

Towards Cognitive Models of Communication and Group Intelligence

David Reitter and Christian Lebriere
Department of Psychology
Carnegie Mellon University

Abstract

In social tasks, communication and task execution are often co-dependent, yet they represent a tradeoff. We present a cognitive model of an experimental task consisting of a collaborative and competitive game played by groups of human participants organized in a small-world graph. In an experiment involving teams of humans playing a cooperative game, the effect of local communication policies on the efficiency and the performance of networked participants was observed. The model follows the ACT-R theory and provides a formalization of the decision-making processes and tradeoffs involved. A simulation of the hypothetical case of unnatural memory decay shows decreased performance and supports a prediction of the thesis that memory limitations have co-evolved with social structure.

Keywords: Communication Networks; Belief Propagation; Cognitive Modeling.

Introduction

Among our most fascinating abilities is communication and cooperation with others. To solve problems, we organize ourselves as teams and exchange information effortlessly. This achievement is the result of both innate and culturally acquired abilities. The combination of language-based communication, social interaction and the emergence of distributed human information systems is ubiquitous, but has seen relatively little empirical exploration. Many of the implicit (subconscious) decisions taken in communication may be the result of innate mechanisms of cognition, specifically memory. Cognition is the result of adaptation to the environment according to the rational analysis of cognition (Oaksford & Chater, 1999). Therefore, mechanisms may have evolved that influence individual's interaction with a social network. Advances in cognitive modeling on the one hand, and in the description of structural properties of networks on the other allow us to correlate universal properties of social structures with the cognitive limitations of the individual. *Performance or bounded rationality* may be a rational optimization that supports an emergent group intelligence, or specifically the externalization of memory through a community of collaborators.

Any computational cognitive model of team communication needs to take into account the simple fact that humans have limited attentional and memory capacities. To illustrate, consider a thought experiment. Assume a network of communicating, cognitive agents that pass questions and answers about distributed information from node to node. Suppose two extremes: in the first, agents have unlimited attentional capacity, that is, their ability to process information in the environment is not constrained. However, agents have no memory to retain information that may be needed later on. In the

second, agents have accurate and infinite memory, but lack the ability to process more than a single piece of information at a time. Which communication policies would be appropriate for the two cases? Perhaps, in the first case, agents would have to re-send questions regularly, while in the second, they would send out questions only once. Human attentional and memory capacities are both limited; hence, they are subject to a strategic tradeoff. The costs of producing and comprehending linguistic communication, of attentional bottlenecks and decaying memory are relevant to a model, as is the structure of the network. One argument in support of co-adaptation could come from rational behavior in a dynamic group context. This means that cognitive parameters such as memory decay and perceptual speed should influence communication and team success, and that cognitively plausible parameter choices in a model correspond to high performance during collaborative and communicative tasks, when other human performance parameters such as speed of communication are realistic.

Lab-based experiments, in which humans are networked artificially, playing an abstract game, represent a class of non-computational simulations that examine real-world phenomena at the group level. In this paper, we describe an experimental paradigm and previously presented data, a computational model and simulations that interpolate between the hypothetical cases of either perfect or absent human memory. The experiment looks beyond one-on-one dialogue to iterated communication in social word-of-mouth networks. It involves a new cooperative foraging game, the *Geo Game*, in which medium-size and larger (20+) teams of players engage in information exchange, exploration of a game landscape, and foraging for goods. The game elicits a cognitive tradeoff between a core task and a communication task, in which we expect the human cognitive apparatus to maintain a purposeful balance. In the experiment that provided the basis for our model, we used different communication policies to explore both sides of this tradeoff. To quantify the success of communication, we use a combined individual and team task as a benchmark. It measures the limits of individual and team performance in dynamically changing and time stressed environments. The cognitive model then explains performance data obtained with the experiment and is ultimately used to show how team performance would be affected if basic memory parameters were different from their known, constrained values. We propose that indi-

vidual cognitive constraints lead to communication and cooperation that optimize outcomes for the individual, but also for the dyad, the group and the community.

Prior Work

Much recent work on social networks has focused on their structural and resulting computational properties, but does so independently of two major aspects of real-world networks: *Humans, and their joint objectives, or task*. The study design presented here uses human participants and employs a task that individuals connected over the network have to execute. We provide a measurable objective and a task that depends both on individual performance and collaboration. We connect to work by Bavelas (1950) and Leavitt (1951), who detailed the effects of network structure with human nodes, arguing that networks with centrality show more stable performance, but increased dependency on those central nodes and decreased flexibility with respect to the integration of information.

We have been engaged in a program of scaling up cognitive models from the case of learning-based adaptation in dyads to a model of domain language evolution (Reitter & Lebriere (2011)), and extrapolating the results to investigate the effect of network structure on such convergence processes (Reitter & Lebriere, 2010b). The influence of structural properties in social tasks is evident even when payoffs are determined by individual performance: Judd et al. (2010); Kennedy (1999); Bhattacharyya & Ohlsson (2010). With our task we intend to also complement readily available datasets with an experimental design that gives us communication data ready to be analyzed in terms of its semantics and its timing. Datasets of language-based communication show the spread of memes (idealized ideas) or opinions (e.g., Twitter datasets), or they represent socializing or debate that is difficult to analyze and operationalize for the purposes of problem-solving research (Klimt & Yang, 2004). Such datasets, however, are not the result of explicit human collaboration in the context of a well-defined task. Tasks, in such datasets, are coincidental, while the task in Geo Game is central to driving communication and provides an objective basis for evaluating the effect of communication on performance.

The Geo Game

The Geo Game is a spatial search game, where all players simultaneously engage in a foraging task. Players are shown a map of several named cities, connected by a road network. At any given time, each player is located in one city or is moving between two connected cities; players are shown their own location, but not that of the others (see Fig. 1). The key features of this game are as follows. **Collaborative problem-solving:** Participants are tasked to find *items* by moving via roads to a city; they find one item at a time (their *goal item*).

Participants can visit a neighboring city (directly connected by a road to his current location) by clicking on its symbol on the map. Each city has a small number of items available; this item set differs for every city. Items located in a city are revealed to the participant only while the participant is “visiting” the city. After finding the goal item, a subsequent item in the list is shown to the participant. Moving from city to city takes time, so players are pressured to rely on their knowledge and that of others to find the city efficiently rather than to merely scavenge for items. Participants are asked to find as many items as possible within the duration of the session; a timer is displayed that indicated the remaining time. Once an item is found, it may be either removed permanently from the inventory of the city or it may be replenished (half the item types are replenishable). This renders the environment dynamic. **Dependency on communication:** The key feature of the game is that players can improve their performance through communication. They may exchange information through natural language, such as requests for an item they need or responses about the whereabouts of items. A chat interface allows each player to broadcast written messages to a fixed set of other players (“neighbors”). A player, receiving a message, may chose to re-broadcast the information since his set of neighbors is likely to differ from the original message sender. Players use this facility to ask about the whereabouts of their goal item, or to tell them about the locations of items. The task was designed so that the crucial decision a participant had to make was whether to send a piece of information or not, possibly based on its relevance. In practice, the language used by participants is simple and easy to decode automatically. Facts could be analyzed as either *Information Requests (Item, Requester)* (“I need a towel!”) or *Fact (Item, Location)* (“The cat is in Pittsburgh”). Players are organized in a graph structure, with participants as nodes, and vertices indicating communication channels. In our experiments, the underlying network topology is one in which the number of connections per node is power-law distributed; this topology is typically called a *small world*, because it is possible to connect any two vertices in the network through just a few links. The graphs used in the present experiment are “re-wired ring lattices” (Watts & Strogatz, 1998): starting with a ring in which each node is connected to exactly two neighbors, long-range links are added to a few pairs of randomly chosen nodes. The communication network embodies principles well-known to participants from online networks such as Facebook. **Individual and collective payoffs:** Players move around the road network until they find a city that can provide the item (the road network is fixed and unrelated to the graph that defines the social/communication network). Players are rewarded for finding a goal item with $r_0 = 1000$ game points. As

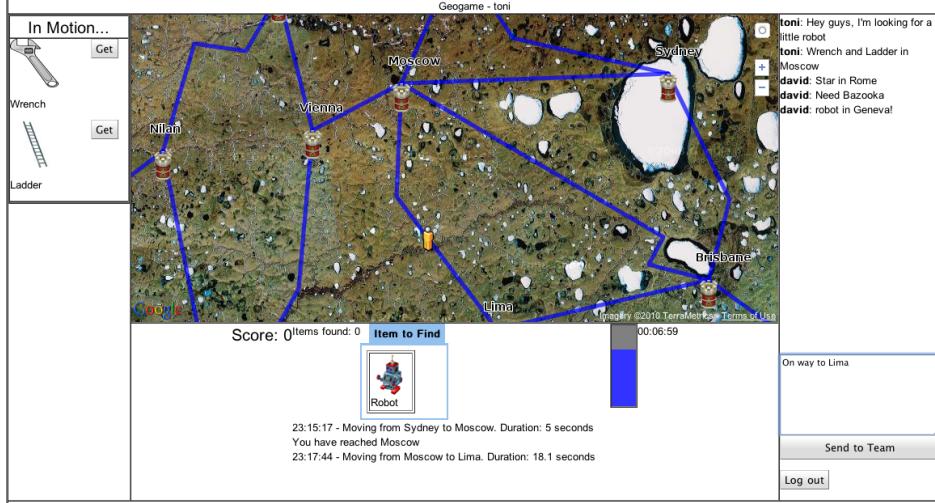


Figure 1: The Geo Game participant interface.

a further incentive to not only ask for information, but also provide useful information to their contacts, we reward each participant (node x) with game points whenever a node y to which they are connected obtains a goal item. The reward is highest for immediate neighbors and is defined as $r_d = \frac{r_0}{2^{d(x,y)}}$ for all y for which the distance between x and y expressed in intermediate nodes along the shortest connection, $d(x,y) \leq 3$. This reward distribution system follows common pyramid schemes and also Pickard et al. (2010). The final game score of a player is the sum of all rewards obtained during one session. **Task success metrics:** The game was designed to give measurable task success metrics, where the communication within the player network would be critical to task success. These include the accumulated payoff (the communicated objective), the number of goal items found, and the average time it takes to receive an answer to one's information request.

In the following, we summarize the experiment and its results (refer to Reitter et al., 2011, for the primary discussion). In the experiment participants were split into two groups with 17 participants each. These teams then played the game during two 30-minute sessions, one for each of two conditions (in permuted order). Participants were instructed to adapt their communication strategy: In the *dump* condition, participants were asked to indiscriminately broadcast a maximum of information available. In the *target* condition, they were asked to request and target information so that only such knowledge was disseminated (and passed on) that was known to be relevant to others in the network through prior requests.

The first question we asked is whether the information propagation policy influences task performance in the community. If either *dump* or *target* condition prove substantially advantageous, then the manipulation would appear to interact with the specific task. Under the assumption that humans can effortlessly integrate information, we would expect that the team directly benefits

from unlimited communication in the *dump* condition. This is not the case: Subjects scored higher in the *target* condition than in the *dump* condition (Table 1). Even when the main effects of message *quantity* and network position are taken into account do we observe benefits of targeted communication policy on residual task performance, i.e., we see an effect of message *quality* (Figure). An analysis of efficiency (score over messages) in different network positions suggested a decreased efficiency advantage from targeting with each additional neighbor.

Human networks can filter information and direct it to where it is needed, provided that humans are encouraged to make decisions about local information distribution. Communication processes are not necessarily mechanistic: people can pay attention to the informational needs of dialogue partners if asked to do so. In networks and in a situation where individuals communicate one-to-one, this appears to be beneficial to their task success.

The increase in performance suggests that maximizing “targeting” of information accommodates attentional limitations at the cost of losing information. Given memory decay, information that is left on a single (or a few) nodes risks being forgotten and have to be rediscovered through experience at possibly significant cost. Conversely, decreased targeting of information may improve the life and utility of information by spreading its availability and hedging against its decay, at the cost of attentional overload and interference with more important knowledge. Finding the optimum requires a computational account of human memory and attention.

Modeling the Geo Game task

The Geo Game model follows the ACT-R architecture Anderson (2007). It is implemented in ACT-UP (Reitter & Lebiere, 2010a) to provide a suitable abstraction of a high-fidelity model. The goal is to achieve a tractable computational model of the perceptual and cognitive processes involved in the foraging task that interact with the multi-agent interaction over the social network. Our

Cond.	Degree	Score	Received	Msgs	Msgs Sent	Items Taken	n
dump	2	19512		198.6	68.7	6.45	20
target	2	22576		111.9	37.3	7.1	18
dump	> 2	27455		299.6	64.4	6.36	11
target	> 2	28000		166.9	40.8	6.5	14

Table 1: Performance and communication: Human participant means in the two conditions for low- and highly connected nodes. The degree describes the number of links a node has; degrees ranged from 2 to 5 (power-law distributed). n shows the number of participants for each degree (networks were randomized between conditions).

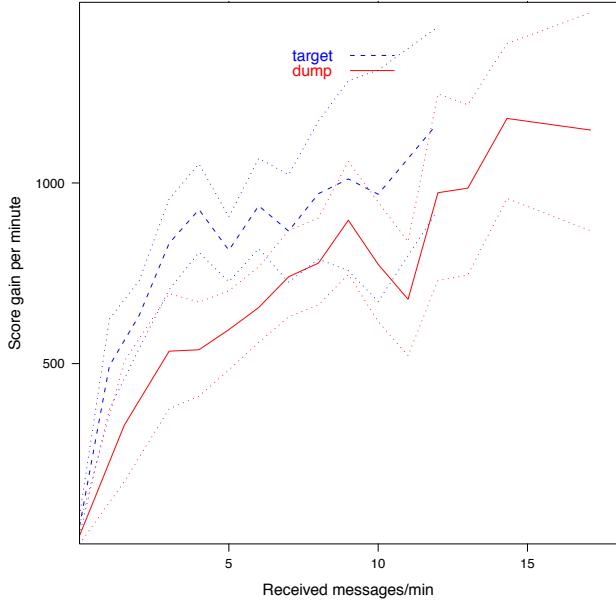


Figure 2: (a) Score gain per minute vs. number of messages received: At the same numbers of received messages, communities gained a higher score in the *target* condition than in the *dump* condition. Note that even in the *dump* condition, sparsely-connected players may, at some points, send very few messages. (Means are over per-user/minute measures for all users. Dotted lines show 95% confidence intervals, assuming bins per user for each minute elapsed and gaussian distributions).

model of the team is made up of a number of individual cognitive models, one for each player, which can communicate via a simulated natural language interface. They exchange information requests and replies to those. The models are organized in a graph as in the social network of the Geo Game experiment; they communicate by broadcasting to their network neighbors.

ACT-R Model

The basic structure of the model is an iterative process that allocates attention over two separate subtasks. The main subtask focuses on navigating around the map while the other involves processing messages received through the chat interface. That structure is similar to that of the model of the AMBR air traffic control simulation (Lebriere, 2005) and many other interactive tasks, suggesting a general, reusable modeling pattern.

The navigation subtask is focused on the target item to collect. When given a new target item, the model attempts to retrieve a known location for that item from memory. If successful, a path to that location from the current city is constructed, and the first step along that path is taken. Otherwise, the model explores the map randomly, selecting any of the directly reachable cities with equal probability. When arriving in a new city, the model first identifies the items available. If one of them matches the target item, the model picks it up and the cycle restarts with a new target item. Otherwise, the model takes its next step, either along a known path to the item's location, or another random step in the graph.

The chat management subtask is focused on handling the information flow to and from the user when not focused on the primary navigation task, that is while in transit between cities. The model reads the next message, processes it appropriately for the experimental conditions, then repeats the cycle if it hasn't yet arrived at its next destination. Processing the message means reading its contents, which can be of two primary types: requests for the location of a target item, and information about the presence of item(s) at a given location. For requests, the model determines if it knows the information needed, and if so replies to the message. In either case, if the model is in the *dump* experimental condition, the model retransmits the message as long as it doesn't remember doing so previously (to avoid overwhelming the network by repeatedly broadcasting the same message).

The two subtasks are not entirely independent, however. Each involves a processing step related to the other. In the navigation subtask, when the model arrives at a new city, in addition to searching the local items for the one that it needs, it will also broadcast the location of the items in the chat if in the *dump* condition, or of a specific item that is known to have been requested if in the *target* condition. Conversely, when processing chat messages, the model will check an information message to see if it contains the location of its current target item, and if so set course for that location.

Model performance

We obtained measurements from 30 model runs of 30 simulation minutes, simulating 17 players each. The network graph generation, the distribution of items over locations and the starting points of each player were

randomized as in the Geo Game experiment; the road network was fixed and the same as in the experiment. The parameters governing memory performance in the ACT-R architecture were left at their defaults (base-level learning decay $bll=0.5$, latency factor $lf=1.0$). The base-level constant blc was set to 5.0.

The model shows behavioral metrics slightly below those of the human teams. Each agent collects, on average, 4.6 items ($dump$: 4.11, $target$: 5.15) and obtains a mean score of 15,750 points. Mirroring the empirical results, agents receive significantly higher scores in the *target* condition than during *dump* ($\beta = -61504$, $SE = 12868$, $p < 0.001$ for score totals, i.e., an effect of 3,617 points per agent for *target*).

While model performance (expressed in game points) is variable and influenced by details of the communicative behavior, model runs during development suggested that the effect of communication policy is robust. We consider the model preliminary, mainly for two reasons. First, the model does not achieve full human performance. Second, the model adopts a specific strategy that involves committing seen item-location facts to memory. It is possible that subjects adapt their strategy depending on whether they expect to remember information. In an information-overload situation, they may refrain from memorizing, or they may lack the time for rehearsal. Thus, players may shift between internal and externalized memorization. The simulation does, however, implement the internal memory and some of the tradeoffs implicated in the experiment, allowing us to cautiously extrapolate to our hypothetical cases.

Model predictions

The model allows us to make predictions for two hypothetical situations described in the introduction. First, we simulate team performance for the case of poor memory performance. Setting a base-level learning decays of 1.5 and 3.0, we manipulate the cognitive architecture into forgetting facts at higher rates¹. Second, we simulate the absence of memory decay: how would the team perform if its members had infinitely-large, infinitely-lasting memory? Having set base-level learning decay to 0.0, the architecture allows models to retrieve information even long after it was last presented. In this condition, memory retrievals no longer prefer recent information. The model was not optimized towards the following results.

The simulation results show significantly decreased performance for the case of high decay $bll=3.0$ ($\beta = -27419$, $p < 0.005$, total points), but also for the case of no memory decay ($bll=0.0$, $\beta = -27612$, $p < 0.006$). The middle condition of increased decay at $bll=1.5$ yields a nonsignificant performance boost ($\beta = 13227$, $p = 0.16$).

¹ Bll represents the negative exponent; with the time since information presentation as base. The resulting activation influences the log-odds of retrieval of the piece of information.

Discussion

The Geo Game task manages to replicate a common real-life observation: that communication with collaborators is fundamental to task success when it is considerate and targeted, but detrimental to success if communication becomes overwhelming. Ignoring many less intuitive results concerning network structure, communication overload does not change our thinking about information exchange, but it helps define the correct cognitive model that explains network-based exchanges. The model we describe performs similarly well as humans and shows the same sensitivity to communication policies and attentional demands. The experiment is designed such that participants seek out information that originates from places far beyond their network neighborhood. Thus, the results show that careful communication practices impact the information state and attentional demands not just among interaction partners, but also further away in the network. A network of basic cognitive models can replicate those demands. The model, just as participants, can employ targeted communication, or *audience design* to address attentional limitations of their interaction partners.

The description of the model has not mentioned the specifics of the cognitive and perceptual operations involved in both subtasks, but they play a fundamental role. Learning is a pervasive process that takes place at almost every step. In the navigation subtask, the items encountered in the current location are not only broadcast to the network using the chat interface, but also automatically committed to memory. Similarly, information read in chat messages, whether requests for information or availability of items at locations, is also internalized in memory. Parameters controlling access to memory information for future use will play a fundamental role in determining the effectiveness of those learning processes, and whether they need to be supplemented with strategic processes such as rehearsal or deliberate organization. Similarly, the parameters of perceptual processes (scanning and reading latency, typing speed) affect how fast external information can be processed, and thus what the optimal tradeoff is in propagating that information over the network.

The information learned decays with practice and is reinforced with subsequent presentations. The resulting activation in memory plays an important role, not only determining which facts can be retrieved at any given point in time, but also which facts take precedence over others. For instance, when attempting to retrieve the location of the current target item, it is possible that multiple such locations are known. One possible strategy would be to retrieve all such locations, then determine which is the closest. However, some of that information might also be obsolete, since some of those items might have been harvested since the information was first com-

municated. Instead, the model simply retrieves the most active of the competing facts, relying on the property that the most recent one is also likely to be the most active, and also corresponds to the item that is least likely to have been harvested. We admit that in some subjects, this basic heuristic could, conceivably, give way to more complex metacognitive strategies, which might control attentional resources and the time allocated to the tasks.

The exploration of activation decay revealed, empirically, the influence memory has on the performance of the team. We show that too little or too much decay will impede targeted communication. Memory decay in the individual provides contextualization and is beneficial to the team. That said, we do not claim that memory decay is the only way to prioritize recent information. More explicit strategies such as mental timestamps associated with pieces of information are conceivable; yet they would be complex and acquired. Our thesis of innate cognitive mechanisms that are optimized towards cooperation within social networks predicts the pattern of results.

Conclusion

The Geo Game allows us to observe performance and communication efficiency in communities of human participants. The game is neither limited to a single group objective (as in joint problem-solving), nor do humans act as individual, adversarial agents. We have presented a cognitive model describing data in which teamwork benefits from communication policies that make use of active information filtering by humans. Human Geo Game players do very well with a strategy of carefully targeting information, which perhaps is a natural or acquired communication maxim.

A relatively simple, high-level cognitive model can account for the attentional trade-offs involved. We demonstrate the influence of individual memory on task success in communication-based tasks. Counter to intuition, better memory does not help humans communicate or perform the task. The benefits of perfect memory are overcome by the lack of contextualization: agents fail to account for a rapidly changing environment, as it presents itself with the Geo Game. This first version of the model and the early Geo Game experiment are a step towards an understanding of the mechanisms involved in the interaction of individual cognition, communication and resulting group phenomena.

Acknowledgements

We would like to thank Jerry Vinokurov, Antonio Juarez for their programming work, and Katia Sycara, Michael Lewis, Paul Scerri for their collaboration in the design of the Geo Game platform. This work was funded by the Air Force Office of Scientific Research (MURI grant FA95500810356).

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford, UK: Oxford University Press.

Bavelas, A. (1950). Communication patterns in task-oriented groups. *J. Acoustical Soc. America*, 22, 725-730.

Bhattacharyya, S., & Ohlsson, S. (2010). Social creativity as a function of agent cognition and network properties: A computer model. *Social Networks*, 32(4), 263-278.

Cannon-Bowers, J. A., & Salas, E. (1997). Teamwork competencies: The interaction of team member knowledge, skills, and attitudes. In J. H. F. O'Neil (Ed.), *Workforce readiness: Competencies and assessment* (p. 151-174). Mahwah, NJ: Lawrence Erlbaum.

Guettzow, H., & Simon, H. A. (1955). The impact of certain communication nets upon organization and performance in task-oriented groups. *Management Sciences*(233-250).

Judd, S., Kearns, M., & Vorobeychik, Y. (2010). Behavioral dynamics and influence in networked coloring and consensus. *PNAS*, 107(34).

Kennedy, J. (1999). Small worlds and mega-minds: effects of neighborhood topology on particle swarm performance. In *Proc. evolutionary computation*.

Klimt, B., & Yang, Y. (2004). The enron corpus: A new dataset for email classification research. In J.-F. Boulicaut, F. Esposito, F. Giannotti, & D. Pedreschi (Eds.), *Machine learning: Ecmi 2004* (Vol. 3201, p. 217-226). Springer Berlin / Heidelberg.

Leavitt, H. (1951). Some effects of certain communication patterns on group performance. *Abnormal and Social Psychology*, 46, 38-50.

Lebiere, C. (2005). Constrained functionality: Application of the act-r cognitive architecture to the ambr modeling comparison. In K. A. Gluck & R. Pew (Eds.), *Modeling human behavior with integrated cognitive architectures*. Mahwah, NJ: Erlbaum.

Oaksford, M., & Chater, N. (Eds.). (1999). *Rational models of cognition*. Oxford University Press.

Pickard, G., Rahwan, I., Pan, W., Cebrián, M., Crane, R., Madan, A., et al. (2010). Time critical social mobilization: The DARPA network challenge winning strategy. *CoRR*, *abs/1008.3172*.

Reitter, D., & Lebiere, C. (2010a). Accountable modeling in ACT-UP, a scalable, rapid-prototyping ACT-R implementation. In *Proceedings of the 10th international conference on cognitive modeling (iccm)* (p. 199-204). Philadelphia, PA.

Reitter, D., & Lebiere, C. (2010b). Did social networks shape language evolution? A multi-agent cognitive simulation. In *Proc. cognitive modeling and computational linguistics workshop (cmcl)*. Uppsala, Sweden: Association for Computational Linguistics.

Reitter, D., & Lebiere, C. (2011). How groups develop a specialized domain vocabulary: A cognitive multi-agent model. *Cognitive Systems Research*, 12(2), 175-185.

Reitter, D., Sycara, K., Lebiere, C., Vinokurov, Y., Juarez, A., & Lewis, M. (2011). How teams benefit from communication policies: information flow in human peer-to-peer networks. In *Proc. behavior representation in modeling and simulation (brims)*.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature*, 393(6684), 440-442.