Kathleen M. Carley, Michael J. Prietula and Zhiang Lin (1998)

Design Versus Cognition:
The interaction of agent cognition and organizational design on organizational performance

Journal of Artificial Societies and Social Simulation **vol. 1, no. 3, <http://www.soc.surrey.ac.uk/JASSS/1/3/4.html>**

To cite articles published in the *Journal of Artificial Societies and Social Simulation,* please reference the above information and include paragraph numbers if necessary

Received: 5-May-98 Accepted: 8-Jun-98 Published: 30-Jun-98

Abstract

The performance of organizations with different structures are examined using multiple computer simulation models, experimental data, and archival data focused on the relation between the way in which the organization is coordinated and its performance. These variations enable the exploration of the role of agent capabilities, and the way in which agent capability and coordination interact to effect performance. Both micro and macro organizational behavior are examined. Results suggest that simpler models of agents are needed at macro levels and more detailed, more cognitively accurate models are needed at micro or small group levels, to generate the same predictive accuracy.

Keywords:
organization theory, agent cognition, validation

Introduction

1.1

Much of organizational research is focused on the determinants of organizational performance. One line of reasoning, falling predominantly within organizational behavior, suggests that organizational performance is tied to the organization being composed of "smart people" (Perrow, 1984; Hastie, 1986; Levitt and March, 1988; Dunbar and Stumpf, 1989; Roberts, 1989). According to this perspective, fill the organization with smarter, better trained, more cognitively capable individuals, overcome the limits of bounded rationality and the organization's performance will improve. Another line of reasoning, falling predominantly in the structural camp, suggests that organizational performance is tied to the organization's design, or as population ecologists and contingency theorists suggest the match between design and task or environment (Woodward, 1965; Lawrence and Lorsch, 1967; Lupton, 1976; Hannan and Freeman, 1987, 1989; Carroll and Hannan, 1990; Houskisson and Galbraith, 1985; Burton and Obel, 1984, 1990). According to this perspective, the organization's structure, procedures, environment, etc. are so constraining that they dictate organizational performance regardless of the behavior of the individuals filling the various roles within the organization. Rarely have these alternative views been contrasted. Rarely has the question been raised as to whether there is an interaction between design and cognition, between structure and individual capability. Recent studies, both computational and experimental suggest that there is an interaction (Lin and Carley, 1993; Carley and Lin, 1995; Carley, forthcoming). In this paper, this interaction is explored in greater detail using a set of organizational designs and varying the "model" of individual cognition.

1.2

In an attempt to move beyond description of organizations as complex systems, organizational theorists have utilized various formalisms to describe organizations and predict behavior. Formal models have been developed using mathematics (Padgett, 1980), simulation (Cyert and March, 1992; Cohen et al., 1972; Masuch and LaPotin, 1989; Lant and Mezias, 1990; Harrison and Carrol, 1991; Carley, 1992), expert systems (Baligh et al., 1990, forthcoming), and formal logic (Beroggi and Wallace, 1993; Salancik and Leblebici, 1988; Leblebici and Salancik, 1989). The complex adaptive nature of organizations makes such formal models a valuable tool for theory development. This is particularly the case when the organization is viewed as an adaptive system composed of agents who are themselves adapting. Such formal models can; a) provide valuable insights into organizational behavior; b) locate errors or gaps in verbal theories; and c) demonstrate which theoretical propositions are logically consistent (i.e., follow from the same base assumptions). Such models let us address issues of scalability (what happens as the size of the organization increases) and response under critical conditions (such as industrial accidents).

1.3

In this paper, computational models, human experiments, and archival data are used to provide insight into the relative impact of cognition and design on organizational performance. This analysis is carried out at both the small group (micro) and the organizational (macro) level. The basic approach is to examine for a small set of organizational designs how changing the way in which the personnel are modeled affects the overall performance of the organization. To orient the reader, a summary of which approaches and which variables are used at both levels is displayed in Table 1.

Table 1: Results used to Illustrate Artificial Organization Project

	CORP			Radar-Soar	Human Experiments	Archival Data
	ELM	SOP	P-ELM			
Micro						
organization structure						
team	O	O	O	O	O	
hierarchy*	O	O	O	O	O	
resource access structure						
segregated						
non-segregated*	O	O	O	O	O	
operating conditions						
errors						
Macro						
organization structure						
team	O	O				O
hierarchy*	O	O				O
resource access structure						
segregated	O	O				O
non-segregated*	O	O				O
operating conditions						
errors	O	O				O

* For the micro study the hierarchy has a single tier - manager and subordinates. For the macro study, the hierarchy may be either a single tier or a multi-tier structure. For the micro study the two non-segregated resource access structures are compared the distributed structure and the blocked structure; whereas, for the macro study segregated access is contrasted with non-segregated access.

1.4

The remainder of this paper is organized as follows. First an overview of the way in which organizations are modeled at both the micro and macro level is provided. Then, the micro level analysis -- models then results -- is presented. This is followed by a presentation of the macro level analysis -- models then results. This in turn is followed by a discussion section.

ORGANIZATIONS AS COLLECTIONS OF ADAPTIVE AGENTS

2.1

Organizations can be characterized as complex adaptive systems composed of intelligent, task-oriented, boundedly-rational, and socially-situated agents and faced with an environment that also has the potential for change (Carley and Prietula, 1994; Prietula and Carley, 1994). As such, task, cognitive agency, and organizational design should be equally important factors in explaining organizational performance. Following is a description of organizations, task, and organizational design, performance and cognition that underlies the following analyses. All organizations modeled fit this description.

Organizations

2.2

Organizations are composed of intelligent adaptive agents. The agent's cognitive capabilities affects whether it learns from experience, what it learns from experience, and how it makes a decision. The organization is faced with a sequence of problems such that each problem is similar to, but not identical to, previous problems faced by the organization.[1] The organization might face externalities such as missing information, erroneous information, turnover, missing personnel, and so forth. Individual agents within the organization gather information about the task, make recommendations to other agents in the organization, and eventually, the organization resolves the problem. In general, the problem is sufficiently complex that no one agent is capable of doing the problem on its own. Thus the task must be done in a distributed fashion.

Radar Task

2.3

Organizational performance is a function of the task being performed. A typical task faced by many organizations, most of the time, is a classification choice task. In a classification-choice task, the decision maker gathers information about the problem, classifies it, and then on the basis of this classification makes a decision (choice) about the task. Typically the number of choices are limited, such as in a buy or sell stock decision, hire or not hire personnel decision, and fund, partially fund, or not fund budgeting decision.

2.4

The specific task used in all experiments in this paper, virtual and real, is the ternary choice task. The ternary choice task is often referred to as a radar task. In this task (which is highly stylized), the organization must determine whether the "blip" on the screen is a hostile aircraft, a flock of geese, or a civilian aircraft. The radar task is a fixed choice task similar to that used by many researchers interested in organizational design (Hollenbeck et al., 1991; Ilgen et al., 1991; Tang et al., 1990).

2.5

In the radar task, there is a single aircraft in the airspace at any one time. This aircraft is uniquely characterized by nine different characteristics (such as speed) such that each characteristic can take on one of three values (low = 1, medium = 2, or high = 3). The organization must determine for each observed aircraft whether that craft is friendly, neutral, or hostile. There are a total of 19683 possible aircraft and 30 problems are chosen randomly without replacement from this set. The true state of the aircraft is determined by simply adding the values of the 9 characteristics. If this sum is less than 17 the true state is said to be friendly, if this sum is greater than 19 the true state is said to be hostile, otherwise the true state is said to be neutral.[2] The true state of the aircraft is not known a priori by the organization or by any of the organizational personnel.

2.6

The radar task is a distributed decision making task such that each problem is similar to (but not identical to) previous problems. The problems are sufficiently complex that no one agent has access to all the information necessary to make the decision. Decisions are made by integrating decisions made by distributed agents on different aspects of the task rather than by consensus (Bond and Gasser, 1988). The organization is assumed to have sufficient personnel to access all information and who can evaluate this information and coordinate with all organizational members to make this determination before the aircraft changes position in the airspace. Each personnel who actually evaluates the aircraft can only evaluate a few characteristics (three). In fact, these agents do not actually observe a radar screen. Rather, they are essentially receiving a report that simply states information of the form "speed is low, range is high, angle is medium." The agent takes a snapshot of the plane, determines its value on each of the three characteristics that it examines, and then makes a recommendation (or prediction) as to whether it thinks the aircraft will be friendly, neutral, or hostile.

Organizational Design

2.7

The recommendations made by organizational personnel are the basis for the organization's decision. Organizations combine and utilize personnel decisions in different ways depending on the organization's design. For example, in teams personnel decisions are treated as votes and the organizational decision is the majority decision. In hierarchies, information is passed up from subordinates to superiors an the organizational decision is the CEO's (chief executive officer's) decision. All superiors only examine the recommendations of their subordinates. When there are multiple tiers in a hierarchy mid-level managers make a recommendation of the state of the aircraft on the basis of the recommendations passed to them by their subordinates.

2.8

Which personnel have access to which information on which of the nine task characteristics is defined by the resource access structure. A typical scheme is the segregated structure where each agent sees information on different characteristics. In contrast, in the distributed structure each analyst sees a completely different set of three characteristics. While the analysts have some information in common they should still have distinct mental models. Another non-segregated structure is the blocked structure in which three analysts all see exactly the same information.

Performance

2.9

Organizational performance has been characterized in a variety of ways including efficiency, effectiveness, and perceived effectiveness. In this paper, performance is characterized as accuracy. For the radar-task, the true state of the aircraft can be either friendly, neutral, or hostile as can the organization's decision. The organization's decision is correct if it is the same as the true state of the aircraft. Performance, for a set of problems is simply the percentage of these problems for which the organization makes the right decision. This measure of performance can be thought of as a prediction about likely it is that the organization will make the right decision in any particular case.

Cognition

2.10

The approach taken herein is that individuals are at least boundedly rational agents (March and Simon, 1958). These bounds are set both by their cognitive architectures and by their position in the organization (Carley and Newell, 1994). Physiological and social constraints are equally important in determining the individual's performance. What the individual knows is a function of the organization's design and the individual's position within the organization. For example, while most individual's in the organization only know a few pieces of information about the specific problem, their manager only knows the recommendations of his or her subordinates. How the individual agents use this information to make a decision depends on the agents cognitive architecture (or as in this paper, the specific "model" of decision making). More details on particular models of agent cognition are provided in the following section.

2.11

Individual agents are autonomous but limited. Agents are autonomous in that their processes for handling information are entirely self contained and they can act on the information that they receive and make decisions without acquiring process from other agents. Agents are limited as no one agent can do the task perfectly, totally on his or her own.

MICRO ANALYSIS

3.1

The first analysis is at the micro or group level. A series of experiments, virtual (using computational models) and real (using humans), were run in which the organizational design and agent cognition were varied (see Figure 1). Five different "models" of the agent were examined each varying in the level or type of cognition (the complexity and the realism of the agent model). These agent models were CORP-ELM, CORP-P-ELM, CORP-SOP (Carley and Lin, 1995; Carley, forthcoming), Radar-Soar (Ye and Carley, 1995) and humans (Carley and Prietula, 1992; Carley, 1993). Four different organizational designs were examined each varying in the organizational structure (who reports to whom) and the resource access structure (who does what). A set of matched organizations facing the same set of 30 tasks, in the same order, with the same organizational design but varying in the agent model were examined. This micro analysis can be thought of as an experiment with a five (agent model) by four (organizational design) design. For the micro-analysis all organizations were operating under ideal, error free, operating conditions. That is, there was nothing going wrong, no missing or erroneous information, no missing personnel, no turnover.

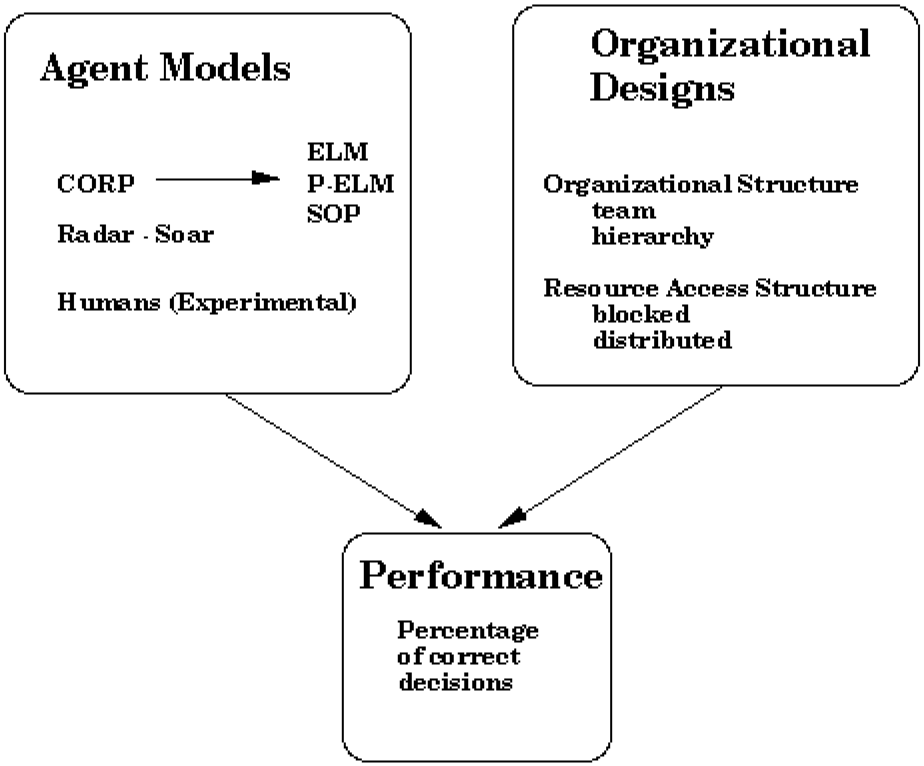

Figure 1. Micro Level Experimental Design

3.2

The general characteristics of the cognitive models will be described. An exhaustive description is beyond the scope of a single paper. Readers are directed to Carley and Lin (1995, forthcoming) for a more detailed description of CORP; to Ye and Carley (1995) for a more detailed description of Radar-Soar, and to Carley and Prietula (1992) for a more detailed description of the human experiments. As the move is made from CORP, to Soar, to human experiments, at the micro level, all details of the task and organizational design are identical.

Alternate Models Of Agent Cognition

3.3

Five alternate "models" of cognition were examined. These were CORP-ELM, CORP-P-ELM, CORP-SOP, Radar-Soar, and Humans. It might seem odd to think of humans as a model of cognition; but within this empirical framework, the issue is whether the procedure for making decisions and utilizing experience interacts with the organizational design to influence organizational performance. As such, human decision making is important to consider. While we do not have a precise description of how humans make these decisions, as we do with the various artificial agents, we still know what decisions they make and like artificial agents, we can place human in various organizations and observe their behavior. Thus we can contrast human behavior with the behavior of various artificial agents. It is in this sense, that humans are a "model" of cognition. These models vary in their complexity and realism such that CORP-SOP is the least realistic and simplest and Humans are the most realistic and complex. The Radar-Soar agents are less realistic than humans and more realistic than the CORP models.

3.4

CORP agents are limited in their intelligence and adaptive procedures to these specific tasks. Further, there are three distinct types of agents in CORP that vary in whether they make decisions by: (1) by following the organizationally prescribed standard operating procedure (SOP), (2) by perfectly following the dictates of their own personal experience (ELM), or (3) by guessing based on a probabilistic estimate of the answer garnered through personal experience (P-ELM). Whereas, Soar agents are thought to be more generally intelligent than the CORP agents as the Soar architecture can be used to solve multiple types of tasks. Finally, humans are expected to be the most generally intelligent and adaptive of all these agents. In addition, SOP agents are not adaptive and ELM agents have the fewest bounds on their intelligence. General intelligence is seen as being affected by the multiplicity of tasks that the agent can do, and the multiplicity of mechanisms for adaptation held by the agent. The more tasks the agent can perform, the more adaptive mechanisms available to the agent, the greater the agents "intelligence." We want to distinguish here between being right, and being capable. That is, more generally intelligent agents may not be more likely to make the right decisions.

Corp

3.5

CORP is a simulation framework for examining the relationship between organizational design and environment in affecting organizational performance. CORP models are artificial organizations composed of complex adaptive agents with task specific abilities. Within CORP the researcher can choose the method of decision making and learning employed by the agents, the type of organizational design, the type of task environment, and the type of task difficulties faced by the organization. The choices available to the research for examining type of organizational design, the type of task environment, and the type of task difficulties faced by the organization are a super-set of those available in the Radar-soar, human experiments, and archival data.

3.6

Within CORP there are three simple agent models: ELM, P-ELM, and SOP. ELM -- the experiential learning model -- is a model of individual cognition such that the individual considers new information, compares it to what it has seen in the past, and then makes as its decision that decision that was most often correct in the past. Each time period the agent increments its memory by altering the frequency with which a particular decision was correct. The experientially trained agents have memories containing an exhaustive list of all incoming patterns of information and the number of problems they have had experience with that fall into each pattern and the percentage of times that they observed that pattern that a particular outcome (friendly, neutral, or hostile) was observed. As the agent gains experience these percentages change. Thus, an agent may know that 50% of the time that it saw pattern xxy the aircraft was truly friendly, 30% of the time it was neutral, and 20% of the time it was hostile. For agents trained to perfectly follow their experience they will make as their decision the choice with the highest probability. Note, this makes these agents, effectively, insensitive to sample size and overconfident in their decisions.

3.7

P-ELM -- the probabilistic experiential learning model -- is a model of individual cognition such that the individual considers new information, compares it to what it has seen in the past, and then determines its decision stochastically using the frequency with which a decision has been right in the past as the probability of making that decision this time. As with ELM, the P-ELM agent increments its memory each time period by altering the frequency with which a particular decision was correct. For agents trained to probabilistically follow their experience they will choose their decision with the probability associated with the likelihood of that decision in the past. (Note, this makes these agents insensitive to sample size and capable of making decisions in contradiction to the historical evidence.) Thus, if an agent sees the pattern xxy and is trained to perfectly follow its experience it will claim that the aircraft is friendly; in contrast, if the agent was trained probabilistically it has only a 50% chance of claiming that the aircraft is friendly.

3.8

SOP -- the standard operating procedure model -- is a model of individual cognition such that the individual simply applies a standard operating procedure (majority rule) to the pattern of new information to determine what its decision will be. Operationally trained agents do not adjust their behavior given feedback. Rather, these operational agents follow a set of standard operating procedures (SOPs) based on a majority rule classification. This operational training allows the agent to employ a set of general purpose static heuristics. Agents trained operationally are assumed to have perfect memory for the SOPs and to execute them perfectly every time. These SOPs are based on the sum of the values associated with those pieces of information available to the agent.

Radar-Soar

3.9

The Soar agents represent an added level of complexity and realism in the model of agency. Soar is a detailed model of cognition that has been shown on numerous individual tasks to behave as humans do, and which characterizes all decision making as search through problem spaces (Laird et al., 1987; Laird et al. , 1986a, 1986b; Newell, 1990).

3.10

Radar-Soar is a simulation system designed to allow the researcher to compare and contrast the performance on the radar task of organizations of Soar agents (Ye and Carley, 1995; Papageorgiou, 1992). Each Radar-Soar agent is built as a separate agent in Soar (Laird et al., 1987; Laird et al., 1986a, 1986b) and learns incrementally on the basis of experience. Soar is arguably a model of general cognition (Newell, 1990). Unlike the experiential agents in CORP, the Radar-Soar agents do not keep track of frequencies. Rather they build models relating a specific problem to the specific feedback they receive for this problem. When a new problem arises they locate all models, find those that match the current problem the best, and then probabilistically choose among the models with the best match. They then stochastically choose a decision from those outcomes that have historically been the most correct.

3.11

Within Soar, agents have problem-spaces, areas of expertise with associated operators (actions they can take) and states (current mental model of the situation in that area). In the current design, the Radar-Soar agent has the following task-related problem spaces including interpret command, make decision, communicate, update models. Typical operators (actions the agents can take) are "ask question", "read scanner", "give command", and "update model".

Human Experiments

3.12

Subjects in the human experiments took part in an organization trying to do the radar task (Carley and Prietula, 1992, Carley and Prietula 1993). The experiment had a staged design. First, data was collected from all subjects at the subordinate position. Second, the results of these subordinate's recommendations were combined in different ways and given to the subjects in the managerial positions. The experiment was run physically by having subjects, individually, log onto a "radar" program on a Macintosh that trained them, provided them with information on the nature of the aircraft, recorded their responses, and provided them with feedback on the accuracy of their recommendations. Results were gathered in such a fashion that only one subject needed to be present at a time.

3.13

Subjects were given the set of problems; of which in this paper our concern is with only the first 30 problems.[3] Each subject acting as subordinate saw information on only three characteristics for each of these 30 problems. The subjects were given information on one aircraft at a time and the information was in the form "speed is low, range is high, angle is medium." This is exactly the same information that the artificial agent would have seen. The subjects were then asked if they think the aircraft was friendly, neutral, or hostile. After the subjects provided their recommendation, they were asked for their confidence in their decision. Each subject acting as a manager saw for these same 30 problems only information on the opinion of his or her nine subordinates. The subjects as managers were given information on one aircraft at a time and the information was in the form "subordinate one thinks the aircraft is friendly, subordinate two thinks the aircraft is neutral" and so on for all nine subordinates. The subjects as managers were then asked if they thought the aircraft was friendly, neutral, or hostile. After the subjects as managers provide their recommendation, they were asked for their confidence in their decision. In all cases, the subjects' decision, confidence, and the time to make that decision were collected automatically.

Alternate Organizational Designs

3.14

The organizational designs had either a team or a single tier hierarchical organizational structure and a blocked or distributed resource access structure. In a team the organizational decision is made by voting and is the majority vote of the personnel. In the hierarchy all personnel report to a single supervisor and the organizational decision is that supervisor's decision. In both cases there are nine personnel (subordinates) who gather and analyze information on the task. In a blocked structure, each subordinate sees three pieces of information and three subordinates all see exactly the same three pieces of information for each problem. In a distributed structure each subordinate sees three pieces of information and no two subordinates see exactly the same three pieces of information for each problem. In both cases each piece of information is analyzed by three different subordinates.

Results

3.15

Does agent cognition interact with organizational design in affecting organizational performance? The answer appears to be yes. As can be seen in Table 2, which organizational design exhibits the highest performance depends on the type of agent. For example, while humans exhibit highest performance in a team with a distributed structure and lowest performance in a hierarchy with a blocked structure, P-ELM agents tend to exhibit highest performance in a team with a blocked structure and lowest performance in a hierarchy with a distributed structure. Secondly, we see that adding cognitive constraints, increasing the general intelligence, or increasing the adaptiveness of the agents tends to decrease organizational performance. Third we see that CORP agents following SOPs are most suited to this task; i.e., performance is highest under all organizational designs. Finally, a detailed analysis of the results suggest that none of the artificial agent models are accurately modeling which specific decisions humans are getting right and or wrong. However, each of the models reflects different aspects of the average trends in human behavior. Thus, while none of these models are adequate for exactly describing the behavior of individuals, they are reasonably adequate for describing the average behavior of collectives of individuals

Table 2: Average Performance, Model by Organizational Design

Agent	**Organizational Design**			
	Team		Hierarchy	
	Blocked	Distributed	Blocked	Distributed
CORP-ELM	88.3%	85.0%	45.0%	50.0%
CORP-P-ELM	78.3%	71.7%	40.0%	36.7%
CORP-SOP	81.7%	85.0%	81.7%	85.0%
Radar-Soar	73.33%	63.33%	63.33%	53.33%
Human	50.0%	56.7%	46.7%	55.0%

MACRO ANALYSIS

4.1

Now let us consider organizational behavior at a more macro level. Here the concern is with whether the formal CORP model is a reasonable approximation of the macro level behavior of real organizations. The particular focus is on organizations faced with an industrial crisis. The research strategy was to take a set of 69 real organizations, characterize them given each of the components in the CORP model, then simulate a matched artificial organization which matched the real organization on all of these components (see Figure 2). Four different "models" of the agent were examined each varying in the way the agents make decisions - CORP-ELM, CORP-SOP, Human-Experiential, and Human-SOP. ELM and experiential agents make decisions on the basis of their personal experience. We can think of organizations composed of agents who follow their experience as experiential organizations. SOP agents make decision by following standard operating procedures. We can think of organizations composed of agents who follow SOPs as organizations of operationally trained agents. Four different organizational designs were examined each varying in both organizational structure (who reports to whom) and the resource access structure (who does what). The organizations had either a team structure or a hierarchical structure. If a hierarchy, the number of levels in the organization was either one or two depending on the level of structure in the real organization. The organizations had either a segregated or non-segregated resource access structure. If segregated, then each member of the organization has a unique job and there is no overlap across jobs. If non-segregated then the degree of overlap (low or medium), and the type of overlap (blocked or distributed) depended the way resources and access to information were shared in the real organization.

4.2

For each organization considered the individuals in the organization were trained to cope with these types of crises, but what is peculiar in each crisis was the information processing constraints that made the operating conditions less than ideal, such as turnover and communication errors. The performance of each artificial organization was estimated using Monte Carlo techniques. Then the performance of the artificial organization was compared with that of the real organization. Implicitly this comparison addressed the question -- is this computational model adequate for predicting overall organizational performance. In this paper, a detailed answer to this question is not provided. For a more detailed answer refer to Lin (1994) and Carley and Lin (1994). Herein, the specific focus is on determining whether there is an interaction between agent cognition and organizational design at this more macro level.

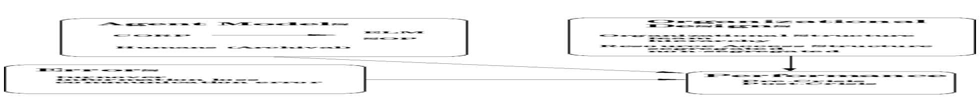

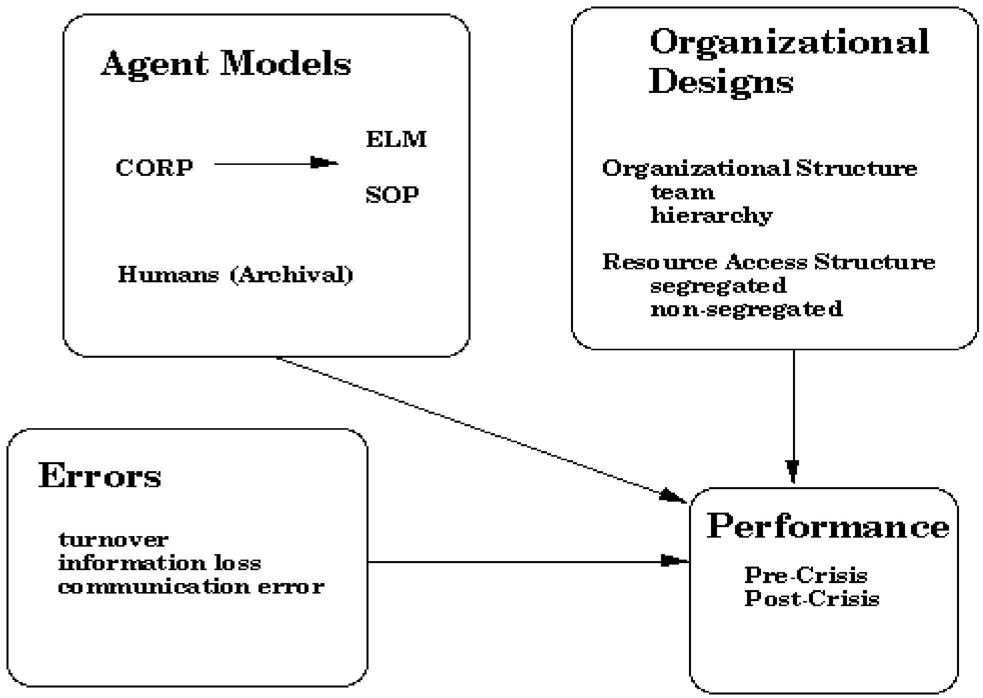

Figure 2. Macro Level Research Design

Archival Data

4.3

The archival data is drawn from numerous academic and news source accounts of 69 different organizations faced with a technological crisis. Examples of these crises are the Exxon Valdez, and the U.S. Vincennes. This data was coded on each of the variables in the CORP framework (Lin, 1994; Lin and Carley, 1993). Specifically, these organizations tend to be facing various information processing constraints such as missing information, erroneous information, or missing personnel. In fact, each of the archival organizations tends to be simultaneously faced with multiple problems; e.g., information and personnel may be unavailable. Each organization was characterized by whether the personnel were directed to act on the basis of their experience or were expected to following standard operating procedures. In addition, each organization was characterized by its dominant organizational structure and resource access structure. Most of these organizations are essentially operating in a biased environment; i.e., certain outcomes are much more likely than other outcomes. Performance during crisis was measured in terms of the socio-economic-environmental impact of the crisis and the chance for avoidance. Low organizational performance (=

1) occurs when the organization did not handle the crisis as well as it could have; i.e., actual severity of the crisis is high and the potential is low or medium. In contrast, high organizational performance (= 3) occurs when the severity of the crisis is low and the potential is medium or high. All other cases are defined to be moderate performance (= 2). For the Vincennes incident, the organizational performance is low because the actual event was severe (many people died) and because the crisis was, to a large extent, avoidable. Performance was measured on a three point scale (1 = low, 2 = medium, and 3 = high).

Simulated Data

4.4

Using CORP a set of 69 artificial organizations were constructed which matched the real organizations in terms of the following characteristics: organizational structure, resource access structure, operating conditions, task environment, and personnel training. There are two main differences between the real and simulated organizations. In the real organizations, the organization was often simultaneously being faced by multiple limits to its operating conditions (e.g., missing information and missing personnel); whereas, in the simulated organization, only the dominant factor limiting the operating environment of the real organization is simulated. The main effect of this difference should be to raise the overall performance level of the artificial organizations. Secondly, performance is measured in a somewhat different fashion. For the real organizations, pre and during crisis performance was measured as an indicator of the general or average accuracy of the organization's decision. Basically, the assumption was made that the further the "actual severity of the crisis" was from the "potential severity" the greater the organization's general accuracy in addressing problems. In the artificial organizations, performance was measured as an ensemble average accuracy across a set of problems. Performance under crisis was measured across just those events defined to be hostile and where there were one or more errors that limited the operating conditions. For each organization, its level of accuracy was remapped into low, medium or high using the following procedure. The number of low, moderate, and high performers in the 69 real organizations was determined (28 low, 31 moderate, and 10 high). Then the 69 artificial organizations were ranked from low to high in their performance. These artificial organizations were categorized as low performers if they were in the bottom 28, the next 31 were categorized as moderate performers, and the top 10 as high performers.

Results

4.5

Does agent cognition interact with organizational design in affecting organizational performance at the macro level? The answer appears to be yes. As seen in Table 3 organizations with experientially and operationally trained agents exhibit substantially different performance. Moreover, organizations with operationally trained agents are less affected by organizational design; i.e., performance is similar under all organizational designs. In this sense, operational training tends to make the organization more impervious to organizational design. There is generally high agreement between the simulated and archival data. Organizations, whether CORP or human, tend to exhibit lower performance if they are following standard operating procedures. Further, teams tend to outperform hierarchies and organizations with non-segregated resource access structures tend to outperform those with segregated structures.

Table 3: Performance Under Crisis Conditions, Model by Organizational Design

	Organizational Structure			
	Team		Hierarchy	
	Experiential	Operational	Experiential	Operational
CORP	3.00	1.50	2.35	1.41
Human Archival	3.00	1.50	2.35	1.41
N	1	4	20	44
	Resource Structure			
	Segregated		Non-Segregated	
	Experiential	Operational	Experiential	Operational
CORP	2.30	1.40	2.45	1.60
Human Archival	2.10	1.42	2.64	1.80
N	10	43	11	5

4.6

This research demonstrates the value of simultaneously considering the organizational structure, the resource access structure, and the training/decision making procedure in explaining general organizational performance. These results suggest that even simple models of organizational performance can, at the macro level, predict general organizational performance under crisis conditions. Whether this model can also predict non-crisis performance, and the relative degradations due to different types of errors is a point for future research.

5.1

The question "Does agent cognition interact with organizational design in affecting organizational performance?" has been addressed at both a micro and macro level. At both levels, both the agent's cognitive style and decision making capability and the organization's design influences organizational behavior. Even more importantly, at both levels, we see that there are complex interactions between cognition and organizational design such that the relative impact of each on performance depends on the other. Consequently, both cognition and design can be thought of as tools in the manager's toolkit that can be manipulated, at least partially, to improve organizational performance. To an extent, organizations can substitute cognitive capability for design or design for cognitive capability.

5.2

The results provided are drawn from the artificial organizations project. This project is relatively unique as it is one of the few studies where computational models are contrasted with both empirical data and other computational models. This research demonstrates that simple models of organizational behavior have the ability to predict and characterize organizational behavior at a macro level. However, these same models, are less accurate as predictors of micro organizational behavior and are nowhere as accurate as predictors of individual behavior. Rather, more detailed and cognitively accurate models are needed to predict micro level organizational behavior. This seems to suggest that different models may be adequate for predicting individual and organizational level performance. Whether in fact a completely accurate model of individual performance would generate accurate organizational performance is not known, and the results shed little light on this question. These results suggest that simple models of agent behavior may not be adequate for accurately predicting organizational performance at the small group or micro level if personnel are undifferentiated; but may be accurate at a more macro level. In particular, simple models of agent behavior, such as CORP-ELM, may be adequate at both the group and the organizational level when we are dealing with more complex organizational structures and when the information is filtered through a chain of command. Finally, these results demonstrate that the development of more veridical organizational models is greatly facilitated by doing both cross model comparisons or docking (Axtell et al., 1996) and comparisons with empirical data or validation.

5.3

This research increases our understanding of the role of agent capabilities in affecting organizational performance. It moves us a step closer to developing artificial social agents. Plural-Soar agents (Carley and Prietula, 1992; Prietula and Carley, forthcoming) were the first instantiation of Soar as a social agent. As Carley et al. (1992) noted, Plural-Soar agents exhibited at a rudimentary level many of the capabilities needed by agents in social and organizational settings. Radar-Soar agents are still rudimentary but are more social in nature than are Plural-Soar agents. In theory, this should move us a step closer in our ability to observe emergent social phenomena (Carley and Newell, 1994). The Radar-Soar agents are also more socially capable than CORP agents. Within the simulations reported there is little reason or opportunity for the agents to communicate, other than to pass opinions. Thus, it is not clear whether the increase in social capability in moving from CORP to Soar actually had an effect on task performance. Just because agents are capable of performing an action does not mean that they take that action. Future research should look more closely at the issue of communication.

5.4

What this research has shown is that in order to explain and predict organizational behavior in terms of the behavior of each individual in the organization, greater attention needs to be paid to the way in which the individual agents are modeled even for relatively simple tasks. Even as organizational performance can be improved by finding the right match between organizational design and task, it may be that organizational performance can be improved by finding the right match between individual cognition and task. Seemingly minor differences in cognition can have important repercussions for social and organizational behavior. Seemingly minor differences in organizational design can have important repercussion for the need for different cognitive capabilities in the agents in the organization (Carley and Svoboda, forthcoming).

5.5

Indeed the strong interaction between the agent's cognitive and decision making capabilities and the organization's design suggests that organizational theorists should attend more to the way in which organizational personnel are characterized and the position (both vis other people and tasks) in which they are placed in the organization. This suggests that organizational researchers should be careful when making attributions regarding organizational performance to determine whether performance differences are attributable to some fundamental feature of the agents decision making style, the way in which they were trained, or to some structural feature of the organization including the task and the organization's design. There are several related ideas here. First, in order to achieve a certain level of organizational performance, design and cognition may be able to substitute for each other. Research in cognitive psychology demonstrates that human performance is affected both by the content of what the individual knows and how that content is structured or organized. This research is suggesting, a somewhat parallel notion, that organizational performance is affected both by the capabilities of the nodes (agents in particular roles) and by the way the nodes are linked together. Second, particular types of cognition may mitigate the impact of design. Recall that when personnel follow SOPs the variance in organizational performance by organizational design was reduced. This suggests a particularly intriguing hypothesis that the impact of organizational design may be mitigated by reducing the cognitive ability and decision making capability of the agent. Third, particular designs may mitigate the impact of cognition. Certain organizational designs may be so cognitively constraining that they reduce all intellective options for the agent and so reduce the impact of different cognitive capabilities. Fourth, organizations at different stages in their development, to achieve the same or better performance may find the need to restructure themselves if they find that they are hiring personnel with different capabilities. By the same token, organizations that begin to employ artificial agents (such as information systems and expert systems) to do certain tasks may also find the need to employ different organizational structures than they do with human agents.

5.6

Most studies in the social sciences point to the role of social factors influencing cognition. In contrast, in this paper, we are suggesting that cognition (particularly memory and the type of reasoning available to the agent) can influence social and organizational phenomena. Moreover, in the case of artificial agents or styles of training human agents, cognition can be "chosen" to optimize the value of other social factors such as organizational design. Cognition, in this sense, becomes a tool that can be manipulated and used by organizations to achieve a particular performance outcome.

5.7

The comparisons described herein demonstrate that computational models based on a structural information processing approach are informative for understanding organizations in general, and the relationship between individual cognition and decision making capabilities and organizational design in particular. There are standard claims that individual's are boundedly rational and that organizations satisfice rather than optimize in part because individuals are cognitively limited. These claims generally equate cognition with individuals acting in a boundedly rational fashion, and focus on internal psychological or physiological limits to rationality. In this paper, the results go beyond these claims by pointing out that organizations composed of personnel with different types of cognitive limitations exhibit different levels of performance and may be more or less advantaged by different organizational designs. On the one hand, our results demonstrate that social artifacts, such as roles and organizational design, place bounds on individual cognition as they limit access to information, what lessons individuals can learn from experience, and the applicability of cognitive rules. On the other hand, our results are suggesting that there is a certain "intelligence" bound up in these social artifacts that can supplement, or possibly be exchanged for, the cognitive capability of the individual agents within the organization.

Acknowledgements

This work was supported in part by the Office of Naval Research (ONR), United States Navy Grant No. N00014-93-1-0793 and by the NSF Grant No. IRI-9111804. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Office of Naval Research, the National Science Foundation, or the U.S. government.

Notes

[1] Tasks such as these have been referred to as quasi-repetitive integrated decision making tasks (Carley, 1991; 1992)

[2] This is referred to as an unbiased (all outcomes equally likely) decomposable (all characteristics are equally weighted and have equal impact in determining the outcome) task (Carley and Lin, %%%)

[3] These 30 problems were the first of four sequences of problems seen by the subjects. Overall, the subjects saw 120 problems such that the first set contained 30 problems with characteristics a, b, c with feedback, the second set contained 30 problems with characteristics a, b, c without feedback, the third set contained 30 problems with characteristics d, e, f with feedback, and the fourth set contained 30 problems with characteristics d, e, f without feedback and with equipment failures (so that the subject misses some incoming information). In this paper, only data from the first 30 problems, those for which feedback was received is considered.

References

AXTELL, R., R. Axelrod, J.M. Epstein, and M. D. Cohen (1996), "Aligning Simulation Models: A Case Study and Results." *Computational and Mathematical Organization Theory*, 1(2): 123-141.

BALIGH, H.H., R.M. Burton and B. Obel (1990), "Devising Expert Systems in Organization Theory: The Organizational Consultant," In Masuch M. (Ed.), *Organization, Management, and Expert Systems*. Berlin: Walter De

Gruyter.

BEROGGI, G. E. G. and W. A. Wallace (1994), "A Decision Logic for Operational Risk Management." In Carley, K. M. and M. J. Prietula (Eds.), *Computational Organization Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

BOND, A. H. and L. Gasser (Eds.) (1988), *Readings in Distributed Intelligence*. San Maeto, CA: Morgan Kaufmann Publishers.

BURTON, R. M. and B. Obel (1984), *Designing Efficient Organizations: Modelling and Experimentation*. New York NY: Elsevier Science.

BURTON, R. M. and B. Obel (1990), "Devising Expert Systems in Organization Theory: The Organizational Consultant," In M. Masuch (Ed.), *Organization, Management, and Expert Systems*, Berlin: Walter De Gruyter.

CARLEY, K. M. (1991), "Designing Organizational Structures to Cope with Communication Breakdowns: A Simulation Model," *Industrial Crisis Quarterly* 5: 19-57.

CARLEY, K. M. (1992), "Organizational Learning and Personnel Turnover," *Organization Science* 3(1): 2-46.

CARLEY, K. M. (1993), "Plural-Soar: Towards the Development of a Cognitively Motivated Theory of Organization," In Proceedings of the 1993 Coordination Theory and Collaboration Technology Workshops: Symposium, National Science Foundation.

CARLEY, K. M. (forthcoming), "A Comparison of Artificial and Human Organizations," *Journal of Economic Behavior and Organization*.

CARLEY, K. M., J. Kjaer-Hansen, M. Prietula and A. Newell. (1992) "Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior." in Masuch M. & M. Warglien (Eds.), *Artificial Intelligence in Organization and Management Theory* . Amsterdam: Elsevier, 87-118. CARLEY, K. M. and Z. Lin (1994), "Relating Shifts in C2 Structure to Performance," Proceedings of the 1994 Symposium on Command and Control Research.

CARLEY, K. M. and Z. Lin (1995), "Organizational Designs Suited to High Performance Under Stress," *IEEE - Systems Man and Cybernetics* 25(1): 221-230.

Carley, K. M. and Z. Lin (1997), "A Theoretical Study of Organizational Performance under Information Distortion."

Management Science. 43(7): 976-997.

Carley, K. M. and A. Newell (1994), "The Nature of the Social Agent." *Journal of Mathematical Sociology* , 19(4): 221-262.

CARLEY, K. M. and A. Newell (1994), "The Nature of the Social Agent," *Journal of Mathematical Sociology* 19(4): 221-262.

CARLEY, K. M. and M. J. Prietula (1992), "Toward a Cognitively Motivated Theory of Organizations," Proceedings of the 1992 Coordination Theory and Collaboration Technology Workshop, Washington D.C.

CARLEY, K. M. and M. J. Prietula (1993), "Plural-Soar: Towards the Development of a Cognitively Motivated Theory of Organizations," in Proceedings of the 1993 Coordination Theory and Collaboration Technology Workshop . Symposium conducted for the National Science Foundation, Washington, D.C.

CARLEY, K. M. and M. J. Prietula (1994), "ACTS Theory: Extending the model of bounded rationality," In Carley, K. M. and M. J. Prietula (Eds.), *Computational Organization Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

CARLEY, K. M. and D. Svoboda (forthcoming), "Modeling Organizational Adaptation as a Simulated Annealing Process," *Sociological Methods and Research*.

CARROLL, G. R. & M. T. Hannan (1990), "On the maturation and aging of organizational populations," in J. V. Singh (ed.) *Organizational evolution: new directions*. Beverly Hills, CA: Sage Publications.

COHEN, M. D., J. G. March and J. P. Olsen (1972), "A Garbage Can Model of Organizational Choice," *Administrative Sciences Quarterly* 17(1): 1-25.

CYERT, R. and J. G. March (1992[1963]), *A Behavioral Theory of the Firm* (2nd ed.). Cambridge, MA: Blackwell Publishers.

DUNBAR, R. and S. A. Stumpf (1989), "Trainings that Demystify Strategic Decision-Making Processes," *Journal of Management Development (UK)*, 8(1): 36-42.

HANNAN, M. T. and J. Freeman (1987), "The Population Ecology of Organizations," *American Journal of Sociology*, 82: 929-964.

HANNAN, M. T. and J. Freeman (1989), *Organizational Ecology*. Cambridge, MA: Ballinger.

HARRISON, J. R. and G. R. Carrol (1991), "Keeping the Faith: A Model of Cultural Transmission in Formal Organizations," *Administrative Science Quarterly*, 36: 552-582.

HASTIE, R. (1986), "Experimental Evidence on Group Accuracy," In Jablin, F. M., L. L. Putnam, K. H. Roberts and L. W. Porter (Eds.), *Handbook of Organizational Communication: An Interdisciplinary Perspective*, Beverly Hills, CA: Sage.

HOLLENBECK, J. R., D. J. Sego, D. R. Ilgen, and D. A. Major (1991), Team Interactive Decision Making Exercise for Teams Incorporating Distributed Expertise (TIDE2): A Program and Paradigm for Team Research. Tech. Rep. No. 91-1, East Lansing MI: Michigan State University, Department of Management and Psychology.

HOUSKISSON, R, E. and C. S. Galbraith (1985), "The Effect of Quantum Versus Incremental M-form Reorganization on Performance: A Time Series Exploration of Intervention Dynamics," *Journal of Management*, 11: 55-70.

ILGEN, D. R., D. A. Major, J. R. Hollenbeck, and D. J. Sego (1991), Decision Making in Teams: Raising and Individual Decision Making Model to Team Level. Tech. Rep. No. 91-2, East Lansing MI: Michigan State University, Department of Management and Psychology.

LAIRD, J. E., P. S. Rosenbloom and A. Newell (1986a), "Chunking in Soar: The anatomy of a general learning mechanism," *Machine Learning*, 1(1): 11-46.

LAIRD, J. E., P. S. Rosenbloom, and A. Newell (1986b), *Universal Subgoaling and Chunking: The automatic generation and learning of goal hierarchies*. Boston MA: Kluwer Academic Publishers.

LAIRD, J. E., A. Newell, and P. S. Rosenbloom (1987), "Soar: An architecture for general intelligence," *Artificial Intelligence*, 33(1): 1-64.

LANT, T. K. and S. J. Mezias (1990), "Managing Discontinuous change: A simulation study of organizational learning and entrepreneurial strategies," *Strategic Management Journal* 11: 147-179.

LAWRENCE, P. R. and J. W. Lorsch (1967), *Organization and Environment: Managing Differentiation and Integration*. Boston MA: Graduate School of Business Administration, Harvard University.

LEBLEBICI H. and G. R. Salancik (1989), "The Rules of Organizing and the Managerial Role," *Organization Studies* 10(3): 301-325.

LEVITT, B. and J. March (1988), "Organizational Learning," *Annual Review of Sociology* 14: 319-40.

LIN, Z. (1994), *Organizational Performance - Theory and Reality*. Pittsburgh, PA: CMU-Heinz, PhD Thesis.

LIN, Z. and K. Carley (1993), "Proactive or Reactive: An Analysis of the Effect of Agent Style on Organizational Decision Making Performance," *International Journal of Intelligent Systems in Accounting, Finance and Management* 2(4): 271-288.

LUPTON, T. (1976), "'Best Fit' in the Design of Organizations," In E.J. Miller (Ed.) *Task and Organization*, New York NY: John Wiley and Sons.

MARCH, J. and H. Simon (1958), *Organizations*. New York NY: Wiley.

MASUCH, M. and P. LaPotin (1989), "Beyond Garbage Cans: An AI Model of Organizational Choice," *Administrative Science Quarterly*, 34: 38-67.

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

PADGETT, J. F. (1980), "Managing Garbage Can Hierarchies," *Administrative Science Quarterly*, 25(4): 583-604.

PAPAGEORGIOU, C. P. (1992), Cognitive Model of Decision Making: Chunking and the Radar Detection Task. Pittsburgh, PA: CMU-CS Bachelors Thesis.

PERROW, C. (1984), *Normal Accidents: Living with High Risk Technologies*. New York NY: Basic Books, Inc., 1984.

PRIETULA, M. J. and K. M. Carley (1994), "Computational Organization Theory: Autonomous Agents and Emergent Behavior," *Journal of Organizational Computing* 41(1): 41-83.

PRIETULA, M. J. and K. M. Carley (forthcoming), "Exploring the Effects of Agent Trust and Benevolence in a Simulated Organizational Task." *AAI Special Issue* "Socially Intelligent Agents".

ROBERTS, K. (1989), "New Challenges to Organizational Research: High Reliability Organizations," *Industrial Crisis Quarterly*, 3(3): 111-125.

SALANCIK, G. R. and H. Leblebici (1988), "Variety and form in organizing transactions: A generative grammar of organization," *Research in the Sociology of Education* 6:1-31.

TANG, Z., K. R. Pattipati, and D. L. Kleinman (1991), "An algorithm for determining the decision thresholds in distributed detection problem,"*IEEE Transactions on Systems, Man, and Cybernetics*, 21: 231-237

WOODWARD, J. (1965), *Industrial Organization: Theory and Practice*. London: Oxford University Press.

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

Return to Contents of this issue