Correspondence

Validating Agent Interactions in Construct Against Empirical Communication Networks Using the Calibrated Grounding Technique

Craig Schreiber and Kathleen M. Carley, *Senior Member, IEEE*

***Abstract*—Validating a computational model is important as it establishes that the model has met its intended purpose of representing the system under study. In this paper, we perform a validation study on Construct, a multiagent network model for the coevolution of agents and the sociocultural environments that they inhabit. In particular, we focus on validating agent interactions produced by the model against empirical communication networks collected in real-world organizations. Validation is performed using our novel calibrated grounding technique. Results show that Construct can produce valid agent interactions. The benefits and implications of the study are discussed.**

***Index Terms*—Multiagent network model, organization theory, simulation, social networks, validation.**

I. Introduction

We develop computational models of organizational systems with the intended purpose of representing the real-world phenomenon [1]–[3]. Computational models are a means to deal with the complex, dynamic, and nonlinear functioning of real-world organizations, which often cannot be adequately reduced to an analytic model [2], [4], [5]. From these computational models, we can obtain many useful ends such as predictive emulation [6], [7], normative analysis [8]–[10], and theory development [11]–[14]. However, how confident can we be in the results obtained from the model given that the model is only an approximation? In other words, how well does the model represent the real-world phenomenon?

Validation is the process of determining how well a computational model matches the organizational system that it represents [3], [8], [15], [16]. Validation of agent-based models, like the one used in this study, has been difficult at best due to many methodological challenges [17], [18]. We use validation, first and foremost, to obtain a level of credibility in the model which gives us confidence in the results that we obtain. However, there are many other benefits to the process of validation as it aids in scientific accumulation through an understanding of the strengths as well as the boundaries of the model [14]. Such knowledge can provide phenomenal understanding, guidance for application decisions, and directions for future research.

Manuscript received January 14, 2010; revised January 15, 2011; accepted May 17, 2011. Date of publication September 12, 2012; date of current version December 12, 2012. This work was supported in part by the National Aeronautics and Space Administration under Grant NAG-2-1569 and Contract NNA04AA14C, by the Army Research Laboratory under Grant ARL/CTA DAAD19-01-2-0009, by the National Science Foundation (NSF) under Grant 0452487, by the NSF Knowledge and Distributed Intelligence under Grant IIS-9980109, by the NSF Integrative Graduate Education and Research Traineeship under Grant 9972762 for research and training in the Center for Computational Analysis of Social and Organizational Systems (CASOS), and by CASOS at Carnegie Mellon University. This paper was recommended by Associate Editor W. Pedrycz.

C. Schreiber is with the Charles M. Snipes School of Business and the Broyhill Institute for Business Leadership, Lenoir-Rhyne University, Hickory, NC 28601 USA (e-mail: craigschreiber.acad@gmail.com).

K. M. Carley is with the Center for Computational Analysis of Social and Organizational Systems, Institute for Software Research, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: kathleen.carley@cs.cmu.edu).

Digital Object Identifier 10.1109/TSMCA.2012.2192104

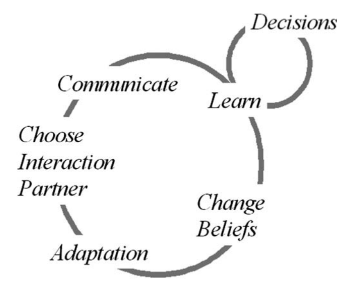

Fig. 1. Construct action cycle.

The goal of this paper is to perform an external validation for Construct, a multiagent network model for the coevolution of agents and the complex sociocultural environments that they inhabit. In particular, we focus on the ability of Construct to produce agent interactions that are representative of communication networks in real-world organizations. To validate Construct, we apply the calibrated grounding technique to nine empirical data sets of real-world organizations. These nine data sets provide a solid foundation for testing Construct.

This paper is organized as follows. First, we briefly describe the Construct model and the importance of validating it against communication networks. Then, in Section III, we describe the data sets and explain the calibrated grounding technique used to validate Construct. Finally, we present and discuss our results and conclude by considering the benefits and implications of the study.

II. Construct

Construct is a multiagent network model for the coevolution of agents and the sociocultural environments that they inhabit [13], [14], [19]–[26]. Agents in the model go through an action cycle; see Fig. 1. In this cycle, agents choose interaction partners, communicate, learn knowledge, change their beliefs about the world, and adapt their networks based on their updated understanding. At the end of the cycle, agents perform tasks based on their current understanding. Outcome measures such as knowledge diffusion, performance accuracy, and consensus are collected.

Agent interactions are the basic foundation on which the output measures depend. These interactions figure prominently on what each agent learns, and agent learning determines the values of the outcome measures. Therefore, the first step in validating Construct is to obtain a reasonable degree of equivalence between agent interactions and real-world communication networks. Such a validation will generate confidence in the model's ability to represent a real-world organization and to originate sound ends, such as theory development. If agent interactions in Construct reasonably represent real-world interactions of a group, then we can say that outcomes of the model could reasonably occur in the real-world phenomenon.

In Construct, there are two main types of core variables that influence agent interactions: organizational representation parameters and interaction processes. Organizational representations are particular networks that characterize the organization. Typical organizational representations that are collected are the task network, knowledge

		Task			
		1	*2*	*3*	*4*
Agent	*1*	0	1	0	1
	2	1	0	0	1
	3	0	1	1	0
	4	0	0	1	0

Fig. 2. Illustrative task network.

		Agent			
		1	*2*	*3*	*4*
Agent	*1*	.00	.18	.05	.09
	2	.10	.00	.06	.01
	3	.04	.08	.00	.19
	4	.11	.02	.15	.00

Fig. 3. Illustrative probability-of-interaction matrix for relative similarity.

network, and cognitive network. We collected the task network, the knowledge network, and two different cognitive networks: cognitive knowledge and cognitive task. Thus, there are four total organizational representations used in this research. These networks are used because it is believed that they influence communication networks. In other words, people will communicate with one another based on the task assignments that people have, the knowledge that people possess, or the perceptions that people hold.

The task network is "who does what" task in the organization. It represents the task assignments that people have. All networks, including the knowledge and cognitive networks described in the following, are represented as matrices in the model. Fig. 2 shows an illustrative task network, where "1" indicates that the agent performs that task and "0" indicates otherwise. An illustrative knowledge network and a cognitive knowledge network for comparison are shown in Figs. 6 and 7 provided in the Appendix.

The knowledge network is "who knows what" in the organization. Knowledge is defined into categories that are relevant to that particular organization. For example, if we were collecting data on an organizational simulation group, we may have knowledge categories such as software development, hardware, organization theory, and statistics. The knowledge network then is simply who possesses what level of expertise in each category.

Cognitive networks are the perception of each person as to "who knows what" or "who does what" in the organization. In other words, they are each person's perception of the knowledge or task network. Therefore, the cognitive knowledge network is a collection of networks, with each network representing a distinct agent's perception of "who knows what." Likewise, the cognitive task network is a collection of networks, with each network representing a distinct agent's perception of "who does what."

The cognitive networks are collected because people interact or make choices based on their perceptions of the world, which vary by person. Therefore, it seems that a cognitive perception network would be a good representation on which one can base interactions. This is obvious for the cognitive knowledge representation, but we can also study network forms of organization [27] where tasks are often not well defined and ambiguous. In this case, task network perception will vary by person, and this could figure significantly into agent interactions. One difficulty with cognitive networks is the time commitment required to reach out to each person in the organization and to get their perception of every other person that they know. This encumbering time commitment often prohibits the actual collection of cognitive network data in real-world organizations. These difficulties are compounded by the way in which such networks change over time.

Interaction processes are based on well-known social processes of human interaction. There are two basic interaction processes in Construct, relative similarity and relative expertise. Relative similarity is based on homophily [28], the finding that people tend to interact with those similar to themselves. Arguments supporting homophily include trust, comfort, communicative ease, and access.

In Construct, agents who are acting on relative similarity will interact more with agents who are similar to themselves than with agents who are dissimilar. For example, agents who work on mostly similar tasks will tend to interact more often than agents who have mostly dissimilar tasks.

Equation (1) shows the calculation for the probability that agent i will interact with agent j based on relative similarity of knowledge. K refers to a set of knowledge bits, and S refers to an agent knowing a specific bit of knowledge within that set. For example, if Sik is binary, then $Sik = 1$ if agent i has knowledge of knowledge bit k; otherwise, $Sik = 0$. The higher the number of knowledge bits that i and j have in common, the higher their relative similarity. This is computed for each communication direction, i.e., to and from, for every pair of agents

$$RSij = \frac{\sum_{k=0}^{K}(Sik * Sjk)}{\sum_{j=0}^{I}\sum_{k=0}^{K}(Sik * Sjk)}. \tag{1}$$

Relative expertise is based on expertise seeking [29]. Arguments supporting expertise seeking are knowledge integration and a need for specialized or nonredundant knowledge.

Construct agents who are acting with relative expertise will interact more with agents who are dissimilar to themselves than with agents who are similar. For example, agents will tend to seek out and interact with other agents who have different knowledge than their own rather than agents who have overlapping knowledge.

Equation (2) shows the calculation for the probability that agent i will interact with agent j based on relative expertise of knowledge. X refers to a specific bit of knowledge that j knows which i does not. The higher the number of knowledge bits that j knows which i does not, the higher j's relative expertise is to i. This is computed for each communication direction, i.e., to and from, for every pair of agents

$$REij = \frac{\sum_{k=0}^{K}(Xjk)}{\sum_{j=0}^{I}\sum_{k=0}^{K}(Xjk)}. \tag{2}$$

The output for either relative similarity or relative expertise is a matrix consisting of the interaction probabilities for every pair of agents. Fig. 3 shows an illustrative probability-of-interaction matrix for relative similarity. Fig. 3 shows a partial matrix and thus does not show all the probabilities. In a full matrix, the probabilities associated with agent i would sum to one over all other agents. Notice that the relative probabilities between pairs of agents are not symmetric. Communication can be initiated from one pairwise direction more often than another due to relative asymmetries.

III. Methodology

A. Data Sets

Construct was tested against nine real-world data sets as shown in Table I.[1] These nine data sets are independent and represent a variety of contexts and group sizes. Group size ranges from a small group of 9 members to a large group of 206 members. The data for

[1]The software company data set courtesy of Ashworth and Carley [30].

TABLE I
REAL-WORLD DATA SETS

Name	Group Size	Organizational Representation
Aeronautics A	13	K, CK
Aeronautics B	10	K, CK
Professional Association	11	K, CK
University	13	K, CK
Consulting Firm	9	K, CK
Concurrent Engineering Team	19	K, CT
Software Company	16	K, T
Battle Command Group A	206	T
Battle Command Group B	156	T

(K=knowledge network, CK = cognitive knowledge network, T=task network, CT = cognitive tast network

People	1	2	3	4
1	0	1	1	1
2	1	0	0	1
3	0	1	0	1
4	1	1	1	0

Fig. 4. Illustrative communication network.

each network were collected during field studies via observations and questionnaires. Different organizational representations were collected at each organization because the data were collected for research purposes other than validation. Consequently, we do not have all four organizational representations for each organization. Unfortunately, validation is often stressed but rarely funded to afford higher levels of effort; thus, one must make do with secondary data.

Although the data were collected for other purposes, we have a reasonable number of cases for each representation, except the cognitive task network. In this study, we test how well these different representations, in general, can be used as input into Construct for producing valid agent interactions. Communication networks from each organization were collected in addition to the organizational representations, and we compared the communication networks against networks derived from Construct's agent interactions for validation.

The communication networks used for statistical comparison are of a similar format to the empirical task, knowledge, and cognitive networks that were previously described. Fig. 4 shows an illustrative communication network. The values correspond to whether or not communication has occurred, where "1" indicates that communication has occurred and "0" indicates otherwise. We can also collect and use data that represent the frequency of communications.

B. Calibrated Grounding

Calibrated grounding is a novel technique that we developed for validating multiagent network models against empirical communication networks; see Fig. 5. This technique is a combination of approaches, specifically initialization grounding and internal calibration. Initialization grounds the model through the use of empirical data as input. This provides the model with a representation of the real-world organization. In this study, internal calibration is the varying of organizational representation and interaction process to obtain different agent interactions. We then test the different agent interactions to see which were validated and to what strength in order to calibrate the model.

The detailed process of calibrated grounding has several steps. First, an organizational representation is inputted into Construct. This representation is a parameter that is believed to influence real-world interactions. Representations include the knowledge network, task network, or cognitive network, all previously described.

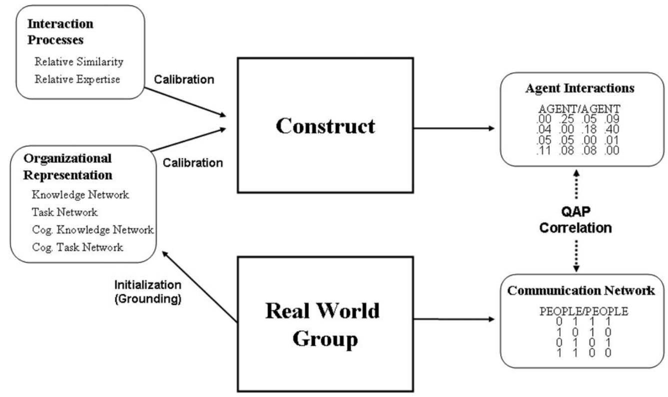

Fig. 5. Calibrated grounding technique uses initialization grounding and internal model calibration to validate agent interactions with real-world communication networks.

Second, an interaction process that drives agent interactions is chosen. The interaction process choices are relative similarity and relative expertise, both previously described.

Third, Construct is run to produce initial agent interactions. Agent interactions are represented by the probability-of-interaction matrix; see Fig. 3. As previously described, the probability-of-interaction matrix is calculated by applying an interaction process equation to an organizational representation.

Fourth, the initial agent interaction network is correlated to the real-world communication network using the quadratic assignment procedure (QAP) [31]. Relational data violate the independence assumption; therefore, usual parametric methods of comparison are not appropriate. QAP correlation is a nonparametric procedure that uses a permutation test to determine significance. First, the Pearson correlation coefficient is calculated by comparing the corresponding cells of the two matrices: the agent interactions in the probability-of-interaction matrix and the real-world communication network (see Fig. 5). Then, the matching rows and columns of one matrix are permuted randomly, and the correlation coefficient is recalculated. For example, when rows 2 and 4 are permuted, then columns 2 and 4 are permuted as well. This preserves the dependences that exist in the relational data. These permutations are repeated thousands of times (2500 in our case), thereby producing a distribution of correlation coefficients. Significance is determined by the location of the original (nonpermuted) coefficient within the distribution. Significance in this study was determined at the standard 0.05 level.

Validation occurs when a significant positive correlation exists between the simulated and real-world networks. A significant positive correlation would indicate that there is significant overlap between the relations of the simulated and real-world networks. This validation process was repeated for every unique combination of organizational representation and interaction process. It should be noted that the value of the correlation coefficient does not offer any indication of degree of correlation. QAP correlation is nonlinear in this regard. Therefore, it cannot be interpreted that one correlation coefficient is stronger than another. The only indicator of significance is whether or not a particular model passes the statistical test at the 0.05 level of significance. We can say that levels of significance greater than 0.05 do provide stronger results.

The focus of this study is on the ability of Construct to produce a valid initial state of interactions. We are not averaging interactions over time periods, and there are no stochastic processes from the model that are involved in this particular study. Additionally, there were no internal mechanism adjustments such as tweaking of equations to get the model to fit the data.

Calibrated grounding as used in this study is a form of distributional equivalence. We use distributional equivalence because we want the

TABLE II
CONSTRUCT VALIDATION RESULTS

				Organizational Representation			
				K	CK	T	CT
Organization	*Aeronautics A*	Interaction Parameter	RS	0.148	0.442**		
			RE	0.173	0.180**		
	Aeronautics B	Interaction Parameter	RS	0.048	0.167		
			RE	-0.028	0.116		
	Professional Association	Interaction Parameter	RS	0.038	0.178*		
			RE	0.056	-0.005		
	University	Interaction Parameter	RS	0.170	0.242**		
			RE	0.195*	0.085		
	Consulting Firm	Interaction Parameter	RS	0.206*	0.427**		
			RE	0.070	0.013		
	Concurrent Engineering	Interaction Parameter	RS	0.174**			0.334**
			RE	0.057			0.199**
	Software Development	Interaction Parameter	RS	0.059		0.179	
			RE	0.023		0.141*	
	BCG A	Interaction Parameter	RS			0.069**	
			RE			0.027**	
	BCG B	Interaction Parameter	RS			0.097**	
			RE			0.038**	
	Number of organizations that validated against an organizational representation/ Total number of organizations			3/7	4/5	3/3	1/1

Pearson correlation coefficients
** -0.05 level of significance, ** -0.01 level of significance*
(RS=relative similarity, RE=relative expertise)
(K= knowledge network, CK = cognitive knowledge network, T=task network, CT = cognitive task network)

distributions of simulated and real-world interactions to match within a reasonable statistical tolerance. Also, our particular approach is a form of parameter and process matching, not tuning as in turning knobs and incrementally changing weights to curve fit. The parameter is varied with a set of fixed organizational representations, and the interaction processes are varied by two distinctions. In other words, there is not a continuous range of values by which the model could be tuned. We are validating a process of how interactions occur. Parameter and process matching is a more meaningful way of validating through calibration [32].

IV. RESULTS AND DISCUSSION

A. Organizational Representation and Validated Agent Interactions

Table II shows the results of the validation. Each organization in Table II is a separate analysis, but they are shown together for analytic clarity. The names of the interaction parameters were abbreviated to fit in the table and are as follows: relative similarity (RS) and relative expertise (RE). The gray-shaded cells indicate the lack of that organizational representation for testing.

There are some interesting results in terms of organizational representations and validated interactions in Construct. The knowledge network is not well validated; significant correlations were found in only three out of seven organizations. This demonstrates a lack of robustness across the various organizational contexts. This was a surprising result as we expected the knowledge network to be validated better. We compared various network characteristics such as group size, knowledge network density, and communication network density to help explain why a few organizations are validated but most are not. Unfortunately, we did not find anything to indicate that network characteristics contributed to this result. We speculate that this result is due to varying levels of granularity with which knowledge was measured.

Also, the condensing of "who knows what" into one matrix that applies to every agent could be oversimplistic. The assumption is that the people in the organization have a shared understanding of the knowledge network and that this representation could be used to derive interactions similar to the real-world interactions. This assumption seemed reasonable given the small size of the organizations for which this representation is employed. This assumption is most likely incorrect.

In contrast, the cognitive knowledge network is well validated; significant correlations were found in four out of five organizations, and this demonstrates a robust representation across the various organizational contexts. The robustness of this representation highlights the fact that people will interact based on their own unique perception of the world. This result indicates that the use of perceptual representation is better than a condensed measure representing a shared understanding, even for small organizations where shared understanding and transactive memory are more easily obtained. We were not able to draw a conclusion why the Aeronautics B organization did not validate the

TABLE III
VALIDATED ORGANIZATIONAL REPRESENTATIONS BY INTERACTION PROCESS

	Organizational Representation			
Interaction Process	Knowledge Network	Cognitive Knowledge Network	Cognitive Task Network	Task Network
Relative Similarity	2/7	4/5	1/1	2/3
Relative Expertise	1/7	1/5	1/1	3/3
Total	3/14	5/10	2/2	5/6

cognitive knowledge representation. This organization did have high task interdependences as compared to the others, with exception of the Professional Association. However, a definitive conclusion cannot be drawn based on fact that the Professional Association was validated. Also, a task network is not available for the Aeronautics B organization to test if the high task interdependences contributed to the results.

The one test of the cognitive task network validated but one test is certainly not conclusive. The cognitive task network is an organizational representation worth additional testing.

On the other hand, the task network does demonstrate extremely good promise. It was validated in all three organizations and in five out of the six tests across interaction parameters. We still need more task networks to test for robustness, but this is an excellent beginning with consistent results. Moreover, the results complement empirical studies which show that task networks significantly influence communication networks in organizations [29].

There is one rather interesting result in terms of interaction processes and validated interactions in Construct. For four organizations, the relative similarity and relative expertise interaction processes were both validated with a particular organizational representation (see Aeronautics A, Concurrent Engineering, BCG A, and BCG B in Table II). What this result likely means is that these particular organizational communication networks have elements of both interacting on similarity, e.g., "I talk to those people who perform the same function as me," and interacting on expertise seeking, e.g., "I talk to those people who have specialized knowledge which I need."

This result is similar in nature to the empirical result obtained by Cross *et al.* [33] where the use of homophily ties or cross-boundary (knowledge-seeking) ties depended on the type of information being sought. For those organizations that validate only one interaction process, we can say with confidence that these organizations interact predominantly with that particular process.

B. Validated Organizational Representations by Interaction Process

Table III shows the number of times that an organizational representation is validated by an interaction process. The totals in the table reinforce the robustness of the cognitive knowledge network, the promise of the task network, and the mixed results for the knowledge network.

A very interesting result occurs with the interaction processes in the cognitive knowledge network. Although robust in total times validated, this robustness is only reflected with the relative similarity interaction process. Relative similarity validated four out of five times as compared to one out of five times for relative expertise. This result suggests that people who interact frequently may believe that they have very similar knowledge or comparable levels of expertise, whether or not this is actually true. It is quite plausible that people who interact with known experts will overrate their own expertise in that area, although this may not necessarily be an intentional act. Regardless, the perception of homophily in the cognitive knowledge network is a strong one.

Another interesting result occurs in the task network. Both interaction processes tend to be validated consistently. Again, the task network results relate to the work by Cross *et al.* which show the influence of the task network on communications [29] and the existence of both homophily and knowledge seeking in communication networks [33]. The reason that we believe that the task network captures both interaction processes where the cognitive knowledge network typically does not is due to the task network being a more objective measure which is not impacted by individual perceptions. The large tendency toward homophilous perceptions in the cognitive knowledge network washes out the effects of knowledge-seeking needs. In other words, cognitive perceptions combine the need for knowledge seeking into the similarity process, thereby distorting the distinction between these two processes.

The task network, on the other hand, preserves the process distinction and the various interactional needs. People need to interact with those who are performing similar functions or tasks as well as interact with others who are performing very different functions or tasks. This is particularly relevant in knowledge-intensive environments. The task network captures the formal task interdependences which contribute to the formation of homophilous groups and communication boundaries.

V. BENEFITS AND IMPLICATIONS

A. Benefits

We performed a validation study on Construct which shows that the model has an ability to represent real-world communication networks. Agent interactions were validated in 15 out of the 32 tests. Of those 15 tests, 11 were particularly strong at the 0.01 level of significance. The overall results of this validation are considerably high for a multiagent model. This demonstrates that computational models need to be at least moderately complex to valuably support operational applications.

These validation results allow us to have reasonable confidence in Construct's representation of real-world communication networks within organizations. We can also have some confidence in the ends obtained from the model when using the validated input. Aside from model credibility, this effort produced other useful insights which highlight additional benefits and implications.

We gained considerable knowledge about the use of various organizational representations in Construct. The knowledge network had mixed results when used to validate Construct and does not display robustness. The cognitive task network had only one organizational representation, and no conclusions could be drawn. The use of cognitive task networks as an organizational representation may prove to be fruitful for organizations that have ambiguous tasks, such as network organizations in high-velocity environments. Additional testing of the cognitive task network is a future research direction.

Validation is robust when using cognitive knowledge networks as organizational representations. Task networks also demonstrate a promise for validation robustness. Future research should conduct more tests using the task network representation to reveal how robust it is. The conclusions that we draw from this study are that the cognitive knowledge network and task network are influential on organizational communications and that these organizational representations offer the best opportunity for producing validated interactions in Construct.

A particularly noteworthy finding is the capability of the task network to preserve the distinction between homophily and knowledge-seeking interactions. This is important if we want to assess the interaction processes in an organization. Through validation using the task network representation, we can evaluate if both processes are present or if only one is present. If both are present, then additional research can be conducted to find out the probabilities associated with each process or even under what conditions they are more likely to occur. Construct can be set accordingly with this additional information. For example, Construct can be set so that agents interact on relative similarity for a certain percentage of the time and on relative expertise for the remaining percentage. This would strengthen the model's ability in representing the real-world organization. If only one interaction process is present, then we can have reasonable confidence that this is the predominate process.

In contrast, the results suggest that the cognitive knowledge network typically does not preserve the distinction between homophily and knowledge-seeking interactions. The perceptions of individuals tend toward a homophily bias which blurs the distinction between these processes and lumps them together under relative similarity. Future research should test the cognitive knowledge and task networks within the same organizations to definitely conclude the difference between these two organizational representations in terms of preserving the interaction process distinction. It is our conjecture that the cognitive knowledge network will maintain the homophily bias while the task network will preserve the interaction process distinction. If this is the case, then the cognitive knowledge network will be an appropriate organizational representation to use when overall interactions without regard to interaction process are the goal of the study.

The cognitive knowledge network and task network also have different data collection implications associated with them. Cognitive networks are very intrusive to collect. Every person must be surveyed or interviewed for extended periods of time in order to obtain their perception about everyone that they know in the organization. Although we have several of these networks collected, the general consensus in organizations is to disincline such time-intensive and intrusive data collection. In fact, the most recent data sets in this study all do not have cognitive knowledge networks.

Alternatively, task networks can be collected by examining organizational documents or by surveying or interviewing a few people. This process is much less intrusive and not intensive in overall time commitment to the organization. The less intrusive the process, the easier it will be to get organizational commitment to collect data. Automated electronic collection would be even more ideal as this is the least intrusive and allows for changes in the organizational network to be collected quickly, almost real time. This can then result in timely analysis and feedback to the organization. Therefore, if the task network is as robust as the cognitive knowledge network, then it would seem to be the ideal organizational representation to pursue.

Another major benefit of this research study is the development of the calibrated grounding technique. This novel technique provides the ability to validate agent interactions in multiagent network models against real-world communication networks. The development and use of this technique are particularly significant given the inappropriateness of usual parametric methods.

Lastly, this validation study is a step toward an applied model. With considerable evidence showing the validity of the agent interactions in Construct, we can now turn to longitudinal predictive result verification. This verification becomes a matter of relational equivalence and can be performed for network change and for outcome measures.

Validating network change as a result of coevolution will require comparing the various model and empirical networks over time. Calibrated grounding can be adapted to perform relational equivalence by running Construct experiments over several time periods and then validating the simulated and real-world networks at various points longitudinally.

Validating network change and outcome measures concurrently using relational equivalence could also be performed and would be considered a more stringent test. Validating both network change and outcome measures for a specific organization should entail the following: 1) performing calibrated grounding as distributional equivalence; 2) running Construct experiments over several time periods; 3) performing calibrated grounding as relational equivalence tests for network change; and 4) validating the outcome measures against real-world outcomes. Validating outcome measures can be performed through a variety of correlation and regression techniques.

Although result verification sounds easy according to the aforementioned description, it is anything but and entails a few hard problems. First, time periods have no relation to real clock time such as minutes, hours, days, weeks, or months. Determining how long to run a virtual experiment and which time periods in the simulation output correspond to data points in the empirical data is a difficult process. Moreover, even if validation occurs for one organization, it is unlikely that the length of run or time period distance between comparisons will correspond across organizations. This is most likely true even for comparisons across organizational representations for the same organization.

Second, the process of result verification will likely highlight changes that need to be made to the model internal functioning. This would send the process back to code changes that can be expensive and time consuming. Once the code changes are complete, then result verification would entail additional data sets that may or may not be on hand. If they are not on hand, then more data collection is required.

All this adds up to an extremely difficult and expensive process which unfortunately often prohibits such valuable validation. Our study demonstrates what validation has to offer, but we also hope to convey that there is immense value yet to be tapped. The mining of this value will depend upon the appropriate funding and effort being afforded.

B. Limitations

An obvious limitation in this validation study is the number of larger groups represented. There are only two large groups, and both are of the same context. We need to increase our validations of large groups and test across various contexts before we can begin to make conclusions about Construct's ability to validate interactions of sizable groups. Also, there is an uneven distribution for organizational representation by group size. This warrants caution about concluding the usefulness of organizational representations to produce valid interactions across group sizes. In defense of the lack of cognitive network representation in large groups, the expense of collecting such data prohibits its presence.

One important note about validation of computational models, in general, is that validation is only a matter of degree [2]. Models are only approximate representations of the complex systems under study. There cannot be any objective proof of a model's validity [34]. We can only have confidence that the model is a reasonable representation of the system [35].

APPENDIX

	Knowledge			
Agent	*1*	*2*	*3*	*4*
1	1	0	0	1
2	1	1	0	1
3	1	0	0	1
4	0	1	1	0

Fig. 6. Illustrative knowledge network. "1" indicates that the agent possesses the knowledge in that category, and "0" indicates otherwise.

Agent 1's perception

Agent	Knowledge *1*	*2*	*3*	*4*
1	1	1	0	1
2	1	0	0	1
3	0	1	0	1
4	0	0	1	0

Agent 2's perception

Agent	Knowledge *1*	*2*	*3*	*4*
1	1	0	0	1
2	1	1	0	1
3	0	0	0	1
4	0	1	1	0

Agent 3's perception

Agent	Knowledge *1*	*2*	*3*	*4*
1	1	1	1	1
2	1	1	0	0
3	1	0	0	1
4	0	0	1	0

Agent 4's perception

Agent	Knowledge *1*	*2*	*3*	*4*
1	1	1	0	1
2	0	1	0	1
3	1	1	0	1
4	0	1	1	0

Fig. 7. Illustrative cognitive knowledge network. Each matrix represents an individual agent's perception of the knowledge network. "1" indicates the perception that an agent possesses the knowledge in that category, and "0" indicates otherwise. The cognitive knowledge network is the collection of individual perceptual matrices.

Acknowledgment

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 National Aeronautics and Space Administration, the Army Research Laboratory, the National Science Foundation, or the U.S. Government.

References

[1] H. Hattori, Y. Nakajima, and T. Ishida, "Learning from humans: Agent modeling with individual human behaviors," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 41, no. 1, pp. 1–9, Jan. 2011.

[2] A. M. Law and W. D. Kelton, *Simulation Modeling and Analysis*, 3rd ed. Boston, MA: McGraw-Hill, 2000.

[3] J. G. Turnley, "Validation issues in computational social simulations," presented at the 3rd Lake Arrowhead Conf. Human Complex Systems, Lake Arrowhead, CA, 1995.

[4] J.-S. Lee, C. Schreiber, and K. M. Carley, "Validation and docking of two organizational simulations using dynamic team performance data," *Comput. Math. Org. Theory*, forthcoming.

[5] K. M. Carley and L. Gasser, "Computational organization theory," in *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, G. Weiss, Ed. Cambridge, MA: MIT Press, 1999, pp. 299–330.

[6] Y. Jin and R. E. Levitt, "The virtual design team: A computational model of project organizations," *Comput. Math. Org. Theory*, vol. 2, no. 3, pp. 171–196, 1996.

[7] R. E. Levitt, G. P. Cohen, J. C. Kunz, C. I. Nass, T. Christiansen, and Y. Jin, "The 'Virtual Design Team': Simulating how organization structure and information processing tools affect team performance," in *Computational Organization Theory*, K. M. Carley and M. J. Prietula, Eds. Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, pp. 1–18.

[8] J. Poropudas and K. Virtanen, "Game-theoretic validation and analysis of air combat simulation models," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 40, no. 5, pp. 1057–1070, Sep. 2010.

[9] H. H. Baligh, R. M. Burton, and B. Obel, "Devising expert systems in organizational theory: The organizational consultant," in *Organization, Management and Expert Systems*, M. Masuch, Ed. Berlin, Germany: Walter De Gruyter, 1990, pp. 35–57.

[10] H. H. Baligh, R. M. Burton, and B. Obel, "Validating the organizational consultant on the fly," in *Computational Organization Theory*, K. M. Carley and M. J. Prietula, Eds. Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, pp. 179–194.

[11] W. Zhu, "Consensus of multiagent systems with switching jointly reachable interconnection and time delays," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 42, no. 2, pp. 348–358, Mar. 2012.

[12] D. A. Pelta and R. R. Yager, "Decision strategies in mediated multiagent negotiations: An optimization approach," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 40, no. 3, pp. 635–640, May 2010.

[13] K. M. Carley, "Group stability: A socio-cognitive approach," in *Advances in Group Processes: Theory & Research*, vol. VII, E. Lawler, B. Markovsky, C. Ridgeway, and H. Walker, Eds. Greenwhich, CT: JAI Press, 1990, pp. 1–44.

[14] C. Schreiber and K. M. Carley, "Going beyond the data: Empirical validation leading to grounded theory," *Comput. Math. Org. Theory*, vol. 10, no. 2, pp. 155–164, Jul. 2004.

[15] O. Balci, "Verification, validation and accreditation," in *Proc. Winter Simul. Conf.*, Washington, DC, 1998, pp. 41–48.

[16] J. P. C. Kleijnen, "Verification and validation of simulation models," *Eur. J. Oper. Res.*, vol. 82, no. 1, pp. 145–162, 1995.

[17] O. Gurcan, O. Dikenelli, and C. Bernon, "Towards a generic testing framework for agent-based simulation models," in *Proc. Fed. Conf. Comput. Sci. Inf. Syst.*, Szczecin, Poland, 2011, pp. 635–642.

[18] P. Windrum, G. Fagiolo, and A. Moneta, "Empirical validation of agent-based models: Alternatives and prospects," *J. Artif. Soc. Social Simul.*, vol. 10, no. 2, p. 8, Mar. 2007.

[19] K. M. Carley, "Growing up: The development and acquisition of social knowledge," in *The Self-Society Dynamic: Cognition, Emotion, and Action*, H. J. and C. P. , Eds. Cambridge, MA: Cambridge Univ. Press, 1991, pp. 72–105.

[20] K. M. Carley, "On the evolution of social and organizational networks," in *Special Issue of Research in the Sociology of Organizations on Networks in and Around Organizations*, vol. 16, S. B. Andrews and D. Knoke, Eds. Greenwhich, CT: JAI Press, 1999, pp. 3–30.

[21] K. M. Carley and V. Hill, "Structural change and learning within organizations," in *Dynamics of Organizations: Computational Modeling and Organization Theories*, A. Lomi and E. R. Larsen, Eds. Menlo Park, CA: MIT Press/AAAI, 2001, ch. 2, pp. 63–92.

[22] C. Schreiber and K. M. Carley, "Key personnel: Identification and assessment of turnover risk," presented at the North American Association Computational Social Organizational Science Conf., Pittsburgh, PA, 2004.

[23] C. Schreiber, S. Singh, and K. M. Carley, "Construct—A Multi-agent Network Model for the Co-evolution of Agents and Socio-cultural Environments," Carnegie Mellon Univ., School of Comput. Sci., Inst. Softw. Res., Int., Tech. Rep. CMU-ISRI-04-109, 2004.

[24] B. R. Hirshman, K. M. Carley, and M. J. Kowalchuck, "Specifying Agents in Construct," Carnegie Mellon Univ., School Comput. Sci., Inst. Softw. Res., Tech. Rep. CMU-ISRI-07-107, 2007.

[25] B. R. Hirshman, K. M. Carley, and M. J. Kowalchuck, "Loading Networks in Construct," Carnegie Mellon Univ., School Comput. Sci., Inst. Softw. Res., Tech. Rep. CMU-ISRI-07-116, 2007.

[26] B. R. Hirshman, M. Martin, and K. M. Carley, "Modeling Information Access in Construct," Carnegie Mellon Univ., School Comput. Sci., Inst. Softw. Res., Tech. Rep. CMU-ISR-08-115, 2008.

[27] N. Nohria, "Is a network perspective a useful way of studying organizations?" in *Networks and Organizations: Structure, Form and Action*, N. Nohria and R. G. Eccles, Eds. Boston, MA: Harvard Business School Press, 1992, pp. 1–22.

[28] P. Lazarsfeld and R. Merton, "Friendship as a social process: A substantive and methodological analysis," in *Freedom and Control in Modern Society*, M. Berger, T. Abel, and C. Page, Eds. New York: Octagon, 1978, pp. 18–66.

[29] R. Cross, R. Rice, and A. Parker, "Information seeking in social context: Structural influences and receipt of informational benefits," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 31, no. 4, pp. 438–448, Nov. 2001.

[30] M. Ashworth and K. M. Carley, "Who you know vs. what you know: The impact of social position and knowledge on team performance," *J. Math. Sociol.*, vol. 30, no. 1, pp. 43–75, 2006.

[31] D. Krackhardt, "QAP partialling as a test of spuriousness," *Social Netw.*, vol. 9, no. 2, pp. 171–186, Jun. 1987.

[32] K. M. Carley, "Validating Computational Models," Carnegie Mellon Univ., CASOS Rep. CASOS Publ., Pittsburgh, PA, Tech. Rep., 1996.

[33] R. Cross, S. Borgotti, and A. Parker, "Beyond answers: Dimensions of the advice network," *Social Netw.*, vol. 23, no. 3, pp. 215–235, Jul. 2001.

[34] J. W. Forrester, *Industrial Dynamics*. Cambridge, MA: MIT Press, 1961.

[35] M. Greenberger, M. Crenson, and B. Crissey, *Models in the Policy Process*. New York: Russell Sage, 1976.