Modelling the Dynamics of Team Sensemaking: A Constraint Satisfaction Approach

Paul R. Smart* and Nigel R. Shadbolt

School of Electronics and Computing Science, University of Southampton, Southampton, SO17 1BJ, UK

**Corresponding author: ps02v@ecs.soton.ac.uk*

Abstract—*An approach to the modelling of team sensemaking is presented that relies on the use of multiple agents integrated into larger communication network structures. Sensemaking is cast as a type of constraint satisfaction problem, and thus the cognitive architecture of each agent within the model is implemented as a constraint satisfaction network. The effect of manipulating a number of communication variables (the frequency of inter-agent communication, the type of information communicated and the point at which inter-agent communication takes place) are explored in three computer simulation studies. The results suggest 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. These results are consistent with other results reported in the distributed cognition literature, and they suggest that the future use of constraint satisfaction network models could have value in terms of improving our understanding of socially-distributed cognition in military coalition environments.*

1. INTRODUCTION

Cognitive processing in military coalition organizations typically involves the coordinated effort of multiple individuals. This raises important questions about how the various features of military coalition environments might affect collective cognitive outcomes (for example, those associated with planning and decision-making). Features that might turn out to be important in this respect include the following:

1. **Communication Networks:** There is an increasing reliance on mobile ad hoc networks (MANETS) within military coalitions, and this may significantly affect the dynamics of inter-agent communication.
2. **Quality of Information:** The information that must be processed by military coalitions is seldom perfect. Information is often uncertain, ambiguous and conflicting, and in some cases it may be deliberately manipulated by hostile agents in order to subvert coalition decision-making and undermine coalition situation awareness.
3. **Cultural Differences:** Cultural differences between the members of a coalition organization may lead to miscommunication and misunderstanding [see 1]. This may stem from differences in language, knowledge, training and experience.
4. **Trust:** Different levels of trust between agents may lead to inefficiencies in information processing. For example, individuals from different groups may fail to adequately integrate available information into decision-making processes as a result of poor trust relationships.
5. **Limited Information Sharing:** The sharing of information within a coalition organization may be limited for a number of reasons. Security constraints may limit information access, communication networks may limit information distribution, and differences in information technology may make information difficult to exploit and integrate into ongoing cognitive processes.

Clearly, this is not an exhaustive list of features; however, it does highlight some of the potential factors that may affect cognitive processing in coalition organizations, especially when that processing is distributed across multiple individuals and culturally disparate groups.

In order to improve our understanding of collective cognition in military coalition environments, it is important to undertake empirical studies that systematically explore the effect of various types of features on the dynamics of collective cognitive processing. Unfortunately, the nature of the coalition environment means that studies with human subjects are both difficult to design and implement. As a result, it may be important to consider the use of computer simulation techniques (particularly those involving multi-agent systems) in which some aspect of collective cognitive processing is studied *in silico*. The problem with many multi-agent simulations, however, is that the agents lack the kind of features that make the models of psychological interest and relevance. In many cases, the agents consist of simple processing units that respond in relatively limited ways to simple and unstructured inputs.

In order to address these concerns, the current work adopts an approach to agent simulation that is grounded in an extensive body of work in the psychological literature. The approach involves the use of constraint satisfaction networks (CSNs), which are used to model the temporal evolution of an agent's cognitive states (beliefs, opinions, attitudes, evaluations and so on) across time. A CSN is apt for problems which involve the simultaneous satisfaction of multiple soft constraints, and many types of psychological phenomena may be seen as involving something like this

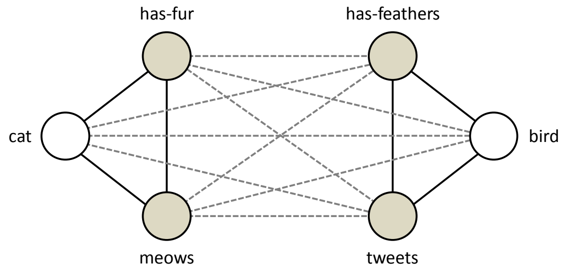

Figure 1. Organization of cognitive units in a single agent. Circles represent cognitive units, each of which consists of two processing nodes. Solid lines symbolize excitatory connections between the units, while broken lines symbolize inhibitory links. Shaded circles represent beliefs about the features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

capability. Indeed, CSNs have been used in simulation studies exploring a range of psychological phenomena, including belief revision, explanation, schema completion, analogical reasoning, causal attribution, discourse comprehension, content-addressable memories, cognitive dissonance and attitude change [2-7]. Furthermore, CSNs have been used to explore the dynamics of socially-distributed cognition [8], and this is precisely the kind of cognitive processing we are interested in when it comes to coalition organizations.

The specific cognitive phenomenon of interest in the current study is sensemaking, which 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" [9]. Sensemaking has also been discussed at the collective or team level as well as at the level of individual agents. Klein et al [10] thus define team sensemaking as "the process by which a team manages and coordinates its efforts to explain the current situation and to anticipate future situations, typically under uncertain or ambiguous conditions." We argue that sensemaking can be approached as a form of constraint satisfaction problem, and CSNs can thus be usefully applied to model the sensemaking efforts of individuals. In addition to this, we suggest that networks of CSNs can be used to model at least some of the constituent processes associated with sensemaking at the collective level (i.e. team sensemaking). In particular, we suggest that the flow of information between individual CSNs in a network of CSNs provides a rough analogue to inter-agent communication in collective sensemaking situations. Inasmuch as this is the case, simulations consisting of multiple CSNs may enable us to explore the effect of different inter-agent communication variables on team sensemaking outcomes.

The experimental studies described in the current paper require agents to perform a particular task. Briefly, agents are presented with information about the features of a particular object, and they then have to form a belief about what the object might be. Agents establish beliefs about the object by integrating presented information with background knowledge. A complicating feature of the task is that agents are not presented with perfect information

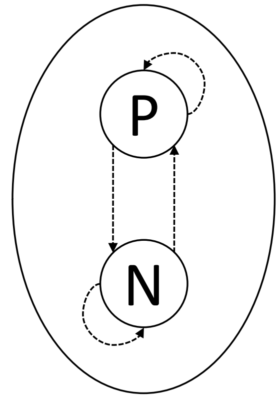

Figure 2. The anatomy of a cognitive unit. Each unit consists of two nodes, one of which is the positive pole (P) and the other is the negative pole (N). P and N are connected via mutual inhibitory links. In addition, each node has an auto-regulatory connection that connects each node to itself. The function of this auto-regulatory link is to dampen the node's activity at each processing cycle [see 4, for details]. All links within the cognitive unit are inhibitory, as is indicated by the broken lines in the figure.

about an object's features. Instead, the information resembles that seen in many military conflict situations; i.e. the information is incomplete, uncertain, ambiguous and conflicting. To keep the discussion as simple as possible and to avoid the introduction of domain-specific terminology, the task used in the current experiments centres on the processing of simple feature sets associated with two types of animals, namely cats and birds.

2. Computational Model

2.1. Architecture

The computational model developed to explore team sensemaking in the current paper is based on the consonance model developed by Schultz and Lepper [4]. Each agent within the model is implemented as a CSN (following the design specification outlined by Schultz and Lepper), and these individual CSNs are connected together to form a network of CSNs (i.e. a network of networks). The nodes which make up each CSN are organized into a number of cognitive units, each of which represents a particular belief held by the agent. For example, the agent in Figure 1 consists of 6 cognitive units, each of which represents beliefs about two types of animals, namely cats and birds. Four of these units represent beliefs about the features typically associated with objects. They are called 'feature beliefs'. Other units represent beliefs about the object itself. They are called 'object beliefs'.

Internally, each cognitive unit consists of two nodes which are connected together in a mutual inhibitory fashion (see Figure 2). One of these nodes is labelled as the 'positive pole' (P), and the other is labelled 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

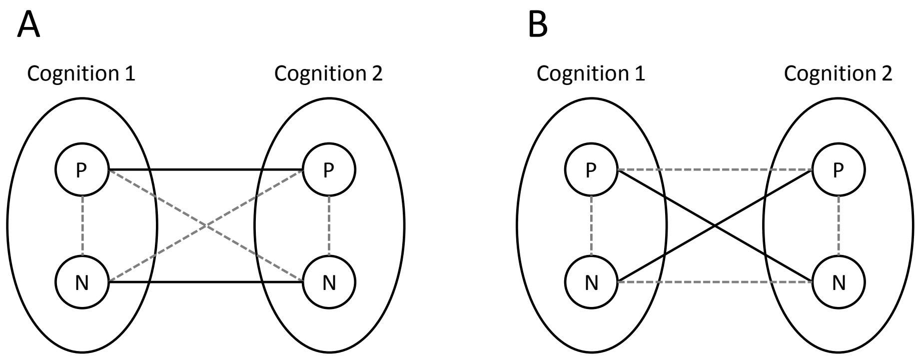

Figure 3. Figure showing the connectivity pattern for both positive (A) and negative (B) linkages between cognitive units. Excitatory connections are symbolized by solid lines, inhibitory connections by dashed lines.

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.

Within an 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 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 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 always connected to the 'tweets' and 'has-feather' units because of the natural association between birds and these features.

Cognitive units that represent incompatible beliefs are connected via inhibitory connections. Thus, in our case, if an agent was to believe that an unknown object corresponded to a cat, it would not make sense to simultaneously believe that the object had a feature naturally associated with a bird. For example, it would not make sense to connect the 'cat' unit and the 'has-feathers' unit with an excitatory connection. To do so would create a situation where an agent held beliefs that conflicted with the structure of the target domain. Agents in our simulations always attempt to make sense of conflicting, ambiguous and uncertain information by changing their beliefs in ways that are consistent with both external information (sensory data) and background knowledge. For this reason, it is important that the pattern of excitatory and inhibitory links between cognitive units (reflecting an agent's background knowledge) coincides with the properties of the domain to which the agent's beliefs apply[1].

The pattern of connectivity between cognitive units for all agents in our simulations is shown in Figure 1. Architecturally, 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, with two connections emanating from each of the nodes in the two cognitive units. In our model, a positive connection between cognitive units is represented by the wiring diagram shown in Figure 3A, and a negative connection is represented by the wiring diagram shown in Figure 3B. With eight connections comprising each inter-cognition linkage, and a connectivity pattern between cognitive units based on that seen in Figure 1, it can be seen that each agent consists of 120 connections between the nodes of different cognitive units. Within each cognitive unit, there are an additional two (inhibitory) connections linking the two nodes, and a further two auto-regulatory connections. The result is that each agent in the model consists of 12 nodes, 6 cognitive units and 144 connections.

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).

Agents are connected into networks via the inclusion of linkages between agents. Each inter-agent connection is in fact a set of bidirectional linkages between the corresponding nodes of each agent. Thus, if agent A is connected to agent B, then bidirectional connections will exist between each of the corresponding nodes of agent A and B. The nodes in the 'cat' unit of agent A will be connected to the nodes of the 'cat' unit in agent B, the nodes of the 'has-fur' unit in agent A will be connected to the nodes of the 'has-fur' unit in agent B, and so on. Based on this organizational scheme, we can see that each inter-agent linkage consists of 24 individual connections (two connections for each node, one going from agent A to agent B and one going from agent B to agent A).

[1] Although it is not a feature of our model, it is obviously possible for cognitive units to be neither compatible or incompatible (i.e. they are not associated in any way). In this case, the cognitive units can be connected by a zero weighted set of links, or the inter-cognition linkage can simply be removed.

As with the connections within an individual agent, the connections between agents can have both a sign and a weighting value. To keep our simulations simple, we use a single weighting value of 0.5 for all inter-agent connections. Clearly, this is an oversimplification relative to the real world since the weighting of the inter-agent connections determines the degree to which one agent influences another, and agents influence one another to *different* degrees based on a variety of factors (e.g. the level of trust that exists between them). One future extension of this work is thus to examine the effect of variable inter-agent connection weights, perhaps allowing these weights to change dynamically throughout the course of a simulation (see Section 4.3).

Note that in the work reported here, we use a small network consisting of only 4 agents, with all agents connected to one another. In this sense, the agent network has a fully-connected topology. Given the small number of agents involved in the simulations it did not make sense to explore alternative kinds of network topology (e.g. random or small-world networks). As with the manipulation of inter-agent connection weights, this dimension of the computational model provides a potential direction for future research (see Section 4.1).

2.2. Initial Activation Vectors

When an agent is presented with a body of environmental information corresponding to the presence or absence of particular features of an object, they will attempt to make sense of the information by systematically changing their beliefs based on the interactions between cognitive units. A processing cycle thus consists of the initial presentation of an activation vector providing information about the presence or absence of particular features in the environment. In our simulations, we assume that the initial information an agent is presented with is based on indirect observation (perhaps via a sensor) and is therefore vulnerable to distortion. In particular, we assume that the information provided to agents resembles that presented to analysts and decision makers in hostile conflict situations (i.e. the information is incomplete, uncertain, ambiguous, and conflicting). For this reason, all the information presented to agents in the current simulations is imperfect. Perfect information would consist of an activation vector that, when processed by the agents, would immediately result in the agents expressing beliefs that perfectly coincided with ground truth. For example, if agents are tasked with making a decision about the type of an unknown object, and the object is a bird, then in a situation of perfect information the agents should express the belief that the object is a bird when presented with the environmental information, and they should express high confidence in this decision outcome. In contrast, in a situation of imperfect information, agents will form beliefs, but they may have little confidence in these beliefs. In addition, some beliefs may be activated at the same time as other, incompatible beliefs.

Node	Agent A	Agent B	Agent C	Agent D	Total
has-fur	0.3	0.3	0.5	0.0	1.1
has-fur	0.0	0.0	0.0	0.0	0.0
meows	0.2	0.2	0.2	0.0	0.6
meows	0.0	0.0	0.0	0.0	0.0
has-feathers	0.0	0.0	0.2	0.5	0.7
has feathers	0.0	0.0	0.0	0.0	0.0
tweets	0.0	0.0	0.0	0.0	0.0
tweets	0.0	0.0	0.0	0.0	0.0

Table 1. Initial activation vectors presented to agents. Shaded cells indicate the vectors used to activate the P nodes of cognitive units. The same set of activation vectors was used at the outset of every simulation reported in the current paper.

In the simulations presented here, agents were presented with imperfect information that resulted in the weak activation of beliefs, some of which were incompatible with other beliefs (the actual activation values presented to agents across all simulations are presented in Table 1). We consider this state to resemble an initial state of confusion or uncertainty that must be resolved over the course of successive processing cycles in order to reach a particular decision outcome. The simultaneous activation of inconsistent beliefs (cognitions) in our model leads an agent to revise their beliefs in ways that satisfy the constraints imposed by both the external information and their internal background knowledge. This constitutes the main driving force for internally-driven cognitive change in our model, and it is the means by which agents attempt to make sense of ambiguous, uncertain and conflicting information.

2.3. Processing Dynamics

In each simulation, the activation vectors presented in Table 1 were used to determine the initial activation levels of the feature nodes of each of the four agents[2]. The activation of each node was then updated across successive processing cycles using the following update 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 unit i at time $t + 1$, $a_i(t)$ is the activation of unit i at time t, *ceiling* is the maximal level of activation of the node[3], *floor* is the minimum activation of the node (zero for all nodes), and net_i is the net input to unit i, which is defined as:

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

where $resist_i$ is a measure of the resistance of unit i to

[2] In addition to these activation values, the positive pole of each of the 'bird' cognitive units within each agent was set to 0.5. This initial activation of the 'bird' cognitive unit was intended to represent the expectations of each agent. Each agent therefore expected to observe a bird, whereas in fact the weight of evidence that agents were presented with suggested that the target object was, in fact, a cat.

[3] In the current simulations, positive poles had a ceiling value of 1.0, whereas negative poles had a ceiling value of 0.5. The reason for this difference in ceiling values is explain in Schultz and Lepper [4].

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 less resistant to change than others; however, in the current simulation, we fixed the $resist_i$ parameter at a value of 0.5 for all nodes.

At each point in the processing cycle, n nodes are randomly selected and updated according to equations 1 and 2, where n corresponds to the number of units in the network (12 in our case). This continues until the pattern of activation in each of the agent networks settles down. Typically, in the case of our simulations, 20 processing cycles were sufficient for a stable pattern of activation to be achieved.

In some simulations, agents were allowed to communicate with one another on particular cycles. When communication was enabled, each talking agent contributed activation to the nodes of listening agents. Each node was associated with a parameter, $comminput_i$, which is the weighted sum of activation received from all talking agents. On cycles where agents communicated, 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:

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

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.

3. Simulations

Based on the kinds of features that may affect collective cognition in military coalition environments (see Section 1), there are clearly a large number of variables that could be the focus of empirical investigations. In the present section, we focus on the communication between agents, and we report the results obtained with three types of experimental manipulation.

3.1. Experiment 1: Communication Frequency

In any form of socially-distributed information processing, it is important to understand the effect that different levels of inter-agent communication have on collective cognitive outcomes. One might assume that a team of individuals will function at its best when they are allowed to communicate to all members of a team at all stages of a problem-solving process. This assumption may not, however, be correct. A number of studies have thus suggested that precipitant forms of information sharing might result in sub-optimal levels of team performance [11].

Method

In order to explore the effect of communication frequency on the temporal evolution of belief states in the aforementioned multi-agent model, we studied groups of agents under three experimental conditions. In the first condition ('No Communication'), we proscribed all inter-agent communication. All agents therefore processed the information that they were initially presented with, and they were not allowed to communicate the intermediate or final results of their processing to other agents. In the second condition ('Low Frequency Communication'), agents were allowed to communicate on every fourth processing cycle. Given that each simulation was run for a total of 20 cycles, this enabled agents to communicate with one another a total of 5 times during the course of the simulation. In the third condition ('High Frequency Communication'), agents were allowed to communicate on every cycle of the simulation. This meant that at every cycle of the simulation, all agents were talking to all other agents and communicating information about their belief states.

For the purposes of this experiment, we did not opt to limit inter-agent communication in any way except with regard to frequency. Thus on every cycle that agents communicated, all agents communicated with *every* other agent. In addition, all agents communicated information about *all* of their beliefs; we did not limit communication to a particular subset of agent beliefs.

Each simulation lasted for 20 cycles, and we ran 50 separate simulations for each of the agents in each of the three experimental conditions (i.e. a total of 50 × 4 × 3 = 600 simulations).

Results

The results of the experiment are presented in Figure 4. Figure 4 shows the net activation of the 'cat' and 'bird' cognitive units for each of the agents in each of the three experimental conditions. As can be seen from the results, when no communication between the agents was allowed, agents A and B both settled on a cognitive state in which the 'cat' belief predominated. This contrasted with the results for agent C in which both the 'cat' and 'bird' beliefs were active. This agent appeared somewhat ambivalent with regard to the object that was presented, and they were not able to 'make up their mind' as to what was the correct solution. Agent D showed no such problems: D settled on a solution by the tenth processing cycle and was unchanged thereafter.

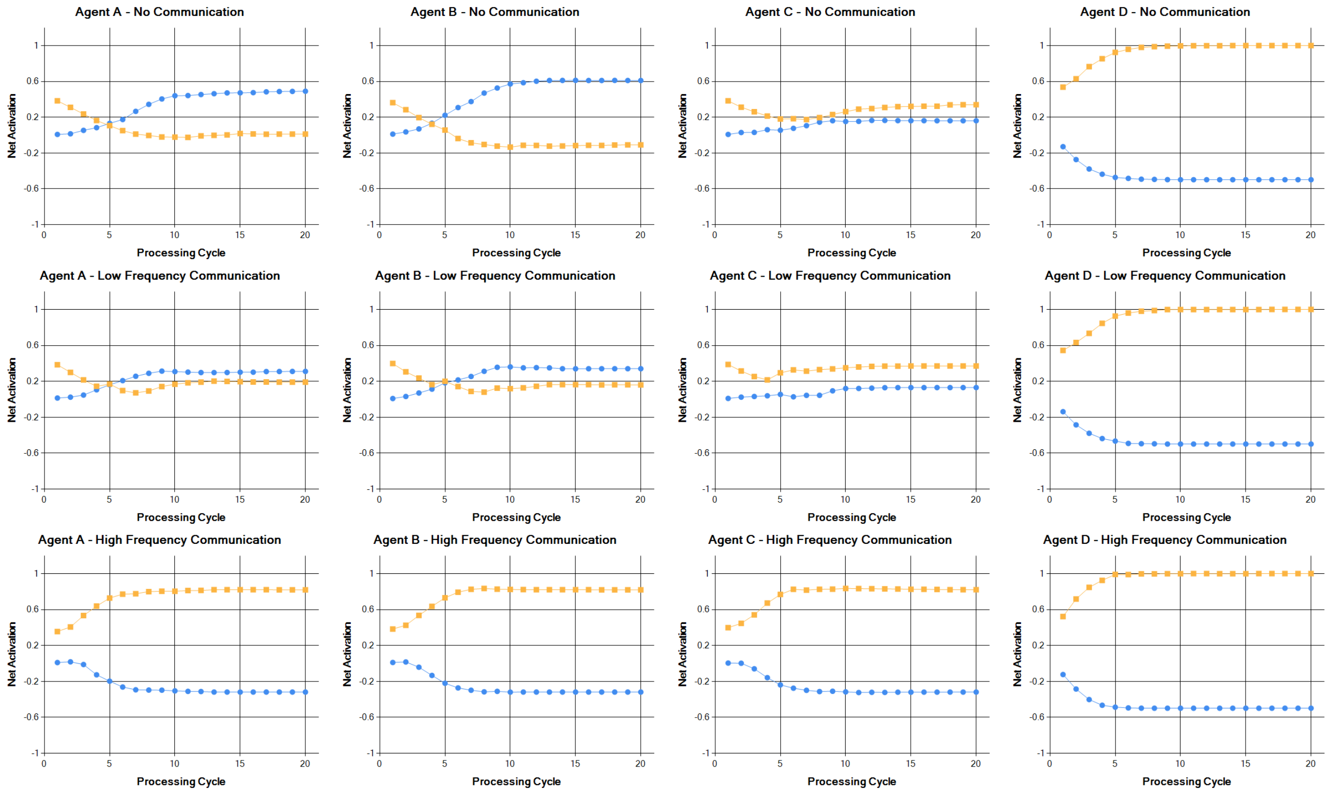

Figure 4. Results for Experiment 1. Each graphic tile shows the results obtained for a particular agent in a particular experimental condition. The first row illustrates the results obtained in the 'No Communication' condition, the second row illustrates the results obtained in the 'Low Frequency Communication' condition, and the third row illustrates the results obtained in the 'High Frequency Communication' condition. The blue data series represents the net activation of the 'cat' cognitive unit (reflecting the agent's belief that a cat is present), while the yellow data series represents the net activation of the 'bird' cognitive unit (reflecting the agent's belief that a bird is present).

A similar pattern of results was obtained in the second condition, which permitted low levels of communication between the agents (second row of tiles in Figure 4). One difference from the results seen in the first condition is that agents A and B appeared less confident about their 'cat' beliefs at the conclusion of the simulations.

Finally, when frequent communication was enabled in the third condition, a different pattern of results was obtained. Now all the agents rapidly developed the belief that the unknown object was a bird. This is despite the fact that the weight of evidence from the body of initial information (see Table 1) suggests that the unknown object was a cat[4].

3.2. Experiment 2: Type of Communicated Information

Experiment 1 varied the frequency with which agents were allowed to communicate, but it did not seek to control the kind of information that was actually communicated. When we look at the kinds of beliefs that agents in the simulation may entertain, we can see two distinct types of beliefs: those corresponding to beliefs about the features of objects ('feature beliefs'), and those corresponding to beliefs about the objects themselves ('object beliefs'). As an extension of Experiment 1, therefore, we might consider the extent to which the *type* of communicated information affects the dynamics of social information processing.

In situations where only the 'feature beliefs' are communicated, agents may be seen as restricting their communication to the information they have received from external sources (e.g. sensor systems) (or at least as limiting their conversation to the 'thoughts' they have about such information). In situations where only the 'object beliefs' are communicated, agents may be seen as communicating their opinions on what the unknown object actually is (i.e. the outcome of the sensemaking process). This latter situation may correspond to a state-of-affairs in which decision outcomes are communicated without any supporting information (or rationale). The former situation may correspond to a state-of-affairs in which received information is subjected to some limited information processing and then communicated without a final decision being reached.

Method

The experimental design was a 2 × 2 factorial design in which one factor (frequency of communication) was derived from the previous experiment. In this case, the factor had two levels: 'Low Frequency Communication' (LFC) and 'High Frequency Communication' (HFC) corresponding to the low and high frequency communication conditions of the previous experiment. The other factor (information type) had two levels based on whether agents communicated

[4] Note that the total activation associated with cat-related feature units in Table 1 (summed across all agents) is higher than that associated with bird-related feature units.

Cognitive Unit	LFC/Feature Beliefs	LFC/Object Beliefs	HFC/Feature Beliefs	HFC/Object Beliefs
Agent A – Cat	0.40	0.61	0.04	-0.32
Agent A – Bird	0.10	-0.11	0.46	0.82
Agent B – Cat	0.43	0.43	-0.02	-0.35
Agent B – Bird	0.07	0.07	0.52	0.85
Agent C – Cat	0.16	0.13	0.07	-0.44
Agent C – Bird	0.34	0.37	0.43	0.94
Agent D – Cat	-0.50	-0.50	-0.50	-0.50
Agent D – Bird	1.00	1.00	1.00	1.00

Table 2. Results of Experiment 2. Table shows the net activation values for the 'cat' and 'bird' cognitive units for each agent in each experimental condition at the conclusion of the simulations (i.e. at the 20th processing cycle of the simulations). Each of the values in the table cells is averaged over 50 simulations. (LFC = Low Frequency Communication; HFC = High Frequency Communication)

information about 'feature beliefs' or 'object beliefs'. As for Experiment 1, each simulation lasted for 20 cycles, and we ran 50 separate simulations for each agent in each experimental condition (i.e. a total of 50 × 4 × 4 = 800 simulations).

Results

The results from Experiment 2 are presented in Table 2. Table 2 shows the net activation values for the 'cat' and 'bird' cognitive units for each agent in each experimental condition at the conclusion of the simulations (i.e. at processing cycle 20).

In the case of 'feature beliefs' the pattern of results obtained for the low frequency communication condition was similar to that obtained for the low frequency communication condition of Experiment 1. In both cases, agents A and B end the simulation with the 'cat' belief predominating, and agents C and D both end the simulation with the 'bird' belief predominating. For the high frequency communication condition, the results are again similar to those seen in Experiment 1: all agents conclude the simulation believing that the unknown object corresponds to a bird.

In the case of 'object beliefs', the pattern of results is also similar to that seen in Experiment 1. The results from these experiments therefore seem to suggest that the distinction between 'feature beliefs' and 'object beliefs' in the current simulation is of little significance when it comes to understanding the potential impact of different types of communicated information on agents' sensemaking capabilities.

3.3. Experiment 3: Timing of Inter-Agent Communication

In addition to the frequency of communication, we can also think about the potential effect of the timing of communication on collective cognitive processing. For example, rather than allow communication at every cycle of the simulation (as was the case in the aforementioned high frequency communication conditions) we could allow agents to communicate with high frequency at particular points in a simulation (for example, at the beginning or the end of the simulation). One reason to think that this may be important is that previous computer simulation studies have suggested that early communication may cause agents to give undue weight to certain features of a problem. In one study, for example, the cognitive anthropologist, Edwin Hutchins, investigated the effect of inter-agent communication on the ability of agents to arrive at an accurate shared interpretation of ambiguous environmental information. What Hutchins [8] discovered was that early forms of interaction led to a situation of confirmation bias in which agents failed to give due weight to information that conflicted with their initial interpretation of some external state-of-affairs.

Method

To understand the effect of early or late communication on the dynamics of agent processing, we ran an experiment in which agents were allowed to communicate on cycles 1-5 or cycles 16-20 of a 30 cycle simulation. Given the apparently insignificant role played by information type in determining cognitive outcomes (see Experiment 2), we allowed agents to communicate information about *all* of their beliefs.

Results

The results of the study are presented in Figure 5. As can be seen from Figure 5, when agents were allowed to communicate at the beginning of the simulation, they rapidly converged on a particular interpretation in which the 'bird belief' predominated. In contrast, when communication was restricted to a later stage of the simulation, agents tended to settle on a pattern of beliefs that resembled that seen when they engaged in no communication at all (see Figure 4 – first row of tiles). This pattern of results is particularly noticeable for agents A and B. The temporal evolution of their belief profiles was completely transformed as a result of participation in early (as opposed to late) forms of inter-agent communication. It is also worth noting that the activation of the 'cat' unit in agent C was higher (relative to the 'bird' unit) in the late communication condition. This contrasted with the pattern of results seen in Experiments 1 and 2 in which the net activation of the 'bird' unit was generally higher than the 'cat' unit for agent C. Given that the weight of evidence suggests the presence of a cat (see Table 1), this might suggest that delayed communication has a positive effect in terms of promoting correct interpretations under highly

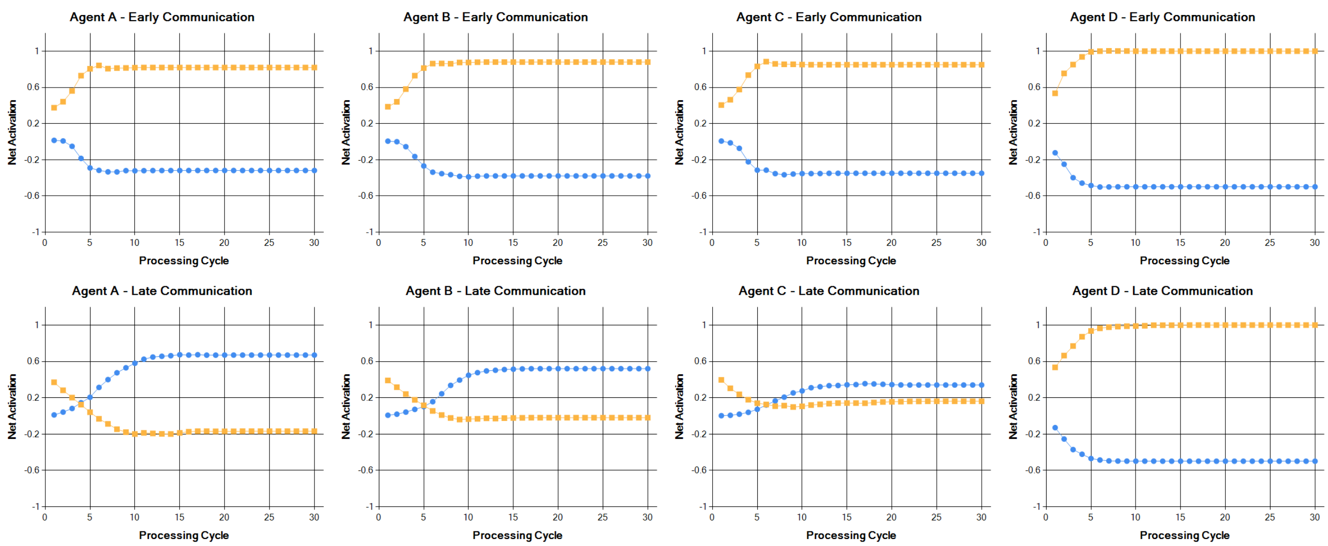

Figure 5. Results for Experiment 3.Each graphic tile shows the results obtained for a particular agent in a particular experimental condition. The first row illustrates the results obtained in the 'Early Communication' condition, while the second row illustrates the results obtained in the 'Late Communication' condition. The blue data series represents the net activation of the 'cat' cognitive unit (reflecting the agent's belief that a cat is present), while the yellow data series represents the net activation of the 'bird' cognitive unit (reflecting the agent's belief that a bird is present).

uncertain or ambiguous conditions – precisely the kind of conditions identified by Klein et al [10] in their analysis of team sensemaking.

3.4. General Discussion

The results of these studies suggest the following:

1. As the frequency of communication increases, agents tend to rapidly converge on a common interpretation of noisy environmental data (see Experiment 1). Thus, when communication was allowed at every cycle of a simulation, agents quickly settled on the belief that a bird was present (see Figure 4 – third row). Unfortunately, this is not consistent with the weight of evidence provided by the initial datasets (see Table 1) which indicate the presence of a cat.
2. The type of information communicated by agents (information about object features or object type) seems to have little effect on the pattern of results obtained in low and high frequency communication conditions (see Experiment 2).
3. Communication that takes place early on in a simulation tends to lead to rapid convergence on a particular interpretation of the data, and agents seem to express high confidence in these interpretations. This result is not seen when communication takes place later on in the simulation. In this case, the pattern of results resembles that seen when agents engage in no communication whatsoever.

The effect of high frequency communication and early communication in these studies is particularly interesting. What seems to happen in these conditions is that undue significance is given to agents' initial expectations about the kind of object they will encounter. Agents started the simulation with the expectation that they would encounter a bird, and this may have been reinforced in situations where agents transmitted information about their initial views and interpretations of the environmental data. This explanation cannot, however, account for all of the simulation results obtained. In Experiment 2, for instance, communication about 'object beliefs' was disabled in one of the experimental conditions, and yet this did not significantly alter the pattern of results obtained.

In addition to the effect of initial expectancies, it may be that precipitant forms of information sharing lead agents to assign undue significance to information that is compatible or consistent with their initial views, especially when the information they receive is ambiguous or uncertain. Thus even in situations where agent communication is restricted to 'feature beliefs', this still allows for the transmission of some information that is consistent with an agent's initial expectations, and, in this situation, the agent may assign greater weight to the 'consistent' information they receive from other agents and discount the information they have acquired or gathered themselves.

4. Model Extensions and Future Work

The simulation results described in the previous section represent an early attempt to develop a computational model for collective sensemaking that draws on the previous use of CSNs to examine psychological phenomena. Previous work has suggested how CSNs could be used to model aspects of individual cognition [e.g. 4], and the current work attempts to extend these efforts by applying CSNs to the problem of collective cognition. Currently, however, the approach features a number of limitations which need to be addressed in future work. In addition, there are a rich range of further simulations studies that could be undertaken based on a consideration of the kind of features that might plausibly affect collective cognition in military coalition environments. Some ideas for future research and development work are detailed below.

4.1. Network Topology

As mentioned in Section 2.1, one extension of the current work is to explore the effect of different network topologies (e.g. random, small-world, and fully-connected) on collective sensemaking. Given the size of the agent teams examined in the current work, it did not make sense to vary the network topology in systematic ways. However, other work has shown that network topology affects the rate at which information flows through a community of agents, and this can sometimes exert effects on collective cognition [11, 12].

Another factor to consider here concerns the difference between dynamic and static network topologies [see 11, 13]. Most multi-agent simulations use communication networks with static topologies; i.e. topologies that are relatively invariant across the course of information processing. In contrast to this situation, many of the networks that are encountered in the real world have topologies that are highly dynamic. A consideration of dynamic networks is particularly important in the context of military coalition environments because of the increasing reliance on MANETS and wireless communication technologies [see 11].

4.2. Agent Networks

The agents in the current model feature a relatively small number of cognitive units (i.e. 6), with uniform absolute weights between the units (i.e. 0.5). One extension of the current model involves the development of networks with a greater number of cognitive units, a greater diversity of cognitive units, and variable weightings between the units.

A greater diversity of cognitive units could be realized by recognizing the existence of different types of cognitions. Schultz and Lepper [4] for example, recognize three types of cognitive unit: justification, evaluation and behaviour units. The significance of this distinction lies in the value assigned to the resistance parameter associated with the nodes within a cognitive unit. In the current study, we use a common resistance parameter of 0.5 for all nodes; however, Schultz and Lepper adopt a different scheme in which some units rely on smaller resistance parameters. This has the effect of making the cognitive units differentially resistant to change, which may have important consequences for cognitive processing[5].

The weightings between cognitive units in our study have fixed absolute values of 0.5 and these weightings are invariant across the course of the simulation. 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 exercise.

The actual values for the weights associated with inter-cognition linkages could be established in a number of ways. They could be subject to manual manipulation (although this clearly becomes unwieldy for simulations involving large numbers of agents with large numbers of cognitive units); they could be subject to some sort of randomized adjustment procedure (although this risks distancing agents from the knowledge-rich contingencies of the relevant task domain); or they could be acquired as a result of some prior learning experience. In respect of this latter possibility, previous work has shown how the connection weights for constraint satisfaction networks might be learned [see 14].

Individual differences in inter-cognition linkages may provide one way in which the current model can be applied to the problem of understanding the impact of cultural differences in military coalitions. 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 [15]. This means that cultural differences could be explored in the current model by varying inter-cognition linkages based on group membership criteria.

Unlike the case with inter-agent linkages, we do not see a major role for inter-cognition linkages that are dynamically updated across the course of a simulation. The reason for this is that we see sensemaking as more a case of applying background knowledge and experience in order to understand some situation or system state, rather than a case of learning new things. Sensemaking (at least in the kinds of cases we are exploring), we suggest, capitalizes on past learning, but it does not involve any new learning.

4.3. Inter-Agent Communication

The simulations presented here used networks in which all agents were connected to all other agents using a standard connection weighting (i.e. 0.5). Since the weight assigned to an inter-agent link affects the influence exerted by one agent over another, it is important to consider models in which such weights have variable values. One interpretation of the weight value is that it represents the degree of trust between agents. Thus high trust between agents is reflected in links that have high weight values, and low trust is reflected in links that have low weight values. In the extreme case, the

[5] This may provide one way of re-examining the effect of type of communicated information on collective sensemaking. One of the reasons Experiment 2 may have failed to yield any interesting results is because of the lack of any real differences between 'feature beliefs' and 'object beliefs'. The use of different resistance parameters could be one way in which differences between these cognitions could be introduced.

weights between two agents may be zero (reflecting a case of zero trust), in which case the effect is the same as if the agents never communicate.

Future work could explore the impact of variable and dynamically updated trust weights. One example of how this scheme might be implemented is provided by Van Overwalle and Heylighen [16]. They present a network-of-networks approach to modelling collective cognition in which each agent is represented by a recurrent auto-associative neural network. The weights associated with inter-agent linkages in their model are updated during the course of a simulation by an updating rule that factors in the similarity of belief states between communicating agents. In essence, agents who share similar beliefs (i.e. express agreement) have their channel of communication strengthened so that they exert a greater influence on one another during future processing cycles. The reverse applies to agents who have dissimilar beliefs.

4.4. Confidence

Another factor for consideration in the context of future work is the confidence that agents have in their belief states and the extent to which this influences the communication of belief states. In the current simulation, confidence is represented by the level of net activation of a cognitive unit: an agent has high confidence in a belief if the net activation of the unit approximates its maximum activation level. Future studies could explore the effect of restricting inter-agent communication to situations in which agents are required to have high confidence in their cognitions. This could be accomplished by incorporating a threshold function for cognitive state propagation.

5. Conclusion

Cognitive processing activities such as problem-solving and decision-making in military coalitions often depend on the coordinated interaction of multiple, distributed agents who communicate with one another via one or more coalition communication networks. In order to begin to understand the dynamics of collective cognitive processing in military coalition environments, we developed a computational model based on the use of multiple interacting agents, each of which was implemented as a CSN. Our results suggest that some aspects of inter-agent communication can affect agents' ability to correctly interpret bodies of ambiguous, uncertain and conflicting information. For example, when agents were allowed to participate in high frequency communication at the outset of a problem-solving task, they tended to converge on an inaccurate interpretation of environmental information. This effect was not observed when agents were allowed to independently 'think about' some body of information and then come to their own conclusion before engaging in communication. In these situations, an agent's beliefs were generally resistant to the effects exerted by agents who had come to different conclusions. These results are generally supportive of other results in the distributed cognition literature that have used CSNs [see 8], and they suggest that the future use of CSN-based models could have value in terms of improving our understanding of socially-distributed cognition in military coalition environments.

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. Poteet, S., et al. *Words Are Mightier Than Swords … and Yet Miscommunication Costs Lives!* in *2nd Annual Conference of the International Technology Alliance (ACITA'08)*. 2008. London, UK.
2. Kintsch, W., *The role of knowledge in discourse comprehension: A construction-integration model.* Psychological Review, 1988. **95**: p. 163-182.
3. Holyoak, K.J. and P. Thagard, *Analogical mapping by constraint satisfaction.* Cognitive Science, 1989. **13**(3): p. 295-355.
4. Schultz, T.R. and M.R. Lepper, *Cognitive dissonance reduction as constraint satisfaction.* Psychological Review, 1996. **103**(2): p. 219-240.
5. Spellman, B.A., J.B. Ullman, and K.J. Holyoak, *A coherence model of cognitive consistency: Dynamics of attitude change during the Persian Gulf War.* Journal of Social Issues, 1993. **49**(4): p. 147-165.
6. Thagard, P., *Explanatory coherence.* Behavioral and Brain Sciences, 1989. **12**(3): p. 435-502.
7. Rumelhart, D., et al., *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, Editors. 1986, MIT Press: Cambridge, Massachusetts. p. 7-58.
8. Hutchins, E., *The social organization of distributed cognition*, in *Perspectives on Socially Shared Cognition*, L. Resnick, J. Levine, and S. Teasley, Editors. 1991, The American Psychological Association: Washington DC, USA.
9. Klein, G., B. Moon, and R.R. Hoffman, *Making sense of sensemaking 1: Alternative perspectives.* Intelligent Systems, 2006. **21**(4): p. 70-73.
10. Klein, G., S. Wiggins, and C.O. Dominguez, *Team sensemaking.* Theoretical Issues in Ergonomics Science, 2010. **11**(4): p. 304-320.
11. Smart, P.R., et al., *Dynamic Networks and Distributed Problem-Solving*, in *Knowledge Systems for Coalition Operations (KSCO'10)*. 2010: Vancouver, British Columbia, Canada.
12. Lazer, D. and A. Friedman, *The Network Structure of Exploration and Exploitation.* Administrative Science Quarterly, 2007. **52**(4): p. 667-694.
13. Huynh, T.D., et al., *The Cognitive Virtues of Dynamic Networks*, in *4th Annual Conference of the International Technology Alliance (ACITA'10)*. 2010: London, UK.
14. Ackley, D.H., G.E. Hinton, and T.J. Sejnowski, *A learning algorithm for Boltzmann machines.* Cognitive Science, 1985. **9**(1): p. 147-169.
15. Sieck, W.R., L. Rasmussen, and P.R. Smart, *Cultural Network Analysis: A Cognitive Approach to Cultural Modeling*, in *Network Science for Military Coalition Operations: Information Extraction and Interaction*, D. Verma, Editor. 2010, IGI Global: Hershey, Pennsylvania, USA.
16. Van Overwalle, F. and F. Heylighen, *Talking nets: A multiagent connectionist approach to communication and trust between individuals.* Psychological Review, 2006. **113**(3): p. 606-627.