Social Institution, Cognition, and Survival: A Cognitive-Social Simulation

Ron Sun
Department of Cognitive Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
rsun@rpi.edu

Isaac Naveh
Department of Computer Science
University of Missouri-Columbia
Columbia, MO 65211

May 25, 2006

Corresponding author: Ron Sun (see the address above)

Abstract

Although computational models of cognitive agents that incorporate a wide range of cognitive functionalities have been developed in cognitive science, most of the work in social simulation still assumes rudimentary cognition on the part of the agents. In contrast, in this work, the interaction of cognition and social structures/processes is explored, through simulating survival strategies of tribal societies. The results of the simulation demonstrate interactions between cognitive and social factors. For example, we show that cognitive capabilities and tendencies may be relevant to what social institutions may be adopted. This work points to a cognitively based approach towards social simulation, as well as a new area of research —exploring the cognitive-social interaction through cognitively based social simulation.

Keywords: cognitive modeling, cognitive architecture, social simulation, institution, survival

1 Introduction

1.1 Social Simulation and Cognition

Computational social simulation has been seen as constituting a new way of exploring social processes. It has undergone significant growth in recent years. Many theoretical arguments were made in support of its role in social theorizing; see, for example, Gilbert (1995), Moss (1999), Sun (2001), Castelfranchi (2001), and Sun (2006).

A particularly pertinent factor is that social simulation can provide support for functionalist explanations of social phenomena. For example, functionalists argued that some specific forms of social structures were functional for society. However, functionalist explanations were often treated as dubious because of the difficulty in verifying such explanations (Gilbert 1995). A particularly important problem with functionalism is that it involves explaining a cause by its effect. It is customary to explain an effect by its cause, and it seems post hoc to explain a cause by its effect. A related problem is that, while focusing on a specific moment in history, it tends to ignore the historical processes leading up to a specific social phenomenon. Social simulation, however, can help to substantiate functionalist explanations, by remedying both of these above two problems. First of all, social simulation focuses on *processes*, and thus it may help to provide some historical perspectives. For example, Cecconi and Parisi (1998) focused on the evolution of survival strategies in tribal societies. Similarly, Doran et al (1994) provided explanations for the increasing complexity of tribal societies in the Upper Paleolithic period. Second, the effect of a cause, which is central to functionalist explanations, can be verified through experimentation using computational social simulation. Consequently, with social simulation, this style of explanation can be better verified or validated, and thus may become more convincing.

However, a significant shortcoming of current computational social simulation is that most of the work assumes very rudimentary cognition on the part of agents. Although agents are often characterized as being "cognitive", there have been relatively few attempts to carefully emulate human cognition. Models of agents have frequently been custom-tailored to the task at hand, often amounting to little more than a restricted set of highly domain-specific rules. Although such an approach may be adequate for achieving some limited objectives of some simulations, it is overall unsatisfactory (see the arguments by Sun and Naveh 2004). It not only limits the realism, and hence the applicability, of social simulations, but also precludes the possibility of tackling the question of the micro-macro link in terms of cognitive-social interaction (cf. Sawyer 2003, Alexander et al 1987).

Computational models of cognitive agents incorporating a wide range of cognitive functionalities (such as various types of memory, various modes of learning, and various sensory motor capabilities) have been developed in cognitive science (e.g., Newell 1994, 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 (e.g., Sun 2002). We have argued elsewhere why such models of cognition can greatly enhance social simulation in general (see, in particular, Sun 2006), which we will not repeat here.

1.2 Revamping An Existing Simple Simulation

To make the same point in a different way, let us look into an existing social simulation as an example. Cecconi and Parisi (1998) created simulated social groups. In these groups, to survive and reproduce, an agent must possess certain resources. A group in which each agent uses only its own resources is said to adopt an individual survival strategy (ISS). 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 (SSS). For instance, the "central store" (CS) is a mechanism to which all the individuals in a group transfer (part of) their resources. The resources collected by the CS can be redistributed to the members of the group (Cecconi and Parisi 1998).

In Cecconi and Parisi (1998), a number of simulations were conducted comparing ISS groups with SSS groups (adopting CS strategies). They used neural networks to model individuals and a genetic algorithm to model evolution. Networks (representing individuals) survive and reproduce differentially based on the quantity of food they are able to consume. Cecconi and Parisi studied how their abilities evolved and also the evolutionary changes in group size. In particular, they explored what conditions determined group survival or extinction. This work is interesting, because it provides a fertile ground for exploring a range of issues, ranging from social institutions to individual behaviors, and from evolution to individual learning, and so on.

However, as is, in this work there is very little in the way of cognition by individual agents in their struggle to survive. We believe that investigation, modeling, and simulation of social phenomena need cognitive science (Sun 2001), because such endeavors need a better understanding, and better models, of individual cognition, only on the basis of which better models of aggregate processes can be developed. Cognitive models may provide better grounding for understanding multi-agent phenomena, by incorporating realistic constraints, capabilities, and tendencies of individual agents in terms of their cognitive processes (which were not present in, e.g., Cecconi and Parisi 1998). This point was argued in Sun (2001). This point has also been made, for example, in the context of cognitive realism of game theory (Kahan and Rapaport 1984, Camerer 1997), or in the context of deeper models for addressing human-computer interaction (Gray and Altmann 2001). In Axelrod (1984), it was shown that even adding a cognitive factor as simple as memory of past several events into an agent model can completely alter the dynamics of social interaction (in the iterated prisoner's dilemma in particular).

In this work, we aim to redress the neglect of cognitive science in social simulation, by incorporating more detailed and more realistic models of cognitive processes into social simulation. The specific model that we adopt is the CLARION cognitive architecture (Sun et al 2001, 2005, Sun 2002, 2003). CLARION has been successful in simulating a variety of psychological tasks. [1] Therefore, we are in a good position to extend the effort on CLARION to the capturing of a wide range of social phenomena through integrating cognitive modeling and social simulation. Some preliminary work has already been done (see, e.g., Sun and Naveh 2004).

1.3 The Context of Our Simulation

In this work, we revamp the simple (not cognitively realistic) simulation of Cecconi and Parisi (1998). The general setup is similar to that of Cecconi and Parisi (1998). The world is of a limited physical dimension, made up of a two-dimensional grid. Within its physical dimensions, food items and agents are randomly distributed among the tiles. The food crops grow by seasons: Every once in a while, food is replenished among the tiles. There are the harsh, medium, and benign conditions, which are distinguished by the agent-to-food ratios. (Social structures and agent behaviors are expected to be different in these different conditions; more on this later.) Agents are of a limited life span, which varies from individual to individual, depending on the energy level of an agent, the maximum lifespan, and other factors. Agents look for and consume food in an effort to prolong their life spans.

[1] These tasks include serial reaction time tasks, artificial grammar learning tasks, process control tasks, categorical inference tasks, alphabetical arithmetic tasks, and the Tower of Hanoi task (see, e.g., Sun 2002). In addition, extensive work has been done on a complex minefield navigation task (Sun et al 2001). Simulations involving motivational structures and metacognitive processes are also under way.

As in Cecconi and Parisi (1998), there is a “central store” in some cases. Agents may be required to contribute to the central store in these cases. However, different from Cecconi and Parisi (1998), we introduced some further social institutions. In the case of mandatory contribution to the central store, we introduce penalty for not contributing to the central store. Penalty for violating the norm may be attributed to two possible sources: (1) social sources, such as a tribal enforcement mechanism that extracts penalty from violators, (2) internal sources, for example, from an internal feeling of guilt.

Most notably, in our work, different from that of Cecconi and Parisi (1998), agents are more cognitively realistic. They are constructed out of a cognitive architecture, which captures a variety of cognitive processes in a psychologically realistic way in detail (see Sun 2002, 2003). Therefore, we expect our simulations of social survival strategies will shed more light on the role of cognition in determining survival strategies and the interaction with social structures (social institutions) and processes. The reason that we introduce these additional factors is to investigate the interaction between social structures/processes and individual cognition (i.e., the micro-macro link). Through experimentation and rigorous data analysis, we hope to arrive at a more precise and more detailed understanding than the previous simulations.

This task is appropriate to our goal of understanding cognitive-social interaction. In a way, social processes may be viewed, in general, as processes for distributing and re-distributing power and wealth. This survival task captures such processes in a microcosm.

In the remainder of this paper, first, a more realistic cognitive architecture, CLARION, will be described, which, among other things, captures the distinction between explicit and implicit learning. This architecture will then be applied to the survival task simulation. The idea here is mainly to substitute more cognitively sophisticated agents, based on CLARION, for the simpler agents used in Cecconi and Parisi (1998). Analysis of results and discussions will be presented. Some brief concluding remarks will then complete this paper.

2 A Cognitive Architecture

Below we introduce the cognitive architecture CLARION to be used in the simulation. In order to justify the architecture itself, we first discuss a major cognitive issue that it addresses, and then we sketch out its major features relevant to our simulation. This cognitive architecture has been extensively described in a series of papers and books (e.g., Sun 1997, Sun 2002, Sun 2003, Sun et al 2005). Given length considerations, we cannot present all the intricate details of the cognitive architecture here. For further information, see Sun (2002, 2003).

A major design goal for CLARION was to have a set of adjustable parameters that correspond to aspects of cognition. This is in contrast to some other models (e.g., Cecconi and Parisi 1998) in which performance depends on a set of variables that are mathematically/technically 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 effects on performance as a function of cognition.

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; Seger 1994; Stadler and Frensch 1998, Tetlock 2002). Although explicit and implicit learning have both been actively studied, the question of the interaction between these two types of processes has rarely been broached. However, despite the lack of study of this interaction, it has become evident (e.g., in Seger 1994) that rarely, if ever, is only one of the two types of learning engaged. Our review of experimental data (e.g., Reber 1989; Stanley et al 1989; Sun et al 2001) 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 et al 2001, 2005), 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 humans (e.g., Stanley et al 1989; Mandler 1992, Sun et al 2001, Sun et al 2005).

2.2 A Sketch of the CLARION Cognitive Architecture

CLARION is an integrative cognitive architecture with a dual representational structure (Sun et al 2001, 2005; Sun 2002, 2003). It consists of two levels: a top level that captures explicit learning, and a bottom level that captures implicit learning (see Figure 1).

At the bottom level, the inaccessibility of implicit learning is captured by subsymbolic distributed representations. This is because representational units in a distributed representation are capable of performing tasks but are generally not individually meaningful (Sun 1995). Learning at the bottom level proceeds in a trial-and-error fashion, guided by reinforcement learning (i.e., Q-learning) implemented in backpropagation neural networks (Sun and Peterson 1998, Sun et al 2001).

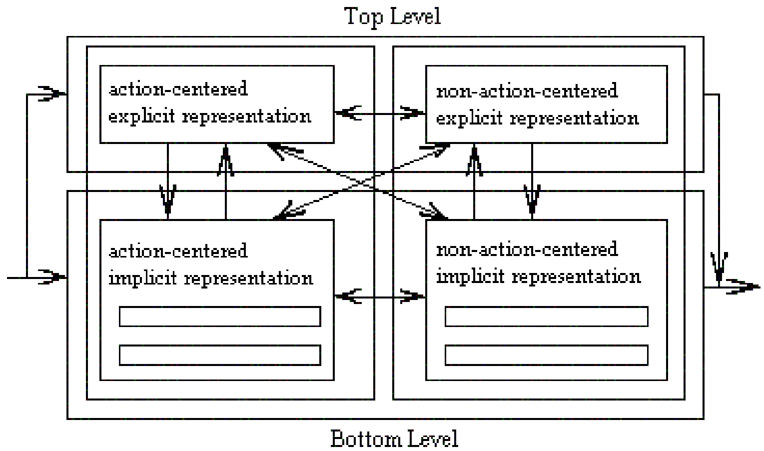

Figure 1: The CLARION architecture

At the top level, explicit learning is captured by a symbolic representation, in which each element is discrete and has a clearer meaning. This accords well with the directly accessible nature of explicit knowledge (Sun 2002). 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 ratios 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 current state x: $Q(x, a_1)$, $Q(x, a_2)$, ..., $Q(x, a_n)$.

3. Find out all the possible actions $(b_1, b_2, ..., b_m)$ at the top level, based on the 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, 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 (reinforcement) information.

7. Update the top level using the *Rule-Extraction-Refinement* algorithm.

8. Go back to Step 1.

Some explanation is in order. 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:

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

where x is the current state, a is the action performed, r is the immediate feedback (reinforcement), y is the next state, and γ is the discount factor. $\Delta Q(x, a)$ 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:

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

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, to form implicit skills (Tetlock 2002).

In the top level, explicit knowledge is captured in a simple prepositional rule form. An algorithm was devised for learning explicit knowledge (explicit rules) using information from the bottom level (i.e., 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. 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 based on Q values at every step when the rule is applied. Specifically, the inequality, $\gamma \max_b Q(y,b) + r - Q(x,a) > threshold_{RER}$, determines the positivity/negativity of a step and the rule matching this step (where a is the action performed in state x, y is the next state, and r is the reinforcement received). The positivity threshold (denoted $threshold_{RER}$ above) corresponds to whether or not an action is perceived by the agent as being reasonably good. Then, based on the positivity of a step, PM (Positive Match) and NM (Negative Match) counts of the matching rules are updated. IG is then 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:

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

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 an analogous fashion, except that a value in 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, Sun 2003). [2]

[2]Due to running-time considerations, the specialization threshold is held constant in all simulations reported here.

To integrate the results from the two levels, levels are chosen stochastically, using a probability of selecting each level (e.g., a probability $Prob_{BL}$ for selecting the bottom level; Sun 2003).

When the outcome from the bottom level is chosen, a stochastic process based on the Boltzmann distribution of Q values is used 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.[3]

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 part. There are also other components, such as working memory, goal structure, drive representation, metacognitive representation, and so on. They are not relevant to the present study.

3 Details of Simulation Setup

Below, details of the simulation based on the CLARION cognitive architecture are presented. Those not interested in technical details may skip this section.

[3]This method is also known as Luce's choice axiom (Watkins 1989). It is found to match psychological data in many domains.

3.1 Agents According to CLARION

Agents are constructed based on CLARION (as described in the previous section). Each agent faces a certain direction (north, south, east, or west). Each agent receives input regarding the location of the nearest food, relative to current position of the agent and its current direction. An agent's visual perception is divided into 4 pizza-slice-shaped quadrants. Thus, if the agent is facing east and the nearest food is in the south quadrant, input to the agent will be "direction: right". Each agent has one action output: either (1) turn 90 degrees right, (2) turn 90 degrees left, or (3) move forward. The reinforcement to the agent is such that upon collecting a food item, a feedback of 1 is provided.

Each agent lives for a maximum of 350 cycles, but it may die early due to lack of food. In our simulation, procreation is asexual (only one parent is required, which is the same as in Cecconi and Parisi's simulation). Procreation occurs if: (1) an agent has reached 120 energy units or more, and (2) there are fewer than the maximum number of agents in the world. The new agent is placed in a random location. The parent hands out 60 energy units to the child upon its birth. The child inherits its parent's internal makeup (including the neural network and the rule set), although when a child is spawned, there is a 10% chance of mutation. If mutation occurs, each of the weights in the neural network have a 20% chance of being randomly decreased or increased by 0.1.

The CLARION parameters were varied in simulation as will be described later.

3.2 The Environment

The world is made up of a 100x100 grid. Each of these 10,000 tiles may or may not contain one food item. Up to 600 tiles (out of a total of 10,000 tiles) contain one food item each. Every 40 cycles, the 100x100 grid is replenished: Randomly selected tiles are restocked with food items, until the grid once again has 600 food tiles. We also tested a more benign condition, in which 900 tiles contain one food item each, and a harsher condition, in which 300 tiles contain one food item each.

The food consumption by agents is as follows: For each agent, capturing a food item increases the agent's energy by 50 units. Each agent begins with 60 units of energy, and each agent consumes one unit of energy per cycle.

The maximum population in this world is 30 agents (for the sake of maintaining a reasonable running speed on the slow PC we used for simulation). There are initially 30 agents to begin with, and the number of agents fluctuates (within the bound of a maximum of 30 agents).

3.3 The Social System

Some simulations were run with agents contributing to (and drawing upon) a "central store" of food (as in Cecconi and Parisi 1998). In these cases, an agent is required to contribute 20 energy units to the central store when it picks up a food item (50 energy units). When a central store is used, at each cycle, 10% of the agent population (randomly selected) receives 5 energy units each from the central store. A variation of this is that only agents with 10 or less energy units receive distributions from the central store.

In some cases, agents are allowed to make their own decisions regarding whether to contribute to the central store or not. Therefore, in these cases, each agent has an extra ("social") action output: indicating whether to contribute to the central store or to cheat. As mentioned earlier, different from Cecconi and Parisi (1998), enforcement (justice) mechanisms are also introduced in some cases in our simulations. If the agent decides not to cheat, it contributes 20 energy units to the central store whenever it picks up a food item. If the agent decides to cheat, there is a 30% chance of being caught (when an enforcement mechanism is in place). If caught, the agent is fined 40 energy units (which are transferred to the central store). Otherwise, the agent keeps all the energy units acquired.

Incorporating the extra decision regarding to contribute or not to contribute, the reinforcement for agents is as follows: If the agent captures one food item and the agent does not cheat, the reinforcement is 0.6. If the agent captures one food item and the agent cheats (without being caught), the reinforcement is 1.0. If the agent captures one food item and the agent cheats and is caught, the reinforcement is 0.2. Note that this reinforcement scheme is proportional to what an agent gets to keep in each possible scenario (based on 30/50, 50/50, or 10/50 energy units being kept, respectively).

3.4 Running the Simulation

A large set of parameters of CLARION can be varied. These parameters were described in detail earlier. The first part consists of fundamental properties of the model, including: (1) learning rate of the neural network, (2) reliance on the top vs. the bottom level, expressed as a probability of choosing each level, (3) temperature (or degree of randomness), and so on. The second part consists of parameters concerning RER learning (rule extraction), including: (1) RER positivity threshold (which must be exceeded to count a step as being positive), (2) RER generalization threshold (which must be exceeded for a rule to be generalized), (3) RER specialization threshold, and so on.

There are also environmental and social variables: for example, (1) survival strategy, (2) food availability, and so on.

Similarly, performance metrics for simulation may include a variety of different measures: (1) number of cycles survived (before population extinction), (2) average food/energy consumed per agent per cycle, (3) average food/energy captured per agent per cycle, (4) average number of agents in a population, (5) average lifespan of agents, and so on.

However, considering the cost of running such complex simulations, in order to simplify the simulations as much as we can, we focused on a smaller set of more important parameters, and we chose a small set of values for each of these parameters (usually 2-3). See Figure 2 for details of these parameters. Dependent variables (performance measures) were also limited to the following: average individual energy acquisition per agent per cycle, and average population size (average number of agents in a population), as well as average lifespan.

All combinations of the above conditions were tested. There were a total of 96 combinations. Each simulation ran for a maximum of 2000 cycles (out of practical considerations concerning running time). Sampling was done 10 times per simulation. Thus we had 960 observations in total.

Probbl (probability of using the bottom level):	
#1	0.25
#2	0.75
Learn (learning rate):	
#1	0.25
#2	0.75
Gen (generalization threshold):	
#1	1.0
#2	3.0
Food (food availability):	
#1	300
#2	600
#3	900
Strat (survival strategy; cs = central store):	
#1	cs/enforcement/choice
#2	cs/no enforcement/choice
#3	cs/no enforcement/no choice
#4	no cs

Figure 2: A list of parameter values varied in simulation.

4 Simulation Results and Analysis

4.1 Findings Regarding Social Survival Strategies

First, let us consider a few performance measures used in our experiments. A reasonable measure of average individual success seems to be the per-capita energy acquisition per cycle. However, sometimes, as population sizes decrease, the performance of survivors actually increases due to less competition for food. For this reason, we simultaneously examine population size as a complementary measure. The two measures should be cross-referenced in reaching any conclusion. Another, more individual measure, average life span, was also looked at, which more accurately measures how much (on average) each individual benefits from food (or suffers from the lack of food), different from the above two more global measures. With these measures in mind, let us look into some results from the experiments.

One of the most important findings is that in terms of average energy acquisition per capita per cycle, strategy matters significantly: ANOVA (analysis of variance) shows $F(3, 864) = 4.108, p < 0.01$. See Figure 3. As indicated in the figure, the worst strategy for this performance measure was that of central store with no free choice. This is probably because agents in this case have no leeway whatsoever in contributing to and drawing from the central store, which leads to less evolutionary pressure on individuals during evolution. The performance of the agents suffers in the long run as a result. Also as expected, the strategy of central store with free choice and some enforcement was significantly better for this performance measure. Central store with free choice but no enforcement was in turn better than central store with free choice and enforcement. The difference between the two was small though ($p > 0.05$). The most interesting finding here is that

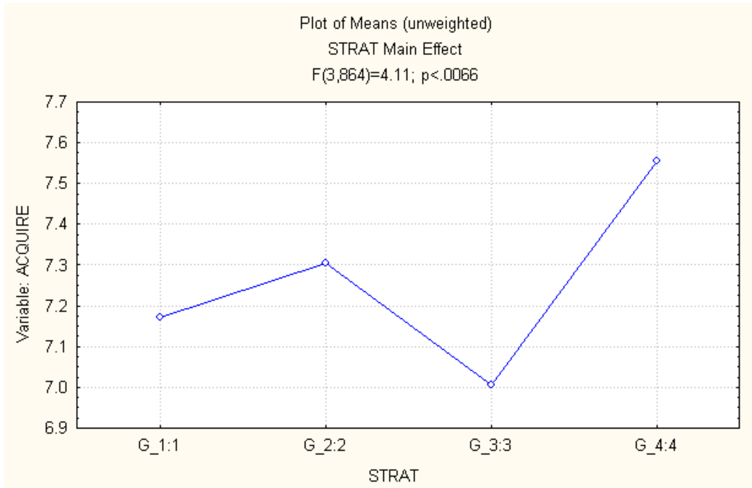

Figure 3: The significant differences in terms of energy acquisition with regard to strategies. Strat: 1 = cs/enforcement/choice, 2 = cs/no enforcement/choice, 3 = cs/no enforcement/no choice, 4 = no cs.

the best strategy was that of no central store. This is not really surprising: for individuals, more evolutionary pressure leads to better individual performance, and the no central store strategy exerts most evolutionary pressure on individuals during evolution. As a result, individuals surviving in this environment fare better in general. The strategy of central store with free choice but no enforcement was the closest to it. This is consistent with Cecconi and Parisi (1998): As demonstrated by Cecconi and Parisi (1998), the strategy of central store with free choice but no enforcement often turns into the strategy of no central store at all.

In addition, in a separate simulation, we also found that with the use of central store, distribution from the central store to the needy only was slightly better than random distribution. But their differences were not statistically significant ($p > 0.05$).

We also found that average rates of energy acquisition per capita per cycle differ a great deal with regard to harsh versus benign environments (low versus high food availability). For example, the rate for a harsh environment was 1.63, while that for a benign environment was 13.97. This difference was statistically significant ($p < 0.01$).

However, in terms of population size (when used as a dependent variable), the choice of strategy had no significant effect. Note that although the previously discussed difference of strategies in terms of average energy acquisition per capita per cycle was statistically significant, the actual difference was small in general. When we switched from this somewhat individual measure to a more global measure, it appears that those small differences became even smaller and thus disappeared in statistical analysis.

However, the two factors, strategy and food availability, together have a significant effect on population size (therefore our result is consistent with the findings of Cecconi and Parisi 1998). Using ANOVA, we indeed found an interaction between strategy and food availability: $F(6, 864) = 4.737, p < 0.01$. This reflects the fact that when food is abundantly available, strategy does not matter very much for maintaining population sizes, since agents can survive regardless (as in Cecconi and Parisi 1998). When food is less available, strategy does make a difference. See Figure 4. This result contrasts with the previous analysis using energy acquisition rate: In the case of using population size as the dependent variable, strategy does not have a universal effect, but has an effect only under low to medium food availability.

In terms of the average lifespan of individuals, there is a significant main effect of strategy: $F(3, 864) = 438.26, p < 0.000$. However, looking into the data itself, we see something interesting: The strategy of central store with no choice fared well, which shows that this mandatory social welfare system

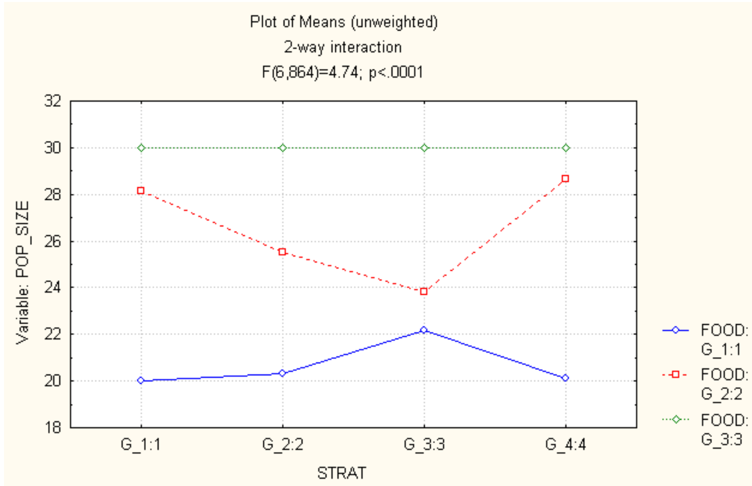

Figure 4: The significant interaction between food availability and strategy. Food: 1 = 300, 2 = 600, 3 = 900; Strat: 1 = cs/enforcement/choice, 2 = cs/no enforcement/choice, 3 = cs/no enforcement/no choice, 4 = no cs.

did help the survival of individuals. The strategy of central store with free choice and enforcement was comparable. These two strategies were the best among the four strategies. However, the strategy of no central store was the worst, for the lack of a social cushion against adverse circumstances that an individual might encounter. Notice that this ordering is the opposite of the ordering resulting from the measure of energy acquisition rate. This is due to the fact that the dependent variable of lifespan is arguably a measure that is most sensitive to individual performance, as it captures how an individual benefits from the availability of food or suffers from the lack of it (on average). So, there are some contrasting effects: While some individuals may die of starvation, the average rate of energy acquisition may be high nevertheless. Conversely, the starvation of some individuals may be prevented, but the population as a whole may be somewhat hurt in the process. That

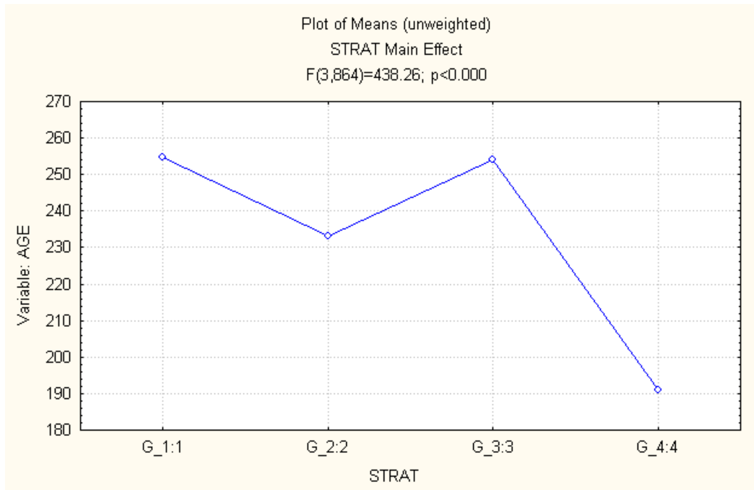

Figure 5: The significant differences of lifespan with regard to strategies. Strat: 1 = cs/enforcement/choice, 2 = cs/no enforcement/choice, 3 = cs/no enforcement/no choice, 4 = no cs.

is, some strategies may lead to benefiting the overall population while at the same time hurting some individuals within that population, and vice versa. This apparent contradiction is a real possibility. See Figure 5. Separately, distribution from the central store to the needy only, as opposed to random distribution, does improve average lifespan significantly ($p < 0.01$).

Overall, our experimental results indicated that strategies do matter for individuals evolved in a particular physical and social environment and surviving in that environment (using both energy acquisition rate and lifespan as measures). In particular, we want to highlight the following findings with regard to strategies: In terms of energy acquisition rate, the strategy of central store with no free choice (that is, the central store strategy with absolute enforcement) performed worse than the strategy of central store with free choice and some enforcement. In other words, overly strict enforcement may be harmful. However, it did fare well in terms of lifespan, which is a measure that is most sensitive to individual performance. Thus, the strategy of central store with no free choice has this contrasting effect on these two measures: It does not help the population as a whole although it does help individuals in it. Similarly, the strategy of no central store also has a contrasting effect on the two measures: It is the best in terms of energy acquisition rate but the worst in terms of lifespan (for essentially the same reason). We also note yet another contrasting effect—that the central store strategy with distributions to the needy only does not improve (global) performance measures significantly in terms of either energy acquisition rate or population size, although it does improve lifespan (an individual performance measure) significantly.

However, these results above, although interesting (and going beyond previous results, by providing more rigorous statistical analysis), are not our main concern here. We are more interested in a detailed analysis of cognitive-social interaction, which was not addressed in the Cecconi and Parisi (1998) study, and which we turn to next.

4.2 Findings Regarding Cognitive Factors

There is more to cognitive modeling and to cognitively based social simulation than merely producing some performance measures. Because CLARION captures a wide variety of cognitive factors, we can vary parameters that correspond to specific cognitive factors, and observe their effects on performance. This approach offers an important advantage over other, more task-specific models, where differences in performance tend to be artifacts of the particular model used and may be of little independent interest. Unlike the Cecconi and Parisi (1998) study, with CLARION, the parameters being altered are the fundamental aspects of cognition, and thus observed differences in performance are far more likely to stem from real, testable differences in individual cognition. Accordingly, in the following simulations, we varied a number of cognitive parameters and observed their effects on performance.

The most important aspect for us to look into is the interaction between cognition and social institution (as well as physical environment): that is, what cognitive parameter settings are suitable for what kind of social institution and physical environment (central store or not, enforcement or not, under the condition of abundant food versus scarce food, etc.) This would again require statistical analysis, and ANOVA in particular.

First let us look into an analysis using the dependent variable of average energy acquisition per capita per cycle. We found some expected, not-so-surprising effects of well known cognitive parameters: for example, probability of using the bottom level, learning rate, generalization threshold, and so on. For instance, we found a significant main effect of probability of using the bottom level as opposed to using the top level: $F(1, 864) = 68.204, p < 0.01$. That is, using explicit rules more at the top level is beneficial (up to a certain extent of course; Tetlock 2002). This is because, as we demonstrated before in our prior work (Sun et al 2001, Sun 2002, Sun et al 2005), using explicit rules more at the top level helps implicit learning at the bottom level and the overall learning performance (see also Stanley et al 1989, Tetlock 2002). Hence there is the performance difference. See Figure 6.

There is also a significant main effect of learning rate: $F(1, 864) = 11.102, p < 0.01$. That is, using higher learning rates is beneficial (up to a certain extent of course). This is because higher learning rates (up to a certain extent) help agents to quickly adapt to situations and exploit their

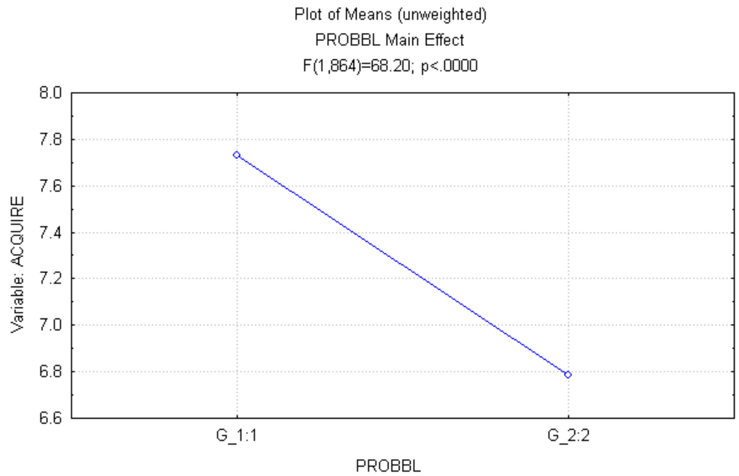

Figure 6: The significant main effect of the probability of using the bottom level. Probbl: 1 = 0.25, 2 = 0.75.

physical and social environments to ensure their own survival. See Figure 7

We did not find any significant effect of generalization threshold ($p > 0.05$). It appears that in this task how readily an individual generalizes acquired explicit knowledge is not very important. This may be attributable to the fact that the effect of this parameter depends on a host of other factors and is thus somewhat insignificant for this performance measure (but see the analysis of interactions later).

Some interesting two-way interactions were also found through ANOVA. First of all, there is an interaction of learning rate and food availability: $F(2, 864) = 5.630, p < 0.01$. Under low food availability, a higher learning rate is better. Under medium or high food availability, it does not matter. This finding may be explained as follows: Under low food availability, that is, under a harsh condition, it is more important to exploit well the environment and the social structures/institutions in order to survive, while in less harsh

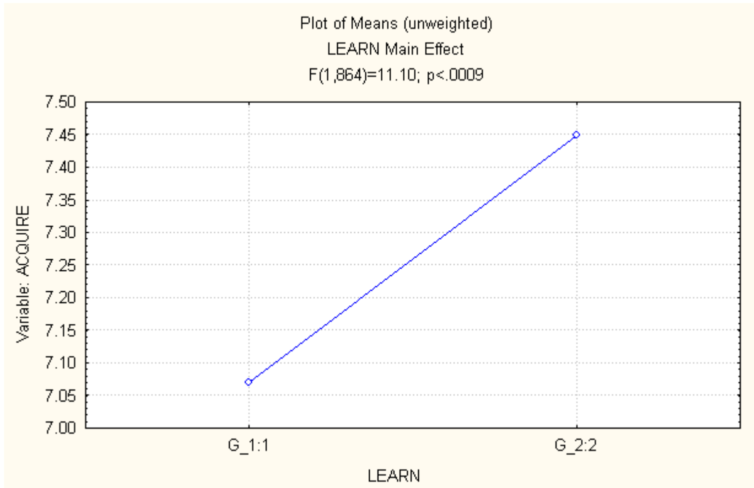

Figure 7: The significant main effect of the learning rate. Learn: $1 = 0.25$, $2 = 0.75$.

conditions, slacking off is less of a problem (see also Sun and Peterson 1998). See Figure 8.

Interestingly, there is the interaction between probability of using the bottom level and strategy: $F(3, 864) = 4.023, p < 0.01$. As indicated by Figure 9, the strategy of no central store is the best strategy when explicit rules are not used much (i.e., when the probability of using the bottom level is high), while it is merely average when rules are heavily used (i.e., when the probability of using the bottom level is low). One possible explanation is that when explicit rules are heavily used (i.e., when cognition is highly explicit), agents are more likely to learn better and explore given situations better (see Sun et al 2001, 2005); note that situations induced by the strategies with central store is more complex than that induced by the no central store strategy and therefore require better learning abilities; however, given a better learning ability, agents may perform better in these more complex

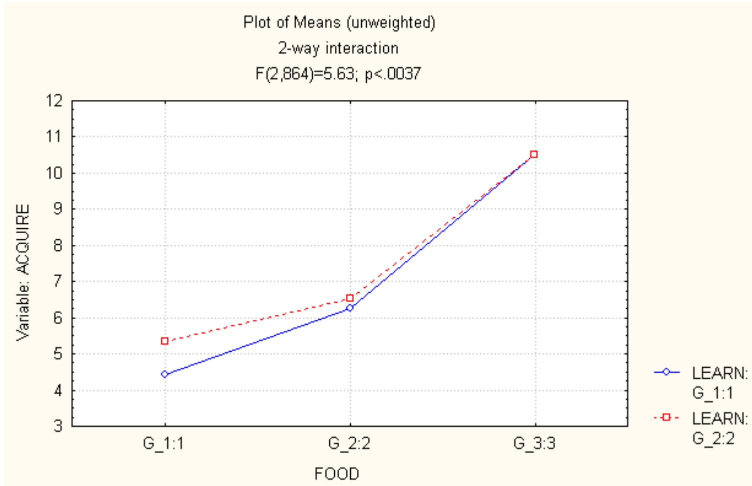

Figure 8: The significant interaction between learning rate and food availability. Learn: 1 = 0.25, 2 = 0.75; Food: 1 = 300, 2 = 600, 3 = 900.

situations by exploiting them better (Cecconi and Parisi 1998); therefore, agents with heavier rule usage (more explicit processing) performed better in these more complex situations (i.e., with central store). In contrast, when explicit rules are less used, agents do not learn as well, and therefore in this case, the no central store strategy—the simplest strategy—appears to be the best. However, for this particular performance measure, the strategy of central store with no free choice is always a bad condition regardless (but not necessarily the worst), because it alleviates evolutionary pressure on individuals by the use of a socialistic system. In general, the differences among strategies are much greater when explicit rules are less used, because agents with poorer learning abilities do not learn to handle more complex situations well.

There are also some interactions between cognitive parameters. For instance, there is the interaction of probability of using the bottom level and

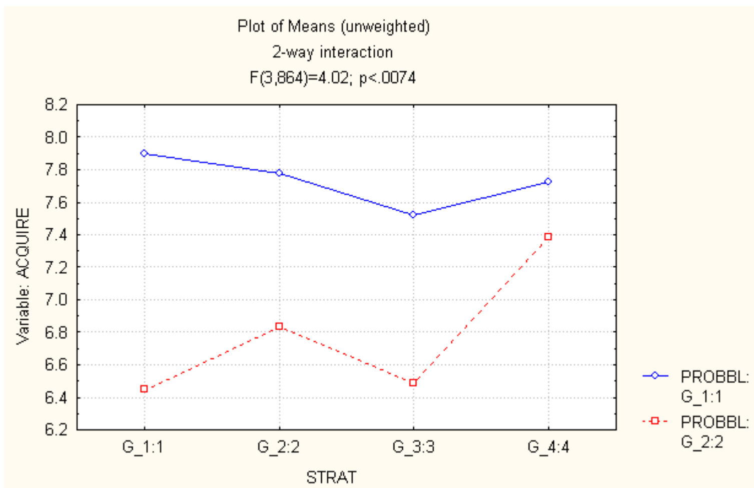

Figure 9: The significant interaction between probability of using the bottom level and strategy. Probbl: 1 = 0.25, 2 = 0.75; Strat: 1 = cs/enforcement/choice, 2 = cs/no enforcement/choice, 3 = cs/no enforcement/no choice, 4 = no cs.

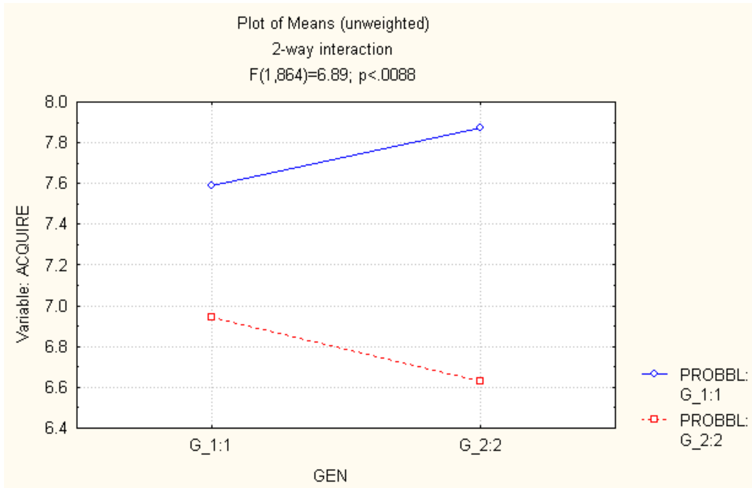

Figure 10: The significant interaction between probability of using the bottom level and generalization threshold. Probbl: 1 = 0.25, 2 = 0.75; Gen: 1 = 1.0, 2 = 3.0.

generalization threshold: $F(1, 864) = 6.890, p < 0.01$. A higher generalization threshold is better when relying more on explicit rules, but worse when relying less on them. This is probably because of the following: relying more on explicit rules (i.e., more explicit processing) leads to frequent rule uses and thus frequent rule generalizations, while relying less on explicit rules leads to less frequent rule uses and thus less frequent rule generalizations; a higher generalization threshold leads to less generalizations, which is fine in the case of frequent rule uses (which entails frequent rule generalizations anyway) but detrimental in the case of less frequent rule uses (because it leads to too infrequent rule generalizations). See Figure 10.

Now let us turn to the analysis using population size as the dependent variable (which is a more global measure of performance). Generally speaking, we found similar effects using population size as the dependent variable, as in the case of using energy acquisition rate as the dependent variable (as described earlier), in terms of significant main effects. For instance, we found a significant main effect of probability of using the bottom level: $F(1, 864) = 7.399, p < 0.01$; that is, using explicit rules more (i.e., more explicit processing) is beneficial (up to a certain extent). This is exactly the same as the previous analysis (using energy acquisition rate as the dependent variable), for the same reason. We also found a significant main effect of learning rate: $F(1, 864) = 7.188, p < 0.01$. A higher learning rate is better (up to a certain extent). This is exactly the same as the previous analysis (using energy acquisition rate as the dependent variable), for the same reason.

However, unlike the analysis of using energy acquisition rate as the dependent variable, in this case, we actually found a significant main effect of generalization threshold: $F(1, 864) = 7.472, p < 0.01$. Using a higher generalization threshold is worse, presumably because it leads to too few explicit rules (cf. Sun et al 2001 for a analogous situation under a dual-task condition). See Figure 11.

In terms of two-way interactions, as before with regard to energy acquisition rate, here we also found an interaction between probability of using the bottom level and strategy: $F(3, 864) = 4.018, p < 0.01$. As before, we notice that no central store is the best strategy under lower usage of explicit rules (less explicit processing), but becomes worse off under heavier rule usage (more explicit processing), which may be attributed to the same reason as speculated before.

As before, there are also some interactions between cognitive parameters, this time with regard to population size. For instance, as in the case of using energy acquisition rate as the dependent variable, there is an interaction

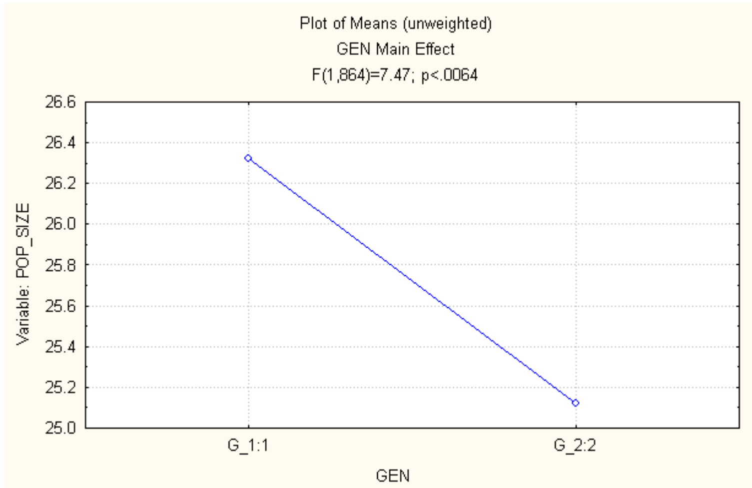

Figure 11: The significant main effect of the generalization threshold. Gen: $1 = 1.0$, $2 = 3.0$.

between probability of using the bottom level and generalization threshold: $F(1, 864) = 6.759, p < 0.01$.

A new finding here is that there is an interaction between learning rate and generalization threshold: $F(1, 864) = 4.06, p < 0.05$. A higher generalization threshold is worse overall, but especially bad for a lower learning rate. This may be due to the fact that the combination of a high generalization threshold and a low learning rate leads to too little knowledge (implicit or explicit) that is in turn detrimental to the performance of maintaining population sizes. See Figure 12.

The analysis using lifespan as the dependent measure led to essentially the same conclusions, and thus details are omitted here.

In all, interactions (1) between cognitive parameters and physical environmental variables (such as food availability) and (2) between cognitive

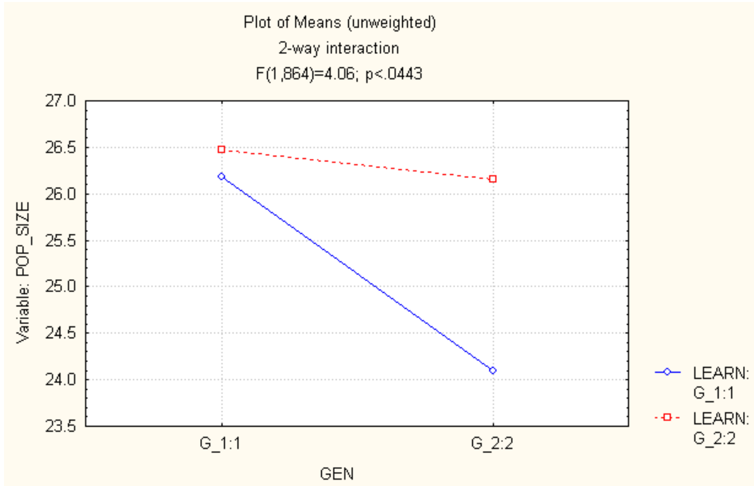

Figure 12: The significant interaction between learning rate and generalization threshold. Learn: $1 = 0.25$, $2 = 0.75$; Gen: $1 = 1.0$, $2 = 3.0$.

parameters and social variables (such as strategy) are fairly common. The interactions indicate the effects cognition has on the survival of agents.

Specifically, the relation between various cognitive parameters and physical environmental variables is such that certain cognitive attributes (such as a higher learning rate) are universally good or bad (for example, a higher learning rate, up to a certain point, is always better for performance), while the effects of some other cognitive attributes (such as probability of using the bottom level) are more dependent on environmental attributes (such as food availability; see the analyses earlier). Existent cognitive attributes may have been selected (through natural evolution) to work within given physical environments, which may be termed the *cognitive-physical dependency* (cf. Cosmides and Tooby 1994, Hutchins 1995). [4]

[4]Note that our simulation so far did not deal with the evolution of cognitive attributes, such as learning rate and so on, which should be tackled in future work.

Similarly, some cognitive attributes (such as a higher learning rate) have universal effects for all social variables (at least in our simulated world), while other attributes (such as probability of using the bottom level) have less universal effects and depend more on social variables (such as strategy). Consequently, the relation between various cognitive parameters and social variables indicates that what social systems, for example, institutions and norms, are adopted may have something to do with cognitive abilities and cognitive tendencies of agents involved (see also Kahan and Rapoport 1984, Conte and Castelfranchi 1995, Tetlock and Lebow 2001, Boyer and Ramble 2001, Atran and Norenzayan 2003, Kluver et al 2005). We may term this relation the *social-cognitive dependency.* This point 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), especially Chapter 1 therein, 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 (e.g., Boyer and Ramble 2001, Lustick 2000, Sun 2006). [5]

The same point can be made of the dependency of social systems (structures and institutions) on attributes of physical environments (e.g., based on the interaction between strategy and food availability, and so on). In turn, we may term this relationship the *social-physical dependency* (Doran et al 1994, Reynolds 1994).

Finally, in the reverse direction of the dependency of social systems on

[5]Note that in this work, we did not deal with the evolution of social institutions in a substantive way. This issue should be tackled in future.

cognitive attributes, cognitive attributes may have been selected (through natural evolution) to work with certain social systems and cultural environments (see Sun 2001, Cosmides and Tooby 1994, Zerubavel 1997, Kluver et al 2003), which may be termed the *cognitive-social dependency.* [6] We may even explore sociocultural principles of cognition (Hutchins 1995, Bourdieu and Wacquant 1992, Durkheim 1962), the opposite of cognitive principles of sociocultural processes as mentioned earlier.

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 (Bar-Yam 1997).

4.3 Summary of Results

In summary, we have clearly shown that, in the context of different social survival strategies and different physical environments, cognition matters. It determines, for instance, which 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 strategy is the best is dependent on the explicitness of cognition of the population. In addition to such social-cognitive dependency, there have also been a host of other dependencies revealed through our analysis. So, even though only very simply representations of sociocultural processes are involved in this work, given the cognitive architecture

[6] But again, our simulation so far did not deal with the evolution of certain important cognitive attributes, such as learning rate and so on.

used, we have nevertheless found significant effects of various interactions.

5 Discussions

5.1 Differences and Similarity

First we may compare our work with that of Cecconi and Parisi (1998). The similarity is obvious. However, significant differences between our work and their work exist. First of all, their model of individual agents was rudimentary. For example, they used simple neural networks, with no individual learning and with no explicit knowledge and explicit reasoning (which humans are certainly known to be capable of; Reber 1989, Sun 2002). Second, their neural networks did not embody realistic cognitive processes and constraints as studied in cognitive science and psychology. Third, although they introduced social norms such as contributing to central stores, we further investigated the involvement of penalty for violating social norms. Fourth, unlike their work, a rigorous statistical analysis was conducted in our work, which revealed some highly interesting points.

The work by Doran and associates, simulating social changes in tribal societies, is also relevant here (Doran and Palmer 1995, Doran et al 1994). They created a simulated landscape over which agents moved. As in our simulation, agents could harvest food resources that were distributed randomly over this space. In their model, each agent was made up of three parts. It had a working memory, in which "facts" were stored, including the current perception. Each agent also had a set of rules, based on which it made decisions. There was a mechanism that matched rules against the current situation as represented by the working memory. Agents could also send each other messages, and through messages, they learned about each other's positions in the space. Agents started with a set of rules for simple actions (such as "if food is adjacent to you, move to it and eat it"). Even though they did not have any pre-conceived notion of groups (or any other social constructs), Doran et al (1994) found that agents formed groups that collectively carried out actions to obtain food through recruiting others to support their own actions. Their simulations were used to explain the emergence of social complexity among pre-historic hunter-gatherers in southwest France during the Upper Paleolithic period about 20,000 years ago (Doran et al 1994). Archaeologists believe that around this time there was a change from societies with a very simple organization of small groups of close relatives to much larger bands with a clearly identified leader. Along with that, there was the development of status differences, as well as the evolution of rituals (Doran et al 1994). In their simulations, these aspects emerged from local interactions among agents, given the deterioration of resource availability.

The differences between their simulation and ours include, first of all, the fact that our cognitive model was significantly more complex and therefore more realistically reflected true human intelligence than their production system inspired agent model. The differences also include the fact that their simulation focused more on one-on-one inter-agent relations, including communications, while our simulations focused more on larger-scale interactions (e.g., through central stores). Their mapping of simulations to historical observations was interesting, although they did not carry out detailed validation (as in, e.g., Sun et al 2005).

Reynolds and associates also simulated tribal societies (see, e.g., Reynolds 1994). For example, one simulation was concerned with the emergence of the Sunay ritual among llama herders in the Peruvian Andes. The herders live at elevation above 4000m in alpine meadow zones, a harsh natural environment. Herding of llamas is the principle economic activity in the region. Many factors, including diseases, predators, and rustlers, threaten the size of existing herds. Sunay happens when one of the herders in the region hosts a ceremony. During the ceremony, a herder may receive the gift of a llama, in return for contributing food and beverages to those present. This is an opportunity for a herder, possibly suffering from a diminishing herd, to acquire fertile females to restore his herd. The question is how this ritual arises, which seems to be disadvantageous to those who give llama to strangers. Their studies showed that there was strong correlation between probability of such cooperation and herd stability. Their simulations provided a detailed explanation for the emergence of the ritual.

Similarly, with an artificial society, Conte and Castelfranchi (1995) explored how the introduction of norms affected macro phenomena in a simple society of food-eater agents. The agents were placed in a two-dimensional environment, with randomly scattered food. Eating food increased an agent's energy, while fighting with another agent over food reduced both agents' energy. When they ran the simulation without norms, all agents acted according to individual utilities, and they frequently attacked each other for food. However, when they ran the simulation with a simple norm "finders keepers", aggressions among agents were dramatically reduced. This led to a higher average strength of agents compared with the simulation without norms. They also found that a society with norms was more equitable, with a smaller variance in the strength of agents. The use of the norm in their simulation was similar in some way to the use of central store strategies in this work. But we also investigated the *maintenance* of such norms, through penalty and other means. In addition, our cognitive model of individual agents (based on CLARION) was much more sophisticated.

Somewhat analogous to the above simulation, Kluver et al (2003) simulated the evolution of social roles—their diversification and maintenance. Their simulations focused on social structures and their effects on social roles of individuals. Likewise, they also simulated the evolution of individual cognition. However, the same as some of the other models reviewed above, their models of cognition were simple and lacked cognitive realism as developed in cognitive science.

5.2 Related Theoretical Issues

Understanding theoretical issues involved in the interaction of cognition and sociality requires computational modeling and simulation, because of the complexity of such an undertaking, and also because of the expressive power of computational models. Unlike mathematical modeling, computational modeling is not limited by traditional mathematical tools. Hence it enjoys greater expressive power. Yet, unlike verbal models, it is precise. It seems to strike a proper balance between rigor and flexibility (expressive power) (see Sun 2006 for more discussions).

Large-scale evolutionary simulation in particular requires computational modeling and simulation, as opposed to mathematical analysis, because of its complexity: It is not just social simulation, or just social simulation with cognitive modeling—it is both plus evolutionary processes on top of them (Wynn 2002, Kenrick et al 2003).

Turning to another issue, 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 needed 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 (as argued by Sun and Naveh 2004).

In addition to offering deeper explanations, cognitive realism can also lead to greater predictive power for social simulation. The ability to produce testable predictions is a vital measure of the usefulness of a simulation. In this regard, there are two significant advantages in using cognitively realistic agents in social simulation. 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 artificial internal parameters of an artificial model or to external measures only.

However, note also that 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 of complex cognitively based simulations (e.g., in terms of precise contributing factors). Complexity also leads to longer running times and hence raises the issue of scalability. Finally, there is the issue of proper choice of theoretical frameworks in relation to cognition, which may hinge on particular ontological conceptions of target phenomena.

In relation to functionalist explanations of social phenomena as alluded to at the beginning, we note that our simulation seemed to have indicated exactly the kind of benefit social simulation provided to functionalist explanations. Beneficial effects of certain social institutions did explain their existence, as our simulations showed how their benefits might have ensured their existence. Thus, certain relevant functional explanations were substantiated. Simulations of evolutionary processes might also provide historical trajectories that led to certain social institutions, also remedying the shortcoming of some functional explanations in terms of lacking historical perspectives. Furthermore, we may even extend functional explanations to cognitive capabilities (e.g., as in evolutionary psychology; Cosmides and Tooby 1994) and to the interaction between cognitive capabilities and social institutions, through an evolutionary perspective.

5.3 Further Issues to Explore

Many more issues need to be explored concerning the interaction of cognition and sociality. First of all, we need to further explore various forms of political systems and their emergence (along the line of Doran et al 1994, but on the basis of individual cognition). That is, we want to explore how environmental factors and cognitive factors give rise to various forms of political systems: for example, in the context of our simulation, (1) the election of a chief versus a hereditary chief, and (2) the authority for strategic decision making (e.g., enactment of law) by the chief, by popular votes, or a combination thereof, (3) means for enforcing social conformity (e.g., through fines and other penalties), as well as other similar issues.

We may also further explore details concerning justice (enforcement) mechanisms: in terms of probability of punishment, amount of punishment, and so on, in case of failing to contribute to the central store when one is supposed to, as well as in other cases.

In these regards, we may consider an additional factor: taking other agents into account explicitly. For example, an agent may perform imitations of other agents, based on observations of others' behaviors.

Beyond the current, much simplified model of evolution (as in, e.g., Cecconi and Parisi 1998, Kenrick et al 2003), we would want to explore more realistic computational models of evolution. For example, we would want more cognitively realistic evolutionary simulation models, that is, evolutionary simulation models that take into account realistic cognitive processes and constraints, as well as their phylogenetic and ontogenetic changes, in the process of capturing social processes. [7] In a more socially realistic simulation of evolution, not only cognitive parameters, but also social institutions, social structures, and culture may be modeled in detail and evolved.

In this process, individual motivational factors should also be taken into full consideration, which may include: following social norms, adhering to ethical values, empathy, need for social acceptance, desire for imitation, and so on. In addition, although CLARION is fully capable of addressing these factors (not covered here, but see Sun 2003 for full details), further developments of cognitive models (e.g., along the line of Carley and Newell 1994) may also be desirable.

Another set of issues concerns validation of simulation findings. We found in our simulation that learning rate, generalization threshold, probability of using the bottom level, and so on matter to performance. Empirical validation of these hypotheses should be attempted. Evidently, it is relatively straightforward in the case of some cognitive factors (e.g., learning rate, which can be plausibly equated with scores on some standardized tests), but trickier in others (e.g., generalization threshold). Future work should address these issues in relation to all of the different types of dependencies as identified earlier.

[7] For example, in CLARION, the cognitive parameters that might be evolved include learning rate, probability of using the bottom level, and so on.

5.4 Concluding Remarks

In this paper, we have aimed for an integration of two separate strands of research: namely, cognitive modeling and social simulation. Such integration, on the one hand, could potentially enhance the predictive accuracy of social simulation models (by taking into account the potentially decisive effects of individual cognition), and on the other hand, could lead to greater explanatory power for these models (by identifying the precise roles of individual cognition in collective social phenomena). In particular, through simulation, we have shown a close interaction between cognitive and social factors, that is, a strong case of the micro-macro link. Cognitive factors may in part determine what social institutions are favored or adopted, and vice versa.

In this work, we tested the approach of cognitively realistic social simulation by deploying the CLARION cognitive architecture. The results from the CLARION based simulations as described above have been encouraging, yielding several findings consistent with the psychological and sociological literatures, as well as hypotheses that are (hopefully) testable in the future.

Acknowledgments

We acknowledge Xi Zhang for for his assistance in conducting this simulation.

References

J. Alexander, B. Giesen, R. Munch, and N. Smelser (eds.), (1987). *The Micro-Macro Link*. University of California Press, Berkeley.

S. Atran and A. Norenzayan, (2003). Religion's evolutionary landscape: Counterintuition, commitment, compassion, and communion. *Brain and Behavioral Sciences*, in press.

R. Axelrod, (1984). *The Evolution of Cooperation.* Basic Books, New York.

Y. Bar-Yam, (1997). *Dynamics of Complex Systems.* Perseus Books, New York.

P. Bourdieu and L. Wacquant, (1992). *An Invitation to Reflexive Sociology.* University of Chicago Press, Chicago.

P. Boyer and C. Ramble, (2001). Cognitive templates for religious concepts: Cross-cultural evidence for recall of counter-intuitive representations. *Cognitive Science*, 25, 535-564.

C. Camerer, (1997). Progress in behavioral game theory. *Journal of Economic Perspectives*, 11 (4). 167-188.

K. Carley and Gasser, L. (1999). Computational organization theory. In: Weiss, G. (ed.), *Multiagent Systems: Modern Approach to Distributed Artificial Intelligence.* The MIT Press, pp. 299-330.

K. Carley and A. Newell, (1994). The nature of social agent. *Journal of Mathematical Sociology*, 19 (4), 221-262.

C. Castelfranchi, (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), 5-38.

F. Cecconi and D. Parisi, (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

R. Conte and C. Castelfranchi, (1995). Understanding the effects of norms in social groups through simulation. In: N. Gilbert and R. Conte (eds.), *Artificial Societies: The Computer Simulation of Social Life*. London, UCL Press.

L. Cosmides and J. Tooby, (1994). Beyond intuition and instinct blindness: Toward an evolutionarily rigorous cognitive science. *Cognition*. 50, 41-77.

J. Doran and M. Palmer, (1995). The EOS project: Integrating two models of Palaeolithic social change. In: N. Gilbert and R. Conte (eds.), *Artificial Societies*. UCL Press, London.

J. Doran, M. Palmer, N. Gilbert, and P. Mellars, (1994). The EOS project: Modeling upper Palaeolithic social change. In: N. Gilbert and J. Doran (eds.), *Simulating Societies*. UCL Press, London.

W. Durkheim, (1895/1962). *The Rules of the Sociological Method*. The Free Press, Glencoe, IL.

N. Gilbert, (1995). Simulation: An emergent perspective. In: *Conference on New Technologies in the Social Sciences*. Bournemouth, UK.

N. Gilbert and R. Conte (eds), (1995). *Artificial Societies*. UCL Press, London.

N. Gilbert and J. Doran, (1994). *Simulating Societies: The Computer Simulation of Social Phenomena*. UCL Press, London, UK.

W. Gray and E. Altmann, (2001). Cognitive modeling and human-computer interaction. In: W. Karwowski (ed.), *International Encyclopedia of Ergonomics and Human Factors* (Vol 1, pp. 387-391). New York: Taylor and Francis, Ltd.

E. Hutchins, (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.

J. Kahan and A. Rapoport, (1984). *Theories of Coalition Formation*. Mahwah, NJ: Erlbaum.

J. Kluver, R. Malecki, J. Schmidt, and C. Stoica, (2003). Sociocultural evolution and cognitive ontogenesis: A sociocultural-cognitive algorithm. *Computational and Mathematical Organization Theory*, 9, 255-273, 2003.

J. Kluver, J. Schmidt, and C. Stoica, (2005). The emergence of social order by processes of typifying: A computational model. *Journal of Mathematical Sociology*, 29, 155-176.

I. Lustick, (2000). Agent-based modeling of collective identity: Testing constructivist theory. *Journal of Artificial Society and Social Simulation*, Vol.3, No.1. http://www.soc.surrey.ac.uk/JASSS/3/1/1.html

J. Mandler, (1992). How to build a baby. *Psychological Review*, 99 (4), pp.587-604.

S. Moss, (1999). Relevance, realism and rigour: A third way for social and economic research. *CPM Report No. 99-56*. Center for Policy Analysis, Manchester Metropolitan University, Manchester, UK.

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

S. Nolfi and D. Floreano, (1999). Learning and evolution. *Autonomous Robots*, 7, (1), 89-113.

T. Ray and J. Hart, (1998). Evolution of differentiated multi-threaded digital organism. In: *Artificial Life*. C. Adami, R. Belew, H. Kitano, and C. Taylor (eds.). MIT Press, Cambridge, MA.

A. Reber, (1989). Implicit learning and tacit knowledge. *Journal of Experimental Psychology: General*. 118 (3), 219-235.

R. Reynolds, (1994). Learning to co-operate using cultural algorithms. In: N. Gilbert and J. Doran (eds.), *Simulating Societies: The Computer Simulation of Social Phenomena*. UCL Press, London, UK.

R. Sawyer, (2003). Multiagent systems and the micro-macro link in sociological theory. *Sociological Methods and Research*, 31, 3, 325-363.

D. Schacter, (1990). Toward a cognitive neuropsychology of awareness: Implicit knowledge and anosagnosia. *Journal of Clinical and Experimental Neuropsychology*, 12(1), pp. 155-178.

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

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

W. Stanley, 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.

R. Sun, (2001). Cognitive science meets multi-agent systems: A prolegomenon. *Philosophical Psychology*, 14 (1), pp. 5-28.

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

R. Sun, (2003). *A Tutorial on CLARION 5.0.* 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*, 25(2), pp. 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*, Vol.7, No.3, June, 2004.

http://jasss.soc.surrey.ac.uk/7/3/5.html

R. Sun and T. Peterson, (1998). Autonomous learning of sequential tasks: Experiments and analyses. *IEEE Transactions on Neural Networks*, Vol.9, No.6, pp.1217-1234.

R. Sun, P. Slusarz, and C. Terry, (2005). The interaction of the explicit and the implicit in skill learning: A dual-process approach. *Psychological Review*, 112 (1), 159-192.

P. Tetlock, (2002). Social-functionalist frameworks for judgment and choice: The intuitive politician, theologian, and prosecutor. *Psychological Review*, 109, 451-472.

P. Tetlock and Lebow, R.N. (2001). Poking counterfactual holes in covering laws: Cognitive styles and historical reasoning. *American Political Science Review*, 95, 829-843.

C. Watkins, (1989). *Learning with Delayed Rewards.* Ph.D Thesis, Cambridge University, Cambridge, UK.

T. Wynn, (2002). Archaeology and cognitive evolution. *Brain and Behavioral Sciences*, 25(3), 389-438.

E. Zerubavel, (1997). *Social Mindscape: An Invitation to Cognitive Sociology.* Harvard University Press, Cambridge, MA.