Journal of Mathematical Sociology, 1989, Vol. 14(2-3) pp. 171-208
Reprints available directly from the publisher
Photocopying permitted by license only

Printed in the United States of America

THE VALUE OF COGNITIVE FOUNDATIONS FOR DYNAMIC SOCIAL THEORY

KATHLEEN CARLEY

Assistant Professor of Sociology
Department of Social and Decision Sciences
Carnegie-Mellon University

6 October 1987

This work was supported in part by the NSF under grant No. IST-8607303.

It is possible to develop models of social behavior that are predicated on detailed mechanical models of cognition. Cognitively based social models are potentially unified theoretical frameworks that can be used to explain a wide variety of social phenomena. Moreover, if a knowledge representation scheme and a knowledge acquisition scheme are specified in the underlying cognitive model then it is possible to produce a dynamic social model. The resulting social model can thus be used to predict and explain not only conditions for specific behaviors but changes in those behaviors over time.

Constructuralism is a theory of social behavior that rests on a cognitive model. The cognitive model specified has a knowledge representation scheme, knowledge acquisition procedures, and control procedures for shifting cognitive attention. The resulting social model is a dynamic model that can be used to explain both conditions for the occurrence of a behavior and social and individual changes that accrue do to a series of behaviors. The explanatory breadth of the model is illustrated by looking at predictions about a variety of social phenomena including: development of shared knowledge, identical behavior by members of the society, foreign language acquisition, clique formation, civil disobedience, and diffusion of innovative information.

Keywords: Attribution, Cognitive Structure, Social Action, Social Structure.

INTRODUCTION

Sociology has much to gain by looking to the area of cognitive science. There are two ways in which sociology can benefit from cognitive science. First, cognitive science has provided us with a new set of analytical tools. Second, the findings in cognitive science make it possible to develop models of individual cognition. Now sociologists can develop models of social behavior that are predicated on cognitive models. The

second of these is potentially more important for sociology. This is because cognitively motivated theories have the potential for providing a unified theoretical framework that can be used to explain a wide variety of social phenomena. This paper illustrates this point by showing how constructuralism, a cognitively motivated theory of social behavior, can be used to explain a wide variety of social phenomena. The phenomena look at include acceptance of innovations, education, and civil disobedience.

The porting of methods from cognitive science and artificial intelligence to sociology will increase the repertoire of analytic tools available to the researcher. The porting of methods has already begun. Event sequence analysis and symbolic logic are being used to analyze social interactions and theoretical conceptions (Heise, 1987). Expert systems have been used to improve the quality of coded verbal protocols (Carley, 1988a). Event ordered knowledge representation systems have been used to explore conditions on social behavior (Goffman, 1963, Goffman, 1974). Production systems have been used to look at role behaviors (Skvoretz, 1980, Skvoretz, 1988). Conceptual modeling has been used to teach history, specifically in the area of revolutions — (Ennals, 1985). This will allow researchers to explore new domains and to increase the precision with which they explore old domains.

For the most part, models of social behavior are not predicated on cognitive models. As Cicourel notes (1974, p. 11):

> "It is commonplace in sociology for writers to acknowledge the ultimate importance of the interacting situation between two or more actors. The assumed relationship between structure and process, however, is often no more than an expression of faith rather than the integration of social process with social structure . . ."

Explicit procedures that relate the individual cognitive world to the social world, explicit cognitive procedures, and a theory of meaning are absent from such formulations of social interaction; e.g., the work of Davis (1949) on status, Parsons (1951) on role expectation, Homans (1961) on social stimulus, Blau, (1967, 1977) on social status, and Granovetter on social ties and mobility (1973, 1974, 1982).

Still, there are a variety of theories that rely on simple cognitive models of individuals. Many of these theories rely on an information processing or limited rationality view. Consider the following two examples. Symbolic interactionism (Blumer, 1969, McCall, 1978, Stryker, 1980) relies on a limited rationality view coupled with an identity maintenance principle to cognitively model the individual. Individuals are not viewed as bringing all of the information that they have to bear on a situation in order to determine their behavior and to predict other's behavior. Rather, the nature of the situation constrains individual cognition by limiting what behaviors are viewed as available for maintaining or altering identities and the definition of the situation. Whereas, Goffman (1963, 1974) rests his theory of interaction on a cognitive model in which the individual has limited rationality and a scripted knowledge representation scheme, referred to as frames. In this case, the structure of the frame defines the situation and limits the individuals subsequent actions. Further, there are a variety of theories that do predicate social behavior on somewhat more detailed cognitive

assumptions. For example, cognitive dissonance theory (Festinger, 1957), social comparison theory (Festinger, 1954), and affect control theory (Heise, 1977, Heise, 1978, Heise, 1979, Heise, 1987), Such theories do try to provide the cognitive principles or rules that underlie social behavior. Theories such as these illustrate the theoretical power that can be derived by bringing even imprecise and relatively simple conceptions of individual cognition to bear.

These theories, despite their cognitive underpinnings, are not dynamic cognitive social theories in the sense that they do not explicitly specify a full procedural cognitive model and they do not admit prediction of continual social change. At least one of the following components is missing from the cognitive model employed by each of these theories: a specific knowledge representation scheme, detailed knowledge acquisition procedures, control procedures for shifting cognitive attention, and mechanisms for generating behavior from the cognitive representation of the task or situation. In all of these theories, the knowledge acquisition procedures are not specified. Thus the theories are somewhat static. That is, these theories focus on predicting or explaining behavior given a situation. They do not specify chains of behavior overtime as the individual engages in a series of similar tasks or is present in similar situations.

Moreover, these theories do not specify how changes to society occur relative to the individuals' cognition and behavior. This is basically Cicourel's point (1974, p. 27) when he argues that "The social analyst's use of theoretical concepts like role actually masks the inductive or interpretive procedures whereby the actor produces behavioral displays which others and the observer label 'role behavior'. Without a model of the actor that specifies such procedures or rules, we cannot reveal how behavioral displays are recognized as 'role taking' or 'role making'." The theories specified above have a sufficient cognitive model for them to begin to explain role taking behavior. The lack of dynamics in these theories, however, prevents them from developing procedural explanations of role making behavior.

If we are to develop dynamic models of social phenomena, that permit us to explore role-making behavior then what we must rest from cognitive science are the knowledge representation schemes, the consequent knowledge acquisition schemes, the control procedures for shifting cognitive attention which allow differentiation among situations and the development of multiple roles, and mechanisms for generating behavior from the cognitive representation which admit precise behavioral predictions.

Constructuralism is a theory about the co-construction of the cognitive world of the individual and the social world. Constructuralism differs from the works previously described in that it is a dynamic theory of social and cognitive evolution whose predictions are a function of the model of individual cognition. The cognitive model used specifies the knowledge representation scheme, the knowledge acquisition scheme, and limited control knowledge. Precise behavior generating mechanisms are not specified. Since a precise behavior generating mechanism is not provided, predictions about the development of specific roles can not be made. Predictions about role development processes and social dynamics can be made.

Given even a simplistic model of individual cognition, constructuralism can be used to explain a wide variety of social phenomena. These phenomena include knowledge

acquisition rates (Carley, 1987a), consensus in knowledge (Carley, 1987a, Carley, 1987b), development of social knowledge (Carley, 1987a), and diffusion of innovative information (Carley, 1987a). With a slightly more realistic cognitive model constructuralism can also be used to explain more complex phenomena such as effective educational procedures (Carley, 1986a), and miscommunication (Carley, 1986a). Further, the more realistic the cognitive model the better constructuralism is at predicting specific behavior, e.g. voting outcomes (Carley, 1987b). In this paper, an extended version of the frame structured cognitive model presented in (Carley, 1984, Carley, 1986b, Carley, 1986a) is used. Consequently, constructuralism is used to explain or make predictions about a variety of complex socially dependent phenomena such as development of shared knowledge, identical behavior by members of the society, foreign language acquisition, clique formation, civil disobedience, and diffusion of innovative information.

A brief description of constructuralism is presented in section 1. For a detailed description of this theory see (Carley, 1988d). The constructural theory is based on a few simple cognitive principles. Section 2 contains a description of a possible cognitive model. Given this cognitive basis constructuralism is then used to explain a variety of social phenomena in Section 3.

1. CONSTRUCTURALISM — AN OVERVIEW

Constructuralism is the theory that the social world and individuals' cognitive worlds co-evolve. The co-evolution occurs as individuals go through their lives, performing tasks, interacting, and acquiring information.[1] The social world is a reflection of the individual cognitive world. The social world is observable only as patterns in individuals' behavior that persist over time.

The social world parallels the cognitive world of the individual — see figure 1. Assume for the moment that the information in the individual's cognitive structure can be divided into three components — general knowledge, interaction propensities, and language. Where, both interaction propensities and individual language are derivable from the general knowledge. In this case three components would be observed for the social world — social knowledge, social interaction structure, and social language. Social structure would be a function of both shared knowledge and social interaction structure. Note: similar to the work by Vygotsky (1962, 1978) for both knowledge and language an inter-mentant (social) and a related, but not equivalent, intra-mentant (individual) dimension are being suggested.

Evolution is a construction process. Individuals by acquiring knowledge are continually constructing their cognitive structure. The social world is constructed by accident as individuals construct their cognitive structures. Social change is an artifact of changes in individuals' cognitive structures. Thus, the rules that govern cognitive development implicitly govern social development.

In order to model this theory a model of the society and a cognitive model of the individual are needed. Figure 2 is a top level view of the necessary components of this model. Following is a brief discussion of some of the major tenets and assumptions needed to model the constructural theory. These can be thought of as a blue-print for a general model sans a specific model of cognition.

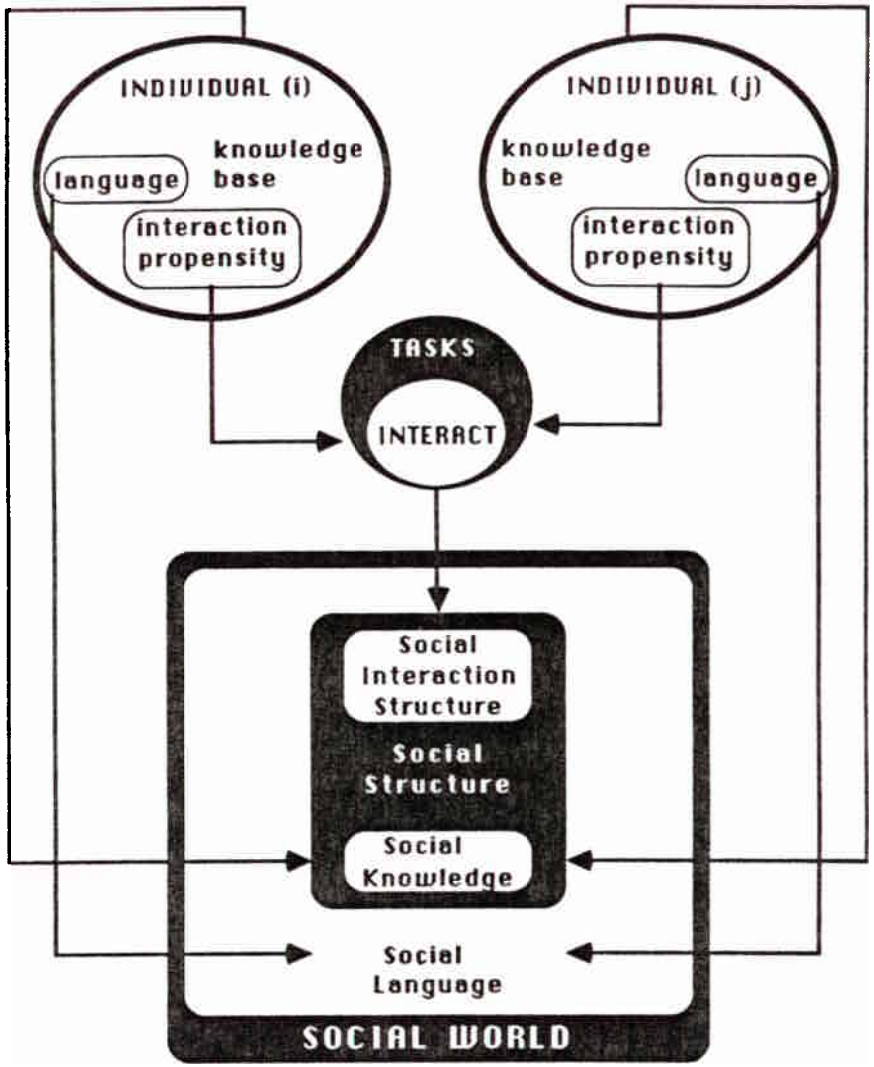

FIGURE 1 Relating the Social World to the Individual's Cognitive Structure

The social world parallels the cognitive world of the individual. Two individual's *i* and *j* are shown at the top of the figure. The components of both individuals' cognitive structures are the same — general knowledge, interaction propensities, and language. The specific values for these components is different for each individual. Both interaction propensities and individual language are derivable from the general knowledge. At the bottom the social world is shown with the three components that correspond to the individuals' cognitive components. These are social knowledge, social interaction structure, and social language. The individuals' knowledge bases collectively form social knowledge. Social language is a function of individuals' languages. The individual's knowledge base determines the individual's interaction propensities. Individuals' interaction propensities, plus the task being performed, determine whether or not two individuals interact. The social interaction structure is the observed interaction structure. It is a function of patterns of interaction that emerge over time. Social structure is a function of both shared knowledge and social interaction structure.

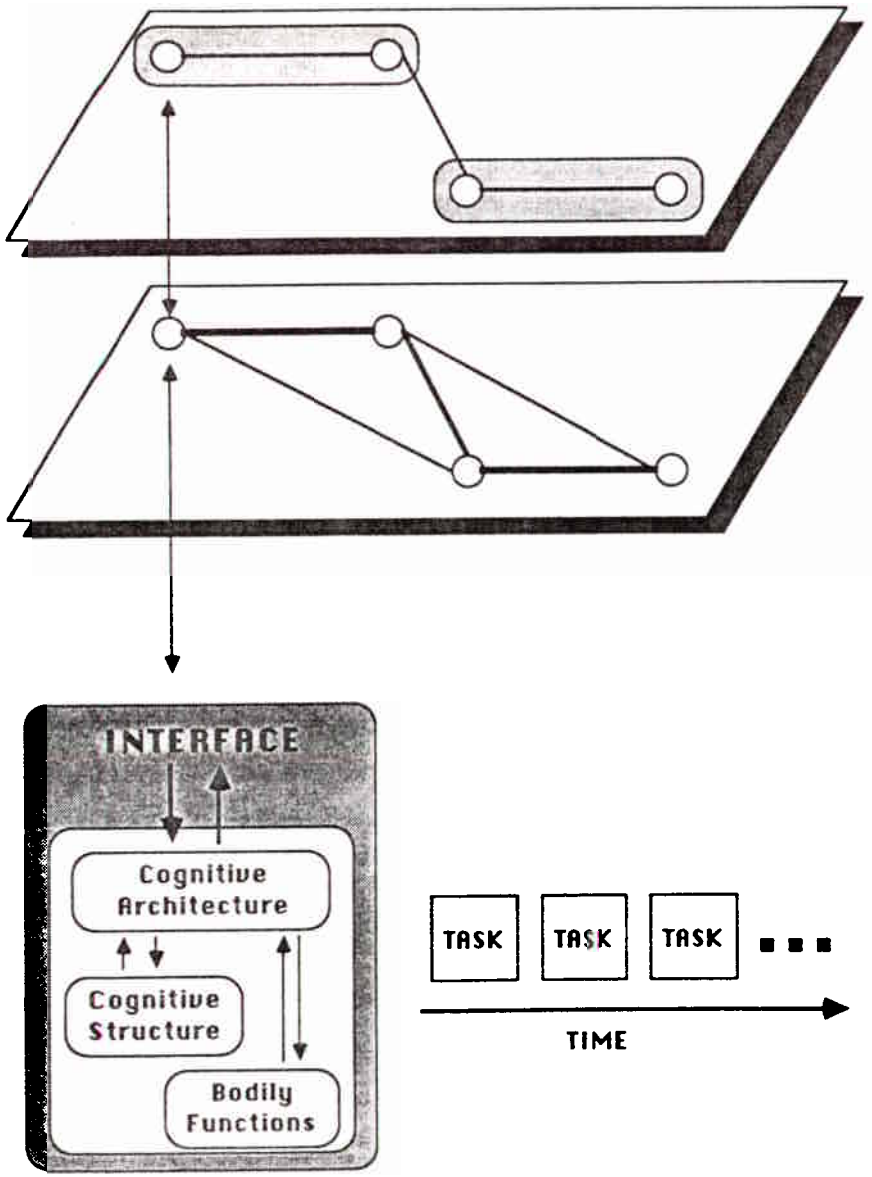

FIGURE 2 Top Level View of Model

The model exists at three levels. At the top, the social level is shown. Here the society is modeled as a set of individuals and the observable ties between them. Social structure can be extracted by locating the groups of individuals that persist over time. A group would be defined as a set of individuals with a certain pattern of ties between them. In the middle, the socio-cognitive level is shown. Here the society is modeled as a set of individuals and the actual cognitive ties between them. These models rest on a third model of individuals performing tasks — bottom level. Here the individual is modeled as having an interface to the external world, a cognitive architecture, a cognitive structure — part of which is currently salient in working memory, and bodily functions. Note the cognitive architecture controls the interface, and the individual's bodily functions affect the individual's cognitive structure through the architecture. Each time period the individual works on a task. These tasks define what part of the environment is currently available to the individual.

The individual is modeled as having a cognitive architecture, a cognitive structure, and a current working memory. By analogy with expert systems terms, the individual is modeled as having an inference engine, a knowledge base, and a working memory (or blackboard) (Hayes, 1983, Davis, 1980, Waterman, 1986). The individual's cognitive architecture is the set of basic mechanical functions that allow the individual to process information. The cognitive architecture does not change and the individual acquires new information.[2] The individual's cognitive structure is the information that the individual knows; e.g. facts, concepts, meanings, interaction propensities, and the relationships between this information. The cognitive structure changes as the individual acquires new information. Moreover, the acquisition of knowledge leads to the construction of meaning (Fauconnier, 1985) and the evolution of language (Whorf, 1956, Carley, 1981). This in turn affects the rate of information transmission during and the effectiveness of future communications. The individual's working memory contains that information that is currently salient that is in or is generable from the individual's cognitive structure.

Individuals are assumed to have the same cognitive architecture, barring physical disabilities. Whereas, individuals are assumed to differ in terms of their cognitive structure. Since some of the information that the individual learns may be information on how to process other information, individuals may appear to have different cognitive architectures. Carbonell (1981) showed that vastly different behaviors could result given the same cognitive architecture but different cognitive structures. That is, individuals with the same mechanisms for processing information but with different information behave in different ways.

In current network tradition, a society is viewed as a set of individuals connected into a multi-dimensional network by multiple ties (White, 1976, Boorman White, 1976, Burt, 1976, Burt, 1977, for example). At any point in time there are a certain number of individuals in the society. This number may change as individuals leave or enter the society. These ties are of two types: observable shared external feature, e.g. they interacted, they are of the same race (Blau, 1977, Breiger, 1978), and shared cognitive feature, e.g. they share the same piece of knowledge (Crane, 1970, Danowski, 1980). The observed social world is a function of these observable shared external features. The social world is constructed as shared cognitive feature ties are developed.

The basic social act is interaction.[3] That is, what is observed as social structure is a function of who interacts with whom and consequently who shares what knowledge with whom. During interactions individuals communicate and through this communication relate to each other and their environments. Interaction and communication make possible the creation and maintenance of social realities (Berger, 1966, Duncan, 1968, Thayer, 1968, Ruben, 1975). If individuals did not interact and communicate the social institutions such as roles, religions, etc. would not develop. Whether or not two individual's interact is a function of their opportunities for interaction and their cognitive structure. See (Carley, 1987b, Carley, 1981, Carley, 1987c) for a possible model. When individuals do interact they communicate, i.e. they send information to each other (Carley, 1986b, Carley, 1986a). What information is communicated, whether or not the information communicated is accepted, and how that information affects the individual's cognitive structure is a function of the

individuals cognitive structure (Carley, 1986a, Carley, 1988d). Thus different models of this cognitive structure will lead to different conclusions. In the following section a possible cognitive model is specified. Given this model the communication and acquisition of information is discussed.

Social structure is a set of groups that span the society and the ties between those groups. These groups are equivalence classes on some subset of ties. Individuals are said to have a shared social position if they are a member of the same structural group. Social structure is an observed phenomena. Consequently, what is measured as the structure depends on the number of individuals looked at, the ties used for determining equivalence classes, and whether actual or partial equivalence is used. The more time periods over which a particular class or structure persists the more institutionalized, or part of the social fabric, or part of the culture that class or structure can be said to be.

The individual's life is modeled as a series of discreet time periods during which the individual processes information and acts upon that information. Each time period is composed of two parts — an information processing cycle and a behavioral or decision making cycle. Note: this is similar to the elaboration and decision cycles in SOAR (Laird, 1986a, Laird, 1986b). The time period or length of the information processing cycle is a function of the cognitive architecture and determines how long the individual has to access a particular fact. It may be that external factors such as emotions and bodily functions may affect individual behavior by affecting the length of the information processing cycle.

What behavior the individual engages in is determined in a mechanical fashion by processing the cognitive structure using the cognitive architecture. What behavior the individual engages in is a function of the individual's cognitive structure. Note, this is similar to the selection of behavior in affect control theory (Heise, 1977, Heise, 1978, Heise, 1979, Heise, 1987). Given an accurate model of the cognitive architecture and the state of the individual's cognitive structure it would be possible to predict the individual's behavior.

At any one time, the individual will be working on a specific task. Different individuals may be working on different tasks. The individual will switch tasks under at least two conditions: 1) they have finished the task e.g. when a decision has been reached, or 2) new information signals a change to a new task. Each time period the individual will interact with someone, perhaps himself.

There are both external, i.e. non cognitive, and internal factors, i.e. cognitive, factors that affect cognitive development and processing. External factors include the physical environment, the individual's bodily state, and the set of tasks being performed. Physical environmental factors, such as where the individual lives, provide opportunities for interaction or make it impossible to interact by placing boundaries on the task being performed. The physical environment may also affect the individual's bodily state. The physical environment does not directly affect the individual's cognitive architecture and ability to perform tasks; rather, the affect is indirect. Note: this is in opposition to those theories that implicitly or explicitly argue that the environment directly affects cognition (Blalock, 1967, Durkheim, 1951, Luria, 1978, Vygotsky, 1962, Vygotsky, 1978). The individual's bodily state sets boundaries on the task and may affect the individual's cognitive architecture. And, the task the individual is performing determines interaction opportunities and determines what part of the cognitive structure is currently salient to the individual (Carley, 1986a, Carley, 1981).

As the individuals perform tasks they interact, and in doing so, they acquire and transmit knowledge. What knowledge is acquired and transmitted is a function of what part of the cognitive structure is currently salient (Carley, 1986a). What knowledge individuals acquire and whom they interact with are thus factors which are intimately related. Task performance is a function of knowledge acquired todate (Laird, 1986a, Laird, 1986c, Carley, 1986, Carley, 1988d).

As individuals acquire knowledge they evolve their own language. Since knowledge is often acquired during interaction language is a socio-semiotic construct (Halliday, 1978, Whorf, 1956, Vygotsky, 1962). Since knowledge is acquired within the context of performing tasks the individual's language contains the articulation of experience. During the interaction both social evolution and language evolution occurs (Vygotsky, 1962, Vygotsky, 1978, Luria, 1978, Luria, 1981). Social knowledge becomes shared individual knowledge. Social language becomes the lossy integration of articulated individual knowledge. Thus language contains embedded social life achieved through consensus of the mass mind (Whorf, 1956).

At the level stated thus far the constructural theory is very vague. In order to increase the precision and predictive power of the theory it is necessary to have a model of individual cognitive structure.

2. A COGNITIVE BASIS FOR BEHAVIOR

In this section a possible model of cognitive structure is presented. This model includes a knowledge representation scheme, principles for modifying the knowledge base, and a set of operators for adding knowledge. A model of cognitive architecture is not presented. Possible models of cognitive architecture include SOAR (Laird, 1986a, Laird, 1986b, Laird, 1986c), the neural networks or parallel distributed processing models (Rumelhart, 1986, McClelland, 1986), and ACT* (Anderson, 1983).

2.1. The Knowledge Representation Scheme

The basis of the model for the cognitive structure is a fact based network structured knowledge representation scheme. In this scheme, information is modeled as divisible into separate pieces which are inter-related by shared concepts — see figure 3. Pattern matching on the concepts forms a network. A variety of representation schemes have been forwarded that share these basic characteristics: conceptual structures (Sowa, 1984), schemes (Anderson J., 1973, Bobrow, 1976), schemata (Rumelhart, 1976, Tesser, 1977, Tversky, 1980), structured frames (Minsky, 1975, Charniak, 1972), dynamic frames (Goffman, 1974), transition networks (Collins, 1975, Clark, 1977, Wyer, 1979, Bobrow, 1969), semantic nets (Simmons, 1973, Schank, 1973), and scripts (Schank, 1977, Abelson, 1976, Abelson, 1969). All such formalisms serve as models of information in memory and as devices for representing articulable

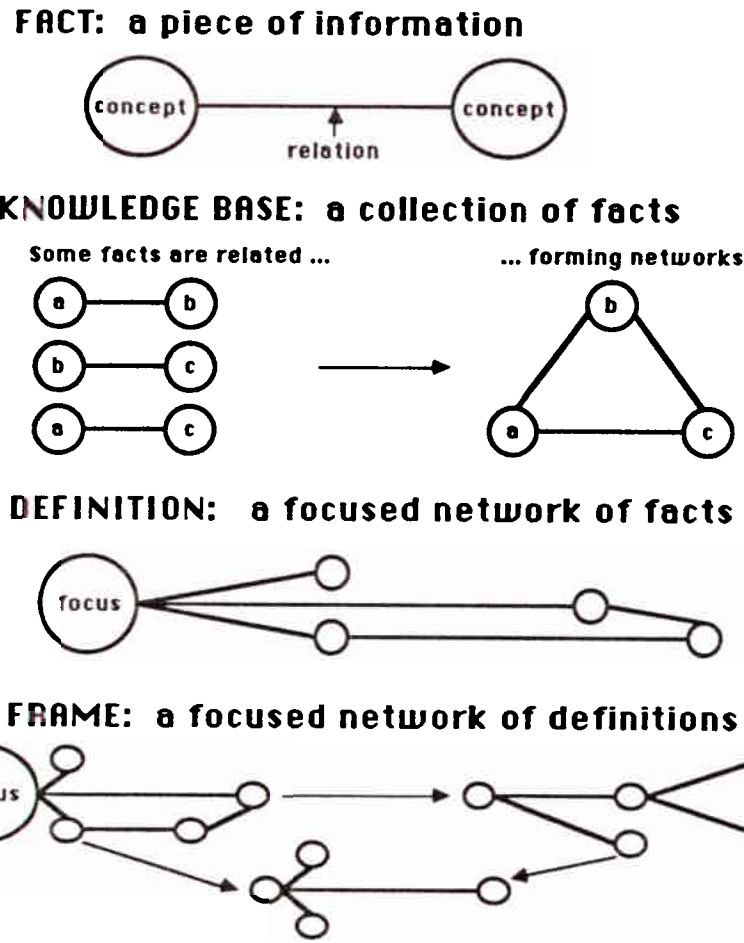

FIGURE 3 Concept, Facts and Frames

The basic item is the concept — shown as a circle. Concepts in isolation are meaningless. Concepts attain meaning by their relationship to other concepts. Two concepts and the relationship between them form a fact (top). Note: the relationship itself is a concept. A set of facts form a knowledge base (upper middle left). Within the knowledge base facts will be related, thus forming networks, because they share concepts (upper middle right). A piece of a network can constitute a definition of a concept (lower middle). The concept being defined is the focus or focal concept. A frame can be thought of as a network of definitions such that these definitions are at different levels of specificity (bottom). The frame serves as the collection of all the information, from some perspective, that is known about the focus or focal concept. The focal concept for a frame is the frame's handle. The focal concept can occur in other frames. In this case the frame of the focal concept is said to be embedded in the other frame.

knowledge. In the representation scheme described below a frame scheme is used due to the emphasis on semantic content rather than on evolving story lines.

A piece of information is modeled as a fact — two concepts and the relationship between them — (Carley, 1984, Carley, 1986b). This model follows from the Minsky notion of a fact (1975, p. 181): "a fact is a relationship together with a few things the relationship ties together in a meaningful way." The general notion of a linked pair of concepts as the representation of information is used by a wide variety of researchers both prior and subsequently to the Minskian articulation. Some examples of facts are a *Third Easter would be a good tutor, Gnerds aren't friendly*, and *Someone who fits in with the hall won't insist on quietness*.

A collection of facts is a *knowledge base*. The individual's knowledge base is that collection of facts that are known by the individual — refer to figure 3. Within a knowledge base, we expect to find that some of the facts will inter-relate; e.g. *Jay loves Ann, Ann loves Greg*, and *Jay and Greg* hate each other. Facts can inter-relate forming a network simply by sharing a concept — refer to figure 3. The base pattern matcher would be part of the cognitive architecture. Thus, pattern matching on concepts leads to the knowledge base being replete with implicit networks. Some networks may be directed or have a hierarchical structure as do inheritance networks.

Some networks of facts are more interesting or have special interpretations, e.g. definitions and frames — refer to figure 3. A definition is a focused network of facts such that the focus is the concept being defined and the other concepts in the network serve to define the focus by their relationship to it and each other. A definition is a particular interpretation of, a set of things that can be said about, or transformations of, the focal concept. A frame is a focused network of facts with a super-network structure imposed on the network of facts. A frame is thus a knowledge representation formalism in which a vast amount of complex information can be encoded in an interconnected fashion and from which that information can be effectively accessed. A frame is a data structure for representing a situation or objects as a structured network of nodes and relations such that the top levels are always true and the lower levels can be filled with the specific instances. Frames can contain multiple definitions where each definition serves as an alternate view of the focal concept. A frame can be thought of as the complete working definition of its focal concept.

The individual's knowledge base is composed of a series of inter-locked frames. Frames may be inter-connected, in that they can share concepts and facts. Each frame has a handle or name. This name is a concept which can be used to evoke that frame. Frames can embed within frames simply by having their handle in the other frames' networks. Thus, at one level, the individual's knowledge may appear as a simple network, and at another level, as a network of networks, and so on.

The individual has a frame for each task performed. This frame contains all of the knowledge that the individual knows relative to this task. There are frames other than task related frames. Such as generalized concept definitions.

Language, i.e. its semantic content, is a function of the individual's knowledge base. Thus, how language evolves is a function of the knowledge representation scheme. For the individual, knowledge and language are inseparable. Language is the articulation of individual knowledge. Thus, for the individual, language is the individual's working lexicon. The semantic content of language is the set of definitions

known by the individual. Because much of the knowledge in the individual's knowledge base is acquired during interactions with others the individual's knowledge base contains vast amounts of social and cultural information. That language contains social and cultural information has previously been argued (Collins, 1977, Charniak, 1972, Whorf, 1956, Vygotsky, 1962, for example). Meaning evolves as knowledge is acquired and these definitions change (Vygotsky, 1962). If knowledge acquisition by interaction is the most common form of knowledge acquisition for physical and social objects then the meaning the individual holds will be a product of social interaction (Blumer, 1969, Ruben and Weimann, 1979). In this way the individual's language will embody his position in the social structure, and the structure of social language will embody social structure (Halliday, 1978, p. 186). Rapidity of communication is a function of shared meaning. Thus, it is a function of the degree of similarity in the two individuals' knowledge bases.

A meaning-nominalist perspective is taken — meaning is a product of one speaker and one articulation (Grice, 1969, Bennet, 1976). Meaning, for the individual, is the explicit and implicit articulation of a concepts definition. All other types of meaning, e.g. the individual's average meaning, social meaning or shared meaning, are derivable from individual meaning.

2.2. Cognitive Principles

There are several basic principles that affect the nature of the knowledge base. These principles are: relative definition, immediate comprehension, saliency dependent access, monotonic changes, local consistency, and autonomy. These principles admit meaning, and place limits on knowledge acquisition behavior. These principles are not tied to the knowledge representation scheme described above.

Principle of Relative Definition:

The principle of relative definition seems necessary for assessing meaning in a network structured knowledge representation scheme. According to this principle:

- Concepts are ideational kernels that in isolation are devoid of meaning.
- Concepts can only be defined by their relationship to each other.
- Meaning is a function of the inter-relationships between concepts in the knowledge base and not the concepts *per se'*.

In terms of social usage and general communication, there are no absolute definitions only relatives ones. We do not go to the dictionary to understand a word; rather, we understand it through the context in which it is being used. The definitions in the dictionary are transient and dependent on the perception of tacit consensus to a set of information that collectively constitutes the meaning of a particular concept. The principle of relative definition leads to individual meaning, shared meaning, etc. being a matter of degree. This is in contrast to Luria's (1981, p. 44) argument that meaning is paradigm — "meaning is a stable system of generalizations represented by a word, a system which is the same for everyone."

Principle of Immediate Comprehension:

The principle of immediate comprehension seems necessary for defining mechanisms of knowledge acquisition. This principle provides a mechanism by which the assimilation and exploitation of new information is permitted by what information the individual already knows. According to this principle:

- Individuals will utilize new information to alter their cognitive structures if they can directly relate that information to information currently in their cognitive structures.

Given the knowledge representation scheme specified above, then according to this principle a new fact can be acquired only if it at least one of the concepts in it is already in the individual's knowledge base. A consequence of applying this principle is that the individual's knowledge base will form a single network. The principle of immediate comprehension, when restated in terms of the knowledge representation scheme, applies through the co-occurrence of concepts in the new fact and the facts in the current knowledge base. That is, if all facts in the current knowledge base are equally easy to access then the principle of immediate comprehension can be restated as:

- If a fact containing concepts α and β is communicated to the individual the individual will add that fact to his knowledge base provided that it is not already there and provided that at least one of the two concepts is currently in the individual's knowledge base.

Principle of Saliency Dependent Access:

The principle of saliency dependent access seems necessary to prevent processing overload. According to this principle:

- Access to information is a function of the degree to which that information is salient.
- The less salient the information the less likely it is to be accessed.

Given the knowledge representation scheme described above this principle can be interpreted. Not all of the information in the individual's knowledge base is always salient. Rather, at a specific point in time only one frame is salient. What task the individual is performing affects what part of the individual's knowledge base is currently salient. Specifically, the frame associated with the current task is the frame that is currently salient. This frame is the current frame. The level of access individuals have to the other frames in their knowledge bases is a function of the degree to which those frames are similar to the current frame.

The similarity between two frames can be measured in terms of the degree to which those frames inter-lock. The number of frames embedded in a frame is defined as the number of concepts in the current frame that have a meaning external to that frame. A concept has a meaning external to the frame in which it is embedded if that concept is a handle for a different frame and if that frame is not a subset of the original frame. Note, a concept will have an external meaning in the knowledge base if there is a fact that includes that concept in that knowledge base that is not in the frame. All frames

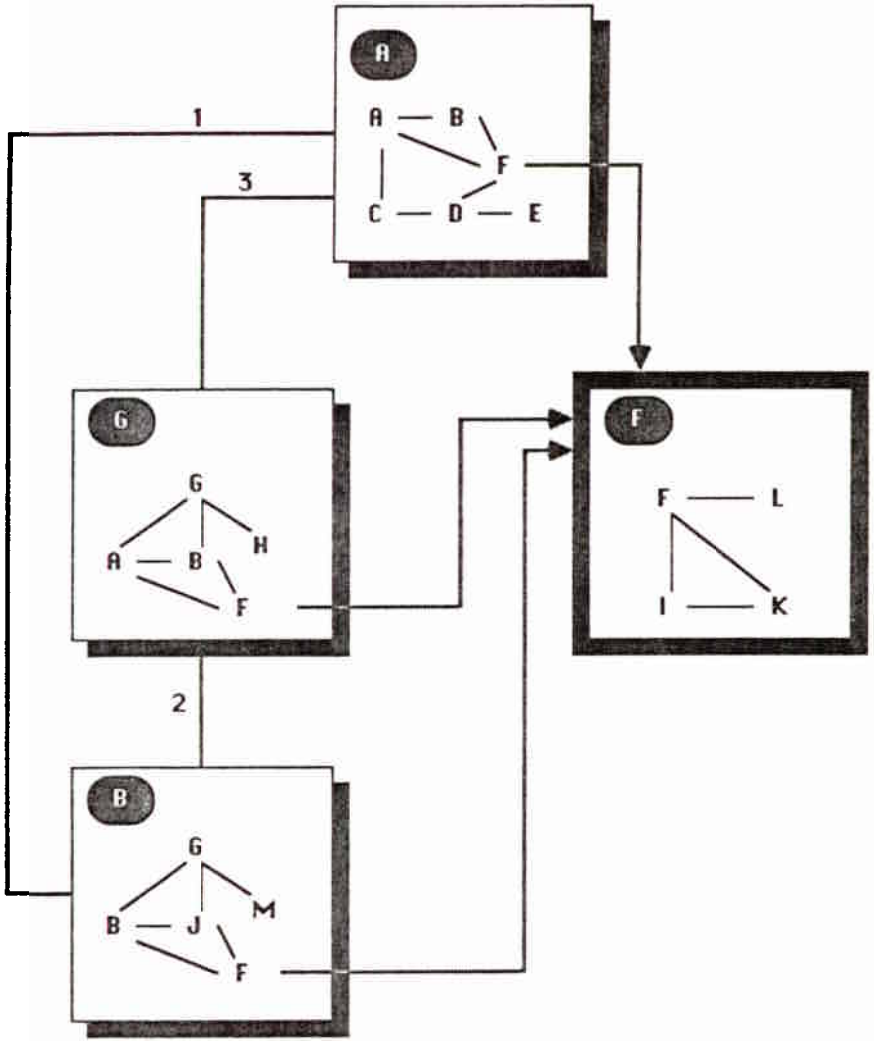

FIGURE 4 Inter-Locked Frames

The individual's knowledge base contains a network of frames. There are fours frames each represented as a square. Each frame has a frame handle. The frame handle is the letter in the top left corner in the shaded region. The frame handle is the focal concept which the frame defines. Frames are inter-related by sharing concepts and by sharing facts. The frame on the right, for concept *F*, is embedded in each of the other 3 frames shown. The shaded lines indicate embedding. That is, the frame on the right is the complete description of *F* and *F* occurs as a concept in each of the other frames. The three frames for, *A*, *G*, and *B*, are inter-locked. The solid lines indicate that the connected frames are inter-locked and the number next to the line is the degree of inter-lock.

embedded in the current frame are directly connected to the current frame. The degree of inter-lock between two directly connected frames is the number of facts that the two frames share — their intersection. Thus, emanating from the current frame is a network of frames, such that, the links in the network are weighted by the degree of inter-lock — see figure 4.

A path between a current frame and a fact is the set of linked frames from the current frame to the frame containing that fact. The strength of the path is a function of both its length and the degree of inter-lock between the frames for each link in this path. For example, strength can be measured as the sum of the inter-lock across all links in the path.

Given this model of frame inter-lock the principle of saliency dependent access can be redefined as:

- Access to a fact depends on what frame is current.
- Access to a fact is proportional to the strength of the path between the current frame and the desired fact.
- All facts in the same frame are equally accessible.

The principle of immediate comprehension limits what the individual *can* accept; however, it does not guarantee acceptance. Whether or not the individual actually accepts a fact is a function of the individual's ability to access those facts needed to apprehend the communicated fact. The individual, as an information processor subject to time and effort constraints is more likely to acquire knowledge, if that knowledge is immediately relevant. The principle of immediate comprehension must be modified by access. The principle of saliency dependent access makes this possible. Hence, this principle is one way in which limited cognition is incorporated into the model.

The individual's propensity for accepting a particular piece of information is proportional to the ease of accessing those facts which share a concept with the new fact. The individual's propensity for communicating a particular piece of information is proportional to the ease of accessing that fact. In a literal sense, the individual's current frame acts as a frame of reference — that body of knowledge that admits understanding and interpretation of new information. The individual's current frame thus serves several functions. The current frame makes it possible for the individual to understand that information which is directly related to it. The current frame inhibits the individual from understanding that information which is not directly related to it. In this way, the internal structure of the individual's knowledge base, i.e. the network of inter-locked frames, serves to inhibit knowledge acquisition.

Principle of Monotonic Changes:

The principle of monotonic changes may be necessary to ensure monotonic learning. According to this principle:

- The individual does not forget.

Loss of memory is viewed as the result of trauma, or biological dysfunction, which actually serves to block off or destroy stored information. The inability to recall

information, barring memory loss, is viewed as a function of access. Thus, at the individual level no mechanisms for deleting information from the knowledge base is proposed. The individual's knowledge and language evolves monotonically. A word attains a more complex meaning as the network of concepts linked to this word increases in density. The repertoire of sentences increases as the number of facts in the knowledge base increases. And so on. Knowledge acquisition thus becomes a heuristic act which irreversibly alters meaning.

Principle of Local Consistency:

The principle of local consistency seems necessary to ensure anything but random behavior. According to this principle:

- For a specific task, if the individual does not acquire new information the individual will perform the task in the same way.
- For a decision task, if the individual does not acquire new information relative to the decision the individual will continue to make the same decision.

A process is consistent if given the same input it produces the same output. The individual's cognitive structure is locally consistent if at two points in time the individual is performing the same task and the frame for that task has not changed the individual performs the task in the same way. This is one way in which at least limited rationality is ensured. Given the knowledge representation scheme suggested this principle means that a frame is always processed in the same way unless the frame itself is changed.

Principle of Autonomy:

The principle of autonomy is necessary to permit individual behavior. According to this principle:

- Knowledge acquisition by one individual does not affect any other individual's cognitive structure.
- Knowledge transformation by one individual does not affect any other individual's cognitive structure.

This basically rules out such factors as a mass mind.

2.3. Operators

Six operators relating to the following areas of individual behavior — knowledge acquisition, knowledge dissemination, and choice of interaction partner — are discussed. These are operators in the sense that they specify basic action mechanisms that individuals can take to modify their knowledge bases or to determine a behavior based on their knowledge bases. The basic procedure underlying these operators are discussed, however, specific implementation procedures are not. Thus the conditions under which this mechanism is viable, and the basic impact of the mechanism are specified, but the mechanism itself is not.

The list of operators discussed is not exhaustive, merely illustrative. For knowledge acquisition four operators are considered — by being told, discovery, generalization, and attribution. Each of these represents a mechanism by which individuals can acquire knowledge. There are other possible mechanisms such as analogy (Moore, 1974, Winston, 1977). There should also be other operators for dealing with choice of task. Such, however, are beyond the scope of this paper. There may also need to be operators that create garbled messages. In the current model, however, information transmission is clean and all miscommunication is a function of the individual's having different meanings and of not being able to comprehend the communicated information within the information processing cycle.

Knowledge Acquisition — By Being Told

Interaction between individuals leads to communication; hence, to the transmittal of information from one individual to another. During the interaction each individual communicates and acquires information from others by being told — see figure 5. Communication leads to frame creation and augmentation and if successful to the development of similar cognitive structures on the part of the communicators. These ideas are similar to Fauconnier's (1985) notions about communication and mental spaces.

If the individual is interacting with others then the individual can acquire knowledge by being told. Knowledge acquisition by being told is assumed to be a dominant way in which the individual acquires new information. Knowledge acquisition is dependent on who interacts with whom and what tasks the individuals are performing. If a piece of information is communicated to the individual the individual will acquire it if it is comprehensible within the information processing cycle. Comprehensibility is determined by the principle of immediate comprehension and access. The longer and weaker the access path the less likely it is that the path will be traversed during the information processing cycle. For knowledge acquisition the information processing style basically determines how long a new fact can be present in the individuals working memory without being a part of the knowledge base.

Knowledge Acqusition — Discovery

If the individual is interacting with himself then the individual can acquire knowledge by discovery. The result of a discovery is that the individual has a new piece of information. This piece of information may be new to the individual and the society or simply new to the individual. Regardles of how the discovery is made, it is assumed that the new piece of information, i.e. the discovery, is linked to information that the individual already knows. This corresponds to Simon's (1986) argument that creativity, or knowledge creation, involves the manipulation of new information as permitted by previously acquired information. Moreover, the only facts that can be discovered during a time period are those that link to facts which are accessible during the information processing cycle.

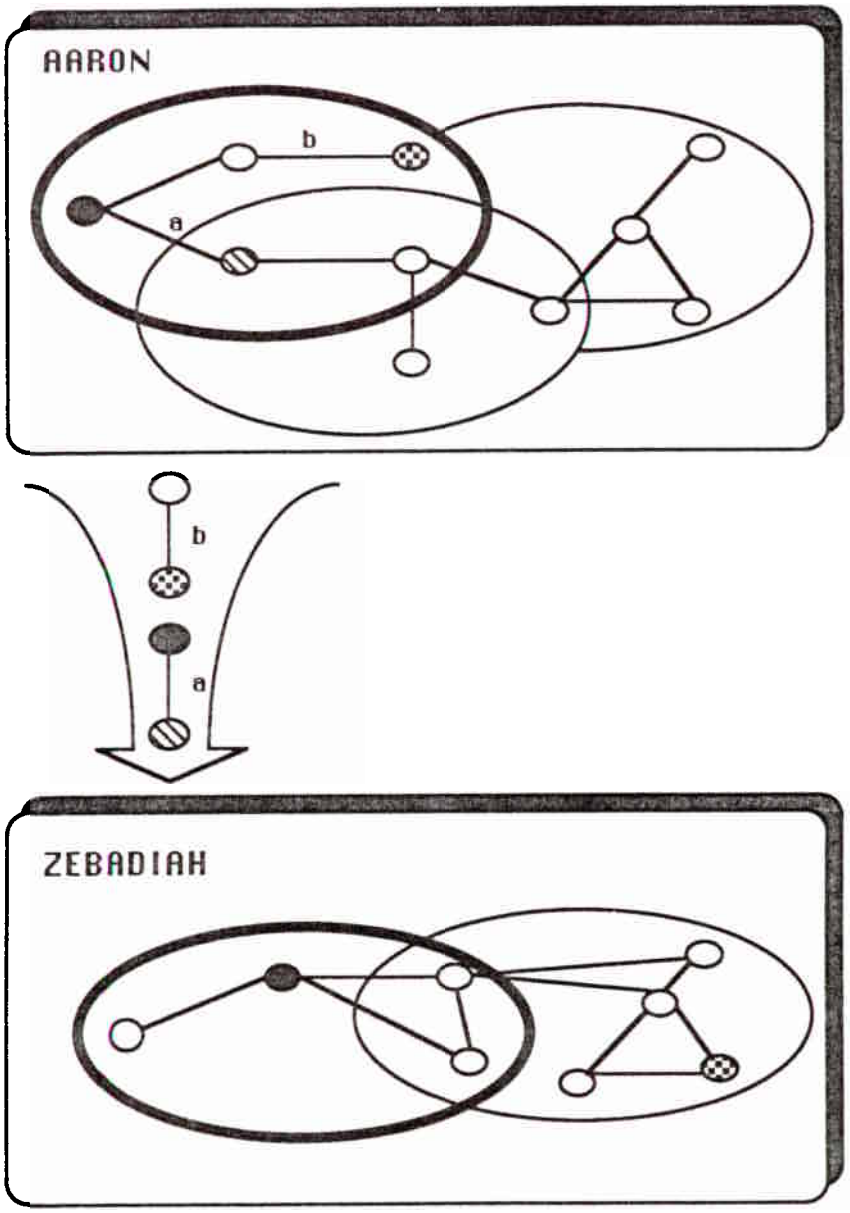

FIGURE 5 As Two Individuals Interact Information Flows

Aaron's knowledge base is shown at the top and Zebadiah's at the bottom. Frames are represented as ovals and the currently salient frame is the darkest oval. Concepts are circles and relations are the lines connecting the circles.

The two individual's, Zebadiah and Aaron, are interacting. Aaron is trying to communicate two different facts, *a* and *b*, to Zebadiah. Zebadiah can comprehend both facts *a* and *b* as one of the concepts in each of these facts is in his knowledge base. However, he is more likely to accept, and hence incorporate into his knowledge base, fact *a* and *b* as it is comprehensible using the information in his currently salient frame.

Knowledge Acquisition — Generalization

If the individual is interacting with himself then the individual can acquire knowledge by generalization. Generalization is a categorization process whereby x on the basis of a limited number of traits, perhaps one trait, is categorized as being similar to y. When a generalization occurs a new frame and a new concept is created for the generalized entity. The concept is the frame handle. A fact links this frame with each of the entities on which the generalization is based. This new frame has embedded within it all of the information known about the original entities. Generalization is a process of embedding one set of knowledge, e.g., a frame, within a new frame. Embedding is powerful as it provides a general framework which can be used to interpret outcomes, and which accurately predict the general or average nature of those outcomes.

Knowledge Acquisition — Attribution

If the individual is interacting with himself then the individual can acquire knowledge by attribution. Attribution is a process of inferencing whereby x because it is judged to be similar to y inherits all of y's properties. Attribution creates knowledge by taking that knowledge which is true for a generalized quantity and attributing that knowledge to those concepts that are special cases of the general. Attribution is thus a process of frame elaboration.

Knowledge Dissemination: As individuals interact, some of the facts in one individual's knowledge base may be communicated to the other individual who may then incorporate those facts into his knowledge base. Knowledge communication is dependent on what task the communicator is performing. Barring time and effort constraints all facts in a frame are equally likely to be communicated. Otherwise, what fact is communicated is proportional to the strength of the path to that fact. The longer and weaker the access path the less likely it is that the path will be traversed during the information processing cycle. For knowledge communication the information processing cycle basically determines which facts can be communicated.

Choice of Interaction Partner:

Who the individual interacts with is a function of both interaction opportunities and the individuals' propensity to interact with each other. The task determines what interaction opportunities are available. Two individuals can not interact unless they have the opportunity to do so. Given that they can interact, individuals are more likely to interact, i.e., they will have a higher propensity to interact, if they are more socio-cognitively similar. How socio-cognitively similar two individuals are is a function of how much information two individuals share relative to how much information they share with everyone else in the society. See (Carley, 1987a, Carley, 1988c) for possible models.

3. IMPLICATIONS AND EXPLANATIONS DRAWN FROM CONSTRUCTURALISM

When the above model of cognition is used as a cognitive basis for the constructural theory it is possible to derive implications for and explanations of a variety of social phenomena. In this section the following phenomena will be looked at: development of shared knowledge, identical behavior by members of the society, foreign language acquisition, clique formation, civil disobedience, and diffusion of innovative information. The purpose of this section is to exhibit the breadth of explanation possible when theories of social behavior are cognitively motivated. This section is not an exhaustive list of the predictions of constructuralism. Where possible, results from other studies that support the predictions are noted.

3.1. Knowledge Acquisition

Recall that the individual can acquire knowledge both through interacting with other individuals and through various cognitive processes such as discovery, generalization, and attribution. A variety of predictions can be made about the state or expected state of the individual's knowledge base. These include:

- Since there is no mechanism in the model by which the individual can forget information; it can be concluded that the individual's knowledge base will increase over time. In (Carley, 1987a) it was shown that on average individual's knowledge bases did increase over time.
- The less the individual interacts overall the less knowlede the individual will share with the typical member of the society.
- If discovery is relatively rare then the less the individual interacts the smaller the individual's knowledge base.

A society is recurrent if there is an interaction path between any two members of the society; i.e. if there is a direct or indirect link between each dyad in the society. In a recurrent society information can be transmitted between any dyad even if that information has to be passed by many third parties. Within a particular time window, i.e. number of time periods, a society is recurrent if it is possible for a piece of information communicated by one individual to reach all other individuals in the society. The rate of turnover in the society is a function of the rate at which individuals enter and leave the society. The length of time until recurrence relative to the rate of turnover and the rate at which information is discovered will determine the level of shared knowledge and consequent homogeneity of behavior in the society. Due to the interaction between these time period, few predictions can be without simulation. It can, however, be predicted that:

- If there is no turnover, no discoveries, and the society is recurrent then the level of shared knowledge will increase over time. In (Carley, 1987a) these conditions were met for the group looked at and their shared knowledge did increase over time and the standard deviation across their knowledge bases decreased over time.

- Since knowledge acquisition and dissemination is a function of what tasks the individuals are performing, shared experiences will lead to an increase in shared knowledge. In (Carley, 1987a) it was shown that the increase in shared knowledge was due to the acquisition of knowledge that was all learned in one particular shared task.
- Since interaction leads to communication and shared knowledge then, *ceterius paribus*, the higher the level of interaction the higher the level of shared knowledge. In (Carley, 1987a) it was shown that on average, individuals with higher interaction rates exhibited higher levels of shared knowledge.

3.2. Identical and Similar Behavior

Given that individuals have identical cognitive architectures and the principle of local consistency it follows that if two individuals have the same knowledge they will behave in the same way. As individual's knowledge bases become more similar their behavior is expected to become more similar. Consider the following predictions:

- The more similar the two individuals' knowledge acquisition histories the more similar their behavior. Identical twins growing up together should exhibit more similarity in behavior than other siblings.
- The smaller the society the greater the chance that all individuals will share the same information. Thus, it is expected that the smaller the society the more similar the behavior of the members of the society.

Given the network structured knowledge representation scheme it follows that if two individuals are performing the same task what behavior they engage in is a function of the inter-relationship among the information in their respective knowledge bases. That is, behavior is not a product of the amount of knowledge known by the individual but of the way that knowledge is linked together. Making the same decision is a type of consensus in behavior. Consensus in behavior should increase as the pattern of knowledge shared by the individuals increases. With respect to decision making the following predictions can be made:

- If two individuals' frames for a decision are identical they will make the same decision.
- The fewer deviations possible in the behavior the more similar the patterns of knowledge need to be to produce similar behavior. That is, if a behavior is binary — do x or don't do x — then for two individuals to produce the same behavior they will need more similar patterns of knowledge then they will if the behavior is higher order. Thus, two individuals need more similar patterns of knowledge to produce the same yes/no decision on a candidate then they do to produce a similar rank ordering of candidates.
- The more two individuals interact, the more knowledge the two individuals will share. However, if the individuals are not performing the same task at the time of interaction this increase in shared knowledge may not lead to similar patterns of shared knowledge. Hence, such increases in knowledge will not necessarily lead to the individuals making the same decision.

- The more two individuals interact relative to a decision the more similar their decision will be. That is, the more two individuals interact while performing the same task the more likely it is that they will develop similar patterns of shared knowledge *vis-a-vis* that task.

In (Carley, 1986b, Carley, 1986a), using a cognitive model similar to that previously described, it was shown that the level of interaction is correlated with the level of shared knowledge. This has also been shown with simpler cognitive models (Carley, 1987a, Carley, 1987b). In (Carley, 1987b) it was shown that the level of interaction is not correlated with whether or not the same decision was made. It was also shown that similarity in the pattern of knowledge known by the individuals is correlated with making the same decision. And in (Carley, 1984) it was shown that the more knowledge individuals share the more likely it is that they will produce similar rank orderings of candidates despite not having the same first choice.

Given that individuals have identical cognitive architectures and the principle of local consistency it follows that if the researcher has an accurate model of the cognitive architecture and knows the state of the individual's cognitive structure at the time that the individual determines a behavior the researcher can reproduce what behavior is chosen mechanically. Decision making is a behavior. Consequently, what decision the individual will make can be determined mechanically given a model of the cognitive architecture and provided that the current state of the individual's cognitive structure is known. In (Carley, 1987b, Carley, 1984) this point was demonstrated.

For a recurrent society mass communication decreases the time window necessary for recurrence. Thus, mass communication leads to more individuals sharing the same information at the same time. This leads to the following predictions:

- In a current society, the higher the average interaction level the greater the level of shared knowledge and homogeneous behavior, *ceterius paribus*.
- In a current society, the greater the level of mass communication the more homogeneous the behavior of the members of the society.
- Mass communication techniques can lead to homogeneous behavior despite high rates of social turnover.
- Mass communication techniques can lead to homogeneous behavior despite high rates of discovery.

3.3. Communication

Given the principle of immediate comprehension it follows that as two individuals' knowledge bases become more similar the propensity of one individual to accept the facts communicated by the other increases.[4] This is a direct function of the number of shared concepts. Note: a communication is considered successful if the receiver adds the information communicated to his knowledge base. The following predictions can be made:

- The higher the level of interaction the more likely it is that communication is successful.
- Communication is more likely to be successful if the individuals are performing the same task.

- If two individuals share few concepts then their initial communications will lead first to the development of a shared vocabulary, then shared facts, and finally to shared definitions.
- If two individuals share a definition then part of the meaning of the focal concept is the same for both individuals. Thus, the higher the level of interaction the larger the number of concepts for which the individuals have shared meanings.

When an individual communicates a fact the definition or meaning of the concepts is implied. Whether or not the same meaning is inferred depends on whether the receiver's salient knowledge base is similar to the sender's. Rapid communication occurs if the sender and receiver have similar meanings and if the sender expects there to be shared meanings. In this case, much of the information is left implicit. Thus, the rapidity of the communication is a function of actual and perceived shared knowledge. These predictions follow:

- If the task requires communication, and if the individual is subject to time and effort constraints, then the individual is more likely to interact with those individual's with whom he expects to share more knowledge.
- Since individuals may have different levels of cognitive similarity for different frames it follows that the rapidity of their communication is a function of what tasks they are currently performing. It follows that if a task is performed over a sequence of time steps the rapidity of communication *vis-a-vis* that task will increase.
- The higher the level of interaction the greater the rapidity of the communication. It is often noted, at least anecdotally, that members of the same profession can communicate vasts amount of information with only a few words when discussing profession related tasks.
- The higher the level of interaction the more likely it is that shared knowledge will not be explicitly communicated. In (Carley, 1984, Carley, 1986b) it was shown that among the members of a highly interactive living group there was a high degree of shared knowledge that was not explicitly communicated.

Given that individuals have identical cognitive structures and the principle of local consistency it follows that what knowledge is communicated is a function of expected shared knowledge. Hence miscommunication is a function of the difference between expected shared knowledge and actual shared knowledge. Consider the following predictions:

- If an individual expects the listener to share a particular piece of knowledge, and if the individual is subject to time or effort constraints on communication, the individual will be less likely to communicate those facts that are expected to be shared.
- Miscommunication will occur if the individuals have different definitions. In (Carley, 1984, Carley, 1986a) it was noted that miscommunication did occur under these circumstances.
- Miscommunication is less expected as interaction levels increase. In (Carley, 1984) the students were surprised that miscommunication did occur. In

particular, those dyads with the highest overall interaction rates were the most surprised when they miscommunicated — i.e. the couples.

- Miscommunication will be more unexpected if the individuals are engaged in a task that is viewed as similar to a task where miscommunication is unlikely to occur by at least one of the individuals. In (Carley, 1986a) it was shown that one couple that had very similar views in terms of education and study habits did not view the need to have a tutor who was academically qualified in the same way. For one of the individuals the task of choosing a tutor was similar to the various academic tasks do to the perceived need to have a tutor who was academically qualified. This couple did not expect the miscommunication that did occur to occur.

3.4. Education

In terms of education, the principle of immediate comprehension suggests that the socratic method, or any other method in which new knowledge is built by reliance on old information will be more effective than a technique based on memorization. Consider the following example — see figure 6. The teacher and student 2 for frame *A* both have facts *A-E*, and *E-D*. The teacher wants the student to know learn fact *B-C*. In order for the teacher to get the student to learn fact *B-C* the teacher will first have to communicate fact *A-B*, or the fact *E-B*. Only after the student knows about *B* can he taught about *C*. Whereas, with student 1 the teacher has the harder task of first getting the student to create a frame for *A*. And then the teacher will need to communicate *A-B* and then *B-C*. Since student 1 already knows about *E* the teacher may try to first establish the connection between *E* and *A* as a way of establishing the frame for *A*.

The principle of immediate comprehension, when applied to the educational process, suggests that education must proceed in an incremental step wise fashion. It also suggests that in order to be maximally efficient, i.e. to maximize the new information and minimize the old information provided, the teacher should begin by determining what information is in the individual's knowledge base. Then, the teacher should contrast the students knowledge base with the desired knowledge base locating the gaps or differences in the two. It is interesting to note that this is the procedure followed in certain computer aided tutoring systems (Stevens, 1978).

Generalization and attribution may lead the student to make the inferences that are incorrect from the teacher's point of view. The way in which information is provided will affect what knowledge is learned by generalization and attribution. Winston (1975, p. 157–209), using a network structured knowledge representation scheme, showed that the order and nature of the examples, i.e. information, provided determined what was learned by generalization and attribution (Winston, 1975).

Under the constructural theory direct communication by the teacher is not the only mode of education. Recall that what facts the individual acquires is a function of with whom the individual interacts. Thus the way the teacher establishes learning or research groups will affect information dissemination. Secondly, shared experience occurs as two individuals perform the same task. A shared experience leads to shared knowledge because it both provides an opportunity for interaction and it makes

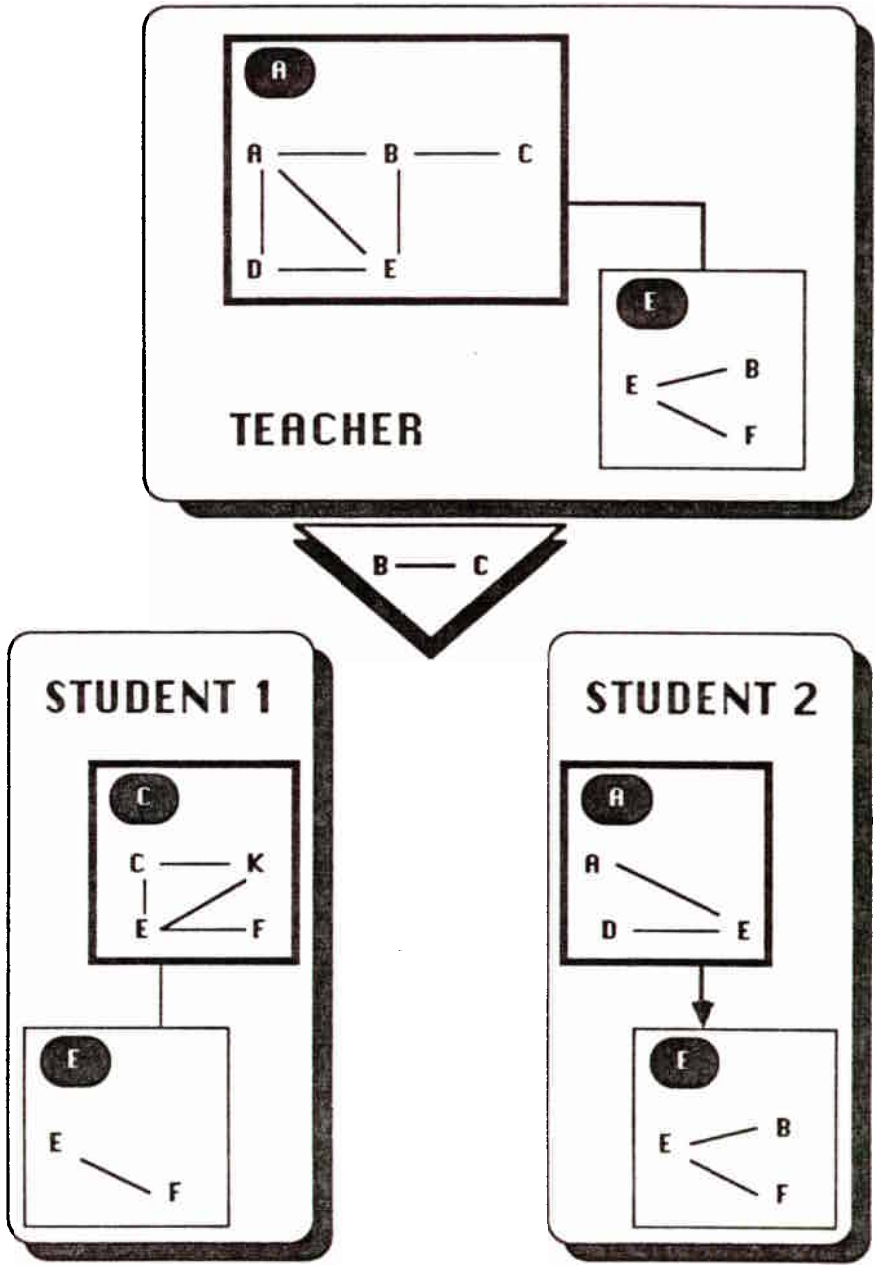

FIGURE 6 Educational Process

The knowledge bases of two students and a teacher are shown. The teacher, top, has two frames, represented as squares. The frame, *A* is the one which is currently salient for the teacher. For student 1 frame *M* is currently salient, and for student 2 frame *A* is currently salient. The teacher is trying to communicate fact *B-C*, in the triangle. Relative to the teacher's knowledge base there are gaps in the students' knowledge bases. That is, the students' are missing certain facts. For example, fact *B-C* is not known by either student and is known by the teacher. Neither student can immediately comprehend the fact *B-C*.

salient for all of the individuals a frame relating to the topic of the experience. While the individuals' frames may not be identical, i.e. they may each view the experience somewhat differently, those experience related frames will share more facts than will any random set of frames. This leads to the following predictions:

- Shared experience increases the likelihood that information communicated will be accepted.
- Individuals performing the same task are more likely to acquire new information from each other than are individuals performing different tasks, provided there is new information to acquire.
- Shared experience leads to high rates of learning.

Given that the higher the rate of interaction between two individuals the more knowledge is shared by the two individuals, and that performing the same tasks increases the likelihood that communication will be successful, it follows that over a set of sequential time periods where the individuals continue to perform the same or similar tasks, the rate at which new knowledge is acquired will increase. Thus an intense experience, performing the same task or a highly similar set of tasks over a large number of sequential time periods will lead to shared knowledge, shared patterns of knowledge, and similar behavior. A rate increase can be achieved due to attribution and the way in which frames can embed within each other. That is, at first the rate of knowledge acquisition is slow as the communication of a fact leads to the acquisition of 1 fact, and perhaps a new concept. However, as definitions are shared, the communication of 1 fact leads to the acquisition of a set of facts linking the definitions. Thus, research groups, think tanks, project groups, and so on, are expected to be educationally effective. If there is a limit to how much knowledge is associated with the task then the rate of knowledge acquisition under these conditions will over time form an ojival function.

The more knowledge that is shared or social the less likely it is that one individual will have to "educate" the other in order to communicate a particular fact. Thus, if the individuals start out with a large amount of shared knowledge, the need for this educational process in order to communicate becomes unnecessary as the requisite facts will already be shared. Thus, within the social unit, education serves to increase the set of shared knowledge, and to decrease the need for explicit communication.

3.5. Language Acquisition

Recall the prediction that: if two individuals share few concepts then their initial communications will lead first to the development of a shared vocabulary, then shared facts, and finally to shared definitions. Social language can be thought of as the articulation of shared knowledge that persists over time (Carley, 1987c). Thus social language is continually constructed as knowledge is articulated. Since people enter and leave the society, thus bringing and taking knowledge and hence language with them, social language construction is a lossy integration process. Lossy integration refers to an integration process in which information is lost, e.g., as when taking a moving average. This leads to the following predictions:

- Social language has embedded within it vasts amount of social and cultural information. This has repeatedly been shown to be the case (Collins, 1977, Charniak, 1972, Whorf, 1956, Schank, 1981, for example).
- Social language, i.e. its semantic content, will evolve as does shared knowledge — first vocabulary, then shared facts, then meaning.
- The principle of immediate comprehension leads to the prediction that an individual can learn the new words of the new language only by relating it to current knowledge.
- Since, different languages have different meanings associated with similar concepts a method of teaching a new language by correlating the words of the new with the words in the old will not be as effective as one which also teaches the meanings of the new words.
- Since knowledge embeds, learning a language involves learning the currently embodied history and culture of the society which produced the language. In this way socialization is a product of learning a language (Halliday, 1978, Luria, 1981). Or as Lock (1980, p. 190) argues "through the guided reinvention of language the child comes to share in his culture's explicit perception of the world."
- The higher the rate of social turnover and the higher the rate of discovery the greater the rate at which language will change.
- A period containing many shared experiences relative to a set of related tasks is an intense educational opportunity. If you couple this with the model of language it suggests that the best way to teach language is by placing the sudent in a culture that speaks only the desired language for a period of time. Note: living in a foreign country is commonly thought to be the best way to learn a foreign language.
- Since shared meaning is constructed as individuals communicate information during task performance intense periods will lead to the development of shared meaning. Thus, individuals involved in such an experience should develop their own language. Note: it is commonly noted that children at camps, highschool groups, scientific sub-specialities campus communities, and so on, develop their own language.
- If there are more students than teachers then new meanings will evolve faster than old meanings will be transferred from teacher to student.
- Because shared experience generates shared patterns of knowledge, those words which are used to describe a historical event that was experienced by the current members of the society will have larger shared meanings than will words that do not correspond to such events. In (Carley, 1987c) it was shown that historically based concepts do stand out as having some of the most dense meanings, i.e. their network position is dense.

Whorf (p. 156) argues that language "limits free plasticity and rigidifies channels of development." Thus, at the individual level language should restrict and perhaps even determine the way in which the individual processes information. At the social level language should limit the development of society and culture. Constructuralism provides a cognitive basis for this effect. Under constructuralism the individual learns

a language by learning the meaning of the concepts in the vocabulary; i.e. the individual learns a language by constructing his cognitive structure. Given the principle of immediate comprehension, what knowledge the individual has limits what knowledge the individual can immediately acquire and determines what order the individual must acquire knowledge in. Since individual language is the articulation of individual knowledge the principle of immediate comprehension leads to the result that language limits and rigidifies cognitive development. That is, if there is no way of expressing a certain piece of information in a language the individual who only knows that language can never acquire that fact.

A secondary point is that while language rigidifies development it does not directly affect cognitive processing. Luria (Luria, 1978) argues that language restricts individual cognitive processing. Under the constructural theory language can not affect the individual's cognitive architecture and basic cognitive processing. Language can, however, affect cognitive processing that is knowledge dependent.

Since language rigidifies cognitive development learning multiple languages should decrease the individual's cognitive rigidity. Thus for the individual who knows multiple languages the following predictions can be made:

- The more languages the individual knows the larger the individual's knowledge base.
- For the individual who knows several languages generalization and attribution will cause the individual's knowledge base to contain more information than simply the sum of the facts needed to express the semantic content of the languages.
- The more languages the individual knows the more likely that the fact being communicated to the individual will be immediately comprehensible by the individual.

The acquisition of a new language leads to the acquisition of the social knowledge and norms of a different culture. Consequently:

- The individual that knows more than one language is less likely to ascribe to the norms of a single culture than is the individual who only knows one language.
- If an individual has a primary and a secondary language then the individual is more likely to ascribe to the norms of the culture associated with the dominant language.

3.6. Expected Behavior

Given that individuals have identical cognitive architectures and the principle of local consistency it follows that if an individual thinks that another individual has the same knowledge as he does the individual will expect the other individual to behave in the same way he does. Given the principle of generalization it follows that an individual will view himself as similar to another individual if they share some trait. Generalization can only take place using the information that is accessible. Thus, individuals will view themselves as similar if they have the opportunity for interaction and are performing the same task. Given the principle of attribution it follows that if an individual views himself as similar to another individual the individual will attribute to the

other individual the knowledge that the individual has. Expected behavior is thus a function of perceived shared knowledge. Consider the following predictions:

- The higher the level of interaction the greater the expectation of shared knowledge.
- The higher the level of interaction the greater the expectation of similar behavior.
- Since interaction leads to shared knowledge, if two individuals interact they are more likely to expect the other to behave similarly to one's self than if the two individuals never interacted.

3.7. Cliques

A clique can be defined as a group of individuals who interact frequently and have a sufficiently high level of shared knowledge *vis-a-vis* a set of tasks that they exhibit shared behavior. Consider a society in which all of the individuals are working on a single task and the individuals in the society can be divided into two groups. In this type of an environment cliques should be a fairly stable and often occurring social formation at least in the short run. Note: in the long run, given a recurrent society, at quiescence all individuals will know everything and the entire society will form a single clique. It is useful to refer to three types of cliques — the clique, the knowledge clique, and the interaction clique. The clique is a network clique formed using both knowledge and interaction ties. The knowledge clique is a network clique formed using knowledge ties. The interaction clique is a network clique formed using interaction ties. The following predictions were derived *via* simulation using a slightly simpler knowledge representation scheme:

- The higher the level of shared knowledge and interaction the more likely the group is to remain a clique.
- If a group has a high rate of interaction but does not have a high level of shared knowledge, i.e. is an interaction clique, the group will over time decrease its cliquishness in interaction.
- If a group has a high level of shared knowledge but does not have a high rate of interaction, i.e. is a knowledge clique, the group will over time increase its cliquishness both in knowledge and interaction. In this sense, cliques are a natural product of cognition.
- Segregation of information will cause individuals to stop interacting, despite previous high levels of interaction.
- Differences in interaction rates across the members of a society will decrease as individuals share more information.
- You can force two distinct cliques to become less cliquish relative to each other by providing both groups with the same information.
- You can promote clique formation by providing different members of the society with different information.
- It takes only a small difference in information, relative to the task, to generate cliques.
- Information diffuses more rapidly within a clique than within non cliques.

- On average, the more cliquish two groups are the less likely information is to diffuse between them within a given time span.
- On average, the more cliquish two groups are the longer it will take for information to diffuse between them.

3.8. Innovation and Diffusion

Innovation is the result of information discovery. Diffusion is the process by which innovations spread through, i.e. are accepted by, the members of the society (Rogers Shoemaker, 1971, Rogers, 1982). The constructural theory can be used to look at innovation in terms of the discovery of new facts and diffusion in terms of the communication of a particular fact. It follows from the theory that since information can only be discovered when the individual is interacting with himself, some isolation is necessary for innovation to occur. Secondly, since information can be communicated between individuals only when they are interacting, interaction is necessary for the diffusion of newly discovered or innovative information. The following predictions can be made:

- Given that all facts in a frame are equally likely to be communicated, if there are no time and effort constraints, then the more the individual works on a particular task the more likely it is that innovative information relative to the task will be communicated.
- The more often two individuals interact, *ceterius paribus*, the more likely it is that a specific piece of information, the innovative piece of information, will be communicated.
- The higher the average rate of interaction the greater the rate of diffusion, *ceterius paribus*. In (Carley, 1987b, Carley, 1987e) it was shown that for a group of scientists that those sub-groups with the highest average interaction rates were in general the first to which the information diffused.
- The higher the average rate of interaction the greater the likelihood that information will diffuse. Proximity has been shown to affect rate of interaction and hence diffusion (Festinger, et al., 1950a, Festinger, et al., 1950b). In (Carley, 1987e, Carley, 1987b) it was shown that for a group of scientists the higher the average level of interaction the more likely information was to diffuse between the scientists.
- If the society is not recurrent complete diffusion can not occur without rediscovery.
- If the society is recurrent, relative to the relevant tasks, complete diffusion will occur.
- If the society is recurrent, mass communication will increase the rate at which diffusion occurs. In (Carley, 1987e) it was shown that the existence of mass communication tools made it possible for information to diffuse at a rate higher than would be expected if such mass communication tools did not exist (Carley, 1987e).

Imagine a society formed of three groups — a central core, a periphery to that core, and a third group of basically unrelated individuals. Let the core be a network clique

in both interaction and shared knowledge. Let the periphery have on average equal shared knowledge and interaction with the core and the third group. And let the third group have low internal levels of shared knowledge and interaction. Under these conditions the constructural theory predicts that:

- For innovations that occur in the third group on average members of the periphery will adopt the innovative information before the members of the core.
- Once innovative information has been adopted by a member of the core that information will rapidly diffuse within the core.

Further predictions about more complex structures should be derivable by simulation.

3.9. The Strength of Invisible Colleges

An invisible college is a group of scientists that have both shared paradigms, i.e. shared knowledge, and person-to-person contact (Price, 1963, Price, 1965, Price and Beaver, 1966, Crane, 1970). Price viewed these informal networks to be a product of the growth of the scientific community. That is, invisible colleges formed due to a communication crisis that occurred as more individuals became involved in scientific pursuits. Under the constructural theory invisible colleges are a cognitive byproduct. Specifically, the feedback relationship between shared knowledge and interaction leads to the development of invisible colleges given the mentor relationships within and the cumulative nature of science. Consider the mentor relationship in PhD programs. The mentor relationship leads to the transfer of information and shared knowledge through intense periods of doing the same or similar tasks. The cumulative nature of scientific research leads to the existence of sequential tasks that share large amounts of information.

The effect that for two individuals the increase in shared knowledge with each other relative to the knowledge they share with other individuals leads to an increase in interaction is exacerbated in science due to the nature of the knowledge learned. That is, in science the information transferred from the mentor to the student is not just a set of independent facts but a set of interconnected facts and facts for organizing facts. Entire frames are transferred. Thus scientific discourse leads to the creation of vast amounts of task related shared knowledge. This in turn leads to high interaction. If the mentor has not one but several students, and if these students are expected take part in shared tasks, e.g. by taking courses together or being part of a research group, then it would follow that the greater the intensity of the shared experience the more likely it is that when these students move to other universities they will maintain their contacts with each other and their mentor thus forming an invisible college. Given the cumulative nature of science, the vast amount of knowledge transferred from mentor to students will be sufficient to maintain initial interaction contact once the students move to other universities. This is because much of the knowledge from former tasks is related to those future tasks in which the individuals will be engaged. Shared future tasks such as joint research projects will serve to maintain the initial link.

The constructural theory thus leads to the following predictions vis-a-vis the development of invisible colleges:

- On average the members of an invisible college are more likely to be related by mentor-student ties than would be a random group of scientists.
- On average the members of an invisible college are more likely to be related by extreme opportunities for contact, e.g. were at the same school, in the same department, etc., than would be a random group of scientists.
- Invisible colleges will form regardless of the existence of a communication crisis. A communication crisis could be defined to exist if the amount of information on the topic around which the invisible college formed was greater than some level.
- Members of an invisible college will develop their own "language" for dealing with the scientific phenomena that they study.
- The more cumulative a scientific specialty the more likely an invisible college will form.

Invisible colleges have been shown to have a strong impact on the diffusion of scientific information (Price, 1963, Price, 1965, Price and Beaver, 1966, Crane, 1970). The basic mechanism postulated rests on the existence of direct and indirect information flows (Crane, 1970):

> It is not necessary to know a particular member of a social circle to be influenced by him. Not only can a scientist be influenced by publications written by authors whom he has never met, but he can also receive information second-hand through conversation or correspondence with third parties.

Such direct and indirect interactions and communications, however, are not unique to invisible colleges.

That invisible colleges should play such a strong role in science can be explained by the constructural theory. An invisible college would be a group of scientists who, relative to a random group of scientists, have on average high interaction levels and high levels of shared knowledge. It is not the existence of direct and indirect ties that lead to diffusion but the rate of interaction and the density of shared ties relative to the shared research area that affects diffusion. If such a group exists in a society the constructural theory would lead to the following predictions about its behavior:

- Innovative information uncovered by a group member will rapidly diffuse within the group. This is true whether the innovative information was actually discovered by the individual or was learned by that individual from an individual outside of the group.
- The higher the level of interaction, relative to the research area, the more rapid the diffusion.
- The higher the level of shared knowledge, relative to the research area, the more rapid the diffusion.

An invisible college basically forms a clique within the boundaries of science. That is, if only interactions and knowledge relative to scientific discourse are counted then an invisible college forms a high shared knowledge high interaction level group, i.e. a clique. Consequently, the predictions about cliques apply.

Note: not all of the consequences of being a member of an invisible college are good. As with cliques, group membership increases the rapidity of diffusion within the group but decreases the likelihood of acceptance of information from outside the group.

3.10. Civil Disobedience and Reinterpretation of Laws

Under the constructural theory a law involves the articulation of a perceived social form. A norm is the normal or average behavior that is performed given a particular task. Norms are constructed as shared knowledge is constructed. Recall that if individuals have identical cognitive structures then given the principle of local consistency it follows that individuals will exhibit the same behavior for a particular task if they have the same pattern of knowledge *vis-a-vis* that task. Thus, norms and the laws corresponding to them will evolve as shared knowledge increases. The law will describe the task and the associated ascribed behavior, however, it will generally not describe the knowledge that leads to that behavior. This leads to the following predictions:

- Laws will be seen as unjust or inappropriate if the perceived norm is not the actual norm.
- Civil disobedience, i.e. not following a law because it is felt to be unjust, is likely to occur if the knowledge supporting the ascribed behavior is no longer shared.
- In order to get individuals to obey laws it is not necessary to educate them as to the law. It is, however, necessary to educate them as to the knowledge relative to the task that will produce that behavior.
- If the society is not recurrent civil disobedience will occur.
- For a recurrent society, the higher the average interaction level the less civil disobedience, *ceterius paribus*.
- In a recurrent society, the greater the level of mass communication the less civil disobedience, *ceterius paribus*.
- The higher the rate of social turnover the more likely civil disobedience is to occur, *ceterius paribus*.
- The higher the rate of social turnover, the higher the rate of discovery, the greater the need to re-interpret, repeal, or add laws. Note if laws are altered by re-interpretation or addition rather than repeal the rate of law creation will be a function of the rate at which social knowledge changes.

3.11. Can Natural Languages be Created

Individuals acquire knowledge while performing tasks. Thus embedded within the individual's knowledge base is the articulated representation of experience. Individual language is the articulation of individual knowledge. Semantic content is a function of the definitions in the individual's knowledge base. Since definitions can accrue in the context of forming tasks semantic content can contain descriptions of experience.

When one thinks of a language, however, one is generally thinking of social language rather than individual language. Social language, i.e., its semantic content, can

be modeled as the lossy integration of the articulation of individual knowledge (Carley, 1987d). Individuals will share knowledge if they participate in similar experiences or if they communicate. Consequently, embedded within the social language are the articulations of shared experience. Typically, constructed languages do not have this feature.

This suggests that natural languages can not be constructed by simply generating a minimum set of words and grammatical rules. Such an artificial language would be doomed to fail due to the lack of experience embedding. This does not mean that artificial languages can not be constructed. A cognitively based simulation system that develops language through shared experience should be able to develop a viable artificial language. The complexity of such a system, however, might need to rival actual social life in order to get the necessary degree of embedding and variety of experience necessary to develop a full language. Once constructed, such artificial languages could be taught and would evolve. It would in general be more feasible to develop such a simulation system to generate a language for a specific task domain.

4. IN CONCLUSION

Constructuralism is a theory of the inter-relationship between individual cognition and social behavior. A possible model of cognitive structure was provided as a basis for this theory. There are three parts to this model: a network structured knowledge representation scheme, a set of cognitive principles (relative definition, immediate comprehension, saliency dependent access, monotonic changes, local consistency, and autonomy) and a set of cognitive operators (knowledge acquisition by being told, knowledge acquisition of discovery, knowledge acquisition by generalization, knowledge acquisition by attribution, knowledge communication, and choice of interaction partner). Constructuralism was used to address a wide variety of social phenomena. These phenomena included development of shared knowledge, identical behavior by members of the society, foreign language acquisition, clique formation, civil disobedience, and diffusion of innovative information. It was thus illustrated that a cognitively based theory can provide a unified basis for explaining a large variety of social phenomena. Fundamental similarities in social behavior across a wide variety of tasks was due to the underlying cognitive model.

For many of the predictions listed no supporting evidence was given. It is a necessary next step to gather data to test these predictions. There are also many other social phenomena that need to be explored or explored in more detail, e.g. role development, concensus formation, the evolution of norms, social stability, and the emergence and value of social support groups. Simulation models, based on constructuralism, can be used to explore these phenomena. That is, the theory is so complex and the phenomena so detailed that simulation is necessary in order to generate meaningful predictions. This approach is being taken to explore social stability and consensus (Carley, 1987b, Carley, 1988c).

In general, cognitively based models are complex. They tend to be non-linear, rife with feedback, and often based on heuristic algorithms. Consequently, it is necessary to turn to simulation in order to explore these models and determine their predictions.

As social theories become cognitively based simulation will become a major tool for theory development. The predictions of simulation models, like analytic models, can and should be tested with data.

The cognitive model described herein might be wrong. At the least, it is incomplete. The cognitive operators are certainly underspecified as is the knowledge representation scheme. Although not inconsistent with recent findings in cognitive science on knowledge acquisition and representation, the schemes employed can not generate all of the knowledge behavior necessary of such schemes. In (Carley, 1987b) it was demonstrated that the better the cognitive model the more accurate the behavioral predictions. This paper illustrates the explanatory power when a moderately realistic model of cognitive structure is used as a basis for social behavioral models. It seems reasonable to postulate that the better the model of the cognitive structure employed by social theory the more social behaviors will be explainable. Thus, for both accuracy of prediction and for breadth of explanation it may behoove social theorists to use better cognitive models than that used herein. Further, since social dynamics are a function of knowledge acquisition the specific area in which improvement in the cognitive models is needed, from a social theory point of view, is in the knowledge acquisition operators and the knowledge representation scheme.

By employing cognitive models as a basis for social models it is possible to specify exactly how interaction and the exchange of information affect behavior. Thus it is no longer necessary to simply assume that interaction and the exchange of information take place and that they somehow lead to social behavior. Rather this fundamental facet of social life, interaction, can be explicitly modeled and used as a link to convert between cognitive structure and social structure. Consequently, by basing social theories on cognitive theories it is possible to create a unified framework for explaining a wide variety of social phenomena at a more detailed and precise level. The point is not that a cognitive bases will admit an understanding of more phenomena, although this may be the case. Rather, the point is that by cognitively basing social theories it will be possible to examine when such effects as knowledge contagion (Coleman, et al., 1966, Crane, 1970), the effect of social groups (Festinger, 1954, Festinger, 1950) etc., will occur. With such a cognitively based framework it will be possible to explicitly model such social phenomena as roles and norms. With such a unified framework it will also be possible to specify the procedures by which, and the conditions under which, social phenomena such as roles and norms come into being and are maintained.

NOTES

1. Note: the term knowledge and information are being used interchangeably. In part, this is because no claims are being made about differences between the two. And, in part, this is to stress the idea that under the proposed theory there is no such thing as absolute value for knowledge; i.e., there are no absolute truths. Rather, validity is a function of the degree of shared tacit consensus to that piece of information (Carley, 1986a, Carley, 1981).
2. The process and theory of knowledge acquisition described are presumed to be relevant to adults gathering information. The early cognitive development of children may actually follow a different process. Whether or not this is the case depends on the cognitive architecture. If the cognitive architecture is not invariant then children and adults may acquire knowledge in different ways. If the cognitive archi-

tecture is invariant across age groups then the knowledge acquisition process would be functionally the same for children and adults. Even in this case knowledge acquisition may appear to be different for adults and children due to the different levels of symbolic encoding, different levels of meaning, different amounts of knowledge, and the different levels of structure or inter-relationship among knowledge.

3. The term interaction is used broadly. It includes both direct face-to-face exchanges as well as indirect exchanges through the media such as electronic mail, articles, books, etc.

REFERENCES

Abelson, R.P. and Reich, C.M. (1969). Implicated Molecules; A method for extracting meaning from input sentences. In Walker, D.E. and L.M. Norton (Eds.), *Proceedings of the International Joint Conference on Artificial Intelligence.:*.

Abelson, R.P. (1976). Script Processing in Attitude Formation and Decision-Making. In Carroll, J.S. and J.W. Payne (Eds.), *Cognition and Social Behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Anderson, J. (1983). *The Architecture of Cognition*. Cambridge, MA: Harvard University Press.

Anderson, J. and G. Bower. (1973). *Human Associative Memory*. Washington DC: Winston-Wiley.

Bennet, J. (1976). *Linguistic Behavior*. New York, NY: Cambridge University Press.

Berger, P. and T. Luckmann. (1966). *The Social Construction of Reality*. Garden City, NY: Doubleday.

Blalock, H.M. Jr. (1967). *Toward a Theory of Minority-Group Relations*. New York, NY: Wiley.

Blau, P.M. (1967). *Exchange and Power in Social Life*. New York, NY: Wiley.

Blau, P.M. (1977). *Inequality and Heterogeneity*. New York, NY: The Free Press of Macmillan Co.

Blumer, H. (1969). *Symbolic Interactionism*. Englewood Cliffs, NJ: Prentice Hall.

Bobrow, D.G. and Fraser B. (1969). An Augmented state transition network analysis procedure. In Bobrow, D.G. and B. Fraser (Eds.), *Proceedings of the First International Joint Conference on Artificial Intelligence*.

Bobrow, D.G. and D.A. Norman. (1976). Some principles of memory schemata. In Bobrow, D.G. and A. Collins (Eds.), *Representation and Understanding: Studies in Cognitive Science*. New York, NY: Academic Press.

Boorman, S.A. and H.C. White. (1976). Social Structure from Multiple Networks. II Role Structures. *American Journal of Sociology*, 81, 1384-1446.

Breiger, R.L. (1978). Career Attributes and Network Structure: A Blockmodel Study of a Biomedical Research Specialty. *American Sociological Review*, 41, 117-135.

Burt, R.S. (1976). Positions in networks. *Social Forces*, 55, 93-122.

Burt, R.S. (1977). Positions in multiple network systems, part one: a general conception of stratification and prestige in a system of actors cast as a social topology. *Social Forces*, 56, 106-131.

Carbonell, J. (1981). POLITICS. In Schank R.C. and C.K. Riesbeck (eds.) (Eds.), *Inside Computer Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Carley, K.M. (1988d). The Social Construction of Knowledge. In Experiments in Epistemology: Learning in People and Machines by Lawler R., Carley K., and G. Drescher Ablex, Norwood New Jersey.

Carley, K.M. (1984). *Constructing Consensus*. Doctoral dissertation, Harvard,

Carley, K.M. (1986). Knowledge Acquisition as a Social Phenomenon, *Instructional Science*, 14, 381-438.

Carley, K.M. (1986). An Approach for Relating Social Structure to Cognitive Structure. *Journal of Mathematical Sociology*, 12(2), 137-189.

Carley, K.M. (1988a). Formalizing the Social Expert's Knowledge. Journal of Sociological Methods and Research, in press.

Carley, K.M. (1987a). Separating the Effects of Structure and Interaction. Social and Decision Sciences Working Paper Series.

Carley, K.M. (1987b). Increasing Consensus Through Interaction and Social Structure. Social and Decision Sciences Working Paper Series.

Carley, K.M. (1987c). Language: Society's Chronicle. Sociological Abstracts.

Carley, K.M. (1988b). Structural Constraints on Communication: The Diffusion of the Homomorphic Signal Analysis Technique through Scientific Fields. Journal of Mathematical Sociology, forthcoming.

Carley, K.M. (1987d). *Construct I — Operational Notes* (Tech. Rep.). Pittsburgh, PA: Carnegie Mellon University,

Carley, K.M. (1988c). Social and Cognitive Stability. Sociological Abstracts.

Charniak, E. (1972). *Toward a Model of Children's Story Comprehension*. Doctoral dissertation, Massachusetts Institute of Technology,

Cicourel, A.V. (1974). *Cognitive Sociology*. New York, NY: The Free Press, Macmillan Publishing Co.

Clark, H.H. and Clark E.V. (1977). *Psychology and language: An introduction to psycholinguistics*. New York, NY: Harcourt Brace Jovanovich.

Coleman, J.S., Katz E., and H. Menzel. (1966). *Medical innovation; A Diffusion Study*. New York, NY: Bobbs-Merrill Co., Inc.

Collins, B.E. and Loftus, E.P. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.

Collins, A., J.S. Brown, and K.M. Larkin. (1977). *Inference in Text Understanding* (Tech. Rep. Report #3684). Bolt Berenek and Newman,

Crane, D. (1970). *Invisible Colleges: Diffusion of knowledge in scientific communities*. Chicago, IL: University of Chicago Press.

Danowski, J.A. (1980). Computer-mediated communication: a network-based content analysis using a CBBS conference. In Burgoon M (Ed.), *Communication Yearbook 6*. Beverly Hills, CA: SAGE.

Davis, R. and D. Lenat. (1980). *Knowledgebase systems in artificial intelligence*. New York, N.Y.: McGraw Hill.

Davis, K. (1949). *Human Society*. New York, NY: Macmillan Co.

Duncan, H. (1968). *Symbols in Society*. London: Oxford University Press.

Durkheim, E. (1951). *Suicide, A Study in Sociology*. Glencoe, IL: The Free Press.

Ennals, R. (1985). *Artificial Intelligence: Applications to Logical Reasoning and Historical Research*. West Sussex, England: Ellis Horwood.

Fauconnier, G. (1985). *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, MA: Bradford Books, MIT Press.

Festinger, L. (1950). Informal social communication. *Psychological Review*, 57, 271-282.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7, 117-140.

Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.

Festinger, L. et. al. (1950). The Study of a Rumor: Its Origin and Spread. *Human Relations*, 1, 464-486.

Festinger, L., S. Schachter, and K. Back. (1950). *Social Pressures in Informal Groups*. New York, NY: Harper and Brothers.

Goffman, E. (1963). *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Glencoe, IL: Free Press.

Goffman, Erving. (1974). *Frame Analysis: An Essay on the Organization of Experience*, New York, NY: Harper and Row.

Granovetter, M.S. (May 1973). The Strength of Weak Ties. *American Journal of Sociology*, 68, 1360-1380.

Granovetter, M.S. (1982). Alienation Reconsidered: The Strength of Weak Ties. *Connections*, 5.2, 4-15.

Grice, H.P. (1969). Utterer's Meaning and Intentions. *Philosophical Review*, 78, 147-77.

Halliday, M.A.K. (1978). *Language as Social Semiotic*. Baltimore, MD: University Park Press.

Hayes-Roth, F., D.A. Waterman, and D.B. Lenat (Eds.). Hayes-Roth, F., D.A. Waterman, and D.B. Lenat (Eds.). (1983). *Building Expert Systems*. Reading, MA: Addison-Wesley.

Heise, D. (1977). Social Action as the Control of Affect. *Behavioral Science*, 22, 163-77.

Heise, D. (1978). *Computer-Assisted Analysis of Social Action* (Tech. Rep.). Chapel Hill, N.C. Institute for Research in Social Science,

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

Heise, D. (1987). *Ethno Documentation* (Tech. Rep.). Indiana University,

Heise, D. (1987). Affect Control Theory: Concepts and Model. *Journal of Mathematical Sociology*, ,.

Homans, G.C. (1961). *Social Behavior: Its Elementary Forms*. New York. NY: Harcourt, Brace, and World, Inc.

Laird, J.E., Newell, A., and P.S. Rosenbloom. (1986). *Soar: An architecture for general intelligence* (Tech. Rep.). Carnegie-Mellon University (also Xerox PARC and KSL Stanford),

Laird, J., Rosenbloom, P., and A. Newell. (1986). *Universal Subgoaling and Chunking*. Boston, MA: Kluwer Academic Publisher.

Laird, J., Rosenbloom, P., and A. Newell. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, 1, 11-46.

Lock, A. (1980). *The Guided Reinvention of Language*. New York, NY: Academic Press.

Luria, A.R. (1978). *Cognitive Development*. Cambridge, MA: Harvard University Press.

Luria, A.R. (1981). *Language and Communication*. New York, NY: John Wiley and Sons.

McCall, G. and J. Simmons. (1978). *Identities and Interactions: An Examination of Human Associations in Everyday Life*, New York, NY: Free Press.

McClelland, J., Rumelhart, D., et. al. (1986). *Parallel Distributed Processing: Explorations in the Microstruccture of Cognition*. Cambridge, MA: MIT Press.

Minsky, M.A. (1975). A framework for representing knowledge. In Winston, P. (Ed.), *The Psychology of Computer Vision*. New York, NY: McGraw-Hill.
Moore, J. and A. Newell. (1974). How can Merlin Understand. In Gregg, L. (Eds.), *Knowledge and Cognition*. Potomac, MD: Lawrence Erlbaum Associates.
Parsons, T. (1951). *The Social System*. New York, NY: The Free Press of Macmillan Co.
Price, D.J. de S. (1963). *Little Science, Big Science*. New York, NY: Columbia University Press.
Price, D.J. de S. (1965). Networks of Scientific Papers. *Science*, 149, 510–515.
Price, D.J. de S. and D. Beaver. (1966). Collaboration in an invisible college. *American Psychologist*, 21, 1011–1018.
Rogers, E.M. (1982). *Diffusion of innovations*. New York, NY: Free Press.
Rogers, E.M., Shoemaker, F.F. (1971). *Communication of Innovations*. New York, NY: Free Press.
Ruben, B.D. (1975). Intrapersonal, Social and Mass Communication Processes in Individual and Mutliperson Systems. In Ruben B.D., Kim J.Y. (Eds.), *General Systems Theory and Human Communication*. Rochelle Park, NY: Hayden.
Ruben, B.D. and Weimann, J.M. (spring 1979). The diffusion of scientific information in the communication discipline. *Communication Quarterly*, 27(2), 47–53.
Rumelhart, D.E. and A. Ortney. (1976). The representation of knowledge in memory. In Anderson, Spiro and Montague (Eds.), *Schooling and the Acqusition of Knowledge*. Lawrence Erlbaum Associates.
Rumelhart, D., McClelland, J., et. al. (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA: MIT Press.
Schank, R. and Colby K. (1973). *Computer Models of Thought and Language*. San Francisco, CA: W.H. Freeman.
Schank, R. and R. Abelson. (1977). *Scripts Plans and Goals and Understanding*. New York, NY: Wiley.
Schank, R.C. and C.K. Riesbeck (eds.) (Eds.). Schank R.C. and C.K. Riesbeck. (1981). *Inside Computer Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.
Simmons, R. (1973). Semantic networks: Their computational use for understanding english sentences. In Schank, R. and Colby K. (Eds.), *Computer Models of Thought and Language*. San Francisco, CA: W.H. Freeman, Publisher.
Simon, H. (1986). What we know about the Creative Process. In Kuhn, R. (Eds.), *Frontiers in Creative and Innovative Management*. :.
Skvoretz, J., Fararo, T. and N. Axten. (1980). Role-programme models and the analysis of institutional structure. *Sociology*, 14, 49–67.
Skvoretz, J. and T. Fararo. (1988). Action programs and sociological action theory. *Journal of Mathematical Sociology*, see this issue.
Sowa, J.F. (1984). *Conceptual Structures*. Reading, MA: Addison-Wesley.
Stevens, A.L. and Collins A. (1978). The Goal and Structure of a Socratic Tutor.
Stryker, S. (1980). *Symbolic Interactionism*. Menlo Park, CA: Benjamin Cummings.
Tesser A. (1977). Toward a theory of self-generated attitude change. In Berkowitz, L. (Ed.), *Advances in Experimental Social Psychology* (10 ed.). New York, NY: Academic Press.
Thayer, L. (1968). *Communication and Communication Systems*. Homewood, IL: Irwin.
Tversky, A. and D. Kahneman. (1980). Causal Schemas in Judgments Under Uncertainty. In Fishbein, M. (Eds.), *Progress in Social Psychology*. Hillsdales, NJ: Erl-Baum.
Vygotsky, L.S. (1962). *Thought and Language*. New York, NY: Wiley.
Vygotsky, L.S. (1978). *Mind in Society*. Cambridge, MA: Harvard University Press.
D.A. Waterman (1986). *A Guide to Expert Systems*. Reading, MA: Addison-Wesley Publishing Co.
White, H.C., S.A. Boorman, and R.L. Breiger. (1976). Social Structure from Multiple Networks. I. Blockmodels of Roles and Positions. *American Journal of Sociology*, 81, 730–780.
Whorf, B.L. (1956). *Language Thought and Reality*. Cambridge, MA: MIT Press.
Winston, P. (1975). Learning Structural Description from Examples. In Winston, P. (Ed.), *The Psychology of Computer Vision*. New York, NY: McGraw-Hill.
Winston, P. (1977). Recognizing Analogies and Learning Simple Concepts. In Winston, P. (Ed.), *Artificial Intelligence*. Reading, MA: Addison-Wesley.
Wyer, R.S. Jr., and D.E. Carlston. (1979). *Social Cognition, Inference, and Attribution*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Journal of Mathematical Sociology, 1989, Vol. 14(2-3) pp. 209–216
Reprints available directly from the publisher
Photocopying permitted by license only

Printed in the United States of America

ON ARTIFICIAL INTELLIGENCE AND THEORY CONSTRUCTION IN SOCIOLOGY

BO ANDERSON

Michigan State University

1. INTRODUCTION

Among the behavioral sciences sociology is less influenced by the developments in Artificial Intelligence and Cognitive Science than the others. For sure, computer simulations have been used by sociologists for a long time: there are computer models of voting behavior, social mobility, decision-making in organizations, urban decay and many other social processes. But most computer simulation is not Artificial Science. Cognitive science is not the same as Artificial Intelligence, although the two fields are closely linked in the minds and procedures of many researchers. The recent history of cognitive science by Howard Gardner (Gardner, 1985) discusses the participation of psychologists, linguists, philosophers, anthropologists and neuroscientists in "the cognitive revolution", but sociologists are not mentioned. This is not the place to investigate the detailed reasons for the reluctance of sociologists to participate in one of the most fruitful movements in contemporary behavioral science. One reason may be that sociologists are, since Durkheim, preoccupied with finding for themselves a *concrete* set of phenomena that is "uniquely sociological". This makes them see certain other sets of concrete phenomena as already claimed for other disciplines, and therefore not the proper study for sociologists. For example, Granovetter has pointed out that sociologists in our time rarely study certain economic phenomena (Granovetter, 1985). The classics, e.g. Max Weber made major contributions to the study of economic life, but contemporary sociologists seem to regard markets, business decisions and transactions and the organization of industries as the province of microeconomics. This is simply confusing *concrete* sociological *subject matter* with the sociological *perspective*. A second reason may be sociologists' eagerness to define structure as the proper sociological concern and, hence, to downplay what happens "inside" the actor. But we also know that sociology does have, in the symbolic interactionisms, a considerable intellectual tradition that is, at least in part, cognitively oriented.

It is time for sociology to break its intellectual isolation and participate in the cognitivist rethinking of human action, and to avail itself of theoretical ideas, techniques and tools that have been developed in AI and cognitive science. This does not mean