Cognitive Social Simulation Incorporating Cognitive Architectures

Ron Sun

June 15, 2007

Abstract

Agent-based social simulation (with multi-agent systems), which is an important aspect of social computing, can benefit from incorporating cognitive architectures, as they provide a realistic basis for modeling individual agents and therefore their social interactions. A cognitive architecture is a domain-generic computational cognitive model that may be used for a broad multiple-domain analysis of individual behavior. In this article, an example of a cognitive architecture is given, and its applications to social simulation described. Some challenging issues in this regard are outlined.

KEYWORDS: social simulation, cognition, cognitive architecture, organization, society, social computing

1 Introduction

Agent-based social simulation, that is, modeling of social phenomena on the basis of the models of autonomous agents, is an important aspect of social computing. We believe that cognitive architectures can play a significant role in it.

Agent-based social simulation is the kind of simulation of social processes that is based on the interaction of autonomous individual entities. Central to such simulation is the concept of an "agent". It is generally agreed that agents are computational entities, typically implemented in software, capable of flexible autonomous action. Agents perceive and influence aspects of their environment (including each other), and they may learn. From their interactions, complex patterns may emerge. Thus, the interactions among agents provide explanations for corresponding social phenomena. Agent-based social simulation has seen tremendous growth in the recent decade. Researchers hoping to go beyond the limitations of traditional approaches to the social sciences have increasingly turned to agent-based social simulation for studying a wide range of theoretical and practical social and economic issues. [1] Issues addressed thus far by social simulation include, for example, social beliefs, social norms, resource allocation, traffic patterns, social cooperation, stock market dynamics, group interaction and dynamics, organizational decision making, and countless others.

Related to that, a cognitive architecture is a domain-generic computational cognitive model that may be used for a broad, multiple-domain analysis of individual behavior. It embodies generic descriptions of cognition in computer algorithms and programs. Agent-based social simulation (with multi-agent systems) in social computing can benefit from incorporating cognitive architectures, as they provide a realistic basis for modeling individual agents.

In this article, an example of a cognitive architecture will be given, and its applications to social simulation will be described. First of all, in section 2, the question of what a cognitive architecture is is answered in detail. In section 3, its importance to cognitive modeling and social simulation is addressed. In section 4, an example cognitive architecture, CLARION, is described (necessarily in a sketchy way). Then, in section 5, its application to social simulation is discussed through several examples. In section 6, the importance and the challenging issues of cognitive social simulation (based on cognitive architectures) are further accentuated. Section 7 concludes this article.

Note that this article is aimed to provide an overview of a body of work and to situate it in the context of social computing. As such, it is not aimed to provide technical details, and not aimed to describe new technical materials. Rather, it interprets and relates current (existing) research issues and ideas, and provides an overall perspective.

[1] Agent-based social simulation differs markedly from traditional (equation-based) approaches to simulation of social phenomena, where relationships among conceptual entities (e.g., social groups or markets) are expressed through a set of mathematical equations. Agent-based modeling has a number of advantages over equation-based modeling (Sun 2006).

2 What is a Cognitive Architecture after all?

As mentioned before, a cognitive architecture is a broadly-scoped, domain-generic computational cognitive model, capturing the essential structure and process of the individual mind, to be used for a broad multiple-domain analysis of cognition and behavior (Newell 1990, Sun 2002). An analogy is appropriate here: The architecture for a building consists of its overall framework and its overall design, as well as roofs, foundations, walls, windows, floors, and so on. Furniture and appliances can be easily rearranged and/or replaced and therefore they are not part of the architecture. By the same token, a cognitive architecture includes overall structures, essential divisions of modules, relations between modules, basic representations, essential algorithms, and a variety of other aspects. In general, an architecture includes those aspects of a system that are relatively invariant across time, domains, and individuals.

In relation to understanding the individual human mind (i.e., cognitive science), a cognitive architecture provides a concrete framework for more detailed modeling of cognitive phenomena, through specifying essential structures, divisions of modules, relations between modules, and so on. Its function is to provide an essential framework to facilitate more detailed modeling and understanding of various components and processes of the mind. Research in computational cognitive modeling explores the essence of cognition and various cognitive functionalities through developing detailed, process-based understanding by specifying computational models of mechanisms and processes. It embodies descriptions of cognition in computer algorithms and programs. Detailed simulations are then conducted based on the computational models. In this enterprise, cognitive architectures may be beneficially employed (for a broad analysis of cognition).

In relation to social simulation, that is, in relation to the work on exploring social phenomena through computational modeling of these phenomena, cognitive architectures lead to a type of computational model of social phenomena that is from the ground up — based on capturing detailed processes of individual cognition. Therefore, the use of cognitive architectures leads to a more detailed (and deeper) computational model of social phenomena. Such models may be used to describe, explain, and/or predict social phenomena, through capturing the cognition of the individuals involved in the social phenomena. They may be run as simulation, and data may be gathered and analyzed on that basis.

3 Why are Cognitive Architectures Important for Social Simulation?

While there are all kinds of cognitive architectures in existence, in this article, I am concerned specifically with psychologically oriented cognitive architectures (as opposed to software engineering oriented “cognitive” architectures). Psychologically oriented cognitive architectures are particularly important because (1) they are “intelligent” systems that are cognitively realistic (relatively speaking) and therefore they are more human-like in many ways, (2) they shed new light on human cognition and therefore they are useful tools for advancing the science of cognition, (3) furthermore, they may (in part) serve as a foundation for understanding collective human behavior and social phenomena.

For cognitive science, the importance of such cognitive architectures lie in the fact that they are useful in terms of understanding the individual human mind. In understanding cognitive phenomena, the use of computational simulation on the basis of cognitive architectures forces one to think in terms of process and mechanistic details. Instead of using vague, purely conceptual theories, cognitive architectures force theoreticians to think clearly. Researchers who use cognitive architectures must specify a cognitive mechanism in sufficient detail to allow the resulting models to be implemented on computers and run as simulations. This approach requires that important elements of the models be spelled out explicitly, thus aiding in developing better, conceptually clearer theories.

Cognitive architectures also provide a deeper level of explanation. Instead of a model specifically designed for a specific task (often in an ad hoc way), using a cognitive architecture forces modelers to think in terms of the mechanisms and processes available within a generic cognitive architecture that are not specifically designed for a particular task, and thereby to generate explanations of the task that is not centered on superficial, high-level features of a task. To describe a task in terms of available mechanisms and processes of a cognitive architecture is to generate explanations centered on primitives of cognition as envisioned in the cognitive architecture, and therefore such explanations are deeper explanations. This type of theorizing is also more likely to lead to unified explanations for a large variety of data and/or phenomena, because potentially a large variety of task data and phenomena can be explained on the basis of the same set of primitives provided by the same cognitive architecture. Therefore, using cognitive architectures leads to comprehensive theories of the mind (Newell 1990, Sun 2002).

Correspondingly, an important development in the social sciences has been that of agent-based social simulation (i.e, ABSS). The use of agents in social simulation mirrors the development of cognitive architectures in cognitive science. So far, however, the two fields of social simulation and cognitive architectures have developed separately from each other (with some exceptions; e.g., Sun 2006). Most of the work in social simulation assumed rudimentary cognition on the part of the agents.

However, for the field of social simulation, the use of cognitive architectures leads to an interesting kind of explanation of social phenomena — cognitively based explanation of social phenomena. This is because social processes ultimately rest on the choices and decisions of individuals, and thus 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 been some work in this area, including, for example, Sun and Naveh (2004), Naveh and Sun (2006), and Sun (2006). In particular, Sun (2006) covers a wide range of current projects in this area by various active research groups.

What this boils down to is cognitive social simulation, or cognitively-based social simulation, as opposed to mere agent-based social simulation (ABSS), as argued for in Sun (2006). Cognitive architectures may be an important centerpiece of this enterprise.

4 An Example of a Cognitive Architecture

4.1 An Overview

Below I will describe a cognitive architecture CLARION (which has been described and justified extensively in Sun et al 2001, Sun 2002, 2003, and so on). CLARION is an integrative architecture, consisting of a number of distinct subsystems, with a dual representational structure in each subsystem (implicit versus explicit representations). Its subsystems include the action-centered subsystem (the ACS), the non-action-centered subsystem (the NACS), the motivational subsystem (the MS), and the meta-cognitive subsystem (the MCS). The role of the action-centered subsystem is to control actions, regardless of whether the actions are for external physical movements or for internal mental operations. The role of the non-action-centered subsystem is to maintain general knowledge, either implicit or explicit. The role of the motivational subsystem is to provide underlying motivations for perception, action, and cognition, in terms of providing impetus and feedback (e.g., indicating whether outcomes are satisfactory or not). The role of the meta-cognitive subsystem is to monitor, direct, and modify the operations of the action-centered subsystem dynamically as well as the operations of all the other subsystems.

Each of these interacting subsystems consists of two levels of representation (i.e., a dual representational structure): Generally, in each subsystem, the top level encodes explicit knowledge and the bottom level encodes implicit knowledge. The distinction of implicit and explicit knowledge has been amply argued for before based on psychological data (see Sun 2002). The two levels interact, for example, by cooperating in actions, through a combination of the action recommendations from the two levels respectively, as well as by cooperating in learning through a bottom-up and a top-down process (to be discussed below). Essentially, it is a dual-process theory of mind. See Figure 1.

4.2 Some Details

4.2.1 The Action-Centered Subsystem

First, let us focus on the action-centered subsystem (the ACS) of CLARION. The operation of the action-centered subsystem may be described as follows:

1. Observe the current state x.
2. Compute in the bottom level the Q-values of x associated with each of all the possible actions a_i's: $Q(x, a_1)$, $Q(x, a_2)$, ..., $Q(x, a_n)$.
3. Find out all the possible actions (b_1, b_2, ..., b_m) at the top level, based on the input x (sent up from the bottom level) and the rules in place.
4. Compare or combine the values of the selected a_i's with those of b_j's (sent down from the top level), and choose an appropriate action b.
5. Perform the action b, and observe the next state y and (possibly) the reinforcement r.
6. Update Q-values at the bottom level in accordance with the *Q-Learning-Backpropagation* algorithm
7. Update the rule network at the top level using the *Rule-Extraction-Refinement* algorithm.
8. Go back to Step 1.

In the bottom level of the action-centered subsystem, implicit reactive routines are learned: A Q-value is an evaluation of the "quality" of an action

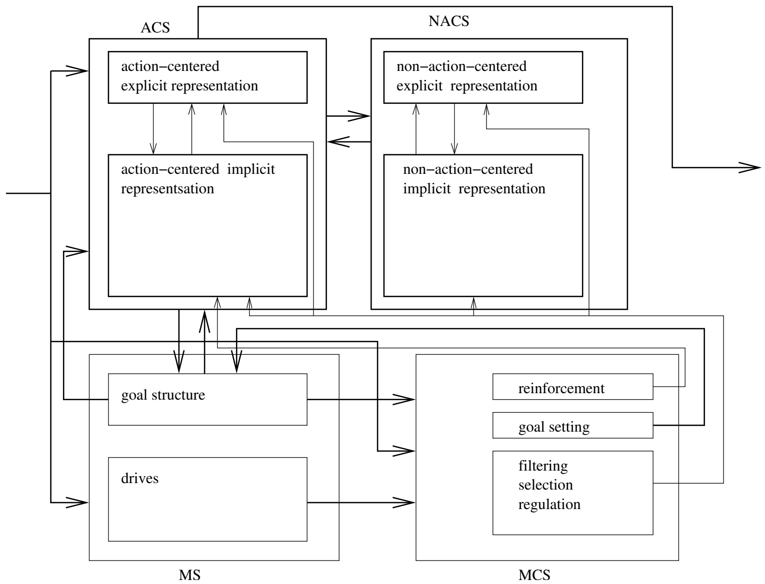

Figure 1: The CLARION Architecture

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 Q-values. To acquire the Q-values, the *Q-learning* algorithm (Watkins 1989) may be used, which is a reinforcement learning algorithm. It basically compares the values of successive actions and adjusts an evaluation function on that basis. It thereby develops (implicit) sequential behaviors (Sun 2003).

The bottom level of the action-centered subsystem is modular; that is, a number of small neural networks co-exist each of which is adapted to specific modalities, tasks, or groups of input stimuli. This coincides with the modularity claim that much processing is done by limited, encapsulated (to some extent), specialized processors that are highly efficient (Sun 2002).

In the top level of the action-centered subsystem, explicit conceptual knowledge is captured in the form of rules. See Sun (2003) for details. There are many ways in which explicit knowledge may be learned, including independent hypothesis-testing learning and "bottom-up learning" as discussed below.

Autonomous Generation of Explicit Conceptual Structures. People are generally able to learn implicit knowledge through trial and error. On top of that, explicit knowledge can be acquired also from on-going experience in the world, through the mediation of implicit knowledge (i.e., the idea of bottom-up learning; Sun et al 2001). 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 that corresponds 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 exclusive of the current case; if the outcome is successful, the agent may try to generalize the rule to make it more universal. [2] After explicit rules have been learned, a variety of explicit reasoning methods may be used. Learning explicit conceptual representation at the top level can also be useful in enhancing learning of implicit reactive routines at the bottom level (e.g., Sun et al 2001).

Assimilation of Externally Given Conceptual Structures. Although CLARION can learn even when no a priori or externally provided knowledge is available, it can make use of it when such knowledge is available. To deal with instructed learning, externally provided knowledge, in the forms of explicit conceptual structures such as rules, plans, categories, and so on, can (1) be combined with existent conceptual structures at the top level (i.e., internalization), and (2) be assimilated into implicit reactive routines at the bottom level (i.e., assimilation). This process is known as top-down learning. See Sun (2003) for more details.

4.2.2 The Non-Action-Centered Subsystem

The non-action-centered subsystem (NACS) is for representing general knowledge about the world, and for performing various kinds of memory retrievals and inferences. Note that the non-action-centered subsystem is under the control of the action-centered subsystem (through its actions).

At the bottom level, "associative memory" networks encode non-action-centered implicit knowledge. Associations are formed by mapping an input to an output. The regular backpropagation learning algorithm can be used to establish such associations between pairs of inputs and outputs.

On the other hand, at the top level of the non-action-centered subsystem, a general knowledge store encodes explicit non-action-centered knowledge. In this network, chunks are specified through dimensional values. A node is set up in the top level to represent a chunk. The chunk node connects to its

[2] The detail of the bottom-up learning algorithm can be found in Sun et al (2001).

corresponding features (dimensional values) represented as individual nodes in the bottom level of the non-action-centered subsystem. Additionally, links between chunks encode explicit associations between pairs of chunks, known as associative rules. Explicit associative rules may be learned in a variety of ways (Sun 2003). In addition to applying associative rules, similarity-based reasoning may be automatically carried out through the interaction of the two levels of representation (Sun 2003).

As in the action-centered subsystem, top-down or bottom-up learning may take place in the non-action-centered subsystem, either to extract explicit knowledge in the top level from the implicit knowledge in the bottom level or to assimilate explicit knowledge of the top level into implicit knowledge in the bottom level.

4.2.3 The Motivational and the Meta-Cognitive Subsystem

The motivational subsystem (the MS) is concerned with drives and their interactions, which leads to actions. It is concerned with why an agent does what it does. Simply saying that an agent chooses actions to maximizes gains, rewards, or payoffs leaves open the question of what determines these things.

Dual motivational representations are in place in CLARION, based on relevant psychological evidence (Sun 2003). The explicit goals (such as "finding food") of an agent (which is tied to the working of the action-centered subsystem) may be generated based on internal drive states (for example, "being hungry"). (See Sun 2003 for details.)

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". While 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 meta-cognitive subsystem (the MCS) is closely tied to the motivational subsystem. The meta-cognitive subsystem monitors, controls, and regulates cognitive processes for the sake of improving cognitive performance. Control and regulation may be in the forms of setting goals for the action-centered subsystem, setting essential parameters of the action-centered subsystem and the non-action-centered subsystem, interrupting and changing on-going processes in the action-centered subsystem and the non-action-centered subsystem, and so on. Control and regulation can also be carried out through setting reinforcement functions for the action-centered subsystem. All of the above can be done on the basis of drive states and goals in the motivational subsystem. The meta-cognitive subsystem is also made up of two levels: the top level (explicit) and the bottom level (implicit).

5 Applications to Cognitive Social Simulation

Three case studies of applying CLARION to social simulation will be sketched below, ranging from a small-scale organizational simulation, to a simulation of the larger-scale phenomenon of academic publishing, and further on to even larger-scale simulations of tribal societies.

One application of CLARION to cognitive social simulation is in understanding organizational decision making and the interaction between organizational structures and cognitive factors in affecting organizational performance (Sun and Naveh 2004).

In terms of organizational structures, there are two major types: (1) teams, in which agents act autonomously, individual decisions are treated as votes, and the organizational decision is the majority decision; and (2) hierarchies, which are characterized by agents organized in a chain of command, such that information is passed from subordinates to superiors, and the decision of a superior is based solely on the recommendations of his/her subordinates. In addition, organizations are distinguished by the structure of information accessible by each agent. Two varieties of information access are: (1) distributed access, in which each agent sees a different subset of attributes (no two agents see the same subset of attributes), and (2) blocked access, in which several agents see exactly the same subset of attributes.

Several simulation models were considered in Carley et al (1998). The human experiments by Carley et al (1998) were done in a 2 x 2 fashion (organization x information access). In addition, the human data for the experiment were compared to the results of the four models (Carley et al 1998). [3] See Figure 2.

In their work, the agent models used were very simple, and the "intelligence" level in these models was low. Moreover, learning in these simulations was rudimentary: there was no complex learning process as one might ob-

[3] 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, and Radar-Soar was a (somewhat) cognitive model built in Soar, which was based on explicit, elaborate search in problem spaces.

Agent/Org.	Team(B)	Team(D)	Hierarchy(B)	Hierarchy(D)
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

Figure 2: Human and simulation data for the organizational decision task. D indicates distributed information access, while B indicates blocked information access. All numbers are percent correct.

Agent/Org.	Team(B)	Team(D)	Hierarchy(B)	Hierarchy(D)
Human	50.0	56.7	46.7	55.0
CLARION	53.2	59.3	45.0	49.4

Figure 3: Simulation data for agents running for 3,000 cycles. The human data from Carley et al (1998) are reproduced here. Performance of CLARION is computed as percent correct over the last 1,000 cycles.

serve in humans. With these shortcomings in mind, it was worthwhile to undertake a simulation that involved more complex agent models that more accurately captured human performance. Moreover, with the use of more cognitively realistic agent models, one might investigate individually the importance of different cognitive capacities and process details in affecting organizational performance (see Sun and Naveh 2004).

Hence, a simulation with CLARION used for modeling individual agents in an organization was conducted. The results (see Figure 3) closely accord with the patterns of the human data, with teams outperforming hierarchal structures, and distributed access proving superior to blocked access. Also, as in humans, performance is not grossly skewed towards one condition or the other, but is roughly comparable across all conditions, unlike some of the simulation results from Carley et al (1998). The match with the human data is far better than in the simulations conducted in Carley et al (1998). The better match is due, at least in part, to a higher degree of cognitive realism in this simulation. See Sun and Naveh (2004) for further details, including observing the interesting effects of varying cognitive parameters.

The use of a cognitive architecture in this simulation enables the exploration of the interaction between cognitive factors/parameters and social/organizational structures. Because CLARION captures a wide range of generic cognitive processes and phenomena, its parameters are generic rather than task-specific Thus we have the opportunity of studying specific issues, such as organizational design, in the context of a general theory of cognition. We performed a detailed statistical analysis of simulation data resulting from varying cognitive parameters in a factorial design (Sun and Naveh 2004). 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. Through this analysis, the importance in social simulation of limiting one's conclusions to the specific cognitive (and other) context in which data were obtained, without over-generalizing the conclusions, was accentuated (Sun and Naveh 2004).

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

In sum, by using CLARION, Sun and Naveh (2004) have been able to more accurately capture organizational performance data and, moreover, to formulate deeper explanations for the results observed. In this way, CLARION might be used in the future to predict human performance in social/organizational settings and, furthermore, to help to improve collective performance by prescribing optimal or near-optimal cognitive abilities for individuals for specific collective tasks and/or organizational structures (Sun and Naveh 2004).

Another application of CLARION to cognitive social simulation is in capturing and explaining the essential process of academic publication and its relation to cognitive processes (Naveh and Sun 2006). Science develops in certain ways. In particular, it has been observed that the number of authors contributing a certain number of articles to a scientific journal follows a highly skewed distribution, corresponding to an inverse power law. In the case of scientific publication, the tendency of authorship to follow such a distribution was known as Lotka's law. Simon (1957) developed a simple stochastic process for approximating Lotka's law. 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 equal to a/i^k, where a is a constant of proportionality. Using Simon's work as a starting point, Gilbert (1997) attempted to model Lotka's law. He obtains his simulation data based

# of Papers	Actual	Simon's estimate	Gilbert's simulation	CLARION simulation
1	3991	4050	4066	3803
2	1059	1160	1175	1228
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

Figure 4: Number of authors contributing to *Chemical Abstracts*.

on 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 non-cognitive and interchangeable; it therefore neglected a host of cognitive phenomena that characterized scientific inquiry (e.g., learning, creativity, evolution of field expertise, etc.).

Using a more cognitively realistic model, one can address some of these omissions, as well as exploring other emergent properties of a cognitively based model and their correspondence to real-world phenomena. The results of the simulation based on CLARION are shown in Figures 4 and 5, along with results (reported by Simon 1957) for *Chemical Abstracts* and *Econometrica*, and estimates obtained from previous simulations by Simon (1957) and Gilbert (1997). The figures in the tables indicate number of authors contributing to each journal, by number of papers each has published.

The CLARION simulation data for the two journals could be fit to the power curve $f(i) = a/i^k$, resulting in an excellent match. The results of the curve fit are shown in Figure 6, along with correlation and error measures.

It is important to note that, in our simulation, the number of papers per author reflected the cognitive ability of an author, as opposed to being based on auxiliary assumptions such as those made by Gilbert (1997). This explains, in part, the slightly greater divergence of our results from the human data: Whereas Gilbert's simulation consisted of equations selected to match the human data, our approach relied on more detailed and lower-level mechanisms—namely, a cognitive agent model that was generic

# of Papers	Actual	Simon's estimate	Gilbert's simulation	CLARION simulation
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

Figure 5: Number of authors contributing to *Econometrica*.

Journal	a	k	Pearson R	R-square	RMSE
CA	3806	1.63	0.999	0.998	37.62
E	418	1.64	0.999	0.999	4.15

Figure 6: Results of fitting CLARION data to power curves. CA stands for Chemical Abstracts and E stands for Econometrica.

rather than task-specific. The result of the CLARION based simulation was therefore emergent, and not a result of specific and direct attempts to match the human data. That is, we put more distance between mechanisms and outcomes, which made it harder to obtain a match with the human data. Thus, the fact that we were able to match the human data reasonably well showed the power of our cognitive architecture based approach.

Yet another application of CLARION to cognitive social simulation is in simulating the survival strategies used by tribal societies under various environmental conditions (this work is yet unpublished, although the technical details of the work, which cannot be covered here, will be published later separately). In the simulation, the world was made up of a two-dimensional grid. Food items and agents were randomly distributed. There were the harsh, medium, and benign conditions, which were distinguished by the agent-to-food ratios. Agents were of a limited life span, which varied from individual to individual, depending on the energy level of an agent. 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 is said to adopt an individual survival strategy. However, in some other groups, resources may be transferred from one individual to another. A group in which there is transfer of resources among agents is said to adopt a social survival strategy. For instance, the "central store" is a mechanism to which all the individuals in a group transfer (part of) their resources. The resources collected by the central store can be redistributed to the members of the group.

Different from previous simulations of tribal societies, in this simulation, agents were more cognitively realistic, and they were constructed out of a cognitive architecture, which captured a variety of cognitive processes in a psychologically realistic way (Sun 2002, 2003). Therefore, this simulation of 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 this simulation was exactly to investigate the interaction between social structures/processes and individual cognition (i.e., the micro-macro link).

First of all, relationships between various cognitive parameters and social variables were found in the simulation, which indicated that what social institutions and norms (such as survival strategies) were adopted might have something to do with cognitive abilities and cognitive tendencies of agents involved. This relation may be termed the social-cognitive dependency, which may have significant theoretical and empirical ramifications: There may be some forms of social systems (structures and institutions) that are suitable for certain cognitive characteristics while unsuitable for certain others. They may not be universally better or worse. It may in fact depend on a host of other factors, and cognitive factors in particular. Sun (2006) contains a fairly substantial discussion of the close relationship between cognitive and social processes in general, and advocates the exploration of cognitive principles of sociocultural processes.

In the reverse direction, it was found that some cognitive attributes might have been selected (through natural evolution) to work with certain social systems and cultural environments, which may be termed the cognitive-social dependency. In this regard, we may explore sociocultural principles of cognition, the opposite of cognitive principles of sociocultural processes as mentioned earlier.

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

Together, these types of dependencies form a complex dynamic system—a system of inter-woven dependencies and interactions. In such a system, it is important to understand not just direct effects of dependencies but also indirect effects that are not obviously related to their causes but are often crucial for discerning the functional structures of the system.

In summary, it has been shown through the simulation that, in the context of different social survival strategies and different physical environments, cognition matters. It determines, for instance, which survival strategy and other social variables are appropriate under what cognitive conditions. Several hypotheses in this regard were generated in the process of simulation, such as, among others, the hypothesis that which survival strategy was the best was dependent on the explicitness of cognition of the population. Even though only very simply representations of sociocultural processes were involved in this work, given the cognitive architecture used, significant effects of various interactions were nevertheless found.

6 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. The two fields can be profitably integrated. This is an important opportunity and challenge.

As has been argued before, social processes ultimately rest on the choices and decisions of individuals, and thus understanding the mechanisms of individual cognition can lead to better theories describing the behavior of aggregates of individuals. Although most agent models in social simulation have been extremely simple, a more realistic cognitive agent model, incorporating realistic tendencies, inclinations, and capabilities of individual cognitive agents can serve as a more realistic basis for understanding the interaction of individuals. Although some cognitive details may ultimately prove to be irrelevant, this cannot be determined *a priori*, and thus simulations are useful in determining which aspects of cognition can be safely abstracted away.

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

The most fundamental challenge in this regard is to develop better ways of conducting detailed social simulation based on cognitive architectures as basic building blocks. This is not an easy task. Although some initial work has been done (e.g., Sun and Naveh 2004, Naveh and Sun 2006, Sun 2006), much more work is needed.

One specific challenge is how to enhance cognitive architectures for the purpose of accounting 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, "motive", "obligation", or "norm") are needed in cognitive modeling of multi-agent interaction? See, for example, Sun (2006) for further discussions.

There is also the challenge of computational complexity and thus scalability that needs to be addressed. Social simulation could involve a large number of agents, up to thousands. Computational complexity is thus already high, even without involving cognitive architectures as agent models. To incorporate cognitive architectures into social simulation, one has to deal with a great deal of added complexity. Thus, scalability is a significant issue.

Finally, it should be noted that 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 mentioned above, complexity may be an issue that prevents wider use of detailed cognitive models in social simulation.

Note also that in social simulation, in many cases, it might be necessary to capture social processes, social institutions, and social mechanisms more extensively and more directly in models, although this aspect is not dealt with here.

7 Concluding Remarks

We can expect that cognitive social simulation with cognitive architectures will have a profound impact both on cognitive science and on social simulation. Such impact may be in the form of better understanding the role of cognition in social interaction and in the form of better understanding the role of sociality in cognitive processes.

The field of cognitive social simulation with the use of cognitive architectures should be considered an essential aspect of social computing, and an important research direction in this emerging field. Correspondingly, a significant amount of collective research effort should be put into it.

Acknowledgments

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

References

K. Carley, M. Prietula, and Z. Lin, (1998). Design versus cognition: The interaction of agent cognition and organizational design on organizational performance. *Journal of Artificial Societies and Social Simulation*, 1 (3).

N. Gilbert, (1997), A simulation of the structure of academic science. *Sociological Research Online*, 2(2). Available online at http://www.socresonline.org.uk/socresonline/2/2/3.html.

I. Naveh and R. Sun, (2006). A cognitively based simulation of academic science *Computational and Mathematical Organization Theory*, 12, 313-337.

A. Newell, (1990). *Unified Theories of Cognition.* Harvard University Press, Cambridge, MA.

H. Simon, (1957), *Models of Man, Social and Rational.* Wiley, NY.

R. Sun, (2002). *Duality of the Mind.* Lawrence Erlbaum Associates, Mahwah, NJ.

R. Sun, (2003). *A Tutorial on CLARION.* Technical report, Cognitive Science Department, Rensselaer Polytechnic Institute.
http://www.cogsci.rpi.edu/∼rsun/sun.tutorial.pdf

R. Sun (ed.), (2006). *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation.* Cambridge University Press, New York.

R. Sun, E. Merrill, and T. Peterson, (2001). From implicit skills to explicit knowledge: A bottom-up model of skill learning. *Cognitive Science.* Vol.25, No.2, 203-244.

R. Sun and I. Naveh, (2004). Simulating organizational decision-making using a cognitively realistic agent model. *Journal of Artificial Societies and Social Simulation*, 7 (3). http://jasss.soc.surrey.ac.uk/7/3/5.html