

Simulating organizational decision-making using a cognitively realistic agent model

Ron Sun
Isaac Naveh

Abstract

Most of the work in agent-based social simulation has assumed highly simplified agent models, with little attention being paid to the details of individual cognition. Here, in an effort to counteract that trend, we substitute a realistic cognitive agent model (CLARION) for the simpler models previously used in an organizational design task. On that basis, an exploration is made of the interaction between the cognitive parameters that govern individual agents, the placement of agents in different organizational structures, and the performance of the organization. It is suggested that the two disciplines, cognitive modeling and social simulation, which have so far been pursued in relative isolation from each other, can be profitably integrated.

Keywords

Cognition, cognitive architecture, cognitive modeling, classification decision making

1 Introduction

1.1 Abstraction vs. Cognitive Realism in Social Simulation

Computational models of cognitive agents that incorporate a wide range of cognitive functionalities (such as various types of memory/representation, various modes of learning, and sensory motor capabilities) have been developed in both AI and cognitive science (e.g., Anderson and Lebiere 1998; Sun 2002). In cognitive science, they are often known as cognitive architectures. Recent developments in computational modeling of cognitive architectures provide new avenues for precisely specifying complex cognitive processes in tangible ways (Anderson and Lebiere 1998).

Artificial intelligence started out with the goal of designing functioning intelligent agents. Although the enormous difficulty of the task has forced a shift in focus to highly restricted domains of intelligence, some researchers have persisted in putting the pieces together with the goal of designing autonomous agents. At the same time, there is a growing interest in multi-agent interactions that address issues of coordination and cooperation among cognitive agents.

In spite of this, however, most of the work in social simulation still assumes very rudimentary cognition on the part of the agents (e.g., Cecconi and Parisi 1998; Gilbert 1997). At the same time, while researchers in cognitive science have devoted considerable attention to the workings of individual cognition (e.g., Anderson 1983; Klahr et al 1987; Rumelhart and McClelland 1986; Sun 2002), sociocultural processes and their relations to individual cognition have generally not been studied by cognitive scientists (with some notable exceptions; e.g., Suchman 1987; Hutchins 1995). The cleavage between the two fields is seen in the different journals dedicated to each area (e.g., *JASSS*, *Emergence* and *CMOT* for social simulations, versus *Cognitive Science*, *Cognitive Systems Research*, and *Cognitive Science Quarterly* for cognitive modeling), as well as in the different conferences (e.g., the Agent 200x series in Chicago and the International Conferences on Social Simulation, versus the International Conference on Cognitive Modeling), organizations (e.g., the North American Association for Computational and Organizational Science and the European Social Simulation Association versus the Cognitive Science Society), and scant overlap of authors in these areas. Most of the commonly available social simulation tools (e.g., Swarm and RePast) likewise embody very simple agent models.

However, we believe that social simulation needs cognitive science, because better models of individual cognition can lead us to a better understanding of aggregate processes involving multi-agent interaction (Moss 1999; Castelfranchi 2001). Cognitive models that incorporate realistic tendencies, biases and capacities of individual cognitive agents (Boyer and Ramble 2001) can serve as a more realistic basis for understanding multi-agent interaction. This point has been made before in the context of social simulation (e.g., Edmonds and Moss 2001) as well as in the context of cognitive realism of game theory

(Kahan and Rapaport 1984) and of multi-agent systems (Sun 2001).

Conversely, by combining social simulation and cognitive modeling¹, we can learn more about individual cognition. Traditional approaches to computational modeling of agents have largely ignored the vital role of socially acquired and disseminated knowledge. By studying cognitive models in a social context, rather than in isolation, we might shed more light on the sociocultural aspects of cognition and on the cognitive processes involved in multi-agent interaction.

Reflecting a growing interest in cognitive realism in social simulation, cognitive agent models have recently been deployed in a variety of social-cognitive systems. Some examples include game playing (West et al, 2003), simulated warfare in urban environments (Best and Lebiere 2003), and design of a virtual 3D world (Maher et al 2003).

1.2 Advantages of a Cognitively Realistic Approach

As noted above, research on social simulation has mostly dealt with simplified versions of social phenomena, involving much simplified agent models (e.g., Gilbert and Doran 1994; Levy 1992). Such agents are clearly not cognitively realistic, and thus may result in important cognition-related insights being left by the wayside. Social interaction is, after all, the result of individual cognition (which includes instincts, routines, and patterned behavior, as well as complex conceptual processes). Therefore, the mechanisms underlying individual cognition cannot be ignored in studying multi-agent interaction. At least, the implications of these mechanisms should be understood before they are abstracted away. An agent model that is sufficiently complex to capture individual cognitive processes should be adopted.

By using cognitively realistic agents in social simulation, we can provide explanations of observed social phenomena based on individual cognitive processes. This allows us to do away with assumptions that are not cognitively grounded. Often, in previous simulations, rather arbitrary assumptions were made, simply because they were important for generating simulations that matched observed data. We instead make assumptions at a lower level. This allows us to put more distance between assumptions and outcomes, and thereby provide deeper explanations.

In this paper, we first describe a more realistic cognitive model, named CLARION, which captures the distinction between explicit and implicit learning. This model has been extensively described in a series of papers (e.g., Sun 1997; Sun 2002). We then apply this model to the problem of organizational design as presented by Carley et al (1998). The idea here is to substitute more sophisticated agents, based on CLARION, for the (mostly) simple agents used in the original task (Carley et al 1998).

This substitution has much in common with the notion of “docking,” or the alignment of models. Docking, as proposed by Axtell and his colleagues, consists of testing whether two models of the same phenomenon can be made to yield equivalent results. This allows both models to be independently validated (Axtell et al 1996). Here, a similar comparison of different models is undertaken, namely between the simpler agent models considered previously (Carley et al 1998) and the more cognitively realistic model CLARION. However, in contrast to docking as envisioned by Axtell et al, the purpose of this study is not to establish a strict equivalence between the models, a task which at any rate is difficult when models are built for different purposes and according to different theoretical conceptions (Takadama et al 2003). Rather, we want to expose the differences between the models. By varying the cognitive factors of our model in a particular organizational task, and comparing the results to those of previous, less cognitively realistic simulations, we can arrive at a better picture of the role of individual cognition in organizational learning.

Finally, previous experiments and simulations left open the question of whether their results were generic or tied specifically to particular settings of the experiments/simulations or to particular assumptions regarding cognitive parameters. Our work is designed to explore a wider range of possibilities and ascertain some answers to the above question.

2 The Model

2.1 Explicit vs. Implicit Learning

The role of implicit learning in skill acquisition has been widely recognized in recent years (e.g., Reber 1989; Stanley et al 1989; Anderson 1993; Seger 1994; Proctor and Dutta 1995; Stadler and Frensch 1998). Although explicit and implicit learning have both been actively studied, the question of the interaction between these two processes has rarely been broached. However, despite the lack of study of this interaction, it has recently become evident (e.g., in Reber 1989; Seger

1994) that rarely, if ever, is only one of type of learning engaged. Our review of experimental data shows that although one can manipulate conditions such that one or the other type of learning is emphasized, both types of learning are nonetheless usually present.

To model the interaction between these two types of learning, the cognitive architecture CLARION was developed (Sun and Peterson 1998; Sun et al 2001), which captures the combination of explicit and implicit learning. CLARION mostly learns in a bottom-up fashion, by extracting explicit knowledge from implicit knowledge. Such processes have also been observed in human subjects (e.g., Willingham et al 1989; Stanley et al 1989; Mandler 1992).

A major design goal for CLARION was to have a set of tunable parameters that directly correspond to aspects of cognition. This is in contrast to some cognitive models in which performance depends on a set of variables that are mathematically motivated (and hence do not translate into mechanisms of individual cognition). We have avoided this, so as to be able to manipulate the parameters of the model and observe the effect on performance *as a function of cognition*.

2.2 A Summary of the CLARION Model

We are now ready to describe our model in greater detail. CLARION is an integrative cognitive architecture with a dual representational structure (Sun 1997; Sun et al 1998; Sun et al 2001; Sun 2002). It consists of two levels: a top level that captures explicit learning, and a bottom level that captures implicit learning (see Figure 1).

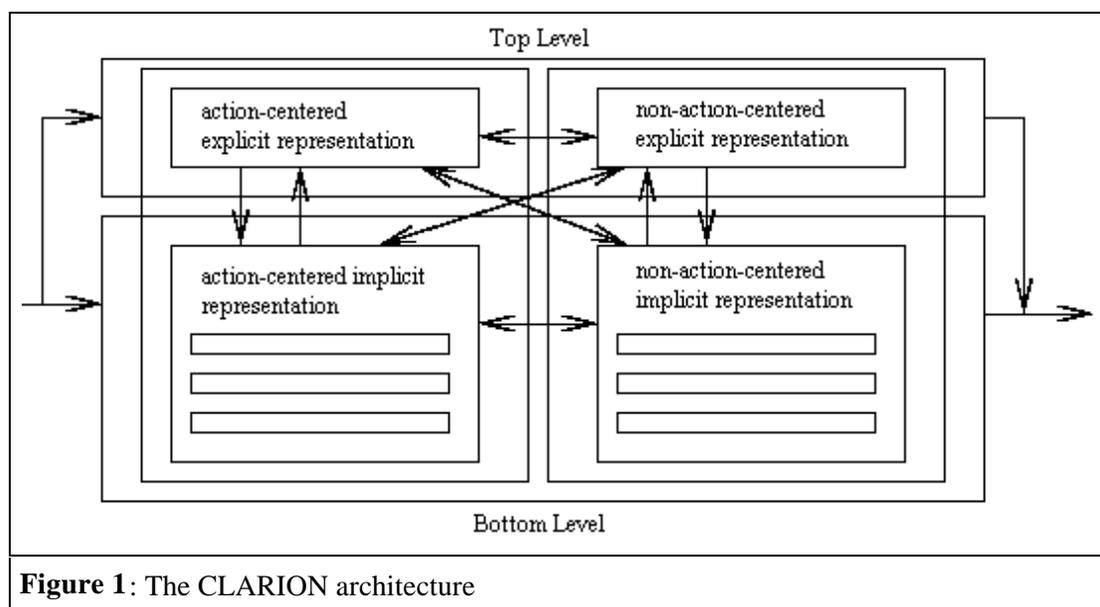


Figure 1: The CLARION architecture

At the bottom level, the inaccessible nature of implicit learning is captured by subsymbolic distributed representations provided by a connectionist network. This is because representational units in a distributed environment are capable of performing tasks but are generally not individually meaningful (Sun 1995). Learning at the bottom level proceeds in trial-and-error fashion, guided by reinforcement learning (i.e., Q-learning) implemented in backpropagation neural networks (Sun and Peterson 1998). Other representations of low-level cognition are also possible, including Bayesian networks (Jensen 1996), hidden Markov models (Rabiner 1989), and models with simulated annealing, the latter of which have been used in an organizational task similar to the one considered in this paper (Carley and Svoboda 1996). However, neural networks are more suitable for representing the low-level structures involved in implicit learning (in contrast to other methods, which strive to approximate cognitive processes at a slightly higher level).

At the top level, explicit learning is captured by a symbolic representation, in which each element is discrete and has a clear meaning. This accords well with the directly accessible nature of explicit knowledge (Smolensky 1988, Sun 1995). Learning at the top level proceeds by first constructing a rule that corresponds to a “good” decision made by the bottom level, and then refining it (by generalizing or specializing it), mainly through the use of an “information gain” measure that compares the success ratio of various modifications of the current rule.

A high-level pseudo-code algorithm that describes the action-centered subsystem of CLARION is as follows:

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

At the bottom level, a Q-value is an evaluation of the “quality” of an action in a given state: $Q(x, a)$ indicates how desirable action a is in state x . Actions can be selected based on Q-values. To acquire the Q-values, Q-learning, a reinforcement learning algorithm (Watkins 1989), is used.

We learn the simplified Q function as follows:

$$\Delta Q(x, a) = \alpha(r + \gamma \max_b Q(y, b) - Q(x, a)) = \alpha(r - Q(x, a))$$

where x is the current state, a is one of the actions, r is the immediate feedback, and $\gamma \max_b Q(y, b)$ is set to zero for the organizational design task that was tackled in this paper, because we rely on immediate feedback in this particular task (details below). $\Delta Q(x, a)$ provides the error signal needed by the backpropagation algorithm and then backpropagation takes place. That is, learning is based on minimizing the following error at each step:

$$err_i = \begin{cases} r - Q(x, a) & \text{if } a_i = a \\ 0 & \text{otherwise} \end{cases}$$

where i is the index for an output node representing the action a_i . Based on the above error measure, the backpropagation algorithm is applied to adjust internal weights of the network.

In the top level, declarative knowledge is captured in a simple prepositional rule form. We devised an algorithm for learning declarative knowledge (rules) using information from the bottom level (the *Rule-Extraction-Refinement*, or RER, algorithm). The basic idea is as follows: if an action decided by the bottom level is successful then the agent extracts a rule (with its action corresponding to that selected by the bottom level and with its conditions corresponding to the current state), and adds the rule to the top-level network. Then, in subsequent interactions with the world, the agent refines the extracted rule by considering the outcome of applying the rule: if the outcome is successful, the agent may try to generalize the conditions of the rule to make it more universal. If the outcome is unsuccessful, the agent may try to specialize the rule, by narrowing its conditions down and making them exclusive of the current state.

The information gain (IG) measure of a rule is computed (in this organizational design task) based on the immediate feedback at every step when the rule is applied. The inequality, $r > threshold_{RER}$ determines the positivity/negativity of a step and the rule matching this step (where r is the feedback received by an agent). The positivity threshold (denoted $threshold_{RER}$ above) corresponds to whether or not an action is perceived by the agent as being reasonably good.

Based on the positivity of a step, PM (positive match) and NM (negative match) counts of the matching rules are updated. IG is calculated based on PM and NM:

$$IG(A, B) = \log_2 \frac{PM_a(A) + c1}{PM_a(A) + NM_a(A) + c2} - \log_2 \frac{PM_a(B) + c1}{PM_a(B) + NM_a(B) + c2}$$

where A and B are two different rule conditions that lead to the same action a , and $c1$ and $c2$ are two constants representing the prior (by default, $c1 = 1$, $c2 = 2$). Essentially, the measure compares the percentages of positive matches under different conditions A and B.

The generalization operator is based on the IG measure. Generalization amounts to adding an additional value to one input dimension in the condition of a rule, so that the rule will have more opportunities of matching input. For a rule to be generalized, the following must hold:

$$IG(C, all) > threshold_{GEN} \text{ and } \max_{C'} IG(C', C) \geq 0$$

where C is the current condition of a rule (matching the current state and action), all refers to the corresponding match-all rule (with the same action as specified by the original rule but an input condition that matches any state), and C' is a modified condition equal to C plus one input value. If the above holds, the new rule will have the condition C' with the highest IG measure. The generalization threshold (denoted $threshold_{GEN}$ above) determines how readily an agent will generalize a rule.

The specialization operator works in analogous fashion, except that an input dimension is discarded, rather than being added. Likewise, a rule must perform *worse* than the match-all rule, rather than better, to be considered for specialization. This process is described in greater detail elsewhere (Sun et al, 2001). Due to running-time considerations, the specialization threshold is held constant in all simulations reported in this paper.

To avoid the proliferation of useless rules, a RER *density* measure is in place. A density of $1/x$ means that a rule must be invoked once per x steps to avoid deletion due to disuse. This corresponds to the agent's memory for rules, necessitating that a rule come up every once in a while in order to be retained.

To integrate the methods from the two levels, a number of methods may be used. Here, we choose levels stochastically, using a probability of selecting each level. Other selection methods are available as well (see Sun et al 2001).

When we use the outcome from the bottom level, we use a stochastic process based on the Boltzmann distribution of Q values for selecting an action:

$$p(a | x) = \frac{e^{Q(x,a)/t}}{\sum_i e^{Q(x,a_i)/t}}$$

where x is the current state, a is an action, and t controls the degree of randomness (temperature) of the process. (This method is also known as Luce's choice axiom (Watkins 1989). It is found to match psychological data in many domains.)

At each level of the model, there may be multiple modules, both action-centered modules and non-action-centered modules (Schacter 1990). In the current study, we focus only on the *action-centered subsystem* (ACS). There are also other components, such as working memory, goal structure, and so on.

2.3 Previously Simulated Tasks

Many well-known cognitive tasks performed by individual cognitive agents have been simulated using CLARION. The tasks include serial reaction time (SRT) tasks, process control (PC) tasks, categorical inference (CI) tasks, alphabetical arithmetic tasks (AA), and the Tower of Hanoi (TOH) task (Sun 2002). Among them, SRT and PC are typical implicit learning tasks (involving mostly reactive routines), while TOH and AA are high-level cognitive acquisition tasks (with explicit processes). In addition, we have simulated a complex minefield navigation (MN) task, which involves complex sequential decision-making (Sun and Peterson 1998; Sun et al 2001). We are now in a good position to extend the modeling effort to capture social cognition processes.

2.4 Implementation

Both CLARION and the current simulation have been implemented as a set of Java packages. The code is available upon request. For more information, see the CLARION web page at <http://www.cogsci.rpi.edu/~rsun/clarion.html>.

3 Organizational Design

The task considered in this paper is taken from organizational research. Research on organizational performance has usually

focused either on an organization's design (i.e., its structure) or on the cognition of its members (i.e., how smart/capable individuals in the organization are). However, the interaction of these two factors -- cognition and structure -- is rarely studied. Carley et al (1998) introduce a classification task involving different types of organizational structures and agents. By varying agent type and structure together, they are able to study how these factors interact with each other. Here, we will build on that research and extend it, with the aim of studying the interaction of cognition and design in the context of a more realistic cognitive architecture (i.e., CLARION).

3.1 Task

A typical task faced by organizations is classification decision making. In a classification task, agents gather information about problems, classify them, and then make further decisions based on the classification. For instance, an organization may classify a company as financially promising or unpromising, and on that basis decide to approve or not approve a loan. In this case, the task is to determine whether a blip on a screen is a hostile aircraft, a flock of geese, or a civilian aircraft (Carley et al 1998). Hence, this is a ternary choice task. It has been used before in studying organizational design (e.g., in Carley and Prietula 1992; Ye and Carley 1995; Carley and Lin 1995).

In each case, there is a single object in the airspace. The object has nine different attributes, each of which can take on one of three possible values (e.g., its speed can be low, medium, or high). An organization must determine the status of an observed object: whether it is friendly, neutral or hostile. There are a total of 19,683 possible objects, and 100 problems are chosen randomly (without replacement) from this set.

The true status of an object is determinable by adding up all nine characteristic values. If the sum is less than 17, then it is friendly; if the sum is greater than 19, it is hostile; otherwise, it is neutral.

No one single agent has access to all the information necessary to make a choice. Decisions are made by integrating separate decisions made by different agents, each of which is based on a different set of information. Each organization is assumed to have sufficient personnel to observe all the necessary information (in a distributed way).

In terms of organizational structures, there are two archetypal structures of interest: (1) teams, in which decision makers act autonomously, individual decisions are treated as votes, and the organization 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 subordinates (Carley 1992). In this task, only a two-level hierarchy with nine subordinates and one superior is considered.

In addition, organizations are distinguished by the structure of information accessible by each agent. There are two varieties of information access: (1) distributed access, in which each agent sees a different subset of three attributes (no two agents see the same set of three attributes), and (2) blocked access, in which three agents see exactly the same attributes. In both cases, each attribute is accessible to three agents.

Several simulation models were considered in the study (Carley et al 1998). Among them, CORP-ELM produces the most probable classification based on an agent's own experience, CORP-P-ELM stochastically produces a classification in accordance with the estimate of the probability of each classification based on the agent's own experience, CORP-SOP follows organizationally prescribed standard operating procedure (which involves summing up the values of the attributes available to an agent) and thus is not adaptive, and Radar-Soar is a (somewhat) cognitive model built in SOAR, which is based on search through problem spaces (Rosenbloom et al 1991).

3.2 Previous Experimental Results

Experiments (in Carley et al 1998) were done in a 2 x 2 fashion (organization x information access). In addition, the human data were compared to the simulations by the aforementioned four models (Carley et al 1998). The data appears to show that agent cognition interacts with organizational design. The human data and the simulation results from this study (Carley et al 1998) were as shown in Table 1.

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

Table 1. Human and simulation data for the organizational design task. D indicates distributed information access, while B indicates blocked information access. All numbers are percent correct.

The human data showed that humans generally performed better in team situations, especially when distributed information access was in place. Moreover, distributed information access was generally better than blocked information access. The worst performance occurred when hierarchical organizational structure and blocked information access were used in conjunction.

It also suggested that which type of organizational design exhibit the highest performance depends on the type of agent. For example, human subjects performed best as a team with distributed information access, while Radar-SOAR and CORP-ELM performed the best in a team with blocked information access. Relatedly, increasing general intelligence, or increasing the adaptiveness of agents, tended to degrade the performance of hierarchical organization. With a non-adaptive agent such as CORP-SOP, there was no difference between the two different organization types.

The above results are interesting because they brought up the issue of the interaction between organizational type and intelligence level. However, the agent models that were used were, to a large extent, fairly simplistic. Therefore, the intelligence level in these models was rather low (including, to a large extent, the SOAR model, which essentially encoded a set of simple rules). Moreover, learning in these simulations was rudimentary: there was no complex learning process as one might observe in human cognition.

With these shortcomings in mind, it is worthwhile to undertake a simulation that involves more complex, more comprehensive agent models that more accurately capture detailed cognitive processes in more realistic ways. Moreover, with the use of more cognitively realistic agent models, we may investigate individually the importance of different cognitive capacities and process details in affecting the performance. In CLARION, we can easily vary parameters and options that correspond to different cognitive capacities and processes and test the resulting performance.

Below, we present four simulations involving the CLARION model. In the first simulation, we use the aforementioned radar task (Carley et al 1998) but substitute a different cognitive model for the one used previously. The second simulation extends the duration of training given to the agents. In the third simulation, we vary a wide range of cognitive parameters of the model in a factorial design. Here, we are interested in observing the interaction of different cognitive factors with organizational design. Finally, in the fourth simulation, we move away from uniform-agent models and investigate organizations in which agents differ in their capabilities and limitations.

4 Simulation I: Matching Human Data

We begin with a simple “docking” simulation (in an abstract sense) -- that is, we use the same setup as in the original study (Carley et al 1998; see Section 3.1), but substitute CLARION-based agents for the simpler agents used previously. Our aim here is to gauge the effect of organization and information access on performance (as in the original study), but in the context of the more cognitively realistic model CLARION.

4.1 Simulation Setup

There are two organizational forms: team and hierarchy. Under the team condition, the input to each agent consists of three of the aircraft’s attributes, selected according to a blocked or distributed information access scheme. The condition where a

hierarchy is used is similar to the team condition, except that a supervisor agent is added. The input to the supervisor corresponds to the outputs of all nine subordinates.

The actions of each agent are determined by CLARION. At the top level, RER rule learning is used to extract rules. At the bottom level, each agent has a single network that is trained, over time, to respond correctly. The network receives an external feedback of 0 or 1 after each step, depending on whether the target was correctly classified. Due to the availability of immediate feedback in this task, simplified Q-learning is used.

All agents run under a single (uniform) set of cognitive parameters², regardless of their role in the organization (later we shall vary the parameters).

4.2 Results

The results of our simulation are shown in Table 2. 4,000 training cycles (each corresponding to a single problem, followed by a single decision by the entire organization) were included in each group. As can be seen, our results closely accord with the patterns of the human data, with teams outperforming hierarchies, and distributed access outperforming 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 better than in the simulations conducted in the original study (Carley et al 1998).

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

Table 2. Simulation data for agents running for 4,000 cycles. The human data from Carley et al (1998) are reproduced here for ease of comparison. Performance for CLARION is computed as percentage correct over the last 1,000 cycles.

To understand these results and their interpretation better, let us examine the curves that compare the learning process. As can be seen in Figure 2, a team organization, using blocked access, quickly achieves a high level of performance. However, thereafter there are very few gains. By contrast, a team using distributed access starts out slowly but eventually achieves a performance nearly as high as that in the distributed condition.

Under the hierarchical conditions, learning is both slower and more erratic. When access is distributed, performance dips in the first few hundred cycles, but afterwards improves steadily. The slowness of the learning process should not surprise us, since two layers of agents are being trained (rather than one), with the output of the upper layer depending on that of the lower layer. In addition, the higher input dimensionality of the supervisor (nine inputs vs. three inputs for a subordinate) greatly increases the complexity of the task, leading to a longer training time for the network and to a slower process of rule refinement. This is analogous to the case of humans, where input dimensionality is known to be one of the chief determinants of task complexity (e.g., Berry and Broadbent 1988). When a hierarchy is combined with blocked access, performance is even worse, with very little learning taking place.

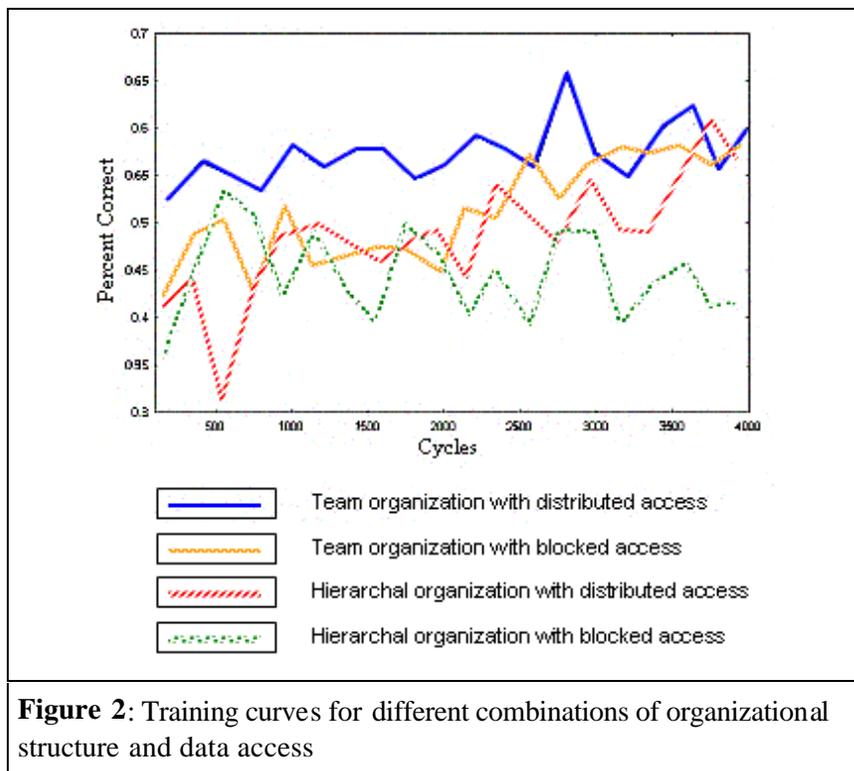


Figure 2: Training curves for different combinations of organizational structure and data access

Throughout the training process, rules were encoded at the top level. A sample rule (randomly selected) extracted by the supervisor in a hierarchal condition was as follows:

$$(3) (3) (3) (1, 2, 3) (3) (3) (3) (3) (3) \rightarrow 3$$

The rule should be read as follows: if input #4 is equal to 1, 2 or 3, and the other inputs are equal to 3, then select action 3 (hostile aircraft).

5 Simulation II: Extending the Simulation Temporally

So far, we have considered agents who trained for only 4,000 cycles. The results were interesting, because they were analogous to those of humans. However, the human data itself was arguably the result of limited training. In the original task, human subjects were presented with only 30 problems (Carley et al 1998). Therefore, it is interesting to see what will happen if we extend the length of the training. In particular, we are interested in knowing if the trends seen above (in section 4.2) will be preserved in the long run. It is extremely important that before we draw any conclusion about human performance, we understand the context and conditions under which data are obtained, and thereby avoid over-generalizing our conclusions (e.g., team vs. hierarchy, blocked vs. distributed; of Carley et al 1998).

Figures 3-6 show learning as it occurs over 20,000 (rather than 4,000) cycles. Previously, the best-performing condition was team organization with distributed information access. As can be seen in Figure 3, this condition continues to improve slowly after the first 4,000 cycles. However, it is overtaken by the team condition where blocked access is used (Figure 4). Thus, it seems that while teams benefit from a diversified (distributed) knowledge base in the initial phases of learning, a well-trained membership with redundant (blocked) knowledge performs better in the long run.

In the hierarchal condition, too, we can see either a reversal or disappearance of the initial trends. Hierarchies using distributed access (Figure 5) now show not only the best, but also the most stable (least variance) performance of any condition. Likewise, a hierarchy with blocked access (Figure 6), previously a weak performer, shows impressive gains in the long run. Thus, while hierarchies take longer to train, their performance is superior in the long run. In a hierarchy, a well-trained supervisor is capable of synthesizing multiple data points with greater sensitivity than a simple voting process. Likewise, the reduced individual variation in blocked access leads to less fluctuation in performance in the long run.

Figure 7 compares the trends discussed in this section and in section 4.2.

There is a serious lesson here: limited data can only allow us to draw limited conclusions -- only with regard to the specific

condition under which the data were obtained. There is a natural tendency for researchers to over-generalize their conclusions, which only be remedied by more extensive investigations. Given the high cost of human experiments, simulation has a large role to play in exploring alternatives and possibilities, especially social simulation coupled with cognitive architectures.

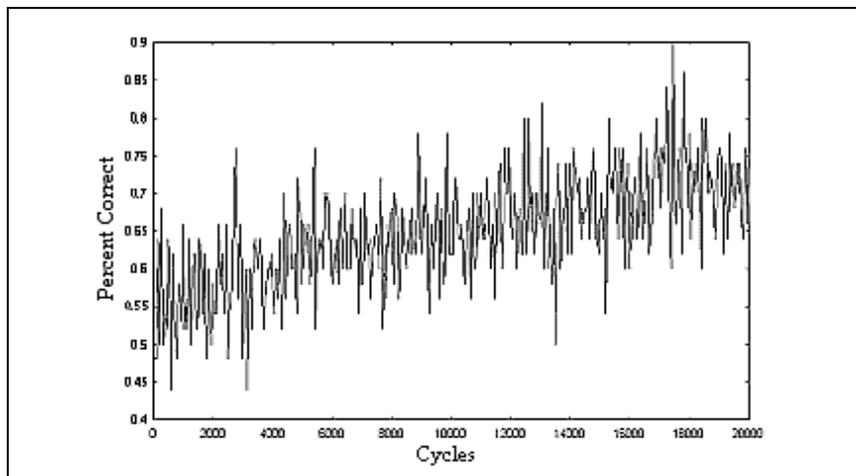


Figure 3: Training curve (team organization, distributed access)

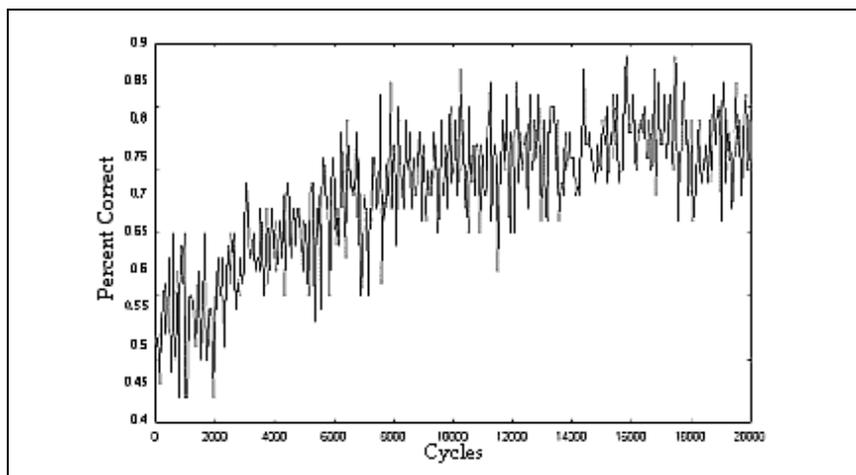


Figure 4: Training curve (team organization, blocked access)

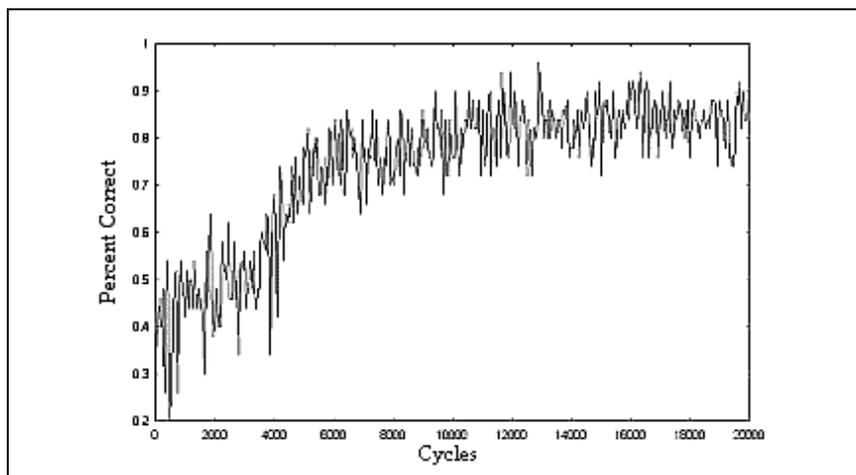


Figure 5: Training curve (hierarchal organization, distributed access)

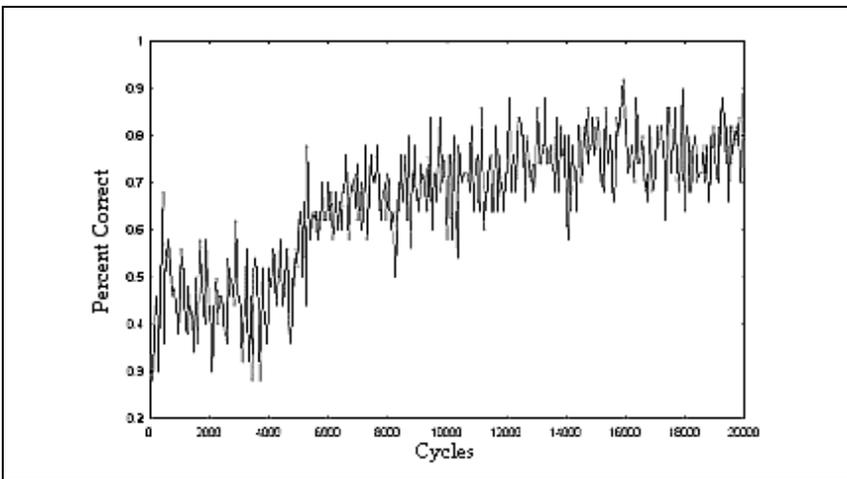


Figure 6: Training curve (hierarchal organization, blocked access)

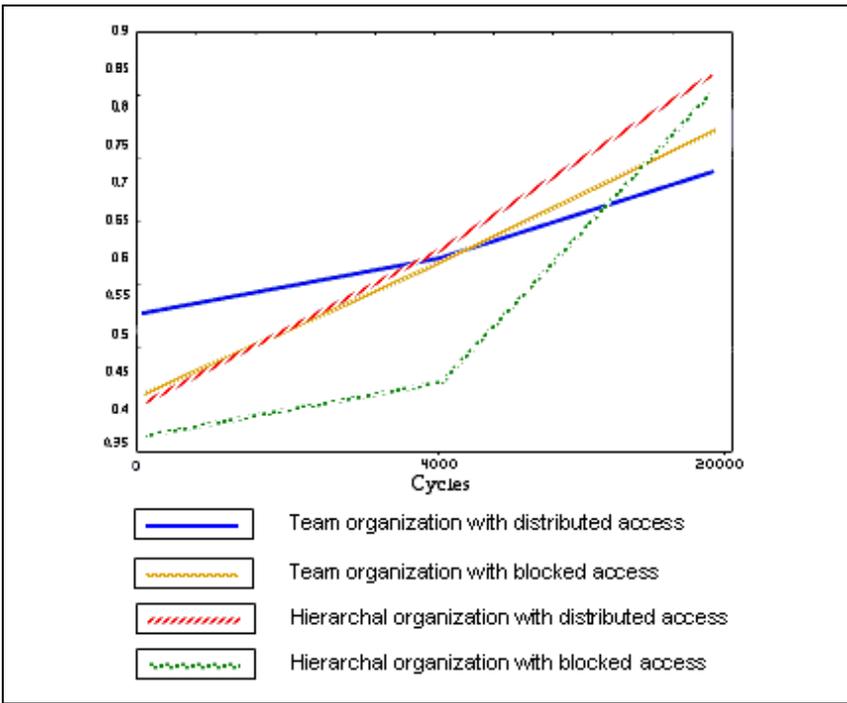


Figure 7: A comparison of performance under different combinations of structure and organization after 100, 4,000 and 20,000 training cycles.

6 Simulation III: Varying Cognitive Parameters

In the two preceding simulations, agents were run under a fixed set of cognitive parameters. Next, let us see what happens when we vary these parameters, analogous to varying the training length earlier. This again allows us to see the variability of results, and thus avoid over-generalization.

As mentioned above, the ability to vary different aspects of cognition is one feature that sets CLARION apart from many other models, especially specialized models that are devised to tackle a specific task. Because CLARION captures a wide range of 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.

In our third simulation, we vary a number of cognitive parameters and observe their effect on performance. Parameters are varied in a full factorial design, such that all combinations of all levels of each parameter are considered. This allows us to study both the influence of individual parameters and their interactions with each other.

6.1 Simulation Setup

Two different sets of parameters of CLARION are separately varied (in order to avoid the prohibitively high cost of varying all parameters simultaneously). These parameters were described in detail in Section 2.2. The first set of parameters consists of fundamental parameters of the model, including: (1) Reliance on the top versus the bottom level, expressed as a fixed probability of selecting each level. (2) Learning rate of the neural networks. (3) Temperature, or degree of randomness.

The second set consists of parameters related to RER and rule extraction, including: (1) RER positivity threshold, which must be exceeded for a rule to be considered “successful.” (2) RER density measure, which determines how often a rule must be invoked in order to be retained. (3) RER generalization threshold, which must be exceeded for a rule to be generalized.

The two sets of parameters above, along with information access and organization, are varied together in a factorial design. For each parameter, 2-3 different levels are tested, resulting in a $3 \times 2 \times 2 \times 2 \times 2$ (probability of using bottom level \times learning rate \times temperature \times organization \times information access) design for the first set of parameters, and a $2 \times 3 \times 2 \times 2 \times 2$ (RER positivity \times RER density \times RER generalization \times organization \times information access) design for the second set.

6.2 Results

Recall that we are interested in observing performance at both ends of the learning curve -- that is, both after a moderate amount of training (since results at that point corresponded closely to the human results) and after extensive training. Therefore, in all conditions of the variable-factor simulation, performance was measured both near the start of the simulation (after 4000 cycles) and at the end (after 20,000 cycles).

An ANOVA confirmed the effects of organization [$F(1, 24) = 30.28, p < 0.001, MSE = 0.05$] and information access [$F(1, 24) = 7.14, p < 0.05, MSE = 0.01$] to be significant. Moreover, the interaction of these two factors with length of training was significant [$F(1, 24) = 59.90, p < 0.001, MSE = 0.73$ for organization; $F(1, 24) = 3.43, p < 0.05, MSE = 0.01$ for information access]. These interactions, which can be seen in Figures 8-9, reflect the trends discussed above: the superiority of teams and distributed information access at the start of the learning process, and either the disappearance or reversal of these trends towards the end. This finding is important, because it shows that these trends persist across a wide variety of settings of cognitive parameters, and do not depend on any one setting of these parameters.

The effect of probability of using the top vs. the bottom level was likewise significant [$F(2, 24) = 11.73, p < 0.001, MSE = 0.02$]. More interestingly, however, its interaction with length of training period was significant as well [$F(2, 24) = 12.37, p < 0.001, MSE = 0.01$]. As can be seen in Figure 10, rule learning is very useful at the early stages of learning, when increased reliance on them tends to boost performance. However, by the 20,000th cycle, this effect disappears. This is because rules are crisp guidelines that are based on past success, and as such, they provide a useful anchor at the uncertain early stages of learning. However, by the end of the learning process, they become too coarse-grained to cover all possible contingencies, and thus are no more reliable than highly-trained networks. This corresponds to findings in human cognition, where there are indications that rule-based learning is widely used in the early stages of learning, but is later increasingly supplanted by similarity-based processes (Palmeri 1997; Smith and Minda 1998) and skilled performance (Anderson and Lebiere 1998). Such trends may partially explain the poor initial performance of hierarchies (see Section 4.2), since the high input dimensionality of a supervisor in a hierarchy slows down the process of rule acquisition (which is essential at the early stages of learning).

Predictably, the effect of learning rate was significant [$F(2, 24) = 32.47, p < 0.001, MSE = 0.07$], and indeed, the groups with the higher learning rate (0.5) outperformed the groups with the lower learning rate (0.25) by between 5-14% (averages are shown in Figure 11). However, there was no significant interaction between learning rate and organization or information access, which suggests that quicker learners do not differentially benefit from, say, being in a hierarchy versus a team. By the same token, the poorer performance of slower learners cannot be mitigated by recourse to a particular combination of organization and information access.

Let us now turn to the parameters related to RER rule extraction. As can be seen in Figure 12, it is unquestionably better to have a higher rule generalization threshold than a lower one (up to a point³). An ANOVA confirmed the significance of this effect [$F(1, 24) = 15.91, p < 0.001, MSE = 0.01$]. Thus, if one restricts the generalization of rules only to those rules that have proven relatively successful, the result is a higher-quality rule set, which leads to better performance in the long run.

Relatedly, while the effect of rule density on performance was insignificant, the interaction between density and generalization threshold was significant [by an ANOVA; $F(2, 24) = 2.93$; $p < 0.05$; $MSE = 0.01$]. As we can see in Figure 13, when rules are of relatively high quality (i.e., under a high generalization threshold) it is advisable to have more of them available (which is achievable by lowering the density). By contrast, when the average quality of rules is lower (i.e., under a low generalization threshold) it is advantageous to have a quicker forgetting process in place, as embodied by a high density parameter.

Finally, the interaction between generalization threshold and organization was significant at the start of the learning process [by an ANOVA; $F(1, 24) = 5.93$, $p < 0.05$, $MSE = 0.01$], but not at the end. This result (shown in Figure 14) is more difficult to interpret, but probably reflects the fact that hierarchies, at the start of the learning process, do not encode very good rules to begin with (due to the higher input dimensionality of the supervisor and the resulting learning difficulty). Thus, generalizing these rules, even incorrectly, causes relatively little further harm.

For the rest of the factors considered above (including temperature and RER positivity threshold), no statistically significant effects were found. Tables 3-4 summarize the results reported in this section.

This simulation confirmed an earlier observation -- namely, that which organizational structure (team vs. hierarchy) or information access scheme (distributed vs. blocked) is superior depends on the length of the training. It also showed that some cognitive parameters (e.g., learning rate) have a monolithic, across-the-board effect under all conditions, whereas in other cases, complex interactions of factors are at work. This illustrates, once again, the importance of limiting one's conclusions to the specific cognitive context in which they were obtained.

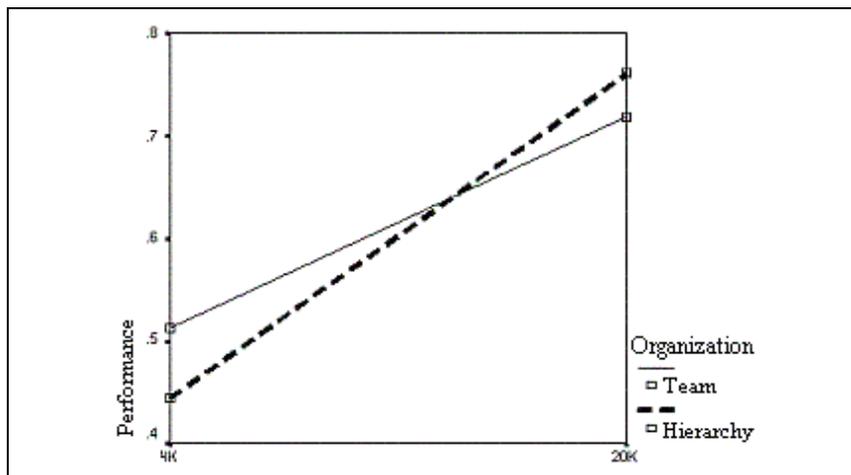


Figure 8: The effect of organization on performance over time

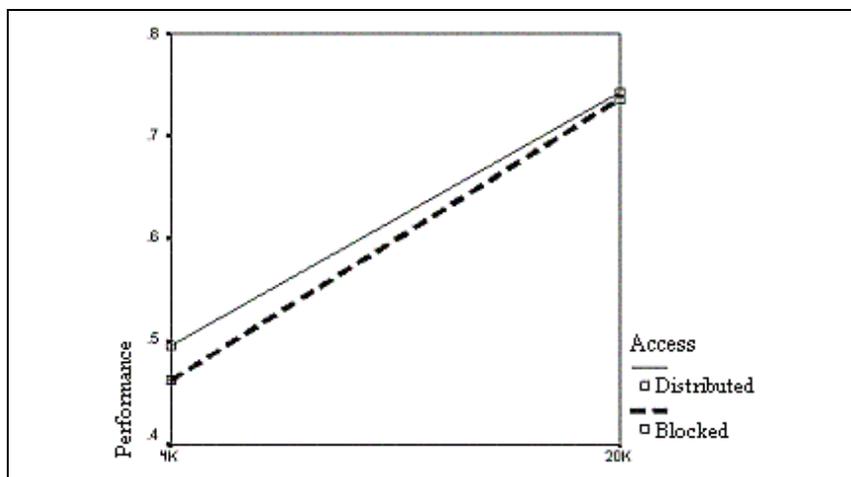


Figure 9: The effect of information access on performance over time

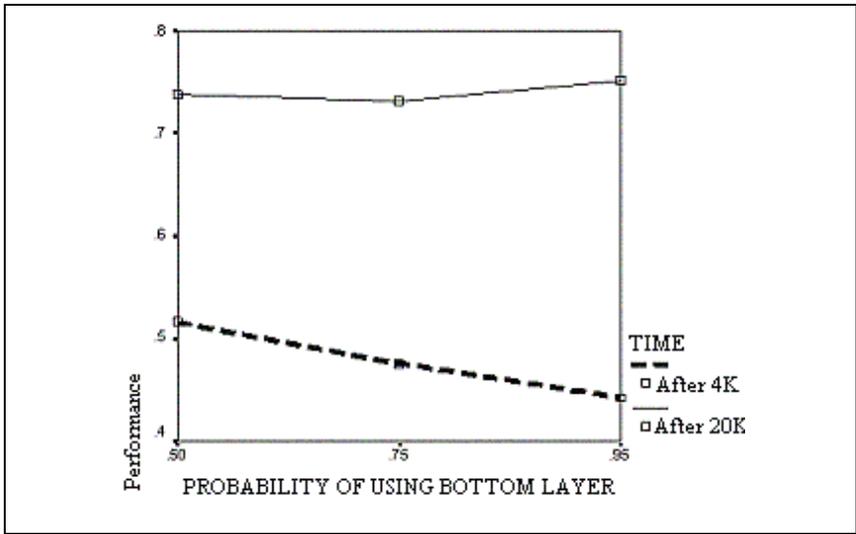


Figure 10: The effect of probability of using the bottom level on performance over time

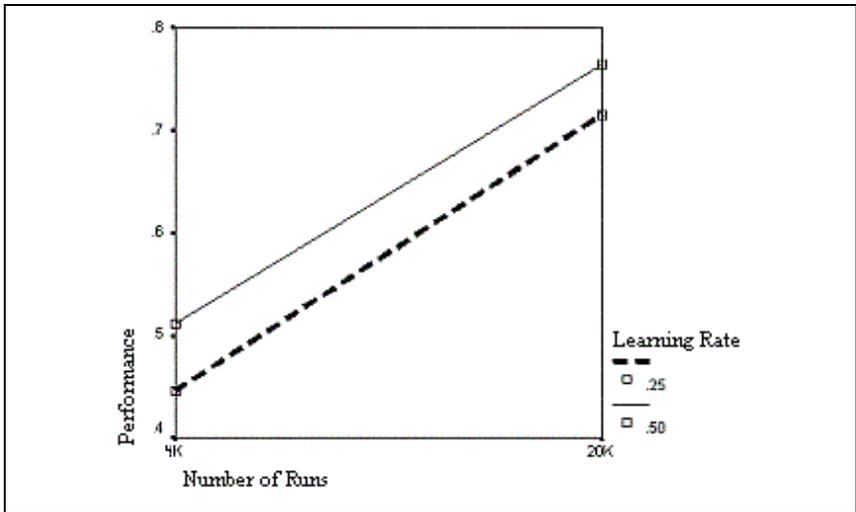


Figure 11: The effect of learning rate on performance over time.

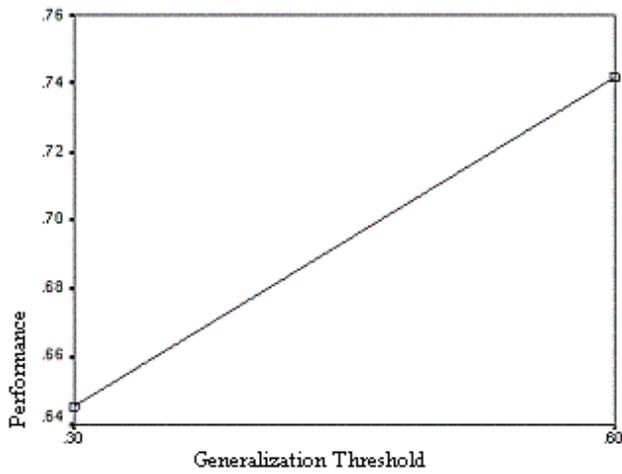


Figure 12: The effect of generalization threshold on the final performance

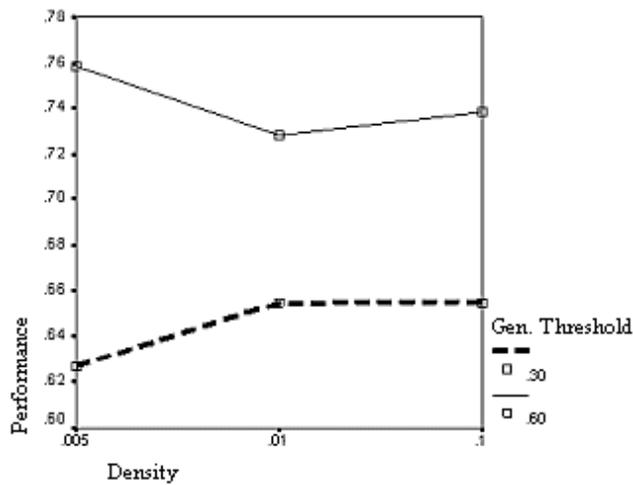


Figure 13: The interaction of generalization threshold and density with respect to the final performance

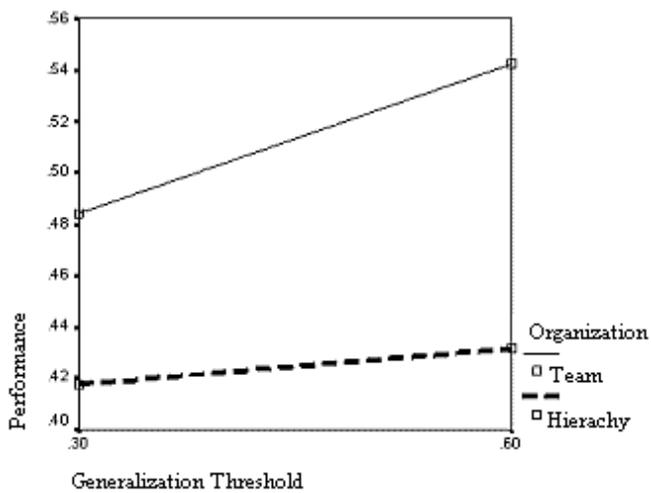


Figure 14: The interaction of generalization threshold and organization with respect to initial performance

	F	df	<i>p</i>
Effect of probability of bottom level usage (PROB_BL)	11.73	2, 24	< 0.001
Effect of learning rate	32.47	2, 24	< 0.001
Effect of temperature	2.89	1, 24	NS
Interaction of PROB_BL and time	12.37	2, 24	< 0.001

Table 3. Simulation results for general parameters of the model. Only statistically significant interactions are shown (for main effects, NS = not significant). Time is computed as a repeated-measures variable at 4,000 and 20,000 cycles.

	F	df	<i>p</i>
Effect of RER positivity threshold	.229	1, 24	NS
Effect of RER density	.094	2, 24	NS
Effect of RER generalization threshold	15.91	1, 24	< 0.001
Interaction of density and generalization threshold	2.93	2, 24	< 0.05
Interaction of generalization threshold and organization after 4,000 cycles	5.93	1, 24	< 0.05

Table 4. Simulation results for parameters related to RER learning.

7 Simulation IV: Introducing Individual Differences

Thus far, we have considered only organizations where the agents were identical in all respects. Real organizations, however, consist of individuals with widely varying cognitive capabilities and capacities. It would be interesting to extend our model to capture this, for two reasons: first, because we can observe how organizations change in response to individual cognitive differences, and second, because we can determine whether certain organizations are better at dealing with such differences.

7.1 A Single Weak Learner

Agents were organized in a hierarchy, and distributed information access was used. All agents were identical, except for one agent, which was a much slower learner than the others (learning rate of 0.01 as compared to 0.25 for other agents).

At the end of the task, the supervisor's network (at the bottom level) was analyzed by summing up the absolute values of the weights corresponding to the outputs of each subordinate agent. This allows us to compare the relative influences of different agents on the supervisor's decision. Under this comparison, the summed weights for the slow-learning agent were shown to be much lower than the other sums (by a factor of at least 2 in all cases). In other words, a supervisor gradually learns to pay less attention to the recommendations of consistently weak performers. Additionally, overall performance for the hierarchy dropped by only 3-4%. Thus, hierarchies are flexible enough to deal with a single weak performer, while showing only a slight degradation in performance.

7.2 Variable Learning Rates

In this simulation, each agent had a different learning rate (instead of having just one agent that differs from the rest). Learning rates for the agents ranged between 0.10 and 0.40, with the average being 0.25. The supervisor had a learning rate of 0.25.

The results of the simulation followed the same trends as reported in Section 5, with hierarchies outperforming teams

(after 20,000 cycles). However, here the margin by which teams were outperformed was significantly greater (by 6-11%) than in the condition where all agents were identical. This is not surprising, since the team's decision-making process -- that majority vote -- is incapable of taking individual differences into account. By contrast, a supervisor can easily emphasize one subordinate over another, based on the past success of their recommendations. A more developed discussion of the impact of weak learners in an organization can be found elsewhere (Carley 1992).

8 Discussion

8.1 Advantages of Cognitive Realism in Social Simulation

This study suggests that a more cognitively realistic simulation, with CLARION, can better capture human performance data in the radar task. Unlike simpler models, which often exhibit specialized intelligence, and thus do very well on some conditions, but poorly on others (for instance, in teams vs. hierarchies), our model, with its more general-purpose learning architecture, performs reasonably well across a variety of conditions. This is consistent with the human results (of Carley et al 1998). Furthermore, after a moderate amount of training, the trends observed closely match the human data. More specifically, teams learn faster and better than hierarchies, due to the simpler structure of teams and the difficulty of training a competent supervisor. Additionally, distributed access is superior to blocked access, showing the advantages of a variegated knowledge base at the early stages of learning. Thus, cognitive realism in social simulation can lead to models that more closely capture human results, even though social and organization simulations tend to be at a higher level and thus often gloss over the details of cognitive processes.

Moreover, by using CLARION, we are able to formulate deeper explanations for the results observed. For instance, based on our observations, one may formulate the following possible explanation: the poorer performance of hierarchies early on (see Section 4.2) may be due, at least in part, to the longer training time needed to encode high-dimensional information for the supervisor, which leads to fewer useful rules being acquired at the top level. This in turn impacts performance, since rule learning is especially important in the early stages of learning (see Section 6.2). Such explanations are only possible when the model is cognitively realistic.

In addition to offering deeper explanations, cognitive realism can lead to greater predictive power for social simulations. As has often been pointed out (e.g., in Carley 1992), the results of social simulations should not be taken as "facts," but rather as predictions that can be empirically verified. The ability to produce testable predictions, then, is a vital measure of the usefulness of a simulation. In this connection, there are two significant advantages to using cognitively realistic agents in social simulations. First, if the model is truly reflective of human cognitive processes, then its predictions will more often prove accurate. Second, predictions that contain references to aspects of human cognition (e.g., explicit vs. implicit learning) will be more illuminating and relevant than ones that refer to the internal parameters of an artificial model (e.g., momentum in a neural network) or to external measures only (e.g., percent correct).

In CLARION, we can vary parameters and options that correspond to cognitive processes and test their effect on performance. In this way, CLARION can be used to predict human performance, and to furthermore help performance by prescribing optimal or near-optimal cognitive abilities for specific tasks and organizational structures. Such prescriptions fall into two general categories. First, some prescriptions may help us to assign agents to organizational roles based on their individual cognitive capacities. For instance, we may learn that a hierarchy's performance hinges crucially on having a quick-learning agent as its supervisor, or alternatively, we may discover that quicker-learning supervisors do not significantly affect the overall performance of the organization. Second, prescriptions generated by CLARION may help us formulate organizational policies. Recall, again, the high importance of rule learning at the beginning of the learning process. Based on this, an organization may decide to emphasize standard operational procedures (i.e., rules) when training new personnel, but to emphasize case studies (i.e., exemplars) when training experienced employees. The value of such prescriptions is contingent on the cognitive realism of the model employed. The more faithfully a model captures aspects of human cognition, the wider the applicability of its predictions and prescriptions.

As stated before, this study shares much of its methodology with docking studies, although the motivations are slightly different. Whereas the primary aim of docking is cross-model validation, the current study seeks primarily to investigate interactions between models at the macro (social) and micro (cognitive) levels. It thus takes a more exploratory approach. This distinction is not clear-cut, however, since docking can also be used in a more exploratory way (Louie et al 2003).

Moreover, future work with complex cognitive models may include their utilization for validation purposes. Phelan and Lin (2001), for instance, use an organizational task very similar to that used by Carley to study different promotion systems. Their model allows for agents with different learning capabilities and thus may be amenable to docking with a cognitive model, such as CLARION, capable of representing such differences (see section 7). Such work may lead to the subsumption (Axtell 1996) of simpler agent models as a special instance of more cognitively realistic models.

Thus far we have argued for embedding cognitive agents within social simulation. However, there are several practical considerations that may limit the applicability of a more cognitive approach. One is the issue of complexity, which can make it difficult to interpret results in terms of their precise contributing factors. It is never an easy task to distinguish between results that genuinely shed light on an issue and ones that are artifactual, and more complex models may exacerbate this difficulty. Complexity also leads to longer running times and hence may raise issues of scalability. Finally, there is the issue of choice of theoretical framework, which may hinge on particular ontological conceptions of the target phenomenon. Goldspink (2000) offers a detailed analysis of these issues.

8.2 Implications for Organizational Design

With greater cognitive realism, social simulations may be able to uncover findings that will have considerable impact on organizational design. It may be suggested that minor differences in cognition may make significant differences in terms of organizational performance. Conversely, seemingly significant differences in cognition may turn out to have little impact on performance. For instance, we discovered that there is an interaction between organization of agents and how readily an agent will generalize a rule (i.e., the generalization threshold). Such a result could not have been predicted without using a simulation. On the other hand, an agent's willingness to experiment (i.e., temperature, or randomness in decision making) had no significant result on performance. Thus, in many cases, there is no *a priori* way of predicting the effects of individual cognitive mechanisms and parameters. Simulation is required in order to generate good predictions. We have found results significantly different from previous ones (in Carley et al 1998) by varying cognitive (and other) parameters. This is only possible with sufficient cognitive realism.

There is a strong possibility that cognition and organizational design may be able to substitute for each other (Carley et al 1998). Thus, if we view an organization as an aggregate entity, we can see that its performance is determined not just by its components but also by the overall structuring. To achieve a certain level of organizational performance, therefore, different combinations of organizational design and individual cognition can be selected. For instance, recall that agents organized in a team fared much better when using a rigorous (high) generalization criterion, whereas in a hierarchy no such results were observed. On this basis, a set of agents with a cautious approach towards formulating new policies (in effect, high generalization threshold) may attempt to improve performance by switching to a team organization. Alternatively, the organization may hire more intelligent personnel (i.e., raising the learning rate, which was shown to improve performance across nearly all conditions). In each case, a different combination of cognition and organizational design is used to generate a better result.

8.3 High-Level vs. Low-Level Modeling

We think of individual cognitive processes as a "lower level" description, and of social phenomena as a "higher level" description. It is therefore evident that the density of descriptions at the higher level is much greater than the density at the lower level. This means that a higher-level description may correspond to a vast multiplicity of lower-level descriptions. Each instance at the higher (or social) level corresponds to a large set of instances at the lower (or cognitive) level.

Theories of social behavior must take this correspondence into account, since any inconsistency between explanations at the higher and lower levels would invalidate the theory. Conversely, if a representative set of tests establishes the consistency of the social and cognitive explanations, the theory is considered valid. Under a successful theory, causal relationships at the cognitive level may produce unexpected, but empirically verifiable predictions at the social level.

The processes that occur at the higher level represent merely a tiny fraction of the ones that could conceivably occur, given a particular combination of entities at the lower level. Although nearly any imaginable high-level process may be described in terms of the low-level entities, the actual high-level processes that occur depend on a particular combination of conditions at the deeper level (in the physical sciences, these are known as "boundary conditions"). There is no *a priori* way

of determining, based on the lower-level entities, which of the higher-level processes will actually occur. Thus, social processes are in this sense “emergent.”

It can be argued that our approach is needlessly reductionist. A higher-level entity may consist of numerous lower-level entities. Likewise, causal relationships at the higher level may be a product of causal relationships at the lower level. Nevertheless, it is possible to describe causal relationships at the higher level without referring to relationships at the lower level. Why, then, is cognitive realism in social simulation necessary? The answer is that an effective theory must be capable, in principle at least, of mapping social phenomena to cognitive attributes. The ability to accurately model high-level phenomena through a high-level theory is a necessary, but not sufficient, condition for validity. Thus, for example, the Ptolemaic method of predicting planetary motion based on epicycles around a series of mathematical points was at least as accurate as the Copernican model of motion when the latter was first proposed. By adding additional epicycles, the Ptolemaic method could be more accurate still. Nonetheless, a theory based on epicycles around a series of theoretical mathematical points could not provide the deeper account offered by the Copernican theory of motion, in which an orbit can be traced to the presence of an astronomically identifiable body in the center of the orbit (Coward and Sun, in press). This is the primary reason why we need to bridge the two levels.

Conclusions

We have tested our approach of cognitively realistic social simulation by deploying the CLARION architecture in an organizational simulation involving multi-agent interaction. The results have been encouraging, yielding several results that accord with the psychological literature, as well as a few testable hypotheses. The empirical verification of these hypotheses should be relatively straightforward in the case of some cognitive factors (e.g., learning rate, which can be plausibly equated with scores on standardized reasoning tests), but admittedly trickier in others (e.g., generalization threshold).

Along the way, we have argued for an integration of two separate strands of research; namely, cognitive modeling and social simulation. Such integration could, on the one hand, enhance the predictive accuracy of social models (by taking into account the potentially decisive effects of individual cognition), and on the other hand, it could lead to greater explanatory power for these models (by identifying the precise role of individual cognition in collective social phenomena).

Acknowledgements

We wish to thank Xi Zhang for implementing the CLARION system and for his assistance in converting this simulation from a previous version of CLARION. Thanks are also due to Jay Hutchinson for implementing a preliminary version of the simulation.

Notes

¹ The term “cognitive modeling” is somewhat ambiguous, since it has also been used to refer to the very simple models used in previous social simulations. Here, we use the term to refer specifically to more developed, more complex cognitive models -- for instance, cognitive architectures.

² The following parameters were used for all agents: *Temperature* = 0.05; *Learning Rate* = 0.5; *Probability of Using Bottom Level* = 0.75; *RER Positivity Criterion* = 0.0; *Density* = 0.01; *Generalization Threshold* = 4.0. See Section 2.2 for a description of the cognitive parameters.

³ If we raise the threshold above a certain point, performance dips and an overall U-shaped curve is observed. The same is true for other parameters.

References

ANDERSON J R (1983) *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.

- ANDERSON J R (1993) *Rules of the Mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- ANDERSON J R and Lebiere C (1998) *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum Associates.
- AXTELL R, Axelrod J and Cohen M (1996) Aligning Simulation Models: A Case Study and Results. *Computational and Mathematical Organization Theory*, 1(2), pp. 123-141.
- BERRY D and Broadbent D (1988) Interactive Tasks and the Implicit-Explicit Distinction. *British Journal of Psychology*, 79, pp. 251-272.
- BEST B J and Lebiere C (2003) Teamwork, Communication, and Planning in ACT-R: Agents Engaging in Urban Combat in Virtual Environments. *Proceedings of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions* (Ron Sun, ed.). Acapulco, Mexico.
- BOYER P and Ramble C (2001) Cognitive Templates for Religious Concepts: Cross-Cultural Evidence for Recall of Counter-Intuitive Representations. *Cognitive Science*, 25, pp. 535-564.
- CARLEY K M (1992) Organizational Learning and Personnel Turnover. *Organizational Science*, 3(1). pp. 20-46.
- CARLEY K M and Lin Z (1995) Organizational Designs Suited to High Performance Under Stress. *IEEE - Systems Man and Cybernetics*, 25(1). pp. 221-230.
- CARLEY K M and Prietula M J (1992) Toward a Cognitively Motivated Theory of Organizations. *Proceedings of the 1992 Coordination Theory and Collaboration Technology Workshop*. Washington D.C.
- CARLEY K M, Prietula M J, and Lin, Z (1998) Design Versus Cognition: The interaction of agent cognition and organizational design on organizational performance. *Journal of Artificial Societies and Social Simulation*, 1(3), <http://www.soc.surrey.ac.uk/JASSS/1/3/4.html>
- CARLEY K M and Svoboda D M (1996) Modeling Organizational Adaptation as a Simulated Annealing Process. *Sociological Methods and Research*, 25(1), pp. 138-168.
- CASTELFRANCHI C (2001) The Theory of Social Functions: Challenges for Computational Social Science and Multi-Agent Learning. *Cognitive Systems Research*, special issue on the multi-disciplinary studies of multi-agent learning (ed. Ron Sun), 2(1), pp. 5-38.
- CECCONI F and Parisi D (1998) Individual Versus Social Survival Strategies. *Journal of Artificial Societies and Social Simulation*, 1(2), <http://www.soc.surrey.ac.uk/JASSS/1/2/1.html>
- COWARD L A and Sun R (in press) Criteria for an Effective Theory of Consciousness and Examples of Preliminary Attempts at Such a Theory. *Consciousness and Cognition*.
- EDMONDS B and Moss S (2001) "The Importance of Representing Cognitive Processes in Multi-Agent Models." In Dorffner G, Bischof H, and Hornik K (Eds.). *Artificial Neural Networks--ICANN'2001*. Springer-Verlag: Lecture Notes in Computer Science, 2130, pp. 759-766.
- GEERTZ C (1973) *The Interpretation of Culture*. New York: Basic Books.
- GILBERT N (1997) A Simulation of the Structure of Academic Science. *Sociological Research Online*, 2 (2).

- GILBERT N and Doran J (1994) *Simulating Societies: The Computer Simulation of Social Phenomena*. London, UK: UCL Press, London.
- GOLDSPINK C (2000) Modeling Social Systems as Complex: Towards a Social Simulation Meta-Model. *Journal of Artificial Societies and Social Simulation*, 3(2), <http://jasss.soc.surrey.ac.uk/3/2/1.html>
- HUTCHINS E (1995) How a Cockpit Remembers Its Speeds. *Cognitive Science*, 19, pp. 265-288.
- JENSEN F V (1996). *An Introduction to Bayesian Networks*. NY: Springer-Verlag.
- KAHAN J and Rapoport A (1984) *Theories of Coalition Formation*. Mahwah, NJ: Erlbaum.
- KLAHR D, Langley P, and Neches R (eds.) (1987) *Production System Models of Learning and Development*. Cambridge, MA: MIT Press.
- LEVY S (1992) *Artificial Life*. London: Jonathan Cape.
- LOUIE M A, Carley K M, Haghshenass L, Kunz J C, and Levitt R E (2003) Model Comparisons: Docking ORGAHEAD and SimVision. *NAACSOS conference proceedings*. PA: Pittsburgh.
- MAHER M L, Smith G J and Gero J S (2003) Design Agents in 3D Virtual Worlds. *Proceedings of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions* (Ron Sun, ed.). Acapulco, Mexico.
- MANDLER J (1992) How to Build a Baby. *Psychology Review*, 99(4), pp. 587-604.
- MATHEWS R, Buss R, Stanley W, Blanchard-Fields F, Cho J and Druhan B (1989) Role of Implicit and Explicit Processes in Learning From Examples: a Synergistic Effect. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, pp. 1083-1100.
- MOSS S (1999) Relevance, Realism and Rigour: A Third Way for Social and Economic Research. *CPM Report No. 99-56*. Manchester, UK: Center for Policy Analysis, Manchester Metropolitan University.
- PALMERI T J (1997) Exemplar Similarity and the Development of Automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, pp. 324-354.
- PANZARASA P and Jennings N (2002) Social Influence, Negotiation and Cognition. In *Simulation Modeling Practice and Theory*.
- PHELAN S and Zhiang L (2001) Promotion Systems and Organizational Performance: A Contingency Model. *Computational and Mathematical Organization Theory*, 7(3), pp. 207-232.
- PROCTOR R and Dutta A (1995) *Skill Acquisition and Human Performance*. Thousand Oaks, CA: Sage Publications.
- RABINER L (1989) A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Proceedings of the IEEE*, 77(2), pp. 257-286.
- REBER A (1989) Implicit Learning and Tacit Knowledge. *Journal of Experimental Psychology: General*, 118(3), pp. 219-235.
- ROSENBLOOM P, Laird J, Newell A, and McCarl R (1991) A preliminary analysis of the SOAR architecture as a basis for

general intelligence. *Artificial Intelligence*, 47(1-3), pp. 289-325.

RUMELHART D and McClelland J (Eds.) (1986) *Parallel Distributed Processing I*. Cambridge, MA: MIT Press.

SCHACTER D (1990) Toward a Cognitive Neuropsychology of Awareness: Implicit Knowledge and Anosagnosia. *Journal of Clinical and Experimental Neuropsychology*, 12(1), pp. 155-178.

SCHUTZ A (1967) *The Phenomenology of the Social World*. Northwestern University Press.

SEGER C (1994) Implicit Learning. *Psychological Bulletin*, 115(2). pp. 163-196.

SMITH J D and Minda J P (1998) Prototypes in the Mist: The Early Epochs of Category Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, pp. 1411-1436.

SMOLENSKY P (1988) On the Proper Treatment of Connectionism. *Behavioral and Brain Sciences*, 11(1), pp. 1-74.

STADLER M and Frensch P (1998) *Handbook of Implicit Learning*. Thousand Oaks, CA: Sage Publications.

STANLEY W, Mathews R, Buss R, and Kotler-Cope S (1989) Insight Without Awareness: On the Interaction of Verbalization, Instruction and Practice in a Simulated Process Control Task. *Quarterly Journal of Experimental Psychology*, 41A(3), pp. 553-577.

SUCHMAN L (1987) *Plans and Situated Actions*. Cambridge, UK: Cambridge University Press.

SUN R (1995) Robust Reasoning: Integrating Rule-Based and Similarity-Based Reasoning. *Artificial Intelligence*, 75(2), pp. 241-296.

SUN R (1997) Learning, Action, and Consciousness: A Hybrid Approach Towards Modeling Consciousness. *Neural Networks*, special issue on consciousness, 10(7), pp. 1317-1331.

SUN R (2001) Cognitive Science Meets Multi-Agent Systems: A Prolegomenon. *Philosophical Psychology*, 14(1), pp. 5-28.

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

SUN R, Merrill E, and Peterson T (1998) A Bottom-Up Model of Skill Learning. *Proceedings of 20th Cognitive Science Society Conference*, pp. 1037-1042. Mahwah, NJ: Lawrence Erlbaum Associates.

SUN R, Merrill E, and Peterson T (2001) From Implicit Skills to Explicit Knowledge: A Bottom-Up Model of Skill Learning. *Cognitive Science*, 25(2), pp. 203-244.

SUN R and Peterson T (1998) Autonomous Learning of Sequential Tasks: Experiments and Analyses. *IEEE Transactions on Neural Networks*, 9(6), pp. 1217-1234.

TAKADAMA K, Suematsu Y L, Sugimoto N, Nawa N E, and Shimohara K (2003) Cross-Element Validation in Multiagent-Based Simulation: Switching Learning Mechanisms in Agents. *Journal of Artificial Societies and Social Simulation*, 6(4), <http://jasss.soc.surrey.ac.uk/6/4/6.html>

WATKINS C (1989) *Learning with Delayed Rewards*. PhD Thesis, Cambridge University, Cambridge, UK.

WEST R L, Lebiere C, and Bothell D J (2003) Cognitive Architectures, Game Playing, and Interactive Agents. *Proceedings*

of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions (Ron Sun, ed.). Acapulco, Mexico.

WILLINGHAM D, Nissen M and Bullemer P (1989) On the Development of Procedural Knowledge. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, pp. 1047-1060.

YE M and Carley K M (1995) Radar-Soar: Towards An Artificial Organization Composed of Intelligent Agents. *Journal of Mathematical Sociology*, 20(2-3), pp. 219-246.