

THE STRUCTURE OF SOCIAL MIND: EMPIRICAL PATTERNS OF LARGE-SCALE KNOWLEDGE ORGANIZATION

David L. Sallach
Department of Computer Science
University of Arkansas

The social model of mind regards intelligence as the result of the interaction of cognitive or subcognitive agents. In recent years, cognitive science research has found the concept of social mind provides a promising model for: 1) the explanation of brain damage (Gardner, 1974; 1983), 2) the design of expert systems (Lee, 1985), and 3) a general model of memory (Minsky, 1981;1985). The premise of the present discussion is that the structure of social mind has research implications for cognitive science that have yet to be explored.

The study of biological systems has yielded genetic algorithms that provide useful insights into the design of classifier systems (Holland, 1986).¹ Like biological systems, social systems are naturally occurring phenomena.² Thus, identification of the principles and organization of social entities also provides insights which may facilitate the construction of computational cognitive systems.

Social entities can be seen as providing a functional architecture through which mind can be instantiated. Accordingly, the study of cognitive systems will benefit by including the social system as a physical instantiation of mind, along with the computer and the human brain. For some cognitive science research questions, social systems are likely to provide a more useful source of data. Further, the models of social mind used by cognitive researchers are likely to be enriched by an investigation of empirical forms of social organization.

FUNCTIONALISM AND SOCIAL ENTITIES

A major side-effect of the modern computer revolution is the growing recognition among philosophers and cognitive scientists that thinking and other mental functions may be instantiated using diverse mechanisms (Dennett, 1981; Loar, 1981; Churchland, 1984; Gardner, 1985). Cognitive research increasingly relies upon both psychological experiments and computer simulation (cf., Kintsch *et al.*, 1984). Cognitive psychology provides insight and hypotheses to AI researchers (Hinton and Anderson, 1981), while computer modeling enriches cognitive psychology (Johnson-Laird, 1983).

One result of cognitive research considering comparative mental mechanisms is the abstraction of function, architecture and representation from physical implementation. Development of the functional perspective in cognitive science has led to the identification of three distinct levels of analysis: the physical level, the functional level and the knowledge or representational level. "The three levels are tied together in an instantiation hierarchy, with each level instantiating the one above" (Pylyshyn, 1984, p. 132).

Figure 1 illustrates the relationship between cognitive science and the disciplines upon which it is based. The highest level Newell calls the knowledge level, Pylyshyn calls the

¹Not surprisingly, genetic algorithms have been shown to provide especially effective models of biological phenomena (Farmer, Packard and Perelson, 1985).

²Social systems differ from the simpler biological systems in that their actions are frequently cognitively penetrable; see Pylyshyn (1984, pp. 130-145) for a discussion of penetrability.

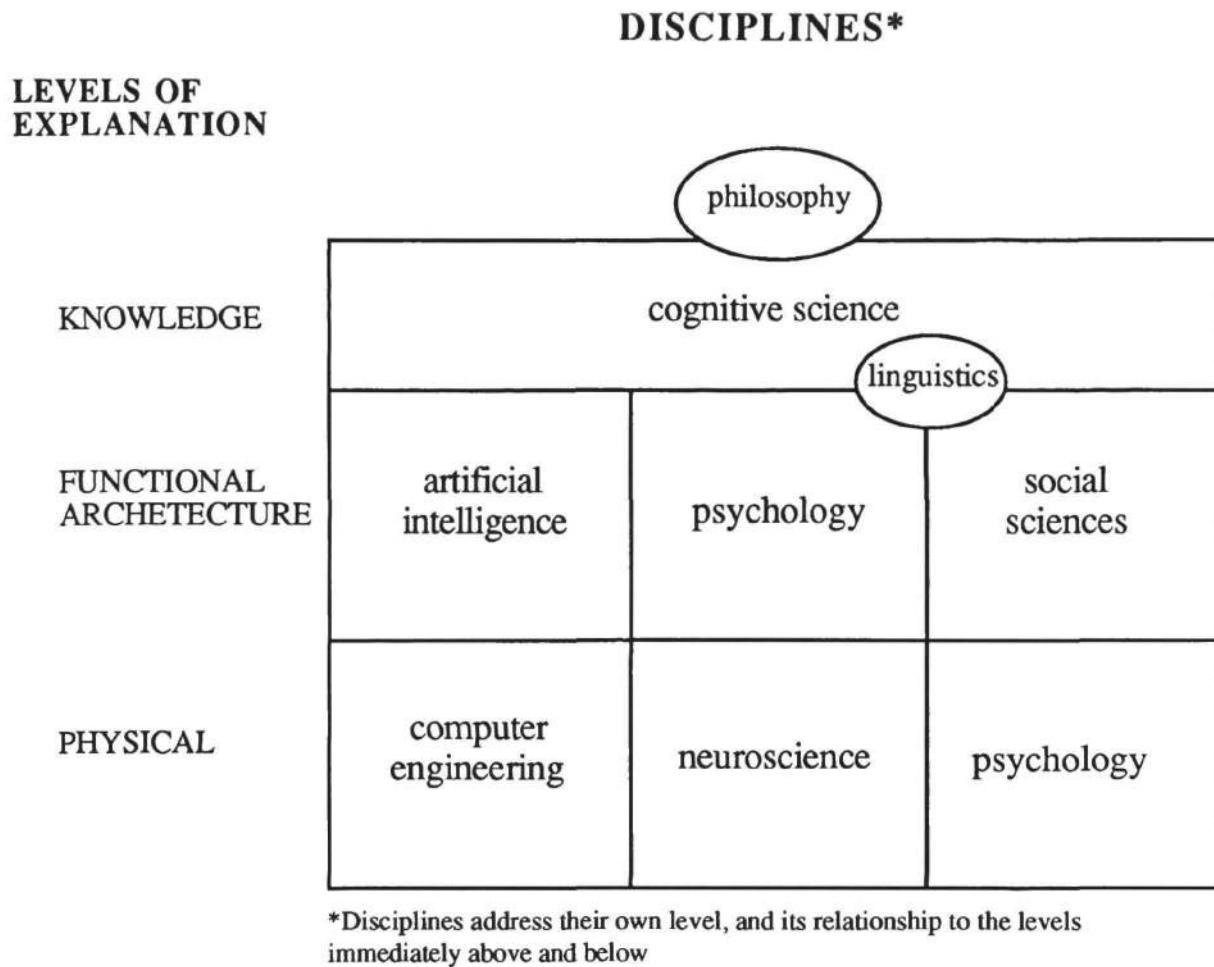


Figure 1. Cognitive Science Disciplines and Levels of Explanation

semantic level and philosophy references as the intentional level. It is the domain of cognitive science.³

The knowledge level may be instantiated using diverse functional architectures. The relationship between the knowledge level and prospective functional architectures is of interest to cognitive science, and to disciplines that study the architectures available through a specific type of physical system. For computers, the discipline is artificial intelligence; for *homo sapiens*, the discipline is psychology. By extension, the relationship between the knowledge level and the functional organization of social entities falls within the domain of the social sciences.

A second, more concrete, focus concerns the types of functional architecture that are available for a specific type of physical system. For computers, this relationship is studied by artificial intelligence and computer engineering (which together constitute a major portion of

³Figure 1 accurately depicts philosophy as addressing levels of analysis more abstract than that of cognitive science (e.g., ontology, epistemology), while philosophy of mind overlaps and interacts with cognitive theory. Figure 1 would be slightly more accurate if it showed the overlap and interaction between philosophy and linguistics.

computer science). For the human individual, this relationship is studied by neuroscience and physiological psychology. For social entities, the relationship between the physical system, and its possible and probable forms of functional organization is studied by the social and psychological sciences.

When applied to social phenomena, the functional theory of mind suggests that mental representations and processes may be attributes of social, as well as psychological and computational, entities. Dennett (1981) recognizes that the very virtues of functionalism⁴ permit a functionalist theory to be instantiated by suprahuman organizations for which it may seem counterintuitive to say they have minds of their own. Intuition, however, is not a reliable indicator. As Dennett further observes, "Inside your skull it is also all darkness, and whatever processes occur in your grey matter occur unperceived and unperceiving."

Even critics of functional theory have (skeptically) suggested that, under functional theory, social entities (e.g., the people of China) might form "a giant brain" (Churchland, 1985, p. 39). While it is more accurate to say "a giant mind", the people of China are hardly the only social entity that acts based upon shared representations and collective goals.

Over the past century, theorists have attributed various levels of 'reality' to social phenomena. The perspective developed in the present discussion is that of social realism. From its earliest articulation, social realism has been anti-reductionist in nature (Durkheim, 1964). In general, social realism holds that: 1) social entities are just as real as psychological entities, but that 2) both are abstract, analytical units, and 3) social phenomena must be explained in terms of a social level level of analysis, not reduced to a psychological one (cf., Warriner, 1956).

Social realism, as described above, allows and requires two caveats: 1) like other phenomena, the behavior of social entities is constrained by the nature of its compositional elements, whether psychological, biological, or physical, and 2) social patterns may be partially or wholly caused by forces which are better conceptualized at a higher, more abstract level of analysis. The latter caveat acknowledges that, insofar as the scientific comprehension of cognitive systems in general grows, social entities may be seen as providing one type of physical system through which a range of functional architectures are implemented. The subsequent review of research is based upon this perspective.

Recognition of the mental dimensions of social phenomena will benefit cognitive science research. Social entities are spatially much larger than the brain, and their communication rates are much slower and (frequently) more observable. It is not clear, however, that the cognitive process of social entities is any less complex than that of the brain. Cognitivists enjoy a unique vantage point on these large-scale cognitive processes: we are on the inside looking out. Incorporating this unique perspective into our research programs is likely to widen the range of cognitive issues that can be addressed.

THE BELIEF SYSTEMS OF MASS PUBLICS

As a preliminary example of how empirical research on the social organization of cognition may provide useful insights into the construction of computational models, consider nature of mass belief systems.⁵ It has long been recognized that opinion formation is a multi-stage process in which opinion leaders form their views based upon public sources of information, and then influence the attitudes of a periphery of opinion followers (Berelson, et al., 1954). This type of layered structure gives support to Minsky's (1981;1985) K-line model in which there is a division of labor between local agents and agents that perform more generalized pattern recognition.

⁴To wit, ". . . abstractness and hence neutrality with regard to to what could 'realize' the functions deemed essential to sentient or intentional systems" (Dennett, 1981, p. 153).

⁵Although the research to be reviewed pertains to the organization of political knowledge mass publics, there is no inherent reason to believe that the structural patterns are restricted to that domain.

Subsequent studies suggest further implications. Converse (1964), in a highly influential study, concluded that the mass public manifests belief systems organized at several identifiable layers of organization, ranging from a coherent organization of information at a high level of abstraction, through interest group identification, to those whose information is very specific, and not coherent or integrated. In Converse's words:

Moving from the top to the bottom of this information dimension the character of the objects that are central in a belief system change. These objects shift from the remote, generic and abstract to the increasingly simple concrete, or 'close to him'.

Converse's findings of differential levels of abstraction and conceptual integration in the belief systems of mass publics have been elaborated by subsequent research. Neuman (1981) has identified two complimentary dimensions of political thought: conceptual differentiation and conceptual integration. Differentiation refers to the number of discrete, concrete elements of information utilized by the actor. Conceptual integration involves the use of abstract concepts to structure discrete elements of information. Neuman posits a spiraling pattern of growing conceptual sophistication in which new elements of information are acquired (differentiation) and then abstracted into a manageable order (integration). The structure of belief then guides and constrains subsequent information acquisition.

Brady and Sniderman (1985) have isolated another aspect of mass belief systems that may be relevant to more generalized models of cognition. Members of mass publics are able to make highly accurate estimates of the political beliefs of strategic groups, and how members of such groups are likely to line up on key issues. This accomplishment is a puzzle since most members of mass publics demonstrate a low level of abstraction and little concrete information. The mechanism by which such estimates are made appears to involve an affective heuristic that serves as an intellectual shortcut. Specifically, respondents appear to combine their own beliefs with their affective response to strategic groups and generate an impressively accurate map of the political landscape. Other research suggests heuristics based upon multiple affective dimensions, for example, evaluation, potency and activity (Osgood, 1962; Heise, 1979).

This cursory review of research on mass belief systems suggests four components that might usefully be incorporated into models of mind based upon social organization. The first component differentiation of agents based upon their location in a knowledge hierarchy varying in level of abstraction and information span. A second component is the use of an affectivity calculus to enhance the application of agent knowledge, extending it at lower levels of the knowledge hierarchy and accentuating it at higher levels. The final component is the dynamic evolution of agents, through a spiral process, from a minimum information span at a low level of abstraction to a broader information span at successively higher levels of abstraction. Agent evolution implies the creation of new agents and the selective retention of mature agents (cf., Holland, *et al.*, 1986).

THE STRUCTURE OF SOCIAL MIND

The social organization of knowledge implies, in addition to agents (individually and collectively), a number of intermediate social entities. Social organizations are composed of smaller organizations which are ultimately composed of groups of agents that interact directly. Thus, the structure of social organizations is defined recursively to an arbitrary depth, further supporting the use of agents that possess extensive modeling capabilities.⁶ In terms of functional architecture, social organizations are massively parallel systems composed of intelligent agents.

⁶ The view developed here posits complex cognitive systems being composed, recursively, of complex cognitive systems. Such structures parallel the universality of scale identified in research on nonlinear dynamics in multiple domains (cf., Cvitanovic, 1986).

The information processing capacity of component agents is not infinite, but it is of sufficient depth that system designers may utilize significant levels of agent complexity.

The computational requirements implied by a system with a recursively nested structure composed of massively parallel agents, each having extensive modeling capabilities, are constrained by a shared definition and utilization of mental models among the multiple agents. Specifically, agents are not discrete entities, but nodes that have access to a larger knowledge network of which they are a part. Non-specialist subgroups and individual agents may thus be conceived as non-autonomous, maintaining simpler local models of objects and object types that draw upon the larger network model(s).

The simple cognitive models allow routine information processing, and serve as associative pointers to the more detailed information available as needed from the higher level agents of which they are a part. The integration of related models at multiple organizational levels is not a result but a task, an ongoing process that is achieved only partially (Schutz, 1967). In social systems, large salient model discrepancies lay the basis for conflict and disorderly change.

A major feature of social mind, emphasized by most theorists who invoke the model, is the use of a functional division of labor. The existence of specialists and generalists suggests a diffuse form of organization connected by another type of (hierarchical or heterarchical) structure. This second (control) structure determines the flow of information and control across multiple levels. The location of an agent within the control structure defines a vertical dimension of social structure. The advantages and disadvantages of alternative control structures constitutes a major topic of research for the social model of mind.

Empirical social systems suggest another type of differentiation: grouping across physical space. Spatial grouping provides communication opportunities and obstacles which, in social systems, result in complex patterns of ethnicity, language, religion, nationality and tradition. Spatial differentiation allows alternative organization and models to develop. Accordingly, it is a potential source of both experimentation and poor social integration. Control structure research might reasonably focus on the identification of structures which benefit from the former while minimizing the latter.

Thus, social models are constituted by a minimum of three dimensions of structure: division of labor, spatial dispersion and stratification of control (cf., Blau, 1977). Each agent, whether individual or social, is defined in terms of each dimension. Models of relevant objects are distributed across the matrix. The composite model of an object may be regarded as a conceptual prototype (cf., Rosch, 1978, ; Lakoff, 1987; Sallach, 1988).

The relationship between the composite model of an object, and the model used by each agent may be regarded as analagous to the multiple levels of database design, where the user view contains the information needed by a specific type of user, while the conceptual view integrates the user views of an enterprise into a single schema (cf., Flavin, 1981). However, in the database analogy, schema definition is performed by the database administrator. In the social model of mind, the integration of the models of multiple agents must be a self-organizing process.

THE REPRESENTATION OF SHARED KNOWLEDGE

The previous section describes a common knowledge network, where the detail of local representation is determined by the functional requirements of the position, and simple local models point to more generalized, or specialized repositories of information. A structure of this type is both efficient (in limiting information redundancy) and problematic (in coordinating the information needs of the agent with the resources of specialist nodes and of the network as a whole). The problem of representing such a knowledge network is largely a problem of the efficient integration of inconsistent knowledge. What is needed is the equivalent of a self-organizing schema for a semantic data model.

From the mass public, through intermediate organizations, to the individual agent, there are multiple levels of information integration to be reconciled in a social model of mind. The general pattern is that agents that are high in the social knowledge system are likely to have highly

integrated information structures. Other agents composing the social knowledge system will vary in the extent of information integration, and in the span of locally available information.

The structure of mass belief systems suggests a strategy for knowledge representation for the social model of mind. Within the social knowledge system as a whole, there will be islands of information that are comprehensive and highly integrated. For these nuclei, coherent schemata are generated. Entities whose information span is narrower but consistent, are defined in terms of a broader schema (the equivalent of a database 'view'). Entities whose information is coherent but organized according to different abstractions are represented as alternate schemata.⁷ Finally, entities and agents that fail to manifest a minimal level of integration are defined in terms of an affective heuristic. This strategy for representation of knowledge shared across a network is neutral as to whether the core schemata are defined in terms of individual agents or organization of agents. If, however, the level of integration is equivalent, the larger social entity would be the unit of preference.

The social knowledge system as a whole would not be presumed to manifest a high level of integration. Rather, knowledge system integration would be used to: 1) to provide a heuristic for the evaluation of alternate schema organization, and 2) as a variable property which constitutes an appropriate subject of investigation.

CONCLUSION

Social models of mind have been useful because of their ability to represent a functional division of labor among semi-autonomous agents. The present discussion maintains that mind may be usefully viewed as instantiated by social systems. This perspective suggests that the empirical investigation of social systems can provide useful insights into the social organization of cognitive systems.

The study of belief systems in mass publics illustrated three principles of the society of mind: the organization of knowledge along hierarchies of conceptual differentiation and integration, 2) the simulation of knowledge among low-level agents by the use of affective heuristics, and 3) dynamic learning through the evolution of agents. The structure of social mind is defined by a minimum of three dimensions (function, control and physical location), which suggests that system knowledge is distributed across a network of agents and organizations. The use of self-organizing schemata as a form of non-redundant knowledge representation was explored.

REFERENCES

- Berelson, B.R., Lazarsfeld, P. and McPhee, W.N. (1954). *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Blau, P.M. (1977). *Inequality and Heterogeneity: A Primitive Theory of Social Structure*. New York: Free Press.
- Brady, H.E. & Sniderman, P.M. (1985). Attitude attribution: A group basis for political reasoning. *American Political Science Review*, 79, 1061-1078.
- Carley, K. (1986). An approach for relating social structure to cognitive structure. *Journal of Mathematical Sociology*, 12, 137-189.
- Carroll, J. and Payne, J.W. (Eds.). (1976). *Cognition and Social Behavior*. Hillsdale, NJ: Lawrence Erlbaum.
- Churchland, P.M. (1984). *Matter and Consciousness*. Cambridge: MIT Press.
- Converse, P.E. (1964). The nature of belief systems in mass publics. In D. Apter (Ed.), *Ideology and Discontent*. New York: Free Press.
- Cvitanovic, P. (1986). *Universality in Chaos*. Bristol, UK: Adam Hilger.

⁷Several authors have developed procedures for assessing cognitive integration (cf., Schoder, Driver & Streufert, 1967; Carley, 1986; Smolensky, 1986).

- Dennett, D.C. (1981). Toward a cognitive theory of consciousness. In *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge, MA: MIT Press
- Durkheim, E. (1964). *The Rules of Sociological Method*. New York: Free Press.
- Farmer, J.D., Packard, N.H. and Perelson, A.S. (1985). The immune system, adaptation and machine learning. *Technical Report LA-UR-85-3151*. Los Alamos National Laboratory.
- Flavin, M. (1981). *Fundamental Concepts of Information Modeling*. Englewood Cliffs, NJ: Yourdon.
- Gardner, H. (1974). *The Shattered Mind*. New York: Vintage.
- Gardner, H. (1983). *Frames of Mind: The Theory of Multiple Intelligences*. New York: Basic Books.
- Gardner, H. (1985). *The Mind's New Science: A History of the Cognitive Revolution*. New York: Basic Books.
- Heise, D.R. (1979). *Understanding Events: Affect and the Construction of Social Action*. New York: Cambridge.
- Hinton, G.E. and Anderson, J.A. (Eds.). (1981). *Parallel Models of Associative Memory*. Hillsdale, NJ: Lawrence Erlbaum.
- Holland, J.H. (1986). Escaping brittleness: The possibilities of general purpose machine learning algorithms applied to parallel rule-based systems. In R.S. Michalski, J.G. Carbonell and T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach*, vol. 2. Los Altos, CA, Kaufmann.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., and Thagard, P.R. (1986). *Induction: Processes of Inference, Learning and Discovery*. Cambridge: MIT Press.
- Johnson-Laird, P. N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference and Consciousness*. Cambridge: Harvard University Press.
- Kintsch, W., Miller, J.R. and Polson, P.G. (Eds.). (1984). *Method and tactics in Cognitive Science*. Hillsdale, NJ: Lawrence Erlbaum.
- Lakoff, G. (1987). *Women, Fire and Dangerous Things: What Categories Reveal about the Mind*. Chicago: University of Chicago Press.
- Lee, R.M. (1985). Bureaucracy as artificial intelligence. In L. B. Methlie & R.H. Sprague (Eds.), *Knowledge Representation for Decision Support Systems*. New York: North-Holland.
- Loar, B. (1981). *Mind and Meaning*. New York: Cambridge University Press.
- Minsky, M. (1981). K-Lines: A Theory of memory. In D.A. Norman (Ed.), *Perspectives on Cognitive Science*. Hillsdale, NJ: Lawrence Erlbaum.
- Minsky, M. (1985). *The Society of Mind*. New York: Simon and Schuster.
- Neumann, W.R. (1981). Differentiation and integration: Two dimensions of political thinking. *American Journal of Sociology* 86, 1236-1268.
- Osgood, C.E. (1962). Studies of the generality of affective meaning systems. *American Psychologist* 17, 10-28.
- Polyshyn, Z.W. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge: MIT Press.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B.B. Lloyd (Eds.), *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum.
- Sallach, D.L. (1988). The representation and use of composite prototypes in problem-solving tasks. *Technical Report CSAS-TR 88-02*. Fayetteville: University of Arkansas.
- Schroder, H.M., Driver, M.J., and Streugert S. (1967). *Human Information Processing*. New York: Holt, Rinehart and Winston.
- Schutz, A. (1967). *The Phenomenology of the Social World*. Evanston: Northwestern.
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of Harmony Theory. In D.E. Rumelhart & J.L. McClelland (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Volume 1. Cambridge: MIT Press.
- Warriner, C. (1956). Groups are real: A reaffirmation. *American Sociological Review*, 21, 549-54.