

Collective Sensemaking and Military Coalitions

Paul R. Smart, *University of Southampton*

Katia P. Sycara, *Carnegie Mellon University*

Military operations must often deal with incomplete, uncertain, ambiguous, conflicting, and inaccurate information—the “fog of war” always confounds attempts at a complete understanding of the battlefield.

Sensemaking is the process that military organizations undertake to deal with

the fog of war. Individuals and teams attempt to develop an understanding against which they can make decisions and formulate plans for action.

At heart, sensemaking is a cognitive activity: it involves the processing of information to yield an outcome that’s recognizably cognitive in nature. This doesn’t mean, however, that only individuals can engage in sensemaking. Appreciation is growing for the prevalence and importance of what we might call *collective sensemaking*—that is, the activities that groups of individuals perform to develop understanding at both the individual and collective levels.¹ Work on collective sensemaking is the focus of an increasing body of empirical and theoretical work in several research communities, and these efforts are paralleled by extensive research into related notions, such as shared/team situation awareness, shared understanding, and shared mental models.²

Sensemaking’s importance to contemporary military organizations is reflected in the fact that sensemaking processes sit at the heart of *network-centric operations* (NCO).

According to the NCO Conceptual Framework (NCO-CF), for example, sensemaking at both the individual and collective levels directly affects decision synchronization, force agility, and mission effectiveness.³ In particular, sensemaking processes are an intervening variable in the NCO value chain: they let military organizations capitalize on the progress made with respect to networking technology and improved information-sharing capabilities.³

Given collective sensemaking’s central role in coalition operations, we must develop a better understanding of the relationship between specific features of the coalition communication environment and aspects of collective sensemaking performance. At present, however, we have little knowledge of these relationships’ precise nature. One approach to improving our understanding is to develop computational models that simulate aspects of the sensemaking process.

Here, we present a computational model of collective sensemaking that we’re developing within the International Technology Alliance (ITA) research program. In addition

Sensemaking is a key capability for military coalitions. It lets both individuals and teams make sense of conflicting, ambiguous, and uncertain information.

Computational modeling provides one means of improving our understanding in this area.

to describing the model, we highlight some factors that might affect sense-making performance in coalition contexts and describe how the model can represent some of them.

Coalition Environments and Collective Sensemaking

Military coalitions exist as complex, socio-technical organizations in which a variety of informational, cognitive, social, and technological factors influence collective cognitive outcomes. At present, our understanding of how these factors affect collective cognitive processes is limited, which makes testing the hypotheses and assumptions associated with network-centric warfare (or network-enabled capability) difficult. The NCO-CF, for example, suggests that better networking, interoperability, and information-sharing capabilities will likely improve collective sense-making abilities.³ This claim certainly has an intuitive appeal, but greater levels of networking and information sharing don't necessarily yield better cognitive outcomes. In the social psychology literature, for example, we encounter the phenomenon of *production blocking*, or the tendency for one individual's contributions to block or inhibit contributions from other group members. In some situations, therefore, the tendency to share information can undermine a group's collective creative potential; instead of stimulating a greater number and diversity of ideas, precipitant forms of information sharing can sometimes impede the creative process.

A poor understanding of how coalition communication environments can affect collective sensemaking also makes it hard to engineer coalition environments in ways that benefit sensemaking abilities. This doesn't just apply to a military coalition's

technological aspects but also its social aspects. For example, military coalitions are composed of individuals from various nation states and military services, so they're in a position to benefit from the diverse knowledge, training, and expertise that individuals bring to shared tasks. However, when it comes to collective cognitive processes, cognitive diversity isn't necessarily a virtue. On one hand, evidence suggests that cognitive heterogeneity is useful in mitigating against the cognitive biases (such as confirmation bias) that are sometimes associated with collective sensemaking.⁴ On the other hand, a team of individuals with different background knowledge and beliefs can present challenges in terms of miscommunication due to linguistic and cultural differences.⁵ The question thus arises as to how we can organize collective sensemaking activities at the social level. Should we form sensemaking teams based on a principle of maximizing or minimizing team members' cognitive diversity?

Homogeneity and diversity issues are also important when considering technology development and use within military coalitions. Different nation states or military services tend to adopt different approaches to representing and storing information, which we often view as a barrier to interoperability and collaboration. Using standardized representational formats as well as common search, retrieval, and storage solutions thus seems like an ideal way to establish cognitively empowering forms of collaboration and engagement. The downside is that we potentially lose any diversity in how information is indexed, retrieved, and presented. This risks exposing all individuals to the same information in the same way, which might affect a team's ability to generate novel ideas and interpretations.

In general, efforts to press maximal cognitive benefit from a coalition's technological and informational assets must be grounded in an understanding of how the coalition communication environment's various features affect collective cognitive processes. Such features can include trust relationships, communication network topologies, the timing and frequency of interagent communication, the extent of information sharing, differential access to specific bodies of information, cognitive diversity, and the potential for miscommunication. Research must assess how these factors affect the dynamics of sensemaking processes and the quality of sensemaking outcomes.

Constraint Satisfaction Model

In general, researchers have used two different approaches to examine collective cognition. Social psychological research tends to observe small groups of individuals in a particular task context. This approach uses real human subjects to test research hypotheses, but is often limited to small groups with minimally complex communication structures, and it focuses on specific tasks that don't necessarily generalize to real-world situations.

The other approach uses computer simulation techniques and is potentially well-suited to studying collective cognitive processes in military coalitions: it avoids the cost (or risk) of running large-scale coalition-based experiments and lets the research community explore "what-if" scenarios that systematically manipulate variables that would be too impractical (or dangerous) to manipulate in real-world situations. Despite these advantages, several problems are associated with using computer simulations to explore collective cognitive processes. 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. Examining the effects of, for example, differences in knowledge and experience on sensemaking processes can thus be difficult.

Recent work in our laboratories has sought to develop a psychologically interesting computational model of collective sensemaking.⁶ The model adopts a network-of-networks approach to cognitive simulation in which we implement each agent as a constraint satisfaction network (CSN) and connect multiple instances of these networks to form an inter-agent communication network. We decided to use CSNs as the basis for individual sensemaking abilities for several reasons.

First, we can usefully cast sensemaking as a type of constraint satisfaction problem. In particular, we suggest that to make sense of information, agents must often use background domain knowledge to form beliefs that are highly consistent or compatible. This is important because CSNs have been used to model psychological processes in which coherence and consistency issues play a major role. For example, one study used CSNs to examine cognitive dissonance, in which the process of cognitive change (for instance, belief modification) is driven by a need for consistency or compatibility between cognitive states.⁷

Second, CSNs let us study knowledge's role in guiding sensemaking




Figure 1. Example constraint satisfaction network (CSN) representing beliefs about two kinds of objects: cats and birds. Solid lines symbolize excitatory connections between the units, whereas broken lines symbolize inhibitory links. Circles represent cognitive units. Shaded circles represent beliefs about the objects' features (feature beliefs), whereas plain circles represent beliefs about object type (object beliefs).

performance. In particular, each node within the CSNs in our model corresponds to a particular belief that an agent might hold. We refer to these nodes as *cognitive units*. They're connected using either excitatory or inhibitory connections, and each such "intercognition" link is associated with a weighting value. The connection pattern between the cognitive units represents an agent's background knowledge or experience in a particular domain. For instance, the CSN that Figure 1 shows has six cognitive units, each of which represents beliefs about two animal types: cats and birds. The cognitive units in this network are connected such that they 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 reflects an agent's (admittedly limited) knowledge about cat and bird objects. If one unit is artificially stimulated (a situation we consider analogous to the presentation of specific kinds of evidence), then other units with positive connections to that unit will have increased activation across successive processing cycles. The result is that agents settle on belief states that are most consistent

with the evidence available to them, as well as their background knowledge about 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, whereas 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, given the agent's background knowledge.

A third reason to use CSNs centers on previous uses of such networks to explore the dynamics of collective sensemaking. In particular, Edwin Hutchins has used CSNs to examine the psychological phenomenon of confirmation bias⁸—that is, people's tendency to ignore or discount evidence that contradicts some initial interpretation of a situation. Using CSNs, Hutchins demonstrated that the timing of interagent communication significantly influences the dynamics of collective sensemaking. In particular, if individual agents communicated with each other from a simulation's outset, then extreme confirmation bias arose. This occurred because each agent, influenced by information from other agents in the social network, was under pressure to discover a shared interpretation of the input data. In other words, the community strove to find a set of activation patterns that best satisfied the internal constraints that interagent communication had established. The result was that agents often failed to give due weight to the evidence from external input data; more often than not, the community

perfectly sensible interpretation of the available evidence, given the agent's background knowledge.




Figure 2. Experimental results. We studied how (a) the communication network structure affected collective sensemaking. Nodes represent agents, and lines indicate communication channels. We also determined (b) the mean activation levels of the “cat” and “bird” cognitive units in each network structure condition. Standard error of the mean (SEM) isn’t shown, but in all cases it was less than 0.03.

exhibited more extreme confirmation bias than did isolated individuals.

Communication Network Structure and Sensemaking Performance

Although CSNs have a long history of use within the computer and psychological sciences, their specific application to collective sensemaking processes has been somewhat limited. In addition, we don’t always know how to accommodate the range of psycho-social and technological factors potentially affecting sensemaking performance in conventional CSN-based simulations. The model of collective sensemaking that we’ve developed relies on a particular form of CSN that’s used to study cognitive dissonance.⁷ By relying on a CSN that was originally developed to model psychological phenomena (and validated against psychological data), we aim to deliver a simulation capability that yields a range of psychologically interesting and relevant results. Furthermore, by incorporating a range of configuration parameters, we hope to represent a

subset of the features that we’d typically find in coalition communication environments—for example, variable levels of interagent trust, differential access to specific bodies of situation-relevant information, and inter-individual differences in background knowledge and experience.

To exemplify our model’s application and use, we recently used it to examine the role different communication network structures played in mediating minority influence under various informational conditions. In one such condition, a minority of agents received strong evidence in favor of one interpretation, whereas a majority received weak evidence in favor of a conflicting interpretation.⁶ We configured the CSNs (each representing an individual sensemaking agent) as indicated in Figure 1, with six cognitive units representing beliefs about cats and birds (note that this configuration represents just one particular instantiation of the model; alternative configurations could have more cognitive units with more complex patterns of intercognitive unit connectivity).

For each simulation, we created 20 agents and connected them to form one of four communication network structure types (see Figure 2a). At the beginning of each simulation, we randomly selected five agents for assignment to the minority group, while assigning the remaining 15 agents to the majority group. We then initialized the agents in each group with activation vectors that established the initial activation levels of cognitive units within the agent. For minority group members, we set the activation of the “has-fur” unit at 0.5 and the activation of all other units at 0.0. For the majority group members, we set the activation of the “has-feathers” unit at 0.1 and all other units at 0.0.

Because the initial activation of cognitive units represents an agent’s beliefs at the simulation’s outset (reflecting, perhaps, exposure to different evidence), we can see that agents in the two groups had different beliefs about the object’s features they were presented with. Over the course of successive processing cycles, we would expect these two initial belief

states to generate different interpretations about the object type. Agents in the two groups also differed with respect to the activation levels associated with cognitive units. In terms of the activation levels' psychological significance, greater activation levels reflect an agent's confidence or certainty in a particular belief. Thus, in our simulations, agents in the minority group had greater certainty that a particular feature was present compared to agents in the majority group.

We studied how these initial informational conditions affected the emergence of belief states using four network structure conditions (see Figure 2a). Our study addressed the question of how the different network structures would affect agents' tendencies to settle on different interpretations of the available evidence. Would agents come to adopt the minority view that a cat object was present, or would they adopt the competing interpretation that a bird was present? At the individual level, each agent attempted to reconcile the information it received from other agents and the environment with its own background knowledge of the domain. Agents arrived at a particular interpretation of the information that manifested itself in the activation pattern across their cognitive units. The communication links thus changed how agents responded to environmental information: agents adjusted their interpretations in ways that considered the information their network neighbors provided.

Figure 2b shows our experimental results, which we obtained by recording the "cat" and "bird" units' activation levels after 20 processing cycles, and by running 50 simulations in each of the four network structure conditions. As the figure suggests, with the "cat" cognitive unit, activation was greatest in the fully

connected network and lowest in the disconnected network; activation in the random and small-world networks was at an intermediate level. With the "bird" cognitive unit, we obtained the reverse pattern: activation was lowest in the fully connected network, highest in the disconnected network, and at intermediate levels in the random and small-world networks. This pattern suggests that communication networks with different structural topologies differentially affect sensemaking performance. In situations in which a minority of agents receives strong evidence in favor of one interpretation and a majority receives evidence in favor of an alternative, competing interpretation, the communication network structure mediated the effect of the minority's influence on collective interpretative outcomes. In particular, fully connected networks enable strong but uncommon evidence to quickly influence all agents' beliefs before weaker, contradictory evidence has had time to contribute to opposing beliefs. With small-world and random networks, weaker evidence has more time to contribute to beliefs that are progressively more resistant to change across successive processing cycles.

Model Parameters and Coalition Environment Features

The aforementioned study focuses on one aspect of military coalition environments: the potential for communication to be influenced by a range of communication networks that differ in their topological structure. Several other factors could affect sensemaking processes in coalition settings, however, including the level of trust between agents, the extent of information sharing, and the potential for miscommunication based on cultural

or linguistic differences. Our model can support simulations that represent at least some of these factors. It can represent trust, for example, via the weighting associated with specific interagent connections, and we can specify different values for the links between the cognitive units located in each connected agent. This reflects the fact that agents might have different trust levels regarding particular beliefs.

Another factor the model can represent concerns individual differences in knowledge and experience. Each link between cognitive units within a single agent represents a psychological implication or association between belief states, with the link's weight reflecting this implication's or association's strength. The set of inter-cognition links for each agent thus represents the background knowledge (including assumptions, stereotypes, and prejudices) that the agent brings to bear in making sense of the information presented. It follows that individual variability in the CSN structure (in terms of the weights associated with inter-cognition links) will reflect differences in background knowledge. This gives us a way to investigate the relative benefits of cognitive homogeneity/heterogeneity in sensemaking teams under different informational conditions. It has been suggested that cultural differences within coalitions reflect statistically significant differences in the cognitive structures associated with the members of different cultural groups.⁹ Thus, varying inter-cognition links according to group membership criteria could let us adapt our model to explore cultural diversity's effect on sensemaking performance.

Other factors our model can represent include the frequency of inter-agent communication, resistance to social influence, and access to specific

bodies of information (for example, where different agents are exposed to different subsets of information). Obviously, these factors will interact with each other in complex ways during a simulation; thus, the results obtained with one particular configuration (such as the experiment exploring the communication network structure) aren't invariant across simulations in which another model parameter (for example, resistance to social influence) is systematically manipulated. We can't use the experimental results described in relation to one particular study to predict model behavior in a different experimental context.

Model Limitations

Although our model can accommodate some features of collective sensemaking in coalition environments, it doesn't capture all the richness and complexity of the sensemaking process as seen in real-world situations. Perhaps one of its most important shortcomings is that it doesn't consider the role knowledge-guided interactions with the real world play in letting agents make sense of conflicting, uncertain, or ambiguous information. Thus, in most cases of real-world sensemaking, human agents are likely to actively engage with the external environment to support their ongoing attempts at pattern completion and recognition. Human agents engaged in sensemaking don't just react passively to the information they receive; they also seek to manipulate their information environments to meliorate their access to hidden patterns, relationships, and contingencies. This emphasis on real-world interaction is evident in many discussions of human sensemaking. David Kirsh, for example, argues that manipulating external representations can transform the cost structure

associated with the inference landscape,¹⁰ improving our access to hidden or implicit information. Clearly, agents in the model presented here don't actively explore or manipulate their information environments, which highlights one focus area for future modeling work.

Our model's shortcomings in this respect obviously limit the range of application contexts in which it can be administered or used. In particular, the model isn't intended to replace real-world empirical investigations involving human agents. Rather, its best use is to guide decisions about what kind of real-world studies with human subjects should be undertaken. Computational models let us run a large number of experiments exploring different combinations of variables—which is difficult or impossible to perform in real-world coalition environments. The results from computer simulation experiments can help generate specific hypotheses concerning collective cognitive performance, and researchers can subsequently evaluate these hypotheses using real-world empirical analyses.

Sensemaking has been described as a macrocognitive function that lets individuals and groups make sense of information and develop the understanding required for effective decision making.¹¹ It's a particularly important topic for research in military coalition contexts because it constitutes a central part of the value chain that leads from higher-quality networking technologies to greater levels of command agility and mission effectiveness.³

The computational model we present aims to improve our understanding of how specific features of the coalition communication environment might affect collective sensemaking

processes. By combining the results from both real-world experiments and computer simulation studies, we hope to derive some insight into how collective cognition is affected by these features. This, in turn, will help guide scientific research and technological development in ways that enable military coalitions to press maximal cognitive benefit from the informational, technological, and human resources they have at their disposal. ■

Acknowledgments

This research was sponsored by the US Army Research Laboratory and the UK Ministry of Defence and was accomplished under agreement number W911NF-06-3-0001. The views and conclusions contained in this document are those of the authors and shouldn't be interpreted as representing the official policies, either expressed or implied, of the US Army Research Laboratory, the US government, the UK Ministry of Defence, or the UK government. The US and UK governments are authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation herein.

References

1. G. Klein, S. Wiggins, and C.O. Dominguez, "Team Sensemaking," *Theoretical Issues in Ergonomics Science*, vol. 11, no. 4, 2010, pp. 304–320.
2. *A Cognitive Approach to Situation Awareness: Theory and Application*, S. Banbury and S. Tremblay, eds., Ashgate Publishing, 2004.
3. D.S. Alberts and J.J. Garstka, *Network Centric Operations Conceptual Framework Version 2.0*, US Office of Force Transformation and Office of the Assistant Secretary of Defense for Networks and Information Integration, US Dept. of Defense, 2004.
4. G. Convertino et al., "Collaborative Intelligence Analysis with CACHE: Bias Reduction and Information Coverage," *Human-Computer Interaction Consortium 2006 Winter Workshop on*

THE AUTHORS

Paul R. Smart is a senior research fellow in the School of Electronics and Computer Science at the University of Southampton. His research interests include cognitive modeling, knowledge engineering, and the Semantic Web. Smart has a PhD in experimental psychology from the University of Sussex. Contact him at ps02v@ecs.soton.ac.uk.

Katia P. Sycara is a research professor in the School of Computer Science at Carnegie Mellon University. Her research interests include autonomous agents, planning, Web services, human-agent interaction, negotiation, and game theory. Sycara has a PhD in computer science from the Georgia Institute of Technology. She's a fellow of IEEE and AAAI. Contact her at katia@cs.cmu.edu.

Collaboration, Cooperation, and Coordination, 2006; www.researchgate.net/publication/228420565_Collaborative_intelligence_analysis_with_CACHE_Bias_reduction_and_information_coverage.

5. S. Poteet et al., "Linguistic Sources of Coalition Miscommunication," *Proc. NATO RTO HFM-142 Symp. Adaptability in Coalition Teamwork*, NATO Research and Tech. Organization, 2008; www.cso.nato.int/Pubs/rdp.asp?RDP=RTO-MP-HFM-142.

6. P.R. Smart, "Studying the Impact of Minority Views in a Computational Model of Collective Sensemaking: The Role of Network Structure," *Proc. 4th Int'l Conf. Advanced Cognitive Technologies and Applications (COGNITIVE 12)*, IARIA, 2012; <http://eprints.soton.ac.uk/id/eprint/337993>.

7. T.R. Schultz and M.R. Lepper, "Cognitive Dissonance Reduction as Constraint Satisfaction," *Psychological Rev.*, vol. 103, no. 2, 1996, pp. 219–240.

8. E. Hutchins, "The Social Organization of Distributed Cognition," *Perspectives*

on Socially Shared Cognition, L. Resnick et al., eds., Am. Psychological Assoc., 1991, pp. 283–307.

9. W.R. Sieck, L. Rasmussen, and P.R. Smart, "Cultural Network Analysis: A Cognitive Approach to Cultural Modeling," *Network Science for Military Coalition Operations: Information Extraction and Interaction*, D. Verma, ed., IGI Global, 2010, pp. 237–255.

10. D. Kirsh, "Interaction, External Representation and Sense Making," *Proc. 31st Ann. Conf. Cognitive Science Society*, Cognitive Science Soc., 2009; <http://csjarchive.cogsci.rpi.edu/proceedings/2009/papers/259/index.html>.

11. G. Klein, B. Moon, and R. Hoffman, "Making Sense of Sensemaking 2: A Macrocognitive Model," *IEEE Intelligent Systems*, vol. 21, no. 5, 2006, pp. 88–92.

IEEE VR 2013

20th IEEE Virtual Reality

16-20 March 2013

Orlando, FL, USA

IEEE VR 2013 is the premier international conference and exhibition on virtual reality. Find the brightest minds, the most innovative research, the leading companies and the most stimulating discussions in the fields of virtual environments, augmented reality, haptics, and 3D user interfaces.

Register today!

<http://conferences.computer.org/vr/2013>

