ELSEVIER

Journal of Economic Behavior & Organization
Vol. 896 (1996) 1-17

JOURNAL OF
Economic Behavior
& Organization

A comparison of artificial and human organizations[1]

Kathleen M. Carley

Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract

Formal models can facilitate the development of organizational theory. The complex, adaptive, non-linear nature of human endeavor makes computational models a particularly useful type of formal model for exploring organizational behavior. Researchers, however, rarely contrast such models with empirical data. Further, researchers tend not to contrast such models with other such models. This paper presents results from the artificial organization project and demonstrates how contrasting models with each other and with empirical data can facilitate the development of more veridical organizational models. Specifically, this paper examines whether certain organizational models, differing only with respect to the cognitive limitations of, and adaptability or general intelligence of the agents are better or worse predictors of the behavior of similar organizations composed of humans. It is shown that, not only do different agent models predict different levels of organizational performance, they also predict different relative standings for different organizational structures. Finally, the adequacy of organizations composed of simple computational agents for predicting the behavior of organizations composed of humans appears to increase as the complexity of the organizational structure increases.

JEL classification: D70; D83; C15

Keywords: Simulation; Organizational performance; Adaptive agents; Organizational design; Computational organization theory

1. Introduction

Formal models[2] of organizational structure and behavior demonstrate that: (a) there is no one optimal organizational design; (b) structural constraints, task constraints, and

[1]This work was supported in part by the NSF Grant No. IRI-9111804 and by the Office of Naval Research (ONR), United States Navy Grant No. N00014-90-J-1664.

[2]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, Baligh et al. 1994), and formal logic (Salancik and Leblebici, 1988, 1993,

SSDI 0167-2681(96)00896-7

cognitive constraints on information processing activities can dictate organizational outcomes; and (c) useful policy implications require the models to contain task information. Further, despite Cyert and March's (1992)[1963] early example and emphasis on relating formal models to empirical observations, few researchers have provided even rudimentary empirical tests of any formal model of organizational structure let alone contrasted alternate models. In this paper, these alternate models of organizations are contrasted with each other and comparable data drawn from human experiments. These models all belong in the CORP[3] simulation framework and they vary on a single parameter – how individual agents make decisions. These alternative formal models involve four distinct representations of decision making: two involving experiential learning strategies, one using a standard operating procedure (SOP), and one in which the model's decision making agent is replaced by human operators.

CORP is a framework for simulating the performance of organizations with different designs, subject to different external environments, and composed of model "agents" with different "cognitive" capabilities. CORP type models have been used to provide insight into various issues of organizational theory such as the role of organizational design in affecting performance (Carley, 1991, Carley, 1992), the role of stress in organizations (Carley and Lin, 1995), the impact of individual's personality style on organizational performance (Lin and Carley, 1993), and the ability of computational models to simulate social and organizational agents (Carley and Prietula, 1994, Prietula and Carley, 1994). This sensitivity analysis provides an insight into which features of the model agents, and of human behavior, are driving the observed organizational results. Further, this analysis demonstrates that by taking a structural information processing approach and critically examining the role of cognition, task, and structure it is possible to link micro-level theories of organizational behavior to macro-level theories.

The work presented here is part of the larger artificial organization project, a multi-year, interdisciplinary project. The goal of the artificial organization project is to develop and test a computational theory of organizations in which: (1) organizations are dynamic entities composed of intelligent adaptive agents each of whom has multiple areas of expertise and is working on one or more complex tasks, (2) the flow of information within the organization and the cost of "running" the organization is determined by the organizational design, and (3) organizational performance is determined relative to the task(s). This project is relatively unique as it is one of the few studies where computational models are contrasted with empirical data and with each other.

2. General approach

In order to compare the behavior of human and artificial organizations a series of experiments, both human and simulation, were run in which the task, organizational

[3]The CORP framework is a simulation testbed designed to allow the researcher to compare and contrast the performance of organizations with different designs subject to different types of operating conditions. 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.

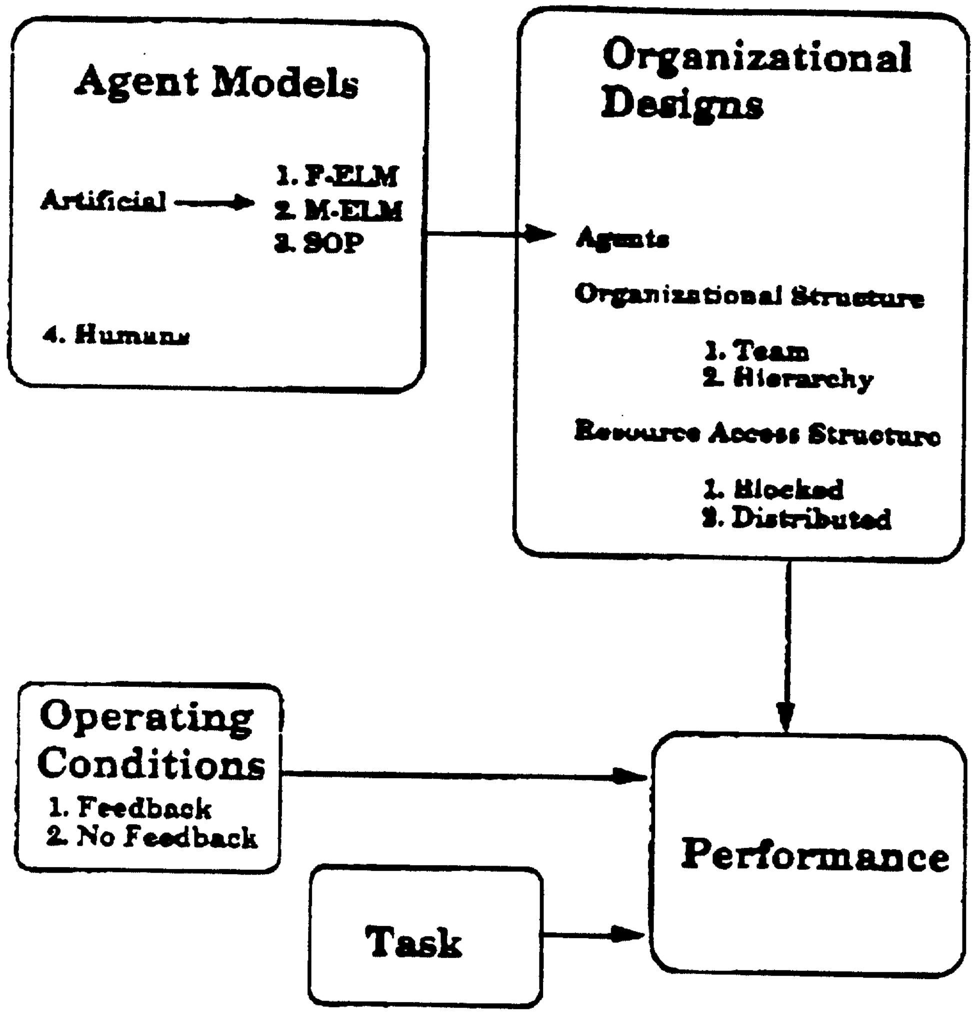

Fig. 1. Description of the experimental process and procedures.

design, and measures of performance were identical provides an orientation to this project. In Fig. 1, for each variable that is varied the categories are enumerated. The basic analysis contrasts three different "agent models" and the human agent experiment in four different organizational designs (characterized by two different organizational structures and two different resource access structures) under two different operating conditions. Across all experiments the focus is on organizational performance in terms of the relative frequency (accuracy) with which the organization makes the right decision.

Performance is potentially affected by the organization's design, the cognitive or information processing capabilities of the agents in the organization, the operating conditions, the task, and the task environment faced by the organization. Typically, in the past, researchers have contrasted the performance and learning behavior of the model with the performance and learning behavior of the "average" human for single person tasks. In this paper, however, the focus is on a multi-person task and the contrast is at the organizational rather than the individual level. Due to this focus multiple organizations (teams and hierarchies, blocked and distributed) are examined under both feedback and non-feedback conditions within each experiment.

The critical difference across the four experiments is the "nature" of the agents. The three model simulation experiments are conducted within the CORP simulation framework using simple agent models. These models are the experiential learning model (ELM) under both complete training (F-ELM) and limited training (M-ELM), and the procedural model (SOP). The fourth experiment is a human experiment. To contrast the different organizations, the organizational designs are held constant and the agents' decisions (artifical or human) are simply substituted for each other.

ELM agents are computational agents who learn from their experience using a simple incremental learning rule. The more experience they have (the more problems they have seen) the greater their training. These agents make their decisions by relying on their experience about what they have seen in the past. Experiential learning is widely recognized as existing in organizations and the particular model used follows from the extensively studied Bush and Mosteller (1955)-type model of individual learning. SOP agents are computational agents who do not learn and make decisions by following standard operating procedures. Many organizations operate exclusively under a procedural system. Additional details about these agents will be provided later. Experiential and procedural (SOP) agents are contrasted as one of the major debates in organizational design is over whether organizations should employ professionals who follow their own experience or should rely on procedurally trained novices (Scott, 1987).

2.1. *Task: The radar task*

The task used is a ternary classification choice task often referred to as a radar task. In the radar 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. 1992).

In the radar task, there is a single aircraft in the airspace at a time. This aircraft is uniquely characterized by nine pieces of information (each representing a different characteristic) such that each piece of information is either a 1, 2, or a 3 (representing a low, medium, or high value on that characteristic). On the basis of these nine characteristics, the organization must determine for each observed aircraft whether that craft is friendly (=1), neutral (=2), or hostile (=3). At each time period the organization is faced with a single aircraft. This craft is randomly chosen from the set of 19683 possible aircraft with replacement. The radar task represents a series of distributed decision making problems such that each problem is similar to (but not identical to) previous problems.[4] The organization is assumed to have personnel who can access raw

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

data (the analysts) and who can evaluate this information and coordinate with other organizational members to make this determination before the aircraft changes position in the airspace. 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, Carley, 1992).

For each time period, i.e. for each aircraft, there is a "true" state and an organizational guess or decision about the nature of that aircraft's state. The true state of the aircraft is algorithmically defined given the nine characteristics of the observed aircraft. The true state of the aircraft is determined by simply adding the values of the nine characteristics. If this sum is less than 17 the true state is friendly, if this sum is greater than 19 the true state is hostile, otherwise the true state is neutral. 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.

Each organization (regardless of its design or the type of agents) is given a set of 60 problems. These 60 problems are divided into two sets of 30. Each set of 30 was selected randomly from the set of 19683 possible aircraft such that one third of those selected problems were from each of the friendly, neutral, and hostile categories. All organizations face the same set of 60 problems. For the first 30 problems the organizations received feedback as to what was the correct answer for that problem. For the second set of 30 problems no feedback was provided.

2.2. *Organizational design*

There are two parts to the organizational design – the resource access structure and the organizational structure (see Fig. 2). The resource access structure defines who has access to which raw information about the observed aircraft. The organizational structure defines who reports to whom within the organization and hence how decisions are reached.

The personnel in the organization who can access raw information on the aircraft are referred to as analysts. Analysts can access information on only three of the nine aircraft characteristics. When analysts access this information they 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." Analysts observe this information and make a recommendation as to whether they think the aircraft is friendly, neutral, or hostile. All managers only examine the recommendations of their subordinates. Managers make a recommendation of the state of the aircraft on the basis of the recommendations passed to them by their subordinates.

Which analyst has access to which information is defined by the resource access structure. In this paper, two different resource access structures are contrasted – the blocked and the distributed. In the blocked structure three analysts see exactly the same information. 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.

Which agent passes information to which other agent and so how the organizational decision is made is defined by the organizational structure. In this paper, two different

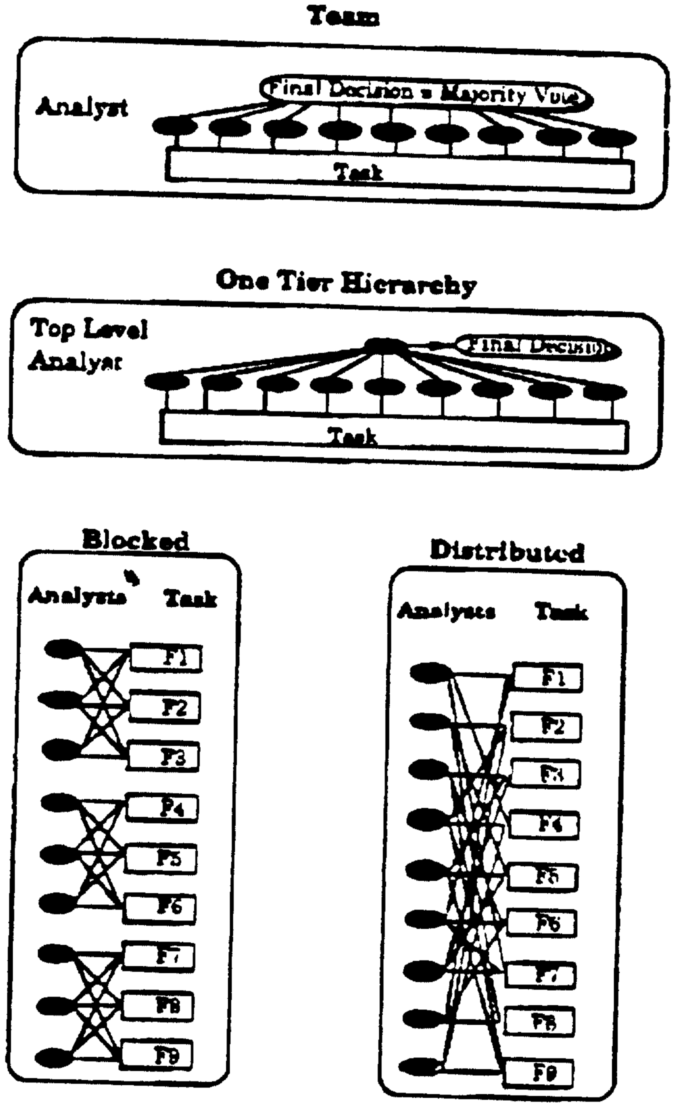

Fig. 2. Organizational structures and resource access structures.

organizational structures are contrasted – team and hierarchy. Teams are characterized by not having a manager and by the organizational decision simply being the majority vote given the opinions of the analysts. Hierarchies are characterized quite simply as a team of analysts reporting to a single manger. In this case, the manager makes the organizational decision.

Overall, four organizational designs are considered – team blocked, team distributed, hierarchy blocked, hierarchy distributed. These four designs, though simple, allow us to examine the impact of altering the agent model at both the staff and managerial level and when consensus should and should not exist. These variations will provide insight into the

Table 1
Organizational decision and true state

True state x organizational decision	Friendly	Neutral	Hostile
Friendly	correct	incorrect (minor)	incorrect (severe)
Neutral	incorrect (minor)	correct	incorrect (minor)
Hostile	incorrect (severe)	incorrect (minor)	correct

conditions under which the nature of the agent actually affects organizational performance.

2.3. Operating conditions

The operating conditions determine the extent to which the organization is faced with externalities such as missing information, erroneous information, turnover, missing personnel, and so forth. In this paper, all organizations are operating under ideal operating conditions in terms of incoming information. That is, nothing is going wrong, there is no missing or erroneous information, no missing personnel, no turnover. The main operating difference is the presence of feedback for the first 30 problems and the absence of feedback for the second 30.

2.4. Performance

As previously noted, performance is defined in terms of accuracy. Two different indicators will be used – accuracy and severity of error. Accuracy is measured as the percentage of the 30 problems for which the organizational decision matches the true decision. Severity of error is measured as the percentage of errors that are severe (see Table 1). Recall that one third of the problems seen by the organization are actually friendly, one third are neutral, and one third are hostile.

3. Agent models

In all organizations, regardless of the design, there are a set of agents. There are nine agents in the team and ten in the hierarchy. All agents in the same organization, regardless of whether they are analysts or managers, are of the same type. Four types of agents are considered. These agents vary in how they make their decision or recommendation and in what information they remember from one problem to the next. Regardless of what type of agent they are, all analysts see exactly three pieces of information and all managers see exactly nine pieces of information. Further analysts only see information about the aircraft's characteristics and managers only see the analysts recommendations.

3.1. Experiential learning model

Agents based on the experiential learning model (ELM) have perfect memories of the frequencies with which each pattern that they have encountered in the past is associated with each of the three outcomes (friendly, neutral or hostile). Analysts, who see three pieces of information, have information about 27 patterns. Mangers, who see nine pieces of information, have information about 19683 patterns. Agents keep track of the percentage of times when they observed that pattern when a particular outcome (friendly, neutral, or hostile) occurred. Thus, an agent may know that 50% of the time when it saw pattern xxy the aircraft was truly friendly, 30% of the time it was neutral, and 20% of the time it was hostile. As the agent gains experience these percentages change. ELM agents perfectly follow their experience. Thus, when faced with new information they will first determine which pattern they are observing, and then they will choose as their decision that outcome which occurred with the highest probability in the past given that pattern. Note, this makes these agents insensitive to sample size and effectively overconfident in their decisions.

For example, imagine that the agent is an analyst and sees information about three of the aircraft's features. The agent may have learned that the speed is low (=1), the direction is medium (=2), and the radar type is hostile (=3). This is the pattern 123. The agent will locate the percentages associated with this pattern. Imagine that, in the past, 20% of the time when this pattern was seen the aircraft was friendly, 50% neutral, and 30% hostile. Thus the agent will, after having seen this pattern, decide that the aircraft is neutral. Further, the agent knows how often it has seen a pattern. Imagine that the agent has seen this pattern 10 times. Now, it has just seen it again. After making its decision the agent is told the true state of the aircraft. Imagine that the aircraft is truly hostile. Now the agent updates its memory from 10 to 11 occurrences and a frequency of 2 friendly, 5 neutral, and 4 hostile. The new probabilities are now friendly at 18.18%, neutral at 45.45% and hostile at 36.36%. Given this learning procedure a fully-trained analyst (one who has seen 19683 tasks) will act as a majority classifier. That is, the probabilities will be such that the agent's decision will be friendly if the majority of the inputs are low, neutral if the majority are moderate, and hostile if the majority are high.

As the ELM agent examines more problems its knowledge base grows. Thus, ELM agents will behave differently depending on the level of their training. Herein two levels of training are considered. In F-ELM each agent has seen all 19683 problems and is operating on the basis of complete training. F-ELM analysts act as majority classifiers. In M-ELM each agent has seen only 10 problems, the same 10 problems that humans face in their training period. Thus, while they are becoming majority classifiers they are not quite there.

3.2. Procedural model

Agents based on the procedural model (SOP) have perfect memories of the rules that they are to apply when faced with a new problem. They do not track their experience. These agents learn a series of standard operating procedures and then always follow them. These agents do not adjust their behavior given feedback. The particular standard operating procedures followed are based on a majority rule classification. The agent makes its decision by contrasting the sum of the values associated with those pieces of information available to it with some cutoff. For analysts the three rules followed are: (1) if the sum is ≤5 then report friendly, (2) if the sum is ≥7 then report hostile, and (3) if the sum is >5 and <7 report neutral. Training does not affect the SOP model. These agents act the same regardless of the number of problems they have faced. The SOP and F-ELM analyst, in the limit, exhibit identical behavior.

3.3. *Human experiments*

The fourth type of agent model is not really a model but humans themselves. Data on human behavior were collected through a series of experiments in which the subjects did the radar task (Carley and Prietula, 1992, Carley and Prietula, 1993). The experiment has a staged design. First, data were collected from all subjects at the analyst position. Second, the results of these analysts' recommendations were combined in different ways and given to the subjects in the managerial positions. The experiment is physically run by having subjects, individually, log into a "radar" program that trains them, provides them with information on the nature of the aircraft, records their responses, and provides them with feedback on the accuracy of their recommendations. Results are gathered in such a fashion that only one subject need be present at a time.

First subjects were trained on 10 problems. Then they faced 30 problems with feedback and then 30 without feedback.[5] Each subject acting as an analyst saw information on only three characteristics for each of 60 problems. The subjects/analysts 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." The subjects/analysts were then asked if they thought the aircraft was friendly, neutral, or hostile. After the subjects/analysts 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 the 9 analysts in his or her group. The subjects/managers were given information on one aircraft at a time and the information was in the form "analyst one thinks the aircraft is friendly, analyst two thinks that the aircraft is neutral, analyst three thinks that the aircraft is hostile." The subjects/managers were then asked if they thought the aircraft was friendly, neutral, or hostile. After the subjects/managers provided their recommendation, they were asked for their confidence in their decision. In all cases, the subject's decision and the time to make that decision were collected automatically. This procedure basically substitutes human decisions for the decisions made by the artificial agents in the CORP framework.

[5]These 60 problems were the first and second 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 are considered.

4. Results

Let us compare the performance of these various organizations as the model of the agent is altered. This comparison addresses two questions. First, does the model of the agent within the organization matter to the organization's performance? And second, which formal model more closely approximates the behavior of organizations of humans.

All three artificial agents are limited in their intelligence to procedures needed for this specific radar task. All agents (human and artificial) are boundedly rational. ELM agents have the fewest bounds on their intelligence. While the SOP agents are the most bounded of the artificial agents, humans are expected be have more bounds on them than any artificial agents as the artificial agents have perfect, albeit limited memories. M-ELM agents have had less training than F-ELM agents and are thus both more flexible than F-ELM agents (their decision structures are not fixed yet) and more limited (less knowledge). Humans are expected to have imperfect and limited memories. F-ELM and M-ELM agents have identical abilities to adapt, although the M-ELM agents have time to learn and so the adaptation will be more noticeable. SOP agents are not adaptive. Humans are expected to be the most generally intelligent, flexible, and adaptive of all agents. General intelligence is seen as being affected by the multiplicity of tasks that the agent can do, the multiplicity of mechanisms for adaptation held by the agent, and the flexibility of the agent. If we rank the agents from low to high in terms of adaptability or general intelligence the order is SOP, F-ELM, M-ELM, Human. Further, humans are much more adaptable and generally intelligent than any of the artificial agents. Given these differences we expect that if the agents capabilities matter to the organization then organizational performance will vary as the type of agent within the organization varies.

Results in terms of accuracy across all 60 problems are shown in Table 2. All pairs of means with a difference greater than 1.7 are statistically significantly different at the 0.005 level using a standard t-test. First, we see that adding cognitive constraints, increasing the general intelligence, or increasing the adaptiveness of the agents tends to decrease organizational performance; i.e. F-ELM and SOP outperform M-ELM which in turn outperforms humans. Second, we see that the artificial SOP agents are very suited to this task (i.e. performance is high under all organizational designs). Third, we see that the impact of organizational structure tends to vary with agent model. There is direct evidence of interaction effects between agent model and organizational design.. For example, organizations composed of humans exhibit their highest performance in a team with a distributed structure and lowest performance in a one-tier hierarchy with a blocked

Table 2
Organization's accuracy by agent model and organizational design

Agent	Training	Team		Hierarchy	
		Blocked	Distributed	Blocked	Distributed
ELM	full	88.3%	85.0%	45.0%	50.0%
ELM	min	78.3%	71.7%	40.0%	36.7%
SOP	full/min	81.7%	85.0%	81.7%	85.0%
Human Experiment		50.0%	56.7%	46.7%	55.0%

Table 3
Severity of error by agent model and organizational design

Agent	Training	Team		Tier Hierarchy	
		Blocked	Distributed	Blocked	Distributed
ELM	full	71.4%	0.0%	6.1%	10.0%
ELM	min	7.7%	23.5%	19.4%	28.9%
SOP	full/min	0.0%	0.0%	0.0%	0.0%
Human Experiment		23.3%	23.1%	15.6%	14.8%

structure. In contrast, organizations composed of M-ELM agents exhibit their performance when the organizational structure is a blocked team, and the lowest when the structure is a distributed hierarchy. Finally, we see that while no model is adequate for explaining the teams of humans, the F-ELM model is comparable to humans when the agents are in a hierarchy.

Now consider the distribution of errors. In Table 3 we see that only organizations composed of SOP agents make no severe errors. In addition, organizations of humans make the most severe errors. As with accuracy, we see that the impact of organizational structure on organizational performance varies with the agent model. Further, when the agents are arranged in a hierarchy, the F-ELM model most closely resembles the human model terms of performance.

As the organization faces these 60 problems several things are happening. First, the agents in the organization are presumably learning from feedback (all except SOP agents who cannot learn). Second, the organizations receive feedback for the first 30 problems and they do not receive feedback for the second 30 problems. These results are shown in Fig. 3.

In examining these figures we see that, in most cases, the computational models exhibit performance improvements even in the absence of feedback. The computational agents are still not learning but are exhibiting better performance due to applying information previously learned. In contrast, we see that organizations of humans do worse when there is no feedback. In certain cases, such as the blocked hierarchy, an artificial organization is a strikingly good model of human behavior, but not in general. As an additional point, F-ELM and M-ELM typically exhibit more severe errors in the feedback condition than in the no-feedback condition; while the opposite is the case for humans. The SOP model has no severe errors in either condition.

Even the SOP model exhibits an increase in performance when there is no feedback. The SOP agents cannot learn. Thus, feedback is actually irrelevant to their behavior. The observed increase in performance is due to the particular problems in the second set of 30 being less ambiguous than the first 30, given the particular sops used. In contrast, any improvement in performance under ELM is due to applying newly learned rules. For organizations of humans, you see a dip in performance when they stop receiving feedback and then later the organization's performance improves as they become accustomed to the lack of feedback. This dip only occurs for the F-ELM model (due to problem ambiguity). In fact, in terms of over-time behavior, F-ELM exhibits the pattern most similar to

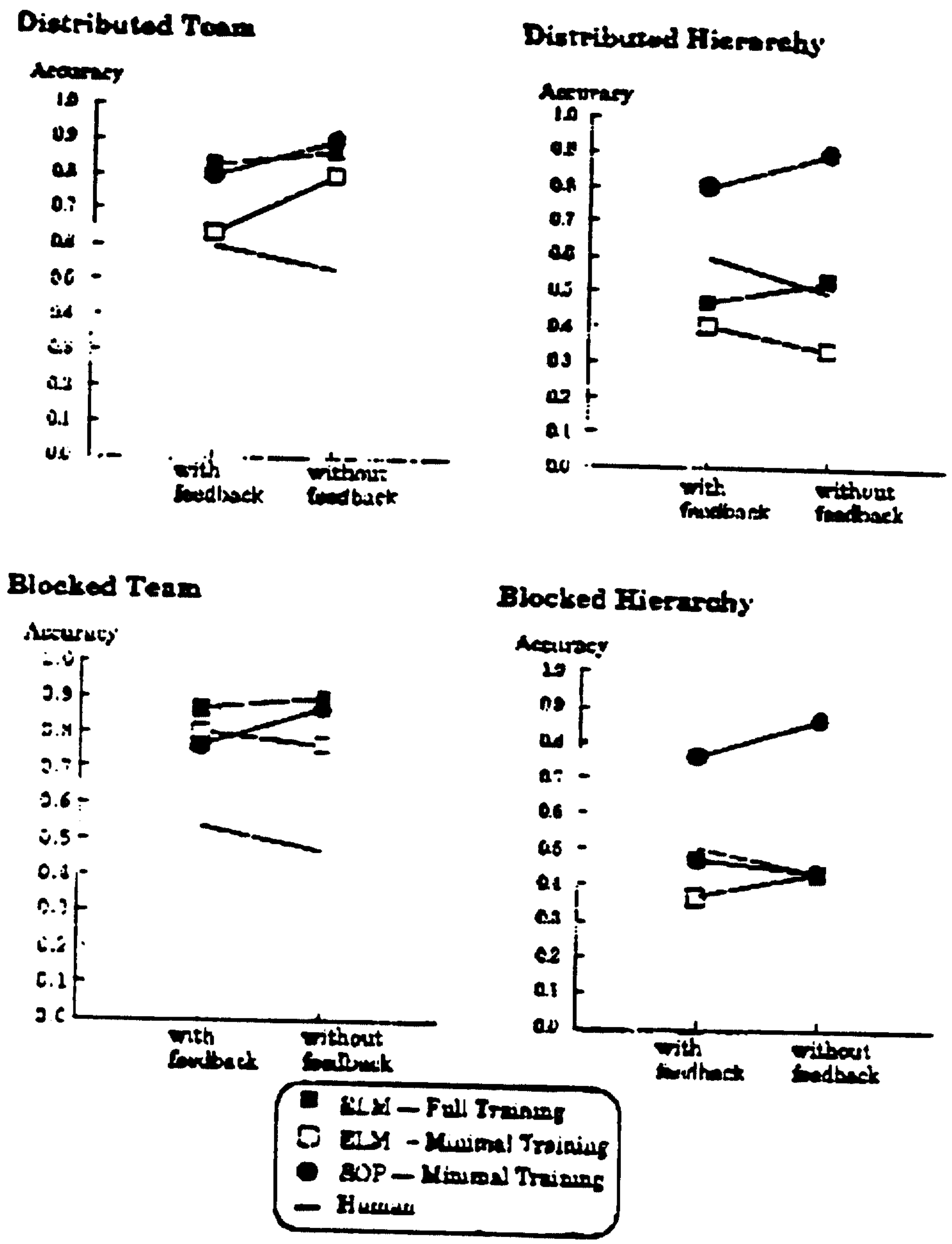

Fig. 3. Impact of feedback on each organizational design.

humans - see Fig. 4. In Fig. 4, the average accuracy across all organizational designs for each agent averaged across five problems is shown. Note, only F-ELM and the human organizations show the initial rise, fall, then three rises, then the dip after problem 30.

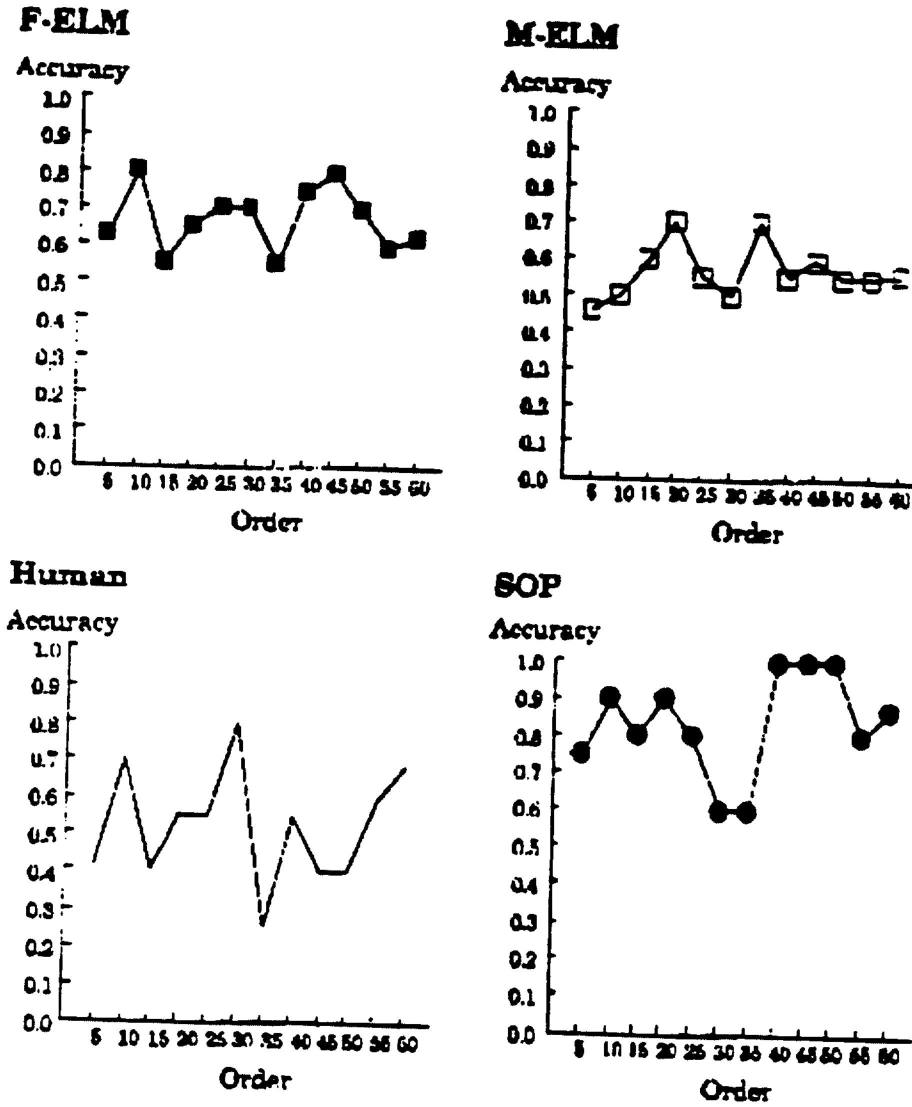

Fig. 4. Over time behavior for organizations.

Finally, the relative performance averaged across all organizational designs for all problems with the same true outcome is examined (Fig. 5). Organizations of artificial agents are more likely to make the correct decision when the aircraft is friendly or hostile. For these models, the characteristics of aircrafts that are neutral are, from an organizational standpoint, ambiguous and so the organizations are more likely to misclassify them. Organizations of F-ELM agents are the best at handling hostile craft. Like humans, the F-ELM agents have a tendency to over predict hostility. Organizations of human agents tend to be as poor at classifying friendly craft as neutral craft. The fact

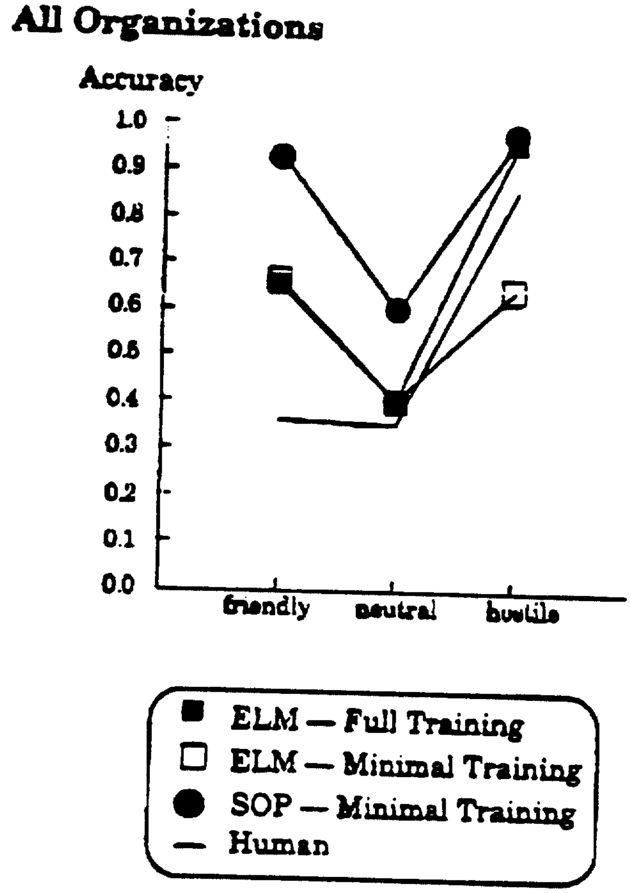

Fig. 5. Accuracy by type of decision.

that organizations of humans do worse than organizations of F-ELM agents is due in part to humans misclassifying friendly aircraft.

5. Discussion

This paper addressed two basic questions. Does the model of the agent within the organization matter to the organization's performance? And, which formal model more closely approximates the behavior of organizations of humans. The answer to the first is yes, and to the second – it depends.

The strong interaction between agent model and organizational design suggests that organizational theorists need to increase their attention to the way in which humans in organizations are characterized. These results have multiple interpretations.[6] First, they suggest that training is a stronger determinant of organizational performance than organizational design. Second, they suggest that simple models of agent behavior may not be adequate for accurately predicting organizational performance at the small group or micro level; but may be accurate at a more macro level. In particular, simple models of agent behavior, such as ELM, may be adequate when we are dealing with more complex organizational structures and when the information is filtered through a chain of command. Third, they suggest that organizational researchers should be careful when making attributions regarding organizational performance. Consider the interaction between hierarchy and the manager model. When the manager is an ELM or human agent and the organization is a hierarchy performance is degraded. This might seem to suggest that a different model is needed for agents who view decisions and agents who view raw data. In actuality, however, this is not the case. The lower performance is due to the manager having, in effect, a harder job. Managers have to learn associations based on nine pieces of information rather than the three that analysts examine. For agents following SOPs the amount of information examined is irrelevant. For ELM agents the more pieces of information the agents must contend with the slower they learn. This results in lower (or lagged) performance for M-ELM agents compared to F-ELM agents. For F-ELM managers more information corresponds to an increased number of opportunities for raw information to be incorrectly summarized by analysts, and hence lower performance. Humans may be suffering from both these effects. Future work might explore how the interaction between formal position in the organization and agent model interact further to affect organizational performance.

These results should be viewed with caution. First, the performance of the computational organizations is contrasted with the performance of a single human organization. The extent to which these human organizations are representative is an open question. A second analysis contrasting four human organizations in the distributed hierarchy suggests that these results are representative as all reported results hold.

An examination of these results suggests ways of modifying agent models so that organizations of such agents more closely resemble the behavior of organizations of humans. Overall, the F-ELM model exhibits behavior more typical of humans than do the other models. Clearly F-ELM and human organizations both exhibit a bias toward correctly interpreting some extreme situations (better at hostile aircrafts). For F-ELM, there is a bias in being better at both extremes. Not so for humans. Clearly organizations of F-ELM agents can outperform organizations of humans. What is interesting is that the difference in behavior is apparently not caused by the level of training. Even M-ELM agents who have identical training to humans outperform humans. Moreover, the M-ELM organizations exhibit behavior more qualitatively different from human organizations than do F-ELM organizations. Clearly organizations of F-ELM agents, unlike humans, are not affected by the lack of feedback. What is interesting is that the absence of feedback may not, in and of itself, be detrimental to performance. One might expect that in the absence of feedback the organization's behavior would remain constant as the individual agents cannot learn. This is what happens in the simulation models. For human organizations, however, performance is initially degraded and then learning continues and performance improves. These differences suggest that a more adequate model of the organizational agent should consider decision biases and agent limitations other than lack of training (e.g. memory loss), and variations in feedback.

These results demonstrate that the development of more veridical organizational models is greatly facilitated by doing both cross model comparisons and comparisons with empirical data. This research increases our understanding of the role of agent capabilities in affecting organizational performance. It moves us a step closer to developing social artificial agents and increases our ability to observe emergent social phenomena (Carley and Newell, 1994). It demonstrates 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 such as the radar task.

References

Baligh, H.H., R.M. Burton and B. Obel, 1990, Devising expert systems in organization theory: The organizational consultant, In: M. Masuch, Ed., Organization, Management, and Expert Systems, (Walter De Gruyter, Berlin) pp.35–57.

Baligh, H.H., R.M. Burton and B. Obel, 1994, Validating the organizational consultant on the fly, In: K.M. Carley and M.J. Prietula, Eds., Computational Organization Theory, (Lawrence Erlbaum, Hillsdale, NJ).

Bond, A.H. and L. Gasser, 1988, Readings in distributed intelligence, (Morgan Kaufmann, San Mateo, CA).

Bush, R.R. and F. Mosteller, 1955, Stochastic models for learning, (Wiley, New York)..

Carley K.M., 1991, Designing organizational structures to cope with communication breakdowns: A simulation model, 1991, Industrial Crisis Quarterly 5 19–57.

Carley, K.M., 1992, Organizational learning and personnel turnover, 1992, Organization Science 3(1) 2–46.

Carley, K. and Z. Lin, 1995, Organizational designs suited to high performance under stress, 1995, IEEE – Systems Man and Cybernetics 25(1) 00.

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

Carley, K.M. and M.J. Prietula, 1992, Toward a cognitively motivated theory of organizations, In: The Proceedings of the 1992 Coordination Theory and Collaboration Technology Workshop, Washington, DC.

Carley, K.M. and M.J. Prietula, 1993, Plural-Soar: Towards the development of a cognitively motivated theory of organizations, In: The Proceedings of the 1993 Coordination Theory and Collaboration Technology Workshop Washington, DC.

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

Cohen, M.D., J.G. March and J.P. Olsen, 1972, A garbage can model of organizational choice, 1972, Administrative Sciences Quarterly 17(1) 1–25.

Cyert, R. and J.G. March, 1992[1963], A behavioral theory of the firm, 2nd edn., (Blackwell, Cambridge, MA).

Harrison, J.R. and G.R. Carrol, 1991, Keeping the faith: A model of cultural transmission in formal organizations, 1991, Administrative Science Quarterly 36 552–582.

Hollenbeck, J.R., D.J. Sego, D.R. Ilgen and D.A. Major, 1991 Team interactive decision making exercise for teams incorporating distributed expertise (TIDE^2): A program and paradigm for team research, Tech. Rep. No. 91-1, East Lansing: Michigan State University, Department of Management and Psychology.

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: Michigan State University, Department of Management and Psychology.

Lant, T.K. and S.J. Mezias, 1990, Managing discontinuous change: A simulation study of organizational learning and entrepreneurial strategies, 1990, Strategic Management Journal 11 147–179.

Lin, Z. and K. Carley, 1993, Proactive or reactive: An analysis of the effect of agent style on organizational decision making performance, 1993, International Journal of Intelligent Systems in Accounting, Finance, and Management 2(4) 271–288.

Masuch, M. and Z. Huang, 1994, A logical deconstruction of organizational action: Formalizing J.D. Thompson's organizations in action in a multi-agent action logic, CCSOM Working Paper 94-120.
Masuch, M. and P. LaPotin, 1989, Beyond garbage cans: An AI model of organizational choice, 1989, Administrative Science Quarterly 34 38–67.
Padgett, J.F., , 1980, Managing garbage can hierarchies, 1980, Administrative Science Quarterly 25(4) 583–604.
Prietula, M.J. and K.M. Carley, 1994, Computational organization theory: Autonomous agents and emergent behavior, 1994, Journal of Organizational Computing 4(1) 41–83.
Salancik, G.R. and H. Leblebici, 1988, Variety and form in organizing transactions: A generative grammar of organization, 1988, Research in the Sociology of Education 6 1–31.
Scott, W.R., 1987, Organizations: Rational, Natural, and Open Systems, (Prentice Hall, Englewood Cliffs, NJ).
Tang, Z., K.R. Pattipati and D.L. Kleinman, 1992, A distributed M-ary hypothesis testing problem with correlated observations, 1992, IEEE Trans on Automatic Control XX 632–636.