

---

WORKSHOP REPORT

---

# Theoretical Issues in Conceptual Information Processing

*James Hendler, B. Chandrasekaran,  
Beth Adelson, Richard Alterman,  
Tom Bylander, and Michael Dyer*

---

*The Fifth Annual Theoretical Issues in Conceptual Information Processing Workshop took place in Washington, D.C. in June 1987. About 100 participants gathered to hear several invited talks and panels discussing the issues relating to artificial intelligence and cognitive science.*

---

## From the Program Chair

In 1981, Bob Abelson delivered a keynote speech to the Cognitive Science Society. In this talk he differentiated between two large segments of the AI community: the neats and the scruffies. The neats liked everything a priori formalized. The scruffies preferred to build systems and experiment with new ideas. The neats tended toward logic, the scruffies toward psychology; the neats worked with paper and pencil, the scruffies with computers; the neats worried about soundness and consistency, the scruffies leaned toward nonlogical representation schemes and used terms such as psychological validity; the neats tended to congregate on the West Coast, the scruffies on the East.

A lot has changed since that talk. As applied AI has gained in popularity, the term scruffy has become associated with all those who build systems, particularly expert systems. The International Joint Conference on Artificial Intelligence and the American Association for Artificial Intelligence began to present two tracks at their conferences: science and engineering. The original scruffies, who had always followed a sort of experimental paradigm (using implementations to explore theoretical ideas) didn't really fit; those still worrying about cognition and cognitive tasks became AI's excluded middle. They started to move into other, more socially acceptable, subfields of the discipline. Some strayed into the logic formalism camp. Some followed the

dollar signs and became corporate tycoons, building expert systems and natural language front ends. Many joined the newest rebel movement and became connectionists.

The annual Theoretical Issues in Conceptual Information Processing Workshop has become a sort of gathering place for those who haven't yielded to temptation and have attempted to live within the old scruffy paradigm: building systems and exploring cognitive tasks. The June 1987 workshop, however, was geared toward bringing new faces into this community. It was aggressively advertised, the call for papers was widely announced, and the program committee (Richard Alterman, Jaime Carbonell, Michael Dyer, and myself) went out of its way to find panels and panelists that would broaden the appeal of the workshop.

I, for one, was amazed by the amount of interest expressed. Over 150 people were invited to attend, and many others, although doing related work, could not be admitted because of space and budget limitations. Participants were active, both in and out of the workshop rooms, and many a debate continued long into the night (I was informed that no fistfights were actually observed). It is impossible to convey all the views expressed at the workshop, but this article includes contributions from those actively involved in organizing the event. These participants describe their panel or talk and add personal observations as necessary.

First, I wish to thank the workshop

sponsors for their support. AAAI provided financial assistance that enabled us to invite many more graduate students than might normally attend such a workshop. The University of Maryland Institute for Advanced Computer Studies matched the AAAI dollars, enabling us to afford the rooms where the workshop was held, feed the participants, and help support the keynote speakers and other invited guests. DARPA, ONR, NASA and NSF also contributed by allowing release time for their attending personnel. A special thanks to Y. T. Chien, Alan Meyrowitz, Mel Montemerlo, and Robert Simpson, who presented a special panel on directions for cognitive science research and funding.

—J. Hendler, University of Maryland

### From the Workshop Chair

AI is commonly said to be preparadigmatic, or prescientific; that is, it is not yet a science with a unified methodology. Stances about the nature of intelligence and how to study it abound. In an earlier workshop in this series, I presented a view of how AI actually works as a discipline and identified three views about how to study and understand intelligence: (1) architectural theories, (2) logical formalisms, and (3) functional theories (Chandrasekaran 1988). The series of TICIP workshops have more or less implicitly brought together workers in the functional theory camp. The goal of these gatherings has been to understand intelligence and cognition as feasible computations as they apply to the construction of performance programs for narrowly defined tasks (expert systems). The feasibility aspect emphasizes organizational issues, which, in turn, strongly color the representational commitments, and leads to an attempt to identify the functional components of intelligence as a process and how each of these components is achieved in a computationally efficient manner.

This workshop, as those which preceded it, explored issues pertaining to knowledge and memory for natural language understanding, planning, problem solving, explanation, learn-

ing, and other cognitive tasks within this paradigmatic perspective. The approaches were limited by a concern with representation, organization, and the processing of conceptual knowledge, with an emphasis on empirical investigations of these phenomena through the experimentation and implementation of computer programs.

Connectionism has been met in the AI community with an interest ranging from mere curiosity to religious conversion. In my talk on connectionism and AI at the workshop, I granted one of its basic claims: It is not to be viewed as a mere implementation-level mechanism for discrete symbolic theories. In principle, a connectionist implementation and a symbolic implementation for the same task can make different representational commitments and, thus, can constitute different theories of information processing for the task. However, as the tasks get sufficiently far from the raw architecture level, that is, for most cognitive phenomena, the differences in representational commitments between connectionist and symbolic realizations become increasingly smaller. Systems for implementing these tasks share the information-processing abstractions needed for the task. These abstractions dominate the differences in the content of the representations. The hard work of theory making in AI will always remain—proposing the information processing-level abstractions needed for the tasks that constitute intelligence. Thus, I argued, connectionism is a corrective to what one might call Turing Imperialism; however, for most of the control issues of intelligence, connectionism is only marginally relevant.

—B. Chandrasekaran, Ohio State

### Reasoning by Analogy

*Panel Chair:* Beth Adelson, Tufts University. *Panel Members:* Jaime Carbonell, Carnegie-Mellon University; Kris Hammond, University of Chicago; Douglas Hofstadter, University of Michigan. *Discussant:* Andrew Ortony, University of Illinois. The panel was concerned with mak-

ing theories of analogic learning powerful in their predictions and explanations. The members of the panel converged on the following paradigm: When modeling a system that builds and uses analogy, the system should be viewed as functioning within a problem-solving context. That is, the system is trying to bring knowledge to bear for the purpose of solving whatever problem is currently at hand. Thus, the knowledge the system has and the particulars of the problem being solved both constrain the system during analogic problem solving. When this theoretical view of problem solving under constraint is taken, the researcher can see the functional constraints under which the system operates. As a result, the system can be formulated at a functional level and gains the desired explanatory and predictive power.

This issue of contextually constrained analogic problem solving addresses the question of how analogies are formed and used. However, our paradigm also leads us to ask why analogies are formed. Clearly, they are formed to solve certain types of problems (those in which a previously solved problem can provide relevant information for a current problem). This answer is important because it leads to the following observation. An analogic problem solver needs to be a subsystem, existing within a larger system for solving all types of problems, only some of which are analogic. Thus, the analogic subsystem has to be formulated in a way that allows it to work not only with a general-purpose memory but also with subsystems which perform other types of problem solving.

The following subsection illustrates the use of contextual constraints as an aid to theory development. (This work was done collaboratively by Mark Burstein and Beth Adelson.)

### Mapping and Integrating Partial Mental Models

The goal of this research is to develop an analogic learning theory that accounts for the formation and use of multiple partial but general purpose-based models (for example, behavioral, causal, and topological models) in complex problem solving.

We are developing the theory by producing a computational model that acquires models by analogy, integrates the models, and answers questions based on the integrated models.

**Mapping Purpose-Based Models.** One aspect of our work concentrates on the development of a detailed theory of purpose-driven analogic mapping that takes into account the type of model which needs to be produced in the current problem-solving context. The distinction between model types needs to be made explicit because it is difficult to find a single analogic model that can adequately describe all aspects of a complex target system and, therefore, can support the solution of all types of problems relevant to the system. We believe that the role of the purpose-based model type in selecting and mapping relevant portions of a base domain is an important and, as yet, poorly understood component of analogic reasoning. In addressing this issue, we are studying and modeling the details of how different types of models are selected and mapped to form models of the same type in a new domain. Our distinction between model types simplifies the description of analogic structure matching and mapping; in our account, models of a given type map to form new models of the same type. (The matching and mapping processes are difficult to describe successfully without this explicit distinction because models of different types use different relations in their descriptions.)

**Integrating Purpose-Based Models.**

A second aspect of our research focuses on the integration of partial purpose-based models in a target domain. Typically, several different types of models are needed to capture all aspects of a complex target domain. As a result, integrating a domain model that was acquired by reasoning from several analogies is an important component of the overall learning process. A detailed theory of the reasoning required to integrate partial models is being developed. These reasoning procedures are general to a learning theory in that they combine partial models which were acquired either by analogy or direct explanation. Additionally, the principles

Old Paradigm	New paradigm
Ahistoric	Memory-based coupled with learning. Associative memory used for retrieval and testing. Case-based
Small chunks	Large chunks
Abstract	Concrete
Situation independent	Situation dependent Situation action Situation matching
Planning/acting independent.	Planning/acting interleaved or indistinguishable. Dynamic situations. Reactive, real time.
complete knowledge	incomplete knowledge
domain independent	domain dependent
Single or few goals	Many goals

Figure 1. *New Directions in Planning Research.*

underlying the procedures are independent of the domains being modeled. They are based on: (1) the kinds of relationships needed to describe each type of model, (2) general knowledge of how these relationships can be used to model various aspects of a system, and (3) knowledge of the correspondences between model types. The principles form the basis of a kind of pattern-matching process in which corresponding elements of two different types of models are identified and then appropriately linked.

We believe that refining the mapping process by including a taxonomy of purpose-based models and specifying a set of domain-independent reasoning mechanisms for merging the resulting models results in a powerful theory of analogic learning.

**New Directions in Planning**

*Panel Chair:* Richard Alterman, Bran-

deis University. *Panel Members:* David Chapman, Massachusetts Institute of Technology; Tom Dean, Brown University; Jim Hendler, University of Maryland; Janet Kolodner, Georgia Institute of Technology; and Edwina Rissland, University of Massachusetts.

Recently, a renewed interest in planning has emerged within the AI and cognitive science communities. The idea behind this panel was to identify some of the major themes, not only by those whose concerns were strictly identified with planning but also those whose interests in the general problem of reasoning seem to contain many of the same motivating concerns. (As Chapman suggested, "Work on planning as an isolated phenomenon will eventually give way to studying planning as one of many sorts of reasoning that go into guiding practical activity.") With this statement in mind, I posed these two deliberately provocative sets of ques-

tions to the panel: (1) What is the old reasoning and planning paradigm? How would you describe the emerging paradigm in reasoning and planning in contrast to the old one? (2) In what direction do you see us moving? How does your own work on reasoning and planning reflect these directions? The panelists exchanged viewpoints before the workshop. During the workshop, each panelist gave a five-minute presentation; an hour-long discussion followed.

The mood of the panel can be characterized by the following concerns and questions: (1) Maybe it is not such a good idea to entirely separate planning and acting. (2) Maybe we need to pay closer attention to the role of situations. (3) How do you plan in situations where routines and habitual activities are the rule? (4) We need to pay closer attention to situations where multiple goals are operative. (5) What is the role of memory in planning? (6) What about dynamic planning situations that require real-time decisions?

The emerging themes can also be characterized by the list of contrasting old and new ideas shown in figure 1. In many cases, the so-called new ideas have been around in one form or another for several years. What does appear to be new is the shift in focus between sets of ideas (or assumptions).

## Diagnostic Reasoning

*Panel Chair:* Tom Bylander, The Ohio State University. *Panel Members:* Shoshana Hardt, the State University of New York at Buffalo; Michael Paz-zani, the University of California at Los Angeles; Jon Sticklen, The Ohio State University; and Roy Turner, Georgia Institute of Technology.

In essence, diagnosis is a problem of explaining failures of expectation. The process of diagnosis starts with observations of a misbehaving system and, if successful, produces a single persuasive explanation of why the system isn't working properly. Theories of diagnosis must then be concerned with the representation of observations and explanations and the process of using diagnostic knowledge to

generate explanations from observations as well as the process of learning how to diagnose.

Diagnostic problems differ along a number of dimensions. Devices generally provide much freedom for observing internal events and changing input; the situation is quite different in medical diagnosis. The form of knowledge can range from statistical associations to a physical model of behavior, from schematic knowledge for each possible malfunction to the generation of possible malfunctions and the rules for evaluating them. In many domains, more than one malfunction is rare, but in others, multiple, causally related malfunctions often occur. Because many types of diagnostic problems exist, it is unlikely that a single theory is sufficient to cover them all. Thus, it is important to understand each theory in terms of the type of diagnostic reasoning that it tries to explain and the inherent limitations of the problem solving it proposes.

Hardt discussed diagnostic reasoning that fits a general understanding process. In this process, features of the domain are organized into meaningful, mostly predefined conceptual structures. The conceptual structures are acquired from a variety of sources with different degrees of abstractness. At opposing extremes, structures might be based on a deep understanding about the domain, or structures might be based on how previous cases were processed and generalized. Reasoning is required to connect different structures, fit a structure to a new situation, or select among various structures in cases of incomplete knowledge. The effectiveness of diagnostic reasoning is measured by the ability to single out, from all the potential structures, only the relevant ones that should be pursued and used to drive the reasoning process. Hardt and her colleagues developed a shell called DUNE (Hardt et al. 1986) that can support feature organization and case-based reasoning. She presented several examples of how these ideas were applied to the domain of psychiatric diagnosis (Hardt and MacFadden 1987).

Turner described his plans to construct MEDIC, a case-based diagnostic

reasoner working in the domain of pulmonology. This project has three goals: (1) to examine the role of experience in medical diagnosis, (2) to examine the human diagnostic reasoning process, and (3) to provide a starting point for a next-generation expert system. MEDIC has four main parts: a long-term memory, a short-term memory, a reasoner, and a controller. The long-term memory is a dynamic memory (Schank 1982) patterned after Kolodner's (1984) CYRUS program. It includes representations of diseases, causes of diseases, doctors, and patients as well as episodic information. The short-term memory is a simple blackboardlike data structure that can contain reminders, hypotheses being entertained about the problem, and expectations of future findings. The reasoner is a case-based opportunistic reasoner. It shares features with case-based reasoners and opportunistic planners (Hayes-Roth 1985): It reasons from previous cases, and it is capable of being interrupted at any time by new information from the memory or the user. The controller handles communication between the memories and the reasoning.

Sticklen (1987) presented MDX2, a problem-solving system that uses a distributed architecture to diagnose cases in a subdomain of clinical medicine. MDX2 is based on the idea that complex knowledge-based reasoning can be decomposed into instantiations of generic problem-solving types (Chandrasekaran 1986). It integrates many of the problem-solving abilities that researchers have found desirable in automated medical diagnostic procedures: data-directed reasoning, question asking, multiple disease perspectives, the ability to handle multiple interacting diseases, the ability to support compiled diagnosis with deep-level reasoning, and abductive problem solving. Sticklen addressed two issues. First, to decompose complex diagnostic problem solving into primitive interacting modules, MDX2 contains an abductive algorithm that directs attention between different disease classes, handles attention shifts as they become necessary because of incoming patient data, and forms composite hypothe-

***This workshop . . . explored issues pertaining to . . . natural language understanding, planning problem solving, explanation, learning, and other cognitive tasks***

ses. Second, to combine compiled problem-solving ability with a deep reasoning component, MDX2 incorporates two types of methods for evaluating a disease: (1) directly matching a prestored pattern with the patient data and (2) deriving a pattern from a functional representation of human body physiology, then carrying out the match.

Pazzani compared several different approaches to learning diagnostic knowledge. In this case, learning to diagnose involves acquiring associations from atypical features (that is, symptoms) and malfunctions. However, in many disciplines, such as diagnosing electronic or mechanical systems, students are taught how these systems operate rather than what the empirical associations are between symptoms and malfunctions. These students acquire empirical associations from problem-solving experiences. Pazzani compared similarity-based and explanation-based learning approaches and failure-driven and success-driven approaches. In contrast to similarity-based systems, which require a large number of examples (Pazzani and Dyer 1987), explanation-based learning can be used to acquire diagnostic knowledge much more effectively if sufficient knowledge of system functionality is available (Pazzani 1986). With regard to failure-driven and success-driven learning of diagnostic knowledge in an explanation-based learning system, failure-driven learning is more effective because it acquires knowledge that better distinguishes between potential diagnoses.

### **Is Connectionism Cognitive Science?**

*Panel Chair:* Michael Dyer, UCLA.

*Panel Members:* B. Chandrasekaran, The Ohio State University; Christopher Cherniak, University of Maryland; Donald Dearholt, New Mexico State University; Stevan Harnad, Princeton University; and Yorick Wilks, New Mexico State University. The TICIP workshops are normally attended by natural language, planning, and cognitive modeling researchers, with a relatively large subgroup consisting of Roger Schank's

students and former students. The approach taken by these researchers is typically symbolic in nature (an approach not normally considered controversial). However, at the most recent TICIP workshop, quite a lot of discussion concerned connectionism. Two major address speakers talked at length on the subject. Chandrasekaran spoke on connectionism and cognitive science. Wilks gave a keynote address entitled "What Does Connectionism Mean for Natural Language Semantics?" Schank, another major address speaker, briefly mentioned connectionism and included it (along with logic and syntax) as another submovement within cognitive science that takes researchers away from the real task of developing content-oriented models of cognitive processing. In addition, I discovered that I had been made chair of the panel entitled "Is Connectionism Cognitive Science?"

Overall, the speakers emphasized how connectionism has to solve essentially the same problems as those facing cognitive modeling researchers; thus, connectionism is not a panacea. Because of the downplaying of connectionism, what is potentially new, useful, or exciting about it did not always get across.

Connectionist Parallel Distributed Processing models are attractive for many reasons: Content addressability and pattern completion are free; the models thrive on noise; they extract statistical regularities from data and are somewhat able to generalize appropriate responses to novel input; they learn through constant repetition and adaptation (versus rule-based induction); they can exhibit some rulelike behavior without any explicit representations of rules or a rule interpreter; they make decisions by massively parallel constraint satisfaction; they provide a closer link to both the brain sciences (for example, they are lesionable) and the physical sciences (through their statistical properties); they implement theories of forgetting based on interference (versus the lost pointer or a weakened chemical trace); they allow intermediate knowledge structures to be generated dynamically (for example, something in between a bedroom and a kitchen); and they provide a way to represent

skill-based (versus declarative) knowledge (Rumelhart, McClelland, and the PDP Research Group 1986).

Another way to ask the question "Is connectionism cognitive science?" is "How far can one push weight matrices?" This question has been taken up by several researchers, for example, Pinker and Prince (1987). At TICIP, I simply pointed out that symbolic and PDP models at this stage seem to live in separate computational-ecological niches. Much of what is currently awkward for symbolic models is easy for connectionist models and vice versa. Here is a brief list of features that are easy for symbolic models but pose difficult problems for PDP models: variables, role bindings, pointers, recursive data structures, separate instantiations, inheritance, stored inference chains, unbounded input sequences, the factoring out of concepts, storage management, theories of complex architectures of control, and portability of knowledge (that is, any procedure can access and apply declarative information produced by any other procedure as long as a shared canonical representational formalism exists).

Shortly after the workshop, I attended (along with 2000 others!) the First International Conference on Neural Networks in San Diego. Here, I overheard more than one attendee state that "AI is dead" and heard more than one major speaker refer to avoiding the "AI trap" (in which claims are made that hard, unsolved problems have been solved by an AI system when the system could never deal with real-world input). At the same time, some commercial fliers at the conference mentioned the total elimination of programmers because of programs that now learn on their own. Some commercial leaders in the new business of neurocomputing have even stated in interviews that their products operate at the stage of a 4-year-old. (A similar category mistake was made years ago by many AI researchers.) Seeing the connectionist field burning with an even hotter and faster flame than AI probably gives some AI researchers pleasure because the field of AI has itself been criticized for excessive hype. However, given the speed and excitement with

which connectionism is growing, it is important that both the correct role and the promise of connectionism (with respect to cognitive science) be elucidated.

So what's the relationship between connectionism and cognitive science? Much connectionist research appears to be bottom up in nature. Many variants on the generalized delta rule are created and their properties observed. Specific networks are built, and experiments are performed to see how well they behave on various tasks. This kind of research is of fundamental importance, but a top-down approach is also needed. It is important to select high-level cognitive tasks, such as language comprehension, and construct complex connectionist architectures capable of performing these tasks (McClelland and Kawamoto 1986; Touretzky and Hinton 1986; Dolan and Dyer 1987; Touretzky 1987). Those connectionists coming from the physical and signal-processing sciences are often unaware of the problems inherent in high-level planning, problem solving, and natural language tasks. Consequently, a need exists for AI researchers to carefully examine this emerging field of connectionism and to contribute their knowledge, insights, and experience. Connectionism is an important part of cognitive science.

#### References

Chandrasekaran, B. 1988. What Kind of Information Processing is Intelligence? A Perspective on AI Paradigms and a Proposal. In *Foundations of Artificial Intelligence: A Source Book*, eds. D. Partridge and Y. Wilks. London: Cambridge University Press. Forthcoming.

———. 1986. Generic Tasks in Knowledge-Based Reasoning: High-Level Building Blocks for Expert System Design. *IEEE Expert* 1(3): 23-30.

Dolan, C., and Dyer, M. G. 1987. Symbolic Schemata, Role Binding, and the Evolution of Structure in Connectionist Memories. In *Proceedings of the Institute of Electrical and Electronics Engineers First Annual International Conference on Neural Networks*.

Hardt, S. L., and MacFadden, D. H. Forthcoming. Computer Assisted Psychiatric Diagnosis: Experiments in Software Design. *Computers in Biology and Medicine*.

Hardt, S. L.; MacFadden, D. H.; Johnson, M.; Thomas, T.; and Wroblewski, S. 1986. The DUNE Shell Manual: Version 1, Research Report, 86-12, Dept. of Computer Science, State Univ. of New York at Buffalo.

Hayes-Roth, B. 1985. A Blackboard Architecture for Control. *Artificial Intelligence* 26:251-322.

Kolodner, J. L. 1984. *Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model*. Hillsdale, N.J.: Lawrence Erlbaum.

McClelland, J. L., and Kawamoto, A. H. 1986. Mechanisms of Sentence Processing: Assigning Roles to Constituents of Sentences. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, vol. 2, eds. J. L. McClelland and D. E. Rumelhart, 272-327. Cambridge, Mass.: MIT Press/Bradford Books.

Pazzani, M. 1986. Learning Fault Diagnosis Heuristics from Device Descriptions. Paper presented at International Meeting on Advances in Learning, Les Arc, France.

Pazzani, M., and Dyer, M. 1987. A Comparison of Concept Identification in Human Learning and Network Learning with the Generalized Delta Rule. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Pinker, S. and Prince, A. 1987. On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition, Technical Report, Occasional Paper #33, Center for Cognitive Science, Massachusetts Inst. of Technology.

Rumelhart, D. E.; McClelland, J. L.; and the PDP Research Group 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, 2 vols. Cambridge, Mass.: MIT Press/Bradford Books.

Schank, R. C. 1982. *Dynamic Memory*. New York: Cambridge University Press.

Sticklen, J. 1987. MDX2: An Integrated Medical Diagnostic System. Ph.D. diss., Dept. of Computer and Information Science, The Ohio State Univ.

Touretzky, D. S. 1987. Representing Conceptual Structures in a Neural Network. In *Proceedings of the Institute of Electrical and Electronics Engineers First Annual International Conference on Neural Networks*.

Touretzky, D. S., and Hinton, G. E. 1986. A Distributed Connectionist Production System, Technical Report, CMU-CS-86-172, Dept. of Computer Science, Carnegie-Mellon Univ.