Artificial Social Intelligence

William Sims Bainbridge; Edward E. Brent; Kathleen M. Carley; David R. Heise;
Michael W. Macy; Barry Markovsky; John Skvoretz

Annual Review of Sociology, Vol. 20 (1994), 407-436.

ARTIFICIAL SOCIAL INTELLIGENCE [1]

William Sims Bainbridge
Sociology Program Director, National Science Foundation, 4201 Wilson Boulevard, Arlington, Virginia 22230

Edward E. Brent
University of Missouri, Columbia, Missouri 65201

Kathleen M. Carley
Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

David R. Heise
Indiana University, Bloomington, Indiana 47401

Michael W. Macy
Brandeis University, Waltham, Massachusetts 02154

Barry Markovsky
University of Iowa, Iowa City, Iowa 52242

John Skvoretz
University of South Carolina, Columbia, South Carolina 29208

Abstract

Sociologists have begun to explore the gains for theory and research that might be achieved by artificial intelligence technology: symbolic processors, expert

systems, neural networks, genetic algorithms, and classifier systems. The first major accomplishments of artificial social intelligence (ASI) have been in the realm of theory, where these techniques have inspired new theories as well as helping to render existing theories more rigorous. Two application areas for which ASI holds great promise are the sociological analysis of written texts and data retrieval from the forthcoming Global Information Infrastructure. ASI has already been applied to some kinds of statistical analysis, but how competitive it will be with more conventional techniques remains unclear. To take advantage of the opportunities offered by ASI, sociologists will have to become more computer literate and will have to reconsider the place of programming and computer science in the sociological curriculum. ASI may be a revolutionary approach with the potential to rescue sociology from the doldrums into which some observers believe it has fallen.

INTRODUCTION

Broadly defined, Artificial Social Intelligence (ASI) is the application of machine intelligence techniques to social phenomena. ASI includes both theory building and data analysis. At the theory building level, ASI includes the computer simulations of societies or groups of organizations in which the individuals are modelled as intelligent actors. At the analysis level, ASI includes the use of AI techniques for intelligently searching and analyzing data.

A decade has passed since the first conference on sociology and artificial intelligence was held (Gilbert & Heath 1985), yet sociologists have made relatively less use of AI than have practitioners of other behavioral sciences (Anderson 1989). Certainly, there is a faddish quality to much AI work, and much publicity has been given to computer programs that achieve only the pretence of intelligence, not its substance. Another factor may be that sociology has placed less emphasis on cognition in recent decades than on social structure, outside of areas like symbolic interactionism and the sociology of knowledge where use of computers is relatively undeveloped. The new interdisciplinary field of Cognitive Science is rooted in AI and draws upon five more traditional fields: psychology, philosophy, linguistics, anthropology, and neuroscience (Heckathorn 1989). For whatever reason, sociology is notably absent from this list.

Traditionally, computer scientists working on AI have ignored the social roots of human intelligence. In the past five years, however, there has been an upsurge of interest among these researchers in the social side in intelligence. Areas such as distributed artificial intelligence (Gasser 1991), coordination theory, and collaboration technology (all with strong roots in engineering or computer science) have begun to look at social issues. To be sure, many of the early programs allowed computers to engage in natural-language conver-

sations with humans, but the programming challenge was always to simulate the behavior of a single human actor. Although AI workers paid some attention to selected schools of thought within psychology, they ignored sociology. Randall Collins (1992) argues, however, that artificial intelligence cannot really be achieved without help from sociologists, and Allen Newell (1990) makes a similar point in his ground-breaking book, *Unified Theories of Cognition.*

We begin with a review of the chief technical approaches that have been developed in artificial intelligence, providing just enough description so the reader can see the possible relevance for sociology; we cite a very few recent references that can provide a deeper introduction. Then we consider the chief areas of theory and empirical research in which AI has relevance for the social sciences, citing sociological work when possible but also identifying accomplishments in neighboring social sciences that may foreshadow future sociological developments.

TECHNIQUES OF ARTIFICIAL INTELLIGENCE

At a first approximation, research in computer intelligence has taken one of two diametrically opposed approaches, which may, somewhat crudely, be called the "top-down" and "bottom-up" strategies. Until recently, most prominent AI researchers have focused on high-level symbolic processes that reflect the complex thought processes of which humans are capable. In contrast, others have attempted to model the basic functioning of nets of biological nerves, like those in relatively dumb simple organisms, with the hope that eventually they could work their way up to the level of human consciousness.

A major difference between the two approaches is in the use of words. Symbolic process models conceptualize the world using a predominantly verbal framework. Procedures (if-then rules) and words are used to describe all behavior and to take all actions. In contrast, "neural models" (and those of similar ilk) conceptualize the world in a predominantly numeric framework. Equations and numbers are used to describe all behavior and to take all actions.

Although some work on these "neural nets" goes back to the 1950s, this approach was eclipsed by the symbolic approach for almost 30 years (Crevier 1993). In part, this was due to work by Minsky & Papert (1969) that demonstrated the limitations of one type of numeric-networked system. Misinterpretation of their results, changes in scientific paradigms, and changes in funding sources all played a role in reducing interest in such systems. Starting around 1986, however, the neural network approach has achieved numerous successes and has grown in popularity as a scientific approach. It is sometimes referred to as "Connectionism" because it asserts that intelligence arises not in the

manipulation of symbols but in the connections between nerves, or between computer components.

We begin consideration of the chief AI techniques with symbolic process-ing—the top-down approach—then discuss its marriage with knowledge bases in what are often called expert systems. Neural networks—the bottom-up approach—comes next. We end this section with a discussion of genetic algorithms and classifier systems which are particular methods of great current interest, combining elements of the main approaches and drawing ideas from biology. Naturally, this very brief overview cannot do full justice to the complexity of this topic, and our aim is merely to provide a reasonably accurate picture of representative methods in each approach.

Symbolic Processors

Since the first conference on artificial intelligence, held at Dartmouth College in 1956, most AI workers have tended to define intelligence in terms of the manipulation of symbols and to write computer programs that could be called symbolic processors (Crevier 1993). While this approach may have limitations, it seemed the best way to study human language and problem-solving. In sociology, not only symbolic interactionists but also practically every re-searcher and theorist interested in norms, values, beliefs, institutions, and organizations gives great significance to shared meanings encoded and com-municated through symbols. Thus, symbolic processors seem the ideal AI technique for analyzing these phenomena, whether by means of abstract com-puter simulations or the empirical computerized analysis of language and meaningful behavior.

At one level, all computer programs operate symbolically. That is, both numbers and words are symbols. Thus all input, output and the computer program itself are expressed in terms of symbols. For example, a symbolic processor cannot actually look at a set of child's blocks and physically arrange them to form a tower. However, the AI researcher can tell the program where the blocks are, define the concept tower for it, and then ask the program to say how to move the blocks into a tower. At another level, a key attribute of symbolic processors is that they do not reason numerically, by counting and analyzing equations, but symbolically. They might measure an emotion in terms of high or low, for instance, rather than measuring it as a ratio variable. In this sense, these AI researchers are intellectual cousins of symbolic inter-actionists.

One way symbolic processors conceptualize intelligent behavior is in terms of problem spaces and production systems (Newell 1990). A problem space is a finite collection of states (situations, arrangements, etc) and operators (actions) that can be represented in the computer. These include the initial state and the desired state that is the goal. Operators can be used to move between

states, to create new states, or to eliminate old states. Intelligent activity is characterized as the search through the problem space for the goal-state, much as one might wend one's way through a maze. There are many different search procedures, some of which are exhaustive, and some of which create intermediate states on the fly. Problem spaces do not require that problems be well defined, but there must exist a goal state.

In a symbolic processing system, everything is done in terms of verbal productions (if-then rules), including description of states and search through the problem space. The set of productions in a system—or production system—is often referred to as its long-term memory. As it searches, the computer creates new productions which represent steps toward the goal. This search can occur in either a forward chaining fashion (from the initial state to the goal state) or a backward chaining fashion (from the goal state to the initial state). Symbolic processors vary in the conditions under which productions are applied, how extensive the match on the left-hand side of the if-then rule needs to be, whether these matches are allowed to be probabilistic, whether the rules can be applied in parallel, and so on.

In some symbolic processors, if there is not sufficient information to reach the goal state, or if some problem occurs during the search procedures, chunking can be used to overcome these "impasses." Chunking is a procedure that creates a new rule of the form, "if you are ever in situation X do Y." This rule is a reduced encapsulation of a chain of reasoning (following a path of many rules across many states and often through many problem spaces) that was done in order to figure out what to do in situation X. Imagine going through a maze, making many wrong turns, and keeping a map. Conceivably the map of a particular maze might be reduced to a single rule, "if you ever encounter this maze again turn left at every juncture." This is a chunk. The chunk itself does not contain all of the rules that were followed, all the missteps that were made, in the original search.

Search is one key element of symbolic systems. The other key component is knowledge representation. In symbolic systems knowledge is often represented using some type of network of linked concepts (Carley 1986). Minsky (1975) refers to the smallest connected unit, two concepts and the relation between them as a fact. A knowledge base can be thought of as a collection of facts or, equivalently, of rules. Sets of facts that are interrelated and serve to define a particular object or situation are often referred to as schemas or frames. State descriptions and problem space descriptions are often done by setting up and searching through particular schemas or frames. Symbolic processing systems vary widely in what search procedures and what knowledge representation schemes they use.

A crucial test for the symbolic processing approach to artificial intelligence is its capacity to handle human language, and one test of a computer pro-

grammer's skill is the ability to write parsers. A parser is a set of rules tied to a dictionary, perhaps explicitly framed as a production system, that is designed to extract meaning from samples of language. It is relatively easy to write a computer program that will respond correctly to keyboard-typed commands like "GO UP," "GO DOWN," "PRINT 'YES,'" or "ADD 2 PLUS 2." Every high-level computing language (BASIC, Pascal, C, FORTRAN, etc) incorporates a parser that translates human language into machine language commands. But the verbiage handled by most parsers is highly stylized, and the human must learn to stay within a fairly small set of linguistic conventions if the computer is to respond correctly.

In the 1960s, computer workers expressed great optimism that they would soon create automatic systems for translating between languages, for example, taking Russian input and producing grammatically correct English output with the same meaning. Manifest failure came in the form of rapidly proliferating sets of rules and the recognition that words may have fluid and multifaceted meanings (Kelley & Stone 1975). Despite these problems, natural language processing by computers has steadily improved, stimulated both by advances in programming and by vastly more powerful computer hardware.

A second challenge for symbolic processors was that if they were to be models of human cognition then they should solve problems as humans do and make the kinds of mistakes humans make. Research on language, as well as research on expert-knowledge, led to the conclusion that human reasoning often employed vast quantities of task specific information. This led some AI researchers to focus on the role of task knowledge in symbolic processing systems. This research grew into the area of expert systems and knowledge engineering. Many expert system programs have been written in LISP or PROLOG, languages developed for symbolic processing, and the rule structures are quite comparable (Clark 1982, Cameron & Dixon 1992).

Expert Systems

Within sociology, expert systems may potentially revolutionize qualitative sociology the way computerized statistical packages have revolutionized quantitative sociology (Brent 1986, Benfer et al 1991, Ohly 1993). Expert systems, for example, can be used to facilitate the coding of texts (Carley 1988), which may be as diverse as the words exchanged by members of a committee, the writings of a sociological theorist, or the public actions of international negotiators. Expert systems have found several applications in social welfare and human services (Schuerman et al 1989, Gingerich 1990, Mutschler 1990), assisting the professional in giving help, and the creation of their knowledge bases is practically equivalent to AI-assisted research on aspects of the profession and the social problem it addresses.

In the editor's statement at the back of every recent issue, the *International*

Journal of Expert Systems says its topic is "knowledge-based approaches to the construction of intelligent artifacts... A system is 'knowledge-based' when its behavior depends largely on information encoded in it or to which it has access, and is a[n] 'expert-system' when this knowledge would be considered expertise in a human." By these criteria, the spell-checker of a word processor could be an expert system. It has information that allows it to duplicate the expertise of a good human speller. However, one might want to reserve the term "expert system" for something a little smarter, that was able to respond in a complex way to different situations. While there is no clear line of demarcation, many would consider a good income tax package to be an expert system.

An expert system can be thought of as a symbolic processor designed to solve some particular problem in a fashion similar to that of a human expert in that area. Current expert systems typically have three parts: an inference engine, a knowledge base, and a user-interface (Gonzalez & Dankel 1993).

The inference engine is a symbolic processor that controls which rules are applied when and how. It contains general knowledge about how to conduct searches given the knowledge representation scheme being used within the expert system. The inference engine is often referred to as the shell. To create an expert system, researchers may purchase a shell and then add their task-specific knowledge base.

The knowledge base is a collection of task-specific (or domain) knowledge. In most current expert systems, it contains some parts of this knowledge in the form of rules and others in the form of auxiliary relational or object-oriented data.

Most current expert systems also have a separate user interface, the part of the package the user interacts with, which facilitates entering rules, generating reports, and managing a variety of ancillary tools. Producers of expert system shells have continually added enhancements in the user-interface area, including graphic ways of visualizing the structure of the knowledge, statistical analysis packages, and even the capacity for the system to learn the habits and priorities of the particular user. Greater use is being made of hypertext, the organization of textual material in a nonlinear manner, allowing the user to call up all kinds of information at any point in the process and to roam the knowledge base at will.

Creation of an expert system by means of a pre-existing shell involves the work of both a knowledge engineer and one or more domain experts. The knowledge engineer is trained in the use of the shell, has experience in eliciting information from other people (usually through interviews), and has experience in converting this information to a set of rules. The domain expert is a person who is thoroughly familiar with the field of knowledge the expert system is being developed for and may know nothing about computers. For example, a

knowledge engineer may be a computer scientist and a domain expert may be a neurosurgeon.

For example, to create an expert system intended to facilitate medical diagnosis in a particular category of diseases, the knowledge engineer will locate and interview a number of senior diagnosticians in that field, perhaps supplementing their interview responses with information from technical publications and other sources. After such a system has been created, it will be distributed to medical personnel who lack expert knowledge of the particular diseases. When they encounter cases they have difficulty diagnosing, they will turn to the expert system, which typically will ask a number of questions about the case, then suggest a diagnosis. Thus, the expert system acts as an expert consultant.

A chief challenge for expert systems is uncertainty. First, the domain of knowledge may be incomplete, probabilistic, problematic, or poorly organized. Second, the questions the user poses to the system may be sketchy and based on insufficient information about the case. These factors make it difficult to locate a deterministic solution. A well-designed system that lacks crucial information about a case will ask the user for it, but there are limits to how well additional information might resolve such ambiguities. Therefore, many expert system shells employ a variety of mathematical techniques to weight different information and provide estimates of its confidence in its conclusions, perhaps listing several possible conclusions with associated confidence scores. These weights can affect which rules are applied to the data in coming to a conclusion. Among these techniques are Bayesian probability measures, certainty factors, and fuzzy logic (Gonzalez & Dankel 1993:232–62).

In practice, commercial expert systems have run into a number of difficulties, chief among them the problem of finding competent human experts, the great cost of all the human labor required to create a worthwhile system, the difficulty in articulating and systematizing knowledge in many domains, and the high cost of updating expert systems as their domains of knowledge develop. In addition, there are social, organizational, and legal constraints on the acceptance and use of expert systems. Entrenched professionals can feel threatened by them. Weaver (1986) predicted that medical expert systems might have the effect of regulating physicians' behavior, undermining their authority and prestige, and leading to a new division of labor among medical professions. This has not yet happened. The public, who might benefit from the presence of expert systems, often feel distrustful of machines and would rather talk to human experts. Finally, there is the issue of who bears legal and moral responsibility when the computerized system makes a costly mistake. For these and other reasons, the term "expert system" is often avoided in commercial applications, in favor of intelligent decision support systems, intelligent advisory systems, consultant systems, and knowledge-based systems.

Shangraw (1987) has identified a number of ways in which the approaches of social scientists and knowledge engineers differ. Sociologists are often concerned with the bases for judgments, and they know that people's views are powerfully shaped by social, cultural, and economic factors. Consequently, sociologists see experts as having "biased" knowledge. Social scientists seek to maximize the validity of judgment. In contrast, knowledge engineers are concerned with duplicating the judgments of the experts rather than criticizing or explaining them. For knowledge engineers, any bias in judgments is important only if it results in too much inter-expert disagreement, and if this is the case, an expert system will not be built. If there is only a limited amount of disagreement, probabilistic weights on the various information will be used to reflect the level of inter-expert agreement. Knowledge engineers seek to replicate and systematize expert judgment. For the social scientist, the knowledge base itself is a source of data for further analysis.

Neural Networks

Artificial neural networks are computer systems involving software (and sometimes hardware) that in some sense emulates the behavior of biological nervous systems (Rumelhart & McClelland 1987, Wasserman 1989, 1993, Karayiannis & Venetsanopoulos 1993). After a very short period of recent development, neural nets have accomplished many useful tasks. For example, it has been demonstrated that a sufficiently large and well-trained neural net can accurately model any continuous mathematical function whatsoever (Smith 1993). Already, neural nets have been used to model economic data, competing successfully with more traditional techniques such as multiple regression, and the same approach can be applied directly to quantitative data collected by sociologists. As is true of symbolic processors, neural nets can model the behavior of intelligent actors in theory-based simulations, thus greatly expanding the tools available to sociological theorists.

Neural nets are radically different from other approaches to artificial intelligence and quite unlike the statistical software familiar to sociologists. Information is stored in the connections between nerve-like structures, in a distributed fashion, so that a particular datum is spread across a number of memory registers that it shares with other data, rather than assigning each datum to a distinct address in the computer's memory. Conventional statistical software reads its data in from a disk in essentially the same form it stores that information in a memory array. Neural networks, in contrast, learn information in training sessions, much as a human might learn, and they store information in forms that bear no resemblance to the raw data that were presented to them. Symbolic processors can also learn during training sessions; the form of the information stored in them may or may not resemble the form of the raw data.

A simple neural net might consist of fewer than a dozen nerve-like units, sometimes called neurons, nodes, or neurodes. There might be three layers of units. The first, or input layer, receives data. The third, or output layer, sends out the net's reactions to the data. Between them lies a layer of "hidden units" that are not directly connected to the outside world. Typically, every input unit is connected to every hidden unit, and every hidden unit is connected to every output unit. Associated with each connection is a connection strength or weight, a particular number that changes as the network learns to respond properly to a set of input data.

The net is trained by presenting it with a series of cases, each of which consists of an input vector and an output vector. The input vector is a set of numbers applied one each to the input units. Then the net produces an output, by using the connection weights to transform the input. The training algorithm compares the net's actual output with the desired output vector, and it follows a complex set of procedures to propagate the error back through the network (back propagation), adjusting all appropriate connection weights so that in future the net should respond with less error. In fact, computer scientists have experimented with many different network structures and training algorithms, but the basic idea is constant. A network of nerve-like units learns to respond in a desired way to a training set of data and then is ready to analyze fresh data of the same kind.

Neural nets are ideally suited for parallel processing, which is a significant revolution in computer technology. Traditional electronic computers consisted essentially of a bank of memory registers and a single central processing unit (CPU). The CPU would be in charge of everything, and all data and programming commands would pass through it. Thus, the speed of a computer was determined by the speed of this CPU, and everything depended upon the efficiency and reliability of a single electronic component. Parallel processing, as the name implies, employs a number of processors, perhaps many thousands of them, operating simultaneously. If problems of coordinating the actions of these processors can be solved, the result is much faster operation, perhaps by several orders of magnitude. Parallel techniques are also important to symbolic processors. In symbolic processors, being able to apply multiple rules in parallel also reduces operation time.

One justification for parallel distributed processing and neural nets is that the human brain must use something very similar. However, severe doubts have been raised whether existing neural nets properly simulate biological nervous systems, in particular whether biological nets make use of anything like back propagation (Hinton 1992). Another justification is simply speed. Doing things in parallel allows far more complex problems to be solved in reasonable time.

Genetic Algorithms and Classifier Systems

Genetic algorithms are an evolutionary approach to the solution of problems with a focus on a population of entities that become progressively better adapted to the problem space. Each entity is a string of symbols that can be read by a classifier system as a set of rules. As time goes by, strings that handle the problems better than others become more numerous in the population, share features of the solution with each other, and develop new features, in processes analogous to biological selection, reproduction, and mutation. While it is clear that this approach can perform some kinds of statistical analysis on empirical data, it has been imported so recently into the social sciences that technical and theoretical issues have been given the greatest attention.

Of special interest for ASI is that the inherent parallelism in the genetic based search procedure puts the intelligence of the procedure at the level of the population. In these systems, it is as though individuals independently and in parallel examine regions of the search space and can never "know" but a small piece of the overall solution, yet the system itself achieves effective search and so makes intelligent decisions. Real human societies possess greater information and capacity to process data than does any given individual member. The population of strings manipulated by a genetic algorithm is analogous to the gene pool of a human population, except that it is possibly cultural rather than biological in nature. Despite the nonsociological origin of its metaphors, thus, the genetic algorithm approach reinforces the sociological principle that real intelligence is in essence social.

The concept of a genetic algorithm arises in Holland's (1975, 1992) general treatise on adaptation in natural and artificial systems. As formulated by Holland (1975:20), a general adaptive problem has three components: (i) a system with an adaptive plan that determines successive changes in structure in response to an environment, (ii) the environment of this system, and (iii) a measure of the performance of different structures in the environment. In each application, the structures that undergo adaptation must be identified and represented, the mechanisms by which structures change must be specified, and the performance criterion affecting a structure's chances of persisting over time must be identified. The language of biology is useful to describe and label these three components. Structures are sets of chromosomes. The mechanisms that change structures are analogous to biological mechanisms of crossover, inversion (reproduction), and mutation, and the performance evaluation represents the fitness of the structure in the environment.

For convenience the process is viewed as a discrete time process, with the system following a trajectory described by the changing probability distribution over attainable structures. At each point in time, the existing structures

are evaluated against the environment in relation to performance criteria. The system's adaptive plan then takes these evaluations and the existing population of structures and produces the next generation's population of structures. Good adaptive plans increase the average performance of structures in the environment. Viewed as a search procedure, the genetic algorithm produces, over time, a concentration of structures in regions of the problem space that have relatively high fitness values.

The second fundamental idea, a classifier system, receives exposition in some recent work by Holland et al (1986), and is similar to a production system. Recall that a production is an if-then rule. In symbolic systems a production is spelled out using words. For Holland and his associates, a classifier is represented as two strings of numbers. Conditions and actions are abstractly represented by strings of symbols from the three letter alphabet "0", "1", and "#". In conditions, "#" is the called the "don't care," meaning that either a "0" or a "1" can occur in a message satisfying the condition; for instance, "0#" interpreted as a condition is satisfied by either of two messages "00" and "01". In actions "#" is interpreted as a place holder passing along to the output message the corresponding "0" or "1" in the input message, for instance, if "11" satisfies the (first) condition of a classifier having "0#" as its action part, then the classifier will post the output message "01".

A classifier K then has the form C1,C2,...,Cr/A where C1,...,Cr designate its r conditions and A designates its action. A classifier system consists of n classifiers denoted K1,K2,...,Kn, a message list, an input interface and an output interface. Execution proceeds by first placing all messages from the input interface on the message list. The message list is then processed by the classifiers for matches to their conditions. Classifiers whose conditions are matched then post their messages to a new message list which replaces the old list. The list is then processed through the output interface to produce the system's overall activity. Control then returns to the first step and execution continues.

Associated to each classifier in the system is a quantity called its strength. This quantity has three functions. First, among the classifiers whose conditions sets are matched and therefore offer competing alternative responses to the current situation, relative strength determines which ones win the competition to post their messages. Second, the strength of a classifier serves as a measure of its usefulness to the system because strength is adjusted based on system performance. The specific mechanism that adjusts strength is called the bucket brigade algorithm. Third, the application of the genetic algorithm to the generation of new classifiers uses classifier strength to choose "parent" rules for the next generation.

Because of the formal simplicity of their construction, classifier systems are natural candidates for evolution by means of a genetic algorithm. Strings

specifying conditions and/or actions can be split and recombined in the cross-over process to produce new and perhaps better classifiers. There are several possible specifications of the conditions under which the genetic algorithm is invoked. It could be invoked at the end of some well-defined performance cycle as in Goldberg's (1983) system to induce expert knowledge with respect to the regulation of gas-pipeline transmission. Holland et al (1986) suggest as general triggering conditions the failure of predictions and the occurrence of unusual events.

APPLICATIONS OF ASI TO THEORY

The growing number of theoretical essays grounded in computer simulations is one indication that social theorists are seeking ways to render their work more rigorous (Hanneman 1988). We suggest that the development of computer-based artificial social intelligence can have as great a positive impact on theory as did computerized statistical analysis on quantitative empirical research. Properly designed ASI programs can assess the logical consistency and completeness of theories, help discover new implications of old ideas, and connect scattered hypotheses into coherent theoretical systems. ASI may inspire altogether new theories, increase our appreciation of classical theories, and help improve and evaluate still-developing theories of social interaction and social structure.

ASI-Inspired Social Theories

Fararo & Skvoretz (1984, Skvoretz & Fararo 1989) have argued that the AI concept of production system can form the basis of a general theory of social institutions and action structures. Since a production is a rule in which a set of conditions demands a particular action, social norms are productions. Institutions and roles (distinctive sets of norms linked into cultural structures) are therefore production systems. Fararo & Skvoretz show how a social interaction is organized by the system of productions that defines the interrelationships of the roles being played, what they call the rolegram. Their work draws ideas from traditional writers, such as Talcott Parsons, and incorporates many conventional sociological concepts, but it places them in dynamic systems that owe much to ASI, even if they need not be realized on a computer.

   In similar fashion, Carley (1989, 1991b) has developed the constructural theory of social change predicated on individual adaptation at the cognitive level. Each individual has a knowledge base, and culture is the distribution of knowledge across individuals. Carley shows how social change, and processes such as diffusion, are controlled by the distribution of knowledge and by actions taken by individuals on the basis of this knowledge. Her work draws on traditional writers such as Homans and Durkheim and on the area of social

network analysis, but it places this work in a dynamic evolving system. Kaufer & Carley (1993) use the constructural theory to examine the social repercussions of print technology, arguing that communication technologies generate artifacts (such as books) that can be effectively modeled as artificial social agents with either more or less cognitive capabilities relative to humans. Using this approach they show the potential of print to change the nature of both the professions and academe.

Some have drawn lessons for theory from the apparent successes of computer intelligence, notably Slezak (1989) who argues that the "strong programme" in the sociology of science must be wrong because symbolic processor AI programs have successfully derived scientific and mathematical laws apparently without being influenced by sociocultural factors. Writing in the *American Journal of Sociology,* Wolfe (1991) draws the opposite lesson from his readings about AI, deciding that humans have a distinctive form of mind that cannot be duplicated by either symbolic processors or neural networks, a conclusion that would support interpretive rather than formal systematic schools of sociological analysis. Kontopoulos (1993) suggests that neural networks offer an appropriate metaphor for understanding social structure, thus incorporating insights from ASI into a general theory that need not be expressed in computational terms.

Expert System Models of Human Theorists

Rule-based expert systems are a very promising tool for theory formalization, and they may be used to analyze the thought of particular theorists. Some traditional sociologists, notably George Homans and Peter Blau, intended their theories to be what today we might call production systems, with each axiom or theorem represented by a production with linkages to others. But one way to study the thought of any social theorist would be to attempt to state his or her arguments in terms of productions. This difficult task might be facilitated by use of a flexible and full-featured expert system shell, taking the role of the knowledge engineer interrogating the writings of the theorist as if they were domain experts.

The Erving programs (Brent et al 1989) are an expert system that simulates Erving Goffman's dramaturgical perspective. Designed as a teaching tool, Erving takes the software's user into a "front-stage" bar, with associated "back-stage" party room and pool room, watching men and women interact and predicting their "impression management" behavior according to Goffman's principles. The user assembles various questions, piece by piece, such as: "How would Diane feel if Dave were to lie about age in the bar." "Would it be disruptive for Dave to make eye contact with members of opposite sex in the pool room?" The computer can test the user's understanding of Goffman's theory and offer explanations to the answers it gives for any question.

Banerjee (1986) wrote production systems in the PROLOG language to simulate sociopolitical theories of Skocpol and O'Donnell. The actors in each system are self-aware social groups with well-developed theories of the interests and possible coalitions in the worlds they inhabit. Skocpol's analysis of China in the decade following 1927, for example, posits the following seven actors: settled peasants, displaced peasants, Communist Party, Kuomingtang Party, gentry, coopted warlords, and independent warlords. In both cases, Banerjee found the predicted result, indicating that the two theories are logically constructed, and if any key assumption was removed, very different results emerged.

Simulations of Markets and the Iterated Prisoner's Dilemma

Computer simulation has a long history in the social sciences (Federico & Figliozzi 1981, Garson 1987), and mathematical models of human learning suitable for use in ASI programs were available decades ago (Bush & Mosteller 1955). However, most sociological computer simulations lacked explicit representations of human intelligence until recently. Today, many studies reported in central sociological journals employ computer models of human learning, decision-making, and social exchange, but they seldom mention any connection to artificial intelligence, even though retrospectively we can identify them as ASI. At their borders, math models and artificial intelligence blend into each other, and neural networks are an especially convenient way of embodying math models in computer programs (Wasserman 1993).

The complexity of social interaction has prompted the increasing use of computer simulation in place of formal mathematical models. The now-classic "prisoner's dilemma" computer tournament organized by Robert Axelrod may have been the turning point (Axelrod 1984). The prisoner's dilemma is a game-theoretic problem that explores the conditions under which cooperation may arise between self-interested economic actors who might gain in the short run if they violated agreements to cooperate (Rapoport & Chammah 1965). Axelrod invited people to submit computer programs that followed various strategies for playing repeated rounds of this game (the iterated prisoner's dilemma or IPD), and his tournament showed that simulation can produce robust yet unexpected results. The winner was one of the simplest contestants, a strategy of "tit-for-tat" that merely cooperated in the first move and thereafter imitated the previous action of its interactant. The simulation results showed how a simple norm of reciprocity could gain a toehold even in a harshly asocial environment and then go on to flourish. This strategy was able to displace much more cognitively sophisticated contestants. In short, the intelligence needed to find a way out of the social trap is not always isomorphic with the cognitive and analytic faculties of the organisms. A key contribution of ASI

is the recognition that problem solving can sometimes depend more on what goes on between organisms than on what goes on within them.

Subsequently, other researchers have held quite complex tournaments, such as one staged by Rust et al (1993) that was a double-auction market, like the one actually conducted in commodities and options by the Chicago Board of Trade. Its winning entries tended to be simple production systems with relatively little intelligence, but one competitor was a mammoth neural net with fully 1262 connection strength and bias parameters, that learned by means of genetic algorithms. For several years, the IPD strategy that drew the most attention from researchers and theorists was tit-for-tat, but a rival called "Pavlov" has recently seized center stage (Nowak & Sigmund 1993). This strategy has the individual actor continue to behave in a given way (keeping bargains or violating them) so long as it wins, and to shift behavior as soon as it loses. Because both tit-for-tat and Pavlov have the simulated person pay attention to what happened in the previous exchange, they connect directly to sociological theories of social learning.

In a series of theoretical papers based on ASI simulations, Macy (1990, 1991a,b) has developed stochastic learning models for the IPD that show how it is possible to "walk" out of social traps. The prisoner's dilemma is a trap, because the contingencies encourage people to act in ways that are not in accord with their own long-term interests. Real human life may be filled with social traps, in which decisions that make sense to each individual aggregate into outcomes that make no sense for all. In a typical run of this series, each individual's probability of cooperating depends upon past experience. The interacting population may be able to escape from a noncooperative equilibrium if a sequence of random events (the proverbial "drunkard's walk") brings them near enough to a cooperative equilibrium for them to settle on it. This research is a critique of rational choice theory, showing how learning theory can solve some of its problems. And like Axelrod's work it demonstrates that social actors may be able to escape the Hobbesian state of nature without the help of a king, a shared set of altruistic values, or even the degree of intelligence required to understand their situation fully.

An amazing variety of excellent work has been based in computer simulations of similar exchange systems, or on experiments with humans that might be further elucidated via simulations. Kollock (1993) has examined the effect of random errors and mistaken perceptions on the relative effectiveness of strategies like tit-for-tat. Orbell & Dawes (1991) explored the evolution of a cooperator's advantage when actors were allowed to withdraw from interaction. Frank (1988, 1993) also allowed exit from the game in his theoretical explorations of rational processes that led to the evolution of displays of emotion, thereby unleashing much simulation work by other researchers. Vila & Cohen (1993) modelled exchanges among individuals who could adopt

either of two strategies, producing wealth or expropriating it, thus exploring a theory of theft based upon earlier work in behavioral population biology. This last study suggests that ASI may have a considerable impact on criminology and the sociology of deviance.

Simulations of Networks, Groups, and Organizations

Many studies have examined social structures by means of ASI simulations, usually without explicitly acknowledging that artificial intelligence was involved. A good example for those who want to learn about ASI in connection with social networks is an article by Markovsky (1987, cf Markovsky 1995), because it focuses on the elemental social structure, the triad, and because it includes the actual source code of the program that was used. This study explored the power associated with position in a three-person social network where person A could interact with persons B and C, who however could not interact with each other. In a round of a typical experiment in the series, each person makes an offer of how 24 points could be divided between himself and another person. Person A compares the offers of B and C, selects one of them, and then the points are divided according to the average of the two persons' offers. After the first round (when the offers are randomly determined), each person adjusts his offer on the basis of what happened last time: if the previous offer was accepted, the new offer will be more demanding; if the previous offer was rejected, the new offer will be less demanding. On this simple basis, Markovsky was able to build a series of nineteen experiments that varied the strategies employed by individuals bargaining with each other. While rudimentary, the decision-making by actors, and their memories of the result of previous exchanges, constitute ASI. Similar work exploring the implications of structure in slightly larger networks has been undertaken by several researchers (Yamagishi et al 1988, Markovsky et al 1993).

Other examples cover a wide sociological territory. Feinberg & Johnson (1988, 1990) simulated the effect of an outside agitator on crowds, moving individuals physically toward the center of a mob and moving them mentally toward the agitator's preferred action. The individuals differed initially in terms of suggestibility and the propensity to move as well as in physical location and action choice. McPhail et al (1992, McPhail & Tucker 1990) modelled the physical movements of individuals as they threaded their way through crowds to reach a destination while remaining with each other in collective locomotion. Hummon (1990) simulated bureaucrats accepting, rejecting, and referring tasks on the basis of their growing experience with different kinds of work, thus creating the division of labor in a network. Anderson (1991) modelled social influence on voting behavior in small groups of union members. Bainbridge (1987, 1995) employed neural networks to simulate actors

with the intelligence to develop schemes for categorizing exchange partners and capable of learning which categories are most rewarding.

There is a long tradition in the use of simulation in organizational theory (Cyert & March 1963, Cohen et al 1972, Harrison & Carrol 1991). Recently, however, these simulations have included the use of ASI concepts (Masuch & LaPotin 1989, Carley et al 1992). As a field, organization theory has begun to recognize that theories of organizations need to take into account the complex interdependencies among people, machines, tasks, and society, the adaptability of both people and organizations, and the processes through which organizational behavior emerges (Masuch & Warglien 1992, Carley & Prietula 1994).

Traditional techniques for theory development are not sufficient for dealing with the difficulties inherent in theorizing about organizations as complex adaptive phenomena. Expert systems that can embody the theories and help the research to apply the theories systematically are making it possible to test cogently the full extent of complex theories (Baligh, Burton & Obel 1990, 1992). Computational models, particularly those in which the agents in the organization are modeled as complex, adaptive, and intelligent, are being used to examine emergent organizational behavior and to develop theory. Illustrative models include Masuch & LaPotin (1989), Cohen (1986), Carley (1991a, 1992), Harrison & Carrol (1991), Lin & Carley, (forthcoming). An area of related interest is distributed artificial intelligence Bond & Gasser (1988).

Recently, organizational models have been built out of collections of Soar agents. Soar is a production system language with built-in procedures for chunking (Laird et al 1987). It has been viewed as a unified theory of cognition (Newell 1990, Carley & Wendt 1991), because it embodies many of the findings on human cognition that have appeared in cognitive science. On numerous tasks it exhibits behavior similar to that of individual humans, and the recent organizational work asks whether in fact Soar can also act like humans in a social context. The first such study examined an organization of agents engaged in a warehouse task (Carley et al 1992). In this case, each Soar agent was run on a different machine, and the organization was formed out of a network of computers. If this vivid metaphor is not considered important, such simulations can be run on a supercomputer, simply allocating different sectors of memory to separate individuals.

Some traditional sociologists might complain that computer simulations inappropriately reduce social interaction to predictable, mechanistic cartoons that fail to capture the complexity and indeterminacy of human affairs. There are two refutations to this uninformed stereotype. First, even simple models often have great predictive power, and simulation results can and do often match empirical data both qualitatively (Carley 1991, Kaufer & Carley 1993) and quantitatively (Carley 1990, Carley & Lin, forthcoming).

Second, computational modeling can often generate surprising results.

Markovsky (1992) found that even very simple models of interaction across networks took the researcher beyond the limits of predictability. Recently, there has been great interest in the role of chance in several of the sciences, and the concept of deterministic chaos has been the subject of many publications of both scholarly and popular kinds (Mandelbrot 1983, Hao 1984, Gleick 1987). Kephart et al (1992) have noted that social behavior can become chaotic, and they ran simulations that showed how some intelligent strategies can reduce chaos, in particular giving actors the capacity to base their actions both on beliefs about the strategies of others and on the observed behavior of the system of agents.

Simulation work employing genetic algorithms has only just begun. Freeman (1993) used one to solve an old problem in social network analysis, namely, the partitioning of members of a group into cliques or subgroups based on members' proximities to one another. The algorithm processes assignments of individuals to subgroups searching for an assignment that maximizes a fitness function that is sensitive to misclassification of individuals and to average proximities. Axelrod (1987) has switched to genetic algorithms to evolve strategies in the Iterated Prisoner's Dilemma. Skvoretz & Fararo (1994) conceptualize social exchange in a game-theoretic context and apply a genetic algorithm to the evolution of mutual aid strategies. The game is similar to the prisoner's dilemma and the analysis is similar to Axelrod's. However, they introduce rudimentary role-differentiation into the problem, as strategies can function either as assistance requesters or providers of help, and they contrast the genetic-algorithm findings with those from a learning model implementation of the problem.

Data-Based Simulations

Conceptually intermediate between theory-driven simulations and quantitative analysis of data are simulations that might be described as data-based. The notable example, with a 20-year history, is the work on affect control theory begun by Heise (1986, 1987, Smith-Lovin 1987). Based on the EPA model of affective meaning developed by Charles Osgood, who employed semantic differential techniques to identify three dimensions of word meaning (Evaluation, Potency, Activity), affect control theory asserts that social events are constructed so as to confirm the meanings of social classifications. Research subjects have provided mean EPA ratings of thousands of words describing social identities, attributes, behaviors, and settings in five nations. Heise and his co-workers have derived a number of mathematical functions to predict how people would rate various combinations of words, and embodied both the formulas and the data in computer programs. Their theoretical agenda particularly stresses differences between social events that confirm or disconfirm sentiments attached to key nouns such as those describing standard social roles,

but their general method could be applied quite widely. Simulations based on data do not qualify as ASI unless they also incorporate a dynamic model of human thought, but because Heise's programs meet this test they give powerful testimony to the great potential of ASI to bring theory and data together in important new ways.

ASI APPLICATIONS TO RESEARCH

AI-assisted empirical sociological research is still in its infancy, so it is difficult not only to predict the range of applications it will have in future, but also to identify the current work that deserves closest attention. However, considerable progress has been achieved in two areas that clearly have great potential and nicely bracket the diversity of techniques we have described: qualitative analysis of verbal or written texts using expert systems or other varieties of symbolic processor, and enhancements to conventional statistical analysis such as substituting neural networks for multiple regression.

Qualitative Research with Symbolic Processors

In the 1960s, researchers found that very convincing interviewing programs could be created with surprisingly limited computing machinery and software. Especially controversial were programs that simulated a psychotherapist conversing with a patient (Colby et al 1966, Weizenbaum 1976). More recent work has suggested that fully computerized interviewing may have distinct advantages for some kinds of research, for example, sensitive topics like sexual behavior where respondents might be embarrassed to answer questions posed by another human being (Binik et al 1989).

Interview programs constructed along the lines of expert systems give sociologists an entirely fresh way of looking at data. Decades of quantitative research have been based on the concept of rectangular data matrices consisting of a large number of cases times a large number of variables, with the assumption that each case has a value (perhaps known or perhaps missing) for each variable. Relational data bases, such as incorporated in many expert systems, are very different in structure, as we have noted. Their topology may be very complex but generally consists of a network of nodes and relationships, with no matrix of cases by variables existing in the computer's memory. Carley has shown that such systems can be used to discover an individual's structure of meanings and then to compare that structure with the cognitive maps of other individuals (1986).

Possibly the greatest research potential for ASI in the coming decade is in computer-assisted analysis of written text. The federal government is increasing its already significant support for development of the National Information Infrastructure (Information Infrastructure Task Force 1993). Whatever the

exact form it takes, the "NII" will involve a tremendous expansion of computer communication networks and on-line databases, with rapid growth of the libraries of text available in electronic form. This includes everything from newspapers, to congressional debates, to (eventually) the entire contents of the Library of Congress. The question then becomes: What software tools will sociologists need if they are to navigate effectively through this ocean of words and analyze selected portions of text in the most effective manner?

Already a number of text-analysis software packages exist for microcomputers. Heise (1992) has shown that much can be accomplished with an ordinary word processor, and software packages like HyperResearch and Ethno have some of the qualities of expert systems, and thus they begin to enter the territory of ASI. "Intelligent" search procedures and modern knowledge representation schemes can help pre-process data (Franzosi 1990a,b) and recode data (Carley 1988) for general content or map analysis procedures. For some of these procedures the "intelligence" is built into the coding mechanism, in the form of "frames" that the researcher must fill in (Roberts 1989, Carley & Palmquist 1992, Carley 1993). These frames, which embody vast quantities of expertise are then used to postprocess and analyze the data.

Within narrative analysis, AI procedures can be used for examining, processing, and generating the story line in the narratives (Abell 1984, 1989). Ethno, for example, allows one to model event structures, including narratives, as production systems (Heise 1989, Griffin 1993). Related approaches are decision-based procedural analysis or protocol analysis, where the goal is to locate the explicit and implicit "rules" that the speaker uses to perform a task such as playing chess (Ericsson & Simon 1984). Gilbert & Heath (1986) have shown how PROLOG can be the basis of an intelligent system to capture public rules from narratives and retrieve the sense of textual items, illustrating the ways this would be done with medical records. Cope is a software system of nodes and linkages designed to produce cognitive maps of texts, thus helping develop grounded theory and capture verbal accounts (Cropper et al 1990). Automatic procedures such as Cirrus (VanLehn & Garlick 1987, Kowalski & VanLehn 1988) and ACM (Langley & Ohlsson 1984) have emerged, providing hope that larger numbers of texts can be analyzed quickly and economically.

Social scientists of politics have used expert system shells to analyze sequences of deeds and words in international relations (Schrodt 1988). For example, Mills (1990) created a rule-based expert system for analyzing negotiations and applied it to three sessions of talks between China and the Soviet Union. As the program runs, it asks the social scientist a set of questions about the behavior of each side at different points in the episode, then it outputs a summary analysis.

Despite the disillusionments of the 1960s, natural language processing has made substantial progress in translating texts and extracting meaning from

them, particularly, in the realm of story understanding (Abelson 1976, Rumelhart 1978, Schank & Riesbeck 1981, Lehnert & Ringle 1982). Plot-based procedural analysis lends itself to automation due to the presence of basic syntactic units (Lehnert 1981, Lehnert & Vine 1987) that make possible the automatic coding of texts.

Often the challenge is to trace linkages between texts, and a prime example is citation analysis in the sociology of science. Recent work in this area has employed search procedures from computer science to locate citation paths giving a history of which researcher cites which (Hummon & Doreian 1989, 1990, Hummon et al 1990, Hummon & Carley 1993, Carley et al 1993). These procedures make it possible not only to identify the main path of scientific development, but to understand the roles played by different types of research.

AI-Enhanced Statistical Analysis

Neural networks can readily substitute for multiple regression and for other multivariate techniques that aim to predict the value of one variable on the basis of the values of other variables. It is claimed that neural nets outperform other methods, chiefly because a sufficiently large neural net can in principle handle any pattern of nonlinearity in the relationships and complex interactions between independent variables. Indeed, for some problems neural nets may represent overkill, and if a net is given too many hidden units, it can overfit the data disastrously, producing an unreasonably complex and contorted curve on the scattergram.

Another disadvantage is that neural nets solve problems in ways that are far from transparent to human users, and they do not automatically generate lucid equations that can be comprehended in terms of explanatory theories. Perhaps this is why neural nets have been employed analytically in the social sciences primarily to make economic predictions (Lin & Lin 1993) when predictive power may be more important than explanatory intelligibility.

Neural nets do readily produce some conventional measures such as mean squared error (Smith 1993), and robust estimates of errors can easily be derived through procedures related to bootstrapping (Dietz et al 1987, 1991). The fact that neural nets have been doing useful statistical analysis for less than five years suggests that a full kit of related tools will require further time and effort. Perhaps one of the best ways to accomplish that is simply to attempt a variety of empirical studies and see what capabilities need to be added to those already possessed by neural nets.

Kimber (1991) compared neural networks with traditional methods and with ID3, a classification algorithm sometimes built into expert systems, to see which technique could best predict the emergence of democracy in nations, on the basis of such variables as urbanization, literacy, and economic resource distribution. Notably, the neural network performed better than did traditional

regression analysis. Similarly, neural nets did well in predicting outcomes of conflict between nations, in a study by Schrodt (1991). Huntley (1991) applied neural networks to analysis of time series data in order to forecast manpower needs in the Navy. Garson (1991) tested neural nets, ID3, and some more traditional techniques on sets of simulated data where the actual relationships between variables could be specified, finding that neural nets did a superior job with several kinds of problem.

Clearly, neural nets are not the only AI technique that may be useful for statistical analysis, and some of the procedures employed in symbolic processors could be applied to quantitative rather than qualitative data. While genetic algorithms have scored successes in econometric modelling (Koza 1992), their potential for analyzing social-scientific data remains largely unexplored.

A range of AI techniques can also be used in automatic security systems that prevent misuse of confidential data while enabling maximum legitimate use by social scientists (Keller-McNulty & Unger 1993). Government agencies and private corporations frequently corrupt datasets before releasing them, to satisfy anonymity and confidentiality regulations. Their chief concern is to prevent users from identifying the record of a particular person, and they do this by deleting cells from a table, truncating or collapsing values, adding random numbers to cells or values, and removing variables. The alternative is to embed the dataset in full-featured statistical software that employs encryption and analysis-monitoring techniques to prevent the user from inspecting the raw data directly and from deducing the identities of particular cases. An AI system could watch the user and prevent any sequence of statistical manipulations that could identify a single case, and they could also link separate datasets about the same people through their names without letting the user see them.

Expert systems and related tools such as hypertext can serve as methodological consultants, whether attached to familiar statistical packages or produced as stand-alone software. Statistical Navigator (Brent et al 1991) is a decision support system that helps a researcher select which techniques of statistical analysis are appropriate for an intended research project, and it can be used to give students an overview of statistical methods commonly used by social scientists. In its "consult mode," the software asks the researcher a series of questions about the aims and assumptions of the research project, and about the intended audience for its publications, using the answers to early questions to decide which other questions to ask. Many require the user to rate various possible objectives of the work on a 0 to 10 scale. After this, the software recommends a short list of methods to the user, rating them in terms of how well they serve the goals, assumptions, and audience for the particular piece of research. The hypertext feature allows the user to see a description of each method, linked to definitions of all technical terms, and the software

can also be run in "browse mode" where the user can roam this network of information at will. A detailed report of how the system arrived at its conclusions can be printed directly or saved onto disk for later editing.

CONCLUSION

Work on artificial intelligence has been subject to innumerable fads and frequently overblown publicity. Early proponents often made highly exaggerated claims for the computer programs they had written and hollow promises about what they would shortly accomplish. But these facts should not deter sociologists from examining the potential of artificial social intelligence, because steady progress in computer science has brought technology to the point where several valuable applications already exist and even modest extrapolations suggest that ASI could be of great significance to sociology.

Sociologists interested in exploring ASI will ask themselves how much they need to learn about computing, as opposed to relying upon computer scientist collaborators for expertise in that field. We think it is essential to be able to program in at least one high-level computing language, preferably in two very different ones such as Pascal and PROLOG, and to inform oneself about a range of recent technical developments. Particularly in the field of artificial intelligence, it is difficult to understand the meaning of the techniques unless one is capable of programming at least some of them from scratch. It is true that ready-made AI software exists that can be used for ASI, notably expert system shells and neural network packages for statistical analysis, but for the foreseeable future most ASI applications in sociology require writing a considerable amount of fresh code. And the fact is that sociologists and computer scientists have great difficulty communicating with each other, neither one generally appreciating the other's assumptions or understanding the simplest things the other says. We are convinced that collaborations between sociologists and computer scientists will be highly fruitful, but the sociologist will have to go the extra mile and enter the world of computers if such projects are to have any chance of success.

Current graduate training does not prepare students to take advantage of ASI. Although some probability theory can be useful, hardly any of the material taught in statistics courses is relevant to the computer techniques described here. Two or three decades ago, there was much debate about requiring students to learn computer programming, even the interesting idea of counting high-level computing languages toward the then-existing foreign language requirements. Of course, today few sociologists write programs from scratch, because they cannot compete with the elaborate and accurate statistical packages available on the market at low cost. But now, we suggest, ASI presents an array of new reasons to become competent in programming, not the least of which

is our belief that ASI skills enhance a person's capacity to think both logically and creatively.

While recognizing the danger of being swept away by excessive enthusiasm, in a field that has been rife with fads, we believe that ASI opens an entirely fresh era for social theory. Sociologists are concerned with the behavior of societies of complex, adaptive individuals. Therefore, rigorous theory building and testing requires the use of mathematics or computer simulations. Even as physics and economics are developing computational components, so should sociology.

The general public, to the extent that it has any opinion about social theory at all, probably considers it to be mere ideology. So long as theories are rambling verbal meditations punctuated with dubious metaphors, there is little defense against this accusation. ASI and mathematical formalism are compatible methods for stating a theory precisely, connecting its concepts in rigorous intellectual structures, and identifying both hidden assumptions and unexpected consequences. Skillfully written simulation programs can be an excellent medium for communication of precise theoretical statements, so long as the intended audience has learned how to read programs.

Whatever the future of various government initiatives in high performance computing, a Global Information Infrastructure is gradually emerging. Internet links the National Science Foundation with literally millions of computer users already, and many archives and libraries are steadily increasing the amount of text and other data available over it. Alternately, the medium for distribution of data may be laser discs (CD ROM) or a successor technology that carries information in physical objects rather than in electronic bit streams. Social scientists need to be involved in the development of these technologies, in part because we can be sophisticated advocates for the general public whose real data and software needs might otherwise have little influence on systems created by engineers and bureaucrats. As the Global Information Infrastructure develops, both social scientists and members of the general public will need ASI agents, specialized computer programs that can be sent into this practically infinite universe of data in search of desired information.

Sociologists have become adept with a wide variety of statistical tools, and one would have thought that our quantitative methodology is thoroughly mature at this point. Thus it is surprising to see neural networks suddenly competing with multiple regression and other well-established methods. Particularly for analysis of texts, ASI techniques may prove superior to other approaches, and it is possible that artificial intelligence will play a prominent role in management and analysis of quantitative datasets, as well.

There is much talk these days about the malaise into which sociology supposedly has fallen. Whether or not this is true, ASI has the potential to revolutionize sociology. To be sure, if unprepared sociologists rushed to ASI

in search of salvation, they would undoubtedly be disappointed, and the sub-stantial real gains that ASI promises might be lost in their disillusionment. Prudent but creative incorporation of ASI methods into sociology could rein-vigorate stagnant subdisciplines, open new fields for exploration, keep sociol-ogists competitive in the study of traditional sociological questions, and pre-vent our discipline from falling behind other social and behavioral sciences that have more enthusiastically exploited the tremendous possibilities offered by computer intelligence.

Literature Cited

Abell P. 1984. Comparative narratives: Some rules for the study of action. *J. Theory Soc. Behav.* 14:309–32

Abell P. 1989. The theory and method of com-parative narratives. *Connections* 12:51–79

Abelson RP. 1976. Script processing in attitude formation and decision-making. In *Cognition and Social Behavior,* ed. JS Carroll, JW Payne. Hillsdale, NJ: Lawrence Erlbaum

Anderson B. 1989 On artificial intelligence and theory construction in sociology. *J. Math. Sociol.* 14:209–16

Anderson JG. 1991. Simulation models of small group behavior. In *Simulation in Health Care,* ed. JG Anderson, pp. 235–242. San Diego, Calif: Soc. Computer Simulation

Axelrod R. 1984. *The Evolution of Coopera-tion.* New York: Basic

Axelrod R. 1987. The evolution of strategies in the iterated prisoner's dilemma. In *Genetic Algorithms and Simulated Annealing,* ed. L Davis, pp. 32–41. London: Pitman

Bainbridge WS. 1987. *Sociology Laboratory.* Belmont, Calif: Wadsworth

Bainbridge WS. 1994. Grand computing challenges for sociology. *Soc. Sci. Computer Rev.* In press

Bainbridge WS. 1995. Neural network models of religious belief. *Sociol. Perspect.* In press

Baligh HH, Burton RM, Obel B. 1990. Devis-ing expert systems in organization theory: The organizational consultant. In *Organiza-tion, Management, and Expert Systems,* ed. M Masuch. Berlin: Walter De Gruyter

Banerjee S. 1986. Reproduction of social struc-tures: An artificial intelligence model. *J. Conflict Resolution* 30:221–52

Benfer R, Brent EE, Furbee L. 1991. *Expert Systems.* Newbury Park, Calif: Sage

Binik YM, Westbury CF, Servan-Schreiber, D. 1989. Interaction with a 'sex-expert' system enhances attitudes toward computerized sex therapy. *Behav. Res. Therapy* 27:303–306

Bond A, Gasser L. 1988. *Readings in Distrib-uted Artificial Intelligence.* San Mateo, Calif: Kaufmann

Brent EE. 1986. Knowledge-based systems: A qualitative formalism. *Qual. Sociol.* 9:256–82

Brent EE, Glazier J, Jamtgaard K, Wetzel E, Hall PM, Dalecki M, Bah A. 1989. Erving: A program to teach sociological reasoning from the dramaturgical perspective. *Teach-ing Sociol.* 17:38–48

Brent EE, Mirielli EJ, Detring E, Ramos F. 1991. *Statistical Navigator Professional: User's Guide and Reference Manual.* Co-lumbia, Mo: Idea Works

Burton RM, Obel B. 1984. *Designing Efficient Organizations: Modeling and Experimenta-tion.* Amsterdam, The Netherlands: Elsevier Sci.

Bush RR, Mosteller F. 1955. *Stochastic Models for Learning.* New York: Wiley

Cameron RD, Dixon AH. 1992. *Symbolic Com-puting with LISP.* Englewood Cliffs, NJ: Prentice-Hall

Carley KM 1986. An approach for relating so-cial structure to cognitive structure. *J. Math. Sociol.* 12:137–89

Carley KM. 1988. Formalizing the social expert's knowledge. *Sociol. Meth. Res.* 17: 165–232

Carley KM. 1989. The value of cognitive foun-dations for dynamic social theory. *J. Math. Sociol.* 14:171–208

Carley KM. 1990. Group stability: A socio-cog-nitive approach. In *Advances in Group Pro-*

cesses, ed. E Lawler, B Markovsky, C Ridgeway, H Walker, pp. 7:1–44. Greenwich, Conn: JAI

Carley KM. 1991a. Designing organizational structures to cope with communication breakdowns. *Industrial Crisis Q.* 5(1):19–57

Carley KM. 1991b. A theory of group stability. *Am. Sociol. Rev.* 56:331–54

Carley KM. 1992. Organizational learning and personnel turnover. *Organ. Sci.* 3(1):2–46

Carley KM. 1993. Coding choices for textual analysis: a comparison of content analysis and map analysis. In *Sociological Methodology*, ed. P. Marsden. Oxford: Blackwell

Carley KM, Hummon N, Harty M. 1993. Scientific influence: An analysis of the main path structure. In *J. Conflict Resolution, Knowledge: Creation, Diffusion, Utilization* 14:417–47

Carley KM, Kjaer-Hansen J, Newell A, Prietula M. 1992. Plural-Soar: A prolegomenon to artificial agents and organizational behavior. In *Artificial Intelligence in Organization and Management Theory*, ed. M Masuch, M Warglien, pp. 87–118. Amsterdam: North-Holland

Carley KM, Lin, Z. 1994. Organizational designs suited to high performance under stress. *IEEE Systems, Man and Cybernetics.* In preparation

Carley KM, Palmquist M. 1992. Extracting, representing and analyzing mental models. *Soc. Forc.* 70:601–36

Carley KM, Prietula M, eds. 1994. *Computational Organization Theory*. Hillsdale, NJ: Lawrence Erlbaum

Carley KM, Wendt K. 1991. Electronic mail and scientific communication: A study of the Soar extended research group. In *J. Conflict Resolution. Knowledge: Creation, Diffusion, Utilization* 12:406–40

Clark KL. 1982. An introduction to logic programming. In *Introductory Readings in Expert Systems,*, ed. D Michie, pp. 93–112. New York: Gordon & Breach

Cohen MD. 1986. Artificial intelligence and the dynamic performance of organizational designs. In *Ambiguity and Command: Organizational Perspectives on Military Decision Making*, ed. JG March, R Weissinger-Baylon. Marshfield, Mass: Pitman

Cohen MD, March JB, Olsen JP. 1972. A garbage can model of organizational choice. *Admin. Sci. Q.* 17(1):1–25

Colby KM, Watt JB, Gilbert JP. 1966. A computer method of psychotherapy. *J. Nervous Mental Dis.* 142:148–52

Collins R. 1992. Can sociology create an artificial intelligence? In *Sociological Insight*, pp. 155–184. New York: Oxford Univ. Press

Crevier D. 1993. *AI: The Tumultuous History of the Search for Artificial Intelligence.* New York: Basic

Cropper S, Eden C, Ackermann F. 1990. Keeping sense of accounts using computer-based cognitive maps. *Soc. Sci. Computer Rev.* 8:345–66

Cyert R, March JG. 1963. *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice-Hall

Dietz T, Frey RS, Kalof L. 1987. Estimation with cross-national data: Robust and nonparametric methods *Am. Sociol. Rev.* 52:380–90

Dietz T, Kalof L, Frey RS. 1991. On the utility of robust and resampling procedures. *Rural Sociol.* 56:461–74

Ericsson KA, Simon HA. 1984. *Protocol Analysis: Verbal Reports as Data.* Cambridge, Mass: MIT Press

Fararo TJ, Skvoretz J. 1984. Institutions as production systems. *J. Math. Sociol.* 10: 117–82

Federico P, Figliozzi PW. 1981. Computer simulation of social systems. *Sociol. Meth. Res.* 9:513-33

Feinberg WE, Johnson NR. 1988. 'Outside agitators' and crowds: Results from a computer simulation model. *Soc. Forc.* 67:398–423

Feinberg WE, Johnson NR. 1990. Radical leaders, moderate followers: Effects of alternative strategies on achieving consensus for action in simulated crowds. *J. Math. Sociol.* 15:91–115

Frank, Robert H. 1988. *Passions within Reason.* New York: Norton

Frank, Robert H. 1993. The strategic role of the emotions. *Rationality Soc.* 5:160–84

Franzosi R. 1990a. From words to numbers: a generalized and linguistic-based coding procedure for collecting textual data. In *Sociological Methodology*, ed. CC Clogg, pp. 263–98. San Francisco: Jossey-Bass

Franzosi R. 1990b. Strategies for the prevention, detection, and correction of measurement error in data collected from textual sources. *Sociol. Meth. Res.* 18:442–72

Freeman LC. 1993. Finding groups with a simple genetic algorithm. *J. Math. Sociol.* 17:227–41

Garson GD. 1987. Computer simulation in social science. In *Academic Microcomputing*, pp. 110–138. Newbury Park, Calif: Sage

Garson GD. 1991. A comparison of neural network and expert systems algorithms with common multivariate procedures for analysis of social science data. *Soc. Sci. Computer Rev.* 9:399–434

Gasser L. 1991. Social conceptions of knowledge and action: DAI foundations and open systems semantics. *Artificial Intell.* 47:107–38

Gilbert GN, Heath C, eds. 1985. *Social Action and Artificial Intelligence.* Aldershot, England: Gower

Gilbert GN, Heath C. 1986. Text, competence

and logic: An exercise. *Qual. Sociol.* 9:215–36

Gingerich WJ. 1990. Expert systems: New tools for professional decision-making. *Computers in Hum. Serv.* 6:219–230

Gleick J. 1987. *Chaos.* New York: Penguin

Goldberg D. 1983. *Computer-aided gas pipeline operation using genetic algorithms and rule learning.* PhD thesis. Univ. Mich.

Gonzalez AJ, Dankel DD. 1993. *The Engineering of Knowledge-Based Systems.* Englewood Cliffs, NJ: Prentice-Hall

Griffin LJ. 1993. Narrative, event-structure analysis, and causal interpretation in historical sociology. *Am. J. Sociol.* 98:1094–1133

Hanneman R. 1988. *Computer Assisted Theory Building.* Newbury Park, Calif: Sage

Hao B, ed. [Or Bai-Lin, Hao] 1984. *Chaos.* Singapore: World Sci.

Harrison JR, Carrol GR. 1991. Keeping the faith: a model of cultural transmission in formal organizations. *Admin. Sci. Q.* 36:552–82

Heckathorn DD. 1989. Cognitive science, sociology, and the theoretic analysis of complex systems, *J. Math. Sociol.* 14:97–110

Heise DR. 1986. Modeling symbolic interaction. *In Approaches to Social Theory,* ed. S Lindenberg, JS Coleman, S Nowak, pp. 291–316. New York: Russell Sage Found.

Heise DR. 1987. Affect control theory: Concepts and models. *J. Math. Sociol.* 13:1–33

Heise DR. 1989. Modeling event structures. *J. Math. Sociol.* 14:139–69

Heise DR. 1992. Computer assistance in qualitative sociology. *Soc. Sci. Computer Rev.* 10:531–43

Hinton GE. 1992. How neural networks learn from experience. *Sci. Am.* 267(3):145–151

Holland JH. 1975. *Adaptation in Natural and Artificial Systems.* Ann Arbor, Mich: Univ. Mich. Press

Holland JH. 1992. Genetic algorithms. *Sci. Am.* 267(July):66–72

Holland JH, Holyoak KJ, Nisbett RE, Thagard PR. 1986. *Induction: Processes of Inference, Learning and Discovery.* Cambridge, Mass: MIT Press

Hummon NP. 1990. Organizational structures and network processes. *J. Math. Sociol.* 15:149–61

Hummon NP, Carley K. 1993. Social networks as normal science. *Soc. Networks* 15:71–106

Hummon NP, Doreian P. 1989. Connectivity in a citation network: the development of DNA theory. *Soc. Networks* 11:39–63

Hummon NP, Doreian P. 1990. Computational methods for social network analysis. *Soc. Networks* 12:273–88

Hummon NP, Doreian P, Freeman L. 1990. Analyzing the structure of the centrality literature between 1948 and 1979. *Knowledge: Creation, Diffusion, Utilization* 11: 459–79

Huntley DG. 1991. Neural nets: An approach to the forecasting of time series. *Soc. Sci. Computer Rev.* 9:27–38

Karayiannis NB, Venetsanopoulos AN. 1993. *Artificial Neural Networks: Learning Algorithms, Performance Evaluation, and Applications.* Boston: Kluwer

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

Keller-McNulty S, Unger EA. 1993. Database systems: Inferential security. *J. Official Statist.* 9:475–99

Kelly EF, Stone PJ. 1975. *Computer Recognition of English Word Senses.* Amsterdam: North Holland

Kephart JO, Huberman BA, Hogg T. 1992. Can predictive agents prevent chaos? In *Economics and Cognitive Science,* ed. P Bourgine, B Walliser, pp. 41–55. Oxford: Pergamon

Kimber R. 1991. Artificial intelligence and the study of democracy. *Soc. Sci. Computer Rev.* 9:381–98

Kollock P. 1993. "An eye for an eye leaves everyone blind:" Cooperation and accounting systems. *Am. Sociol. Rev.* 58

Kontopoulos KM. 1993. Neural networks as a model of structure. In *The Logics of Social Structure,* pp. 243–267. New York: Cambridge Univ. Press

Kowalski B, VanLehn K. 1988. Cirrus: inducing subject models from protocol data. In *Proc. Tenth Ann. Conf. Cognitive Sci. Soc.,* ed. V. Patel. Hillsdale, NJ: Erlbaum

Koza JR. 1992. A genetic approach to econometric modelling. In *Economics and Cognitive Science,* ed. P Bourgine, B Walliser, pp. 57–75. Oxford: Pergamon

Laird J, Newell A, Rosenbloom P. 1987. Soar: an architecture for general intelligence. *Artificial Intell.* 33:1–64

Langley P, Ohlsson S. 19884. Automated cognitive modeling. In *Proc. Fourth Natl. Conf. Artificial Intelligence.* Austin, Tex: Morgan-Kaufman

Lehnert W. 1981. Plot units and narrative summarization. *Cognitive Sci.* 5:293–331

Lehnert W, Ringle M, eds. 1982. *Strategies for Natural Language Processing.* Hillsdale, NJ: Lawrence Erlbaum

Lehnert W, Vine E. 1987. The role of affect in narrative structure. *Cognition Emotion* 1: 299–322

Lin M, Lin FC. 1993. Analysis of financial data using neural nets. In *Neural Network Special Report,* pp. 33–37. San Francisco: Miller Freeman

Lin Z, Carley, KM. 1994. Proactive or reactive: An analysis of the effect of agent style on organizational decision making performance. *Int. J. Intelligent Systems in Accounting, Finance, Manage.* In press

Macy MW. 1990. Learning theory and the logic of critical mass. *Am. Sociol. Rev.* 55:809–26

Macy MW. 1991a. Chains of cooperation: Threshold effects in collective action. *Am. Sociol. Rev.* 56:730–47

Macy MW. 1991b. Learning to cooperate: Stochastic and tacit collusion in social exchange. *Am. J. Sociol.* 97:808–43

Mandelbrot BB. 1983. *The Fractal Geometry of Nature.* San Francisco: WH Freeman

Markovsky B. 1987. Toward multilevel sociological theories: Simulations of actors and network effects, *Sociol. Theory* 5:101–17

Markovsky B. 1992. Network exchange outcomes: Limits of predictability, *Soc. Networks* 14:267–86

Markovsky B. 1995. Developing an exchange network simulator, *Sociol. Perspect.* In press

Markovsky B, Skvoretz J, Willer D, Lovaglia MJ, Erger J. 1993. The seeds of weak power: An extension of network exchange theory. *Am. Sociol. Rev.* 58:197–209

Masuch M, LaPotin P. 1989. Beyond garbage cans: An AI model of organizational choice. *Admin. Sci. Q.* 34:38–67

Masuch M, Warglien M, eds. 1992. *Artificial Intelligence in Organization and Management Theory.* Amsterdam, The Netherlands: Elsevier Sci.

McPhail C, Powers WT, Tucker CW. 1992. Simulating individual and collective action in temporary gatherings, *Soc. Sci. Computer Rev.* 10:1–28

McPhail C, Tucker CW. 1990. Purposive collective action. *Am. Behav. Scientist* 34:81–94

Mills, W. deB. 1990. Rule-based analysis of Sino-soviet negotiations. *Soc. Sci. Computer Rev.* 8:181–95

Minsky M. 1975. A framework for representing knowledge. In *The Psychology of Computer Vision,* ed. P Winston. New York: McGraw-Hill

Minsky M, Papert S. 1969. *Perceptrons.* Cambridge, Mass: MIT Press

Mutschler E. 1990. Computer assisted decision making. *Computers in Hum. Serv.* 6:231–50

Newell A. 1990. *Unified Theories of Cognition.* Cambridge, Mass: Harvard Univ. Press

Nowak M, Sigmund K. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* 365:56–58

Ohly HP. 1993. Knowledge-based systems: Another data approach for social scientists? *Soc. Sci. Computer Rev.* 11:84–94

Orbell J, Dawes R. 1991. A "cognitive miser" theory of coooperators' advantage. *Am. Polit. Sci. Rev.* 85:515–28

Rapoport A, Chammah AM. 1965. *Prisoner's Dilemma: A Study in Conflict and Cooperation.* Ann Arbor: Univ. Mich. Press

Rumelhart DE. 1978. Understanding and summarizing brief stories. In *Basic Processes in Reading,* ed. D LaBerge, SJ Samuels. Hillsdale, NJ: Lawrence Erlbaum

Rumelhart DE, McClelland JL. 1987. *Parallel Distributed Processing.* Cambridge: MIT Press

Rust J, Palmer R, Miller JH. 1993. Behavior of trading automata in a computerized double auction market. In *The Double Auction Market: Institutions, Theories, and Evidence,* ed. J Friedman, J Rust. New York: Addison Wesley

Schank RC, Riesbeck CK. Eds. 1981. *Inside Computer Understanding.* Hillsdale, NJ: Lawrence Erlbaum

Schrodt, PA. 1988. Artificial intelligence and formal models of international behavior. *Am. Sociologist* 19:71–85

Schrodt PA. 1991. Prediction of interstate conflict outcomes using a neural network. *Soc. Sci. Computer Rev.* 9:359–80

Schuerman JR, Mullen E, Stagner M, Johnson P. 1989. First generation expert systems in social welfare. *Computers in Hum. Serv.* 4:111–22

Shangraw RF, Jr. 1987. Knowledge acquisition, expert systems, and public management decisions. *Soc. Sci. Microcomputer Rev.* 5:163–73

Skvoretz J, Fararo TJ. 1989. Action structures and sociological action theory. *J. Math. Sociol.* 14:111–37

Skvoretz J, Fararo TJ. 1994. The evolution of systems of social interaction. In *Current Perspectives in Social Theory,* Vol. 14, ed. B. Agger. Greenwich, Conn: JAI. In press

Slezak P. 1989. Scientific discovery by computer as empirical refutation of the strong programme. *Soc. Stud. Sci.* 19:563–600; Responses and replies 19:601-69

Smith M. 1993. *Neural Networks for Statistical Modeling.* New York: Van Nostrand

Smith-Lovin L. 1987. Affect control theory: An assessment. *J. Math. Sociol.* 13:171–92

VanLehn K, Garlick S. 1987. Cirrus: an automated protocol analysis tool. In *Proc. Fourth Int. Wkshop. Machine Learning,* ed. P. Langley. Los Altos, Calif: Morgan-Kaufman

Vila BJ, Cohen LE. 1993. Crime as strategy: testing an evolutionary ecological theory of expropriative crime. *Am. J. Sociol.* 98:873–912

Wasserman PD. 1989. *Neural Computing: Theory and Practice.* New York: Van Nostrand Reinhold

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

Weaver RR. 1986. Some implications of the emergence and diffusion of medical expert systems. *Qual. Sociol.* 9:237–55

Weizenbaum J. 1976. *Computer Power and Human Reason.* San Francisco: WH Freeman

Wolfe A. 1991. Mind, self, society, and computer: Artificial intelligence and the sociology of mind. *Am. J. Sociol.* 96:1073–96

Yamagishi T, Gillmore MR, Cook KS. 1988. Network connections and the distribution of power in exchange networks. *Am. J. Sociol.* 93:833–51