substitutes for human experimental studies, they can at least help to generate specific hypotheses and identify interesting combinations of experimental factors. These serve to support the design of costly human experiments. As Smart and Sycara [12, p. 55] note, one the purposes of computational modeling and computer simulation “is to guide decisions about what kind of real-world studies with human subjects should be undertaken...The results from computer simulation experiments can help generate specific hypotheses concerning collective cognitive performance, and researchers can subsequently evaluate these hypotheses using real-world empirical analyses.”

tures rely on human data to validate simulation results, in the absence of such data, cognitive architectures provide our best chance of understanding the effect of specific experimental manipulations. In some cases, research with human subjects is difficult due to the cost or risk of running experiments. In other cases, it may be difficult to recruit subjects with the desired characteristics. For example, in one study using a team-based variant of the SOAR cognitive architecture (referred to as Team-SOAR), Kang²⁹ sought to study the effect of personality characteristics on team performance in a naval command and control task. In the absence of cognitive social simulation, the study would have required the recruitment of human subjects with the appropriate personality traits, and this would have presented a serious overhead in terms of running the experiment. Similarly, in cases where we wish to develop a better understanding of the functional role played by cognitive properties in collective performance, it may be necessary to artificially manipulate the cognitive properties of interest. One example of this approach is provided by Reitter and Lebiere³⁰ using the ACT-R architecture. Reitter and Lebiere³⁰ used artificial manipulations of a specific cognitive property, namely, mnemonic decay rate, in order to develop a better understanding of the role that memory plays in collaborative information foraging[§].

- Finally, cognitive architectures make it possible to formulate explanations in terms of readily identifiable cognitive capabilities, such as recall ability. This makes it easier to evaluate the predictions generated by simulations in psychological experiments. As noted by Sun and Naveh³¹ “predictions that contain references to aspects of human cognition (e.g., explicit vs. implicit learning) will be more illuminating and relevant than ones that refer to the internal parameters of an artificial model (e.g., momentum in a neural network) or to external measures only (e.g., percent correct).”

Although there are a number of reasons motivating the adoption of cognitive architectures in social simulation, this does not mean that cognitive architectures are appropriate in all cases. In some cases, the complexity of the task being performed, or the number of agents being modeled, may make it impossible to use cognitive architectures for reasons of scalability and performance. In addition, not all of the details of individual cognitive functioning are necessarily relevant to social simulations. In some situations, it may be perfectly appropriate to abstract away from the details of individual cognitive functioning. Gilbert³² reminds us that the characteristics of phenomena at a higher level (e.g., the social level) often emerge from the behavior of phenomena at a lower level (e.g., the cognitive level), and that not all the details at the lower level are relevant when considering phenomena at the higher level. Cognitive phenomena, for example, may be considered to be emergent with respect to the behavior of neural circuits within the brain. This does not mean, however, that in implementing a cognitive model we need to pay attention to *all* the details of neurophysiological processing. Clearly, for the purposes of computational cognitive modeling, many of the details of neural processing (e.g., the biophysics of single neuron spiking behavior) are deemed to be irrelevant; it is sufficient to develop a model that captures at least some of the coarse-grained functionality implemented by neural systems such as the basal ganglia or the hippocampal formation. Similarly, when it comes to the modeling and simulation of social behavior, it is not always clear that individual agents need to implement *all* the details of human cognitive processing. In some cases, it may be perfectly adequate to use models that abstract away from the details of learning, attention and memory that are implemented at the level of individual agents. To illustrate this, consider the use of multi-agent simulations to model the dynamics of social influence. In many cases, such models make only the crudest assumptions regarding the cognitive capabilities of individual agents, and this is despite the fact that individual cognition is certainly relevant to how social influence takes place. Nevertheless, in many cases, all these models need to assume is that individuals will respond in particular ways to (e.g.) the opinions of their network neighbors (for example, that they will adopt the opinions of their network peers once a certain threshold is reached). The details regarding how these forms of cognitive and behavioral change take place is largely irrelevant in terms of understanding the role that particular social network structures play in influencing opinion dynamics.

Of course, decisions about whether cognition is important to social simulation are far from straightforward. Although Gilbert³² is right that some of the cognitive details may not matter to social phenomena, this is

[§]Studies that manipulate cognitive performance and agent personality characteristics can be useful in prescribing near-optimal cognitive abilities in specific task contexts. In particular, they enable us to specify the kinds of individuals are most likely to benefit team performance in particular situations. This is clearly useful in terms of recruitment, training and team composition.

not an argument against the need for cognitive social simulation. As Sun¹⁶ rightly notes, it is often difficult to determine whether specific cognitive parameters are relevant to social phenomena in the absence of cognitive social simulation. In addition, some of the assumptions about the role played by cognitive factors in social simulation studies may be overly simplistic. When it comes to social influence, for example, Simpkins et al³³ suggest that pre-existing cognitions play an important role in understanding the dynamics of idea propagation. They propose the notion of ‘cognitive advantage’[¶] to account for the differential impact that ideas may have in different communities. Just as the success of a new species in a particular ecological niche is determined by an existing network of inter-species relationships, so the acceptability of a new idea, Simpkins et al³³ suggest, is determined by an existing nexus of beliefs and values that characterize (and indeed define) culturally-circumscribed communities. This is thus a clear case where an understanding of forces and factors at the cognitive level is relevant to simulations dealing with the dynamics of social influence.

3. COGNITIVE SOCIAL SIMULATION AND THE ELICIT EXPERIMENTATION FRAMEWORK

In order to support our understanding of the factors that affect cognitive performance in team-based task settings, we are developing a cognitive social simulation capability, called the ACT-R/CSSC, which is based on the ACT-R architecture. This work is being undertaken as part of our efforts in the International Technology Alliance (ITA) research program, which aims to undertake fundamental research in the network and information sciences in order to support military coalition operations. Our specific aim in the context of the ITA program is to provide a generic capability for performing cognitive social simulation experiments, and then test the capability against a specific team-based task. The task we have chosen to use is the ELICIT task, which is a collaborative problem-solving task that has been used, among other things, to investigate the effect of different organizational structures on team performance.^{34,35} The current status of our implementation efforts in respect of the ACT-R/CSSC are described in Section 3.1 (further information about the ACT-R/CSSC and the cognitive model developed to support simulations of the ELICIT task can be found in Smart et al³⁶ and Smart and Sycara,³⁷ respectively). Section 3.2 provides an overview of the ELICIT task that is the focus of our current modeling efforts and the target of future simulation experiments with the ACT-R/CSSC. Finally, Section 3.3 outlines the focus of future research and development efforts concerning the ACT-R/CSSC.

3.1 ACT-R Cognitive Social Simulation Capability (ACT-R/CSSC)

The ACT-R/CSSC is a Lisp-based cognitive social simulation capability that is built on top of the core ACT-R cognitive architecture (also implemented in Lisp). The core ACT-R architecture includes a number of modules that are intended to represent the functionality of different parts of the human cognitive system. For example, the ‘declarative’ module implements the memory system of the agent. It stores information in the form of what are called ‘chunks’ (these are the basic data structures that are manipulated by all ACT-R modules), and these chunks are associated with context-variant activation levels that determine the probability of an agent recalling a specific piece of information as part of their performance of a cognitive task. The ACT-R/CSSC extends the functionality of the core ACT-R architecture via the incorporation of a number of custom modules. These modules support the functionality required to run cognitive social simulation experiments in which agents interact with one another to solve a common problem:

- **Messaging Module.** This module enables agents to send messages to other agents and receive messages posted by other agents.
- **Language Module.** The language module enables agents to ‘interpret’ the text content of messages received via the messaging module. It processes the text content of messages and creates chunks in the agent’s memory that reflect the semantic content of the message.
- **Web Module.** The web module enables agents to post messages to shared repositories and to also retrieve information from such repositories.

[¶]The cognitive advantage of an idea is, in broad terms, the acceptability of an idea to a particular community based on the prevailing beliefs, attitudes and values adopted by that community.

- **Self Module.** This module stores information about agent properties and characteristics. For example, the self module can be used to store information about an agent’s personality characteristics, attitudes and behavioral predispositions.

In addition to these modules, the ACT-R/CSSC provides access to a number of memory-resident Lisp databases and a Lisp function library. These support the setup and configuration of cognitive social simulation experiments, and they are also used to record and analyze the result of simulation experiments. The architecture and operation of the ACT-R/CSSC is described in more detail in Smart et al.³⁶

In order to support the development of the ACT-R/CSSC, it helps to have a specific team-based task to focus on. The task that is the focus of both our current cognitive modeling efforts and future experimental simulation work is a task called the ELICIT task, which features as part of the ELICIT experimentation framework (see Ruddy³⁸). A brief overview of this task is presented in the following section.

3.2 ELICIT

ELICIT denotes an ongoing effort to provide a common experimental framework to investigate issues in team-based problem solving.³⁸ The ELICIT experimental framework consists of the ELICIT experimental platform, which corresponds to the software environment responsible for running ELICIT experiments; the ELICIT experimental task, which is the actual task performed by experimental participants; and a set of information items, called ‘factoids’, which are the actual pieces of task-relevant information that are processed by experimental participants during the course of task execution. The ELICIT experimental platform consists of the various computer interfaces presented to human subjects during the course of the ELICIT task. These interfaces comprise what we refer to as the task environment. Their importance from the perspective of future cognitive modeling work is discussed in Section 3.3.

The ELICIT task has been the main focus of our recent modeling efforts in the ITA program (see Smart and Sycara³⁷). The task involves the selective presentation of factoids to experimental subjects. Each factoid provides a limited amount of information about a situation, and the aim of the subject is to assimilate enough information in order to make decisions regarding the features of an impending terrorist attack (these features are typically referred to as the ‘dimensions’ of the ELICIT task). The particular features the subject needs to resolve are as follows:

- **who:** the group that will attempt to perform the attack
- **where:** the country in which the attack will take place
- **what:** the kind of target the attack will be against (e.g., an army base)
- **when:** the date and time of the attack (specifically, the day, month and hour of the attack)

In a typical ELICIT experiment, a team consisting of 17 human subjects^{||} is tasked with generating answers to the aforementioned dimensions of the ELICIT task. The individual subjects work in separate locations and communicate using a computer interface that also provides access to a subset of the total set of factoids used for the experiment (the total set of factoids used for an ELICIT experiment is typically referred to as a ‘factoidset’). A limited number of factoids are distributed to subjects at specific intervals in the experiment, and these are referred to as ‘distribution waves’. Assuming that the experiment consists of 17 subjects, two factoids are distributed to each subject in the first distribution wave, with a further two factoids distributed to each subject

^{||}Although a synthetic agent – called abELICIT – has also been developed to support multi-agent simulations of the ELICIT task,³⁹ the agent is not based on a cognitive architecture that is modeled after the human cognitive system. This makes it difficult to run simulation experiments in which the role of specific cognitive variables is systematically investigated (for example, it is difficult to run experiments that explore the effect of different cognitive strategies or different mnemonic capabilities). A key aim of our research is to support the implementation of ELICIT experiments in such a way that the effect of different cognitive variables can be examined. This underlies our decision to use cognitively-sophisticated agents based on the ACT-R architecture as opposed to the existing abELICIT implementation.

in two successive distribution waves. The result is that by the third distribution wave, all factoids ($N = 68$) in the relevant factoidset will have been made available to the human team. The factoids are delivered to each team member's 'inbox', which acts as a repository of all the factoids that an individual team member can access via the aforementioned computer interface.

In order to solve the ELICIT problem, subjects need to share factoids with one another. This is because the four factoids that are delivered to each subject via the distribution waves do not contain sufficient information to solve the problem (i.e., generate answers to the four dimensions of the problem). Subjects therefore need to share the factoids they have access to (i.e., the factoids in their inbox) by either sending factoids to other experimental participants (via a 'send action') or posting factoids to one of a number of 'websites' (via a 'post action'). The result of a send action, in this case, is that the factoid sent by the originating agent gets placed in the inbox of the receiving agent. In the case of a post action, the factoid is posted to a shared repository that can be accessed by any team member that has sufficient access rights to the repository (an ELICIT website thus serves as a mechanism for making factoids available to a larger community of team members than is possible with a single send action). Team members gain access to the factoids associated with websites via what are called 'pull actions'. These actions retrieve the factoids from the relevant website and place the factoids in the team member's inbox.

The profile of information sharing in the task is a function of both agent characteristics (e.g., the agent's willingness to share information), as well as the structure of the organizational environment in which the task is performed. For example, a team member's access to particular websites can be controlled through the setting of 'access rights'. In addition, the structure of the communication network can be configured to limit the channels of communication between team members. In the case of a fully-connected communication network, each team member can send messages to every other team member directly; however, more restrictive communication policies can be created to limit the profile of inter-agent communication. In general, the accessibility of the websites (website access policy) and the structure of the communication network (communication policy) are configured so as to investigate the effect of different organizational environments on the team's ability to solve the ELICIT problem (see, for example, Manso and Ruddy³⁴).

3.3 Modeling the ELICIT Task Environment

The range of actions that users can perform while engaged in the ELICIT task (e.g., send action, post action, and so on) are supported by the task environment of the ELICIT experimentation platform (i.e., the computer interfaces that are presented to subjects during the course of the task). The custom modules that were developed as part of the ACT-R/CSSC (see Section 3.1) are intended to emulate at least some of the features of this task environment; for example, the web module enables agents to interact with ELICIT websites, while the messaging module enables agents to send and receive messages containing information about particular factoids. As discussed by Smart et al,³⁶ future work needs to accommodate additional features of the ELICIT task environment. For example, while performing the ELICIT task, human subjects are able to visualize all the factoids they have access to via the interface of the ELICIT experimentation platform. This arguably serves to reduce the cognitive burden imposed on human subjects because they are provided with continuous (visual) access to task-relevant information and do not need to rely on internal mnemonic encodings of such information in order to perform the task. This is unlike the situation confronted by ACT-R agents in the ACT-R/CSSC. In particular, the ACT-R agents need to memorize whatever factoids they have received (through the aforementioned messaging module), and they then need to recall these factoids when they are required to make specific inferences and decisions. The result is that ACT-R agents are placed at a cognitive disadvantage when performing the task, relative to their human counterparts. In order to address this shortcoming, we plan to develop 'embodied ACT-R agents' (see Schoelles and Gray⁴⁰) that can interact with computational analogues of the ELICIT task environment. This will require us to make greater use of a range of perceptual and motor modules that form part of the core ACT-R architecture.¹⁹ It will also require the implementation of a virtual environment that emulates more of the features of the ELICIT task environment (e.g., the presence of lists of factoids that agents can access during the course of problem solving).

4. RESEARCH FOCUS AREAS

The purpose of cognitive social simulation is to highlight the functional role played by a range of cognitive properties in team performance. One of the immediate aims of our research effort in the context of the ITA program is to replicate a series of experiments that have been performed with the ELICIT task investigating the effect of different military organizational environments on collective problem-solving behavior (see Manso and Ruddy³⁴). Beyond these initial studies, we aim to undertake simulation experiments that explore the interaction between agent-level cognition and features of the social and technological environment in which team-based processing occurs. Subsequent sections describe a number of focus area for future research. Each of these areas is intended to capitalize on the availability of cognitively-sophisticated agent implementations in order to run experiments that would be difficult to perform with conventional (non-cognitive) synthetic agents.

4.1 Agent Communication Strategies

When agents are engaged in a collaborative task, they need to make a number of decisions concerning communicative behavior. For example, if an agent receives a request for information while working on a task, they need to make a decision about whether to suspend their current activity and respond to the request. One option is to temporarily suspend their activity and respond to the request; another is to ignore the request, continue working on their own task, and then (perhaps) respond to the request at a later time. These options reflect two kinds of communication strategy that could be countenanced by agents in responding to requests for information. Other options will relate to decisions about when to relay, re-post or forward information that is received from other agents; when to initiate requests for information; and how much time to spend engaged in communication (as opposed to other task-relevant activities).

Examples of previous work relating to communication strategies can be found throughout the literature on cognitive social simulation. In one study, Kang²⁹ sought to study the effect of agent personality on collective decision making behavior. In undertaking this research, Kang²⁹ made a distinction between ‘active’ and ‘passive’ agents, as well ‘selfish’, ‘neutral’ and ‘cooperative’ agents. These agent personality types are, in fact, all reflected in the communicative propensities of agents. Active agents thus tend to communicate information irrespective of whether or not they are aware of a need for it to be communicated. Passive agents, in contrast, adopt a more conservative approach to communication. They will communicate information but only when they are aware of a need for that information to be communicated within the larger agent community. The distinction between selfish, neutral and cooperative agents is also manifest in terms of the adoption of specific communication strategies: selfish agents assign greater priority to their own tasks at the expense of responding to requests for information; cooperative agents assign greater priority to communication and suspend their own tasks, if necessary, in order to respond to requests for information; and neutral agents adopt a strategy that strikes a balance between cooperative and selfish strategies. Different communication strategies are also evident in the work of Reitter et al⁴¹ and Reitter and Lebiere.³⁰ Reitter et al studied the effect of two different communication strategies – a ‘dump’ and ‘target’ strategy – on collective performance in a social information foraging task. These two strategies resemble the strategies adopted by active and passive agents in the context of the study by Kang;²⁹ the ‘dump’ strategy, in this case, resembles that used by passive agents, whereas the ‘target’ strategy resembles that used by active agents.

All the strategies used by Kang²⁹ and Reitter et al⁴¹ might be described as ‘fixed communication strategies’. They are strategies that agents are required to adopt throughout the course of a task, and they are unable to switch between strategies as the task progresses. For the most part, these sorts of strategies do not require cognitively-sophisticated agents, and they can thus be implemented using conventional social simulation techniques. The same cannot be said of situations where agents adopt more flexible approaches to communicative behavior. Reitter and Lebiere,³⁰ for example, describe a cognitive social simulation experiment in which communicative behavior was subject to a form of ‘metacognitive control’, with the outcome of previous communication decisions being used to guide the choice between alternative communication strategies. Here, issues such as the ability of agents to draw on their past experiences in order to adaptively regulate communication-relevant decisions is of crucial importance. In the context of Reitter and Lebiere’s³⁰ study, the value of particular communication strategies was evaluated with respect to an agent’s past experience of how specific communication

strategies affected access to bodies of task-relevant information. In making a decision about communicative behavior, each agent aimed to minimize the time it took to access information (either by engaging in self-directed exploratory behavior or by relying on the social network as a source of information). If previous experiences indicated that the agent was better off relying on its own exploratory efforts (i.e., exploration resulted in faster access to accurate information), then it opted to refrain from communication. Clearly, such decisions are delicately aligned with the informational dynamics of the task that is being performed. In situations where the information environment is relatively static, then the reliability of messages passed around the social network is likely to be high. This is likely to result in faster access to relevant information compared to situations where the information environment is highly dynamic and messages may contain outdated information that is no longer an adequate reflection of ground truth. This serves to highlight some of the complex interdependencies that exist between the adaptive control of communicative behavior via (memory-based) cognitively-grounded mechanisms, the nature of the information environment (in terms of the dynamics of a constantly evolving situation) (see Section 4.4) and the quality of the information that flows around a networked communication environment.

Decisions relating to communicative behavior may also be influenced by past experiences of interaction with specific individuals. If an agent recalls that a particular agent has consistently provided inaccurate information (or failed to respond in a timely manner) on previous occasions, then this may alter an agent's perceptions regarding the value of communicating with that agent on future occasions. This draws attention to the possible relationships that may exist between social trust (see Section 4.3) and agent communication strategies.

4.2 Agent Communication and the Social Shaping of Memory

Another focus of research attention relates to what might be called the cognitive consequences of communication. The psychological literature has identified a number of ways in which social communication may impact memory processes.^{20, 42–46} In communicating with another person, for example, we are required to retrieve information from memory, and this may constitute a form of socially-mediated rehearsal that serves to enhance the subsequent probability of recalling specific items of information.^{43, 45} Communicative exchanges may also require agents to process information to a greater depth than would otherwise be the case. By talking about a shared experience in a social context, we are obliged to engage in a deeper kind of analysis of the past: exploring its significance for the actors involved, assessing its relationship to other experiences, and reflecting on its relevance to past, present and future goals. All of this is likely to entail a deeper kind of analysis of the past compared to that which would be encountered in the case of solitary reminiscences. The end result should, perhaps, be seen as a form of socially-mediated 'elaborative rehearsal'^{**} that improves the availability of specific memories for future recall.²⁰ Agent communication may also support a form of collaborative recall in which the successful recall of specific memories occurs as a result of the progressive generation and elaboration of mnemonic cues across successive communicative exchanges.⁴⁶ Finally, inter-agent communication may contribute to a form of retrieval-induced forgetting⁴⁸ in which the selective recall of particular information serves to promote the forgetting of related information (see Hirst and Echterhoff⁴²).

The potential interactions that occur between memory and communication constitute an important focus area for future research attention. Cognitive social simulation studies have not tended to look at the social shaping of memory processes; however, given the role that memory plays in many task situations, it may be crucial to incorporate these effects in simulation studies of team problem-solving behavior. One issue of particular interest is whether the response to requests for information makes particular bodies of information more active (and thus more accessible) in memory. If so, then agent communication within a team environment may serve to influence the recall of particular items of information, and this may exert significant effects on team-based processing.

^{**}Elaborative rehearsal is one of two forms of rehearsal identified by Craik and Lockhart⁴⁷ – the other form being maintenance rehearsal. Maintenance rehearsal is a kind of rehearsal that involves the simple repetition of to-be-remembered information. Elaborative rehearsal, in contrast, involves a much deeper kind of analysis: one in which the to-be-remembered information is often subjected to a more in-depth semantic analysis. Crucially, Craik and Lockhart⁴⁷ suggest that elaborative modes of rehearsal provide a more effective route to long-term memory as compared to the simple repetition of information (i.e., maintenance rehearsal).

4.3 Social Trust and Influence

As part of any team-based activity, agents share information in order to influence the beliefs held by other team members. One of the factors that determines the level of influence associated with transmitted information is the level of trust that the receiving agent has in the sending agent. When trust is high, the influence of the transmitted information is assumed to be greater than when trust is low, and thus the transmitted information can be expected to feature to a greater extent in the receiving agent’s belief formation and belief revision processes.

The level of trust that exists between agents, and the associated impact this has on the importance of transmitted information, has emerged as a significant factor in controlling the opinion dynamics of large agent communities. In one study, for example, Ginton et al⁴⁹ examined the effect of social influence parameters in a large team-based information sharing task. In Ginton et al’s model, a large number of (non-cognitive) agents (e.g., 1000) work together to form beliefs concerning the state of a specific feature of the environment. Only a small proportion of the agents (e.g., 5%) are connected to sensors which provide information about the environment, and this leads those agents to form initial beliefs, which are then propagated to other agents. Importantly, the sensors used in the model are noisy, only producing correct observations according to a predefined accuracy. This creates uncertainty for the agents. Ginton et al discovered that the importance that each agent assigns to its neighbor’s opinions has a dramatic effect on the ability of agents to converge on a consistent set of beliefs about the environment. They found that, across a range of model parameters, there is typically an optimal level of influence that enables agents to converge on the correct interpretation. This work, as well as work by Prymak et al,⁵⁰ suggests that mechanisms supporting the adaptive regulation of social influence may help a community of agents deal with uncertain information and establish accurate shared beliefs about the state of the environment. The level of trust that exists between agents may be one means by which this influence is regulated.

These studies serve to highlight the potential role of social trust mechanisms in regulating team-based performance outcomes. When it comes to cognitive social simulation, the implementation of social trust presents a particular challenge. While some previous work has addressed the issue of trust representation in cognitive social simulations,⁵¹ there has been little work to date on representing trust (and social influence parameters, more generally) within ACT-R (although see Kennedy and Krueger⁵²). Further progress in this area is important because social trust not only determines the extent to which received information is factored into an agent’s internal belief formation and revision processes, it is also likely to influence the dynamics of agent interaction, with greater levels of agent interaction being observed between agents with the strongest trust relationships (see Section 4.1). A particular area of interest, and one that extends previous agent-based studies of social trust processes, is the possibility of adapting social trust with respect to an agent’s memory of past social interactions. Such experience-dependent modifications of social trust may be of great importance because they may serve to adaptively regulate social influence over the course of successive task performances. One of the consequences of memory-dependent social trust mechanisms may be to support an effective ‘rewiring’ of the communication network with respect to the channels of influence that exist between agents. Although a particular kind of communication network may support a particular profile of information flow and influence within an agent community, the tendency of agents to disregard the contributions of particular individuals across time (via the dynamic modification of social trust) means that the topology of the influence network is no longer the same as the communication network across which social influence is exerted. Given that previous research has indicated that dynamic communication network topologies may yield certain advantages over their more statically configured counterparts when it comes to collective problem solving,⁵³ a greater understanding of the role that cognitive capabilities (e.g., memory) play in dynamically modulating social influence is a topic of considerable research interest for future social simulation studies into team cognition.

4.4 Dynamic Information Environments

One of challenges facing many teams, particularly those working in military environments, concerns the rapidly changing nature of the external environment. The tempo of military operations, for example, must keep pace with events as they unfold in real-time, and this places limits on the temporal window within which analytic outcomes, as well as associated decisions and military plans, are valid. Due to an ever-changing situation picture, military planners and decision-makers must continuously revise and update their interpretations, hypotheses and

understanding of the current situation in the light of newly received information. Any attempt to model such processes is likely to require the inclusion of dynamic information streams that cause agents to extend and revise the beliefs they have about particular situations.

Dynamic information environments open up a range of interesting issues for computational models of team cognition. Previous work (using non-cognitive agents) has shown that when it comes to information sharing in large teams, the degree of trust between agents plays an important role in determining the extent to which the team converges on accurate and consistent beliefs about the external environment.^{49,50} Importantly, however, such models tend to rely on the use of static facts, where the facts that are observed by agents do not change across the course of the simulation. Recently, Eck and Soh⁵⁴ have shown that when we switch to the assumption that agents are operating in a dynamic information environment, and static facts are replaced with dynamic ones, social trust no longer functions to adequately control belief convergence. They observe a phenomenon that they refer to as ‘institutional memory’ in which team members converge to the initial value of a fact, but then fail to properly revise their belief states in the face of changing fact values. Because this phenomenon appears to be immune to the influence of social trust, it may be important to look for alternative mechanisms that support the convergence to accurate belief states. One option here is to rely on the cognitive properties of agents in the context of cognitive social simulation studies. In particular, once agents are endowed with an ability to forget information, the role that outdated information plays in maintaining collective interpretations of the environment may be undermined (see Reitter and Lebiere³⁰). This possibility serves to highlight the potentially positive functional role played by an apparent shortcoming of the human cognitive system, namely, our tendency to forget information. In dynamic information environments, a capacity to forget earlier information may result in more recent information having greater salience and therefore greater control over ongoing cognitive processes. This serves to highlight the value of using cognitively-sophisticated agents in social simulation experiments. By factoring in the kinds of constraints and capabilities exhibited by the human cognitive system, we are able to assess the role that cognitive-level variables play in supporting (or subverting) task performance. Conventional agent-based simulations often fail to yield this sort of understanding precisely because they fail to adequately represent many of the features associated with human cognitive processing.

5. CONCLUSION

The aim of cognitive social simulation is to improve our understanding of the complex inter-play between factors that are spread across the cognitive, social, informational and technological domains. This makes cognitive social simulation techniques particularly appealing as a means to undertake experiments into socially-distributed cognition. The current paper reports on the results of an ongoing effort to develop a cognitive social simulation capability that can be used to undertake studies into team cognition using the ACT-R cognitive architecture. This capability is intended to support simulation experiments using a team-based problem solving task, called the ELICIT task, that has been used to explore the impact of different organizational environments on collective problem solving performance. One of our immediate objectives with the ACT-R-based simulation capability is to replicate the results of previous studies that have used the ELICIT task with human subjects; this will enable us to test the integrity of the simulation capability and evaluate the fidelity of computational cognitive models of task-relevant behavior. Beyond the replication of previous results, we also aim to apply the cognitive social simulation capability to a variety of research focus areas. In particular, we aim to study the role of cognitive processes in determining aspects of communicative behavior, as well as the impact of communicative behavior on the shaping of task-relevant cognitive processes (e.g., the social shaping of agent memory as a result of communicative exchanges). We also aim to study the interaction between features of the information environment (such as whether the veracity of information is fixed for the duration of the task or whether the veracity of earlier information can be negated by the receipt of subsequent information) and features of the cognitive system (such as the ability of agents to forget outdated information). Finally, we aim to investigate ways to incorporate experience-dependent social trust mechanisms into cognitive social simulation experiments. This will enable us to assess the extent to which agents are able to use social trust as a means of adaptively regulating their exposure to particular bodies of task-relevant information. Research in these areas will, we suggest, help to elucidate some of the complex interactions that exist between cognitive mechanisms, social processes and features of the task environment in the context of team-based problem-solving activities.

ACKNOWLEDGMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

- [1] Hinsz, V., Tindale, R., and Vollrath, D., “The emerging conceptualization of groups as information processors,” *Psychological Bulletin* **121**, 43–64 (1997).
- [2] O’Hara, K., Contractor, N., Hall, W., Hendler, J. A., and Shadbolt, N., “Web science: Studying the emergence of macro-level features on the World Wide Web,” *Foundations and Trends in Web Science* **4**(2-3), 103–267 (2013).
- [3] Malone, T. W., Laubacher, R., and Dellarocas, C., “The collective intelligence genome,” *MIT Sloan Management Review* **51**(3), 21–31 (2010).
- [4] Chi, E. H., Pirolli, P., Suh, B., Kittur, A., Pendleton, B., and Mytkowicz, T., “Augmented social cognition,” in [AAAI Spring Symposium on Social Information Processing], Stanford, California, USA (2008).
- [5] Smart, P. R. and Shadbolt, N. R., “Social machines,” in [Encyclopedia of Information Science and Technology], Khosrow-Pour, M., ed., IGI Global, Hershey, Pennsylvania, USA (in press).
- [6] Bernstein, A., Klein, M., and Malone, T. W., “Programming the global brain,” *Communications of the ACM* **55**(5), 41–43 (2012).
- [7] Law, E. and von Ahn, L., “Human computation,” *Synthesis Lectures on Artificial Intelligence and Machine Learning* **5**(3), 1–121 (2011).
- [8] Kearns, M., “Experiments in social computation,” *Communications of the ACM* **55**(10), 56–67 (2012).
- [9] Kraut, R., Maher, M. L., Olson, J., Malone, T. W., Pirolli, P., and Thomas, J. C., “Scientific foundations: A case for technology-mediated social-participation theory,” *Computer* **43**(11), 22–28 (2010).
- [10] Cooke, N. J., Gorman, J. C., and L, W. J., “Team cognition,” in [Handbook of Applied Cognition], Durso, F. T., Nickerson, R. S., Dumais, S. T., Lewandowsky, S., and Perfect, T. J., eds., John Wiley & Sons, Chichester, UK, 2nd ed. (2007).
- [11] Cooke, N. J., Gorman, J. C., Myers, C. W., and Duran, J. L., “Interactive team cognition,” *Cognitive Science* **37**(2), 255–285 (2013).
- [12] Smart, P. R. and Sycara, K., “Collective sensemaking and military coalitions,” *Intelligent Systems* **28**(1), 50–56 (2013).
- [13] Klein, G., Wiggins, S., and Dominguez, C. O., “Team sensemaking,” *Theoretical Issues in Ergonomics Science* **11**(4), 304–320 (2010).
- [14] Diehl, M. and Stroebe, W., “Productivity loss in brainstorming groups: Toward the solution of a riddle,” *Journal of Personality and Social Psychology* **53**, 497–509 (1987).
- [15] Mason, W., Conrey, F., and Smith, E., “Situating social influence processes: Dynamic, multidirectional flows of influence within social networks,” *Personality and Social Psychology Review* **11**(3), 279–300 (2007).
- [16] Sun, R., “Cognitive social simulation incorporating cognitive architectures,” *Intelligent Systems* **22**(5), 33–39 (2007).
- [17] Thagard, P., “Cognitive architectures,” in [The Cambridge Handbook of Cognitive Science], Frankish, K. and Ramsey, W. M., eds., 50–70, Cambridge University Press, Cambridge, UK (2012).
- [18] Anderson, J. R., [How Can the Human Mind Occur in the Physical Universe?], Oxford University Press, Oxford, UK (2007).
- [19] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y., “An integrated theory of the mind,” *Psychological Review* **111**(4), 1036–1060 (2004).

- [20] Smart, P. R., “Extended memory, the extended mind, and the nature of technology-mediated memory enhancement,” in [*1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*], Maryland, USA (2010).
- [21] Clowes, R. W., “The cognitive integration of e-memory,” *Review of Philosophy and Psychology* **4**(1), 107–133 (2013).
- [22] Clark, A., “Intrinsic content, active memory and the extended mind,” *Analysis* **65**(285), 1–11 (2005).
- [23] Laird, J. E., Newell, A., and Rosenbloom, P. S., “SOAR: An architecture for general intelligence,” *Artificial Intelligence* **33**(1), 1–64 (1987).
- [24] Rosenbloom, P., Laird, J., Newell, A., and McCarl, R., “A preliminary analysis of the SOAR architecture as a basis for general intelligence,” in [*Foundations of Artificial Intelligence*], Kirsh, D., ed., 289–326, MIT Press, Cambridge, Massachusetts (1992).
- [25] Laird, J. E., [*The SOAR Cognitive Architecture*], MIT Press, Boston, Massachusetts, USA (2012).
- [26] Sun, R., “The CLARION cognitive architecture: Extending cognitive modeling to social simulation,” in [*Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Interaction*], Sun, R., ed., 79–99, Cambridge University Press, New York, New York, USA (2006).
- [27] “ACT-R.” URL: <http://act-r.psy.cmu.edu/> [accessed: 2014-04-20].
- [28] Morrison, J. E., “A review of computer-based human behavior representations and their relation to military simulations,” Tech. Rep. IDA Paper P-3845, Institute for Defense Analyses (2003).
- [29] Kang, M., “The effects of agent activeness and cooperativeness on team decision efficiency: A computational simulation study using Team-Soar,” *International Journal of Human-Computer Studies* **65**(6), 497–510 (2007).
- [30] Reitter, D. and Lebiere, C., “Social cognition: Memory decay and adaptive information filtering for robust information maintenance,” in [*26th AAAI Conference on Artificial Intelligence*], Toronto, Canada (2012).
- [31] Sun, R. and Naveh, I., “Simulating organizational decision-making using a cognitively realistic agent model,” *Journal of Artificial Societies and Social Simulation* **7**(3) (2004).
- [32] Gilbert, N., “When does social simulation need cognitive models,” in [*Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Interaction*], Sun, R., ed., 428–432, Cambridge University Press, New York, New York, USA (2006).
- [33] Simpkins, B., Sieck, W. R., Smart, P. R., and Mueller, S. T., “Idea propagation in social networks: The role of cognitive advantage,” in [*1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*], Maryland, USA (2010).
- [34] Manso, M. and Ruddy, M., “Comparison between human and agent runs in the ELICIT N2C2M2 validation experiments,” in [*18th International Command and Control Research and Technology Symposium (ICCRTS)*], Alexandria, Virginia, USA (2013).
- [35] Manso, M., “N2C2M2 validation using abELICIT: Design and analysis of ELICIT runs using software agents,” in [*17th International Command and Control Research and Technology Symposium*], Fairfax, Virginia, USA (2012).
- [36] Smart, P. R., Richardson, D. P., Sycara, K., and Tang, Y., “Towards a cognitively realistic computational model of team problem solving using ACT-R agents and the ELICIT experimentation framework,” in [*19th International Command and Control Research Technology Symposium (ICCRTS’14)*], Alexandria, Virginia, USA (2014).
- [37] Smart, P. R. and Sycara, K., “Cognitive social simulation and collective sensemaking: An approach using the ACT-R cognitive architecture,” in [*6th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE’14)*], Venice, Italy (2014).
- [38] Ruddy, M., “ELICIT – the experimental laboratory for investigating collaboration, information sharing and trust,” in [*12th International Command and Control Research and Technology Symposium (ICCRTS)*], Newport, Rhode Island, USA (2007).
- [39] Ruddy, M., Wynn, D. M., and McEver, J., “Instantiation of a sensemaking agent for use with ELICIT experimentation,” in [*14th International Command and Control Research and Technology Symposium*], Washington D.C., USA (2009).

- [40] Schoelles, M. J. and Gray, W. D., “Argus: A suite of tools for research in complex cognition,” *Behavior Research Methods, Instruments, & Computers* **33**(2), 130–140 (2001).
- [41] Reitter, D., Sycara, K., Lebiere, C., Vinokurov, Y., Juarez, A., and Lewis, M., “How teams benefit from communication policies: Information flow in human peer-to-peer networks,” in [*20th Conference on Behaviour Representation in Simulation and Modeling (BriMS,11)*], Sundance, Utah, USA (2011).
- [42] Hirst, W. and Echterhoff, G., “Remembering in conversations: The social sharing and reshaping of memories,” *Annual Review of Psychology* **63**(1), 55–79 (2012).
- [43] Wegner, D., Giuliano, T., and Hertel, P. T., “Cognitive interdependence in close relationships,” in [*Compatible and Incompatible Relationships*], Ickes, W. J., ed., Springer-Verlag, New York, New York, USA (1985).
- [44] Roediger, H., Meade, M., and Bergman, E., “Social contagion of memory,” *Psychonomic Bulletin and Review* **8**(2), 365–371 (2001).
- [45] Neisser, U., “The ecological study of memory,” *Philosophical Transactions of the Royal Society B: Biological Sciences* **352**(1362), 1697–1701 (1997).
- [46] Sutton, J., Harris, C., Keil, P., and Barnier, A., “The psychology of memory, extended cognition, and socially distributed remembering,” *Phenomenology and the Cognitive Sciences* **9**(4), 521–560 (2010).
- [47] Craik, F. I. M. and Lockhart, R. S., “Levels of processing: A framework for memory research,” *Journal of Verbal Learning and Verbal Behaviour* **11**, 671–684 (1972).
- [48] Anderson, M. C., Bjork, R. A., and Bjork, E. L., “Remembering can cause forgetting: Retrieval dynamics in long-term memory,” *Journal of Experimental Psychology: Learning, Memory and Cognition* **20**, 1063–1087 (1994).
- [49] Grinton, R., Sycara, K., and Scerri, P., “Exploiting scale invariant dynamics for efficient information propagation in large teams,” in [*International Conference on Autonomous Agents and Multi-Agent Systems*], Toronto, Canada (2010).
- [50] Pryymak, O., Rogers, A., and Jennings, N. R., “Efficient opinion sharing in large decentralised teams,” in [*International Conference on Autonomous Agents and Multi-Agent Systems*], Valencia, Spain (2012).
- [51] Pollock, S. S., Alt, J. K., and Darken, C. J., “Representing trust in cognitive social simulations,” in [*4th International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*], Maryland, USA (2011).
- [52] Kennedy, W. G. and Krueger, F., “Building a cognitive model of social trust within ACT-R,” in [*2013 AAAI Spring Symposium: Trust and Autonomous Systems*], Palo Alto, California, USA (2013).
- [53] Smart, P. R., Huynh, T. D., Braines, D., and Shadbolt, N. R., “Dynamic networks and distributed problem-solving,” in [*Knowledge Systems for Coalition Operations (KSCO’10)*], Vancouver, British Columbia, Canada (2010).
- [54] Eck, A. and Soh, L.-K., “Dynamic facts in large team information sharing,” in [*International Conference on Autonomous Agents and Multi-Agent Systems*], Saint Paul, Minnesota, USA (2013).