Using ACT-R to Model Collective Sensemaking in Military Coalition Environments

Darren P. Richardson*†, Paul R. Smart*, Katia Sycara‡, Paul Stone§, Cheryl Giammanco¶ and Gavin Powell‖
*Electronics & Computer Science, University of Southampton, Southampton, SO17 1BJ, UK
‡Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA
§Emerging Technology Services, IBM United Kingdom Ltd., Hursley Park, Winchester, Hampshire, SO21 2JN, UK
¶U.S. Army Research Laboratory, Aberdeen Proving Ground, MD 21005-5425, USA
‖Innovation Works, EADS UK, Newport, Gwent, NP10 8FZ, UK

†Corresponding author – dpr1g09@ecs.soton.ac.uk

***Abstract*—Cognitive social simulations, enabled by cognitive architectures (such as ACT-R), are particularly well-suited for advancing our understanding of socially-distributed and socially-situated cognition. As a result, multi-agent simulations featuring the use of ACT-R agents may be important in improving our understanding of the factors that influence collective sensemaking. While previous studies demonstrate the feasibility of using ACT-R to model collective cognition, as well as sensemaking processes at the individual level, the development of an ACT-R model of collective sensemaking in a coalition environment presents a range of relatively novel methodological, technological and modeling challenges. Such challenges include the need to equip ACT-R agents with communication capabilities, the need to deal with highly dynamic information environments, the need to support intelligent information retrieval capabilities, and the need to represent inter-agent cognitive differences. These challenges shape the nature of research and development efforts to create a multi-agent simulation capability that can be used to explore the impact of different sociotechnical interventions on collective sensemaking processes. In this paper, we discuss the research efforts being undertaken to address these challenges in the context of the International Technology Alliance (ITA) research program. We also discuss the motivations for using ACT-R to model collective sensemaking processes and outline some opportunities for model application and empirical evaluation.**

I. Introduction

In recent years, there has been a growing interest in the socially-distributed or socially-situated nature of human cognition across a number of scientific disciplines [1, 2, 3, 4]. Cognitive processes that were typically studied at the level of individual agents, such as memory, are now being re-examined within a more social context [5], and increasing attention is being paid to the factors that enable groups to function as the processors of information (see [6]). This interest in the social dimension of cognition is, in part, a reflection of the growing popularity of embodied, extended and situated approaches within the sciences of the mind [7, 8, 9]. However, the research is also motivated by an attempt to engineer systems that harness the collective cognitive potential of groups of individuals. The advent of global information and communication networks, such as the World Wide Web, has clearly been one of the drivers in such research, with notions such as collective intelligence [10], augmented social cognition [11] and social machines [12] serving as some of the conceptual anchors for ongoing research efforts. However, systems that support socially-distributed cognition are also important in more restricted organizational contexts. This is particularly so as advances in sensor technology lead to a significant expansion in the scale and scope of available data assets. As organizations move into this 'Big Data' era, so they are under increasing pressure to distribute cognitive effort and harness the collective cognitive potential of their workforces.

Advances in sensor and networking technology pose particular challenges for military coalitions. As organizations that are under continual pressure to make sense of their environment and act accordingly, it is imperative that they make best use of their available information processing resources by appropriately distributing cognitive effort. However, it is not always clear how this distribution of cognitive effort should be achieved. When it comes to collective or team sense-making, for example, should the individual team members be allowed to engage in frequent communication with one another in order to coordinate their interpretations, or should more restrictive communication policies be enforced. While it may be natural to assume that full connectivity and frequent communication is a virtue – 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 [13]. Cognitive biases are also a major concern when it comes to collective cognition, especially since cognitive biases are sometimes more extreme at the group level [6]. In one study using computer simulation techniques, Hutchins [14] showed that the temporal profile of inter-agent communication had a pronounced impact on the tendency for confirmation bias at the collective level, with the early and more rapid exchange of information leading to greater levels of collective confirmation bias. Such studies highlight the importance of undertaking empirical studies that examine the effect of a variety of factors on the dynamics of collective cognition within specific task contexts.

One approach to the study of collective cognition is to rely on multi-agent simulation techniques. 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 effect 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 cognitive architectures in performing cognitive social simulations. Cognitive architectures are frameworks that make particular commitments about the kind of mental representations and computational procedures that are sufficient to explain important aspects of human cognition, such as problem solving, memory and learning [17]. They often serve as a framework for the development of computational cognitive models that are then validated with respect to human experimental data. Although a cognitive architecture can be implemented using connectionist schemes, some of the most influential cognitive architectures, such as ACT-R [18, 19] and SOAR [20] rely on rule and symbol forms of processing. Of particular note is ACT-R, which has been the focus of a sustained research and development effort for more than 30 years, and which has been applied to a broad range of cognitive processing contexts.

In order to advance our understanding of socially-distributed cognition, we advocate the use of cognitive architectures to perform cognitive social simulation, as suggested by Sun [16]. As part of our work in the International Technology Alliance (ITA) – a consortium of academic, industrial and government partners undertaking fundamental research in the network and information sciences – we propose to deploy ACT-R in a multi-agent environment, with distinct ACT-R models serving as individual cognitive agents. Our particular focus of attention is collective sensemaking within military coalition environments. In this paper, we describe the motivation for using ACT-R to model collective sensemaking. We also outline some of the areas of research interest associated with the modeling effort.

II. Collective Sensemaking in Military Coalitions

Sensemaking has been the focus of sustained research attention over the past 10-20 years [21, 22, 23, 24]. It has been defined as a "motivated, continuous effort to understand connections (which can be among people, places, and events) in order to anticipate their trajectories and act effectively" [21]. It is essentially the activity that individuals engage in in order to explain and predict the features of some object, event or situation.

Sensemaking is, at heart, a cognitive activity: it is an activity that involves the processing of information in order to yield an outcome (i.e., understanding) that is recognizably cognitive in nature. This does not mean, however, that sensemaking is an activity that only individuals engage in. There is a growing appreciation of the prevalence and importance of what might be called 'collective sensemaking' [3] or 'team sensemaking' [25]; i.e., the activities that are performed by groups of individuals in order to develop understanding at both the individual and collective levels. Work on collective sensemaking is the focus of an increasing body of empirical and theoretical work within a number of research communities, and these efforts are paralleled by an extensive body of work into related notions, such as shared situation awareness, shared understanding and shared mental models.

Collective sensemaking is a phenomenon of considerable importance in a number of different task contexts, such as intelligence analysis [22, 26], military planning [27] and healthcare provision [28]. It tends to emerge in any situation where a group or team of individuals is required to pool their cognitive resources in an attempt to interpret complex, incomplete and uncertain bodies of data. As with other types of distributed cognition, there are a number of factors that motivate the transition from individual to collective sensemaking. One reason is that the complexity of the available information may be such that no one individual has the relevant knowledge or expertise to interpret it properly. Different individuals may possess specialist knowledge and expertise, and the involvement of these individuals may be necessary in order for a group to make sense of some larger body of information. Another reason why collective sensemaking is important is that different individuals may have access to different bodies of information. In a military coalition environment, for example, different individuals may have access to different sources of information (e.g., particular sensor systems or databases) by virtue of their position in organizational structures. Individuals may also have different abilities when it comes to the effective probing of the information environment and the elicitation of cues which serve to guide interpretational processes. In intelligence analysis, for example, different individuals may have different expertise in retrieving information from particular intelligence sources. A final reason why individuals may resort to collective sensemaking strategies is because it builds redundancy into the sensemaking process. Any errors or omissions that might be made by one individual have a greater chance of being detected or compensated for by the efforts of others.

Sensemaking is a key capability for military coalitions, enabling both individuals and teams to make sense of conflicting, ambiguous and uncertain information. This importance is reflected in the vision of network-centric operations (NCO), where sensemaking is seen as a key factor in enabling coalitions to respond in an adaptive manner to complex and dynamic situations. According to the NCO Conceptual Framework (NCO-CF), for example, sensemaking at both the individual and collective levels has a direct impact on decision synchronization, force agility and mission effectiveness [29]. In particular, sensemaking processes are seen as an intervening variable in the NCO value chain: they enable military organizations to capitalize on the progress made with respect to networking technology and improved information sharing capabilities [29]. This is something that makes research into collective sensemaking of vital importance to military coalition organizations (see [3]). Indeed, the extension of sensemaking into the social domain is something that is explicitly recognized by the NCO-CF. Within the NCO-CF, collective (or shared) sensemaking is seen as the collective counterpart of individual sensemaking, and it is regarded as something that is strongly influenced by social interaction and social networks [29].

The military coalition environment presents a number of challenges to collective sensemaking. Some of the features that are likely to complicate collective sensemaking include

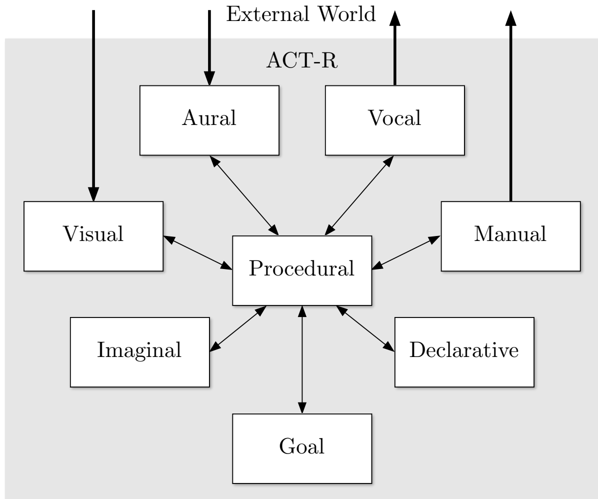

Fig. 1. The core modules of the ACT-R v.6 cognitive architecture.

(but are not limited to) variable trust relationships, dynamic communication network topologies, the extent of information sharing, differential access to specific bodies of information, the presence of cognitive diversity and the potential for mis-communication. Given the centrality of collective sensemaking to coalition operations, it is important that we develop a better understanding of the relationship between specific features of the coalition communication environment and aspects of collective sensemaking performance. Attempts at cognitive computational modeling should clearly aim to accommodate these features as part of their attempt to explain and predict the effect of specific socio-technical interventions on the dynamics of collective sensemaking.

III. ACT-R

A cognitive architecture is a model of how cognition can occur within the physical structures of the human brain [18]. ACT-R is one of the most popular and influential cognitive architectures, and it has been widely used within the cognitive science research community. ACT-R allows researchers to create models of cognitive performance in various task contexts by providing a representational framework and set of computational procedures, both of which are modeled on the human cognitive system.

ACT-R is comprised of several modules, where each module represents a particular brain function. Communication between the modules is limited to very small amounts of information, which are stored in buffers, and a single module – the procedural module – is responsible for inter-module coordination. As shown in Fig. 1 there are eight core modules in the most recent version of ACT-R, released in 2005.

One of the factors that motivates the application of ACT-R to collective sensemaking is that considerable effort has already been invested in applying ACT-R to the modeling of sensemaking processes at the level of individual human agents [30]. Using a geospatial intelligence analysis task, researchers working within the Intelligence Advanced Research Projects Activity (IARPA) Integrated Cognitive-neuroscience Architectures for the Understanding of Sensemaking (ICArUS) program have used ACT-R to model several of the core information foraging and hypothesis-updating processes associated with human sensemaking. Their research has correctly predicted both the presence and degree of four main cognitive biases, namely, confirmation bias, anchoring and adjustment, representativeness, and probability matching. The ACT-R model also exhibits a behavioral profile that is similar to human subjects in terms of information selection and the allocation of military resources based on probability estimates (see [30] for a review).

IV. Research Areas

Although ACT-R has typically been used to study cognitive processes at the level of individual human agents, the growing interest in the social dimension of cognition motivates a consideration of its use in collective processing contexts. The feasibility of using ACT-R to model socially-distributed processes has already been demonstrated in a number of studies [31, 32]. In one recent study, for example, Reitter and Lebiere [32] used a multi-agent simulation involving multiple ACT-R agents to investigate the effect of differential rates of decay in individual memory on performance in a simulated information foraging task. Their results suggest that average task performance improves with increasing rates of individual memory decay. This highlights the value of using a cognitively-rich agent model to investigate the potential interactions between individual cognitive function and collective task performance.

While previous studies demonstrate the feasibility of using ACT-R to model collective cognition, the development of an ACT-R model of collective sensemaking in a coalition environment presents a range of relatively novel methodological, technological and modeling challenges. These challenges serve to ground the focus of our research efforts within the ITA program. In subsequent sections, we provide an overview of a number of areas where research and development efforts are currently being focused.

A. *Agent Communication*

As with any collaborative process, collective sensemaking often entails a capacity for inter-agent communication. Agents may need to initiate requests for information and communicate information in a form that other agents can understand. While a number of approaches to inter-agent communication could be explored, our research efforts in the ITA are centered on the use of a controlled variant of natural language, called Controlled English (CE). As with other controlled natural languages, CE imposes a set of grammatical constraints on linguistic expressions that serve to reduce semantic ambiguity and support human-machine interaction. By using CE as a vehicle for communication, agents can provide information to other agents in a form that is semantically-interpretable, irrespective of the actual domain in which sensemaking occurs. In other words, CE provides us with a domain-neutral mode of information representation, something that is of critical significance in terms of the broader applicability of the model to other domains of interest (see Section V). A number of other factors motivate a consideration of the use of CE within the current context. Firstly, CE supports the expression of queries, which can be used to provide question-answering capabilities. Secondly, because CE is human-readable, it can be used as a vehicle for communication between human and machine agents. This is important because it allows human observers to understand the communicative transactions that are made between agents in the context of a simulation. It also allows human agents to participate in the simulation by acting as an information source. In addition, previous work has explored the incorporation of language-enabled ACT-R agents into team training simulations where they act as synthetic teammates [33]. The use of CE could enable ACT-R agents to play a similar role in the context of collective sensemaking tasks without the overhead incurred by the need to implement full natural language processing capabilities. Thirdly, CE can serve as a knowledge representation language. This makes it of potential value in terms of furnishing ACT-R agents with the knowledge and conceptual structures necessary to deal with particular domains. The integration of computational ontologies with ACT-R serves a similar objective [34].

B. Social Trust and Influence

As part of collective sensemaking efforts, agents share information into order to influence the beliefs held by other agents. 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 heavily 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, Glinton et al [35] examined the effect of social influence parameters in a large team information sharing task. In Glinton et al's model, a large number of 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. Glinton 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 Pryymak et al [36], 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.

One of the challenges associated with the use of ACT-R to support socio-cognitive simulations of sensemaking is thus to develop a means of representing social trust. While some previous work has addressed the issue of trust representation in cognitive social simulations [37], there has been little work to date on representing trust (and social influence parameters, more generally) within ACT-R. 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[1].

C. Dynamic Information Environments

One of challenges facing sensemakers in military coalitions concerns the rapidly changing nature of the operational environment. The tempo of military operations 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, coalition sensemakers must continuously revise and update their interpretations, hypotheses and understanding in light of newly received information. Any attempt to model collective sensemaking in coalition contexts is thus likely to require the inclusion of dynamic information streams that cause agents to extend and revise the beliefs they have about the current situation.

Dynamic information environments open up a range of interesting issues for computational models of collective sensemaking. We have already seen 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 belief states [35, 36]. However, such models rely on the use of static facts, where the facts that are observed by agents do not change across the course of the simulation. In these situations, any variability in belief states that arises due to the observation of the environment comes about as the result of inaccuracies in the sensors that provide information to the agent team. Recently, Eck and Soh [39] 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 such mechanism may rely on the cognitive properties of agents. 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

[1] It should also be noted that social trust is not something that need be fixed throughout the course of a simulation. Agents may revise their trust in one another based on their past interaction experiences. For example, if one agent is the source of information that turns out to be consistently incorrect, then trust in that agent may deteriorate over time. This is a potentially important feature for socio-cognitive simulations because dynamic trust relationships may contribute to an effective 'rewiring' of the inter-agent communication network across the course of a simulation. Such dynamically configured networks may yield performance profiles that are not seen in the case of their more statically configured network counterparts (see [38]).

may be undermined. This possibility serves to highlight the importance of factoring cognitive variables into multi-agent simulations. By using cognitively-sophisticated agents, we can observe complex interactions that subtend the cognitive (e.g., memory), social (e.g., social trust) and technological (e.g., communication network structure) domains. In some cases, this presents us with additional opportunities to control or influence team performance outcomes.

D. Active Probing of the Information Environment

In making sense of a situation, a human agent does not sit passively waiting to be informed of relevant information. Instead, sensemaking is often an active process. Human agents engaged in sensemaking do not simply react passively to the information they receive, they also seek to manipulate their information environments in ways that meliorate their access to hidden patterns, relationships and contingencies.

What this means in terms of the development of ACT-R models is that ACT-R sensemaking agents should be equipped with an ability to actively probe the information environment in order to seek out new information. Such activities can be seen as a form of information foraging, which is a key component of many sensemaking models (e.g., [26]). One means by which these active probing capabilities could be implemented is by enabling ACT-R agents to interact with semantic information repositories (e.g., RDF triple stores) that can be queried using semantic query languages such as RDQL or SPARQL. Although previous work has demonstrated the feasibility of interfacing ACT-R agents with relational databases (as a means of extending declarative memory) [40], the nature of the interaction mechanisms required for active probing capabilities are somewhat more complex. At a minimum, ACT-R agents need to be able to formulate queries and process query resultsets, and this is probably best accomplished via the development of a specialized query module within ACT-R. The LarCT-R extension of jACT-R (a Java implementation of ACT-R)[2] attempts to combine ACT-R with Semantic Web technology, and some of the work associated with this past initiative may be relevant to the current modeling effort.

Of course, the ability to formulate queries is only one of the challenges associated with the active probing of the information environment. Another challenge concerns the implementation of cognitive mechanisms that initiate and direct the query formulation process. Both the decision to construct queries, as well as the nature of the queries that get constructed, are likely to be knowledge-intensive processes worthy of independent cognitive analysis and modeling.

In addition to an ability to interact with external information repositories, communication with other agents is also an important consideration when it comes to the active probing of the information environment. Other ACT-R agents may be the source of relevant information, and requests for information to other agents may be assumed to be a key feature of many collaborative processes. This highlights the need for a consideration of agent communication and interaction mechanisms as part of collective sensemaking processes (see Section IV-A).

[2]See http://larctr.sourceforge.net/.

E. Cognitive Diversity

One of the advantages of military coalitions is that they bring together individuals from different cultural backgrounds, with different levels of expertise and experience. From a sensemaking perspective, such differences in knowledge and experience may be important in terms of both hypothesis generation and the mitigation of biases. In the case of face-to-face groups, for example, Schulz-Hardt et al [41] have shown that confirmation bias is exacerbated in homogenous groups and attenuated in heterogeneous ones. Similarly, Convertino et al [42] showed that heterogeneous groups exhibited less confirmation bias than homogenous groups when using a tool developed to support collaborative intelligence analysis. Such studies suggest that the cognitive diversity of a team may be an important factor in determining the team's resistance to factors (e.g., cognitive bias) that subsequently limit the team's performance.

Issues of cognitive diversity could be approached in a number of ways using ACT-R. Firstly, agents could be initialized with different bodies of background knowledge (asserted in declarative memory). This would help to reflect the differences in expertise and experience that individual analysts bring to the sensemaking task within particular domains. Secondly, individual agents could be developed with different bodies of procedural knowledge (implemented as production rules in ACT-R). Again, this would help to introduce diversity in the agent community by virtue of the different inferences that agents make as part of the sensemaking process. Finally, agents could be led to favor different hypotheses by manipulating their exposure to initial information items. This was, in fact, one of the strategies used by Convertino et al [42] in creating heterogeneous groups: they began the experiment by exposing group members to evidence favoring different hypotheses. Interestingly, this is perhaps one example where a specific form of cognitive bias, namely anchoring, could prove of adaptive relevance to group-level cognitive processes; it essentially provides a means by which differences in opinion can be established within a team in order to mitigate against other forms of cognitive bias.

F. Differential Information Access

Information access and sharing in coalition environments is often limited by security constraints. Information that may be easily accessed by one analyst may not be accessible to another, and, even if it is accessible, it may be degraded due to obfuscation strategies. There is considerable interest in what effect this differential information access has in terms of collective sensemaking. Although, we might expect restricted information access to undermine collective sensemaking efforts, it is also possible that security policies might adaptively regulate the flow of information through a community, forcing different analysts to work on different subsets of information and thereby ensuring greater levels of information coverage. Exposure to different bodies of information may also lead to greater diversity regarding initial hypotheses and interpretations, which may serve to mitigate group-level biases (see [41]). Finally, differential information access may reduce the extent to which analysts are exposed to common bodies of information, thereby limiting the impact of shared information on group decision making [43, 44].

G. Collaborative Technologies

The development of collaborative technologies to support sensemaking in a number of task areas is the focus of both recent and ongoing development efforts [42, 45]. In the case of collaborative intelligence analysis, for example, an environment, called CACHE, has been developed to support analysts in processing evidentiary information and sharing information with other analysts. Studies suggest that the use of this environment can be used to mitigate against cognitive biases, such as confirmation bias, although factors such as the initial cognitive diversity of the team (i.e., the extent to which all team members are biased towards the same solution) may affect the extent of bias mitigation [42].

CACHE provides a number of features that may be important to incorporate in cognitive computational models of collective sensemaking processes. Firstly, CACHE features the use of representations (e.g., matrices) that support analysts in evaluating hypotheses with respect to particular bodies of evidence. These representations can be viewed by other users in order to promote awareness of what information is being processed and how this information is contributing to analytic outcomes. Recent extensions of CACHE, in the form of a Bayes Community version of CACHE (or CACHE-BC), support additional features, such as the ability to treat other analyst's hypotheses as part of the evidence base for one's own analytic efforts [45].

The features of collaborative technologies change the nature of the sensemaking process at both an individual and collective level, and it is thus important to consider these in the context of computational modeling efforts. The sharing of local representations (containing original evidence, analyst inferences and emerging hypotheses) may be particularly important in terms of enabling a team to adequately distribute the effort associated with the processing of large bodies of relevant information. Within ACT-R such 'workspaces' could be implemented as agent-level modules that are accessible to other agents within the immediate network neighborhood.

One motivation for incorporating the features of technological environments into an ACT-R model of collective sensemaking is that simulation studies using the model can be used to help both inform the design of future interventions and evaluate the impact of existing solutions. They can also help to reveal complex interactions between factors that subtend the cognitive, social, informational and technological domains. In this respect, a number of studies have shown that the benefits of using technology in collaborative situations are often linked to other factors such as information load, opinion distribution and the extent of initial biasing [42, 46, 47].

V. Approach and Motivation

As mentioned in Section II, the requirement to draw on the resources of the social environment in making sense of complex situations is something that is common to many spheres of activity, including conventional intelligence analysis, business intelligence and scientific analysis. In spite of this commonality, however, it is likely that any model of collective sensemaking will need to account for the specific mix of technological, social, informational and cognitive resources that are recruited in specific organizational and task contexts. Collective sensemaking is likely to be a process that is heavily influenced by a number of factors, such as the nature of technological support, the organizational context in which sensemaking occurs, and the goal of the sensemaking process. In view of this, we do not expect to develop a *universal* model of collective sensemaking. Rather, our aim is to develop a framework that incorporates the kind of capabilities alluded to in Section IV (e.g., the capacity to instantiate teams of ACT-R agents that must deal with dynamic information environments and differential information access policies) and then demonstrate the application of the framework in a specific organizational and task context, namely coalition-based collective sensemaking. We suggest that the different sociotechnical ecologies of collective sensemaking processes bear enough in common for the framework to be of generic relevance to a range of different application areas. For instance, we would expect the framework to support modeling activities in respect of medical sensemaking, forensic crime analysis, and business intelligence analysis, thereby expanding the applicability of the current work beyond the military domain.

The value of using cognitive architectures in the context of multi-agent simulations is that they increase the cognitive fidelity and sophistication of the simulations being performed. This is clearly important in situations where collective behavior and team performance outcomes are influenced by cognitive processes. By using an influential and widely exploited cognitive architecture, such as ACT-R, we aim to develop predictive models that will highlight the potential impact of different sociotechnical interventions on collective sensemaking performance. One of the main points of interest, here, concerns the role that specific factors play in accentuating or mitigating cognitive biases (i.e., systematic departures from an ideal standard of reasoning). As mentioned above, previous work using ACT-R has focused on cognitive biases [30], and this is an important source of information regarding our understanding of cognitive biases at the individual level. Our own work seeks to extend this body of existing research by examining how factors at the informational, social, and technological levels interact with cognitive mechanisms in order to influence the expression of cognitive biases at the collective level[3].

VI. Model Evaluation

Ultimately, we aim to evaluate the outcomes of the modeling effort within a military coalition context. One of the focus areas of interest, in this respect, concerns the Warfighter Associate (WA) intelligent agent decision support tool, which is being developed by Veloxiti[4] [48]. WA technology provides a knowledge base to support the maneuver, intelligence, fires, and communications domains; intelligent agents to monitor services, such as Command Post of the Future (CPOF) and the Publish and Subscribe Server (PASS); and natural language processing capabilities to filter text from multiple chat rooms [48]. The WA intelligent agent capability alerts the maneuver, intelligence, fires, and communications staff to domain specific information consistent with the Commander's intent and

[3]In addition to the nature and extent of cognitive biases, other dependent variables of interest include the extent of information coverage, correspondence of sensemaking outcomes with ground truth, the temporal profile of cognitive convergence and (perhaps) the degree of cognitive load experienced by individual agents in particular experimental contexts.

[4]see http://www.veloxiti.com/.

then portrays this information within a Common Operational Picture specific to their staff roles. The WA and its data analysis component have been used in human-in-the-loop experimentation for the US Army Tactical Human Integration of Networked Knowledge program to capture operator behaviors. The system could also be used in experimentation efforts to validate the ACT-R models and multi-agent simulations consistent with the US Army Research Laboratory Model-Test-Model paradigm. It may also be possible to apply the WA knowledge base to facilitate the development of ACT-R models of individual cognitive functions and collective task performance. Finally, the WA's natural language processing capabilities for filtering chat messages may be augmented by CE facts, represented as declarative memory in ACT-R, in order to assist humans with sensemaking activities (see Section IV-A).

VII. Conclusion

Cognitive social simulations are particularly well-suited to advancing our understanding of cognition in socially-distributed or socially-situated contexts. Although agent-based models have been applied to the military domain previously, the use of cognitive social simulations, enabled by cognitive architectures, in this area is quite new. One of the advantages of such models is that they enable us to explore the complex relationships that exist between factors at the informational, technological, social and cognitive levels. We can, for example, in the case of sensemaking, investigate the impact of specific aspects of the sociotechnical ecology on sensemaking performance at both the individual and collective levels. Cognitive social simulations also enable us to explore the role that specific cognitive mechanisms (such as mnemonic decay rates - see Section IV-C) play in influencing collective behavior and team performance outcomes.

When applied to the domain of coalition-based collective sensemaking, a number of challenges confront the attempt to develop a model that can be used in cognitive social simulations. These challenges motivate our research into how to develop a framework that incorporates various features of the military coalition environment into a multi-agent system based around the use of ACT-R agents. The development and use of models created with this framework should shed light on the impact that different sociotechnical environments have on the dynamics of collective sensemaking processes within a variety of application domains.

Acknowledgement

Research was sponsored by US Army Research laboratory and the UK 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 authors and should not be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the U.S. Government, the UK Ministry of Defence, or the UK Government. The US and UK Governments are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] N. J. Cooke, J. C. Gorman, C. W. Myers, and J. L. Duran, "Interactive team cognition," *Cognitive Science*, vol. 37, no. 2, pp. 255–285, 2013.

[2] G. R. Semin and E. R. Smith, "Socially situated cognition in perspective," *Social Cognition*, vol. 31, no. 2, pp. 125–146, 2013.

[3] P. R. Smart and K. Sycara, "Collective sensemaking and military coalitions," *Intelligent Systems*, vol. 28, no. 1, pp. 50–56, 2013.

[4] G. Theiner, C. Allen, and R. L. Goldstone, "Recognizing group cognition," *Cognitive Systems Research*, vol. 11, pp. 378–395, 2010.

[5] J. Sutton, C. Harris, P. Keil, and A. Barnier, "The psychology of memory, extended cognition, and socially distributed remembering," *Phenomenology and the Cognitive Sciences*, vol. 9, no. 4, pp. 521–560, 2010.

[6] V. Hinsz, R. Tindale, and D. Vollrath, "The emerging conceptualization of groups as information processors," *Psychological Bulletin*, vol. 121, pp. 43–64, 1997.

[7] R. Menary, *The Extended Mind*. Cambridge, Massachusetts, USA: MIT Press, 2010.

[8] P. Robbins and M. Aydede, *The Cambridge Handbook of Situated Cognition*. New York, USA: Cambridge University Press, 2009.

[9] L. A. Shapiro, *Embodied Cognition*. Abingdon, Oxfordshire, UK: Routledge, 2011.

[10] T. W. Malone, R. Laubacher, and C. Dellarocas, "The collective intelligence genome," *MIT Sloan Management Review*, vol. 51, no. 3, pp. 21–31, 2010.

[11] E. H. Chi, P. Pirolli, B. Suh, A. Kittur, B. Pendleton, and T. Mytkowicz, "Augmented social cognition," in *AAAI Spring Symposium on Social Information Processing*, Stanford, California, USA, 2008.

[12] J. Hendler and T. Berners-Lee, "From the Semantic Web to social machines: A research challenge for AI on the World Wide Web," *Artificial Intelligence*, vol. 174, pp. 156–161, 2010.

[13] M. Diehl and W. Stroebe, "Productivity loss in brainstorming groups: Toward the solution of a riddle," *Journal of Personality and Social Psychology*, vol. 53, pp. 497–509, 1987.

[14] E. Hutchins, "The social organization of distributed cognition," in *Perspectives on Socially Shared Cognition*, L. Resnick, J. Levine, and S. Teasley, Eds. Washington DC, USA: The American Psychological Association, 1991.

[15] W. Mason, F. Conrey, and E. Smith, "Situating social influence processes: Dynamic, multidirectional flows of influence within social networks," *Personality and Social Psychology Review*, vol. 11, no. 3, pp. 279–300, 2007.

[16] R. Sun, "Cognitive social simulation incorporating cognitive architectures," *Intelligent Systems*, vol. 22, no. 5, pp. 33–39, 2007.

[17] P. Thagard, "Cognitive architectures," in *The Cambridge Handbook of Cognitive Science*, K. Frankish and W. M. Ramsey, Eds. Cambridge, UK: Cambridge University Press, 2012, pp. 50–70.

[18] J. R. Anderson, *How Can the Human Mind Occur in the Physical Universe?* Oxford, UK: Oxford University Press, 2007.

[19] J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin, "An integrated theory of the mind," *Psychological Review*, vol. 111, no. 4, pp. 1036–1060, 2004.

[20] J. E. Laird, A. Newell, and P. S. Rosenbloom, "SOAR: An architecture for general intelligence," *Artificial Intelligence*, vol. 33, no. 1, pp. 1–64, 1987.

[21] G. Klein, B. Moon, and R. R. Hoffman, "Making sense of sensemaking 1: Alternative perspectives," *Intelligent Systems*, vol. 21, no. 4, pp. 70–73, 2006.

[22] D. T. Moore, "Sensemaking: A structure for an intelligence revolution," National Defense Intelligence College, Tech. Rep., 2011.

[23] P. Pirolli and D. M. Russell, "Introduction to this special issue on sensemaking," *Human-Computer Interaction*, vol. 26, no. 1-2, pp. 1–8, 2011.

[24] K. E. Weick, *Sensemaking in Organizations*. Thousand Oaks, Califfornia, USA: Sage Publications, 1995.

[25] G. Klein, S. Wiggins, and C. O. Dominguez, "Team sensemaking," *Theoretical Issues in Ergonomics Science*, vol. 11, no. 4, pp. 304–320, 2010.

[26] P. Pirolli and S. Card, "The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis," in *International Conference on Intelligence Analysis*, McLean, Virginia, USA, 2005.

[27] E. Jensen, "Sensemaking in military planning: A methodological study of command teams," *Cognition, Technology & Work*, vol. 11, no. 2, pp. 103–118, 2009.

[28] S. A. Paul and M. C. Reddy, "Understanding together: Sensemaking in collaborative information seeking," in *ACM Conference on Computer-Supported Cooperative Work*. Savannah, Georgia, USA.

[29] D. Alberts and J. Garstka, "Network centric operations conceptual framework version 2.0," U.S. Office of Force Transformation and Office of the Assistant Secretary of Defense for Networks and Information Integration, U.S. Department of Defense, Tech. Rep., 2004.

[30] C. Lebiere, P. Pirolli, R. Thomson, J. Paik, M. Rutledge-Taylor, J. Staszewski, and J. R. Anderson, "A functional model of sensemaking in a neurocognitive architecture," *Computational Intelligence and Neuroscience*, in press.

[31] D. Reitter and C. Lebiere, "Towards cognitive models of communication and group intelligence," in *33rd Annual Meeting of the Cognitive Science Society*, Boston, Massachusetts, USA, 2011.

[32] ——, "Social cognition: Memory decay and adaptive information filtering for robust information maintenance," in *26th AAAI Conference on Artificial Intelligence*, Toronto, Canada, 2012.

[33] J. Ball, C. Myers, A. Heiberg, N. J. Cooke, M. Matessa, M. Freiman, and S. Rodgers, "The synthetic teammate project," *Computational and Mathematical Organization Theory*, vol. 16, no. 3, pp. 271–299, 2010.

[34] A. Oltramari and C. Lebiere, "Extending cognitive architectures with semantic resources," in *4th International Conference on Artificial General Intelligence*. Mountain View, California, USA.

[35] R. Glinton, K. Sycara, and P. Scerri, "Exploiting scale invariant dynamics for efficient information propagation in large teams," in *International Conference on Autonomous Agents and Multi-Agent Systems*, Toronto, Canada, 2010.

[36] O. Pryymak, A. Rogers, and N. R. Jennings, "Efficient opinion sharing in large decentralised teams," in *International Conference on Autonomous Agents and Multi-Agent Systems*, Valencia, Spain, 2012.

[37] S. S. Pollock, J. K. Alt, and C. J. Darken, "Representing trust in cognitive social simulations," in *4th International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*, Maryland, USA, 2011.

[38] P. R. Smart, T. D. Huynh, D. Braines, and N. R. Shadbolt, "Dynamic networks and distributed problem-solving," in *Knowledge Systems for Coalition Operations (KSCO'10)*, Vancouver, British Columbia, Canada, 2010.

[39] A. Eck and L.-K. Soh, "Dynamic facts in large team information sharing," in *International Conference on Autonomous Agents and Multi-Agent Systems*, Saint Paul, Minnesota, USA, 2013.

[40] S. Douglass, J. Ball, and S. Rodgers, "Large declarative memories in ACT-R," in *International Conference of Cognitive Modeling*. Manchester, UK, 2009.

[41] S. Schulz-Hardt, D. Frey, C. Lüthgens, and S. Moscovici, "Biased information search in group decision making," *Journal of Personality and Social Psychology*, vol. 78, no. 4, pp. 655–669, 2000.

[42] G. Convertino, D. Billman, P. Pirolli, J. Massar, and J. Shrager, "The CACHE study: Group effects in computer-supported collaborative analysis," *Computer Supported Cooperative Work*, vol. 17, no. 4, pp. 353–393, 2008.

[43] D. Gigone and R. Hastie, "The impact of information on small group choice," *Journal of Personality and Social Psychology*, vol. 72, no. 1, pp. 132–140, 1997.

[44] G. Stasser and W. Titus, "Hidden profiles: A brief history," *Psychological Inquiry*, vol. 14, no. 3-4, pp. 304–313, 2003.

[45] J. Shrager, D. Billman, G. Convertino, J. P. Massar, and P. Pirolli, "Soccer science and the Bayes community: Exploring the cognitive implications of modern scientific communication," *Topics in Cognitive Science*, vol. 2, no. 1, pp. 53–72, 2010.

[46] A. R. Dennis, K. M. Hilmer, N. J. Taylor, and A. Polito, "Information exchange and use in GSS and verbal group decision making: Effects of minority influence," in *30th Hawaii International Conference on System Sciences (HICSS)*, Maui, Hawaii, 1997.

[47] R. Hightower and L. Sayeed, "The impact of computer-mediated communication systems on biased group discussion," *Computers in Human Behavior*, vol. 11, no. 1, pp. 33–44, 1995.

[48] M. Gresham, J. Pilipovic, E. Drucker, and L. Lafferty, "S-6 associate: A unified approach to building and managing the Network Operating Environment within the context of tactical missions and other warfighting functions," in *32nd Annual Military Communications Conference*, San Diego, California, USA, in press.