Studying the Impact of Minority Views in a Computational Model of Collective Sensemaking: The Role of Network Structure

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

***Abstract*—A series of experiments were performed in order to explore the effect of communication network structure on collective sensemaking under a variety of informational conditions. A multi-agent computational model of collective sensemaking was used in which each agent was implemented as a constraint satisfaction network. Within the simulations, agents were tasked with the interpretation of information indicating the presence of a particular object, and they were allowed to share information with other agents while performing this task subject to the constraints imposed by the structure of a communication network. In all simulations, a minority of agents (5) received evidence in favor of one interpretation, while a majority of agents (15) received evidence in favor of a conflicting interpretation. Communication networks with four types of topological structure (i.e., disconnected, random, small-world and fully-connected) were used in the experiments. The results suggest that network topology influences the extent to which minority views are able to influence collective cognitive outcomes. 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.**

***Keywords*-sensemaking; distributed cognition; social influence; network science; social information processing.**

I. Introduction

The emergence of network science as a scientific discipline in recent years has focused attention on how the various features of social and communication networks can affect aspects of human thought and action. The study of social networks, for example, has yielded a number of important findings regarding the effect of network structure on the adoption of specific ideas or technological innovations (see [1]), and such findings have important implications for those who seek to influence the spread of specific beliefs, attitudes and behavioral patterns throughout a target community. Other studies within the network science literature have attempted to shed light on how particular features of networks (for example, time-variant changes in network structure) affect the dynamics of cognitive processing, either at the individual or social (collective) levels [2], [3].

In general, there are two approaches to studying the relationship between cognitive processes and the features of social/communication networks. The first of these involves the use of human subjects who are observed in a specific experimental context. This is the approach typically adopted by members of the cognitive and social psychology communities. A second approach involves the use of multi-agent simulation environments and computational models of cognitive processing. This approach has the advantage of enabling researchers to explore scenarios that would be difficult or impossible to explore with human subjects; however, it is an approach that has been criticized in terms of the psychological plausibility and relevance of the computational methods used to simulate aspects of human cognition [4], [5]. Clearly, any computational model needs to make some simplifying assumptions about the real-world; otherwise it loses the elements of computational tractability and explanatory concision that make it useful as an aid to both analysis and comprehension. However, the problem with many multi-agent simulations is that the agents are too rigid and simplistic to be even approximate simulacra of their real-world human counterparts. In many cases, the agents are represented by single, time-variant numerical values, and they lack any kind of internal cognitive processing capability. This is arguably a crucial limitation, and one that needs to be addressed in the context of future work on multi-agent cognition (see [6]).

The current series of experiments seeks to further our understanding of the relationship between network structure and the dynamics of collective cognition. It relies on the use of multiple, inter-connected constraint satisfaction networks (CSNs) in order to provide a model of what is called collective sensemaking. Collective sensemaking is an extension of sensemaking abilities at the level of individual agents. At the individual level, sensemaking has been defined as "a motivated, continuous effort to understand connections (which can be among people, places, and events) in order to anticipate their trajectories and act effectively" [7]. The notion of collective sensemaking simply extends this ability to the realm of multiple agents. It refers to the effort of multiple agents to coordinate their individual sensemaking capacities in order to accurately interpret some body of environmental information, which is typically ambiguous, conflicting or uncertain in nature. Collective sensemaking is therefore a specific form of socially-distributed cognition in which cognitive processing can be seen as involving the information processing efforts of multiple, interacting individuals [8].

The use of CSNs to provide a model of sensemaking abilities is justified on the grounds that 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 set of beliefs subject to the constraints imposed by background knowledge, initial expectations and received information. In addition to this, CSNs have been used to model a wide variety of psychological phenomena, including belief revision, explanation, schema completion, analogical reasoning, causal attribution, discourse comprehension, content-addressable memories, cognitive dissonance and attitude change [9], [10], [11], [12], [13]. This helps to address some of the issues of psychological plausibility and relevance that are often associated with computer simulation studies of collective cognition.

By representing individual agents as CSNs, each engaged in the process of making sense of environmental information, a model of collective sensemaking can be implemented by allowing agents to exchange information with one another via communication links. These links enable agents to share information about their beliefs with other agents, and they support investigations into how factors like network structure might affect collective sensemaking abilities. CSNs have been used in precisely this way to explore the dynamics of collective sensemaking in a couple of previous studies [14], [8]. In one study, Smart and Shadbolt [14] used CSNs to examine the way in which collective sensemaking was affected by manipulations that altered the dynamics of inter-agent communication. The current work extends these initial observations in a number of ways. Firstly, the current work uses a larger number of agents, which are connected together in a greater variety of ways. All of the experiments described by Smart and Shadbolt [14] involved the use of only 4 agents, and these agents were connected together using a fully-connected network topology. The current study uses 20 agents in each simulation, and these agents are connected together using a variety of network topologies. Secondly, the study by Smart and Shadbolt [14] aimed to investigate the effect of a number of communication variables (e.g., communication frequency) on collective sensemaking using a fixed body of environmental information. This contrasts with the current work where the aim is to examine the effect of different network structures on agents' ability to make sense of different bodies of environmental information.

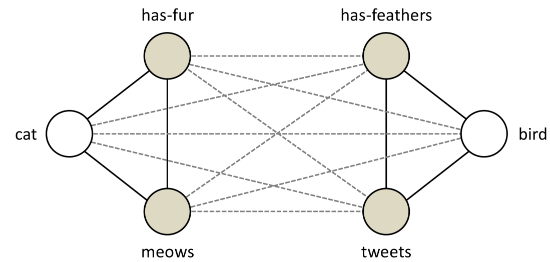

Figure 1. Organization of cognitive units in a single agent. Nodes represent cognitive units, each of which consists of two processing units. Solid lines represent excitatory connections between the units, while broken lines represent inhibitory links. Colored circles represent beliefs about the features of objects (feature beliefs), while white circles represent beliefs about the object type (object beliefs).

The specific aim of the current work is to examine the way in which different communication network structures (i.e., communication networks with different structural topologies) affect collective sensemaking performance under a variety of informational conditions. Three experiments were performed in which a minority of agents were presented with evidence supporting one interpretation and a majority of agents were presented with evidence supporting an alternative, conflicting interpretation. The evidence presented to agents at the outset of the simulation caused agents to adopt different beliefs or views, and these were subsequently subject to modification across successive processing cycles. In the absence of any social influence (i.e., input from other agents), agents developed cognitive states that were consistent with the initial evidence they were provided with. However, in situations where agent communication was enabled, agents were forced to factor in the views of other agents into their emerging belief states. Communication networks with different structural topologies might be expected to differentially influence the dynamics of collective cognitive processing in this situation; however, it is unclear what the nature of this influence is at present. The three experiments reported in this paper aim to shed light on this issue.

II. Method

A. Computational Architecture

Details of the computational model used in the current study are described in Smart and Shadbolt [14]. The CSNs used in the current study are based on a model developed by Schultz and Lepper [15], called the consonance model. This model was used by Schultz and Lepper to replicate the findings associated with a number of studies purporting to study the phenomenon of cognitive dissonance [16].

Each agent within the current 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 or view held by an agent. The agent depicted in Figure 1 shows the internal structure of all the agents used in the current study. As can be seen from Figure 1, each agent consists of 6 cognitive units, and each of these units represents beliefs about two types of animals, namely cats and birds. Four of the units represent beliefs about the features typically associated with objects, while the other two units represent beliefs about the object itself. For convenience, the former are called 'feature beliefs', while the latter are referred to as 'object beliefs'.

The pattern of connectivity between the units in Figure 1 reflects the compatibility or consistency between different types of beliefs. Cognitive units can be connected to other cognitive units via inhibitory or excitatory links, and these reflect the background knowledge or experience that an agent has of a particular domain. 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. Thus, in our simulations, agents are presented with the task of making a decision about the type of an object (an animal) based on limited information about the presence of its associated features (e.g., whether it has feathers or fur). The result is that cognitive units are always connected together in a way that reflects the association of particular animals with particular features. For example, the 'cat' cognitive unit is always positively connected to the 'meows' and 'has-fur' units because if an agent believes that a cat is present then they will also believe in the presence of cat-related features. Similarly, the 'bird' cognitive unit is positively connected to the 'tweets' and 'has-feathers' units because of the natural association between birds and these features.

Note that it is a feature of these models that mutually reinforcing sets of beliefs will tend to emerge across the course of successive processing cycles. Thus, if the 'has-feathers' unit was activated at the beginning of a simulation, then both the 'bird' and 'tweets' units would also become active at later stages during the simulation. All these units would then reinforce one another's activation throughout the remainder of the simulation. By the end of the simulation, the agent could be said to believe that the object was both a bird and that it had tweeting features, even though no evidence for these particular beliefs was provided at the outset of the simulation (e.g., in the form of an initial activation vector). The process of forming beliefs in the absence of evidence reflects the inferential capabilities of an agent. These inferences occur against a backdrop of background knowledge, which is encoded in the organizational structure (i.e., pattern of inhibitory and excitatory connections) of the CSN.

In addition to a sign, indicating whether a connection exerts an excitatory or inhibitory influence on its target node, each connection has a weighting that determines the amount of influence it exerts. Although these weights could assume a variety of values, in the current study we limit all weights to values of either 0.5 (excitatory) or -0.5 (inhibitory).

B. Computational Processing

Computational processing in each CSN proceeds by the activation of particular cognitive units at the beginning of a simulation. This initial activation is deemed to represent an agent's beliefs or views at the outset of the simulation. In the case of feature beliefs, the activation of each cognitive unit reflects the agent's beliefs or views in response to information about the features of different types of objects (in our case, either cats or birds). Computational processing then occurs via the spreading of activation between the cognitive units of the CSN following the pattern of excitatory and inhibitory linkages between the units (see Figure 1). 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, $floor$ 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 (as mentioned above, w_{ij} assumes values of either 0.5 or -0.5), 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; in the current simulation, however, we fixed the $resist_i$ parameter at a value of 0.5 for all nodes.

At each point in the simulation, n nodes were randomly selected and updated according to equations 1 and 2, where n corresponds to the number of nodes in the (agent-level) CSN. Agents were then allowed to communicate information to their connected peers (i.e., their immediate neighbors in the communication network). Communication involved each agent contributing activation to connected agents based on the activation levels of their own constituent nodes.

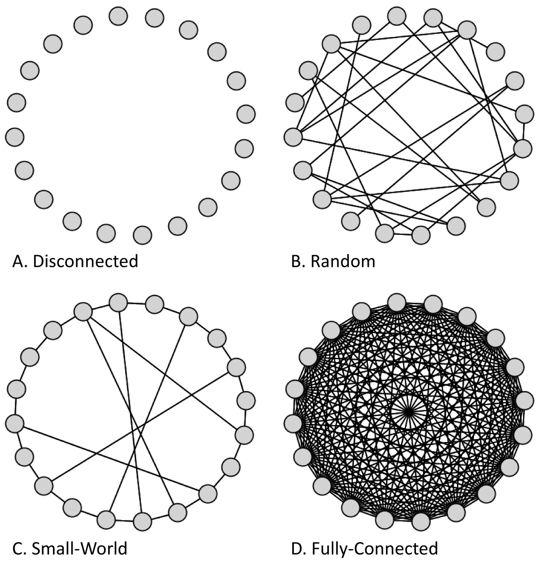

Figure 2. Examples of the different network structures used in the experiments. Nodes represent agents and lines indicate channels of communication.

Each node was associated with a parameter, $comminput_i$, which is the weighted sum of activation received from all talking agents. This parameter was 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, $comminput_i$ was incorporated into the activation equations by extending equation 3 as follows:

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

Once the communicated activation had been incorporated into the node's current activation level, $comminput_i$ was reset to zero in order to avoid repetitive presentation of the same communicated information across successive processing cycles. Processing continued until the pattern of activation in each of the agent networks had settled down. Typically, in the case of our simulations, 20 processing cycles were sufficient for a stable pattern of activation to be achieved.

C. Communication Networks

In order to examine the effect of network structure on collective performance, agents were organized into communication networks with different structural topologies. Examples of these network structures are shown in Figure 2. In the case of disconnected networks, no communication links were included between the agents. Experimental conditions involving this type of network structure allowed us to examine the performance of agents when inter-agent communication was effectively disabled and each agent functioned autonomously. The second type of network structure was the random network. Networks of this type were created following the same procedure as that described in Mason et al [3]. Bidirectional links between agents were added at random between the agents until a specific number of bidirectional links (i.e., 1.3 times the number of agents) had been created. Given that all our simulations involved 20 agents, the number of bidirectional links added to random network configurations was (1.3 $\times$ 20 =) 26. An additional constraint used in the creation of random networks was that every agent could be reached from every other agent (i.e., the network had a single component). The procedure for generating the small-world network was also the same as that reported by Mason et al [3]. In this case, agents were initially connected into a ring structure. Six agents were then selected at random and each of these randomly selected agents was connected to another randomly selected agent subject to the constraint that connected agents were at least 6 agents apart in the ring topology. Finally, in the case of the fully-connected network, all agents were connected to all other agents.

D. Procedure

Three experiments were performed in order to explore the effect of different networks structures on the processing of minority views. These experiments used the four types of network structure described in Section II-C; the main difference between the experiments was with respect to the initial activation vectors used to activate nodes at the start of each simulation. Table I shows the activation vectors that were used in each experiment.

At the beginning of each simulation, 20 agents were created and configured according to the description in Section II-A. All agents were created with identical cognitive architectures. Agents were then configured into one of four communication network structures using the procedures described in Section II-C.

At the beginning of each simulation 5 agents were selected at random and assigned to a 'Minority' group; the remaining (15) agents were assigned to a 'Majority' group. The agents in each group were then initialized with the activation vectors shown in Table I. In the case of Experiment 1, the minority group were presented with weak evidence that favored a cat interpretation, while the majority group were presented with weak evidence that favored a bird interpretation. The aim of this experiment was to examine the effect of different network structures on sensemaking performance when only weak evidence was available to all

Table I
INITIAL ACTIVATION VECTORS USED IN THE EXPERIMENTS.

Experiment	Agent Group	cat	has-fur	meows	bird	has-feathers	tweets
Experiment 1	Minority Group (5 agents)	0.0	0.1	0.0	0.0	0.0	0.0
	Majority Group (15 agents)	0.0	0.0	0.0	0.0	0.1	0.0
Experiment 2	Minority Group (5 agents)	0.0	0.2	0.0	0.0	0.1	0.0
	Majority Group (15 agents)	0.0	0.1	0.0	0.0	0.2	0.0
Experiment 3	Minority Group (5 agents)	0.0	0.5	0.0	0.0	0.0	0.0
	Majority Group (15 agents)	0.0	0.0	0.0	0.0	0.1	0.0

the agents. In Experiment 2, conflicting information was presented to agents in both the minority and majority groups. The aim of this experiment was to examine the impact of communication network structures under conditions of high uncertainty or ambiguity. The third experiment examined the effect of communication network structures under conditions where a minority of agents received strong evidence in favor of one interpretation, and a majority of agents received weak evidence in favor of a conflicting interpretation.

After the initial activation levels had been established for each agent, the simulation commenced and continued for 20 processing cycles. The activation level of both the 'cat' and 'bird' cognitive units was recorded from each agent throughout each simulation. The activation levels of these units at the end of the simulation (i.e., at the 20th processing cycle) was subjected to statistical analysis.

The design for all three experiments was a two-way (4 $\times$ 2) factorial design with a between subjects factor of Network Structure (with levels reflecting the types of networks tested: Disconnected, Random, Small-World and Fully-Connected) and a within subjects factor of Belief Type (with two levels - Cat and Bird - each corresponding to the activation of the 'cat' and 'bird' cognitive units, respectively). The data were analyzed using Analysis of Variance (ANOVA) procedures. For each experiment, significant two-way interactions were explored by running separate one-way ANOVAs at each level of the Belief Type factor (i.e., separate ANOVAs were performed for both the 'cat' and 'bird' cognitive unit data). Comparisons between the activation levels obtained for cognitive units across the 4 network structure conditions were made using Tukey's HSD test.

Fifty simulations were run for each of the different network structure conditions. This yielded a total of (50 $\times$ 4 =) 200 simulations for each experiment. Given that there were 20 agents in each network and we recorded from two cognitive units, each experiment yielded a total of (20 $\times$ 200 =) 8000 data points for each experiment.

III. RESULTS

A. Experiment 1: Minority and Majority Views Based on Weak Evidence

The results from Experiment 1 are shown in Figure 3A. ANOVA revealed significant main effects of Network Structure ($F_{(3,3996)}$ = 11.908, $P < 0.001$) and Belief Type ($F_{(1,3996)}$ = 2879.846, $P < 0.001$), as well as a significant two-way interaction ($F_{(3,3996)}$ = 73.720, $P < 0.001$). The interaction was explored by running two reduced one-way ANOVAs at each level of the Belief Type factor. These analyses revealed significant differences between the network structures for both the 'cat' ($F_{(3,3996)}$ = 73.713, $P < 0.001$) and 'bird' ($F_{(3,3996)}$ = 73.726, $P < 0.001$) cognitive units. Post hoc comparisons using Tukey's HSD test revealed that, in the case of the 'cat' cognitive unit, all the activation levels were the same, except for the fully-connected network in which the level of activation was significantly below that seen with other network types. Similar results were obtained in the case of the 'bird' cognitive unit: activation levels were the same across all network types with the exception of the fully-connected network. In the case of the 'bird' cognitive unit, activation in the fully-connected network was significantly above that seen with other network types.

B. Experiment 2: Minority and Majority Views Based on Conflicting Evidence

The results from Experiment 2 are shown in Figure 3B. As is suggested by Figure 3B, the activation of the 'bird' cognitive unit was, in general, higher than that of the 'cat' cognitive unit (main effect of Belief Type: ($F_{(1,3996)}$ = 2965.586, $P < 0.001$)). ANOVA also revealed a significant main effect of Network Structure ($F_{(3,3996)}$ = 4.820, $P < 0.01$) and a significant two-way interaction ($F_{(3,3996)}$ = 4.820, $P < 0.01$). ANOVAs at each level of the Belief Type factor revealed significant differences across the network conditions for both 'cat' ($F_{(3,3996)}$ = 4.817, $P < 0.01$) and 'bird' ($F_{(3,3996)}$ = 4.824, $P < 0.01$) cognitive units. Post hoc analyses revealed that the activation of the 'bird' cognitive unit was significantly higher in the fully-connected network versus the disconnected network. The reverse result was seen in the case of the 'cat' cognitive unit (i.e., activation levels in the fully-connected network condition were significantly below that seen in the disconnected network condition). No other differences between the network structure conditions were observed at either level of the Belief Type factor.

C. Experiment 3: Minority Views Based on Strong Evidence versus Majority Views Based on Weak Evidence

Figure 3C shows the results for Experiment 3. As with the other experiments, ANOVA revealed significant main effects of Network Structure ($F_{(3,3996)}$ = 11.038, $P < 0.01$), Belief

 ISBN: 978-1-61208-218-9

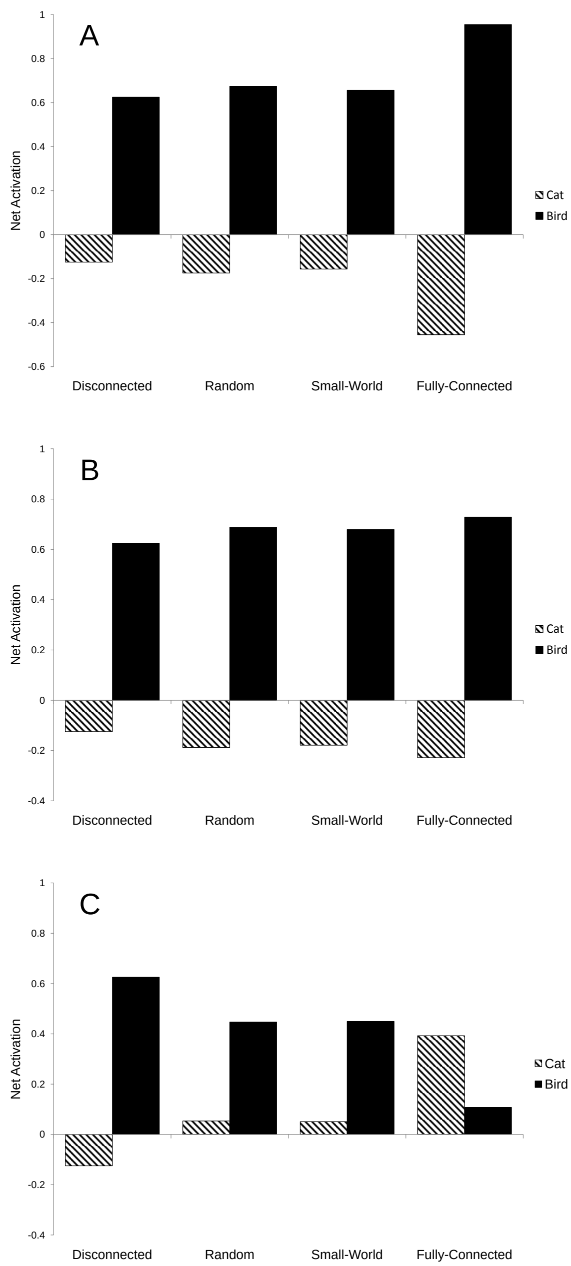

Figure 3. Mean activation levels of ‘cat’ and ‘bird’ cognitive units in each of the four network structure conditions for Experiment 1 (A), Experiment 2 (B) and Experiment 3 (C). Standard error of the mean (SEM) is not shown. In all cases, SEM was less than 0.03.

Type ($F_{(1,3996)}$ = 196.247, $P < 0.001$) and a significant two-way interaction ($F_{(3,3996)}$ = 93.131, $P < 0.001$). Separate one-way ANOVAS at each level of the Belief Type factor revealed significant differences between the network structure conditions for both the ‘cat’ ($F_{(3,3996)}$ = 93.140, $P < 0.001$) and ‘bird’ ($F_{(3,3996)}$ = 93.121, $P < 0.001$) cognitive units. Post hoc analyses revealed significant differences between the activation levels of the ‘cat’ cognitive unit across all the network conditions with the exception of the small-world and random networks, which did not differ from each other. The same pattern of results was seen in the case of the ‘bird’ cognitive unit (i.e., significant differences were observed across the different network structures with the exception of the small-world and random networks). As is suggested by Figure 3C, in the case of 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 between these two extremes. In the case of the ‘bird’ cognitive unit, the reverse pattern of results was obtained: activation was lowest in the fully-connected network, highest in the disconnected network, and at intermediate levels in the random and small-world networks.

IV. Discussion

The results of this study suggest that communication networks with different structural topologies differentially affect performance in a simulated sensemaking task under a variety of informational conditions. Some of the most interesting results were obtained with fully-connected networks. When minority and majority groups were presented with weak evidence (Experiment 1), a performance profile emerged in which the majority view predominated. This is reflected in the higher average activation of the ‘bird’ cognitive unit in the fully-connected network condition relative to that seen with other network types. It was also the case that the minority view (reflected in activation of the ‘cat’ cognitive unit) had less influence in fully-connected networks relative to other networks (this is reflected in the fact that activation of the ‘cat’ cognitive unit in the fully-connected network condition was below that seen with other types of network). Fully-connected networks therefore seem to result in the discounting of minority views in favor in majority opinions when weak evidence is presented to all agents. In other words, when weak evidence is presented to all agents and agents are configured into communication networks with fully-connected topologies, then the evidence available to minority groups has less influence on final collective judgements compared to other types of communication network structure (i.e., networks with random and small-world topologies). This particular result may be attributable to the greater speed at which information is shared between agents in fully-connected networks. Because all agents receive information from all other agents in fully-connected networks, there is a tendency for minority views to be swamped by the weight of initial majority opinion. Other types of network, such as the small-world and random networks tend to support information propagation rates that are slower than those seen in fully-connected networks, and thus there is greater chance that minority views will have time to become established before majority influence begins to take effect.

Experiment 3 differed from Experiment 1 in that it examined the effect of network structure on performance in cases where minority group members were presented with strong rather than weak evidence. In this situation, fully-connected network topologies yielded a collective outcome in which 'cat' beliefs predominated. This contrasted with the results obtained with all other network types in which 'bird' beliefs predominated. The results seem to indicate that fully-connected networks are particularly effective at integrating minority views into collective judgements when the evidence in favor of the minority view is high and the evidence in favor of the majority view is weak. Random and small-world networks were not as effective as fully-connected networks in producing this effect, although they were better than the situation observed in the disconnected network condition. These results may be interpreted in terms of the nature of the dynamics of social influence in fully-connected networks. Fully-connected networks enable strong, but uncommon, evidence to quickly influence the beliefs of all agents before weaker, contradictory evidence has had time to contribute to opposing beliefs. In the case of small-world and random networks, weaker evidence has longer to contribute to beliefs that are progressively more resistant to change across successive processing cycles.

The profile of results seen in Experiment 3 is particularly interesting when compared to the results obtained in Experiment 1. In Experiment 1, fully-connected networks were the most effective in terms of attenuating minority influence; the same networks, in Experiment 3, were the most effective in terms of promoting the influence of minority views. The difference between these results stems from the relative differences in the initial strength of minority versus majority opinion.

Experiment 2 studied the effect of conflicting information that was presented to both minority and majority groups. Notwithstanding the significant differences between fully-connected and disconnected networks, the results from this experiment suggest that network topology has little effect on collective sensemaking in this particular informational condition. In all cases, inter-agent communication seemed to result in a performance profile in which majority views predominated.

V. Conclusion and Future Work

The current work explored the effect of different network structures in a simulated version of a collective sensemaking task. This work extends earlier work that has used CSNs to explore the dynamics of collective cognition [14], [17]. The results suggest that network topology influences the extent to which minority information is able to influence collective cognitive outcomes. 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.

There are variety of ways in which the current work could be extended. One direction for future research is to explore models in which agents have more complicated belief structures; for example, agents could have a greater number of beliefs arranged in more complex configurations. In the current study, all cognitive units were configured so that positive (excitatory) connections had a weight of 0.5 and negative (inhibitory) connections had a weight of -0.5. One extension of the current work is thus to examine the effect of variable weightings between cognitive units. Since each linkage between cognitive units represents a psychological implication or association between belief states, the weighting associated with inter-cognition linkages may be deemed to reflect the strength of this implication or association. We assume that inter-cognition linkages are acquired as a result of prior learning, experience or training, and that they reflect the background knowledge (including assumptions, stereotypes and prejudices) that an agent brings to bear on a particular problem-solving activity. Inasmuch as this is true, we can see individual variability in the inter-cognition linkages as reflecting differences in the background knowledge that was acquired before the simulation. Such manipulations may have value in terms of shedding light on how individual differences in background knowledge and experience can influence the dynamics of collective cognition in a variety of network-mediated communication contexts.

Acknowledgment

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

References

[1] D. J. Watts, *Six Degrees: The Science of a Connected Age*. London, UK: William Heinemann, 2003.

[2] P. R. Smart, T. D. Huynh, D. Braines, and N. R. Shadbolt, "Dynamic networks and distributed problem-solving," in *Knowledge Systems for Coalition Operations*, Vancouver, British Columbia, Canada, 2010.

[3] W. A. Mason, A. Jones, and R. L. Goldstone, "Propagation of innovations in networked groups," in *27th Annual Conference of the Cognitive Science Society*, Stresa, Italy, 2005.

[4] 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.

[5] P. R. Smart, T. D. Huynh, D. Braines, K. Sycara, and N. R. Shadbolt, "Collective cognition: Exploring the dynamics of belief propagation and collective problem solving in multi-agent systems," in *1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*, Maryland, USA, 2010.

[6] R. Sun, "Cognitive science meets multi-agent systems: A prolegomenon," *Philosophical Psychology*, vol. 14, no. 1, pp. 5–28, 2001.

[7] 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.

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

[9] K. Holyoak and P. Thagard, "Analogical mapping by constraint satisfaction," *Cognitive Science*, vol. 13, no. 3, pp. 295–355, 1989.

[10] W. Kintsch, "The role of knowledge in discourse comprehension: A construction-integration model," *Psychological Review*, vol. 95, pp. 163–182, 1988.

[11] B. Spellman, J. Ullman, and K. Holyoak, "A coherence model of cognitive consistency: Dynamics of attitude change during the persian gulf war," *Journal of Social Issues*, vol. 49, no. 4, pp. 147–165, 1993.

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

[13] D. Rumelhart, P. Smolensky, J. McClelland, and G. Hinton, "Schemata and sequential thought processes in pdp models." in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 2*, D. Rumelhart and J. McClelland, Eds. Cambridge, Massachusetts: MIT Press, 1986, pp. 7–58.

[14] 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.

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

[16] L. Festinger, *A Theory of Cognitive Dissonance*. Stanford, USA: Stanford University Press, 1957.

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

 ISBN: 978-1-61208-218-9