A Constraint Satisfaction Model of Collective Sensemaking: Insights into the Effect of Coalition Communication Environments on the Dynamics of Collective Cognition

Paul R. Smart
School of Electronics & Computer Science,
University of Southampton,
Southampton,
SO17 1BJ, UK
Email: ps02v@ecs.soton.ac.uk

Katia Sycara
Robotics Institute,
Carnegie Mellon University,
5000 Forbes Ave., Pittsburgh,
15213-3890, USA
Email: katia@cs.cmu.edu

Tom Klapiscak
Emerging Technology Services,
IBM United Kingdom Ltd,
Hursley Park, Winchester,
SO21 2JN, UK
Email: klapitom@uk.ibm.com

Abstract—**The socially-distributed nature of cognitive processing in military coalitions means that various features of the coalition communication environment (e.g., communication network topology) have the potential to influence collective cognitive outcomes. At the present time, however, we have little understanding of how the specific features of coalition communication environments influence the dynamics of collective cognitive processing. In order to address this issue, a computational model of one particular aspect of cognitive processing, namely collective sensemaking, is proposed. The computational model provides the basis for computer simulations that can be used to provide some initial insight into how collective cognition is affected by a variety of psychosocial and technological variables. In addition to a description of the model, some specific proposals for future simulation work are outlined. The results of experiments using this model can be used to guide decisions about what kind of real-world observational and experimental studies to perform. In addition, when combined with empirical studies that seek to validate model outputs in specific situations, the model can be used to support the development of techniques and technologies that maximize the collective cognitive potential of military coalition organizations.**

I. INTRODUCTION

Military coalitions are complex socio-technical organizations in which a variety of cognitive processes (e.g., those associated with planning and decision-making) can be seen as both socially- and technologically-distributed [1]. This view has important implications for how we think about cognition in military coalition contexts, and it also influences the approach we adopt with respect to the analysis and engineering of coalition communication environments. In particular, once we see military coalitions as instances of distributed cognitive systems (see [2]), it becomes clear that the effort to enhance cognitive performance in such environments cannot concern itself solely with forces and factors that reside inside the heads of individual human agents; instead, the project needs to broaden its scope to include those features of the coalition environment that influence cognitive processing at the collective level.

One way in which we can develop a better understanding of collective cognition in military coalition environments is by undertaking empirical studies that systematically explore the effect of various types of features (e.g., frequency of communication, communication network topology, information distribution, and so on) on particular kinds of collective cognitive outcome (e.g., the ability to make sense of uncertain, conflicting or ambiguous information). Unfortunately, the nature of the coalition environment means that studies with human subjects are both difficult to design and implement. As a result, it may be important to consider the use of computer simulation techniques (particularly those involving multi-agent systems) in which some aspect of collective cognitive processing is studied *in silico*. To this end, the current paper describes a computational model that has been developed to support empirical investigations into how collective cognition may be affected by a variety of psychosocial and technological factors. This model has been used to further our understanding of how one particular kind of collective cognitive processing, namely collective sensemaking (see Section II), might be influenced by factors such as communication network structure [3] and the frequency of inter-agent communication [4]. The current paper provides an overview of the computational model and explains how it can be adapted to explore specific research questions concerning the relationship between coalition communication environments and collective sensemaking abilities.

II. COLLECTIVE SENSEMAKING AND COALITION OPERATIONS

As mentioned in Section I, the cognitive phenomenon of interest in the current work is referred to as collective sensemaking. As its name suggests, collective sensemaking is regarded as the collective form of individual sensemaking in which multiple individuals strive to make sense of particular bodies of information, typically under situations of uncertainty or ambiguity. At the individual level, sensemaking has been defined by Klein et al [5] 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". The notion of collective sensemaking is simply the extension of this concept into the social domain. It refers to the attempts of multiple individuals to coordinate their individual processing activities in order to interpret some body of conflicting, ambiguous or uncertain information.

Sensemaking is an important focus of research attention for military coalitions because much of what military coalitions do can be regarded as a form of sensemaking. Thus, the attempt to make sense of information as a prelude to planning and decision-making is a major preoccupation of coalition forces. Sensemaking is also a concept that is central to the conceptual framework for military network-centric operations (NCO) as espoused by the U.S. Department of Defense [6]. In this case, the advent of a more "robustly networked" fighting force, resulting from the introduction of more sophisticated communication and information networks, is supposed to facilitate sensemaking at the collective level and thereby contribute to better decision-making capabilities in a joint/coalition military environment [6]. This hypothesized relationship between communication networks, collective cognitive processing and decision-making capability certainly has intuitive appeal; however, empirical studies are required in order to evaluate the hypothesis and improve our understanding of how the specific features of military coalition environments affect collective sensemaking processes.

III. UNDERSTANDING THE IMPACT OF COALITION COMMUNICATION ENVIRONMENTS ON COLLECTIVE SENSEMAKING: RESEARCH ISSUES

There are a variety of factors that may affect sensemaking performance in a coalition setting. These include:

1) **Network Structure**: There is an increasing reliance on mobile ad hoc networks (MANETS) within military coalitions, and this may significantly affect the dynamics of inter-agent communication. We need to know what effect different communication network structures have on collective sensemaking. Do some types of network structure yield better performance than others? What about situations that feature time-variant changes in communication network topology as might be encountered in MANET environments?
2) **Communication Frequency**: What is the effect of communication frequency on collective sensemaking? Is it better to enable continuous modes of inter-agent communication, or is some more restrictive communication policy to be preferred?
3) **Trust**: Different levels of trust between agents may lead to inefficiencies in information processing. For example, individuals from different groups may fail to adequately integrate available information as a result of poor trust relationships. We need to develop a better understanding of how the dynamics of inter-agent trust affect collective cognitive processing.
4) **Cultural Differences**: How should we represent cultural differences in a computational model of collective sensemaking, and what effect do these cultural differences have on collective sensemaking performance? Should sensemaking teams consist of individuals from the same cultural background, or should more culturally-diverse groupings be countenanced?
5) **Confidence**: Information transmission may be affected by the confidence individuals have in their judgements. For example, an analyst may limit communication to situations in which they have a certain level of confidence in their conclusions. How does confidence-based communication affect sensemaking performance? Should information always be communicated irrespective of confidence, or are more conservative schemes to be preferred?
6) **Information Sharing**: The sharing of information within a coalition organization may be limited for a number of reasons. Security constraints may limit information access, communication networks may limit information distribution, and differences in information technology may make information difficult to exploit and integrate into ongoing cognitive processes. What effect does partial information sharing have on sensemaking abilities?
7) **Information Distribution**: The information that is available to a coalition may be distributed across different information sources. This means that some agents (and groups of agents) may have access to bodies of information that are not immediately accessible to other agents. When agents collaborate in order to make sense of some external state-of-affairs, we may expect information sharing policies to interact with the distribution of information and affect collective sensemaking performance accordingly. It is important to understand the nature of this effect.
8) **Miscommunication**: What happens to sensemaking performance when miscommunication occurs during information transmission? How detrimental is miscommunication to collective sensemaking?
9) **Quality of Information**: Military coalitions often have to deal with information that is imperfect in a number of ways. For example, information may be incomplete, conflicting, ambiguous, and uncertain. How does collective sensemaking fare under these different informational conditions?
10) **Deception**: Coalitions often operate in an environment where hostile agencies attempt to undermine coalition decision-making. This means that there is a high risk of misinformation and deception. How vulnerable is group performance to the effects of deception, and how is this vulnerability affected by different aspects of the social, informational and technological environment?
11) **Expectations, Assumptions and Biases**: How do initial

> expectations and assumptions affect subsequent judgements in collective sensemaking? Can we gain some insight into the kind of cognitive biases that might emerge in collective sensemaking scenarios by performing computer simulation studies?

As is clear from this list, there are a broad range of research issues to consider in the context of sensemaking research. These issues extends across the technical, social, cognitive, cultural, and organizational dimensions of a coalition organization. In attempting to address these issues, we have adopted an approach that is grounded in the use of a computational model of sensemaking. Section IV provides details of this computational model, and Section V shows how the model can be adapted to address specific issues.

IV. A COMPUTATIONAL MODEL OF COLLECTIVE SENSEMAKING

The computational model developed to explore the dynamics of collective sensemaking is based on the consonance model of Schultz and Lepper [7]. The consonance model was originally developed to explore the psychological phenomenon of cognitive dissonance; however, we suggest that the model can also serve as the basis for computational analyses of collective sensemaking. There are two reasons for this.

Firstly, the consonance model is based on the use of constraint satisfaction networks (CSNs). CSNs are networks of processing units that rely on patterns of spreading activation in order to find solutions to constraint satisfaction problems (or problems in which solution outcomes need to be derived against a backdrop of positive and negative constraints) [8]. Typically, CSNs are developed so that different solution outcomes are represented by the activation levels of nodes in the network, and constraints are represented by linking the nodes together with positive or negatively weighted links. The initial state of the problem is encoded as a pattern of activation across the nodes of the network, and solutions that maximally satisfy the constraints emerge across the course of successive processing cycles. The solution to the problem is encoded in the pattern of activity across the nodes at the end of the simulation.

The use of CSNs to study constraint satisfaction problems makes them of interest in the current context because we suggest that sensemaking can be usefully viewed as a form of constraint satisfaction problem. Sensemaking is a process in which an agent attempts to form a consistent set of beliefs subject to the constraints imposed by background knowledge, initial expectations, available evidence and existing interpretations. In many ways, sensemaking bears much in common with the psychological notion of coherence, which has also been viewed as a constraint satisfaction problem (and researched using CSNs) [9]. In their analysis of the concept of coherence, Thagard and Verbeurgt [10] thus write:

> "When we make sense of a text, a picture, a person, or an event, we need to construct an interpretation that fits with the available information better than alternative interpretations. The best interpretation is one that provides the most coherent account of what we want to understand, considering both pieces of information that fit with each other and pieces of information that do not fit with each other." (pg. 2)

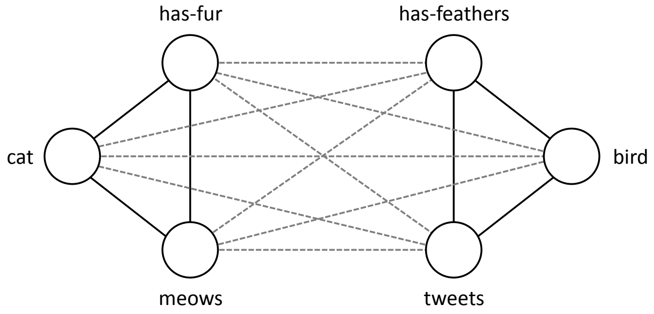

Fig. 1. Example organization of cognitive units in a single agent. The circles represent cognitive units, each of which consists of two processing nodes. Solid lines represent excitatory connections between the units, while broken lines represent inhibitory links.

A second reason why the consonance model is attractive as a sensemaking model is that dissonance reduction is seen as the primary motivating mechanism for cognitive change within the model. This aligns itself with work suggesting that cognitive consistency may be useful in understanding the way in which multiple forces (operating at the individual, social and cultural levels) combine to influence the dynamics of collective cognition [11].

The following subsections describe the various components of the consonance-based computational model of collective sensemaking.

A. Agent Cognitive Architecture

Each agent within the model is implemented as a separate CSN, an example of which is illustrated in Figure 1. The nodes in these networks are referred to as cognitive units because each unit represents a particular cognition (e.g., belief) that an agents can have[1]. The agent depicted in Figure 1 consists of six cognitive units, and it is thus capable of entertaining six beliefs.

Internally, each cognitive unit consists of two nodes which are connected together in a mutual inhibitory fashion (see Figure 2). One of these nodes is labeled as the 'positive pole' (P), and the other is labeled as the 'negative pole' (N). The difference in activation between these two nodes determines the net activation of the cognitive unit, which itself reflects the strength and direction of the cognition represented by the cognitive unit. In the case of beliefs, the polarity of cognitive unit activation indicates whether an agent believes something to be true or false, and the level of activation indicates the degree of certainty or confidence that the agent has with respect to the belief. In the case of other cognitions (e.g., attitudes),

---

[1] The term 'cognition' can be used to refer to both the beliefs and attitudes of an agent; however, in our work, we deal only with beliefs, and thus each cognitive unit represents a distinct belief.

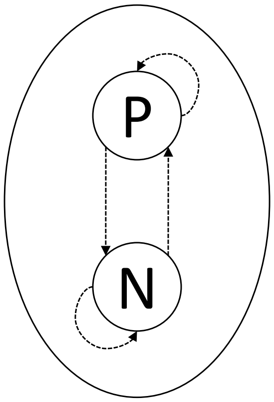

Fig. 2. The anatomy of a cognitive unit. Each unit consists of two nodes, one of which is the positive pole (P) and the other is the negative pole (N). P and N are connected via mutual inhibitory links. In addition, each node has an auto-regulatory connection that connects each node to itself. The function of this auto-regulatory link is to dampen the node's activity at each processing cycle (see [7], for details). All links within the cognitive unit are inhibitory, as is indicated by the broken lines in the figure.

the direction of activation may indicate whether something is positively or negatively evaluated, and the activation level will indicate the strength of the evaluation[2].

The positive and negative poles of the cognitive unit are connected via mutual inhibitory links, and this forms the basis for competitive interactions between the nodes of the cognitive unit[3]. The strength of cognitive unit activation (i.e., net activation) is thus a function of both the 'P' and 'N' nodes, and both of these nodes may be activated or inhibited by nodes external to the cognitive unit. As noted by Schultz and Lepper [7], this representational scheme allows for some degree of ambivalence in agent cognitions. For example, there might be some evidence supporting a particular belief and other evidence that undermines the same belief. Similarly, when it comes to evaluations and attitudes, something may be both liked and disliked at the same time. Ambivalence will tend to be reduced across successive processing cycles due to the inhibitory connections between the 'P' and 'N' nodes; however, ambivalence could persist if both the 'P' and 'N' nodes receive strong support from other cognitions.

Within an agent, cognitive units are connected to other cognitive units via inhibitory or excitatory links (see Figure 1). These links represent the relationship between two cognitions, where the notion of a relationship subsumes things like logical implications, causal relationships, expectations and associations. The decision about whether to connect cognitive units with a positive or negative link is based on the degree of compatibility or consistency between the cognitions represented by the cognitive units. In dissonance theory, two cognitions can be said to be dissonant when one follows from the obverse of the other; consonant cognitions, in contrast, occur when one cognition implies the other. In the consonance model, this notion of dissonant and consonant cognitions is used to configure the pattern of connections between cognitive units: negative implications between cognitions are represented by negatively-weighted (inhibitory) connections, and positive implications are represented by positively-weighted (excitatory) connections (unrelated cognitions are represented by zero-weighted connections)[4].

The role that inter-cognition linkages play in giving rise to consistent or compatible belief states can be exemplified with respect to the CSN shown in Figure 1. Here we see that the cognitive units are connected together in such a way as to reflect the natural association of particular features with particular objects. Thus, the 'has-feathers', 'tweets' and 'bird' units are all connected with excitatory connections. This organization is intended to reflect an agent's (admittedly limited) knowledge about cat and bird objects. If one of the units is artificially stimulated (a situation we consider as analogous to the presentation of specific kinds of evidence), then the activation of units that have positive connections to that unit will be increased across successive processing cycles. The end result is that agents will settle on belief states that are most consistent with the evidence made available to them, as well as their background knowledge of the domain in question. If we artificially stimulate the 'has-feathers' unit, for example, then the activity of the 'bird' and 'tweets' units will increase, while that of the 'cat', 'has-fur' and 'meows' units will decrease. This reflects the agent's belief that the unidentified object is a bird, which seems a perfectly sensible interpretation of the available evidence.

In addition to a sign, representing the nature of the relationship between two cognitions, each connection between cognitive units has a weighting that determines the amount of influence one cognitive unit exerts over another. These weights could assume a variety of values, and there is no reason why the weights need to be uniform either within a single agent or across multiple agents. The same two cognitions could therefore be associated with the same or different weightings in different agents. In general, we assume that the pattern of connectivity between cognitions at the level of individual agents reflects an agent's background knowledge, training or experience in some domain. This opens up the possibility that agents could have different patterns of connectivity between cognitive units based on different prior experiences. In

---

[2] The pattern of activity across the cognitive units at any point in time corresponds to the cognitive state of the agent, and each of these states can be seen as a particular point in the 'cognitive landscape' of the agent. In essence, the cognitive units define the dimensions of a multi-dimensional state space, and specific cognitive states correspond to points within this space. As the pattern of activity in cognitive units changes across successive processing cycles, the cognitive evolution of an agent can be tracked as a trajectory through multi-dimensional space.

[3] The basis for this organizational scheme is tied to reviews of the neuroscientific literature (see [7]). In particular, Schultz and Lepper [7] argue that: "neurons are sometimes organized into excitatory and inhibitory camps that respond in opposite ways to input, one group being excited and the other group being inhibited by the same input" (pg. 221).

[4] It is important to note that the nature of these connections is deemed to reflect the psychological implication between two cognitions as *perceived by the agent*. Thus, an agent might perceive there to be a logical or causal relationship between two states of the environment, whereas in reality no such relationship exists.

previous simulation work, we assumed that all agents had identical cognitive architectures, and all the inter-cognition linkages were configured by hand [3], [4]. However, it is also possible to imagine situations in which the kind of cognitions agents have, and the manner in which they are connected, is determined by some prior learning experience. In respect of this possibility, previous work has shown how the connection weights for CSNs might be learned through experience [12], and this could constitute one means by which the current model could be adapted for situations in which we want to explore the effect of individual and group-level (e.g., cultural) differences on sensemaking performance.

B. Agent Communication

Agents are organized into communication networks via the inclusion of linkages between agents. These linkages represent channels of communication and influence between agents, and they enable us to explore the role of social interaction in mediating cognitive change. Each inter-agent connection is in fact a set of bidirectional linkages between the corresponding cognitive units of communicating agents. Thus, if agent A communicates with agent B, then connections will exist between the respective 'P' and 'N' nodes of the corresponding cognitive units in each agent. This pattern of connectivity means that when agents communicate, they exchange information about their beliefs along belief-specific channels of communication.

As with the connections within an individual agent, the connections between agents can have both a sign and a weighting value. In previous work [3], [4], we used a single weighting value of 0.5 for all inter-agent connections; however, it is possible for more complex patterns of communication to be envisaged. For example, particular communities of agents may be endowed with strong intra-group and weak extra-group channels of communication. Since the weighting of the inter-agent connections determines the degree to which one agent influences another, and agents influence one another to different degrees based on a variety of factors (e.g., the level of trust that exists between them), we can use the connectivity patterns between agents to represent situations in which agents exert variable influences over one another based on a variety of criteria (e.g., cultural similarity). We can also allow the weights to change dynamically throughout the course of a simulation in order to model situations where communication and influence grow progressively stronger (or weaker) over time. Such situations are of particular interest in light of the finding that dynamic communication networks may yield a performance profile that differs from that seen in the case of more statically-configured networks [13].

The channels of communication between agents determines the topological structure of the communication network. For example, if all agents communicate with all other agents, then the communication network structure will be fully-connected. By varying the linkages between agents, therefore, we can begin to examine how factors like the structure of communication networks might affect collective sensemaking performance (see [14]). In fact, the situation is somewhat more complicated than this because there is no requirement for agents to communicate *all* information about their cognitive states to other agents; in some cases, agents may limit communication to a particular subset of their cognitions. This means that different cognitions within the same agent may be differentially accessible to different communities of agents. Some cognitions may be totally private (i.e., an agent never establishes communication links with an agent about these cognitions), while others may be publicly accessible to different groups of agents (e.g., members of the same social group). The fact that cognitions may be shared to a greater or lesser extent means that we can study situations in which information sharing is restricted to certain kinds of information. This reflects the situation in military coalitions where information sharing may be limited for technological, social, or security-related reasons. It is even possible to imagine situations in which some agents may deliberately try to subvert the sensemaking abilities of other agents by broadcasting false information about their cognitive states.

In most simulations, agents will be assumed to have the same kind of cognitive units, and any actual or potential communication will be assumed to occur through linkages that connect cognitive units of the same type. By connecting corresponding cognitive units together, we are effectively creating a situation where agents always understand the content of messages transmitted by another agent. Even in situations where the talking agent is deliberately misinforming the listening agent, we are still assuming that listening and talking agents both understand what is being talked about. In real-world situations, however, the content of specific messages may be misunderstood as a result of linguistic or cultural differences [15]. Such misunderstanding can be represented in the current model by including linkages between different cognitive units (i.e., cognitive units representing different kinds of beliefs).

C. Computational Processing

At the beginning of a simulation, the activation levels of certain nodes within each agent can be specified via an initial activation vector. This initial activation is deemed to represent an agent's cognitive state at the outset of the simulation; i.e., it establishes the pre-existing beliefs and attitudes held by an agent. In most collective sensemaking scenarios, the activation level of certain nodes may also be influenced by 'environmental information' that is external to the agent community. This information represents the information that agents receive from a variety of information sources (e.g., sensor systems). In previous work, such information has always been supplied at the outset of a simulation (e.g., [4]); however, information could also be supplied at certain points during the course of a simulation, and, indeed, this may provide a better approximation to the situation faced by military analysts in real-world sensemaking scenarios.

Once the activation levels of nodes have been established, computational processing occurs via the spreading of activation between the cognitive units of the CSN following the pattern of excitatory and inhibitory linkages between the units. At each processing cycle in the simulation, the activation of each node in the CSN is updated according to the following rules:

$$a_i(t+1) = a_i(t) + net_i(ceiling - a_i(t)) \tag{1}$$

when $net_i \geqslant 0$, and

$$a_i(t+1) = a_i(t) + net_i(a_i(t) - floor) \tag{2}$$

when $net_i < 0$.

In these equations, $a_i(t+1)$ is the activation of node i at time $t+1$, $a_i(t)$ is the activation of node i at time t, $ceiling$ is the maximal level of activation of the node (Schultz and Lepper [7] use a value of 1.0 for 'P' nodes and a value of 0.5 for 'N' nodes), $floor$ is the minimum activation of the node (zero for all nodes), and net_i is the net input to node i, which is defined as:

$$net_i = resist_i \sum_j w_{ij} a_j \tag{3}$$

where a_j is the activation of node j that is connected to node i, w_{ij} is the weighting associated with the connection between i and j, and $resist_i$ is a measure of the resistance of node i to having its activation changed. In general, the smaller the value of this parameter, the greater the resistance to activation change, and thus the greater the resistance to cognitive change. One possible use of this parameter is to make certain cognitions more or less resistant to change than others.

At each point in the simulation, n nodes are randomly selected and updated according to Equations 1 and 2, where n corresponds to the number of nodes in the CSN. Agents are then allowed to communicate information to their connected peers (i.e., their immediate neighbors in the communication network). Communication involves each agent contributing activation to connected agents based on the activation levels of their own constituent nodes. Each node is associated with a parameter, $comminput_i$, which is the weighted sum of activation received from all talking agents. This parameter is then updated according to the following equation:

$$comminput_i = \sum_j W_{ij} A_j \tag{4}$$

where A_j represents the activation value of a node in the talking agent and W_{ij} represents the weight of the connection from node j (in the talking agent) to node i (in the listening agent).

At the next processing cycle, $comminput_i$ is incorporated into the activation equations by extending Equation 3 as follows:

$$net_i = resist_i(\sum_j w_{ij} a_j + comminput_i) \tag{5}$$

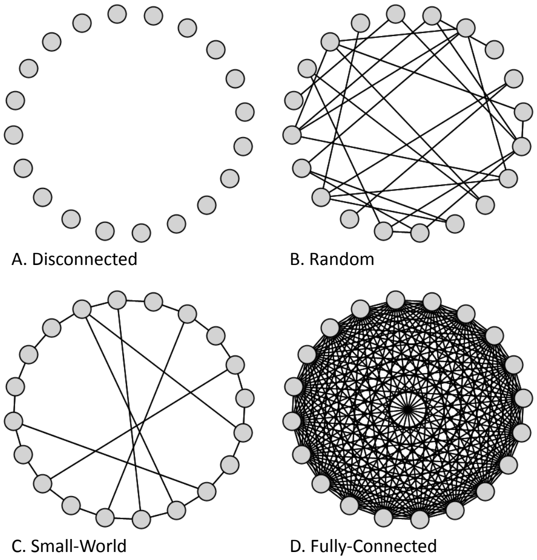

Fig. 3. Examples of the network structures used in the experiments by Smart [14]. Nodes represent agents and lines indicate channels of communication.

Once the communicated activation has been incorporated into the node's current activation level, $comminput_i$ is reset to zero in order to avoid repetitive presentation of the same communicated information across successive processing cycles.

V. EXPLORING THE DYNAMICS OF COLLECTIVE SENSEMAKING IN COALITION ENVIRONMENTS

Section III highlighted a number of issues that could be the target of research into collective sensemaking. Table I indicates how these research issues could be explored using the computational model described in Section IV.

To date, the model described here has been used to examine the effect of factors such as communication frequency and communication network structure on collective sensemaking [4], [3], [14]. In one study, for example, Smart and Shadbolt [4] examined the effect of a number of communication variables on sensemaking performance. The results obtained from this study suggested that precipitant forms of information sharing may result in agents assigning undue significance to information that is largely consistent or compatible with preexisting or prevailing cognitions.

In another study, Smart [14] explored the effect of different communication network structures on minority influence in collective sensemaking situations. This study used four types of communication network topology, examples of which are illustrated in Figure 3. The results from this study suggested that network topology can influence the extent to which minority views are integrated into collective judgements. In particular, fully-connected networks deliver a performance profile in which minority influence is minimized in situations where both minority and majority groups are exposed to weak evidence. However, the same networks serve to maximize minority influence when minority group members are selectively

TABLE I
EXPLORING THE DYNAMICS OF COLLECTIVE SENSEMAKING IN COALITION ENVIRONMENTS.

Research Issue	**Experimental Approach**
Network Structure	Given that the connections between the cognitive units in different agents represent channels of communication, the pattern of connectivity between agents can be used to study the effect of communication network structure on sensemaking abilities. We can thus examine sensemaking with different network structures (e.g., small-world, random and fully-connected networks).
Communication Frequency	Irrespective of the communication network structure that exists in any given simulation, communication need not be enabled on every processing cycle. Instead, we can limit communication to particular cycles or stages of a simulation. We can thus look at collective sensemaking performance under conditions of both low and high frequency communication.
Trust	Trust is represented by the weighting associated with specific inter-agent connections, and different values can be specified for each channel of communication between two agents (i.e., agents may have different levels of trust in respect of particular kinds of beliefs). This means that we can vary the weighting between agents in order to examine the effect of trust on sensemaking processes.
Cultural Differences	Cultural differences can be represented by creating groups of agents with different kinds of connectivity between cognitive units. For example, a strong positive connection between two cognitive units in agents of one cultural group may exist as a weak negative connection between agents in a different cultural group. This will lead agents to process information in different ways and come to different conclusions given the same body of external information. This may or may not be beneficial to collective sensemaking performance. Ideally, experimental simulations would use cultural groups in which the weighting associated with cognitive unit connections had been acquired as a result of exposure to different bodies of training data.
Confidence	Since the level of activation of a cognitive unit can be seen as providing a measure of confidence, the effect of communication under different levels of confidence can be studied by introducing a threshold for communication. For example, we may impose a constraint that confidence (i.e., net activation) must exceed a certain threshold before communication can take place.
Information Sharing	The effect of information sharing can be studied by varying the connections that exist between particular groups of agents. For example, we could have a situation where the connections between two agents are limited to certain types of cognitive units. This corresponds to a situation where the agents communicate freely about certain cognitions but not about others. We could also explore situations in which information sharing is influenced by group membership. For example, communication between groups may be more restrictive compared to that that exists within groups.
Information Distribution	The effects of information distribution can be studied by presenting different agents with different kinds of information (i.e., external activation) at the beginning and throughout a simulation. This distribution could be organized with respect to group membership criteria such that agents in different groups are selectively exposed to different bodies of information.
Miscommunication	Miscommunication can be represented in the model by connecting different cognitive units in different agents together either on a permanent or intermittent basis. Under normal circumstances, connections should exist between cognitive units that represent the same kind of cognition. If a connection exists between cognitive units that represent different kinds of cognition, then communication will cause the listening agent to have a cognitive state that is different from that intended by the talking agent.
Quality of Information	Quality of information is reflected in the kind of information agents receive from the environment. When agents have access to perfect information about the environment, they will receive external activation that coincides with reality (i.e., ground truth). More often than not, however, the information agents receive will not be perfect; it will be distorted in various ways. Agents may thus be exposed to information on some processing cycles that conflicts with that presented on earlier cycles, or they may be presented with information that conflicts with that presented to other agents. Agents may also be presented with ambiguous information that is consistent with two or more interpretations. In all these situations, information quality is determined by the profile of activation vectors that are supplied to agents at the beginning and throughout the simulation.
Deception	In some situations, agents may attempt to misinform other agents in order to subvert the sensemaking process. This situation can be studied in the computational model by creating agents whose communicative output is subject to distortion. For example, the polarity of the activation in some cognitive units might be reversed whenever communication takes place with certain other agents.
Expectations, Assumptions and Biases	Assumptions, expectations and biases can be represented by pre-activating certain cognitive units at the beginning of a simulation. We might expect this to lead to particular biases in cases where ambiguous information is subsequently presented (e.g., agents may converge on an interpretation that is consistent with their initial expectations). The notion of cognitive resistance (as represented by the resistance parameter - see Section IV-C) is likely to be important here. By varying the extent to which agents are willing to change their beliefs in the face of conflicting evidence, we may be able to model the discounting of contradictory information that takes place in at least some sensemaking scenarios.

exposed to strong evidence. These results suggest that fully-connected networks differentially regulate minority influence based on the kinds of evidence presented to both minority and majority group members.

Obviously, the results that are obtained from these sorts of simulation experiments should be treated with some caution. In particular, we should avoid generalizing from the results of computer simulation studies to the real-world without first validating model outputs using human-based experiments. The main advantage of the computational model is that it enables us to run a large number of experiments each exploring different combinations of variables, and this is something that is obviously difficult or impossible to perform in real-world coalition environments. The results from computer simulation experiments can be useful in terms of guiding decisions about the kind of real-world observational or experimental studies to perform. In particular, they can be used to generate specific hypotheses concerning collective cognitive performance, and these hypotheses can subsequently be evaluated in the context of real-world empirical analyses. By combining the results from both real-world experiments and computer simulation studies we can hope to derive some insight into how collective cognition is affected by the features of coalition communication environments. This, in turn, will help to guide research and development in other areas (e.g., the development of physical communication networks, the development of sensor information processing and delivery mechanisms, etc.) in ways that enable military coalitions to optimally configure their informational, technological and human resources from a collective cognitive processing perspective. This is really the *sine qua non* of military coalition research at the organizational level: it is an attempt to understand how best to 'wire up' the informational, technological and human elements of a military coalition in order to maximize collective cognitive capabilities.

VI. CONCLUSION

Cognitive processing in military coalitions is often a collective endeavor. More often than not, activities such as planning and decision-making require the coordinated effort of multiple individuals working in conjunction with a variety of props, aids and artefacts. The socially-distributed nature of cognitive processing in military coalitions means that various features of the coalition communication environment (e.g., communication network topology) have to potential to influence collective cognitive outcomes. At the present time, however, we have little understanding of how cognition may be influenced by these social and technological factors. The current paper has presented a computational model of one particular aspect of cognitive processing, namely collective sensemaking. This modeling effort provides the basis for computer simulations that can be used to provide some initial insight into how collective cognitive outcomes are influenced by features of the coalition communication environment. When combined with empirical studies that validate the results of specific simulation efforts, the output from this work can be used to support the effort to develop supportive techniques and technologies that maximize the collective cognitive processing potential of military coalition organizations.

ACKNOWLEDGMENT

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] P. R. Smart, P. C. Engelbrecht, D. Braines, M. Strub, and C. Giammanco, "The network-extended mind," in *Network Science for Military Coalition Operations: Information Extraction and Interaction*, D. Verma, Ed. Hershey, Pennsylvania, USA.: IGI Global, 2010.

[2] E. Hutchins, *Cognition in the Wild.* Cambridge, Massachusetts, USA: MIT Press, 1995.

[3] P. R. Smart, "A computational model of collective sensemaking: Differential effects of communication network structure on collective sensemaking abilities," in *11th International Conference on Cognitive Modeling (ICCM'12)*, Berlin, Germany, 2012.

[4] P. R. Smart and N. R. Shadbolt, "Modelling the dynamics of team sensemaking: A constraint satisfaction approach," in *Knowledge Systems for Coalition Operations*, Pensacola, Florida, USA, 2012.

[5] 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.

[6] 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.

[7] T. R. Schultz and M. R. Lepper, "Cognitive dissonance reduction as constraint satisfaction," *Psychological Review*, vol. 103, no. 2, pp. 219–240, 1996.

[8] J. J. Hopfield, "Neural networks and physical systems with emergent collective computational abilities," *Proceedings of the National Academy of Sciences*, vol. 79, pp. 2554–2558, 1981.

[9] P. Thagard, "Explanatory coherence," *Behavioral and Brain Sciences*, vol. 12, no. 3, pp. 435–502, 1989.

[10] P. Thagard and K. Verbeurgt, "Coherence as constraint satisfaction," *Cognitive Science*, vol. 22, no. 1, pp. 1–24, 1998.

[11] P. R. Smart, W. R. Sieck, D. Braines, T. D. Huynh, K. Sycara, and N. R. Shadbolt, "Modelling the dynamics of collective cognition: A network-based approach to socially-mediated cognitive change," in *4th Annual Conference of the International Technology Alliance (ACITA'10)*, London, UK, 2010.

[12] D. Ackley, G. Hinton, and T. Sejnowski, "A learning algorithm for boltzmann machines," *Cognitive Science*, vol. 9, no. 1, pp. 147–169, 1985.

[13] T. D. Huynh, P. R. Smart, D. Braines, N. R. Shadbolt, and K. Sycara, "The cognitive virtues of dynamic networks," in *4th Annual Conference of the International Technology Alliance (ACITA'10)*, London, UK, 2010.

[14] P. R. Smart, "Studying the impact of minority views in a computational model of collective sensemaking: The role of network structure," in *4th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE'12)*, Nice, France, 2012.

[15] S. Poteet, P. Xue, J. Patel, A. Kao, C. Giammanco, and I. Whiteley, "Linguistic sources of coalition miscommunication," in *NATO RTO HFM-142 Symposium on Adaptability in Coalition Teamwork*, Copenhagen, Denmark, 2008.