

CHAPTER 4

Computational Approaches to Sociological Theorizing

KATHLEEN M. CARLEY

INTRODUCTION

Sociologists are concerned with explaining and predicting social behavior. Theorists are not concerned with providing a description of society at a particular point in time. Rather, they are in the business of providing insight into how societies change, into the dynamics underlying behavior, into the processes that result in the observed correlations, and into the way in which multiple factors come together to create specific social situation. In other words, the sociological theorizing involves explaining how multiple factors interact in complex, often nonlinear ways to affect social behaviors and in explaining the dynamics by which social agents, groups, teams, organizations, societies, cultures evolve and coevolve. The focus of theory on complex nonlinear and dynamic systems makes computational analysis a natural methodological choice for theorizing.

The use of formal techniques¹ in general and computational analysis in particular is playing an ever-increasingly important role in the development of sociological theory. Computational analysis has been used to theorize about a large number of social behaviors: organizational exploration and exploitation (March, 1996), cooperation (Macy, 1991a,b), coordination (Carley & Prietula, 1994), diffusion and social evolution (Carley, 1991; Kaufer & Carley, 1993), organizational adaptation (Carley & Svoboda, 1996; Levinthal, 1997; Sastry, 1997), change in social networks (Markovsky, 1987) and exchange networks (Yamagishi et al., 1988; Markovsky et al., 1993), collective action (Feinberg & Johnson, 1988, 1990; McPhail & Tucker, 1990), and the fundamental nature of the social agent (Carley & Newell, 1994). One reason for this movement to computational social theory is the growing recognition that social

¹Formal techniques include logic modeling (e.g., Skvoretz & Fararo, 1989, 1994), mathematical modeling (Lave & March), and simulation (Bainbridge et al., 1994; Carley & Prietula, 1994; McKelvey, 1997). In this chapter the focus is on simulation; however, the reader should recognize that these other formal techniques play complementary roles in sociological theorizing.

KATHLEEN M. CARLEY • Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213.

Handbook of Sociological Theory, edited by Jonathan H. Turner. Kluwer Academic / Plenum Publishers, New York, 2002.

processes are complex, dynamic, adaptive, and nonlinear; that social behavior emerges from interactions within and between ecologies of entities (people, groups, technologies, agents, etc.); and that the relationships among these entities are critical constraints on and enablers of individual and social action. Another reason for the movement to formal approaches is the recognition that societies and the groups and institutions within them are inherently computational since they have a need to and the ability to scan and observe their environment, store facts and programs, communicate among members, and transform information by human or automated decision making. In general, the goal of this formal research is to build new concepts, theories, and knowledge about complex systems such as groups, organizations, institutions, and societies. Using formal techniques, theorists search for fundamental social objects, processes, and the mathematical formalism with which to describe their behavior and interactions. Another goal of this research is to discover the most reasonable basis from which, at least in principle, theories of all other processes and behaviors can be derived.

Computational analysis is an ideal way to meet this goal. Indeed, the literature is replete with claims about the value and use of simulation. Fundamentally, these claims rest on the fact that computational analysis enables the theorist to think systematically and thoroughly about systems that are larger, more complex, have more interactions, and have more underlying dynamics than can be thought through without the aid of such automated accounting devices. For example, Markovsky (1992) showed that even very simple models of social interaction across networks were sufficiently complex that researchers are unable to predict their behavior. In a complex process there typically are many interacting objects (e.g., people or procedures in the society or institution), and it is rarely possible to proceed to a complete mathematical solution. Systems in which there are complex processes often exhibit nonlinear behavior and phase changes in behavior and often reach dramatically different end states given only minor changes in initial conditions. Such underlying nonlinearities make it nontrivial to think through the implications of the dynamic processes and multiactor situations. There is a general recognition that the nonlinear dynamics that characterize social systems are not mathematically tractable; hence, simulation is needed. Computational analysis can be used to track and analyze the detailed behavior within and among people, groups, teams, organizations, institutions, and societies. Computational models enable the analysis of groups far larger in size and over longer time periods than can be analyzed in a field or laboratory setting. Computational analysis enables the theorist to address issues of scalability; that is, do behaviors remain the same, do our theories hold, as we move from groups of two or three to thousands? Through simulation, we can gain some insight into whether scale matters to the nonlinear dynamics that underlie fundamental sociological processes such as reciprocity. This is particularly important as we move into a world where technology is making organizations of unprecedented size and distribution possible and giving people unprecedented access to larger numbers of others, ideas, technologies, and resources.

Computational analysis enables the theorist to think through the possible ramifications of complex systems and to develop a series of consistent predictions. Computational models, because they can track and emulate learning, adaptive, and evolutionary behavior are ideally suited to the examination of dynamic systems, to examining emergent behavior, and to suggesting the long-term impacts of new technologies and policies. Consequently, computational models can be and have been used in a normative fashion to generate a series of hypotheses by running virtual experiments. The resultant hypotheses then can be tested in other empirical settings. Other reasons for using computational analysis to do social theorizing include the facts that the resultant analysis can provide sufficiency explanations, can be used to

demonstrate gaps in extant verbal theories, and can be used to check the consistency of the predictions made using verbal theories.

In other words, computational models are most useful when there are a large number of variables, substantial nonlinearities, multiple interactions, complex interactions (such as three way or higher interactions, dynamics, and cases where there are more than a small number of actors. As such, simulations are tools for doing theory development. This is not to say that other methods cannot also enable the theorist to reason about complex dynamic systems, but that computational analysis is an important tool in the theorist's toolkit. Computational models, however, are not a panacea. A disadvantage is that such models cannot be used to conclusively demonstrate what people do in novel situations. There are limitations to their usefulness and there are areas where they are more useful than others.

Computational theorizing has led to a new way of thinking about theory and theory building, and four elements of this approach will be described: the model as theory, virtual worlds, empirically grounded theory, and hypotheses generation. After this, some of the major theoretical breakthroughs will be described, then the relation of computational reasoning about social systems to complexity theory.

THE COMPUTER MODEL AS THEORY

One of the first insights of the novice computational social theorist is that the verbal theory is incomplete compared to the level of detail needed in a computer model. Generally, moving from the verbalization to a formal representation, either mathematical or computational, requires the theorist to think through a series of relations among the component processes and entities lying at the core of the theory. This "filling in the detail" is part of the process by which theories are developed. Another part is determining which of these details are irrelevant and which are critical relative to the outcomes of interest. In other words, the computational model becomes the theory.

This same argument underlies the Turing test, a method for testing and validating computational models. Computational models do the task they seek to explain. Consequently, the model itself is substitutable for the entity that is being modeled. For example, if the model is of a human agent, then in principle in an experiment the computational model can substitute for a human or a machine. In this case if recognizable behavior emerges and if the observer cannot discern a relevant difference, then the model passes the Turing test. Turing tests vary in the degree of rigor, use of quantitative data, and use of statistics. Turing tests typically have been employed in simulations of machines or of single humans. Since the computational model does the task it seeks to explain the model itself is the theory.

From a social perspective, however, the Turing test is insufficient. Carley and Newell (1994) suggest that when the computational model is meant to act as a social agent or a group of social agents it is more appropriate to use a revised version of the Turing test that they refer to as the social Turing test. Unlike the classic Turing test, the social Turing test takes the group context into account and expects social not just behavioral responses. To do a social Turing test the researcher follows these three steps:

1. First, a collection of social agents is constructed according to the hypotheses and placed in a social situation, as defined by the hypotheses. Recognizably social behavior should emerge from the computational model.

2. Second, there will be many aspects of the computational model that are not specified by the hypotheses. That is, to create a working computational model one often has to develop the model at a level of detail well beyond that of verbal theorizing. In general, such aspects should be set based on known human data or handled using Monte Carlo techniques.
3. Third, the behaviors that emerge from the computational model can vary widely with such specification, but the behavior should remain recognizably social. The behavior of a computational model is recognizably social if it meets the following criteria. The group's behavior (two or more agents) should take into account social knowledge and the actions and interactions among the agents. Group behavior cannot be generated from the behavior of a single agent. Finally, the behavior of the computational model must match, within some predefined level of fit, the behavior of a human group in terms of qualitatively or quantitatively fitting known data on a human group or fitting a subject matter expert's description of the human group's behavior. When these criteria are met the social Turing test is met.

THE COMPUTER MODEL AS VIRTUAL WORLD

Virtual worlds are computer-simulated worlds of particular environs, social, or physical. Within a virtual world there is a particular set of physical, temporal, social or cultural laws that must be obeyed. These laws need not match those of earth. For example, in a virtual world agents could have paranormal capabilities, no emotions, or the ability to breathe under water. The worlds are populated by artificial agents. Examples range from game systems such as SimCity to research worlds for understanding the complexities of biological and social life. Virtual worlds become testbeds in which researchers can grow artificial life and communities, develop new procedures and hypotheses, and reason about the impacts of interaction among agents. Virtual worlds also can be used as a learning environment in which the student can engage in solving the problems within the scope of the world described.

One of the earliest examples of a virtual world in sociology is Bainbridge's (1987) sociological laboratory. Bainbridge created worlds in which the student could explore classic theories of social behavior. Within these worlds, the student could run virtual experiments to see the impact of changing norms of social behavior.

Artificial life, or A-life, studies are done within virtual worlds. One example of such a world is Sugarscape (Epstein & Axtell, 1997). In Sugarscape, multiple simple agents engage in social interaction in the process of consuming sugar, moving across planes, giving birth, and dying. Sugarscape is a virtual world where agents can live, eat, die, engage in social interactions, and so forth. In its simplest form the sugarscape world is a torus made by wrapping a grid of 50 by 50 squares such that sugar grows in some of the squares and not others. This world is populated by a number of agents each of whom can move one square on the von Neuman grid (NSEW) each time period as long as they do not occupy a square occupied by another agent or they can consume sugar. Consumed sugar translates into energy. Agents can eat or store sugar. Energy is needed to move, reproduce, and so forth. Without energy the agents die. Agents placed randomly on the grid over time will develop a collective intelligence that moves them toward the fields of sugar. Additional rules, such as inheritance produce more social-like behavior. Indeed, the more factors and rules that are added, the more social

behaviors can be explored. For example, introducing a second resource, spice, leads to the emergence of trading and an economic market.

Epstein and Axtell (1997) argue that they are using computers to do “bottom-up” social science. Given a collection of very simple agents with very simple rules, they try and grow complex social behavior. One of the key insights is that complex social outcomes need not have complex causes. Simple behaviors on the part of individual agents can have major social consequences.

The value of virtual worlds is that they enable the researcher and student to reason about social behavior in a controlled setting like a human experiment. However, unlike a human experiment, these worlds can be quite large. Thus issues of scale, emergence, and time varying behavior can be addressed.

THE COMPUTER MODEL AS EMPIRICALLY GROUNDED THEORY

Computer models can be empirically grounded. A key example here is Heise’s affect control theory (Heise, 1986, 1987; Smith-Lovin, 1987). Charles Osgood used semantic differential techniques to identify three dimensions of word meaning: evaluation, potency, and activity. These dimensions became known as the EPA model. Heise, in developing affect control theory, asserts that individuals construct social events to confirm the meanings of social classifications. This is done through a cognitive emotional process in which actors respond to each other at both a sociophysical and affective level. Ongoing interactions lead to alterations both in what the actors know about each other and in how they feel about themselves, the situation, and others. This affective response in turn affects the actors choice of possible actions.

Heise and co-workers have derived a number of mathematical functions to predict how people would rate various combinations of words describing social identities, attributes, actions, and situations and how these ratings would affect the actors choice of words, situations, and actions. These formula have embodied the formulas in the affect control computer model. In addition, Heise and colleagues have gathered substantial EPA data from men and women in various walks of life and in various cultures. Research subjects have provided mean EPA ratings of hundreds of words describing social identities, attributes, actions, and situations. These data also are embedded in the computational model. This empirically grounds the theory and moves the EPA ratings from being theoretical abstractions to concrete and measurable predictions.

Affect control theory, although conceptually straightforward, is sufficiently complex and nonlinear that humans have difficulty determining the full range of implications of the theory. Researchers and students can use the model to reason about the impact of emotions on social outcomes and the effect of actions, situations, and roles on emotions in dynamic situations. Using this model, the researcher can derive hypotheses from the theory to test in the field or laboratory, or can emulate an observed interaction to see whether the behavior observed matches that predicted by the model.

In general, many computational modelers have the desire to develop empirically grounded theory. In many cases, the appropriate data are not amenable to being placed in a large-scale quantitative database. Rather, the type of data needed for empirically grounding the model is ethnographic. In building a model of social and organizational processes, the

computational theorist often has a large number of questions of the form “if this happens, then what?” or “what are the space of possible human actions?” Often experimental and survey data have too little detail to answer questions at the depth posed by the modeler, whereas a single protocol or an ethnographic study typically provides the kind of rich detail necessary for addressing the modelers concerns. Although it is ironic to those who view computational theorizing as the most quantitative of the formal theorizing techniques, in actuality the ethnographer is the computational theorists best friend.

THE COMPUTER MODEL AS HYPOTHESIS GENERATOR

The process of designing, building, and analyzing a computational model is a process of theory building. An important component of this process is hypothesis generation. Computer models can in fact be viewed as hypothesis generation machines. To generate hypotheses the researcher first uses the model to conduct a virtual experiment. A virtual experiment is an experiment in which the data for each cell in the experimental design are generated by running a computer simulation model. In designing this virtual experiment, standard principles of experimental design should be followed.

There are several key ways in which virtual experiments differ from traditional human laboratory experiments. The first of these is scale. That is, it is possible in the simulated environment to run experiments with more subjects (agents), for longer periods of time, without subject fees, and to have the subjects engage in more activities, in larger groups, and so forth. The computer basically mitigates the physical, temporal, coordinating, and monetary constraints. The result is that it is very easy to design too extensive a virtual experiment. That is, it is relatively easy to generate more data than any current statistics package can handle. The constraints on the scale of the virtual experiment are computer storage space, processing speed, and the size constraints of the statistical analysis package.

Another way in which the virtual experiment differs from the traditional human laboratory experiment is in the role of statistical significance. For the human laboratory experiments, access to subjects, monetary and statistical power concerns determine the number of repetitions per cell. In this case, for an outcome variable of interest, such as organizational performance, statistically significant differences in this variable across two or more conditions may signal support for a hypothesis. The lack of a statistically significant difference signals the lack of support. In this case, the fact that there are observable differences is not interesting, whereas the fact that a difference is statistically significant is theoretically interesting and suggests that a real difference has been found. Many virtual experiment have at their core a Monte Carlo experiment, i.e., an experiment in which samples are generated randomly across some parameter space. For virtual experiments, the number of repetitions per cell are influenced by computational processing speed, storage space, and the number of repetitions per cell needed to determine whether an observed difference is statistically significant. When computational storage space and time are not at a premium, the computational theorist will simply run sufficient repetitions to guarantee that observed differences are significant. In this case, the fact that there are observable differences is theoretically interesting, the fact that they are statistically significant differences simply means the theorist had lots of computing power.

This does not mean that statistics does not have a role in analyzing the results of a virtual experiment. Rather, it means that the way in which the statistics are used is different than the way they are used in analyzing the data from a human laboratory experiment. Indeed, a critical

step in generating hypotheses from a virtual experiment is to statistically analyze the results. The results of that analysis are the hypotheses that can be examined using data from human laboratory experiments, live simulations, games, field studies, or archival sources. In conducting a virtual experiment and generating a series of hypotheses the followings stages are gone through. First, identify key inputs. Second, explore the input parameter space to determine the values of the key variables that you want to explore. Third, the nonkey inputs, such as control and secondary variables, should be set. These may be set to predefined values or chosen randomly in a Monte Carlo fashion. Fourth, the virtual experiment is conducted, i.e., a number of simulations are to be run. Fifth, to analyze the output a series of statistical analysis should be run. Sixth a number of hypotheses should be generated. This is done by taking the statistical findings and converting them into statements. For example, imagine that when a regression is run related to uncover the relation between two inputs, such as group size and level of turnover, on some outcome such as performance you find that the beta coefficient for size is .2 and for turnover is .4. Then that finding about the relative strength of the regression beta coefficients of two input variables could be converted into a statement of the form turnover has twice the impact of size on performance. These hypotheses can then be tested in other venues.

THEORETICAL BREAKTHROUGHS: THE NATURE OF THE SOCIAL AGENT

Computational models can be used in a number of ways. Several of these have been discussed: model as theory, model as virtual world, and model as hypothesis generator. Theorists using computational models have used them in these ways and several other ways (e.g., to fine-tune a human experiment). Computational models have a long history in the development of social theory and methodology (Federico & Figliozzi, 1981; Garson, 1987). Results derived from computational models have led to a number of important theoretical breakthroughs that collectively generate a more complete understanding of the social agent.

One of the key uses of computational models is to demonstrate how fundamental behavior can arise from principles other than those taken for granted and to call into question various paradigms. One of the earliest and most profound uses of computational models was to develop the theory of bounded rationality. Prior to the 1960s, most formal theories of social and organizational behavior assumed rational actors with complete information and total insight. Arising out of the Carnegie School, Herb Simon and others argued that humans were boundedly rational, i.e., social structure limits their access to data and human cognition limits their ability to process that information. As a result, decisions are made by satisficing and not by trying to locate the optimal decision. In *A Behavioral Theory of the Firm* (Cyert & March, 1963/1992) and in the garbage can model of organizational choice (Cohen, March, & Olsen, 1972) the authors demonstrate through computational based theorizing that human limitations affect what choices are made when and how. Further, by taking such bounds into account, theoretical propositions better match actual observations. This work was instrumental in revolutionizing theoretical and empirical work on group and organizational behavior. The resulting information processing view is now an integral part of many social theories.

Computer tournaments such as the now-classic “prisoner’s dilemma” computer tournament organized by Robert Axelrod (1984) have served to focus the attentions of many researchers on a particular area of behavior, leading to rapid theoretical advancement. The prisoner’s dilemma is a game-theoretic problem that explores the conditions under which cooperation may arise between self-interested actors who have the potential to gain in the short

run if they violate agreements to cooperate (Rapoport & Chammah, 1965). In a computer tournament, the organizer invites people to submit computational models that do the “task” of the tournament. Results from tournaments demonstrate repeatedly that computational modeling can produce robust and sometimes unexpected results and advance theory. Tournaments enable theories to “compete” with each other, leading to better theoretical understanding of their relative efficacy, common features, and differences. For Axelrod’s prisoner’s dilemma tournament the programs submitted were models of actors who followed various strategies for playing the iterated prisoner’s dilemma game. The “agent” that won was one of the simplest: the agent employing the “tit-for-tat” that cooperated on the first move and then imitated the previous action of its partner. This demonstrated that agents following a simple reciprocity norm could succeed even in an environment that was largely asocial. There were two central theoretical contributions: (1) the primacy of the reciprocity norm, and (2) the recognition that social outcomes are as affected by social interaction as by human cognitive architecture. Today, numerous researchers are using multiagent systems to understand the dynamics and contexts under which cooperation evolves.

Another theoretical breakthrough is in the area of chaos. The notion of deterministic chaos has captured the imagination of scholars and the public (Mandelbrot, 1983; Hao, 1984; Gleick, 1987). Kephart et al. (1992) noted that social behavior can become chaotic. This in itself is not the novel theoretical proposition. However, they ran a series of simulations that demonstrated that intelligent strategies can reduce chaos. The level of chaos in societies is reduced when the intelligent agents in these societies have the capacity to base their actions on beliefs about others’ strategies and on the observed behavior of the collection of agents. Mental models, even rudimentary ones, about others and knowledge of the collective behavior are sufficient to generate nonchaotic behavior. This suggests that in order for recognizable social behavior to emerge the content of individual cognition needs to contain mental models of others and knowledge of others actions. One of the key questions here is to what extent does this knowledge need to be of specific others versus the generalized other, and to what extent does this knowledge need to be accurate.

A related breakthrough is in the area of social knowledge. Wegner (1995), using a computer system as a metaphor for human memory, developed the powerful idea of transactive memory. Transactive memory (Wegner, 1987; Wegner et al., 1991) refers to the ability of a group to have a memory system exceeding that of the individuals in the group. The basic idea is that knowledge is stored as much in the connections among individuals as in the individuals. Wegner argues that factors that are relevant in linking computers together such as directory updating, information allocation, and coordination of retrieval also are relevant in linking the individuals’ memories together into a group memory. Empirical evidence provides some conformation and suggests that for a group, knowledge of who knows what is as important as knowledge of the task. Transactive knowledge can improve group performance (Moreland, Argote, & Krishnan, 1996). Thus, in order for recognizable social behavior to emerge, part of the content of an individual’s mental models needs to be knowledge of who knows what and presumably knowledge of who knows who.

These findings suggest a paradigmatic view in which social agents are information processors and interactors with internal mental models containing a model of self and others (specific and generalized) and others’ interaction that can be used to predict others performance, to determine who to interact with, and to select among alternative actions. This model of others and their interaction includes knowledge such as who knows who and who knows what. Norms become rules of behavior that link knowledge of self and other to specific interaction behaviors. An important element of this paradigm is that social behavior realizes

social outcomes as the result of changes in interaction among agents and agents who themselves both in their mental models and in their behavior take into account what others know, who they know, and what they are likely to do. This sociocognitive perspective on agents is partially summarized by the Carley–Newell model social agent (1994).

COMPLEXITY AND EMERGENT BEHAVIOR

One of the tremendous values of using computational models is that you can reason about dynamic systems and complex systems. Questions such as what is the impact of learning, socialization, enculturation, and so on can be addressed. Moreover, the coevolution of groups, groups and technology, two cultures, and so on can be addressed. This ability to use computational models to look at change on multiple fronts at once is a very powerful theoretical device.

The ability to use simple computational models to examine dynamics has fired the theoretical imagination leading to a vocabulary and set of tools that can be used for sociological theorizing and collectively referred to as complexity theory. Complexity analysis provides us with a means for rethinking and extending social theories (Axelrod & Cohen, 1999; Morel & Ramanujam, 1999). The vocabulary and tools are deriving in part from advances in biology and physics. Clearly the history of social theory is replete with work on change and on societies as complex, adaptive, evolutionary systems. The work today, however, uses a distinct vocabulary, has a level of formalism not heretofore possible given the ubiquitousness of computing, and has a focus on the pattern of outcomes relative to the pattern of desired outcomes. From a paradigm standpoint there is now an increase of interest in complexity, landscapes, emergence, and coevolution.

Social theory in general and organizational theory in particular has a tradition of looking at societies and organizations as complex systems. A common feature of this work and modern complexity theory is that systems are seen as more complex when there are more parts and more connections among the parts. So, for example, when there are more individuals in the society or more activities in the organization, the system is more complex. Further, when there are more ties among individuals or more constraints among organizational activities then again there is more complexity. To this the formal work on complexity theory would add that complex systems typically have internal change, adaptation, or evolutionary mechanisms that result in behavior that might appear random but that actually has an underlying order (Holland, 1998). The underlying order is attributable to multiple agents, often operating on very simple principles.

The outcomes are complex for two reasons. First, there may be an appearance of randomness. Second, what outcomes are reached can be quite divergent depending on the initial conditions the history of the agents activities (Kauffman, 1995). Complex systems have the ability to self-organize (Bak, 1996). Self-organization means in part that group behavior occurs that is distinct from a simple average of individual agent behaviors. For example, imagine a group of agents moving through a grid. Even though individual agents may only move N, S, E, or W, the group *qua* group can move NE. Moreover, even though individual agents may be trying to optimize their outcome, the overall outcome for the group may be less than optimal. For example, in a group-level repeated dilemma game even though each individual is trying to maximize their return, the overall pattern of cooperation may result in the return to the group being suboptimal. Having a group-level outcome that is less than the optimal can be thought of as a complexity catastrophe (McKelvey, 1999).

The research on complex systems varies in whether or not a time-invariant fitness

function is assumed. A fitness function defines the relation of inputs to outputs. A landscape is a graphic depiction of the relation of inputs to the output of concern, of the independent variables to the dependent variable (Levinthal, 1997). In this case there is a known a priori fitness function. The landscape is fixed and features of the landscape may make it harder or easier for individuals and groups to locate the optimum point in this landscape. In contrast, the fitness function may not be known a priori and/or the relation of inputs to outputs is time variant (dancing landscapes). In this case features of the landscape and of the way it changes determine the ease or difficulty of moving about this space. Changes in the landscape are commonly attributable to factors such as learning, innovation, population, turnover, birth and death processes, and so on. Agents and groups move about on this landscape. For some theorists, the agents move about because they are trying to optimize their outcomes. For other theorists, the agents are seen as trying to satisfice or to move randomly about on this landscape. While there is growing agreement that social theories need to be dynamic, there is not widespread theoretical agreement about the shape of the underlying landscapes, the nature of the fitness function, whether or not the fitness function is fixed over time, and whether agents optimize or satisfice.

In complex systems things emerge. As previously noted this emergence may take the form of group outcomes that are distinct from individuals, or they may take the form of new procedures, groups, ideas, and so forth. The theoretical issues include what enables new things to emerge, to become stable, to become accepted and whether the history of the system alters what can and does emerge. In Kauffman's models (Kauffman, 1993) emergent structures derived from two sources: (1) forcing functions, and (2) homogeneity bias. When only one input can force a particular outcome state, there is a forcing function. An example here is the relation between age and drug usage. If it is the case that if you have never used drugs by the time you are 35, then you never will use drugs, then there is a forcing function. There will be a resultant structure to the relation between age and drugs that remains fixed across populations and groups such that there will be a larger than chance number of individuals over 35 who do not use drugs. The second factor is a homogeneity bias. Homogeneity biases may be processes or limits on variables that reduce the number of actual outcomes. For example, if the outcome were clothing color and only two dyes were available, that is a stronger homogeneity bias than if the number of dyes were unlimited. Other theorists working with complex systems have noted that learning (Carley, 1990, 1991), mimicry (Macy, 1991a,b), and temporal and physical constraints (Carley & Prietula, 1994) all lead to a certain level of emergent structuration. The basic nature of dynamic systems is that things tend to emerge. Thus the issue is not do patterns emerge, but what patterns emerge when and under what conditions.

The work on organizational adaptation is a classic example. Here, computational theorists have shown that individual learning and strategic choice do enable organizations to adapt. Change, in fact, is pretty much inevitable. Some of those changes are adaptive, i.e., they enable certain outcomes to be maintained or improved. Moreover, this work demonstrates that the emergent patterns are path dependent (history matters) and a complex function of interactions among agents, knowledge, and task. Interdependence among agents and among knowledge and tasks is a fundamental feature of social and organizational systems that affects group and organizational learning, performance, and adaptability (Levinthal, 1997; Sorenson, 1997; Krackhardt & Carley, 1998).

A related argument surrounds the notion of coevolution. Work in the area of ecology led to the view that populations of agents evolve through processes involving Darwinian selection of the fittest and trait propagation. This view was translated to the study of organizations where population ecologists (Hannan & Freeman, 1977, 1989; Hannan & Carroll, 1992), proponents

of organizational evolution (Aldrich, 1979; Baum & Singh, 1994b), and economists (Nelson & Winter, 1982) all put forward the view that selectionist processes drive out the less fit firms, leaving “order” to be explained as the consequence of the survival of the more fit firms. Factors such as population density and niche overlap were used, in conjunction with evolutionary principles to explain the varying states populations of firms. Kauffman (1993) suggests that complexity effects may thwart selectionist effects under some circumstances. In particular he argues that the extent of interdependence within and among firms when coupled with regression to the mean can generate order even when the most fit firms do not survive. Kauffman (1993) suggests that *internal* and *external* interdependencies (K and C , respectively) create a level of complexity that can explain order in any world that can be thought of as a multiple groups with linkages within and between the groups. His analysis suggests that aggregate economic order (such as the distribution of size of firms that survive or the distribution of survival times) could be at least as much determined by intrafirm and interfirm dependencies as it might stem from external selection.

These interdependencies ensure that firms coevolve. This line of reasoning, although novel to economics is the familiar social network explanation long part of sociological theorists toolkit. However, from a sociological perspective the core assumptions are obviously nonsocial. For example, a limitation of Kauffman’s model from a grounded theory perspective is that it assumes that the number of links within and between groups (or in the example above, within and among organizations) is fixed over time and identical across group members. In other words, the shape of the underlying social network is fixed with uniform centrality per agent. Thus, social order is explained in this work without an appeal to differentiation in the network. In contrast, the Carley (1990, 1991) constructural model, the social network, and the knowledge network coevolve, and initial network differentiation is a key determinant of order.

Not all work in complexity assumes away the importance of the pattern of relations. A parallel line of theorizing, derived from information processing theory and cognitive science, argues that the relations within and among groups evolve over time through a process of individual learning. Moreover, these links are seen as embedding knowledge. In particular, organizational learning theorists (Levinthal, 1997; Carley & Hill, 2001) suggest that even when there is not population-level evolution, organizational learning is sufficient to generate order. In this case, learning determines both the degree of interdependence and the value of that interdependence. For Carley (1990, 1991), the degree of interdependency within and among groups is dynamic and coevolves with the knowledge network (who knows what). Natural leaning and communication processes coupled with barriers to communication and to learning and the advent of new people, technology, and inventions prevent or enable groups from collectively moving toward optimal positions in the landscape and produce order. Coevolution occurs on multiple fronts: individuals coevolve with each other, individuals and groups coevolve, groups coevolve with each other, culture and social structure coevolve, and patterns of individual knowledge and interaction coevolve.

This work has led to a neoinformation-processing perspective in which social outcomes emerge from network dynamics over a set of interlocked networks. Traditional social theory has looked at networks in terms of the relations among people, or at an interorganizational level—the relations among organizations. Linguists and learning theorists have looked at networks in terms of the relations among knowledge and people. Operational researchers at networks as relations among tasks. This can be brought together in a unified metanetwork perspective where the relations within and among agents, knowledge, tasks, and organizations are seen as constraining and enabling social behavior (see Fig. 4.1). Moreover, these networks are seen as coevolving through dynamics such as learning, innovation, and population growth.

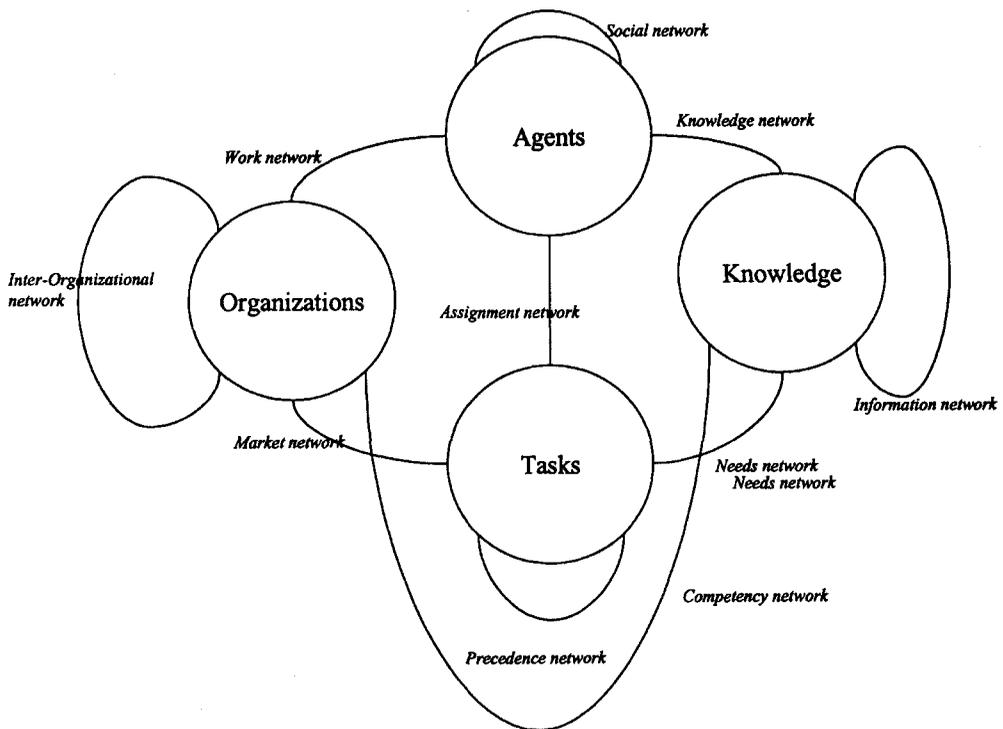


FIGURE 4.1. Embedding social behavior in a metanetwork.

A GLIMPSE OF THE FUTURE

Traditionally there has been a strong tie between the study of machine and human intelligence. For example, mathematical models of human learning (Bush & Mosteller, 1955) were among the first “learning” algorithms used in artificial agents. Advances in cognitive science have come hand-in-hand with advances in computer science. Today, one of the links between social theory and computer science is in the area of machine learning. Machine learning techniques inform and are informed by our understanding of both individual cognition and group or organizational learning. One of the areas of greatest potential for their application is in the area of social agents, another is in the area of organizational theory.

Social agents, such as personalized webbots, avatars, and socially sophisticated data bases are becoming increasingly technologically feasible. Artificial agents can generate and respond to rudimentary emotions, do tasks, and answer questions. As more of these agents are created the issue arises as to how these agents can and should interact with each other and human beings. New questions that need to be addressed include “do these agents need to be social to communicate effectively with humans?” and “will a society of webbots be more or less effective if it follows the same norms of behavior as a human society?” For the social theorists the creation of societies of these agents provides an important opportunity for addressing fundamental questions about the nature and value of socialness and the relation of sociality to other human characteristics such as emotions and cognitive processing.

Computational analysis and theorizing is playing an increasingly important role in the

development of social and organizational theory. There is a growing recognition that complex adaptive processes underlie social life. There are three elements to this view. First, social and organizational processes are complex, dynamic, adaptive, and nonlinear. Second, organizational and social behavior emerges from interactions within and between ecologies of agents, resources, knowledge, tasks, and other organizations. Third, the relationships among and within these entities are critical constraints on and enablers of individual and organizational decision making and action. There is a parallel recognition that societies in general and organizations in particular are inherently computational, the reason being that they have a need to scan and observe their environment, store information and procedures, communicate, and transform information through human or artificial agents. Computational theories are providing the social theorist with both a new toolkit for examining social phenomena and new insights into the fundamental nature of the social agent.

ACKNOWLEDGMENTS. This work was supported in part by the Office of Naval Research ONR N00014-97-1-0037 and by the National Science Foundation NSF IRI9633 662.

REFERENCES

- Aldrich, H. (1979). *Organizations and environments*. Englewood Cliffs, NJ: Prentice Hall.
- Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
- Axelrod, R. M. (1997). *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton studies in complexity. Princeton, NJ: Princeton University Press.
- Axelrod, R. M., & Cohen, M. D. (1999). *Harnessing complexity: Organizational implications of a scientific frontier*. New York: Free Press.
- Bainbridge, W. S. (1987). *Sociology laboratory*. Belmont, CA: Wadsworth.
- Bainbridge, W., Brent, E., Carley, K., Heise, D., Macy, M., Markovsky, B., & Skvoretz, J. (1994). Artificial social intelligence. *Annual Review of Sociology*, 20, 407–436.
- Bak, P. (1996). *How nature works: The science of self-organized criticality*. New York: Copernicus.
- Baum, J. A. C., & Singh, J. V. (1994a). Organization-environment coevolution. In J. A. C. Baum & J. V. Singh (Eds.), *Evolutionary dynamics of organizations* (pp. 379–402). New York: Oxford University Press.
- Baum, J. A. C., & Singh, J. V. (Eds.). (1994b). *Evolutionary dynamics of organizations*. New York: Oxford University Press.
- Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. New York: John Wiley and Sons.
- Carley, K. (1990). Group stability: A socio-cognitive approach. In E. Lawler, B. Markovsky, C. Ridgeway, & H. Walker (Eds.), *Advances in group processes: Theory and research*, Vol. VII (pp. 1–44). Greenwich, CT: JAI Press.
- Carley, K. (1991). A theory of group stability. *American Sociological Review*, 56(3), 331–354.
- Carley, K. (1995). Communication technologies and their effect on cultural homogeneity, consensus, and the diffusion of new ideas. *Sociological Perspectives*, 38(4), 547–571.
- Carley, K. M. (1999). On the evolution of social and organizational networks. In S. B. Andrews & D. Knoke (Eds.), *Networks in and around organizations* (pp. 3–30). Stamford, CT: JAI Press.
- Carley, K., & Hill, V. (2001). Structural change and learning within organizations. In A. Lomi (Ed.), *Dynamics of organizational societies: Models, theories and methods*. Boston: MIT Press/AAAI Press/Live Oak.
- Carley, K., & Newell, A. (1994). The nature of the social agent. *Journal of Mathematical Sociology*, 19(4), 221–262.
- Carley, K. & Prietula, M. (1994). ACTS theory: Extending the model of bounded rationality in K. M. Carley and M. Prietula (Eds.), *Computational Organizational Theory* (pp. 55–87). Hillsdale, NJ: Lawrence Erlbaum.
- Carley, K. M., & Svoboda, D. M. (1996). Modeling organizational adaptation as a simulated annealing process. *Sociological Methods and Research*, 25(1), 138–168.
- Cohen, M. D., March, J. G., & Olsen, J. P. (1972). A garbage can model of organizational choice. *Administrative Sciences Quarterly*, 17(1), 1–25.
- Cyert, R., & March, J. G. (1963/1992). *A behavioral theory of the firm*, 2nd ed. Cambridge, MA: Blackwell Publishers.

- Epstein, J., & Axtell, R. (1997). *Growing artificial societies*. Boston: MIT Press.
- Federico, P., & Figliozzi, P. W. (1981). Computer simulation of social systems. *Sociological Methods and Research*, 9, 513–533.
- Feinberg, W. E., & Johnson, N. R. (1988). Outside agitators and crowds: Results from a computer simulation model. *Social Forces*, 67, 398–423.
- Feinberg, W. E., & Johnson, N. R. (1990). Radical leaders, moderate followers: Effects of alternative strategies on achieving consensus for action in simulated crowds. *Journal of Mathematical Sociology*, 15, 91–115.
- Garson, G. D. (1987). Computer simulation in social science. In *Academic microcomputing* (pp. 110–138). Newbury Park, CA: Sage.
- Gleick, J. (1987). *Chaos*. New York: Penguin.
- Hannan, M. T., & Carroll, G. R. (1992). *Dynamics of organizational populations*. New York: Oxford University Press.
- Hannan, M. T., & Freeman, J. (1977). The population ecology of organizations. *American Journal of Sociology*, 83, 929–984.
- Hannan, M. T., & Freeman, J. (1989). *Organizational ecology*. Cambridge, MA: Harvard University Press.
- Hao, B. (Ed.). (1984). *Chaos*. Singapore: World Scientific. (NOTE: some libraries list this under Bai-Lin, Hao)
- Heise, D. R. (1986). Modeling symbolic interaction. In S. Lindenberg, J. S. Coleman, & S. Nowak (Eds.), *Approaches to sociological theory* (pp. 291–316). New York: Russell Sage Foundation.
- Heise, D. R. (1987). Affect control theory. Concepts and model. *Journal of Mathematical Sociology*, 13(1–2), 1–33.
- Holland, J. (1998). *Emergence: From chaos to order*. Reading, MA: Perseus Books.
- Kaufman, D. S., & Carley, K. M. (1993). *Communication at a distance: The effect of print on socio-cultural organization and change*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kauffman, S. A. (1993). *The origins of order: Self-organization and selection in evolution*. New York: Oxford University Press.
- Kauffman, S. A. (1995). *At home in the universe: The search for laws of self-organization and complexity*. New York: Oxford University Press.
- Kephart, J. O., Huberman, B. A., & Hogg, T. (1992). Can predictive agents prevent chaos? In P. Bourgin & B. Walliser (Eds.), *Economics and cognitive science* (pp. 41–55). Oxford: Pergamon Press.
- Krackhardt, D., & Carley, K. M. (1998). A PCANS model of structure in organization. In *Proceedings of the 1998 International Symposium on Command and Control Research and Technology* (pp. 113–119). Monterey, CA: EBR.
- Lave, C. A., & March, J. G. (1975). *An introduction to models in the social sciences*. New York: Harper and Row.
- Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*, 43, 934–950.
- Levinthal, D. A., & March, J. B. G. (1993). The myopia of learning. *Strategic Management Journal*, 14, 95–112.
- Macy, M. W. (1991a). Chains of cooperation: Threshold effects in collective action. *American Sociological Review*, 56, 730–747.
- Macy, M. W. (1991b). Learning to cooperate: Stochastic and tacit collusion in social exchange. *American Journal of Sociology*, 97, 808–843.
- Macy, M., & Skvoretz, J. (1998). The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(October), 638–660.
- Mandelbrot, B. B. (1983). *The fractal geometry of nature*. San Francisco: W. H. Freeman.
- March, J. G. (1996). Exploration and exploitation in organizational learning. In M. D. Cohen & L. S. Sproull (Eds.), *Organizational learning* (pp. 101–123). Thousand Oaks, CA: Sage.
- Markovsky, B. (1987). Toward multilevel sociological theories: Simulations of actors and network effects. *Sociological Theory*, 5, 101–117.
- Markovsky, B. (1992). Network exchange outcomes: Limits of predictability. *Social Networks*, 14, 267–286.
- Markovsky, B., Skvoretz, J., Willer, D., Lovaglia, M. J., & Erger, J. (1993). The seeds of weak power: An extension of network exchange theory. *American Sociological Reviews*, 58, 197–209.
- McKelvey, B. (1997). Quasi-natural organization science. *Organization Science*, 8, 351–380.
- McKelvey, B. (1999). Avoiding complexity catastrophe in coevolutionary pockets: Strategies for rugged landscapes. *Organization Science*, 10, 294–321.
- McPhail, C., & Tucker, C. W. (1990). Purposive collective action. *American Behavioral Scientist*, 34, 81–94.
- Morel, B., & Ramanujam, R. (1999). Through the looking glass of complexity: The dynamics of organizations as adaptive and evolving systems. *Organization Science*, 10(3), 278–293.
- Moreland, R. L., Argote, L., & Krishnan, R. (1996). Socially shared cognition at work: Transactive memory and group performance. In J. L. Nye & A. M. Brower (Eds.), *What's social about social cognition? Research on socially shared cognition in small groups* (pp. 57–84). Newbury Park, CA: Sage.
- Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Cambridge, MA: Belknap/Harvard.

- Rapoport, A., (1911). *Mathematical models in the social and behavioral sciences*. Translated from the German *Mathematische Methoden in den Sozialwissenschaften*. New York: Wiley.
- Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma: A study in conflict and cooperation*. Ann Arbor: University of Michigan Press.
- Sastry, M. A. (1997). Problems and paradoxes in a model of punctuated organizational change. *Administrative Science Quarterly*, 42(2), 237–275.
- Skvoretz, J., & Fararo, T. J. (1989). Action structures and sociological action theory. *Journal of Mathematical Sociology*, 14, 111–137.
- Skvoretz, J., & Fararo, T. J. (1994). Adaptive processes, grammars of action and genetic algorithms: Toward a synthesis. In B. Agger (Ed.), *Current perspectives in social theory*, Vol. 14. (pp. 275–299). Greenwich, CT: JAI Press.
- Smith-Lovin, L. (1987). Affect control theory: An assessment. *Journal of Mathematical Sociology*, 13, 171–192.
- Sorenson, O. (1997). The complexity catastrophe in the evolution in the computer industry: Interdependence and adaptability in organizational evolution. Unpublished Ph.D. dissertation, Sociology Department, Stanford University, Stanford, CA.
- Tuma, N. B., & Hannan, M. T. (1984). *Social dynamics: Models and methods. Quantitative studies in social relations*. Orlando, FL: Academic Press.
- Wegner, D. M. (1987). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185–208). New York: Springer-Verlag.
- Wegner, D. M. (1995). A computer network model of human transactive memory. *Social Cognition*, 13(3), 319–339.
- Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology*, 61, 923–929.
- Yamagishi, T., Gillmore, M. R., & Cook, K. S. (1988). Network connections and the distribution of power in exchange networks. *American Journal of Sociology*, 93, 833–851.