Cognitive Social Simulation Incorporating Cognitive Architectures

Ron Sun, *Rensselaer Polytechnic Institute*

The CLARION cognitive architecture illustrates why cognitive social simulation with cognitive architectures may have a significant impact on both cognitive science and social simulation.

Agent-based social simulation—modeling social phenomena on the basis of models of autonomous agents—has grown tremendously in recent decades. Researchers use this approach to study a wide range of social and economic issues, including social beliefs and norms, resource allocation, traffic patterns, social cooperation, stock market dynamics, group interaction and dynamics, and organizational decision making. Agent-based social simulation can go beyond the limitations of traditional (equation-based) approaches, which express relationships among conceptual entities (for example, social groups or markets) through a set of mathematical equations.[1]

However, agent-based social simulation with multiagent systems in social computing can benefit from incorporating cognitive architectures. A *cognitive architecture* is a domain-generic computational cognitive model that captures essential structures and processes of the individual mind for the purpose of a broad (multiple-domain) analysis of cognition and behavior.[2,3] Such architectures embody generic descriptions of cognition in computer algorithms and programs and provide a realistic basis for modeling individual agents. More importantly, incorporating cognitive architectures into agent-based social simulation results in *cognitive social simulation*. Here, I discuss the importance and challenges of cognitive social simulation and present CLARION, a cognitive architecture, and its applications to cognitive social simulation.

What are cognitive architectures?

To better understand the concept of a cognitive architecture, consider the following analogy: An architecture for a building consists of its overall framework and design as well as the roof, foundation, walls, windows, floors, and so on. Furniture and appliances are not part of the architecture; they can be easily rearranged or replaced. Likewise, a cognitive architecture includes essential divisions of modules, relations between modules, basic representations, essential algorithms, and so on—that is, those aspects of a system that are relatively invariant across time, domains, and individuals.

A cognitive architecture should provide a framework for facilitating detailed modeling of various components and processes of the mind. Research in computational cognitive modeling explores the essence of cognition and various cognitive functionalities through a detailed, process-based understanding developed by specifying computational models of mechanisms and processes. A cognitive architecture embodies descriptions of cognition in computer algorithms and programs. Researchers then conduct detailed simulations according to the computational models and gather and analyze data on that basis. Thus, cognitive architectures can be beneficial for detailed analysis of cognition.

Social simulation involves exploring social phenomena through computational modeling of such phenomena. Cognitive architectures lead to social-

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

simulation models that are based on capturing detailed processes of individual cognition. Therefore, such social-simulation models tend to be deeper and more detailed. They can help describe, explain, and predict social phenomena, through capturing the cognition of the individuals involved in these phenomena.

Why are they important for social simulation?

Although there are all kinds of cognitive architectures, this article focuses on *psychologically oriented* (as opposed to software engineering oriented) cognitive architectures, which are particularly important for the following reasons:

- they are more cognitively realistic than other types of models, so they are more humanlike in many ways;
- they shed new light on human cognition, so they are useful for advancing cognitive science; and
- they may also constitute a foundation for understanding collective human behavior and social phenomena.

The importance of such cognitive architectures for cognitive science lies in the fact that they are useful in terms of understanding the individual human mind. In understanding cognitive phenomena, using computational simulation based on cognitive architectures forces theoreticians to think clearly in terms of process and mechanistic details rather than use vague, purely conceptual theories. Researchers who use cognitive architectures must specify a cognitive mechanism in sufficient detail to implement the resulting models on computers and run them as simulations. This approach requires explicitly spelling out important elements of the models, leading to better, conceptually clearer theories.

Cognitive architectures also provide a deep level of explanation. They force modelers to think in terms of the mechanisms and processes available within a generic cognitive architecture that is not specifically designed for a particular task. They can then generate explanations of that task that aren't centered on its superficial, high-level features. Describing a task in terms of a cognitive architecture's available mechanisms and processes involves generating explanations centered on primitives of cognition as envisioned in the cognitive architecture, and such explanations tend to be deeper. Moreover, because it's possible to explain a large variety of data and phenomena on the basis of the same set of primitives provided by the same cognitive architecture, this type of theorizing is also more likely to lead to unified explanations for a large variety of data and phenomena. Therefore, using cognitive architectures leads to comprehensive theories of the mind.[2,3]

Likewise, agent-based social simulation is an important development in the social sciences, and the use of agents in social simulation mirrors the development of cognitive architectures in cognitive science. So far, however, these two fields have developed separately (with some exceptions[1]). Most of the work in social simulation assumes rudimentary cognition on the part of the agents.

> Using computational simulation based on cognitive architectures forces theoreticians to think clearly in terms of process and mechanistic details.

However, using cognitive architectures in social simulation leads to cognitively based explanations of social phenomena. This is because social processes ultimately rest on the choices and decisions of individuals, so understanding the mechanisms of individual cognition can lead to better theories describing the behavior of aggregates of individuals. A realistic cognitive agent model, incorporating realistic tendencies, inclinations, and capabilities of individual cognitive agents can serve as a solid basis for understanding the interaction of individuals. There has already been some work in this area.[4,5] In particular, I presented a set of current projects in the book *Cognition and Multi-Agent Interaction*.[1]

What this all boils down to is cognitive social simulation, or cognitively based social simulation, as opposed to mere agent-based social simulation.[1] Cognitive architectures could be an important centerpiece of this enterprise.

The Clarion cognitive architecture

Clarion is an integrative cognitive architecture consisting of several distinct subsystems.[3,6,7] These include the action-centered subsystem (ACS), the non-action-centered subsystem (NACS), the motivational subsystem (MS), and the metacognitive subsystem (MCS). The ACS controls actions, whether for external physical movements or internal mental operations. The NACS maintains general knowledge, either implicit or explicit. The MS provides underlying motivations for perception, action, and cognition in terms of impetus and feedback (for example, indicating whether outcomes are satisfactory or not). The MCS monitors, directs, and modifies the operations of the ACS dynamically, as well as the operations of all the other subsystems.

Each of these interacting subsystems has a dual-representational structure, meaning it consists of two "levels" of representation: The top level encodes explicit knowledge; the bottom level encodes implicit knowledge. (Detailed arguments about the distinction between implicit and explicit knowledge on the basis of psychological data are available elsewhere.[3]) The two levels interact by cooperating in action through a combination of each level's action recommendations and by cooperating in learning through a bottom-up process and a top-down process. Thus, Clarion is essentially a dual-process theory of cognition. Figure 1 illustrates the Clarion architecture.

The action-centered subsystem

The ACS operates as follows:

1. Observe current state x.
2. Compute at the bottom level the Q-values of x associated with each of all the possible actions a_i: $Q(x, a_1)$, $Q(x, a_2)$, …, $Q(x, a_n)$.
3. Determine all possible actions (b_1, b_2, …, b_m) at the top level, based on input x (sent up from the bottom level) and the rules in place.
4. Compare or combine the selected a_i values with those of b_j (sent down from the top level) and choose an appropriate action, b.
5. Perform action b and observe the next state y and (possibly) the reinforcement r.
6. Update the Q-values at the bottom level in accordance with the Q-learning

algorithm (implemented with backpropagation).
7. Update the rule network at the top level using the rule-extraction-refinement algorithm.
8. Go back to step 1.

The bottom level of the ACS learns implicit reactive routines. A Q-value evaluates the quality of an action in a given state: $Q(x, a)$ indicates how desirable action a is in state x (which consists of some sensory input). The agent may choose an action in any state based on the Q-values. The Q-learning algorithm (implemented with backpropagation) can help acquire the Q-values.[8] This reinforcement learning algorithm basically compares the values of successive actions and adjusts an evaluation function on that basis. It thereby develops implicit sequential behaviors.[7]

The bottom level is modular; that is, several small neural networks coexist, each adapted to specific modalities, tasks, or groups of input stimuli. This coincides with the modularity claim that much cognitive processing is done by limited, encapsulated (to some extent), specialized processors that are highly efficient.[3]

The top level captures explicit conceptual knowledge in the form of rules.[7] There are many ways to learn explicit knowledge, including independent hypothesis-testing learning and bottom-up learning.

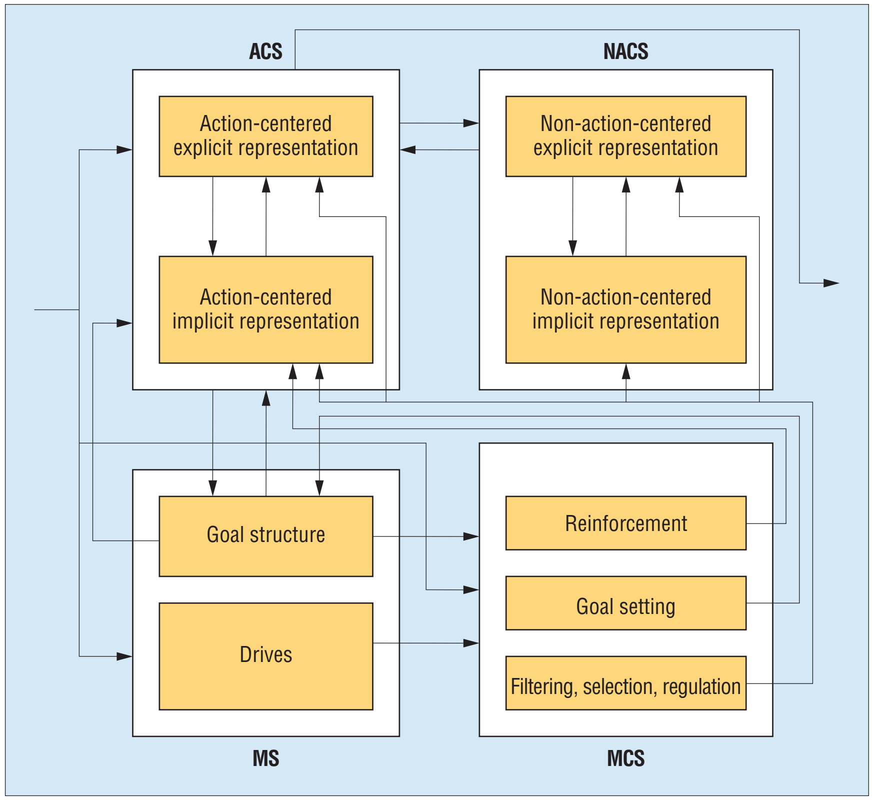

Figure 1. The Clarion architecture, including its four subsystems: the action-centered subsystem (ACS), the non-action-centered subsystem (NACS), the motivational subsystem (MS), and the metacognitive subsystem (MCS).

Autonomous generation of explicit conceptual structures. People generally learn implicit knowledge through trial and error. They may acquire explicit knowledge also from ongoing experience in the world, possibly through the mediation of implicit knowledge (thus, the idea of bottom-up learning[6]). The basic process of bottom-up learning is as follows: If an action implicitly decided by the bottom level is successful, then the agent extracts an explicit rule corresponding to the action selected by the bottom level and adds the rule to the top level. Then, in subsequent interaction with the world, the agent verifies the extracted rule by considering the outcome of applying the rule: If the outcome is not successful, then the rule should be made more specific and should exclude the current case; if the outcome is successful, the agent can try to generalize the rule to make it more universal.[6] After the top level of the ACS has learned explicit rules, a variety of explicit reasoning methods can be used. Learning explicit conceptual representation at the top level can also be useful for enhancing learning of implicit reactive routines at the bottom level.

Assimilation of externally given conceptual structures. Although Clarion can learn even when no a priori or externally provided knowledge is available, it can use such knowledge when it is available. For instructed learning, the ACS can combine externally provided knowledge (in the form of explicit conceptual structures such as rules, plans, categories, and so on) with existing conceptual structures at the top level—that is, *internalization*; the ACS can also assimilate this knowledge into implicit reactive routines at the bottom level—that is, *assimilation*. This process of internalization and assimilation is known as top-down learning.[7]

The non-action-centered subsystem

The NACS represents general knowledge about the world and performs various kinds of memory retrievals and inferences. The NACS is under the control of the ACS (through the latter's actions).

At the bottom level of the NACS, associative-memory networks encode non-action-centered implicit knowledge. The backpropagation learning algorithm can establish associations by mapping an input to an output. At the top level of the NACS, a general knowledge store encodes explicit non-action-centered knowledge. The NACS specifies chunks (that is, concepts) through dimensional values. The NACS sets up a node at the top level to represent a chunk. The chunk node connects to its corresponding features (dimensional values), represented as individual nodes at the bottom level. Additionally, links between chunks encode *associative rules*, which are explicit associations between pairs of chunks. There are various ways to learn explicit associative rules.[7] In addition to such rules, similarity-based reasoning is possible through the interaction of the two levels of representation.[7]

As with the ACS, the NACS allows top-down or bottom-up learning, either to extract explicit knowledge of the top level from the implicit knowledge at the bottom level or to assimilate explicit knowledge from the top level into implicit knowledge at the bottom level.

Table 1. Human and simulation data for the organizational decision task, comparing blocked and distributed information access for teams and hierarchies.

Human or model	Percent correct: Team (blocked)	Team (distributed)	Hierarchy (blocked)	Hierarchy (distributed)
Human	50.0	56.7	46.7	55.0
Radar-Soar	73.3	63.3	63.3	53.3
CORP-P-ELM	78.3	71.7	40.0	36.7
CORP-ELM	88.3	85.0	45.0	50.0
CORP-SOP	81.7	85.0	81.7	85.0

The motivational and metacognitive subsystems

The MS deals with drives and their interactions, which lead to actions. It is concerned with why an agent does what it does. Simply saying that an agent chooses actions to maximize gains, rewards, or reinforcements leaves open the question of what determines these things.

The MS contains dual motivational representations (based on relevant psychological evidence).[7] An agent's explicit goals (for example, "finding food"), which partly determine the operation of the ACS, can be generated according to internal drive states (for example, "being hungry").[7]

Beyond low-level drives concerning physiological needs, there are also higher-level drives. Some of them are primary, in the sense of being hard-wired. Whereas primary drives are built in and relatively unalterable, there are also derived drives, which are secondary, changeable, and acquired mostly in the process of satisfying primary drives.

The MCS is closely tied to the MS. It monitors, controls, and regulates cognitive processes to improve cognitive performance. Control and regulation can be in the form of setting goals for the ACS, setting essential parameters of the ACS and the NACS, interrupting and changing ongoing processes in the ACS and the NACS, and so on. Control and regulation can also be in the form of setting reinforcement functions for the ACS. All of these depend on drive states and goals in the MS. The MCS also has two levels: top (explicit) and bottom (implicit).

Applications to cognitive social simulation

Here, I discuss three case studies of applying Clarion to social simulation. These range from small-scale organizational simulations to large-scale simulations of the phenomenon of academic publishing, and further on to even larger-scale simulations of tribal societies.

Case study: Simulating organizational decision making

One application of Clarion to cognitive social simulations is in understanding organizational decision making and the interaction between organizational structures and cognitive factors in affecting organizational performance.[4] There are two major types of organizational structures:

- *teams*, in which agents act autonomously, individual decisions are treated as votes, and the organizational decision is the majority decision; and
- *hierarchies*, in which agents are organized in a chain of command such that information passes from subordinates to superiors, and superiors' decisions are based solely on their subordinates' recommendations.

Another way to distinguish organizations is according to the structure of information accessible to each agent. Two varieties of information access are

- *distributed access*, in which each agent sees a different subset of attributes (no two agents see the same subset); and
- *blocked access*, in which several agents see exactly the same subset of attributes.

Kathleen Carley, Michael Prietula, and Zhiang Lin conducted human experiments in a 2 × 2 fashion (organization × information access),[9] and they compared the human data from the experiments to the results of four simulation models. Among them, CORP-ELM produced the most probable classification based on an agent's own experience. CORP-P-ELM stochastically produced a classification in accordance with the estimate of the probability of each classification based on the agent's own experience. CORP-SOP followed the organizationally prescribed standard operating procedure, which involved summing up the values of the attributes available to an agent, and thus was not adaptive. Radar-Soar was a somewhat cognitive model built in Soar, which was based on explicit, elaborate searches in problem spaces. See table 1.

In their work, Carley, Prietula, and Lin used very simple agent models with a low intelligence level. Moreover, learning in these simulations was rudimentary; there was no complex learning process as in humans. With these shortcomings in mind, Isaac Naveh and I conducted simulations involving more complex agent models that more accurately captured human performance. Moreover, by using more cognitively realistic agent models, we could individually investigate the importance of different cognitive capacities and process details affecting organizational performance.[4] Hence, we conducted simulations using Clarion to model individual agents in an organization.

As table 2 shows, the results closely accord with the patterns of the human data, with teams outperforming hierarchal structures, and distributed access proving superior to blocked access. Also, as with humans, performance was not grossly skewed toward one condition or another but was roughly comparable across all conditions—unlike some of the simulation results from Carley, Prietula, and Lin.[9] The Clarion simulations' far better match with the human data was due, at least in part, to a higher degree of cognitive realism. (Further details are available elsewhere,[4] including interesting effects from varying cognitive parameters.)

The use of a cognitive architecture in these simulations enabled the exploration of the interaction between cognitive factors and parameters on the one hand and social and organizational structures on the other. Because Clarion captures a wide range of generic cognitive processes and phenomena, its parameters are generic rather than task specific. Thus, it's possible to study specific issues, such as organizational design, in the context of a general theory of cognition. Isaac Naveh and I performed a detailed statistical analysis of simulation data resulting from varying cognitive parameters in a factorial design.[4] We found that some cognitive parameters had a monolithic, across-the-board effect, whereas in other cases complex interactions of factors were at work. This illustrated the advantages of using cognitive architectures in social simulation, whereby cognitive parameters could be easily varied. This analysis accentuated the importance in social simulation of limiting conclusions to the specific cognitive (and other) context in which the data was obtained, without overgeneralizing the conclusions.[4]

Cognitive social simulation can help generate new theories and hypotheses in this regard. Moreover, it can reduce the need for costly (and sometimes impossible) human experiments or at least can make such experiments more focused on testing specific hypotheses generated by social simulation based on cognitive architectures.

In summary, by using CLARION, we more accurately captured organizational performance data and formulated deeper explanations of the observed results.[4] In the future, CLARION might help predict human performance in social or organizational settings. It might also help improve collective performance by prescribing optimal or near-optimal cognitive abilities for individuals, for specific collective tasks or organizational structures.[4]

Table 2. CLARION simulation. "Human" refers to the human data from Carley, Prietula, and Lin. Performance of CLARION was computed over the last 1,000 cycles of the 3,000-cycle simulations.

	Percent correct			
Human or model	Team (blocked)	Team (distributed)	Hierarchy (blocked)	Hierarchy (distributed)
Human	50.0	56.7	46.7	55.0
CLARION	53.2	59.3	45.0	49.4

Table 3. Number of authors contributing to *Chemical Abstracts*.

No. of articles published	No. of authors contributing			
	Actual[10]	Simon's estimate[10]	Gilbert's simulation[11]	CLARION simulation[5]
1	3,991	4,050	4,066	3,803
2	1,059	1,160	1,175	1,228
3	493	522	526	637
4	287	288	302	436
5	184	179	176	245
6	131	120	122	200
7	113	86	93	154
8	85	64	63	163
9	64	49	50	55
10	65	38	45	18
11 or more	419	335	273	145

Case study: Simulating academic publishing

Another application of CLARION to cognitive social simulation is in explaining the essential process of academic publication and its relation to cognitive processes.[5] Science develops in certain ways. In particular, the number of authors contributing a certain number of articles to a scientific journal tends to follow a highly skewed distribution, corresponding to an inverse power law. In the case of scientific publications, the tendency of authorship to follow such a distribution is known as Lotka's law. Herbert Simon developed a simple stochastic process for approximating Lotka's law.[10] One of the assumptions underlying this process is that the probability that a paper will be published by an author who has published i articles is a/i^k, where a is a constant of proportionality. Using Simon's work as a starting point, Nigel Gilbert attempted to model Lotka's law.[11] He obtained his simulation data on the basis of some simplifying assumptions and a set of mathematical equations. To a significant extent, Gilbert's model was not cognitively realistic. The model assumed that authors were noncognitive and interchangeable; it therefore neglected a host of cognitive phenomena that characterized scientific inquiry (for example, learning, creativity, and so on).

Using a more cognitively realistic model, Naveh and I addressed some of these omissions and explored other emergent properties of a cognitively based model and their correspondence to real-world phenomena.[5] Tables 3 and 4 compare the actual results reported by Simon for two journals (*Chemical Abstracts* and *Econometrica*),[10] with CLARION simulation results[5] and estimates obtained from previous simulations by Simon and Gilbert.[10,11] Both tables indicate the number of authors contributing to each journal in accordance with the number of articles they published.

The CLARION simulation data for the two journals fit the power curve $f(i) = a/i^k$ with an excellent match. Table 5 gives the results of the curve fit along with correlation and error measures.

In our simulations, the number of articles per author reflected the cognitive ability of an author, as opposed to being based on auxiliary assumptions such as those made by Gilbert.[11] This partly explains the slightly greater divergence of our results from the human data: Whereas Gilbert's simulation consisted of equations selected specifically to match the human data, our approach relied on more detailed, lower-level mechanisms—namely, a cognitive agent model that was generic rather than task specific. The CLARION simulation results were, thus, emergent, and not from specific, direct attempts to match the human data. That is, we put more distance between mechanisms and outcomes, making it more difficult to obtain a match with the human data. Thus, the fact that we were able to match the human data reasonably well shows the power of our approach.

Case study: Simulating survival strategies of tribal societies

Yet another application of CLARION to cognitive social simulation is in simulating the survival strategies of tribal societies under various environmental conditions. (Space limitations prevent me from covering the technical details of the work here, but they will be published later separately). These simulations represented the world as a 2D grid, and randomly distributed food items and agents. There were harsh, medium, and benign conditions, distinguished by the agent-to-food ratios. Agents were of a limited life span—

Table 4. Number of authors contributing to *Econometrica*.

No. of articles published	No. of authors contributing: Actual[10]	Simon's estimate[10]	Gilbert's simulation[11]	Clarion simulation[5]
1	436	453	458	418
2	107	119	120	135
3	61	51	51	70
4	40	27	27	48
5	14	16	17	27
6	23	11	9	22
7	6	7	7	17
8	11	5	6	18
9	1	4	4	6
10	0	3	2	2
11 or more	22	25	18	16

Table 5. Results of fitting Clarion data to power curves $f(i) = a/i^k$. (RMSE: root-mean-square error).

Journal	No. of authors contributing: a	k	Pearson R	R-square	RMSE
Chemical Abstracts	3,806	1.63	0.999	0.998	37.62
Econometrica	418	1.64	0.999	0.999	4.15

varying for each individual agent, depending on its energy level. Agents looked for and consumed food in an effort to prolong their life spans.

A tribe in which each agent uses only its own resources has adopted an *individual survival strategy*. A tribe in which resources can be transferred from one individual to another has adopted a *social survival strategy*. For example, the *central store* is a mechanism to which all the individuals in a tribe transfer part of their resources. The resources collected by the central store can be redistributed to the members of the tribe.

The agents in this simulation, unlike previous simulations of tribal societies, were cognitively realistic. The Clarion cognitive architecture from which Naveh and I constructed these agents captured a variety of cognitive processes in a psychologically realistic way.[3,7] Therefore, simulating social survival strategies could shed more light on the role of cognition in determining survival strategies and its interaction with social structures (social institutions) and processes. The major motivation behind these simulations was to investigate the interaction of individual cognition with social structures and processes—that is, the micro-macro link.

On the one hand, we found relationships between various cognitive parameters and social variables, indicating that the social institutions and norms (such as survival strategies) adopted might have something to do with the cognitive abilities and cognitive tendencies of the agents involved. This relation, which we may call the *social-cognitive dependency*, could have significant theoretical and empirical ramifications. For example, some forms of social systems (structures and institutions) might be suitable for certain cognitive characteristics but unsuitable for others. Thus, such social systems might not be universally better or worse. Rather, a host of other factors—cognitive factors, in particular—might affect which social system is best in each situation. (An earlier work of mine contains a fairly substantial discussion of the close relationship between cognitive and social processes,[1] and advocates the exploration of cognitive principles of sociocultural processes.)

On the other hand, some cognitive attributes might have been selected (through natural evolution) to work with certain social systems and cultural environments, which we may call the *cognitive-social dependency*. In this regard, we may explore sociocultural principles of cognition—the opposite of cognitive principles of sociocultural processes.

We also found that the relation between various cognitive parameters and physical environmental variables was such that certain cognitive attributes were universally good or bad, whereas the effects of some other cognitive attributes were more dependent on environmental attributes. Cognitive attributes might have been selected (through natural evolution) to work within given physical environments, which we may call the *cognitive-physical dependency*.

Together, these three types of dependency form a complex dynamic system of interwoven dependencies and interactions. In such a system, it's important to understand not only direct effects of dependencies but also indirect effects that are not obviously related to their causes but are often crucial for discerning the system's functional structures.

In summary, these simulations show that in the context of different social survival strategies and different physical environments, cognition matters. For instance, cognition can help determine which survival strategies and other social variables are appropriate under particular cognitive conditions. Even though we used only very simple representations of sociocultural processes in this work, we nevertheless found significant effects of various interactions, thanks to the cognitive architecture we used.

The challenges of cognitive social simulation

The development of agent-based social simulation (as a means for computational study of societies and social phenomena) has been mirroring the development of cognitive architectures in cognitive science. Integrating these two fields opens up an important opportunity but also involves challenges.

Social processes ultimately rest on the choices and decisions of individuals. Hence, understanding the mechanisms of individual cognition can lead to better theories describing the behavior of aggregates of individuals. Most agent models in social simulation have been extremely simple. However, a cognitive agent model that incorporates realistic tendencies, inclinations, and capabilities of individual cognitive agents could serve as a more realistic basis for understanding the interaction of individuals. Although some cognitive details might ultimately prove irrelevant, it's not possible to determine this a priori. Thus, simulations are useful for determining which aspects of cognition can be safely abstracted away.

At the same time, integrating social simulation and cognitive modeling can lead to a better understanding of individual cognition. Traditional approaches to cognitive modeling have largely ignored the potentially decisive effects of the social aspects of cognition (social beliefs, norms, and so on). By modeling cognitive agents in a social context, we can learn more about the sociocultural processes that influence individual cognition.

The most fundamental challenge is to develop better ways of conducting detailed social simulations based on cognitive architectures as basic building blocks. Although there has been some initial work in this area,[1,4,5] much more work is still needed.

One specific challenge concerns how to enhance cognitive architectures to account for sociality in individual cognitive agents. There are many questions in this regard. For example, what are the characteristics of a proper cognitive architecture for modeling the interaction of cognitive agents? What additional sociocultural representations (for example, motives, obligations, or norms) are needed in cognitive modeling of multiagent interaction? For further discussions, see my earlier work.[1]

There is also the challenge of computational complexity and thus scalability. Social simulation could involve many agents—up to thousands. Computational complexity is thus already high, even without involving cognitive architectures as agent models. Incorporating cognitive architectures in social simulation will add considerably more complexity. Thus, scalability is a significant issue.

Finally, whether or not to use detailed cognitive models in social simulation is a decision that should be made on a case-by-case basis. There are many reasons for using or not using detailed cognitive models in social simulation. As just mentioned, complexity may be an issue that prevents wider use of detailed cognitive models in social simulation. Also, in many cases, it might be necessary to capture social processes, social institutions, and social mechanisms more extensively and more directly in models.

Cognitive social simulation with cognitive architectures will have a profound impact on both cognitive science and social simulation. For example, it might help us better understand the role of cognition in social interaction and the role of sociality in cognitive processes. The area of cognitive social simulation with the use of cognitive architectures is an essential aspect of social computing and an important research direction in this emerging field. Therefore, we should put a significant amount of collective research effort into this area.

Acknowledgments

This work was supported in part by ARI grants DASW01-00-K-0012 and W74V8H-04-K-0002 (to Ron Sun and Bob Mathews). Isaac Naveh conducted much of the simulation work described here.

References

1. R. Sun, ed., *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*, Cambridge Univ. Press, 2006.
2. A. Newell, *Unified Theories of Cognition*, Harvard Univ. Press, 1990.
3. R. Sun, *Duality of the Mind*, Lawrence Erlbaum Associates, 2002.
4. R. Sun and I. Naveh, "Simulating Organizational Decision-Making Using a Cognitively Realistic Agent Model," *J. Artificial Societies and Social Simulation*, vol. 7, no. 3, 2004, http://jasss.soc.surrey.ac.uk/7/3/5.html.
5. I. Naveh and R. Sun, "A Cognitively Based Simulation of Academic Science," *Computational and Mathematical Organization Theory*, vol. 12, no. 4, 2006, pp. 313–337.
6. R. Sun, E. Merrill, and T. Peterson, "From Implicit Skills to Explicit Knowledge: A Bottom-Up Model of Skill Learning," *Cognitive Science*, vol. 25, no. 2, 2001, pp. 203–244.
7. R. Sun, *A Tutorial on Clarion 5.0*, tech. report, Cognitive Science Dept., Rensselaer Polytechnic Inst., 22 July 2003; www.cogsci.rpi.edu/~rsun/sun.tutorial.pdf.
8. C.J.C.H. Watkins, "Learning with Delayed Rewards," doctoral dissertation, Psychology Dept., Univ. of Cambridge, 1989.
9. K. Carley, M. Prietula, and Z. Lin, "Design Versus Cognition: The Interaction of Agent Cognition and Organizational Design on Organizational Performance," *J. Artificial Societies and Social Simulation*, vol. 1, no. 3, 1998, pp. 1–19.
10. H. Simon, *Models of Man: Social and Rational*, Wiley, 1957.
11. N. Gilbert, "A Simulation of the Structure of Academic Science," *Sociological Research Online*, vol. 2, no. 2, 1997, www.socresonline.org.uk/socresonline/2/2/3.html.

The Author

Ron Sun is a professor in the Cognitive Science Department at Rensselaer Polytechnic Institute in Troy, New York. His research interests include cognitive architectures, learning and skill acquisition, cognitive social simulation and multiagent systems, everyday commonsense reasoning, computational studies of consciousness, and connectionist models and hybrid systems. Sun received his PhD in computer science from Brandeis University. Contact him at Cognitive Science Department, Rensselaer Polytechnic Inst., 110 Eighth St., Troy, New York 12180; rsun@rpi.edu.