

Computational organization science: A new frontier

Kathleen M. Carley*

Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213-3890

Synthetic adaptation is the process whereby any entity composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent. Because of synthetic adaptation, organizations, like the agents of which they are composed, are inherently computational. We can gain insight into the behavior of groups, organizations, and societies by using multiagent computational models composed of collections of intelligent adaptive artificial agents. CONSTRUCT-O and ORGAHEAD are examples of such models whose value for social, organizational, and policy analysis lies in the fact that they combine a network (social and knowledge approach) with a multiagent approach to effect more realistic behavior. The results from a series of virtual experiments using these models are examined to illustrate the power of this approach for social, organizational, and policy analysis.

Social, organizational and policy analysts have long recognized that groups, organizations, institutions, and the societies in which they are embedded are complex systems. For example, the thought of an organization conjures up a series of interlocked images of people working on tasks, managers coordinating personnel, alliances, mergers, practices for hiring, layoffs and technology adoption, strategic and operational behavior, and so forth. In these complex systems, part of the complexity emerges from the fact that humans are a key component. Humans are the canonical adaptive intelligent agent, complete with computational (i.e., bounded information processing) capabilities. Because of synthetic adaptation, groups, organizations, institutions, and societies, like the humans of which they are comprised, are intelligent, adaptive, and computational. Specifically, synthetic adaptation is the process whereby any entity composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent. The very nature of these systems suggests that multiagent models composed of intelligent adaptive and computational agents may be valuable in predicting, understanding, and interpreting behavior. Further, because computational models enable the user to run large scale virtual experiments without altering the actual socio-technical world in which we live, such models become an important tool for organizational and policy analysis.

Computational analysis using multiagent models is beginning to significantly impact the way groups, organizations, institutions, and societies are managed and the way individual, organizational, market, technology, and policy decisions are made and evaluated. Computational analysis is reshaping the way we think and reason about complex phenomena where emergence plays a central role. Some of the areas most affected are organizational design and strategy, commerce, organizational learning, information technology impacts, innovation, and diffusion. Multiagent models used in a “what-if” fashion are improving our understanding of how different technologies, decisions, and policies influence the performance, effectiveness, flexibility, and survivability of complex social systems.

Computational organization science is a new perspective on groups, organizations, societies, and policy systems that has emerged in the past decade in response to the need to under-

stand, predict, and manage system level change including, but not limited to, change that is motivated by new technology (1). Computational organization science is a neo-information processing approach to the study of social, organizational, and policy systems that combines social science, computer science, and network analysis. All agents, synthetic or not, human or artificial, are viewed as knowledge consumers and producers embedded in an ecology of networks linking entities such as agents, resources, knowledge, organizations, and tasks. Learning and change practices, such as innovation, turnover, etc., are viewed as linking the micro (individual agent behavior) to the macro (multiagent behavior). Hence, computational organization scientists see emergence and adaptivity as resulting from these change processes being enacted to maintain, alter, or transform networks. Multiagent models are used to enable theory building (2).

To see the revolutionary sweep of this idea, it is important to recognize that traditional organizational theory was concerned with manufacturing firms, traditional policy analysis with rational actors, and traditional social science with collectives of independent actors. Yet, we know that many modern organizations trade in knowledge, not just goods, in producing services and information and not just physical devices. In today's world, information processing, communication, and knowledge management have become key. Changes in computational power, telecommunications, and information processing are affecting when, where, and how work is done. Further, we know that humans and the systems in which they are embedded are not rational—in the traditional economic sense; rather, they satisfy, are cognitively limited, act emotionally, and so on. These bounds on human behavior will not be eliminated by technology. Finally, we know that actors interact and have affiliations that form a network constraining and enabling behavior. Technology, although it might affect the size and density of this network, will not eliminate the role of such networks. Computational organization science builds on these fundamental findings.

In this paper, this new scientific paradigm and its power for fundamentally changing the way we approach and understand social, organizational, and policy issues are illustrated. A series of computational multiagent models will be used to generate these illustrations. In particular, we will focus on results from the ORGAHEAD (3, 4) and CONSTRUCT-O (5, 6) models.

Computational Socio-Knowledge Perspective

Recent advances in social networks, cognitive science, computer science, and organization theory have led to a new perspective on groups, organizations, institutions, societies, and policy systems that takes into account the computational nature of organizations and the underlying social and knowledge networks. Socio-technical systems such as organizations are viewed complex, computational and adaptive (1). Such systems are

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling,” held October 4–6, 2001, at the Arnold and Mabel Beckman Center of the National Academies of Science and Engineering in Irvine, CA.

*E-mail: kathleen.carley@cmu.edu.

composite synthetic agents composed of other complex, computational, and adaptive agents constrained and enabled by their position in a social and knowledge web of affiliations linking agents, knowledge, and tasks. The use of computational models enables the generation of meaningful insights and the evaluation of policies and technologies. This result is true, whether the system is a collection of people, artificial agents, synthetic agents, or some combination of these. Because different agents will have different cognitive and communicative abilities, different capabilities for acquiring, processing, storing, retrieving, and communicating information, and the behavior of the socio-technical systems will depend on the capabilities of the constituent agents. The capabilities of the agents will define what types of “social” behaviors emerge (7).

The basis for this argument is a “bodiless” view of knowledge. Knowledge is a complex structure of ideational kernels and the connections among them and is not directly linked to the physical body of any one agent. The relations among information can occur within the mind of an agent or between agents, such as “shared ideas” or the “I know that you know” linkages. Knowledge exists within and between individual agents and groups of agents. Synthetic adaptation derives from the structural nature of intelligence. Change in this knowledge is synthetic adaptation. Synthetic adaptation occurs in a composite agent as the member agents change or the connections within and among them change. Because of synthetic adaptation, groups and organizations are more than the simple aggregate of the constituent personnel.

Learning leads to bits of information being learned and forgotten (nodes added and dropped) and connections among information being learned and forgotten (relations added and dropped). By extrapolation, any agent that can add or drop nodes or relations in the knowledge space is learning. Any entity composed of intelligent, adaptive, and computational agents is also an intelligent, adaptive, and computational agent. In this sense, organizations are intelligent, adaptive computational agents in whom learning and knowledge are distributed (2, 8–10). Consequently, organizations can take action distinct from member agents. This synthesis process involves many complex nonlinear processes and is not a simple aggregation across member agents.

Agents, resources, knowledge, tasks, and organizations are connected by and embedded in an ecology of networks (6, 11, 12). We can think of a set of core corporate entities—agents, knowledge, resources, tasks, and organizations—that define a set of networks. The network among agents, the social network, entails all relations by which agents interact, communicate, and exchange goods, services, and information. The network linking agents to bits of information is the knowledge network. This network defines who knows what. As agents learn, as agents interact, these networks change dynamically (13). These networks, when coupled with the change processes, provide a detailed model of the socio-technical system. Aspects of the socio-technical system that have traditionally been examined by noncomputational researchers can now be precisely measured when formulated in this fashion. For example, the authority structure and the communication structure are defined in terms of the interaction network, the culture is defined in terms of the knowledge network, and the potential data in terms of the information network.

Information processing research (14) demonstrates that limits to an agent’s information processing capabilities effect outcomes in socio-technical systems. Taking such limitations into account leads to more accurate prediction of organizational performance (15). For agents, what actions they can take depends on what knowledge and resources they have, what they are assigned to do, and what knowledge is currently salient. What knowledge is salient is in part a function of the agent’s cognitive architecture

(7). Humans’ cognitive architecture makes them at least boundedly rational (16, 17). In fact, they are actually even more limited. Such cognitive limitations create the need for social behavior (7). Social interaction, and more specifically, the structure of the socio-technical system, places a limit on agent rationality and hence behavior (18). Thus, what information is available and salient to the individual is a function of the agent’s position in this set of interlocked networks. Additionally, there is an interaction between knowledge (e.g., training, what agents know, and their information processing capabilities) and structure in effecting organizational performance (19, 20). Consequently, all actions are embedded in an ecology of networks that curtail the creation, use, and acquisition of information. In this sense, it makes more sense to think of agents not as boundedly rational, but as structurally rational.

For individual agents, decision-making is often characterized as search (21); so too, for organizations (22, 23). For organizations, there is a relationship between their organizational architecture and its performance (20). The networks in which agents, individual or synthetic, are embedded constrain and enable their search (8, 24) and limit performance.

The organization’s architecture, despite being this set of networks, is typically characterized in terms such as size, span of control, density of connections among personnel, workload, and so on. If we take the set of all organizations for a particular task, there is a performance surface that characterizes maximum performance achievable, given this set of characteristics.

The organization, as a synthetic agent, is trying to move through this space searching for the structural form that enables higher performance. The design of socio-technical systems thus becomes a strategic exercise in establishing and managing these networks to conduct search in this space (25). Various change agents or processes in the organization, such as the CEO, do the actual search. As structural changes are made, the socio-technical system moves about in this space. For an organization, structural changes include such activities as hiring new personnel, eliminating divisions, redesigning the organization so that who is reporting to whom changes, and retasking individuals so that who is doing what changes. This search may be done locally—looking only at organizations or agents who are nearby in the performance space and so have similar characteristics—or globally—looking at the vast panorama of what architectural forms are possible. This search may be done in an exploratory manner—looking at totally new forms—or in such a way to exploit known competencies—for example, moving along only a single dimension.

All of the search strategies that humans are capable of are also available to organizations. This finding does not mean that organizations can search and change in any way they wish. Rather, the adaptation of agents (human, artificial, or organization) is constrained, often by forces beyond the agents immediate control. Laws, financial constraints, technology, etc. limit where in the performance space organizations can move. For example, the tenure system effectively curtails the search through the performance space for universities. In general, the organization will have a set of change strategies that will determine how it moves through this space.

Changes in the knowledge network are experiential learning. Changes in the social, assignment and organizational networks are structural learning. Moving an organization from one location to another may eliminate knowledge, or render it less useful. Experiential learning is fragile. For example, if the organization changes its architecture, the agents within the organization may find that their experience is not valuable. The organization can move to a position that structurally enables a higher potential maximum performance but in which, at least initially, the organization performs worse. As an example, installing new technologies and new databases in organizations, which can be

thought of as altering the structure of the organization, can temporarily degrade performance because the knowledge people had about how to operate the old technology is no longer valid (26). Experiential learning and structural learning clash when the lessons of experiences become invalid as the agent is given a totally new task to do, or a new group of people to interact with. Experiential and structural learning are only two of the change processes in these systems. The networks affect the rate of learning, thus determining how fast the organization attains its full potential.

This depiction of socio-technical systems represents a new scientific paradigm. This perspective, drawn from a large number of empirical studies, is consistent with arguments of distributed cognition, transactive memory, and the social construction of knowledge. Viewing socio-technical systems in this way makes it clear that multiagent models can be meaningfully used in the development and explication of organization theory.

Socio-Technical Systems as Networked Multiagent Structures

We can embody these ideas in computational models. Two such models are ORGAHEAD and CONSTRUCT-O. Both of these simulations are multiagent models in which the agents are structurally rational, task oriented, and embedded in networks of interaction and knowledge. Agents are engaged in doing a classification choice task. The distinction between the models is that ORGAHEAD focuses on organizational adaptation and all communication occurs through the formal organizational networks; whereas CONSTRUCT-O focuses on information diffusion and the impact of technology at the informal or inter-organizational level. Both models have been previously described in the literature in detail so here only a brief description is provided.

ORGAHEAD is a computational framework for examining the behavior of individuals and organizations as they learn, interact, and perform tasks (3, 4). Issues relating to organizational learning, adaptation, design, restructuring, training, agent ability, and so on can be examined by using ORGAHEAD. ORGAHEAD has been successfully used to examine and predict strategic change in small organizations of 5–45 agents. Components of ORGAHEAD have been validated at various levels.

ORGAHEAD characterizes organizations at two levels: operational and strategic. At the operational level, the organization is characterized as a multiagent system in which the agents can learn and have a position in the organization's architecture that constrains whom they communicate/report to, what resources they have access to, and what subtasks they are assigned to. At the strategic level, the organization is characterized as a purposive actor; i.e., there is a CEO or executive committee that tries to forecast the future and decides how to change the organization's structure to meet anticipated changes in the task environment or to improve general performance. As a result of this characterization, ORGAHEAD is comprised of a series of inter-linked models: the operational agent model, the CEO model, and the task model, all of which are linked by the organization's architecture. The architecture is defined in terms of the number of personnel, resources/subtasks and the connections among these entities.

CONSTRUCT-O is a computational framework for examining the coevolution of social structure and culture under different technological and demographic (5, 6, 8, 27, 28) conditions. Issues relating to information diffusion, impact of new telecommunication technologies, friendship formation, alliances, public goods, and so on can be examined by using CONSTRUCT-O. CONSTRUCT-O has been successfully used to examine and predict the flow of information and the assimilation of groups. Components of CONSTRUCT-O have been validated at various levels.




Fig. 1. Performance profiles over time.

Agents in CONSTRUCT-O learn through interaction. They concurrently engage in a communication learning cycle where whom one interacts with changes with what one knows and vice versa. Information technologies, as artificial agents, take part in this cycle. Differences among agents—humans vs., e.g., databases—are made on the basis of their information processing characteristics. For example, humans can initiate interaction, can communicate to one or many others, can forget, and can learn from only one other at a time whereas a database cannot initiate interaction, can communicate to one or many others, cannot forget, and can learn from only one other at a time. Subtle differences in capabilities can have profound consequences (29).

Addressing Organizational and Policy Issues

Multiagent models of complex systems such as ORGAHEAD and CONSTRUCT-O can be used to address a large number of design and policy issues. In a sense, such models are the new decision aids. The decision-maker comes up with a “what if” question. Then, a virtual experiment is run by using the model in which a number of scenarios are examined. Each scenario is run hundreds if not thousands of times, and the results are statistically analyzed. Mean and variance of the underlying distributions produced by the model provide guidance on the likely impact on outcomes of interest to the decision-maker given the scenarios examined. The potential of this approach will be demonstrated by looking at a set of typical findings. These results should be treated as illustrative, with each topic requiring more in-depth analysis.

Can Any Organization Succeed? Using ORGAHEAD, a virtual experiment was done to look at the tendency to industry diversification in terms of performance. A set of 100 organizations were simulated for 40,000 time periods in both stable and changing environments. In the changing environment, there is a dramatic change in the task at time 20,000. Initial architectural forms were chosen randomly. Results indicate that industry diversification is inevitable. Moreover, the final behavior of any one organization is due to organizations locking into the wrong strategies, and failing to respond to major environmental changes, and not because of the specific set of tasks in a particular time period.

In Fig. 1, the behavior of two prototypical organizations faced with the changing environment is shown. Oscillations in performance occur as structural changes result in knowledge losses as personnel are laid off, given new tasks, or placed in new divisions. Clashes between structural and experiential learning increase diversity in architectural form and performance. Maladaptive organizations lock into strategies of change that are

counterproductive (such as oscillating bouts of hiring and firing). In contrast, adaptive organizations lock into strategies that enable continued flexibility (such as repeated tuning through retasking and redesign). However, behavior and so performance can be turned around. Major changes such as changing the CEO and top management team at the time of an environmental shift can be quite effective. As we see in Fig. 1, one organization struggles with various forms altering its performance and happens to be in a high performance state when the environment changes and is able to learn new strategies of change that are better suited to this new environment. In contrast the original top performer has learned how to be very top in the first environment, but is so inflexible that when the environment shifts it cannot learn new appropriate strategies. Organizations that start out as top performers can fail and vice versa as is indicated in this figure. Path dependence, initial conditions, fortune, and response to triggering events all affect long run behavior.

How Do You Design for Success? A well-established finding in organization theory is that there is no one right organizational design. Rather, the optimal design depends on a variety of factors. A key is congruency or match. Organizations whose structure (reporting, task assignment, knowledge network, capabilities network) is matched to the demands of the task (needs network, requirements network) tend to outperform those that are less congruent. What kinds of changes can we expect in congruent organizations that enable them to maintain or improve performance. Will congruent organizations remain congruent? Using ORGAHEAD, 1,000 organizations varying in initial organizational architecture, training, and agent capabilities were simulated for 10,000 time periods (at one task per time period). These organizations were all adapting to a stable environment. Data on performance, strategic changes, agent training, and the initial and final architecture of the organization were recorded.

The relative behavior of the top and bottom performing organizations are examined separately. The top performing organizations or adaptive organizations are defined as the 10% of the 1,000 organizations with the highest average performance in the last 500 tasks whereas the low performance or maladaptive organizations are the 10% with the lowest average performance during this same period. Despite comparable initial organizational architectures and levels of congruence, the adaptive and maladaptive organizations end up with dramatically different architectures. No organization is perfectly congruent. Initial choices lead toward improvement or error. Over time, organizations become increasingly differentiated. Adaptive organizations become larger and less dense whereas maladaptive organizations become smaller and denser.

Policies that inhibit growth may well stifle the long run ability of the organization to be a top performer. Results also indicate that the combination of structural learning (through redesign, retasking, hires, and fires) and individual learning (training and memory), as well as initial conditions, influences performance and ultimate structure. Adaptive organizations, i.e., those that exhibit higher sustained performance, learn to be flexible in whom they assign to what (retasking), grow in size, and employ personnel who are more capable (more memory). Organizations learn to be large if they learn not to redesign (change who is reporting to whom) and if their personnel are not over-trained initially.

What Enables or Inhibits Adaptation? Over time, however, organizations can vary both in the type of structural change in which they engage and the amount of such change. We can think of a sudden increase in the amount of structural change as “shaking” the organization. Examples of shaking would be rapid down-

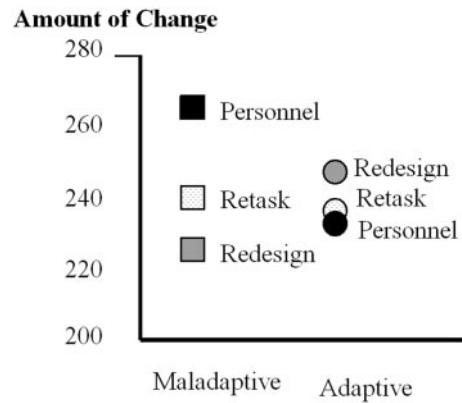


Fig. 2. Adaptive organizations tune rather than shake.

sizing or change in the CEO and top management team. We can think of gradual changes in who is reporting to whom and who is doing what as tuning. With ORGAHEAD, we examine the relative impact of tuning and shaking on organizational performance.

In the stable environment, adaptive organizations tend to engage in much more tuning than do maladaptive organizations (see Fig. 2). In particular, they tend to bring individual agents in and then expend their effort moving the agents about until the reporting structure stabilizes. Tuning, both redesign and retasking, tends to facilitate adaptation. For high performance organizations, they learn that the optimum type of structural learning is redesign, then retasking, then hiring, and then firing whereas the maladaptive organizations lock into a cycle of hire-fire-hire-fire. This simulation result suggests that humans that have been in high performance firms will prefer to change their organization by tuning even if they are told that a shake will result in an optimized organizational structure. Such behavior is observed in real groups where overworked individuals prefer to shed tasks than to change personnel.

There are many other factors that also promote adaptation. Simulation results suggest that adaptability is facilitated by agents having sufficiently complex positions that they need to interact with others, and have been sufficiently cognitively challenged in the past that they have developed a wide enough transactive memory to adapt. Increasing redundancy, so that there are sufficient resources and tasks that change is possible, and pushing the power to handle exceptions (least upper bound-ness), as low in the team as possible, also enhance adaptability, although not as strongly as these other factors. In contrast, performance is enhanced by designing a team that is tuned for the specific set of tasks, is large, and has a low span of control, low cognitive load, and little redundancy—all factors that promote rapid but narrow learning. As can be seen in Fig. 3, factors



Fig. 3. Factors affecting performance and adaptability

promoting adaptation may work against performance and vice versa. This result suggests that it is difficult, and perhaps impossible, to design for both adaptability and high performance. The process distinction between performance and adaptability hinges on flexibility and learning. Rapid learning enhances high performance. Low span of control and low cognitive load enable rapid learning. Teams learn rapidly and are therefore good structures for novel situations or rapid response situations such as fire fighting.

In contrast, adaptability is enhanced by the agents having complex situation awareness and excess (redundant) access, both of which enable flexibility. High cognitive load and redundancy slow learning, yet, over time, increase insight into different situations, thus enabling broader knowledge and greater generalization. Training in teams and creating more complex tasks that require more discussion among personnel would facilitate adaptation.

Consider for a moment the opposite question: what inhibits adaptation? Groups, after all, are easier to manage and respond to if they are not adaptive. The implication of this work is that you can never totally prevent new behaviors from emerging. Learning is ubiquitous and occurs in all agents, and the effects of learning lead to cascades of change across the metamatrix, for example, as individuals learn new ideas that have the potential to change who they are likely to interact with, which resources they can use, and which tasks they can or are assigned to do. Further, if you want to make a team or organization less adaptive, give them little to do, divide the tasks so that there is no need for personnel to communicate to get the task done, isolate personnel, inhibit communication, eliminate resources (e.g., through tight budgets), and so on.

Still another approach to minimizing adaptation may be to convince others to learn only from their successes. Consider the following application. Organizations are often faced with industrial accidents to which they need to respond. In such crisis situations organizations often respond by becoming more rigid (30). Data on 69 organizations faced with industrial accidents were collected. Then, each of these organizations was simulated for their behavior both before and during the crisis. The model predictions fit the data remarkably well in terms of predicting performance. But, with simulation you can go one step further. Not all organizations adapted their architecture when faced with an accident. Simulation was used to predict what would have happened if the organizations that had changed had not, and if those that had not changed had. We found that organizations that do not shift their structure when faced with a threat actually outperform those who do shift. However, shifting structure does significantly improve performance more than the performance gain by not shifting. Left to themselves, these organizations might learn that it pays to shift. However, by using simulation, we ask what would have happened had they not shifted. In this case, the analysis suggests that performance would have been even better had they not shifted.

Will Technology Help Retain Expertise? We often find that technology may not have the expected impact (27, 28). In general, it is difficult to think ahead about the impact of technology because there are so many complex interactions. Here, we use multiagent models to look at the impact of databases. Many organizations are worried that they are losing expertise. As individuals forget, age, and so retire, or move from one company to the next, information is lost. In many cases the information may be critical to the core competencies of the company, thus reducing their long run ability to succeed. One potential solution is to develop databases that the experts fill with their knowledge. The move to Lotus Notes is such a process. But will these databases really serve as an auxiliary memory? By using construct, a virtual experiment was run to look at whether such databases could




Fig. 4. Impact of database on information diffusion.

overcome the data loss because of forgetting. The database was treated as a public good with little information initially but whose value increases as more people contribute their ideas.

The results are shown in Fig. 4. Notice that, even when people forget, eventually everyone will learn new or novel information. However, once a database is put into the corporation, several things happen. First, the rate at which the new idea diffuses slows down. In other words, contributing to the database takes time away from learning novel concepts. Second, rather than preventing information loss, the database may guarantee it. Notice that, when there is a database, some people never learn the new idea.

Will Technology Increase the Unity of Perspective? Another reason for using databases is that they provide a common reference. They should increase shared knowledge within the group, enable more complex and accurate team mental models, and so on. Simulation results, however, suggest that databases may have the opposite effect. Asking individuals to put their information in the database reduces the time they spend with each other conversing. Consequently, it will take longer for shared knowledge and these team mental models to form. Moreover, in the long run, with a database less should be shared (Fig. 5). Remember, having a database guarantees that information does not diffuse completely. One reason for this result is that, because many people can access the database at once, the frequency of known information being repeated increases, thus interfering with novel information diffusing. As a consequence, the level of shared knowledge and the ability to generate a unified perspective may be compromised. These effects hold even when individuals are seeking information.

Conclusion

Herein, a new theoretical perspective on organizations was described that draws on issues of computation, social networks, and knowledge networks. This perspective argues that organizations are synthetic agents (complex, computational, adaptive, and multileveled in their own right) whose behavior is a function

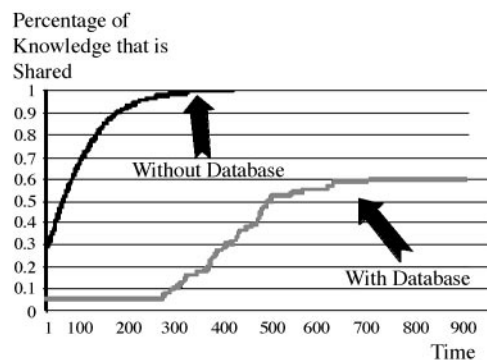


Fig. 5. Impact of database on shared information.

of the webs of affiliation linking tasks, resources, knowledge, and member agents who are themselves complex, computational, and adaptive. It follows from this perspective that multiple types of learning are possible. There are structural learning (learning as reflected by changes in the social network) and experiential learning by individuals, to name just two. Clashes and synergies between structural learning within the synthetic agent and experiential learning within member agents can cause the organization to lock into specific change strategies (to engage in metalearning) that may be detrimental to adaptivity.

Another feature of this perspective is that many technologies can themselves be agents. Databases, for example, are artificial agents with a position in the networks in the metamatrix. When such agents are part of a socio-technical system, there are unintended consequences. In particular, the benefits of such technologies may not be realized. In particular, time spent interacting with these artificial agents impacts time spent interacting with humans.

For every result described herein, it is possible to imagine conditions where the results might not hold. For example, it may be that information has only a window of time where it is useful and capturing such knowledge in a database may be worthwhile. It may be that it is good to lose some information, and such loss may actually improve performance. It is important to note that,

regardless of the methodological approach used by the researcher, there may be alternative explanations for the results observed. There may be reasons to question the applicability of the findings. This result is nothing unique to the multiagent approach. The key point is that these models enable the analyst to look at complex systems and reason about such scenarios in a more systematic fashion than many other forms of theorizing. In many cases, simply building the model brings value in and of itself to the policy maker or manager as it lays bare hidden assumptions and potential limitations in the current system. As we move toward more realistic agents and tasks, the value of these models can only grow.

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-97-1-0037, National Science Foundation (NSF) IRI9633 662, Army Research Labs, NSF ITR/IM IIS-0081219, NSF KDI IIS-9980109, NSF Integrative Graduate Education and Research Traineeship (IGERT): Computational Analysis of Social and Organizational Systems (CASOS), and the Pennsylvania Infrastructure Technology Alliance, a partnership of Carnegie Mellon, Lehigh University, and the Commonwealth of Pennsylvania's Department of Economic and Community Development. Additional support was provided by ICES (the Institute for Complex Engineered Systems) and CASOS center at Carnegie Mellon University (<http://www.casos.ece.cmu.edu>).

1. Carley, K. M. & Gasser, L. (1999) in *Distributed Artificial Intelligence*, ed. Weiss, G. (MIT Press, Cambridge, MA), pp. 299–330.
2. Epstein, J. & Axtell, R. (1997) *Growing Artificial Societies* (MIT Press, Cambridge, MA).
3. Carley, K. M. & Svoboda, D. M. (1996) *Soc. Methods Res.* **25**, 138–168.
4. Carley, K. M. & Lee, J. (1998) in *Advances in Strategic Management*, ed. Baum, J. (JAI, Greenwich, CT), pp. 269–297.
5. Carley, K. (1990) in *Advances in Group Processes: Theory & Research*, eds. Lawler, E., Markovsky B., Ridgeway C. & Walker H. (JAI, Greenwich, CT), pp. 1–44.
6. Carley, K. (1991) *Am. Soc. Rev.* **56**, 331–354.
7. Carley, K. & Newell, A. (1994) *J. Math. Soc.* **19**, 221–262.
8. Carley, K. M. (1999) in *Research in the Sociology of Organizations: On Networks In and Around Organizations*, eds. Andrews, S. B. & Knoke D. (JAI, Greenwich, CT), pp. 3–30.
9. Hutchins, E. (1991) in *Perspectives on Socially Shared Cognition*, eds. Resnick, L. B., Levine, J. M. & Teasley, S. D. (Am. Psychol. Assoc., Washington, DC), pp. 283–307.
10. Hutchins, E. (1995) *Cognition in the Wild* (MIT Press, Cambridge, MA).
11. Carley, K. M. & Prietula, M. (1994) in *Computational Organization Theory*, eds. Carley, K. M. & Prietula, M. (Lawrence Erlbaum Associates, Hillsdale, NJ), pp. 55–87.
12. Krackhardt, D. & Carley, K. M. (1998) in *Proceedings of the 1998 International Symposium on Command and Control Research and Technology* (Evidence Based Research, Vienna, VA), pp. 113–119.
13. Carley, K. M. & Hill, V. (2001) in *Dynamics of Organizations: Computational Modeling and Organizational Theories*, ed. Lomi, A. & Larsen, E. R. (MIT Press/AAAI Press, Live Oak, CA), pp. 63–92.
14. Cyert, R. M. & March, J. G. (1963) *A Behavioral Theory of the Firm* (Prentice—Hall, Englewood Cliffs, NJ).
15. March, J. G. & Simon, H. (1958) *Organizations* (Wiley, New York).
16. Simon, H. (1955) *Q. J. Econ.* **69**, 99–118.
17. Simon, H. (1956) *Psych. Rev.* **63**, 129–138.
18. Burt, R. S. (1992) *Structural Holes: The Social Structure of Competition* (Harvard Univ. Press, Cambridge, MA).
19. Masuch, M. & LaPotin, P. (1989) *Adm. Sci. Q.* **34**, 38–67.
20. Carley, K. M., Prietula, M. J. & Lin, Z. (1998) *J. Artif. Soc. Soc. Simul.* **1**, <http://www.soc.surrey.ac.uk/JASSS/1/3/4.html>.
21. Newell, A. (1990) *Unified Theories of Cognition* (Harvard Univ. Press, Cambridge, MA).
22. March, J. G. (1996) in *Organizational Learning*, eds. Cohen, M. D. & Sproull, L.S. (Sage, London).
23. Levinthal, D. A. & March, J. G. (1981) *J. Econ. Behav. Org.* **2**, 307–333.
24. Contractor, N. & Eisenberg, E. (1990) in *Organizations and Communication Technology*, eds. Fulk, J. & Steinfield, C. (Sage, London), pp. 145–174.
25. Burton, R. M. & Obel, B. (1998) *Strategic Organizational Design: Developing Theory for Application* (Kluwer, Boston).
26. Barley, S. (1990) *Adm. Sci. Q.* **35**, 61–103.
27. Carley, K. (1995) *Soc. Perspect.* **38**, 547–571.
28. Carley, K. M. (1996) *Technol. Soc.* **18**, 219–230.
29. Kaufer, D. S. & Carley, K. M. (1993) *Communication at a Distance: The Effect of Print on Socio-Cultural Organization & Change* (Lawrence Erlbaum Associates, Hillsdale, NJ).
30. Staw, B. M., Sanderlands, L. E. & Dutton, J. E. (1981) *Adm. Sci. Q.* **26**, 501–524.