The potential linkages between artificial intelligence and sociology are growing. This growth is due to importation of artificial intelligence techniques into methodological tools for data analysis, a growing interest among researchers in artificial intelligence in the socially situated agent, and a growing interest among sociologists in using artificial intelligence techniques for theorizing about social phenomena. Increasingly, researchers are addressing concerns of traditional importance within sociology, such as the bases for cooperation, the role of structure in affecting individual agency, and interaction using computational models of intelligent adaptive agents. This article provides an overview of the role that artificial intelligence currently plays within sociology.

Artificial Intelligence Within Sociology

KATHLEEN M. CARLEY
Carnegie Mellon University

Traditional artificial intelligence (AI) focuses on the agent in isolation (e.g., logic theorem provers and chess programs). Much of this work led into current cognitive psychology. Even these single-agent studies often came up against the bulwark of society. For example, Drescher (1986) in his work on genetic AI found it necessary, even as Piaget did, to rely on social knowledge to understand individual action. Researchers in expert systems have found that important gains are possible by including general social knowledge in expert systems and when using findings from the sociology of knowledge (Collins 1990). R. Collins (1992) has argued that AI cannot really be achieved without help from sociologists—a point echoed by Carley and Newell (1994) in their discussion of what it takes to create a model social agent. Today, a growing amount of research in the AI community is focusing on the situated agent—the agent in the group, organization, or society. Currently, computer scientists and engineers working in the areas of distributed artificial intelligence (DAI) and complex adaptive agents are using computational techniques to address questions of traditional interest to social scientists, such as the emergence of cooperation and the role of structure in influencing action.

SOCIOLOGICAL METHODS & RESEARCH, Vol. 25 No. 1, August 1996 3-30
© 1996 Sage Publications, Inc.

The potential for fruitful linkage between sociology and AI is growing. Increasingly, sociologists are asking questions that are appropriately addressed using techniques from AI and computer science. Increasingly, researchers in AI are working on topics of specific interest to sociologists. Areas in which sociologists should expect to see the infiltration of AI and computer science include textual analysis, network analysis, theory examination, social evolution, and organizational theory. Of particular import to social theorists is the recent work using complex adaptive agent models.

Importantly, the contributions of AI to sociology are at both the methodological (Blank, McCartney, and Brent 1989) and the theoretical level (Castelfranchi and Werner 1992; Cohen, 1986; Gilbert and Heath 1985; Wolfe 1991). For a general review, see Bainbridge et al. (1994) and Carley (1994b). In terms of methodology, AI techniques and findings are being used in the analysis of texts, network analysis, and theory evaluation. In terms of theory, computational models in which agents are intelligent are increasingly being used to address issues of social and organizational change, development of social knowledge, and the linkage between individual agency and action in a multiagent environment. Within political science, there is also a growing link with the work in AI (Hudson 1991). Researchers interested in issues of theory evaluation, cooperation, and textual analysis will find additional information in the political science literature. This article discusses several areas in which AI and computational techniques are infiltrating or have the potential to change the way sociologists do research. These areas include textual analysis, network analysis, expert systems, emotions, and multiagent models.

TEXTUAL ANALYSIS

Advances in linguistics and AI point to a future in which intelligent systems will exist for parsing and coding texts (Bechtel forthcoming; Golumbic 1990; Zock and Sabah 1988; Gazdar and Mellish 1986; Winghart 1988). Such systems will enable the researcher to move beyond content analysis to more relational modes of text analysis (Schank and Abelson 1977; Sowa 1984; Roberts 1989, forthcoming-b). These techniques enable researchers to analyze societal shifts not just in terms of difference in word usage but in terms of differences in meaning (Roberts forthcoming-b; Carley and Palmquist 1992; Carley and Kaufer 1993). Because these techniques enable more automated coding of texts, they make possible the analysis of larger quantities of texts. This will make it possible for researchers to address important questions in a number of areas of sociological inquiry, including symbolic interactionism, sociology of science, sociology of knowledge, culture, and social change (Roberts forthcoming-b; Carley 1994a). Researchers are also exploring attaching these coding tools directly into retrieval and browsing facilities for electronic text. This will dramatically affect the sociology of knowledge (Lanham 1994).

Such techniques often focus on the text as a representation of the agent's mental model (Johnson-Laird 1983). As such, they focus on not just words or concepts but the relationships between them. These techniques are also often more tuned to the underlying semantics and rarely just process text on the bases of its syntax (Roberts forthcoming-a). Texts as networks of concepts can then be analyzed using standard social network techniques. Additionally, by focusing on relationships, researchers can examine texts for patterns of behavior and can trace the logic of arguments. Consequently, narratives or stories can be analyzed as event sequences (Abell 1984, 1988; Heise 1991). This enables researchers to address the relationship between agency and action. Finally, AI techniques for searching can be used not only to locate patterns in texts but also to locate appropriate texts to analyze. This will enable researchers to explore issues of organizational rhetoric and to locate appropriate historical and archival documents to analyze.

Researchers engaged in textual analysis, however, must take care in using the various new techniques to make sure that the assumptions built into the software match those that the researchers are making (Tesch 1989; Carley 1993). Importantly, AI knowledge acquisition and expert system techniques for knowledge acquisition can be used in the development of textual analysis tools to extract the researchers' assumptions and use these assumptions to generate partially customized text-coding software (Carley and Palmquist 1992; Carley 1993). Expert systems with social knowledge about concept meaning can also be used in the coding process. Such systems can decrease reliance on coders and increase interrater reliability (Carley 1988). Indeed, whether and how coders will be needed in the future of text analysis is a matter of debate (Shapiro forthcoming).

NETWORK OR PATTERN ANALYSIS

Search techniques from AI can be used to analyze social and citation networks (Hummon, Doreian, and Freeman 1990; Hummon and Carley 1993; Carley, Hummon, and Harty 1993). Hummon and Doreian (1989, 1990) developed a network search technique to locate the main path through scientific literature that identifies the connections between the key intellectual developments in scientific fields. These techniques can also be used to evaluate the development of a network through time. Search and pattern-matching techniques can also be used to analyze event sequences and to explore the internal logic of proposed verbal theories (Heise 1991).

Social network evolution and organizational change can also be studied using the techniques of neural networks (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986; Wasserman 1989, 1993; Karayiannis and Venetsanopoulos 1993), genetic algorithms and classifier systems (Holland 1975, 1992; Holland and Miller 1991; Holland et al. 1986), and simulated annealing (Kirkpatrick, Gelatt, and Vecchi 1983; Rutenbar 1989). Neural networks are meant to emulate the behavior of biological nervous systems. In a neural network, information is stored in the connections between nodes that are typically arranged in sequential layers (often three layers) such that there are connections between nodes in contiguous layers but not within a layer. Genetic algorithms draw from biologic theories of evolution. A genetic algorithm simulates the evolutionary process by allowing a population of entities to adapt over time through mutation and/or reproduction (crossover) in an environment in which only the most fit survive. Simulated annealers draw from the process of metal or chemical annealing. Simulated annealers search for the best solution by first proposing an alternative, seeing if its fit is better than the current system, adopting it if it is better, and otherwise adopting even the "bad" or risky move with some probability. The probability of accepting the bad move decreases over time as the "temperature" of the system cools. All three of these approaches are procedures that can be thought of as optimization techniques for classifying objects or locating solutions. They can also all be thought of as techniques for modeling complex adaptive agents (because such agents must be capable of solving problems through a search procedure).

An example of using one of these techniques for studying organizational change is found in the work by Carley. For example, Carley (1992) used a model, similar to a neural network, to examine how organizational structure constrains the ability of organizations to take advantage of the experiential lessons learned by the agents in the organization. As another example, Macy (1991b) uses evolutionary techniques to examine cooperation in social groups. Kontopoulos (1993) suggests that neural networks are an appropriate metaphor for understanding social structure, and Eccles and Crane (1988) suggest that annealing is an appropriate metaphor for organizational change. To date, most of the computational work using these techniques have examined networks of individuals that are largely undifferentiated in terms of their structural position. As a result, most of this work can only speak to the impact of making the underlying social network more or less dense. However, the possibility exists to combine these models with models of organizational or social structure to begin to examine how specific structures limit agent adaptive behavior. Work in this direction is the constructural model by Carley (1991b, see also Kaufer and Carley 1993) and the work by R. J. Collins (1992) on spatial constraints on evolution and social learning. Early results suggest that the existence of spatial or social structure may actually increase the effectiveness of individual agent learning and may increase the robustness and stability of the collectivities' ability to problem solve in the face of change among the constituent members.

Optimization and pattern-matching routines coming out of the work in AI such as classifier systems and simulated annealing can be used for partitioning graphs. The area of graph partitioning has received a great deal of attention by social scientists (Breiger, Boorman, and Arabie 1975; White, Boorman, and Breiger 1976) and computer scientists (Bodlaender and Jansen 1991). One such approach is the optimization approach proposed by Batagelj, Doreian, and Ferligoj (1992) for locating regular equivalence partitions. They use both a simple hill-climbing routine for one algorithm and an exhaustive search for another. Freeman (1993) uses genetic algorithms to partition members of a group into cliques or subgroups based on the proximity between individuals. In a related vein, Krackplot (Krackhardt, Blythe, and McGrath 1994) uses a simulated annealing technique to draw networks with minimal line overlap based on a routine developed by Eades and Harel (1989). This procedure aids in data interpretation.

EXPERT SYSTEMS

Expert systems make decisions by deriving conclusions from vast quantities of task-specific knowledge (Hayes-Roth, Waterman, and Lenat 1983; Gonzales and Dankel 1993). Expert systems use heuristic-based reasoning, sometimes in conjunction with probabilistic-based reasoning, to generate solutions, narrow the search for solutions, and determine critical problem features. Expert systems are typically developed in areas in which experts are rare, there is a clear distinction in the performance of experts and novices, and the task is sufficiently complex that no analytical solution is known. This is often the case for qualitative data analysis. Indeed, social scientists have heralded expert systems as being as important to qualitative data analysis as statistical packages are to quantitative data analysis (Brent 1986; Benfer, Brent, and Furbee 1991). Importantly, from a sociological point of view, the knowledge bases of these systems implicitly contain the norms, beliefs, preferences, and attitudes of the experts whose knowledge is encoded. Thus comparison of the knowledge bases for different expert systems in the same area provides insight into the social differences in that area.

Probably one of the greatest gains for social theorists to be had from the use of expert systems is in the area of theory development. Many social and organizational theories are often specified only in terms of rules of action. When possible, statements of these theories using predicate logic enable the researcher to determine logical consistency and to derive testable propositions. Salancik and Leblebici (1988; see also Leblebici and Salancik 1989) used this approach to examine the variety and form of organizational design. Sometimes, it is possible to specify only specific propositions and not the entire grammar. In this case, expert systems, rather than formal logic, provide a means for determining the logical consistency of the rules and their predictions.

These rules can be represented within an expert system. Then, using the expert systems techniques for applying rules, the researcher can determine which of the rules (or findings) are incompatible and can determine whether the specified rules are logically sufficient for making the conclusions claimed by verbal theorists. This approach has been used by Baligh, Burton, and Obel (1987, 1990, 1994), who created an expert system containing the many and varied rules (findings) of contingency theory. Similarly, Brent et al. (1989) represent Goffman's dramaturgical perspective as an expert system in which agents interact and behave according to Goffman's principles. Banerjee (1986) developed an expert system for the Skocpol and O'Donnel sociopolitical theory of coalitions and interests and used it to empirically demonstrate the plausibility and logical consistency of their argument. Fararo and Skvoretz (1984; Skvoretz and Fararo 1995) have suggested that production systems can be the basis for a general theory of action and social institutions. Drawing from the work of Talcott Parsons, they note that social norms can be thought of as rules linking situations to particular actions that organize social interaction.

Expert systems have also been used as models of reasoning and used to examine organizational and managerial decisions (Masuch 1990). Simon (1981a, 1981b) has repeatedly argued that any physical symbol system has the necessary and sufficient means for intelligent action. Expert systems fit this conceptualization. Indeed, expert systems are thought to emulate the behavior of experts, acting as they would on a "good" day and under ideal conditions. Most expert system architectures, however, have a rigid procedure for applying search processes. Thus they cannot make claims to being generally intelligent. Newell, Laird, and Rosenbloom (Newell 1990; Laird, Newell, and Rosenbloom 1987; Rosenbloom et al. 1991) developed a system referred to as Soar. Soar is a symbol-level system in which all action is goal oriented and characterized as search through problem spaces (limited knowledge domains). Social knowledge can be encoded as a series of preferences about the desirability and acceptability of various actions. Soar, like traditional expert system architectures, uses heuristic-based reasoning. However, unlike other expert system architectures, the decision procedure is not rigid. Rather, Soar employs a technique called universal subgoaling, which enables it to employ any and all search techniques whenever any of these procedures are called for.

Further, Soar has a built-in procedure for learning (creating new rules) called *chunking*. Because of these features, Soar has been characterized as exhibiting general intelligence, and not, like so many expert systems, intelligence specific to a particular task. It has also been demonstrated on numerous individual tasks that Soar behaves like human agents; thus, it can be reasonably argued that Soar is a unified theory of cognition (Newell 1990). This makes Soar a better candidate architecture for a model social agent (Carley and Newell 1994).

EMOTION

One of the areas receiving increased attention in both sociology and AI is the study of emotions. Researchers in AI point to the need to represent emotions in order for an agent to be taken as social agents (Carley and Newell 1994). As part of the Oz project, Bates, Loyall, and Reilly (1994) have developed a prototype architecture, called Tok, that supports agent reactivity, goals, action, emotion, and social behavior for multiple interacting agents. In this case, emotions aid in generating recognizable social behavior. Sloman and Croucher (1981) demonstrate that the understanding of motives is largely related to an understanding of emotions. Researchers in sociology point to the necessity of understanding emotion to predict social action (Heise 1979).

Repeatedly, researchers interested in story and text understanding demonstrate the importance of interpreting emotions to understand the story (Frijda 1993; Lehnert 1981; Lehnert and Vine 1987). For example, Dyer (1983:214) notes that emotional states of story characters are represented as knowledge structures referred to as AFFECT. These structures are used to follow the emotional states of story characters. Dyer found that affects are important to narrative comprehension for a variety of reasons: Emotions reveal underlying goals, emotions are not tied to a specific goal, emotions signal the occurrence of expectation failures (e.g., Joe is worried about *x*), emotions indicate the status of interpersonal relations, and emotions influence what thematic structures become instantiated in episodic memory (e.g., Is this a romance or a mystery?). On the basis of this analysis, Dyer argues that in terms of textual analysis, emotions are important for both processing and memory:

> Experience with BORIS strongly supports the claim that AFFECT processing, for narratives at least, is highly cognitive in nature. That is, the importance of emotions in narratives arises from the cognitive structures they signify. These structures reveal the goals of the characters, the goal outcomes they expect versus what actually occurs, and the long-term relationships which are active between the characters themselves. (P. 241)

MULTIAGENT MODELS

Researchers in AI recognize that people do not live in isolation but in societies. As such, multiagent computational models are relevant both to cognitive science and to theories of social action (Craig 1994b). Within AI and computer science, there is a growing interest in the exploration of multiagent models. These models range from the more symbolic DAI models to the models using one of the various complex adaptive agent techniques such as genetic algorithms (Holland et al. 1986), neural networks (Rumelhart and McClelland 1986; McClelland and Rumelhart 1986), simulated annealing (Kirkpatrick et al. 1983; Rutenbar 1989), and chunking (Laird, Rosenbloom, and Newell 1986a, 1986b). Some of this work rests on, or is related to, mathematical models of distributed teams (Pete, Pattipati and Kleinman 1993, 1994; Tang, Pattipati, and Kleinman 1993) and social psychological experimental work on teams (Hollenbeck, Ilgen, Sego, et al. 1995; Hollenbeck, Ilgen, Tuttle, and Sego 1995).

DAI is the application of AI techniques to the area of distributed decision making and behavior (Bond and Gasser 1988; Gasser and Huhns 1989). In such an environment, multiple agents are expected to interact and jointly function in a loosely coupled but coordinated fashion. Often, the focus is on increasing the effectiveness, efficiency, or performance of the various members of the society or organization and on increasing the overall effectiveness and efficiency of the society or organization. Much of this work employs what are often termed multiagent models—that is, models in which there are multiple intelligent and often adaptive agents. Many DAI systems are symbolic models that perform specific stylized tasks such as surveillance, navigation, or traffic control. DAI researchers address issues of the relationship between agent and group or social knowledge (Halpern and Moses 1990), socially shared cognition (Hutchins 1990, 1991), coordinated action (Durfee 1988; Durfee and Montgomery 1991; Gasser and Majchrzak 1992; Ishida, Gasser, and Yokoo 1992), planning (Gasser and Majchrzak 1994), organizational design (Majchrzak and Gasser 1991), and problem solving (Gasser and Ishida 1991; Ye and Carley 1995). Some DAI research employs general models of cognition that make very strong claims about the nature of agent intelligence (e.g., the work using Soar).

Within the various multiagent models, agents are treated very generically, and there is an implicit assumption that the results are generalizable to networks of computers, networks of software programs, and networks of people. Researchers have argued that the mind is reasonably conceptualized as a society of actors (Minsky 1988; Kornfeld and Hewitt 1981), that society is reasonable conceptualized as a collection of artificial agents (Carley 1991a, 1991b; Macy 1990; Glance and Huberman 1993), and that the society metaphor is useful for artificial societies (Moses and Tennenholtz 1991). However, the mapping between social and artificial agents, particularly mechanical agents such as robots, is not perfect. As Moses and Tennenholtz (1990, abstract) argue:

> Given the many differences between people and robots, existing social systems cannot simply be copied as is for an artificial "society." Rather, such systems should be studied in order to provide a good understanding of their underlying structure, and the problems they successfully address. These will form the basis for the design of artificial social systems.

Nevertheless, artificial social systems composed of multiple artificial agents have many uses in the development of social theory (Castelfranchi and Werner 1992; Gilbert and Heath 1985). At issue is how social these agents need to be and how the structure of the artificial society interacts with agent capabilities in affecting long- and short-term social dynamics. Answers to these questions, whether or not they help in the design of societies of robots, will be valuable in increasing our understanding of human social behavior.

THE SOCIALNESS OF ARTIFICIAL AGENTS

Socialness derives both from limitations to an agent's cognitive capabilities and from acquisition of multiple types of knowledge (Carley and Newell 1994). Agents that are too capable cognitively know too much and so have no need for social interaction or learning. Social interaction and learning are the basis for social life (Carley 1991b). Complex social and organizational phenomena emerge from the interactions among even simple agents (Shoham and Tennenholtz 1994). Social dynamics are determined, at least in part, by agent cognitive capabilities such as learning (Oliveira 1992; R. J. Collins 1992; Carley 1989, 1991b).

Multiagent models with more restricted cognitive capabilities exhibit a greater variety of social behaviors. For example, Plural-Soar agents (Carley et al. 1992) use a cognitive agent that is more restricted than the boundedly rational agent used in AAIS (Masuch and LaPotin 1989). Further, the AAIS model effectively contains more types of social information than does the Plural-Soar model. Combining these two models led to an agent that was capable of exhibiting a greater range of social behaviors (Verhagen and Masuch 1994).

Most DAI models employ symbolic agents that are either boundedly rational agents or simple cognitive agents. In most of these models, although agents may interact, their interaction is rarely based on a detailed model of the other agents that takes into account the social position of the other agents and the position of the ego or self agent in the society. Craig (1994a) used a symbolic interactionist perspective in developing the self-and-acquaintance multiagent models and analyzes these models as devices for making predictions about, and explanations of, agents' behavior. Craig's DAI agent models thus have more, and different, social knowledge than do traditional DAI models. He suggests that the symbolic interactionist view is superior to what he characterizes as the "ad hoc, static and somewhat low-level approach" currently adopted by DAI researchers. Craig's agents exhibit more social-like behavior than traditional DAI models.

These two lines of research suggest that increasing the realism of the agent—by restricting its cognitive capability and/or increasing the amount or type of knowledge available to the artificial agent—enables collections of these agents to increasingly exhibit behavior that is social in nature. An important caveat, however, is that having a model social agent (Carley and Newell 1994) that is maximally capable of producing all social behaviors is not necessary for studying many social phenomena. Indeed, many social behaviors, such as cooperation, play, and diffusion, can be adequately studied with less social agents. Particularly when the researcher is interested in macro social behavior, it may be possible to make due with, for example, agents who are only boundedly rational and with only minimal cultural and historical knowledge (Carley forthcoming).

THE IMPORTANCE OF CONCURRENCY FOR ADAPTIVE AGENTS

Recently, the work in the area of complexity—specifically, that on artificial life—has captured the imagination of a number of social scientists (Langton 1992, 1994). Much of the work in the area of artificial life is not technically in the AI area; however, from the point of view of the development of social theory, that is irrelevant. These artificial life models follow from a biological metaphor. The aspect of this movement that ties in most directly with AI, and interestingly with sociology, is the work on evolution (Jefferson 1990). Many of these models use some form of genetic algorithm or neural network algorithm.

Most sociologists and AI researchers interested in social evolution point out that socially organized agents must be adaptive by nature in order for social structure to be dynamic. In fact, many theorists argue that the dynamic of social and organizational life emerges from the concurrent adaptive behavior of the individual agents in the society or group. Shoham and Tennenholtz (1994) refer to this process behavior as co-learning—"a process in which several agents simultaneously try to adapt to one another's behavior so as to produce desirable global system properties." Manthey (1990) points to the importance of synchronization as a basis for both evolution and development of hierarchy. Carley (1991b) argues that concurrent interaction is necessary for the emergence of social stability and consensus. Moreover, concurrent interaction among agents when combined with access to different forms of communication media can effect radical changes in the ability of subgroups to acquire novel information and to be socialized (Carley 1995a).

The agents in these complex adaptive multiagent models are nondeterministic and undergo a coevolutionary process. During their lifetimes, they may move through and interact with the environment, sexually reproduce, consume resources, age, learn, and die. In some systems, such as Enact (Oliveira 1992), the agents are contiguous cells in which one single cell can be thought of as the agent's genotype and others as phenotypes that can change as the agent interacts with the environment. A result is that the process of coevolution is both genotypic and mimetic. In this system, emergent social phenomena and the evolutionary dynamics depend on the rate at which the agents age.

Some, though by no means all, of the work on evolution is carried out using genetic algorithms. Traditional genetic algorithms are driven by a goal optimization process. In contrast, the artificial evolution genetic algorithms try to achieve greater biological realism by providing a clear separation between genotype and the information encoded in the genotype, and by determining genotype fitness by decoding and executing the program, perhaps in an environment that is shared by the other members of the population (R. J. Collins 1992). The genotype is typically represented as a string. Recombination and mutation operators act randomly on the string, ignoring both syntactic and semantic structure. Realistic population dynamics may be achieved by simulating large populations (tens of thousands of agents) and by having the agent's position in the spatial structure influence its selection and mating process. Much of this work is on generic agents or attempts to model low-level biological organisms such as parasites or ants (R. J. Collins 1992). For example, R. J. Collins (1992) uses an artificial neural network encoding scheme to evolve ant-like behavior by putting the connection strengths and the connectivity pattern under genetic control.

Within organizational theory, researchers have found that modeling organizations as collections of intelligent adaptive agents acting more or less concurrently is critical for understanding issues of organizational learning and design. Organizational learning focuses on performance improvement and adaptation to the organization's external and internal environment. Elofson and Konsynski (1993) apply AI and machine learning techniques to the analysis of organizational learning for the purpose of monitoring and analyzing decisions relative to organizational structure and for monitoring organizational changes as part of the organizational learning and adaptation cycle. Their analysis demonstrates that increased flexibility is possible by knowledge caching, which provides a means to realize an explicit organizational memory when information and processing capabilities are distributed among the organizational members. Carley (1991a, 1992) used an approach akin to neural networks to represent hierarchies and demonstrated that when organizational learning was embedded in the relationships between agents and not just in the agents, the organization was more robust in the face of "crises" such as turnover and erroneous information. Organizational evolution can be examined by using genetic algorithms to simulate the behavior of populations of organizations evolving their forms over time. Here, the concurrency across multiple organizations is key to determining the dynamics of organizational survival. Crowston (1994, forthcoming) has used this approach to examine Thompson's theory of organizational forms and the evolution of novel forms. Typically, however, within organizations, all agents cannot act in a completely concurrent fashion, as one agent may not be able to begin a particular task until another agent has finished a different task. Thus issues of scheduling and coordination are paramount. Coordination of these intelligent agents can be thought of as a search process through a hierarchical behavior space (Durfee 1988; Durfee and Montgomery 1991). The organizational design can also act as a socially agreed upon coordination schedule constraining agent opportunities for action. This approach is taken in the virtual design team in which pert chart techniques are combined with hierarchical organizational structures to enable the researcher to examine issues of information flow and organizational design (Levitt et al. 1994; Cohen 1992). For an overview of the role of AI and computational modeling in organizations and institutions, see Ennals (1991) and Carley (1995b).

GENERAL FINDINGS

Many researchers interested in adaptive agents have focused on the emergence of conventions and/or the evolution of cooperation. These studies demonstrate that seemingly simple rules, such as trying to attain the highest cumulative award, often give rise to interesting and nontrivial social dynamics (Shoham and Tennenholtz 1994). One problem that lends itself to computational modeling is the tragedy of the commons and, relatedly, the prisoner's dilemma (Axelrod 1987; Skvoretz and Fararo 1995; Turner 1993). Turner (1993) argues that this problem of resource consumption exists for artificial societies and that computational studies about how to avoid or mitigate the problem will have value both to the design of computer systems and to societies.

Much of the adaptive agent research leads to the following conclusions. Social dynamics, both equilibrium and nonequilibrium, depend on the rate of agent learning (adaptation or evolution; e.g., Oliveira 1992; R. J. Collins 1992; Carley 1991b). Spatial or organizational structure enables faster and more robust social or organizational learning (R. J. Collins 1992; Carley 1992). Social chaos is reduced by having intelligent adaptive agents determine their actions using strategies based on observations and beliefs about others (Kephart, Huberman, and Hogg 1992). This suggests that cognition and structure play a dual defining role on emergent phenomena. In some sense, this argument is not new. Simon, in proposing his bounded rational agent model, argued both that the agent itself was boundedly rational (Simon 1955) and that the environment and the interrelations between the environment and the agent set bounds on, and therefore constrain, the agent (Simon 1956). Simon (1956) suggested that the agent's mental models employ simplifications of reality that "may depend not only on the characteristics—sensory, neural, and other—of the organism, but equally upon the structure of the environment" (p. 130). However, the recent work goes beyond this claim in several important ways. Specificity of cognition, specificity of structure, and, importantly, specificity of task result in specific claims and the modeling of actual biological organisms. Indeed, it can be argued that with neither a detailed task nor agent model, communication becomes featureless, actions become uniform, and agents become skill-less. A consequence is that many important social and organizational phenomena cannot be studied without specifying agent, structure, and task.

Importantly, these more detailed models have shown that complex social and organizational behavior emerge from even simple interactions among simplistic agents and that this emergence is discontinuous. For example, simple interaction among individuals has a strong and nonlinear impact on effectively generating social sanctions, cues, and norms (Macy 1990; Horgan 1994); avoiding or getting out of social traps (Glance and Huberman, 1994; Macy 1990, 1991a, 1991b); and achieving social stability and information diffusion (Carley 1991b, forthcoming). Further, it has been repeatedly demonstrated that environmental and institutional factors such as payoffs, population dynamics, and population structure influence the evolution of cooperation in a discontinuous fashion (Axelrod and Dion 1988); that structure can, under certain conditions, mitigate the effect of individual learning and turnover on organizational performance (Carley 1992); and that simple interactions can lead to the development of and change in beliefs (Ballim and Wilks 1991). These models demonstrate the ability of complex social phenomena to emerge from simple processes (Dutton and Starbuck 1971).

As agent models have become increasingly cognitively realistic, it has become clear that many actions are dictated by the social norms and beliefs of the agents. For example, in Plural-Soar, how an agent responds to a communication of another agent depends on the norm the agent follows when choosing between otherwise equivalent actions (Carley et al. 1992). Essentially, human action is viewed as a result of a goal-directed search process through the set of possible actions. What task the agent is engaged in and the cognitive capabilities of the agent constrain what actions are available for the agent to choose among. However, these factors do not completely eliminate all alternatives. Thus it is left to the agent's preferences (or embedded social norms) to make the final determination.

DISCUSSION

The potential for linkage between sociology and AI, and computational analysis more generally, is growing. Computational techniques from AI are making possible new methodologies for the analysis of data and for theory building. One of the important benefits of computational modeling is showing proof of concept and doing theory building. For example, Craig (1994a) points out that in the multiagent symbolic interactionist model, the self model is an important precursor for the acquaintance model. This finding is illustrative of how modeling helps develop theory by forcing the researcher to make assumptions explicit and demonstrating the plausibility or impossibility of basic assumptions.

The potential for fruitful linkages between sociology and AI is growing. To turn this potential into a reality, sociologists will need to be trained in the science of computational methods. This will require learning not just how to program but also how to develop models, design virtual experiments, dock models, and evaluate computational data. It will also require noncomputational modelers to develop an understanding of how to appropriately evaluate, review, and use the findings from computational analyses. Virtual experiments are comparable to human experiments in terms of the difficulty in running them and in analyzing the results. A set of canonical tasks may be useful for developing this area. And finally, it will be necessary to developing a better understanding of what it means to be social and the types of knowledge necessary for a social agent.

Insightful research on social action requires surprisingly simple tasks. Task simplicity can be advantageous for the researcher, as it clarifies the relationship between individual action, the emergence of social behavior, and the structural and cultural constraints. Within many scientific disciplines, as multiple researchers attended to a small number of tasks that could be performed either in the human or in the computational laboratory, scientific knowledge accumulated. Examples are the Tower of Hanoi and the Eight Puzzle in cognitive psychology and the Prisoner's Dilemma in political science. Social and organizational theory can also benefit from such "Drosophila tasks." Clearly, such canonical tasks can be limiting when one is concerned with adaptive agent behavior (Cliff 1994). For example, the lessons learned may not necessarily generalize or scale to large social systems with tens, hundreds, or even thousands of agents. Thus issues of scalability are critical. However, the adaptive behavior of artificial worlds that appear on the surface to be "toy"-like, and so may ill afford advances in the social sciences, still demonstrates the importance of synchronicity, and there are some suggestions that behavioral patterns will scale (Prietula and Carley 1994). Thus the jury is still out on whether or not all social behavior will emerge from looking at a set of canonical tasks, particularly for very large societies. However, such tasks, particularly those requiring interaction among a minimum of three agents (so that issues of coalition can emerge), should be valuable for advancing our understanding of society.

Simon (e.g., 1981a, 1981b) has argued repeatedly that a computer science is worth developing because it can reveal insights into the operation of the human mind. Essentially, the computer and the human mind are both symbolic processors—that is, they must both use symbols to designate concepts, objects, other agents, and relations, and they must both develop processes for interpreting, executing, storing, and retrieving these expressions in symbolic form. These common features lead Simon to conclude that any physical symbol system has the necessary and sufficient means for intelligent action. However, if the agent has sufficient means for intelligent action, that does not necessarily guarantee that the agent has sufficient means for social action (Carley and Newell 1994). Carley and Newell (1994) have argued that to be social, the agent must have limited cognitive capabilities and must be embedded in a sociocultural-historical environment and have at least limited knowledge about the various aspects of this environment. Essentially, they argue that agents who are too cognitively capable have no need to engage in certain actions that social agents engage in. Further, agents in an impoverished situation with correspondingly impoverished knowledge have no need for all the actions in which social agents engage. Figure 1 shows the additional actions possible for an agent with particular cognitive capabilities and in a particular situation. Moving from the upper left corner to the lower right corner, an agent in a particular cell is capable of all the actions above and to the left of its position.

Also in Figure 1, some of the various systems that have been discussed in this article are shown in italics. These artificial agents (or collections of artificial agents) should, according to the Carley and Newell (1994) theory, be capable of all the actions above and to the left of their position. This chart becomes a way of assessing the social capabilities and limitations of proposed social agents and of determining the minimal agent necessary to explore a particular social action. As can be seen, there is no candidate agent that is completely socially capable. Further, explaining a phenomenon such as ritual maintenance from the vantage of a collection of adaptive agents requires agents who are much more social than those currently in existence.

Knowledge

Increasingly Rich Situation in Which Agent Is Embedded and Has Corresponding Knowledge of →

Processing: Increasingly Limited Capabilities ↓	Nonsocial Task	Multiple Agents	Real Interaction	Social Structural	Social Goals	Cultural Historical
Omnipotent Agent	goal directed models of self produces goods uses tools uses language	models of others turn taking	face-to-face timing constraints	socially situated class differences	social goals organizational goals	historical motivation
Rational Agent	*Logic thm prover* reasons acquires information	learns from others education general diffusion	scheduling	social ranking social mobility competition	disillusionment	*Craig (1994a)* social inheritance social cognition
Boundedly Rational Agent	*ACTION, HITOP-A* *Expert Systems* *Simulated Annealing* *Genetic Algorithms* satisfices task planning adaptation	*Chess Programs* group making	*VDT* social planning coercion priority disputes mis-communication	*AAIS* *Constructuralism* restraints on mobility uses networks for information corporate intelligence	*Axelrod (1987)* *Macy (1990, 1991b)* *Artificial Life, Enact* party line voting delays gratification moral obligation cooperation altruism	*Erving* gate keeping controlled diffusion etiquette deviance roles sanctions
Cognitive Agent	*Neural Network* *Genetic Neural Nets* *Soar* compulsiveness lack of awareness interruptability automatic action	group think	*Plural-Soar* crisis response	*TASCCS* *Radar-Soar* automatic responce to status cues	*R. Collins (1992)* clan wars power struggles	*ACTS theory* develop language role development institutions
Emotional Cognitive Agent	intensity habituation variable performance	*TOK* protesting courting	mob action play rapid emotional response	campaigning conformity	nationalism patriotism	*Model Social Agent* norm maintenance ritual maintenance advertising social play

Figure 1: Artificial Agents and Their Expected Action Sets

Simon (1981a, 1981b) has also argued that rationality can be used to explain organizational and social traits and so can predict social behavior even when there is not a detailed understanding of the specific decision processes in which agents engage. By rationality, he means that the agent will bring to bear in solving a problem all the knowledge that the agent currently has. This moves the scientific debate from understanding rationality per se to defining what information is known by the agent. Turner's (1988) work on social interaction points in this direction. Research from a variety of traditions suggests that it is probably not possible, or even reasonable, to try to list all social knowledge; however, it may be a valuable intellectual exercise for social theorists to attempt to create a categorization scheme for the types of knowledge necessary for an agent to be social. The categories by Carley and Newell (1994) are a move in this direction. Having an understanding of the possible categories will guide development of social theory by providing a basis for judging the adequacy of the theory. Even more important will be developing an understanding of the sociocognitive mechanisms for creating, attaining, elaborating, forgetting, and altering social knowledge; mechanisms for creating, maintaining, and dissolving social institutions; and social constraints on knowledge acquisition and dissemination. This is where much of the work on artificial social agents is directed.

REFERENCES

Abell, Peter. 1984. "Comparative Narratives, Some Rules for the Study of Action." *Journal for the Theory of Social Behavior* 14:309-32.

Abell, Peter. 1988. "The Structuration of Action: Inference and Comparative Narratives." In *Actions and Structures*, edited by N. G. Fielding. London: Sage.

Axelrod, Robert M. 1987. "The Evolution of Strategies in the Iterated Prisoner's Dilemma." Pp. 32-41 in *Genetic Algorithms and Simulated Annealing*, edited by L. Davis. London: Pitman.

Axelrod, Robert M. and Doug Dion. 1988. "The Further Evolution of Cooperation." *Science* 242:1385-90.

Bainbridge, William, Edward Brent, Kathleen Carley, David Heise, Michael Macy, Barry Markovsky, and John Skvoretz. 1994. "Artificial Social Intelligence." *Annual Review of Sociology* 20:407-436.

Baligh, Helmy H., Richard M. Burton, and Borge Obel. 1987. "Design of Organizational Structures: An Expert System Method." Pp. 177-82 in *Economics and Artificial Intelligence*, edited by J. L. Roos. Oxford, UK: Pergamon.

———. 1990. "Devising Expert Systems in Organization Theory: The Organizational Consultant." Pp. 35-57 in *Organization, Management, and Expert Systems*, edited by M. Masuch. Berlin: Walter De Gruyter.

———. 1994. "Validating the Organizational Consultant on the Fly." Pp. 179-94 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.

Ballim, Afzal and Yorick Wilks. 1991. *Artificial Believers: The Ascription of Belief.* Hillsdale, NJ: Lawrence Erlbaum.

Banerjee, Saibal. 1986. "Reproduction of Social Structures: An Artificial Intelligence Model." *Journal of Conflict Resolution* 30:221-52.

Batagelj, Vladimir, Patrick Doreian, and Anuska Ferligoj. 1992. "An Optimizational Approach to Regular Equivalence." *Social Networks* 14:121-36.

Bates, Joseph, A. Bryan Loyall, and W. Scott Reilly. 1994. "An Architecture for Action, Emotion, and Social Behavior." Pp. 55-68 in *Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-Agent World.* New York: Elsevier North-Holland.

Bechtel, Robert. Forthcoming. "Developments in Computer Science With Application to Text Analysis." In *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences From Texts and Transcripts*, edited by C. W. Roberts. Mahweh, NJ: Lawrence Erlbaum.

Benfer, Robert A., Edward E. Brent, and N. Louanna Furbee. 1991. *Expert Systems*. Newbury Park, CA: Sage.

Blank, Grant, James L. McCartney, and Edward E. Brent, eds. 1989. *New Technology in Sociology: Practical Applications in Research and Work.* New Brunswick, NJ: Transaction.

Bodlaender, H. L. and Klaus Jansen. 1991. "Restrictions of Graph Partition Problems: Part 1." Research Paper No. RUU-CS-91-44, Department of Computer Science, Utrecht University, The Netherlands.

Bond, Alan and Les Gasser. 1988. *Readings in Distributed Artificial Intelligence*. San Mateo, CA: Kaufmann.

Breiger, Ronald L., Scott A. Boorman, and Phipps Arabie. 1975. "An Algorithm for Clustering Relational Data With Applications to Social Network Analysis and Comparison With Multidimensional Scaling." *Journal of Mathematical Psychology* 12:328-83.

Brent, Edward E. 1986. "Knowledge-Based Systems: A Qualitative Formalism." *Qualitative Sociology* 9:256-82.

Brent, Edward E., Jack Glazier, Keith Jamtgaard, Edward Wetzel, Peter Hall, Michael Dalecki, and Abdoulaye Bah. 1989. "Erving: A Program to Teach Sociological Reasoning From the Dramaturgical Perspective." *Teaching Sociology* 17:38-48.

Carley, Kathleen M. 1988. "Formalizing the Social Expert's Knowledge." *Sociological Methods & Research* 17:165-232.

———. 1989. "The Value of Cognitive Foundations for Dynamic Social Theory." *Journal of Mathematical Sociology* 14:171-208.

———. 1991a. "Designing Organizational Structures to Cope With Communication Breakdowns: A Simulation Model." *Industrial Crisis Quarterly* 5:19-57.

———. 1991b. "A Theory of Group Stability." *American Sociological Review* 56:331-54.

———. 1992. "Organizational Learning and Personnel Turnover." *Organization Science* 3:20-46.

———. 1993. "Coding Choices for Textual Analysis: A Comparison of Content Analysis and Map Analysis." Pp. 75-126 in *Sociological Methodology*, vol. 23, edited by P. Marsden. Oxford, UK: Blackwell.

———. 1994a. "Extracting Culture Through Textual Analysis." *Poetics* 22:291-312.

———. 1994b. "Sociology: Computational Organization Theory." *Social Science Computer Review* 12:611-24.

———. 1995a. "Communication Technologies and Their Effect on Cultural Homogeneity, Consensus, and the Diffusion of New Ideas." *Sociological Perspectives* 38:547-71.

———. 1995b. "Computational and Mathematical Organization Theory: Perspective and Directions." *Computational and Mathematical Organization Theory* 1:39-56.

———. Forthcoming. "A Comparison of Artificial and Human Organizations." *Journal of Economic Behavior and Organization*.

Carley, Kathleen M., Norman Hummon, and Martha Harty. 1993. "Scientific Influence: An Analysis of the Main Path Structure in the Journal of Conflict Resolution." *Knowledge: Creation, Diffusion, Utilization* 14:417-47.

Carley, Kathleen M. and David. S. Kaufer. 1993. "Semantic Connectivity: An Approach for Analyzing Semantic Networks." *Communication Theory* 3:183-213.

Carley, Kathleen M., Johan Kjaer-Hansen, Michael Prietula, and Allen Newell. 1992. "Plural-Soar: A Prolegomenon to Artificial Agents and Organizational Behavior." Pp. 87-118 in *Artificial Intelligence in Organization and Management Theory*, edited by M. Masuch and M. Warglien. Amsterdam: Elsevier Science.

Carley, Kathleen M. and Allen Newell. 1994. "The Nature of the Social Agent." *Journal of Mathematical Sociology* 19:221-62.

Carley, Kathleen M. and Michael Palmquist. 1992. "Extracting, Representing and Analyzing Mental Models." *Social Forces* 70:601-36.

Castelfranchi, Cristiano and Eric Werner, eds. 1992. *Artificial Social Systems: 4th European Workshop on Modeling Autonomous Agents in a Multi-Agent World*. Cimino, Italy: MAAMAW.

Cliff, David. 1994. "AI and A-Life: Never Mind the Blocksworld." Cognitive Science Research Paper No. CSRP-323, University of Sussex at Brighton, UK.

Cohen, Geoff P. 1992. "The Virtual Design Team: An Information Processing Model of Coordination in Project Design Teams." Ph.D. dissertation, Department of Civil Engineering, Stanford University, Stanford, CA.

Cohen, Michael D. 1986. "Artificial Intelligence and the Dynamic Performance of Organizational Designs." Pp. 53-70 in *Ambiguity and Command: Organizational Perspectives on Military Decision Making*, edited by J. G. March and R. Weissinger-Baylon. Marshfield, MA: Pitman.

Collins, Harry M. 1990. *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge, MA: MIT Press.

Collins, Randall. 1992. "Can Sociology Create an Artificial Intelligence?" Pp. 155-84 in *Sociological Insight*. New York: Oxford University Press.

Collins, Robert James. 1992. "Studies in Artificial Evolution." Research Paper No. CSD-920037, Computer Science Department, University of California, Los Angeles.

Craig, Iain. 1994a. "Agents That Model Themselves." Research Paper No. RR266, Department of Computer Science, University of Warwick, Warwick, UK.

———. 1994b. "A Perspective on Multi-Agent Systems." Research Paper No. RR273, Department of Computer Science, University of Warwick, Warwick, UK.

Crowston, Kevin. 1994. "Evolving Novel Organizational Forms." Pp. 19-38 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.

———. Forthcoming. "An Approach to Evolving Novel Organizational Forms." *Computational and Mathematical Organization Theory*.

Drescher, Gary L. 1986. "Genetic AI: Translating Piaget Into Lisp." Memo 890, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

Durfee, Edmund H. 1988. *Coordination of Distributed Problem Solvers*. Boston, MA: Kluwer.

Durfee, Edmund H. and Thomas A. Montgomery. 1991. "Coordination as Distributed Search in a Hierarchical Behavior Space." *IEEE Transactions on Systems, Man, and Cybernetics* 21:1363-78.

Dutton, John M. and William H. Starbuck. 1971. *Computer Simulation of Human Behavior*. New York: John Wiley.

Dyer, Michael. 1983. "The Role of Affect in Narratives." *Cognitive Science* 7:211-42.

Eades, Peter and D. Harel. 1989. "Drawing Graphs Nicely Using Simulated Annealing." Technical Report No. CS89-13, Department of Applied Mathematics Computer Science, Weizmann Institute of Science, Rehovot, Israel.

Eccles, Robert G. and Dwight B. Crane. 1988. *Doing Deals: Investment Banks at Work*. Boston, MA: Harvard Business School Press.

Elofson, Gregg S. and Benn R. Konsynski. 1993. "Performing Organizational Learning With Machine Apprentices." *Decision Support Systems*, 10:109-119.

Ennals, John Richard. 1991. *Artificial Intelligence and Human Institutions*. London, New York: Springer-Verlag.

Fararo, Thomas J. and John Skvoretz. 1984. "Institutions as Production Systems." *Journal of Mathematical Sociology* 10:117-82.

Freeman, Linton C. 1993. "Finding Groups With a Simple Genetic Algorithm." *Journal of Mathematical Sociology* 17:227-41.

Frijda, Nico H., ed. 1993. *Cognition and Emotion*. Hillsdale, NJ: Lawrence Erlbaum.

Gasser, Les and Michael N. Huhns, eds. 1989. *Distributed Artificial Intelligence*. Vol. 2. Los Altos, CA: Kaufmann.

Gasser, Les and Toru Ishida. 1991. "A Dynamic Organizational Architecture for Adaptive Problem Solving." Pp. 185-90 in *Proceedings of the Ninth National Conference On Artificial Intelligence*. Anaheim, CA.

Gasser, Les and Ann Majchrzak. 1992. "HITOP-A: Coordination, Infrastructure, and Enterprise Integration." Pp. 373-78 in *Proceedings of the First International Conference on Enterprise Integration*. Hilton Head, SC: MIT Press.

———. 1994. "ACTION Integrates Manufacturing Strategy, Design, and Planning." Pp. 133-6 in *Ergonomics of Hybrid Automated Systems IV*, edited by P. Kidd and W. Karwowski. Amsterdam, The Netherlands: IOS Press.

Gazdar, Gerald and Christopher S. Mellish. 1986. "Computational Linguistics." Cognitive Science Research Paper No. CSRP-058, Cognitive Studies Programme, School of Social Sciences, University of Sussex at Brighton, UK.

Gilbert, Nigel G. and Christian Heath. 1985. *Social Action and Artificial Intelligence*. Brookfield, VT: Gower.

Glance, Natalie S. and Bernardo A. Huberman. 1993. "The Outbreak of Cooperation." *Journal of Mathematical Sociology* 17:281-302.

———. 1994. "Social Dilemmas and Fluid Organizations." Pp. 217-40 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.

Golumbic, Martin Charles, ed. 1990. *Advances in Artificial Intelligence: Natural Language and Knowledge-Based Systems*. New York: Springer-Verlag.

Gonzalez, A. J. and Douglas D. Dankel. 1993. *The Engineering of Knowledge-Based Systems*. Englewood Cliffs, NJ: Prentice Hall.

Halpern, Joseph Y. and Yoram Moses. 1990. "Knowledge and Common Knowledge in a Distributed Environment." *Journal of Association for Computing Machinery* 37(3):549-87.

Hayes-Roth, Frederick, Donald A. Waterman, and Douglas B. Lenat, eds. 1983. *Building Expert Systems*. Reading, MA: Addison-Wesley.

Heise, David. 1979. *Understanding Events: Affect and the Construction of Social Action*. New York: Cambridge University Press.

———. 1991. "Event Structure Analysis: A Qualitative Model of Quantitative Research." Pp. 136-63 in *Using Computers in Qualitative Research*, edited by N. G. Fielding and R. M. Lee. Newbury Park, CA: Sage.

Holland, John H. 1975. *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.

———. 1992. "Genetic Algorithms." *Scientific American* 267 (July): 66-72.

Holland, John, Keith Holyoak, Richard Nisbett, and Paul Thagard. 1986. *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

Holland, John H. and John Miller. 1991. "Artificial Adaptive Agents in Economic Theory." *American Economic Review, Papers and Proceedings* 81:365-70.

Hollenbeck, John R., Daniel R. Ilgen, Dale B. Tuttle, and Douglas J. Sego. 1995. "Team Performance on Monitoring Tasks: An Examination of Decision Errors in Contexts Requiring Sustained Attention." *Journal of Applied Psychology* 80:685-96.

Hollenbeck, John R., Daniel R. Ilgen, Douglas J. Sego, Jennifer Hedlund, Debra A. Major, and Jean M. Phillips. 1995. "Multi-Level Theory of Team Decision Making: Decision Performance in Teams Incorporating Distributed Expertise." *Journal of Applied Psychology* 80:292-316.

Horgan, J. 1994. "Sex, Death and Sugar." *Scientific American* 271 (5): 20-24.

Hudson, Valerie M., ed. 1991. *Artificial Intelligence and International Politics*. Boulder, CO: Westview.

Hummon, Norman and Kathleen Carley. 1993. "Social Networks: As Normal Science." *Social Networks* 15:71-106.

Hummon, Norman P. and Patrick Doreian. 1989. "Connectivity in a Citation Network: The Development of DNA Theory." *Social Networks* 11:39-63.

———. 1990. "Computational Methods for Social Network Analysis." *Social Networks* 12:273-88.

Hummon, Norman P., Patrick Doreian, and Linton Freeman. 1990. "Analyzing the Structure of the Centrality Literature Between 1948 and 1979." *Knowledge* 11:459-79.

Hutchins, Edwin. 1990. "The Technology of Team Navigation." Pp. 191-220 in *Intellectual Teamwork*, edited by J. Galegher, R. Kraut, and C. Egido, Hillsdale, NJ: Lawrence Erlbaum.

———. 1991. "The Social Organization of Distributed Cognition." Pp. 238-307 in *Perspectives on Socially Shared Cognition*, edited by L. B. Resnick, J. M. Levine, and S. D. Teasley. Washington DC: American Psychological Association.

Ishida, Toru, Les Gasser, and Makoto Yokoo. 1992. "Organization Self-Design of Distributed Production Systems." *IEEE Transactions on Knowledge and Data Engineering* 4:123-34.

Jefferson, David. 1990. "Evolution as a Theme in Artificial Life: The Genesys/Tracker System." Research Paper No. CSD-900047, Computer Science Department, University of California, Los Angeles.

Johnson-Laird, Philip N. 1983. *Mental Models: Toward a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, MA: Harvard University Press.

Karayiannis, Nicolaos B. and Anastasios N. Venetsanopoulos. 1993. *Artificial Neural Networks: Learning Algorithms, Performance Evaluation, and Applications*. Boston, MA: Kluwer.

Kaufer, David S. and Kathleen M. Carley. 1993. *Communication at a Distance: The Effect of Print on Socio-Cultural Organization and Change*. Hillsdale, NJ: Lawrence Erlbaum.

Kephart, Jeff O., Bernardo A. Huberman, and Tad Hogg. 1992. "Can Predictive Agents Prevent Chaos?" Pp. 41-55 in *Economics and Cognitive Science*, edited by P. Bourgine and B. Walliser. Oxford: Pergamon.

Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi. 1983. "Optimization by Simulated Annealing." *Science* 220:671-80.

Kontopoulos, K. M. 1993. "Neural Networks as a Model of Structure." Pp. 243-67 in *The Logics of Social Structure*. New York: Cambridge University Press.

Kornfeld, William A. and Carl E. Hewitt. 1981. "The Scientific Community Metaphor." *IEEE Transactions on Man, Systems, and Cybernetics* 1:24-33.

Krackhardt, David, Jim Blythe, and Cathleen McGrath. 1994. "KrackPlot 3.0: An Improved Network Drawing Program." *Connections* 17:53-55.

Laird, John E., Allen Newell, and Paul S. Rosenbloom. 1987. "Soar: An Architecture for General Intelligence." *Artificial Intelligence* 33:1-64.

Laird, John E., Paul S. Rosenbloom, and Allen Newell. 1986a. "Chunking in Soar: The Anatomy of a General Learning Mechanism." *Machine Learning* 1:11-46.

———. 1986b. *Universal Subgoaling and Chunking: The Automatic Generation and Learning of Goal Hierarchies*. Boston, MA: Kluwer.

Langton, Christopher G., ed. 1992. *Artificial Life II: Proceedings of the Workshop on Artificial Life Held February 1990 in Santa Fe, New Mexico*. Santa Fe Institute Studies in the Sciences of Complexity Proceedings, vol 10. Redwood City, CA: Addison-Wesley.

———, ed. 1994. *Artificial Life III: Proceedings of the Workshop on Artificial Life, Held June 1992 in Santa Fe, New Mexico*. Santa Fe Institute Studies in the Sciences of Complexity Proceedings, vol. 17. Reading, MA: Addison-Wesley.

Lanham, Richard. 1994. "The Implications of Electronic Information for the Sociology of Knowledge." *Leonardo* 27:155-64.

Leblebici, Huseyin and Gerald R. Salancik. 1989. "The Rules of Organizing and the Managerial Role." *Organization Studies* 10:301-25.

Lehnert, Wendy. 1981. "Plot Units and Narrative Summarization." *Cognitive Science* 5:293-331.

Lehnert, Wendy and Edward Vine. 1987. "The Role of Affect in Narrative Structure." *Cognition and Emotion* 1:299-322.

Levitt, Raymond E., Geoff P. Cohen, John C. Kunz, Chris I. Nass, Torre Christiansen, and Yan Jin. 1994. "A Theoretical Evaluation of Measures of Organizational Design: Interrelationship and Performance Predictability." Pp. 1-18 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula . Hillsdale, NJ: Lawrence Erlbaum.

Macy, Michael W. 1990. "Learning Theory and the Logic of Critical Mass." *American Sociological Review* 55:809-826.

———. 1991a. "Chains of Cooperation: Threshold Effects in Collective Action." *American Sociological Review* 56:730-47.

———. 1991b. "Learning to Cooperate: Stochastic and Tacit Collusion in Social Exchange." *American Journal of Sociology* 97:808-43.

Majchrzak, Ann and Les Gasser. 1991. "On Using Artificial Intelligence to Integrate the Design of Organizational and Process Change in U.S. Manufacturing." *Artificial Intelligence and Society* 5:321-38.

Manthey, Michael. 1990. "Hierarchy and Emergence: A Computational View." Research Paper No. 0106-0791, R 90-25, Department of Mathematics and Computer Science, Institute for Electronic Systems, University of Aalborg, Denmark.

Masuch, Michael, ed. 1990. *Organization, Management and Expert Systems: Models of Automated Reasoning*. Berlin, New York: Walter de Gruyter.

Masuch, Michael and Perry LaPotin. 1989. "Beyond Garbage Cans: An AI Model of Organizational Choice." *Administrative Science Quarterly* 34:38-67.

McClelland, James and David Rumelhart. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 2. Cambridge, MA: MIT Press.

Minsky, Marvin Lee. 1988. *The Society of Mind*. New York: Simon and Schuster.

Moses, Yoram Ofer and M. Tennenholtz. 1990. "Artificial Social Systems. Part I, Basic Principles." Technical Report No. CS90-12, Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel.

———. 1991. "On Formal Aspects of Artificial Social Systems." Technical Report No. CS91-01, Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel.

Newell, Allen. 1990. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Oliveira, Pedro Paulo Balbi de. 1992. "Enact: An Artificial-Life World in a Family of Cellular Automata." Cognitive Science Research Paper No. CSRP-248, School of Cognitive and Computing Sciences, University of Sussex at Brighton, UK.

Pete, Andras, Krishna R. Pattipati, and David L. Kleinman. 1993. "Distributed Detection in Teams With Partial Information: A Normative Descriptive Model." *IEEE Transactions on Systems, Man and Cybernetics* 23:1626-48.

———. 1994. "Optimization of Detection Networks With Multiple Event Structures." *IEEE Transactions on Automatic Control* 39:1702-1707.

Prietula, Michael J. and Kathleen M. Carley. 1994. "Computational Organization Theory: Autonomous Agents and Emergent Behavior." *Journal of Organizational Computing* 41:41-83.

Roberts, Carl W. 1989. "Other Than Counting Words: A Linguistic Approach to Content Analysis." *Social Forces* 68:147-77.

———, ed. Forthcoming-a. *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences From Texts and Transcripts*. Mahweh, NJ: Lawrence Erlbaum.

———. Forthcoming-b. "Semantic Text Analysis: On the Structure of Linguistic Ambiguity in Ordinary Discourse." In *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences From Texts and Transcripts*, edited by C. W. Roberts. Mahweh, NJ: Lawrence Erlbaum.

Rosenbloom, Paul S., John E. Laird, Allen Newell, and R. McCarl. 1991. "A Preliminary Analysis of the Soar Architecture as a Basis for General Intelligence." *Artificial Intelligence* 47:289-325.

Rumelhart, David and James McClelland. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 1. Cambridge, MA: MIT Press.

Rutenbar, Rob A. 1989. "Simulated Annealing Algorithms: An Overview." *IEEE Circuits and Devices Magazine* 5:12-26.

Salancik, Gerald R. and Huseyin Leblebici. 1988. "Variety and Form in Organizing Transactions: A Generative Grammar of Organization." *Research in the Sociology of Organizations* 6:1-31.

Schank, Robert C. and Robert P. Abelson. 1977. *Scripts Plans and Goals and Understanding*. New York: John Wiley.

Shapiro, Gilbert. Forthcoming. "The Future of Coders: Human Judgments in a World of Sophisticated Software." In *Text Analysis for the Social Sciences: Methods for Drawing Statistical Inferences From Texts and Transcripts*, edited by C. W. Roberts. Mahweh, NJ: Lawrence Erlbaum.

Shoham, Yoav and M. Tennenholtz. 1994. "Co-Learning and the Evolution of Social Activity." Research Paper No. STAN-CS-TR-94-1511, Department of Computer Science, Stanford University, Stanford, CA.

Simon, Herbert. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69:99-118.

———. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63:129-38.

———. 1981a. *The Sciences of the Artificial*. 2d ed. Cambridge, MA: MIT Press.

———. 1981b. "Studying Human Intelligence by Creating Artificial Intelligence." *American Scientist* 69:300-309.

Skvoretz, John and Thomas J. Fararo. 1995. "The Evolution of Systems of Social Interaction." *Current Perspectives in Social Theory* 15:275-99.

Sloman, Aaron and Monica Croucher. 1981. "You Don't Need a Soft Skin to Have a Warm Heart: Towards a Computational Analysis of Motives and Emotions." Cognitive Science Research Paper No. CSRP-004, Cognitive Studies Programme, School of Social Sciences, University of Sussex at Brighton, UK.

Sowa, John F. 1984. *Conceptual Structures*. Reading, MA: Addison-Wesley.

Tang, Zhuong-Bo, Krishna R. Pattipati, and David L. Kleinman. 1993. "Optimization of Detection Networks: Part II—Generalized Tree Structures." *IEEE Transactions on Systems, Man and Cybernetics* 23:211-21.

Tesch, Renat. 1989. "Computer Software and Qualitative Analysis: A Reassessment." Pp. 141-54 in *New Technology in Sociology: Practical Applications in Research and Work*, edited by G. Blank, J. L. McCartney, and E. E. Brent. New Brunswick, NJ: Transaction.

Turner, Jonathan H. 1988. *A Theory of Social Interaction*. Stanford, CA: Stanford University Press.

Turner, Roy M. 1993. "The Tragedy of the Commons and Distributed AI Systems." Technical Report Series No. 93-01, Department of Computer Science, University of New Hampshire, Durham, NH.

Verhagen, Harko and Michael Masuch. 1994. "TASCCS: A Synthesis of Double-AISS and Plural-SOAR." Pp. 39-54 in *Computational Organization Theory*, edited by K. M. Carley and M. J. Prietula. Hillsdale, NJ: Lawrence Erlbaum.

Wasserman, Philip D. 1989. *Neural Computing: Theory and Practice*. New York: Van Nostrand Reinhold.

———. 1993. *Advanced Methods in Neural Computing*. New York: Van Nostrand.

White, Harrison C., Scott A. Boorman, and Ronald L. Breiger. 1976. "Social Structure From Multiple Networks: I. Blockmodels of Roles and Positions." *American Journal of Sociology* 81:730-80.

Winghart, Olivier. 1988. "Roles, Events and Saying Events in Expository Discourse." Research Paper No. AI88-85, Artificial Intelligence Laboratory, University of Texas at Austin.

Wolfe, Al. 1991. "Mind, Self, Society, and Computer: Artificial Intelligence and Sociology of Mind." *American Journal of Sociology* 96:1073-96.

Ye, Mei and Kathleen Carley. 1995. "Radar-Soar: Towards an Artificial Organization Composed of Intelligent Agents." *Journal of Mathematical Sociology* 20:219-46.

Zock, Michael and Gerard Sabah, eds. 1988. *Advances in Natural Language Generation: An Interdisciplinary Perspective*. London: Pinter.

Kathleen M. Carley is an associate professor of sociology and organizations at Carnegie Mellon University. She examines how learning, structure, and culture influence social and organizational behavior. She has developed methodological techniques for both textual and network analysis. Her current projects include examining the impact of organizational design, strategies, and member skills on organizational performance in a changing environment; the interplay of social structure, cognition, and technology in affecting information diffusion and group formation; mental models analysis; and statistical analysis of social networks and their evolution over time.