Defining Factors of Interest for Large-scale Socio-cognitive Simulations

Ryan Kaulakis, Changkun Zhao, Jonathan H. Morgan, Jeremiah W. Hiam, Joseph P. Sanford, Frank E. Ritter
(rmk216@psu.edu, cuz111@psu.edu, jhm5001@psu.edu, jwh189@psu.edu, jps314@psu.edu, frank.ritter@psu.edu)
The College of Information Sciences and Technology
The Pennsylvania State University

Abstract

We examine how several environmental factors influence cognition and the emergence of social networks in artificial societies. Using a large-scale socio-cognitive simulation (VIPER), we generated social networks consisting of 20, 40, and 60 agents. We tested the impact that environmental factors such as population size, map configuration, and run time, particularly on the formation of social ties in memory. We analyzed 1,080 ego-nets across 27 conditions, measuring the number of links and average degree of each network. While all these factors influenced these network measures, our results suggest that population size has the largest influence. In addition, we examined what impact activation values and retention parameters have on the construction of social networks in memory, finding that a shift in activation values resulted in a loss of links, thus indicating a cognitive foundation for Dunbar's Number (Dunbar, 1998) for the maximum number of social ties stored in memory.

Keywords: ACT-R; socio-cognitive network; memory; network formation

Introduction

We seek to show two things here: first, that several commonly used parameters in network science have an impact and need to be reported, but which are currently under-reported; and, secondly, that cognitive architectures used in network science will have different results than intelligent agents lacking cognitive plausibility. So, in this paper, we examine three common environmental factors that influence the construction of social networks: population size, environmental configuration, and run duration. Later work uses these results to examine how these factors influence specific network relations such as Dunbar's Number (Dunbar, 1998) in (Zhao, Kaulakis, Morgan, Hiam, Sanford, et al., 2012; Zhao, Kaulakis, Morgan, Hiam, & Ritter, 2012).

This work is motivated by a desire to better understand how socio-cognitive processes influence the development of persistent patterns of relations, represented in this paper as network topologies. By socio-cognitive processes, we refer to both those cognitive resources and mechanisms necessary to create and sustain social ties, as well as those group-level factors known to moderate human decision-making as modeled by (Morgan, Morgan, & Ritter, 2010). We, also examine memory retention because it seems foundational to understanding the construction and maintenance of social networks within cognitive architectures.

We refer to these networks as socio-cognitive networks. We differentiate socio-cognitive networks from other social networks by their means of interaction, or the degree to which the networks structure, modes of communication, and resources derive from an external medium (e.g., cell phone networks or web-based friend networks). We believe a deeper understanding of the constraints imposed by memory decay on network formation is important because such an understanding may allow us to refine our predictions regarding the likely structures and capabilities of socio-cognitive networks. In particular, we believe an agent-based approach can allow us to deepen our understanding of the relationship between the carrying capacity (in this case an agents ability to recall its friends network) of a networks nodes and its topology.

To explore these questions, we introduce a set of cognitive models and experiments that vary across three factors. The outputs of this model are ideal networks (whole networks representing the total number of agent interactions that occurred within a single run) and ego-nets (declarative representations of the agents friends network). For any one run, there is then a single ideal network and as many ego-nets (networks from egocentric points of view) as there are agents. We tested the effects of environmental conditions on the construction of the agents ego-nets by comparing multiple ideal networks with their related ego-nets. We also analyzed to what extent these factors (coupled with memory constraints) influenced the constructed networks, measured by differences in the number of links and average node degree.

Our model is unusual in that we model social processes using a full cognitive architecture (ACT-R). To our knowledge, (Carley, 1991, 1992; Carley & Newell, 1994) were the first to implement social network models using a cognitive architecture (Plural Soar) to study organizations. More recently, (Gonzalez, Lerch, & Lebiere, 2003), (Lebiere, Gonzalez, Dutt, & Warwick, 2009), (Reitter & Lebiere, 2010), and (Juvina, Lebiere, Martin, & Gonzalez, 2011) have used cognitive architectures to model human decision making in collaborative tasks. While our work builds upon these efforts, we focus here on the formation of social networks. Further, while cognitive architectures bring great power, they are computationally expensive. We thus, must address both questions of utility and theoretical subsumption. In other words, what do these architectures uniquely bring to the table that we need? In the next section, we will address these questions by orienting ourselves with respect to past work in social modeling before describing our model more fully.

Computational Social Models

We briefly describe a cross-section of related work. We will move from simple non-cognitive models to models that center on cognition and cognitive modeling. All of these approaches are agent-based, or refer to predictions based on individual decision-making.

In many models, the role of actors is often described using closed-form mathematical formulas that offer parsimony, but often prove difficult when modeling complex stochastic systems. In contrast, computational models typically try to model the effects of a wider range interrelated factors using more complex but bounded actors (Axelrod & Hammond, 2003). These effects often include both environmental, as well as cognitive factors.

Drawing from work in environmental and social psychology (Kraut, Fussell, Brennan, & Siegel, 2002; Allen, 1977), we found that actor proximity fundamentally influenced the evolution of network topologies by determining the interaction frequencies of actors across the network. Allen (Allen, 1977) demonstrated that the probability of two people communicating in an environment could be defined by a decreasing hyperbolic function of the distance between them. After a certain distance, the probability that two people will communicate decreases rapidly, making link formation unlikely. We thus choose to focus on factors that directly affect agent proximity or inter-agent distance: population size, run duration, and map configuration.

We expect that larger populations acting over longer periods in fully connected rooms will result in the richest declarative network structures. We also expect that layouts that afford greater distances will result in interaction networks that consist of multiple components, leading to smaller ego-nets. We also expect that map configurations characterized by nexus points will exhibit behaviors similar to the water-cooler effect. We, however, are less certain where we might see thresholds in network formation, where for instance population growth no longer has an effect or run time is no longer relevant. We expect these thresholds will provide us a better understanding of the cognitive dimension behind the shifts in group behavior associated with changes in group size.

Drawing from previous work in cognitive science (Simon, 1984; Prietula & Carley, 2001), we were interested in how bounded rationality influenced network formation by constraining network construction in memory. We believe that these constraints will result in interesting and sometimes unexpected aggregate behaviors that may provide insights for future efforts in multi-level modeling. In particular, we believe the concept of nodal carrying capacity, or the number of agents any one agent can retain in its social declarative representation, may be helpful for predicting the capabilities and structure of a network of interest. To that end, we examine how shifts in activation values and retention parameters, as well as differences in environmental factors contribute to the consolidation and retention of social ties in memory. This concept is similar to those of Dunbar (Dunbar, 1998) and the Bernard-Killworth median (McCarty, Killworth, Bernard, Johnsen, & Shelley, 2001). Dunbar's number arises from the limitations associated with the neocortex. Because maintaining a stable relationships requires repeated memory activations in the human neocortex to identify not only one-on-one relationships but also third party relationships (i.e., the knowledge that my friend is also friends with other actors who I, in some senses, monitor), the cognitive load associated with maintaining this set of relationships in memory rises exponentially as group size increases (Dunbar, 1998, p.63). Based on retrospective empirical studies, (Dunbar, 1998, p.65–78) argues that this ratio between cognitive load and group size underlies the small-world effect observed by Milgram and others.

Nodal carrying capacity: The effect of agents memory and space

Thus far, we have categorized the factors that influence the construction of social networks into two groups: factors that influence interaction frequency and factors that influence retention. In this section, we will discuss each of these categories, and give a general prediction on how the factors associated with each effect the construction and retention of cognitive ego-nets.

Interaction Frequency

We examine three factors that influence the interaction frequencies associated with a given social network. These factors include: population size, duration of contact (run time), and environment map connectivity (defined by a grid ratio, or the number of links over the total number of links possible in a similar grid).

Population size: We suspect that population density is the most influential factor governing interaction frequency. Because we are comparing map configurations consisting of the same number of rooms, we manipulate population density by testing three different population sizes (20, 40, and 60 agents).

Length of simulation: The time period that agents interact directly influences the structure of the simulated social network because more time allows for more interaction opportunities, making it more likely that agents will establish a stable network. Consequently, determining the run times necessary for a network to reach a stable state under a given set of conditions is important for accurately representing the formation of a group of interest. We use total degree with respect to time as a measure of network stability. Modeling memory decay, however, seems essential for determining a meaningful notion of network stability. Otherwise, we suspect simulated networks will tend to achieve arbitrary and inflated levels of connectivity ending with complete connectivity at infinite time.

Environment configuration: We believe that the configuration of the rooms of the environment influences the structure of the simulated social network. We measure the relative connectivity of our three map configurations by defining its grid ratio. The grid ratio is the ratio of the number of edges over the total number of edges possible for a rectangular grid containing the same number of rooms.

We tested three map configurations. The first configuration is a full 5x5 grid with grid ratio 1.0. We expect this environment will result in relatively high connectivity. The second configuration (shown in Figure 1a) is a two-hallway configuration with grid ratio 0.6. This configuration should lead to low connectivity due to the large distances between agents. The third configuration (shown in Figure 1b) has a central area with grid ratio 0.75. We believe this central meeting point will lead to agent behavior between 1 and 2.

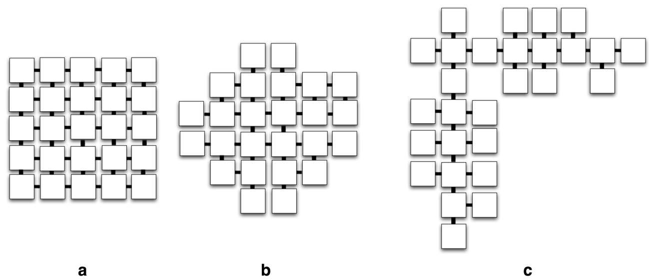

Figure 1: a: Full Grid, b: Central Area, c: Hallway

Cognitive factors

To better simulate the construction of social networks, it is necessary to consider the behavior patterns of agents at the cognitive level. In this paper, we particularly focus on memory decay. We examine memorys effect on tie formation by using Anderson's activation theory (Anderson et al., 2004) to model the construction of social knowledge in declarative memory. In our model, the number of friends depends on the number and size of active long-term memory chunks associated with the agents social relationships. The number of active memory chunks can be influenced by several factors, including initial memory activation, retrieval threshold, memory decay speed, time of retrieval, and practice time.

Experiment Environment

To model multi-agent social behavior using cognitive architectures, we constructed a simulation environment called VIPER.

The VIPER Server

VIPER is a lightweight, extensible, text-based simulation environment. It uses the telnet protocol to handle plain text communication between agents and the simulation environment and forces a separation of environment and agent. VIPER records agent behaviors within the environment using a logging system that provides a detailed list of every action taken by all agents chronologically ordered to the tenth of a second. VIPER resolves events in either real or accelerated time: the networks speed and frequency of communication is determined by its component agents with no queue of events being enforced within the environment.

Within VIPER, agents are situated in maps of interconnected rooms. In each room, the agents can see and communicate locally. Agents can walk freely around the rooms, and can interact with objects in their environment.

To connect ACT-R to VIPER, we implemented a VIPER client, the Telnet Agent Wrapper (TAWA) for ACT-R, in Common Lisp. It handles everything from logging in, waiting for synchronization, logging, halting, and writing results to CSV files automatically. Additionally, it provides functions to examine the environment, speak, listen, move, and otherwise control virtual bodies in VIPER.

When an ACT-R model is wrapped by TAWA, any execution of model code is delayed until a privileged administrator actor inside of VIPER signals synchronization. Additionally, any error conditions are caught by TAWA and used to return standard UNIX error codes instead of dropping into the debugger. For example, a successful run returns zero to the parent process, while any error (e.g., network errors like the server being unreachable) causes a non-zero return value. Returning error codes allows automated error checking in large scale experiments.

All of our experiments were conducted on a 2GHz eight-core Linux 2.6.31 under Ubuntu 11.04 server with 8GB of RAM, with SBCL 1.0.52 as our Lisp run-time. We use ACT-R 6 described in (Anderson et al., 2004).

Synchronization

Because memory decay and networks are strongly temporal, we paid special attention to time. When designing our experiments, we developed VIPER and TAWA to use a synchronization process. During an experimental run, TAWA delays the evaluation of the model code until synchronization, this means that no Agent experiences time before the synchronization signal. Further, all ACT-R models are set to run in real-time, and for the full time, using the run-full-time function from standard ACT-R with :real-time enabled. All agents run for the same amount of real time, so they all halt after the same perceived period after starting. Thus, the total time experienced is the same for all agents.

Scalability

Early benchmarks showed that ACT-R processes took up about 80MB per process. Because we had only 8GB of RAM, we would only have been able to run about 100 processes on a single machine before swapping. To reduce the per-process footprint, a number of optimizations were implemented. Basic space reductions were achieved by using the DECLARE operator, as well as by compiling all libraries, removing the debugger, and saving the whole system (sans the ACT-R agent model) as a system image. This reduced our per-process memory footprint somewhat, but they were not the biggest contributions towards memory usage reduction.

In SBCL version 1.0.52, the –merge-core-pages flag was recently added. This flag enables Kernel SamePage Merging (Arcangeli, Eidus, & Wright, 2009) under recent versions of Linux. This optimization flags shared areas of memory as merge-able unless modified. Because a significant percentage of our agents were replicated, we could reduce the per-process memory footprint as low as 8MB per process in benchmarks. Thus, the only things that increase the size of this footprint during run-time are changes to memory done by the ACT-R agent model. Benchmarks have shown reasonable performance with 700 test ACT-R agents.

ACT-R Agents

We built an ACT-R model to conduct a *random walk* in VIPER. The model contains two declarative memory types: a goal type containing current location, remaining steps, total friends counts; and a friend type to store a friends name. The model consists of four basic components with 9 productions. First, the *walking* component selects an available direction randomly, and moves itself through VIPER. Second, the *waiting* component utilizes temporal buffer of ACT-R to wait two real-time seconds when the agent comes into a new room, to simulate how long it might take a normal person to enter a room. Third, the *checking* component (with 3 productions) checks if the current room is empty. Finally, the *memorizing* component (with 3 productions) will check its declarative memory first to check if it can recall the fact that the agent it has encountered is a friend. If it is not, the model will create a new friend chunk, and store the encountered agents name in it using the imaginal buffer.

Experiment and Results

In this section, we discuss the results of our analysis of 27 system logs and 1,080 ego-net logs.

Experiment Parameters

In the first experiment, we have 27 runs of our simulation to test three environmental factors: population size, run time, and map configuration. To manipulate room configuration, we use a 5x5 grid map and two other maps shown in Figure 1. All parameters are shown in the Table 1.

Table 1: Experiment Parameters

Factors	**Testing Values**
Population Size	20, 40, 60
Running time (seconds)	125, 250, 500
Map Configuration (grid ratio)	Full Grid (1.00), Central (0.75), Hall (0.60)

To test cognitive parameter, we make all agents output the activation value of each friend chunks as a ego-net log file, which contains friend name and location of the last meeting. In ACT-R, the activation value represents the memory retention of an object or an event. With the activation value of each relation chunk, we could easily find the weight of each friend tie in memory.

Results

As noted in section 4, our simulation generates two types of network data: log data extracted from Viper directly, and egocentric data stored in the agents declarative memory. In this section, we will present samples of the data and some related network measures.

Ideal Network Figure 2 shows a sample interaction-network. The nodes in the figure are the agents in the simulation (20 in total); the links in the figure are unweighted, only representing the co-occurrence of two agents in a given room during the run.

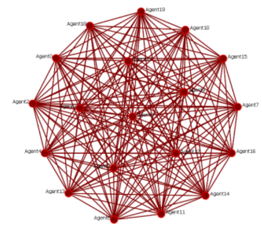

Figure 2: The network of log data for 20 agents running 125s in the Hallway.

Table 2 shows the measure comparison between 7 runs. We found that the population size and running time influence the network measures, reflected in Average Node Degree and Degree Centrality. We find that the Total Links and Average Degree increase when population size increases. Total Links and Average Standard Degree also increase when doubling run-time.

Table 2: Measurements of seven sample runs, grouped by variable (italicized); N is Population Size, T is Run-time is in seconds, R is Grid Ratio, Links is Total Links, Degree is Average Node Degree, Centrality is Degree Centrality.

N	T	R	Links	Degree	Centrality
20	125	0.60	324	0.436	0.627
40	125	0.60	1410	0.784	0.227
60	125	0.60	3801	0.370	0.651
20	*125*	0.60	324	0.436	0.627
20	*250*	0.60	357	0.503	0.490
20	*500*	0.60	360	0.555	0.481
20	125	*0.60*	324	0.436	0.627
20	125	*0.75*	354	0.507	0.546
20	125	*1.00*	360	0.569	0.476

Table 2 shows that all three factors influence measures of the network. As each factor increases, the total links and average node degree increase correspondingly, but, the degree centrality decreases. Of the three factors, the population size has the most influence on these measures.

Egocentric Network The egocentric network is extracted from the declarative memory of each ACT-R agent. Figure 3 shows Agent-0s egocentric point-of-view. All relationships are weighted based on the agents activation values. Figure 4 shows a combined egocentric network of all the agents egocentric networks. We considered either a tie of $n_i \rightarrow n_j$ or $n_j \rightarrow n_i$ sufficient for inclusion in the combined network. We, thus, expect that the activation values found in each ego-net for ties $n_i \rightarrow n_j$ and $n_j \rightarrow n_i$ to differ. Nevertheless, when we compare Figure 2 and Figure 4, they have similar structures.

When we increase the activation threshold, we do see that links are eliminated as the minimal activation value neces-

Figure 3: Egocentric network of and individual agent (name = Agent-0, population size = 20, running time = 125, configuration = Hallway)

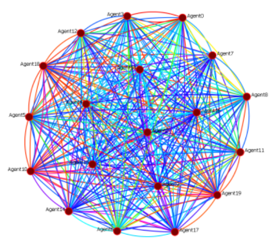

Figure 4: Egocentric network of all agents (population size = 20, running time = 125, configuration = Hallway)

sary for inclusion in the network increases, i.e. we see a layered network consisting of a strong set of relationships at its core and more casual ones at the periphery. Figure 5 shows the changes to the network when we increase the activation threshold. It appears that Figure 5 part c has a smaller but strong network because the memory retention between the existing nodes are relatively high.

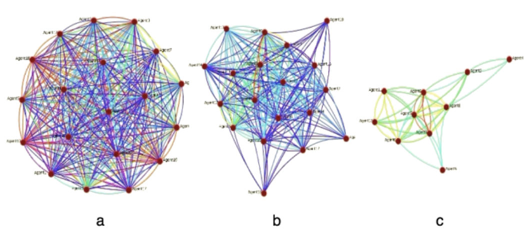

Figure 5: Egocentric networks with applied filters of a: -3.5, b: 0.0, c: 1.0

Example Pairwise Comparison

To examine the effect of cognitive constraints, we compared the activation levels of declarative memory chunks representing other agents in memory.

In a simple case (from a test run), Agent-0 remembers Agent-1 at about 0.2466, while Agent-1 remembers Agent-0 at about 0.5342. The higher the activation value, the better one agent remembers another, and here, we can see that Agent-0 does not remember Agent-1 as well as Agent-1 remembers Agent-0. This asymmetry can be due to Agent-0 not really paying attention to Agent-1, despite the simple fact that they had met.

In fact, looking at the ground truth from the system logs, Agent-0 and Agent-1 met each other 32 times in 125 seconds. Thus, despite meeting about once every four seconds, Agent-0 did not remember Agent-1 as well as it might have. As a matter of fact, when a minimal memory threshold of 0.0 was applied (n.b. activation values can be negative), Agent-0 had no memory of Agent-1. A threshold analysis, which we discuss next, clearly shows the asymmetry of ties in memory (in this case, "A knows B but B does not know A").

Based on this kind of asymmetric behavior, we suggest that the main impact of this kind of analysis on social network research will be in the realm of asymmetric relations. Ultimately, we expect to see that there is a distinction that needs to be made between knowledge of a relationship (among many possible relationships) and the attentional importance of that relationship.

Example Activation Cutoff

In Figure 5, we find another interesting story. When we set the threshold at 0.0, Agent-10 loses many of its connections between other agents. When the threshold is equal to -3.5, Agent-10 has relations with all other agents (19 relations in total); but Agent-10 agent only has 9 relations left after the threshold was applied. After checking the log file, we found that the Agent-10 had multiple interactions (at least 13 times) with every agent. Most of these interactions, however, took place at the beginning of the simulation after this initial period Agent-10 was isolated at end of a hallway. The activation value between Agent10 and the other agents continued to decay to values frequently below 0.0, with values ranging between -1.07 and 0.11. This case directly illustrates that not only does the frequency of interaction influence memory activation or the ties strength in memory but also the time and sequence of interactions.

Discussion and Conclusions

These results show how several common effects of cognition often influence network growth and shape. In this study, we created a multi-agent social network simulation that provides us a very flexible platform to examine factors that influence the development and maintenance of social networks in memory. Based on a review of the literature, we identified and modeled ecological (population size, map configuration, and run time) and cognitive factors, with the cognitive factors represented using memory activation parameters.

We conducted 27 runs of our simulation to test our model. The results indicate that all three factors influence the networks total links, average degree, and degree centrality. As each factor increases, the total links and average node degree increase correspondingly, but the as expected, the networks degree centrality decreases. Of these three factors, the population size has the most influence on the network measures. The effect of running time is not as significant as we expected, and shows plateauing after the 250s run. The large running time also weakens the effect of map configuration, because it provides agents enough time to travel around the whole map.

Taking advantage of the ACT-R memory mechanism, we were able to model the ego-nets of 1080 agents, and combine those networks to compare the agents declarative representation of their friends network to the ideal network (or ground truth) generated from VIPERs log. We found the structure of ideal network and merged egocentric network to be similar (see 2 and 4). However, this similarity ends when we apply thresholds to the link weights in the merged egocentric networks (see Figure 5), where higher thresholds result in less connected networks that bear little resemblance to the ideal network. Semantically, this difference shows that memory limits how much of the ideal network an agent can remember well.

Future avenues of work will build upon some of the more interesting issues. First, we would do further analysis of normalized activation thresholds to see if these reliably effect either the network topology of the interaction network, or the topology of the agents declarative representation. Second, we would run more agents, because our test systems were kept deliberately small. Finally, we would analyze the effects of cognition on network measures analogous to Dunbar's Number.

Acknowledgments

DTRA (HDTRA1-09-1-0054) supported this work. Jaeyhon Paik provided technical support and Geoffery P. Morgan useful comments. CASOS developed ORA which proved invaluable as an analysis tool.

References

Allen, T. (1977). Managing the flow of technology: Technology transfer and the dissemination of technological information with the R&D organization.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060.

Arcangeli, A., Eidus, I., & Wright, C. (2009). Increasing memory density by using KSM. In *Proceedings of the linux symposium* (p. 1928).

Axelrod, R., & Hammond, R. (2003). The evolution of ethnocentric behavior. In *Midwest political science convention, chicago, il.*

Carley, K. (1991). A theory of group stability. *American Sociological Review*, 331–354.

Carley, K. (1992). Organizational learning and personnel turnover. *Organization Science*, 20–46.

Carley, K., & Newell, A. (1994). The nature of the social agent. *Journal of mathematical sociology*, *19*(4), 221–262.

Dunbar, R. (1998). *Grooming, gossip, and the evolution of language*. Harvard Univ Pr.

Gonzalez, C., Lerch, J., & Lebiere, C. (2003). Instance-based learning in dynamic decision making. *Cognitive Science*, *27*(4), 591–635.

Juvina, I., Lebiere, C., Martin, J., & Gonzalez, C. (2011). Intergroup prisoners dilemma with intragroup power dynamics. *Games*, *2*(1), 21–51.

Kraut, R., Fussell, S., Brennan, S., & Siegel, J. (2002). Understanding effects of proximity on collaboration: Implications for technologies to support remote collaborative work. *Distributed work*, 137–162.

Lebiere, C., Gonzalez, C., Dutt, V., & Warwick, W. (2009). Predicting cognitive performance in openended dynamic tasks a modeling comparison challenge. In *Proceedings of the 9th international conference on cognitive modeling, manchester, uk.*

McCarty, C., Killworth, P., Bernard, H., Johnsen, E., & Shelley, G. (2001). Comparing two methods for estimating network size. *Human Organization*, *60*(1), 28–39.

Morgan, J., Morgan, G., & Ritter, F. (2010). A preliminary model of participation for small groups. *Computational & Mathematical Organization Theory*, 1–25.

Prietula, M., & Carley, K. (2001). Boundedly rational and emotional agents: Cooperation, trust and rumor. *Trust and Deception in Virtual Societies. Edited by Cristiano Castelfranchi and Yao-Hua Tan. Kluwer Academic Publishers*, 169–194.

Reitter, D., & Lebiere, C. (2010). A cognitive model of spatial path-planning. *Computational & Mathematical Organization Theory*, 1–26.

Simon, H. A. (1984). *Models of bounded rationality: Economic analysis and public policy*. The MIT Press.

Zhao, C., Kaulakis, R., Morgan, J. H., Hiam, J. W., & Ritter, F. E. (2012). Modeling a cognitively limited network in an agent-based simulation. In *Proceedings of the 34th conference of the cognitive science society.* Sapporo, Japan. (SUBMITTED)

Zhao, C., Kaulakis, R., Morgan, J. H., Hiam, J. W., Sanford, J. P., Ritter, F. E., et al. (2012). Socio-cognitive networks: Modeling the effects of space and memory on generative social structures. In *Proceedings of the 21st conference on behavior representation in modeling and simulation.* Amelia Island, FL. (ACCEPTED)