

Coordination Theory and Collaboration Technology Workshop

National Science Foundation

July 8-10, 1993
Washington, DC

Plural-Soar: Towards the Development of a Cognitively Motivated Theory of Organizations

NSF Grant # IRI-9111804

Kathleen M. Carley
Carnegie Mellon University

Introduction

This research examines the role of cognition in affecting organizational behavior. We address the question, "Given that organizations are composed of intelligent agents, how intelligent do those agents need to be to achieve reasonable levels of organizational performance?" We also address a related question, "how intelligent do our models of humans need to be to generate organizational behavior comparable to that observed in human organizations?" We address these questions by comparing the performance of organizations that are composed of either humans or artificial agents. This research takes the ACTS approach (Carley & Prietula, forthcoming). In this meso-level approach, organizations are modeled as collections of agents who are cognitively restricted, engaged in tasks and socially situated within a particular organizational design and culture.

We expanded Plural-Soar so that the agents were more social in character (Carley, Park and Prietula, 1993). This expansion included adding the ability to lie and detect falsehoods in others. This made it possible for the agents to have meaningful social models of other agents. This task (the warehouse task), however, was too constraining. Therefore, due to task constraints, we have moved from the warehouse task used in Plural-Soar to a stylized radar task. Using this radar-task we then began to examine our fundamental question.

Agent Models

We consider two types of artificial agents varying in their level of general intelligence—ELM agents (with only task specific intelligence), and Soar agents (with general intelligence). Agent intelligence is defined simply as the extent to which the cognitive architecture can be adapted to different tasks.

ELM, the experiential learning model, are agents that learn incrementally on the basis of their experience (Carley, 1990, 1991, 1992; Carley & Lin, 1993). They keep a perfect record of all problem types that they see and the frequency of particular outcomes given this type of problem. When faced with a new problem, they simply determine its type and then choose as their decision that outcome that was historically the most frequently correct.

Radar-Soar agents can also learn incrementally on the basis of their experience (Ye and Carley, 1993). Each Radar-Soar agent is built in Soar (Laird, Newell, and Rosenbloom, 1987; Laird, Rosenbloom, and Newell, 1986a, 1986b) that is arguably a model of general cognition (Newell, 1990). Unlike ELM agents 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.

We analyzed the behavior of a set of organizations such that each organization was composed of a set of organizations all of whom are of the same type (ELM, Soar, or Human). Each organization is faced with a stylized radar task in which the true decision for each problem is related to the information in an unbiased and decomposable fashion (Carley & Lin, 1993). In the radar task, there is a single aircraft in the airspace at a time and it is uniquely characterized by nine pieces of information (each representing a different aircraft characteristic) such that each piece of information is either a 1, 2, or a 3 (low, medium, or high on that characteristic). The organization must determine for each aircraft it sees whether it is friendly, neutral, or hostile. The organization sees a sequence of 60 problems. For the first 30 feedback is provided. For the second 30 no feedback is provided. These organizations vary in their structure and task decomposition scheme. We examined two structures—team with voting and team with manager. In both cases the team (sans manager) is composed of 9 agents. We examined two task decomposition schemes—blocked and distributed. In both cases each agent (sans manager) sees only a portion of the problem—three pieces of information.

Results

The main result is that in moving from ELM, to Radar-Soar, to Humans, the general intelligence level of the agents is increasing and the organizational performance is decreasing (see Figure 1). Compared to the ELM model, Soar moves us a step closer to Human behavior. ELM agents, however, are better suited to the radar task. Soar agents are better suited to modeling Humans. This pattern is consistent across different organizational designs. Regardless of their structure, organizations of more generally intelligent agents tend to perform worse.

An implication of these findings is that it is necessary to use better models of cognitive agents in the study of human organizations in order to achieve organizational performance results akin to those which an organization of humans would achieve. Another implication is, it is possible that in certain organizations, such as those faced with a choice task in which feedback is possible, performance can be improved by substituting artificially intelligent agents for humans in the mundane process of locating the decision.




Figure 1. Tradeoff in General Intelligence and Performance on Radar Task.

This research increases our understanding of the role of agent capabilities in affecting organizational performance. It also moves us a step closer to developing social artificial agents. Plural-Soar agents (Carley, Kjaer-Hansen, Prietula, and Newell, 1992; Prietula and Carley, forthcoming) were the first instantiation of soar as a social agent. As Carley et al. noted, Plural-Soar agents exhibited at a rudimentary level many of the capabilities need of 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 & Newell, 1993). The Radar-Soar agents are also more socially capable than ELM agents. In this research, we did not address the social capability of Radar-Soar' agents. There was little reason or opportunity for the agents to communicate with each other. Thus, it is not clear whether the increase of social capability in moving from ELM to Soar actually had an effect on task performance. Future research should examine this question.

We have demonstrated that by increasing the general intelligence of agents in the organization the organization's performance can decrease. We have also demonstrated that the model of the agent makes a difference in assessing organizational performance. This suggests that organizational researchers should be careful when making attributions regarding organizational performance as in whether performance differences are attributable to some fundamental feature of the agents or to some other aspect of the organization.

References

- Carley K.M. (1990). "Trading Information Redundancy for Task Simplicity." Proceedings of the 23rd Annual Hawaii International Conference on System Sciences.
- 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., J.Kjaer-Hansen, M. Prietula, and A. Newell (1992). "Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior." P. 87-118 in Masuch M. and M. Warglien (Eds.) *Distributed Intelligence: Applications in Human Organizations* Amsterdam, The Netherlands: Elsevier Science Publications.
- Carley K.M. and Z. Lin (1993). "Maydays and Murphies: A Theoretical Study of Organizational Performance Under Stress." SDS Working Paper, CMU.
- Carley K.M. and A. Newell (1993). "On the Nature of the Social Agent." SDS Working Paper, CMU.
- Carley K., D. Park and M. Prietula (1993). "Agent Honesty, Cooperation and Benevolence in an Artificial Organization" In Prietula M., K.Carley, L. Gasser and D. King (Eds.) Workshop Notes for the AI and Theories of Groups and Organizations: Conceptual and Empirical Research, Eleventh National Conference on Artificial Intelligence, Washington D.C.
- Carley K.M. and M.J. Prietula (forthcoming). "ACTS Theory: Extending the model of bounded rationality." In *Computational Organization Theory* edited by Carley, K.M. and M.J. Prietula. Lawrence Erlbaum Associates.
- Laird J. E., Newell, A., and P. S. Rosenbloom (1987). "Soar: An architecture for general intelligence." *Artificial Intelligence*, 33:(1), 1-64.
- Laird J. E., Rosenbloom, P. S., and A. Newell (1986a). "Chunking in Soar: The anatomy of a general learning mechanism." *Machine Learning*, 1:(1), 11-46.
- Laird J. E., Rosenbloom, P. S., and A. Newell (1986b). *Universal Subgoaling and Chunking: The automatic generation and learning of goal hierarchies*. Boston, Massachusetts: Kluwer Academic Publishers.
- Prietula M.J. and K.M. Carley (forthcoming). "Computational organization theory: Autonomous agents and emergent behavior" *Organizational Computing*.
- Newell A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.
- Ye M. and K.M. Carley (1993). "Radar-Soar: An Artificial Organization Composed of Generally Intelligent Agents" SDS Working Paper, CMU.

Project Publications

- Carley K.M. and M.J. Prietula (forthcoming). "ACTS Theory: Extending the model of bounded rationality." In *Computational Organization Theory* edited by Carley, K.M. and M.J. Prietula. Lawrence Erlbaum Associates.
- Prietula M.J. and K.M. Carley (forthcoming). "Computational organization theory: Autonomous agents and emergent behavior" *Journal of Organizational Computing*.
- Carley K., D. Park and M. Prietula (1993). "Agent Honesty, Cooperation and Benevolence in an Artificial Organization" In Prietula M., K.Carley, L. Gasser and D. King (Eds.) Workshop Notes for the AI and Theories of Groups and Organizations: Conceptual and Empirical Research, Eleventh National Conference on Artificial Intelligence, Washington D.C.
- Carley K. and M. Prietula (1992). "Toward a Cognitively Motivated Theory of Organizations." In the Proceedings of the 1992 Coordination Theory and Collaboration Technology Workshop, Washington D.C.

Project Researchers

Kathleen M. Carley, Associate Professor of Sociology and Information Systems, Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, 1-412-268-3225, 1-412-268-6938, kathleen.carley@centro.soar.cs.cmu.edu

Michael J. Prietula, Associate Professor of Information Systems, Graduate School of Industrial Administration, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, 1-412-268-5067, 1-412-268-7357, mp2j+@andrew.cmu.edu