

Exploring the Dynamics of Collective Cognition Using a Computational Model of Cognitive Dissonance

Paul R. Smart^a, Katia Sycara^b and Darren P. Richardson^a

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

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

ABSTRACT

The socially-distributed nature of cognitive processing in a variety of organizational settings means that there is increasing scientific interest in the factors that affect collective cognition. In military coalitions, for example, there is a need to understand how factors such as communication network topology, trust, cultural differences and the potential for miscommunication affects the ability of distributed teams to generate high quality plans, to formulate effective decisions and to develop shared situation awareness. The current paper presents a computational model and associated simulation capability for performing *in silico* experimental analyses of collective sensemaking. This model can be used in combination with the results of human experimental studies in order to improve our understanding of the factors that influence collective sensemaking processes.

Keywords: collective cognition, social information processing, distributed cognition, cognitive dissonance, social influence, constraint satisfaction network, collective sensemaking, sensemaking, team sensemaking

1. INTRODUCTION

The socially-distributed nature of tasks such as intelligence analysis and military planning means that there is increasing scientific interest in how collective cognitive capabilities are affected by the features of collaborative work environments. In military coalitions, for example, there is a need to understand how factors such as communication network topology, trust, cultural differences and the potential for miscommunication affect the ability of distributed teams to generate high quality plans, to formulate effective decisions and to develop shared situation awareness.

Typically, research into collective cognition has been undertaken using two very different approaches. Firstly, social psychological research has tended to focus on small groups of human individuals who are observed in a particular task context. The advantage of this approach is that it uses real human subjects to test research hypotheses. However, the research also suffers from a number of disadvantages: it is often limited to small groups with minimally complex communication structures, and it focuses on tasks that do not necessarily generalize to real-world situations. The other approach to collective cognition research has focused on the use of computer simulation techniques. This approach is potentially well-suited to the study of collective cognitive processes in organizational settings because it avoids some of the difficulties associated with the running of large-scale social experiments. It also enables the research community to explore ‘what-if’ scenarios involving the systematic manipulation of variables that would be difficult or impossible to manipulate in real-world situations. Despite the advantages, however, there are a number of problems associated with the use of computer simulations to investigate collective cognition. One problem concerns the psychological interest and relevance of the computational models used in the studies. In many cases, for example, the agents used in simulation studies are represented by single, time-variant numerical values, and they lack any kind of internal cognitive processing complexity. This has encouraged some to argue in favor of *socio-cognitive* simulations in which agents are endowed with psychologically-realistic cognitive processing capabilities.¹

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

Paul R. Smart: E-mail: ps02v@ecs.soton.ac.uk, Telephone: +44 (0)23 8059 6669

In the current paper, we describe a computational model of collective cognition that is grounded in an extensive body of work in the psychological literature. The model features the use of constraint satisfaction networks (CSNs) to model one particular form of collective cognition, namely collective sensemaking. Collective sensemaking is a phenomenon of considerable importance in a number of different task contexts, such as intelligence analysis,² military planning³ and healthcare provision.⁴ It tends to emerge in any situation where a group or team of individuals is required to pool their cognitive resources in the attempt to interpret conflicting, complex, incomplete and uncertain bodies of data. Collective sensemaking is of particular importance to military coalitions because sensemaking processes sit at the heart of the vision of network-centric operations (NCO). According to the NCO Conceptual Framework, for example, sensemaking at both the individual and collective levels has a direct impact on decision synchronization, force agility and mission effectiveness.⁵ This makes research into collective sensemaking of vital importance to military coalition organizations.⁶

The main aim of the current paper is to provide an overview of the aforementioned computational model of collective sensemaking (see Section 3). The computational model is based on a previous CSN-based model, which has been used to explore the psychological phenomenon of cognitive dissonance.⁷ The relevance of both cognitive dissonance and CSNs to the modeling of sensemaking processes is discussed in Section 2. The paper also presents a simulation capability that has been developed to support users in configuring and executing experimental simulations (see Section 4). This capability can be used by researchers to undertake their own investigations into collective sensemaking. Finally, Section 5 describes a number of ways in which the computational model could be used to investigate the impact of factors such as trust, communication network topology, miscommunication and cultural differences on collective sensemaking processes.

2. COLLECTIVE SENSEMAKING, COGNITIVE DISSONANCE AND CONSTRAINT SATISFACTION NETWORKS

2.1 Collective Sensemaking

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’⁶ or ‘team sensemaking’;⁸ i.e., the activities that are performed by groups of individuals in order to develop understanding at both the individual and collective levels. Collective sensemaking is an example of socially-distributed cognition, or collective cognition, in which the processes that yield an understanding of some state-of-affairs are distributed across multiple individuals. As with other cases of distributed cognitive processing, there are a number of reasons why collective sensemaking may be important in an organizational context. One reason is that the complexity of the available information may be such that no one individual has the relevant knowledge or expertise to properly interpret it. 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 medical diagnosis, for example, one individual may rely on another to provide specialist diagnostic data, while in intelligence analysis, 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.

2.2 Constraint Satisfaction Networks

Our approach to the computational modeling of sensemaking processes is based on the use of CSNs, which are typically used to study constraint satisfaction problems. Using CSNs, a constraint satisfaction problem can be modeled as a set of inter-connected nodes, with the weight assigned to the linkages between nodes indicating the

relative importance of constraints that must be satisfied as part of the problem-solving process. The problem-solving process itself typically assumes the form of spreading activation between the nodes, and the solution to the problem is given as the activation vector across all nodes after a specified number of processing cycles has elapsed.

There are a number of reasons that motivate the use of CSNs to study sensemaking processes. One reason relates to the fact that sensemaking (or at least some parts of the sensemaking process*) can be cast as a form of constraint satisfaction problem. In particular, the ability to interpret (make sense of) information (especially when that information is incomplete, uncertain, ambiguous or conflicting) can be seen as a form of constraint satisfaction in which an agent attempts to establish a consistent or coherent set of beliefs subject to the constraints imposed by background knowledge, initial expectations and received information. The notion of cognitive consistency or coherence is important here because CSNs have been used to model psychological processes in which issues of coherence and consistency play a major role.⁹ As an example, CSNs have been used to study the phenomenon of cognitive dissonance, in which the process of cognitive change (e.g., belief modification) is driven by a need for consistency or compatibility between particular cognitions (e.g., beliefs).⁷

A second reason to countenance CSNs relates to the previous use of such networks in psychological modeling contexts. CSNs have been used to model a wide variety of psychological phenomena, including belief revision, impression formation, explanation, schema completion, analogical reasoning, causal attribution, discourse comprehension, content-addressable memories, cognitive dissonance and attitude change.¹⁰⁻¹⁵ This grounding in psychological research helps to improve the psychological plausibility and relevance of simulation results, which is an issue of considerable importance for those interested in developing computational models of socially-situated or socially-distributed cognition.^{1,16} Sun,¹ for example, claims that “Most of the work in social simulation assumes rudimentary cognition on the part of agents” (pg. 34). He advocates the use of psychologically-rich models to improve the cognitive fidelity of computer simulations and improve our understanding of real-world collective behavior and social phenomena.

A third reason motivating the use of CSNs centers on the previous use of such networks to explore the dynamics of collective sensemaking. In particular, Hutchins¹⁷ has used CSNs to examine the psychological phenomenon of confirmation bias; i.e., the tendency to ignore or discount evidence that contradicts some initial interpretation of a situation. Using CSNs, Hutchins was able to show that the timing of inter-agent communication exerts a significant influence on the dynamics of collective sensemaking. In particular, if the individual agents were allowed to communicate with one another from the outset of a simulation, then extreme levels of confirmation bias arose. This occurred because each agent, under the influence of information provided by other agents in the social network, was under pressure to discover a shared interpretation of the input data. In other words, the community of agents strove to find a set of activation patterns that best satisfied the internal constraints established by inter-agent communication. The result was that agents often failed to give due weight to the evidence provided by external input, and thus, more often than not, the community of agents tended to exhibit more extreme forms of confirmation bias than was the case with isolated individuals.

2.3 Cognitive Dissonance and Sensemaking

In using CSNs to model sensemaking processes, a variety of decisions need to be made concerning the specific kind of CSN to use for modeling purposes. The CSNs used in the current model (see Section 3) are based on those developed by Schultz and Lepper⁷ as part of their work into the phenomenon of cognitive dissonance.^{18,19} Cognitive dissonance is characterized as a state of psychological discomfort or tension that arises as a result of a mismatch or conflict between cognitions and observed behavior. In order to resolve the state of psychological

*It should be noted that the current model does not capture all the richness and complexity of the sensemaking process as seen in real-world situations. In particular, it does not consider the role of knowledge-guided interactions with the real-world in order to make sense of conflicting, uncertain or ambiguous information. In most cases of real-world sensemaking, human agents actively engage with the external environment in order to support their ongoing attempts at pattern recognition and pattern completion. Human agents engaged in sensemaking do not just 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.

tension, the subject is motivated to either change their behavior or (more typically) cognitions in order to re-establish consistency.

The notion of cognitive dissonance has been used to account for the cognitive changes observed in subjects who engage in counterattitudinal behavior. Linder et al,²⁰ for example, asked college students to write a forceful essay supporting a ban on communist speakers - a position with which the students strongly disagreed. The subjects were paid either small (\$0.50) or large (\$2.50) amounts of money to write this essay (the study was performed in 1967). Dissonance theory predicts that when justification for engaging in counterattitudinal behavior is weak (as is the case when subjects are paid a small amount of money), then cognitive dissonance will be high, and subjects will, as a result, be inclined to engage in attitudinal change as a means of providing additional justification for their actions. Consistent with this prediction, Linder et al observed that subjects who had been given relatively small amounts of money for writing the essay exhibited an attitudinal shift in favor of banning communist speakers.

As part of their attempt to develop a computational model of cognitive dissonance, Schultz and Lepper used CSNs to model the conditions associated with a number of cognitive dissonance experiments, including those seen in the Linder et al study. They showed that CSNs could be used to model the profile of cognitive changes seen across a number of cognitive dissonance research paradigms, including prohibition (where a subject is prevented from engaging in a desired activity), initiation (where a cost is imposed on group membership) and forced compliance (where a subject is coerced or induced to engage in counterattitudinal behaviors).

Although Schultz and Lepper's model correctly predicts the profile of cognitive changes seen in cognitive dissonance experiments, its relevance to sensemaking processes may not be readily apparent. One reason why the model was adopted as a computational model of sensemaking in the current case is because the kind of cognitive changes seen in cognitive dissonance studies could be described as a response to situations in which a subject is trying to make sense of observed behavior in terms of an underlying set of cognitions. In this case, the subject is trying to account for their actions in terms of the cognitive factors that must have given rise to those behaviors - the subject is essentially trying to make sense of their own behavior with respect to a set of casually-significant beliefs, attitudes and values. This may be seen as a specific example of the kind of sensemaking that occurs when we try to make sense of the behavior of other people via (e.g.) causal attribution mechanisms. Thagard and Kunda²¹ thus suggest that in making sense of both ourselves and others we rely on cognitive mechanisms for maximizing coherence, and such mechanisms can be understood in terms of parallel constraint satisfaction processes. Inasmuch as the subjects in cognitive dissonance experiments can be seen as trying to make sense of their own behavior, then the success of Schultz and Lepper's model in accounting for the results obtained in human experimental studies is clearly something that motivates its use in other sensemaking contexts.

A second reason why Schultz and Lepper's model is attractive as a means of studying sensemaking phenomena is because it sees cognitive change as motivated by the need to establish consistency between a set of variables representing cognitions. When we focus on sensemaking, we can characterize cognitions in terms of the beliefs that subjects hold about the features of some external state-of-affairs, such as the identity of a disease (in the case of medical diagnosis) or the nature of a situation (in the case of intelligence analysis). When subjects have partial access to this state-of-affairs, either because the available information is incomplete or uncertain, they must make inferences that are consistent with observed information as well as background knowledge and experience. Beliefs that are inconsistent or incompatible with one another, based on a subjects background knowledge and experience, will serve as the impetus for cognitive change (i.e., belief revision). From the perspective of CSNs, sensemaking processes can be modeled as the attempt to generate a profile of belief states (represented as the activation of nodes across the network) that are maximally consistent with one another, subject to the profile of inhibitory and excitatory links that connect the nodes. In the context of the sensemaking model outlined below, nodes represent the beliefs that individual sensemakers can have within a domain of interest, links represent the psychological association between specific beliefs (and reflect the background knowledge and experience that a sensemaker has in a domain), the activation of nodes at the start of a simulation represents the initial information that sensemakers are provided with, and the profile of activation-related change in the nodes represents the attempt of sensemakers to come to some cognitively-consistent interpretation of the initial information based on their background knowledge. The profile of node activation at the conclusion of the simulation represents

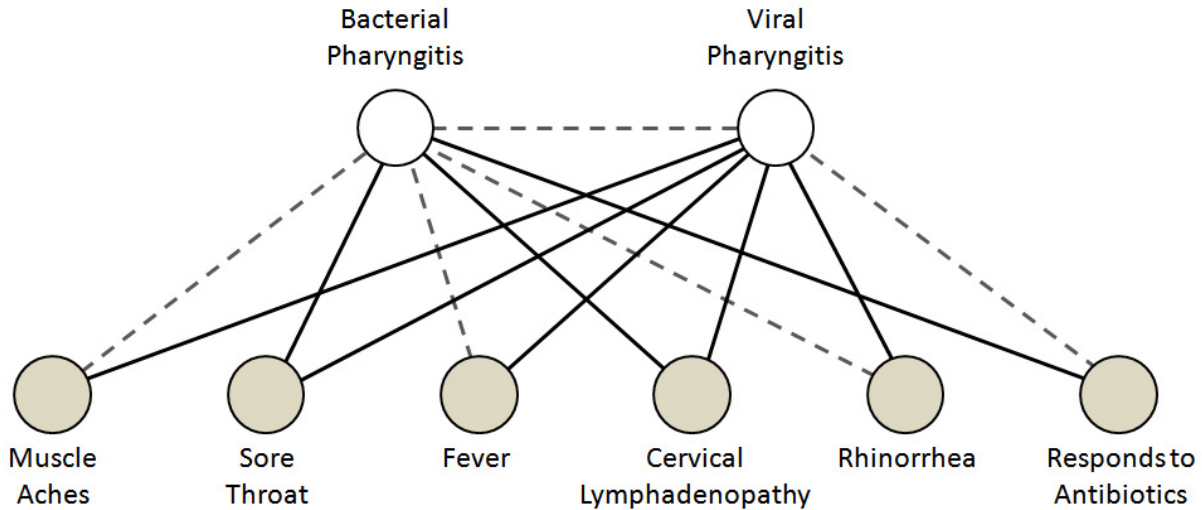


Figure 1. Organization of cognitive units in a single agent. Nodes in this diagram represent cognitive units, each of which represents a belief held by the agent. Solid lines represent excitatory connections between the units, while broken lines represent inhibitory links (all links are bidirectional). Colored circles, in this case, represent beliefs about the features of a disease, while white circles represent beliefs about the disease type.

the cognitive outcome of the sensemaking process. The activation of some nodes will, in this case, be positive (indicating that the subject holds the belief represented by the node), while the activation of other nodes will be negative (indicating that the subject does not hold the belief represented by the node). In general, the set of positive activations will represent the most cognitively-consistent or compatible set of cognitions the subject could have generated relative to the initial body of information they were provided with.

3. COMPUTATIONAL MODEL

As stated in Section 2.3, the computational model developed to explore collective sensemaking in the current paper is based on the model used by Schultz and Lepper.⁷ The current section provides an overview of the computational model. Further details can be found in Schultz and Lepper⁷ and Smart and Shadbolt.²²

Each agent within the model is implemented as a CSN based on the design specification outlined by Schultz and Lepper. The nodes which make up each CSN at the agent level are organized into a number of cognitive units, each of which represents a particular belief held by the agent. Fig. 1 illustrates the cognitive units that might be created for an agent in the domain of medical sensemaking. In this case, the agent is furnished with cognitive units representing the agent's beliefs about the presence of different symptoms, as well as the agent's beliefs about the diseases responsible for those symptoms.

Within an individual agent, cognitive units can be connected to other cognitive units via inhibitory or excitatory links. Whether the connection between two cognitive units is excitatory or inhibitory in nature depends on the compatibility or consistency of the beliefs represented by the cognitive units. In the case of Fig. 1, we see a profile of connectivity reflecting an agent's knowledge about the relationship between particular diseases and their associated symptoms. Positively weighted links thus exist between the 'Bacterial Pharyngitis' cognitive unit and the units representing the presence of a sore throat and the presence of cervical lymphadenopathy (swollen lymph nodes), while negatively-weighted links exist between the 'Bacterial Pharyngitis' unit and units representing the presence of a fever and rhinorrhea. Also note that a negative link connects the 'Bacterial Pharyngitis' and 'Viral Pharyngitis' units. This reflects the incompatibility between these beliefs in the context of single diagnostic session. As a result of the negative association between these units, an agent presented with a set of symptoms will be inclined to adopt the belief that the underlying disease is one of either bacterial or viral pharyngitis.

In addition to a sign, indicating whether a connection exerts an excitatory or inhibitory influence on its target, each connection has a weighting that determines the amount of influence it exerts. The weighting assigned to




Figure 2. The connectivity pattern for both positive (A) and negative (B) linkages between cognitive units. Excitatory connections are symbolized by solid lines, inhibitory connections by dashed lines.

a link will determine the level of influence a source cognitive unit has on a target unit, and this can be used to indicate the strength of the association between different cognitions. In the case, of the CSN shown in Fig. 1, for example, a stronger weighting could be assigned to the link connecting the ‘Cervical Lymphadenopathy’ and ‘Bacterial Pharyngitis’ units as compared to the link connecting the ‘Cervical Lymphadenopathy’ and ‘Viral Pharyngitis’ units. This would reflect the fact that this particular symptom is more commonly associated with bacterial forms of pharyngitis.

Cognitive units in Schultz and Lepper’s model have a degree of internal complexity. Thus, rather than each cognitive unit being represented by a single node, each unit in fact consists of two nodes that are connected together in a mutual inhibitory fashion (see Fig. 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 extent to which an agent holds the belief represented by the cognitive unit. Thus, if the activation of the positive pole is high relative to the negative pole, then the net activation for the cognitive unit will be positive and the agent can be said to possess the belief represented by the cognitive unit. Conversely, if the activation of the negative pole is high relative to the positive pole, then the net activation for the cognitive unit will be negative and the agent cannot be said to possess the belief represented by the cognitive unit.

The links between cognitive units are also architecturally complex. Each inter-cognition linkage is represented by connections between the constituent nodes of the cognitive unit. Thus, an inter-cognition linkage is not a single connection; instead, it consists of a total of eight connections (assuming bidirectional connectivity between the units), with two connections emanating from each of the nodes in the two cognitive units. A positive connection between cognitive units is represented by the wiring diagram shown in Fig. 2A, and a negative connection is represented by the wiring diagram shown in Fig. 2B.

In addition to the connections between cognitive units within a single agent, connections can also exist between different agents. These linkages are used to represent channels of communication between the agents in a given model; they allow agents to communicate their current beliefs about a situation to other agents. The nature of these inter-agent linkages resembles that those of inter-cognition linkages (see Fig. 2). For example, a positively-weighted inter-agent connection will be implemented according to the wiring diagram seen in Fig. 2A. In most cases, inter-agent connections connect cognitive units that represent the same kind of cognition. For example, if a channel of communication is established between two agents with the cognitive architecture seen

[†]Each node also has an ‘auto-regulatory connection’ that connects each node to itself (not shown in Fig. 2). The function of this auto-regulatory link is to dampen the node’s activity at each processing cycle (see Schultz and Lepper⁷ for details).

in Fig. 1, then the ‘Fever’ unit in the source agent will be connected to the ‘Fever’ unit in the target agent, the ‘Sore Throat’ unit in the source agent will be connected to the ‘Sore Throat’ agent in the target unit, and so on.

Computational processing in each CSN starts with the activation of particular cognitive units at the beginning of a simulation. This initial activation is deemed to represent an agent’s beliefs at the outset of the simulation (we assume that these beliefs are derived from some initial body of information that is provided to the agent). Computational processing then occurs via the spreading of activation between the cognitive units of the CSN according to the pattern of excitatory and inhibitory linkages. At each processing cycle in the simulation, the activation of each node in the CSN is updated according to the following rules:

$$a_i(t + 1) = a_i(t) + net_i(ceiling - a_i(t)) \quad (1)$$

when $net_i \geq 0$, and

$$a_i(t + 1) = a_i(t) + net_i(a_i(t) - floor) \quad (2)$$

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 maximum level of activation of the node (usually 1.0 for positive pole nodes and 0.5 for negative pole nodes), $floor$ is the minimum level of activation of the node (usually zero for all nodes), and net_i is the net input to node i , which is defined as:

$$net_i = resist_i \sum_j w_{ij} a_j \quad (3)$$

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 types of beliefs more or less resistant to change than others (see Section 5).

In each processing cycle of the simulation, n nodes are (typically) selected at random within each agent and updated according to equations 1 and 2, where n corresponds to the number of nodes in each agent (the total of number of nodes selected during a specific processing cycle is thus equal to the number of nodes in the model). Agents are then allowed to communicate information to their connected peers (i.e., their immediate neighbors in the communication network). Communication involves each agent sending a message to connected agents, which contains information about the activation of nodes in the source cognitive unit. When all messages have been sent, each agent processes the list of received messages and calculates the value of a parameter, called $comminput_i$, which is associated with every node. A variety of message processing functions are available to support the calculation of $comminput_i$; for example, the ‘MostTrustedAgent’ function selects messages from the most trusted communicating agent (as determined by the weight assigned to the connection between the agents) and sets the value of $comminput_i$ to the activation value of nodes in this agent[‡].

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

$$net_i = resist_i \left(\sum_j w_{ij} a_j + comminput_i \right) \quad (4)$$

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.

[‡]Other message processing functions can be selected by the user when setting up a simulation experiment (see Section 4.4).

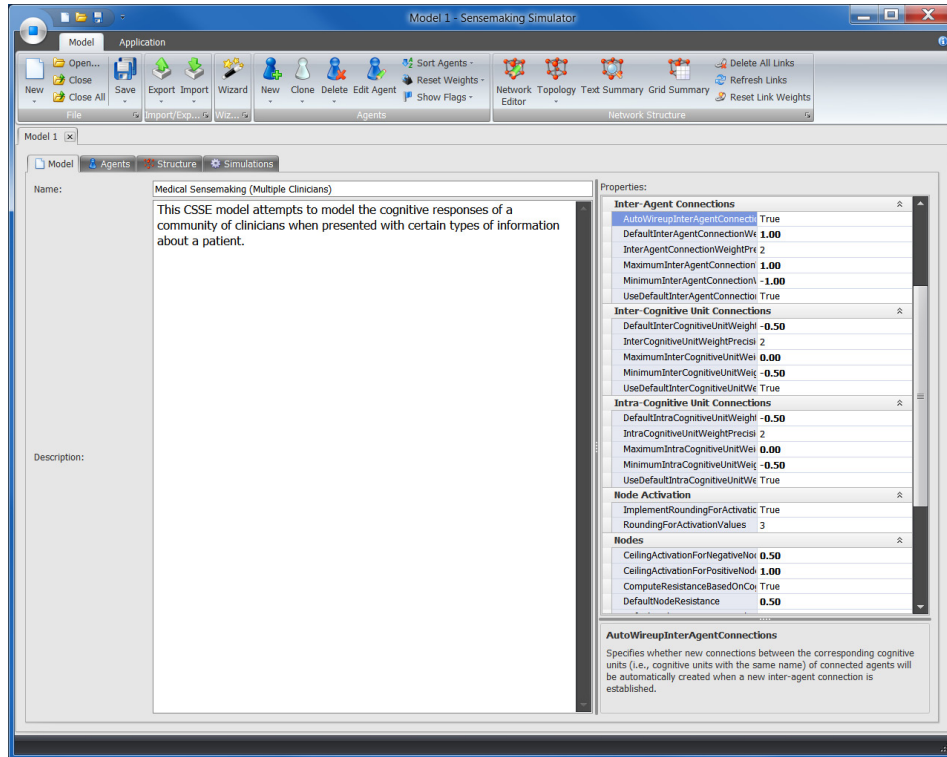


Figure 3. The ‘Model’ component of the Model Form as it appears within the CSSE. This part of the Model Form is used to configure model parameters, a subset of which are shown in the property grid control on the right hand side of the window.

4. SIMULATION CAPABILITY

In order to support experimental simulations into collective sensemaking using the aforementioned computational model (see Section 3), we have developed a Windows-based desktop application called the Collective Sensemaking Simulation Environment (CSSE). The CSSE is implemented using Microsoft’s .NET framework, and it features a graphical user interface that supports end-users with all aspects of the model configuration process. It also enables end-users to setup and run simulation experiments, as well as visualize and export simulation results. Subsequent sections describe the capabilities of the CSSE in more detail.

4.1 Model Creation

Fig. 3 shows the CSSE when a model is first loaded into the application. As can be seen from Fig. 3, the CSSE features a ribbon bar control along the top of the main application window. This provides access to common model manipulation functions; for example, the controls in the ‘File’ ribbon group on the left of the ribbon bar enable the user to create new models, open existing models and save opened models to disk. Once a model is loaded into the CSSE, it appears as a tabbed MDI child form within the main application window. This is called the Model Form, and it is organized into four tabs, named ‘Model’, ‘Agents’, ‘Structure’, and ‘Simulations’. These tabs enable end-users to configure the model and setup simulation experiments according to their research interests.

The first tab, named ‘Model’, is shown in Fig. 3. It allows users to provide a name and description for the model, and it also provides access to model properties via a property grid control. These properties are common to all models created using the CSSE. They enable users to, among other things, specify default weights for inter-agent and inter-cognitive unit connections, set values for node resistance and control how cognitive units are wired up to other cognitive units. A description of each of the properties appears in a text field below the property grid control whenever a model property is selected by the user.

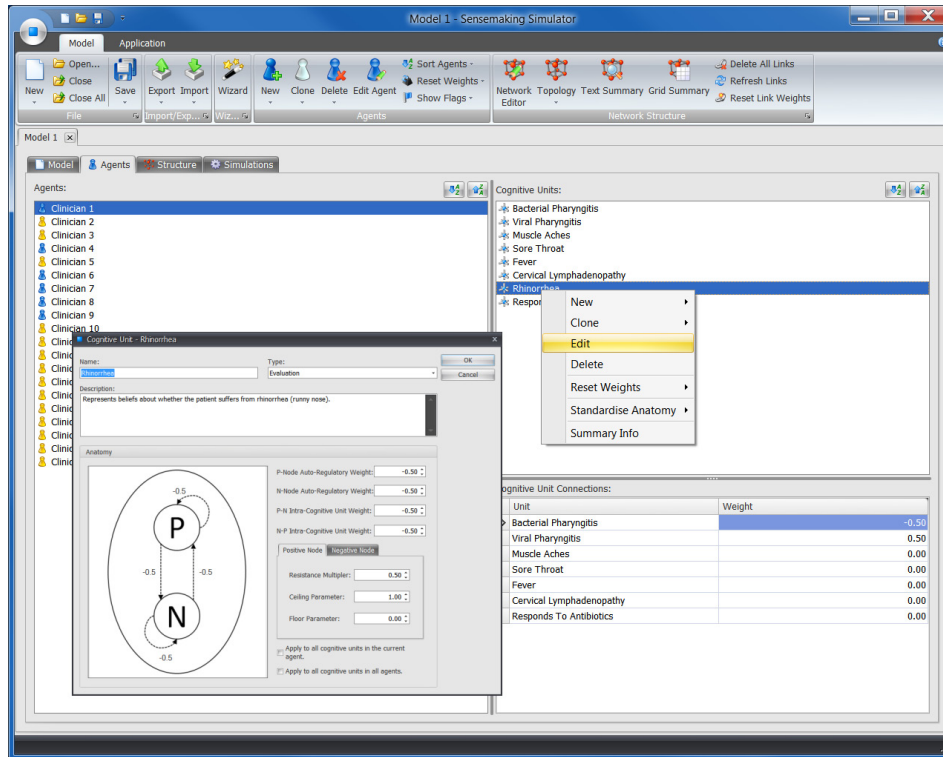


Figure 4. The ‘Agent’ component of the Model Form as it appears within the CSSE. This part of the Model Form is used to create the agents that are included in a computational model specification. The ‘Agent’ tab also allows the user to create, edit and configure each agent’s cognitive units.

4.2 Agent Creation and Cognitive Unit Configuration

The second tab of Model Form, named ‘Agents’, is shown in Fig. 4. This tab enables users to create, edit and delete the individual agents that appear within the model. Agents that are currently defined in the model are listed in a listbox control - the Agent Listbox - on the left hand side of the window. Whenever a specific agent is selected from the Agent Listbox, the cognitive units for that agent are listed in another listbox - the Cognitive Unit Listbox - on the right hand side of the window. The user can create, edit and delete the cognitive units for any given agent by selecting items from a popup menu that appears when the user right clicks within the bounds of the listbox (see Fig. 4). When the user selects the ‘Edit’ menu item from the popup menu, they are presented with the dialog box shown in Fig. 4. This enables the user to configure all the properties associated with an individual cognitive unit; for example, the weight assigned to intra-cognitive unit connections. The properties established for a single cognitive unit can be copied to all cognitive units within the same agent (or all agents within the same model) by selecting options from within the dialog box (the same functionality is also provided by the ‘Standardise Anatomy’ menu option, which is accessible from the aforementioned popup menu).

The bottom right hand side of Fig. 4 shows a datagrid control, which is used to configure the connectivity of cognitive units within a single agent. Whenever a user selects a cognitive unit from the Cognitive Unit Listbox, the datagrid shows all other cognitive units within the same agent that are connected to the selected cognitive unit. The user can then edit the weights assigned to these connections using the datagrid control. As mentioned before, the weights assigned to the inter-cognitive unit connections indicate the strength of a psychological implication or association between cognitions (e.g., beliefs). If the user wants to specify that no psychological implication or association exists between two cognitive units, as might be the case in domains where some beliefs are independent of one another, then the weight assigned to the inter-cognitive unit connection can be set to zero.

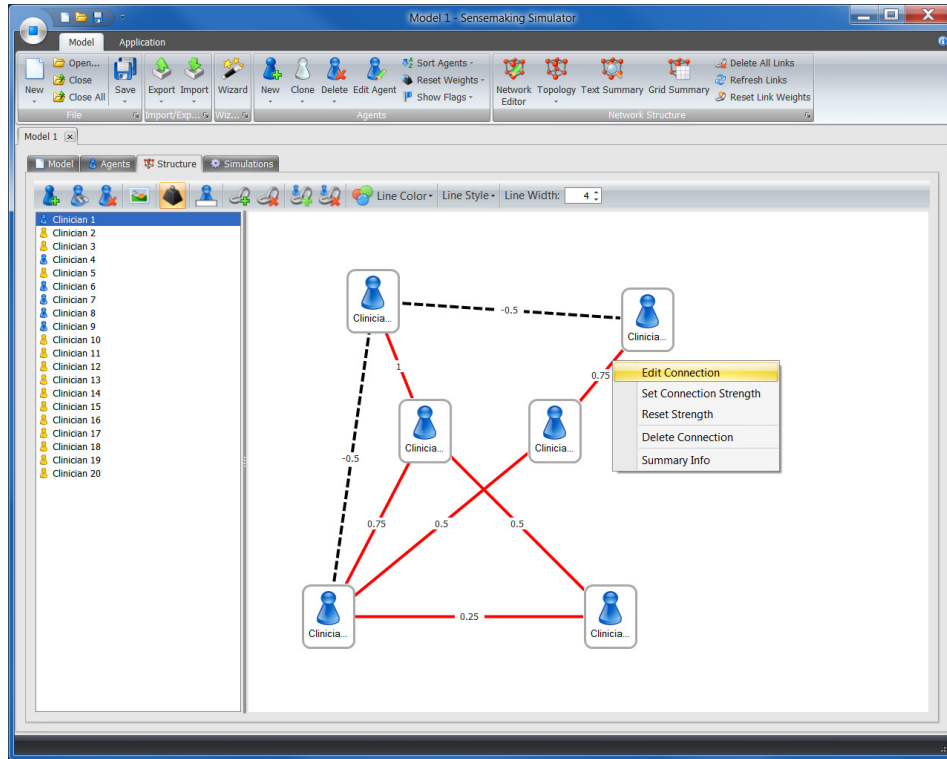


Figure 5. The ‘Network’ component of the Model Form as it appears within the CSSE. This part of the Model Form is used to configure the communication network that allows agents to communicate with one another during the course of a simulation. The ‘Network’ tab features the Network Structure Control, which provides a drawing canvas, enabling users to manually configure the communication links that exist between agents.

4.3 Communication Network Configuration

Fig. 5 shows the third tab of the Model Form, which is named ‘Network’. This tab enables the user to configure the structure of the communication network that exists between agents using what is called a Network Structure Control. The agents currently defined within the model are listed in a listbox on the left hand side of the window, and users can select agents from this list and add them to the drawing canvas on the right hand side of the window. Once visible on the drawing canvas, the user can establish connections between agents by selecting the source and target agents using the mouse[§]. Existing connections can be deleted by right-clicking on the connection and selecting the ‘Delete Connection’ option from a popup menu (see Fig. 5). The CSSE provides a number of shortcuts to support the configuration of communication networks. Firstly, users have the option of connecting all agents that are visible on the drawing canvas. Secondly, users can use a drop-down menu option available from the main ribbon bar control, which allows them to create four types of network structure topology, namely fully-connected, small-world, random and ring topologies. In this case, the network structure is established over all agents defined within the model, rather than just those agents shown on the drawing canvas.

As mentioned in Section 3, communication links between agents assume the form of connections between the corresponding cognitive units in each agent (e.g., a ‘Fever’ unit in the source agent will establish a connection with the ‘Fever’ unit in the target agent). The user can edit the structure of these inter-cognitive unit linkages, as well as the weights assigned to the linkages, by selecting the ‘Edit Connection’ option that appears in the popup menu associated with a link on the drawing canvas (see Fig. 5). This will display a dialog box to the

[§]Whenever agents are connected to at least one other agent, the color of the icon representing the agent in the CSSE changes from yellow to blue. This provides the user with a visual cue as to what agents are incorporated into the communication network.

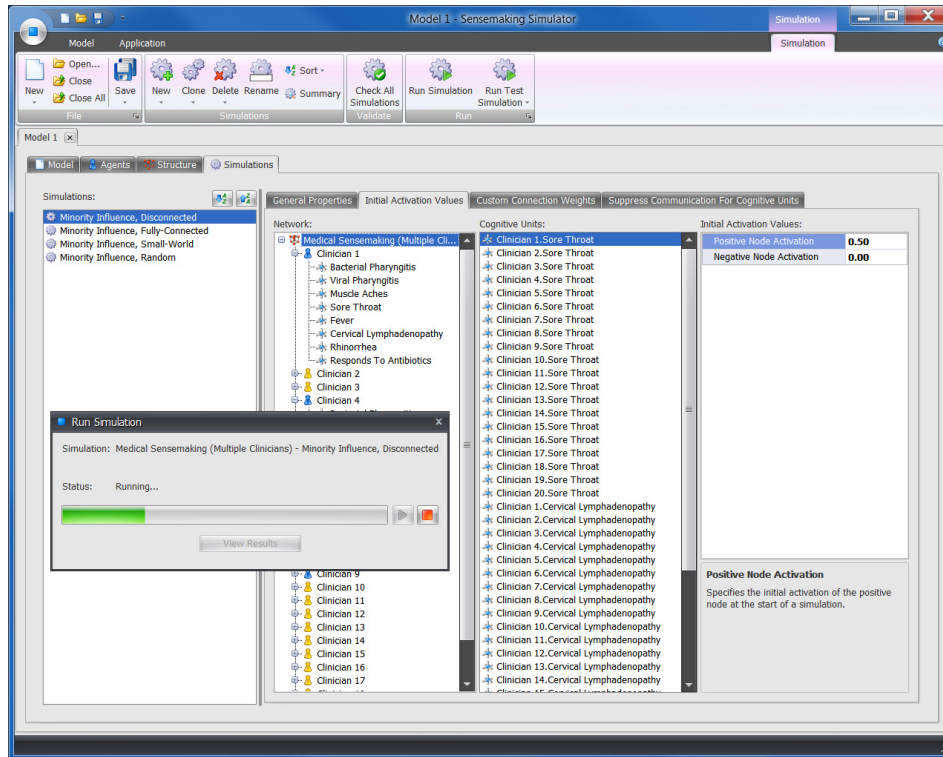


Figure 6. The ‘Simulation’ component of the Model Form as it appears within the CSSE. This part of the Model Form is used to create, configure and execute simulation experiments.

end user, which provides complete control over how cognitive units are connected in the context of an inter-agent communication link. Specific connections could be deleted, for example, in order to prevent inter-agent communication about particular cognitions (this could be used to represent situations in which an agent wishes to keep some beliefs private - see Section 5). Alternatively, cognitive units could be wired up in a such a way that different kinds of cognition are connected (e.g., a ‘Fever’ cognitive unit in the source agent could be connected to the ‘Rhinorrhea’ unit in the target agent). This state-of-affairs could be used to implement a primitive form of miscommunication between the agents where one agent intends to communicate information about a particular belief but this is misinterpreted by another agent as providing information about a different belief (see Section 5).

In addition to supporting the configuration of the communication network, the Network Structure Control also provides a number of visual formatting options. The width, color and style of the connections between agents can be edited using options available from a toolbar situated above the drawing canvas, and the value of these properties can be specified for particular types of connection (e.g., the user could render excitatory and inhibitory connections between agents using different colors). Snapshots of the network structure diagram can be exported as images for use in other applications.

4.4 Simulation Setup and Execution

The fourth tab of the Model Form, named ‘Simulation’, is shown in Fig. 6. This tab enables the user to create, configure and execute simulation experiments. Individual simulation experiments are listed on the left hand side of the window in a listbox control - the Simulations Listbox. When a particular simulation experiment is selected, the user can edit its properties via the tabbed control on the right hand side of the window. These properties are organized into a number of categories, which are indicated by the labels assigned to the tabs within the tabbed control. These categories are as follows:

- **General Properties.** This set of properties includes the number of trials to perform as part of the simulation and the number of processing cycles to implement as part of each trial. Other properties in

this category determine whether agents are allowed to communicate on every cycle of the simulation and whether specific kinds of communication network topology (e.g., a small-world topology) is created anew at the beginning of each trial. The user can also select from a list of ‘message processing functions’ in order to specify how the messages sent from communicating agents are processed by the receiving agent.

- **Initial Activation Values.** These properties determine the initial activation values assigned to specific cognitive units at the beginning of each trial within a simulation. As indicated in Fig. 6, the user can select cognitive units from particular agents and specify an initial activation value for both the positive and negative poles of the cognitive unit.
- **Custom Connection Weights.** These properties specify connection weights between the cognitive units of agents that are connected within a communication network. These weights override those that are established using the Network Structure Control (see Section 4.3). The reason for allowing users to override connection settings in the context of a particular simulation is due to the fact that a user may want to adjust the communication links across multiple simulations that are executed against the same model. If this feature was not provided, the user would have to create multiple copies of the same model and distribute simulations across multiple model files.
- **Suppress Communication for Cognitive Units.** These properties allow the user to suppress communication for particular cognitive units across all cycles of a simulation. This may be important in cases where the user wants to explore the effect of restricted communication profiles, in which some subset of beliefs is kept private.

Once a simulation has been created and configured, it can be executed by clicking the ‘Run Simulation’ button that appears on the main ribbon bar control. This will bring up a dialog box that enables the user to control the execution of ongoing simulations. Using this dialog box, users can suspend, resume and cancel the execution of ongoing simulations.

Simulations are run in a separate process from the CSSE in order to allow the user to continue working on simulation models. When a simulation has completed, the results are stored in a user-specified directory on the host machine.

4.5 Results Visualization and Data Export

The results of simulation experiments can be viewed using a secondary application that ships with the CSSE. The main purpose of this application is to allow the user to view the results of simulations, and to this end, it provides access to a datagrid control that supports a rich range of filtering, grouping, sorting, searching and data manipulation capabilities. It also enables users to view the activation profile of individual cognitive units or groups of cognitive units across the course of a simulation. Whenever a group of cognitive units is selected, the grid control shows the mean, standard deviation, and standard error of the net activation values for the cognitive units within the group.

In order to support more advanced forms of data visualization and analysis, the user can export the simulation results to another application. At the present time, three kinds of export format are supported: Excel workbooks, comma-separated value (CSV) files and HTML. In the case of the Excel export option, the data from the simulation is stored in a series of spreadsheets that form part of the workbook. The user can either use the built-in analysis and visualization options available within the Excel environment or they can export the data to another compatible application for further analysis, e.g., SPSS.

5. EXPLORING THE EFFECT OF SOCIO-TECHNICAL FACTORS ON COLLECTIVE SENSEMAKING: FOCUS AREAS FOR EXPERIMENTAL SIMULATION

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.^{22,23} In one study, for example, Smart and Shadbolt²² 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 pre-existing or prevailing cognitions.

In another study, Smart²³ explored the effect of different communication network structures on minority influence in collective sensemaking situations. This study used four types of communication network topology, namely, disconnected, random, small-world and fully-connected topologies. 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 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.

Aside from studies investigating factors such as communication network structure, communication frequency and minority influence, a number of other factors are important foci of research attention for those interested in collective sensemaking. Many of these factors relate to features of the sociotechnical environment in which sensemaking occurs. For example, in military coalition contexts we are often interested in how factors such as cultural differences, inter-agent trust and miscommunication affect collective sensemaking processes. The following list provides an overview of how the computational model presented here could be applied to support research into the effect of these (and other) factors.

- **Cultural Differences.** One way of viewing cultural differences is to see them as reflecting statistically-significant differences in the cognitive structures associated with the members of different cultural groups.²⁴ Given that inter-cognition linkages within a single agent reflect that agent's knowledge and experience of a domain, the introduction of inter-agent differences in the way (intra-agent) cognitive units are connected may provide one way to investigate the impact of cultural differences on collective sensemaking processes. Cultural differences could thus 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 (perhaps) come to different conclusions given the same body of external information.
- **Inter-Agent Trust.** One interpretation of the weight value assigned to inter-agent linkages is that it represents the degree of trust between agents. High trust between agents is thus reflected in links that have high weight values, and low trust is reflected in links that have low weight values. By introducing variable weights for the linkages that connect agents, we can vary the amount of influence they exert over one another during the course of a simulation. Different values could also be specified for each channel of communication between two agents. This would correspond to a situation in which agents had different levels of trust in one another with respect to particular kinds of beliefs.
- **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.
- **Selective Communication.** Agents do not have to communicate the state of all their beliefs during the course of inter-agent communication. They could, for example, restrict their communication to particular kinds of beliefs. This would enable some agents to maintain a degree of privacy with regard to some beliefs. One case where this may be important is in situations where (e.g.) security constraints prohibit the communication of information to other agents. The notion of privately held beliefs and opinions is also something that is important in computer simulation studies of social influence.¹⁶ In terms of the current model of collective sensemaking, this mode of selective communication can be established by removing the

communication links associated with particular cognitive units (or by setting the weight assigned to the link to zero).

- **Resistance to Social Influence.** In Section 3, we encountered a parameter, called $resist_i$, which is a measure of the resistance of node i to having its activation changed. The value of this parameter can be changed in order to control the susceptibility of an agent to socially-mediated cognitive change. Studies manipulating this parameter could be useful for investigating issues of social influence in situations where beliefs are differentially resistant to change. Previous research in this area has investigated the effect of cognitive resistance in minority influence. Xie et al,²⁵ for example, show how the prevailing majority opinion in a population can be rapidly reversed by a minority of agents who consistently proselytize the opposing opinion and are immune to social influence.

6. CONCLUSION

Cognitive processing in organizational settings is often a collective endeavor. More often than not, activities such as planning and decision-making require the coordinated effort of multiple individuals working together. The socially-distributed nature of cognitive processing in these task contexts means that a variety of factors have the potential to influence cognitive outcomes. 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 social, technological and organizational factors. When combined with empirical studies using human subjects, the output from this work can be used to support the effort to develop techniques and technologies that maximize the collective cognitive processing potential of individuals engaged in sensemaking activities.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Sun, R., “Cognitive social simulation incorporating cognitive architectures,” *Intelligent Systems* **22**(5), 33–39 (2007).
- [2] Pirolli, P. and Card, S., “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).
- [3] Jensen, E., “Sensemaking in military planning: A methodological study of command teams,” *Cognition, Technology & Work* **11**(2), 103–118 (2009).
- [4] Paul, S. A. and Reddy, M. C., “Understanding together: Sensemaking in collaborative information seeking,” in [*ACM Conference on Computer-Supported Cooperative Work*], Savannah, Georgia, USA (2010).
- [5] Alberts, D. and Garstka, J., “Network centric operations conceptual framework version 2.0,” Technical Report, U.S. Office of Force Transformation and Office of the Assistant Secretary of Defense for Networks and Information Integration, U.S. Department of Defense, USA (2004).
- [6] Smart, P. R. and Sycara, K., “Collective sensemaking and military coalitions,” *Intelligent Systems* **28**(1), 50–56 (2013).
- [7] Schultz, T. R. and Lepper, M. R., “Cognitive dissonance reduction as constraint satisfaction,” *Psychological Review* **103**(2), 219–240 (1996).
- [8] Klein, G., Wiggins, S., and Dominguez, C. O., “Team sensemaking,” *Theoretical Issues in Ergonomics Science* **11**(4), 304–320 (2010).

- [9] Thagard, P. and Verbeurgt, K., “Coherence as constraint satisfaction,” *Cognitive Science* **22**(1), 1–24 (1998).
- [10] Holyoak, K. and Thagard, P., “Analogical mapping by constraint satisfaction,” *Cognitive Science* **13**(3), 295–355 (1989).
- [11] Kintsch, W., “The role of knowledge in discourse comprehension: A construction-integration model,” *Psychological Review* **95**, 163–182 (1988).
- [12] Spellman, B., Ullman, J., and Holyoak, K., “A coherence model of cognitive consistency: Dynamics of attitude change during the Persian Gulf War,” *Journal of Social Issues* **49**(4), 147–165 (1993).
- [13] Thagard, P., “Explanatory coherence,” *Behavioral and Brain Sciences* **12**(3), 435–502 (1989).
- [14] Kunda, Z. and Thagard, P., “Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory,” *Psychological Review* **103**(2), 284–308 (1996).
- [15] Rumelhart, D., Smolensky, P., McClelland, J., and Hinton, G., “Schemata and sequential thought processes in PDP models,” in [*Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2*], Rumelhart, D. and McClelland, J., eds., 7–58, MIT Press, Cambridge, Massachusetts (1986).
- [16] Mason, W., Conrey, F., and Smith, E., “Situating social influence processes: Dynamic, multidirectional flows of influence within social networks,” *Personality and Social Psychology Review* **11**(3), 279–300 (2007).
- [17] Hutchins, E., “The social organization of distributed cognition,” in [*Perspectives on Socially Shared Cognition*], Resnick, L., Levine, J., and Teasley, S., eds., The American Psychological Association, Washington DC, USA (1991).
- [18] Festinger, L., [*A Theory of Cognitive Dissonance*], Stanford University Press, Stanford, USA (1957).
- [19] Cooper, J., [*Cognitive Dissonance: Fifty Years of a Classic Theory*], Sage Publications Ltd, London, UK (2007).
- [20] Linder, D. E., Cooper, J., and Jones, E. E., “Decision freedom as a determinant of the role of incentive magnitude in attitude change,” *Journal of Personality and Social Psychology* **6**, 245–254 (1967).
- [21] Thagard, P. and Kunda, Z., “Making sense of people: Coherence mechanisms,” in [*Connectionist Models of Social Reasoning and Social Behavior*], Read, S. J. and Miller, L. C., eds., Lawrence Erlbaum Associates, Mahwah, New Jersey, USA (1998).
- [22] Smart, P. R. and Shadbolt, N. R., “Modelling the dynamics of team sensemaking: A constraint satisfaction approach,” in [*Knowledge Systems for Coalition Operations (KSCO’12)*], Pensacola, Florida, USA (2012).
- [23] Smart, P. R., “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).
- [24] Sieck, W. R., Rasmussen, L., and Smart, P. R., “Cultural network analysis: A cognitive approach to cultural modeling,” in [*Network Science for Military Coalition Operations: Information Extraction and Interaction*], Verma, D., ed., IGI Global, Hershey, Pennsylvania, USA. (2010).
- [25] Xie, J., Sreenivasan, S., Korniss, G., Zhang, W., Lim, C., and Szymanski, B. K., “Social consensus through the influence of committed minorities,” *Physical Review E* **84**(1) (2011).