AI Magazine Volume 3 Number 2 (1982) (© AAAI)

## Research in Progress

# AI Research at Rutgers

Edited by:
Tom M. Mitchell

*Department of Computer Science*
*Rutgers University*
*New Brunswick, New Jersey 08903*

RESEARCH by members of the Department of Computer Science at Rutgers, and by their collaborators, is organized within the Laboratory for Computer Science Research (LCSR). AI and AI-related applications are the major area of research within LCSR, with about forty people—faculty, staff and graduate students—currently involved in various aspects of AI research.

One project which has had a major impact on the stimulation and growth of AI research at Rutgers is the Research Resource on Computers in Biomedicine, which was started by Saul Amarel in 1971, and has been supported by NIH since that time. Casimir Kulikowski is Co-Principal Investigator of the Rutgers Research Resource. The emphasis of research in this project has been on applications of AI to problems of knowledge organization and reasoning of the type encountered in medicine and in psychological research. It has been the philosophy of the project from the start to combine collaborative research on the construction of high performance expert systems with core research on AI methodologies and with general system development activities that enhance the environment for AI research. The Rutgers Resource involves a large number of research subprojects and collaborations with researchers from other institutions. The Rutgers Resource is part of the national AIM (AI in Medicine) project which has been supported by NIH since 1974.

More recently, a number of newer AI research projects have grown within LCSR. Research on AI applications to Digital System Design is being conducted under DARPA sponsorship, with Tom Mitchell and Saul Amarel as Co-Principal Investigators. Research on AI approaches to the study of Legal Reasoning is supported by NSF, with Thorne McCarty and N. S. Sridharan as Co-Principal Investigators. Research on building AI consultation systems for Naval mission planning is supported by NRL, with C. V. Srinivasan as Principal Investigator. Research on Machine Learning of Heuristics is supported by NSF, with Tom Mitchell as Principal Investigator. Research on expert systems for oil exploration is supported by Amoco, with Casimir Kulikowski and Sholom Weiss as Co-Principal Investigators. And basic research on the Representation of Knowledge in AI and Database Theory is being conducted with Ray Reiter as Principal Investigator.

Thus, current research at Rutgers covers a broad range of AI concerns—from the development of expert systems in a variety of domains to the study of basic issues of representation and inference. In the following, we summarize each of the research projects mentioned above.

This article is based on research project summaries provided by the following people: Saul Amarel, Casimir Kulikowski, L Thorne McCarty, Tom Mitchell, Raymond Reiter, Charles Schmidt, N S Sridharan, C V Srinivasan, and Sholom Weiss Jo Ann Gabinelli provided valuable assistance in organizing and formating the articles

## Expert Systems Research

*Participants:* Casimir Kulikowski, Sholom Weiss, Peter Politakis, John Kastner, George Drastal, Chidanand Apte

*Software Support:* Kevin Kern, Michael Uschold

*Research Support:* NIH (BRP), Amoco Production Company

**Background.** Research in the area of expert systems has developed from our experience in building consultation programs in a number of application domains (Weiss, Kulikowski, and Safir, 1978; Lindberg et al., 1980; Kulikowski, Weiss, and Galen, 1981; Kulikowski, 1980).

The EXPERT system (Weiss and Kulikowski, 1979) is a generalized scheme for building expert reasoning models, exercising them with individual problems, testing and analyzing their performance on large numbers of problem-types, and improving them by knowledge base refinement techniques. The system has been operational on DEC 10/20 computers since 1978; versions also exist on VAX and IBM computers This system has been used by specialists in medicine, biomedical modeling, oil exploration, and chemistry to build models that capture their expertise in problem solving. In 1981 we completed an interesting technology transfer experiment in which a model for the interpretation of serum protein electrophoresis patterns was automatically translated from its EXPERT representation into algorithmic form, and then automatically translated into assembler code for running on a microprocessor (Weiss, Kulikowski, and Galen, 1981).

The EXPERT system is unusual among knowledge-based AI systems in that efficiency is a major design goal. This permits rapid analysis of cases, evaluation of performance over large numbers of cases, and working in real time with large size models. The system is also easily portable, being written in FORTRAN The types of consultation problems best suited for EXPERT are classification problems. These are problems which have a predetermined list of potential conclusions which the program can combine into interpretation or advice sequences. PROSPECTOR and EMYCIN are also knowledge-based systems which specialize in forms of classification problems.

**Collaborative research projects in medicine.** EXPERT is being used extensively in the development of several medical consultation models in collaboration with clinical investigators in rheumatology, ophthalmology, clinical pathology, and with researchers in biomedical modeling.

Problems in rheumatology are particularly important in health care, given the high prevalence and chronic nature of arthritis and related disorders. Drs. D. Lindberg and G. Sharp head the research team at the University of Missouri at Columbia which is developing the knowledge base in rheumatology using the EXPERT scheme. The experience with the design of the rheumatology model (Lindberg et al , 1980) has shown that the knowledge engineering tools (EXPERT and SEEK) and know-how developed so far make it possible to move incrementally and rapidly in the construction of a new medical knowledge base in collaboration with expert clinical researchers. Moreover, this experience is leading us to the development of a methodology for guiding the interaction of medical and computer science researchers in model building. The sequence of developments of consultation models follows a natural progression, aided at each step by an interplay between the clarification of medical concepts and the application of logical methods of model design. Our work in this area is contributing to a better understanding of a central problem in the application of Artificial Intelligence to the design of expert computer-based systems; namely, what are the representations, the processes, and the interface facilities that are needed to acquire, augment, and refine knowledge bases of different types by interacting with specialists in a domain.

A number of improvements in the EXPERT scheme have come about from the work in this domain. Because of the need to facilitate the acquisition of reasoning rules, we developed a criterion-based scheme for defining the diagnostic rules. A knowledge-acquisition front end for EXPERT prompts the model builder for the major and minor findings, necessary and exclusionary conditions for the presence of a disease at various levels of certainty.

In ophthalmology (Weiss, Kulikowski, and Safir, 1978), research includes a neuro-ophthalmological model developed by Dr. W. Hart at Washington University, a glaucoma model by Dr. Y. Kitazawa at Tokyo University, and an infectious eye disease model by Dr. Chandler Dawson at the University of California at San Francisco. The elaboration of a specialized treatment strategy and explanation model has received strong support through the infectious eye disease project.

In clinical pathology, we have collaborated with Dr. Robert Galen of Columbia University and Overlook Hospital in the development of models for the interpretation of serum protein electrophoretic patterns, thyroid tests, and LDH/CPK isoenzymes (Kulikowski, Weiss, and Galen, 1981), which are some of the most frequently used laboratory tests

**Expert Systems in Oil Exploration.** In this project we are concerned with problems of expert reasoning in oil exploration, concentrating on problems of well-log interpretation. Work is proceeding in close collaboration with experts at the Amoco Research Labs in Tulsa, Oklahoma. Well log interpretation is of critical importance in exploration and production in the oil industry. Once a well has been drilled, instruments are placed downhole and readings are recorded providing information on the geology below. This research takes advantage of existing Amoco software for graphically interpreting the logs. The application illustrates several fundamental characteristics of practical expert systems, including the interpretation of large amounts of signal data, the interrelationship of mathematical models and production rules in a consultation program, and the interfacing of advanced applications software to an interpretive system. In this system, the user is guided through an interpretation with low-keyed advice, being restrained from carrying out a particular analysis only if a necessary antecedent step has been omitted. However, should the intermediate results show inconsistencies, or perhaps other methods appear more appropriate, the system will point out where he could have done things differently and potentially arrive at more accurate results. This conception of an expert system is a much broader one than that commonly taken: not only does the knowledge base have to include the specifics of interpretations and actions in the domain, but it must also contain information about the methods for solving subproblems, and the sequencing of subproblems that constitutes expert utilization of existing analysis programs. The expert system then becomes a kind of master programmer that guides and advises the user.

**Methodologies and techniques for building expert systems: Developing Microprocessor Based Expert Models For Instrument Interpretation.** This project involves the development of expert interpretive systems and transferring them to a microprocessor environment (Weiss, Kulikowski, and Galen, 1981) The scheme has been successfully implemented and tested by producing a program for interpreting results from a widely used medical laboratory instrument: a scanning densitometer. Specialists in the field of serum protein electrophoresis analysis provided the knowledge needed to build an interpretive model using EXPERT. By constraining a few of the structures used in the general model, it was possible to develop procedures for automatically translating the model to a specialized application program and then to a microprocessor assembly language program Thus, the model development can take place on a large machine, using established techniques for capturing and conveniently updating expert knowledge structures, while the final interpretive program can be targeted to a microprocessor depending on the application. Our experience in carrying out the complete process illustrates many of the requirements involved in taking an expert system from its early development phase to actual implementation and use in a real world application.

**A System for Empirical Experimentation with Expert Knowledge.** SEEK is a system which has been developed to give interactive advice about rule refinement during the design of an expert system (Politakis and Weiss, 1981). The advice takes the form of suggestions for possible experiments in generalizing or specializing rules in an expert model that has been specified based on reasoning rules cited by the expert. Case experience, in the form of stored cases with known conclusions, is used to interactively guide the expert in refining the rules of a model. The design framework of SEEK consists of a tabular model for expressing expert-modeled rules and a general consultation system for applying a model to specific cases. This approach has proven particularly valuable in assisting the expert in domains where the logic for discriminating two diagnoses is difficult to specify, and we have benefited primarily from experience building the consultation system in rheumatology.

**A precedence scheme for selection and explanation of therapies.** A general scheme for treatment selection and explanation has been developed (Kastner and Weiss, 1981). This involves a topological sorting procedure within a general production rule representation. The procedure is used to choose among competing therapies on the basis of precedence rules keyed to various selection criteria. This approach has a degree of naturalness that helps in automatically producing explanations of the choices made. A system has been implemented using this approach which plans therapies for patients diagnosed as having ocular herpes simplex The reasoning model in this domain is undergoing testing, and the system has been generalized to handle models in other domains.

### Representation and Inference Facilities in AIMDS

*Participants:* N. S. Sridharan, Brian Lantz, John Bresina, John Roach

*Local User Group:* L T. McCarty, D Nagel, J. Goodson, C. F. Schmidt

*Remote User Group:* A Barbedienne, J-F Cloarec, J-F Cudelou

*Research Support* NIH (BRP), CNET

AIMDS is a programming environment (language, editors, display drivers and file system) in which several Artificial Intelligence programs are being constructed. The facilities offered in AIMDS are for specifying explicit meaning representation for concepts in a chosen domain in which we wish to construct reasoning programs The language has been under development for about five years, closely guided by its two main applications—psychological models of common sense reasoning and cognitive models of legal argumentation.

The main facilities offered by AIMDS include:

An assertional data base accessed by object descriptions;

Multiple data bases, supporting branching historical states;

Automatic consistency checking and automatic update;

User-callable deductive facilities, using concepts from the class calculus, relational algebra, and first-order logic;

A Description formalism, that provides a rich and interesting partial order among descriptions;

A RuleSet formalism, that allows associating descriptions with descriptions; Retrieval from a RuleSet permits inheritance over the partial order relation;

An Action formalism for invoking state transitions, checking preconditions and postconditions;

A variety of advanced programming constructs, including flexible data- and control-backtracking;

User customization of system facilities is one of the design objectives, assuring prospective and actual users of freedom from straightjackets

We feel the need to scale up the sizes of the knowledge bases and data bases that we use in our applications by an order of magnitude or more. The challenge of allowing such scale-up while enhancing performance, calls for innovations at the hardware level and at the operating system level. We are now engaged in the design of a multiprocessor for parallel execution of AIMDS/LISP.

## Understanding of Action Sequences and Computational Traces

*Participants:* Charles Schmidt, N. S. Sridharan

When one person observes another performing a purposeful sequence of physical actions, the observer is typically able in real-time to move from these fine-grained observations to an easily remembered and communicable understanding of the actor's plan. Our work has resulted in the development of a computational model of the plan recognition process as well as facilities for representing plans, actions, changes, states, and for simulating action execution (Schmidt, Sridharan, and Goodson, 1976; Schmidt, Sridharan, and Goodson, 1978). This work has been carried out within, and contributed to the development of, the AIMDS system framework (see above summary). As a result of this work, we are in a position to carry out detailed empirical studies of how the human information processing system accomplishes tasks which involve the understanding of rapidly occurring processes. The main focus of this research is to develop a methodology, and then use this methodology to empirically test detailed predictions about the human plan recognition process.

The second focus is to heuristically apply the predictions of this theory of plan recognition to the domain of program trace understanding. Within this second domain, the same methodology is being used to empirically investigate how persons move from the observations provided by a trace of a computation to an understanding of the computational process itself.

## The TAXMAN Project

*Participants:* L Thorne McCarty, N. S Sridharan, Donna Nagel, David Raab

*Research Support:* NSF

The TAXMAN project is an experiment in the application of Artificial Intelligence techniques to the study of legal reasoning and legal argumentation, using corporate tax law as an experimental problem domain In our earlier work (McCarty, 1977), in a system called TAXMAN I, we were able to construct computer models of the facts of corporate tax cases and the concepts of the United States Internal Revenue Code, so that the system could produce an "analysis" of the tax consequences of a given corporate transaction. Our current research is concerned with some theoretical questions which were left open in the earlier study. It is clear that the TAXMAN I system is inadequate as a model of legal reasoning and legal argumentation since it provides no facilities for representing the "open-texture" of most legal concepts and no facilities for modeling the "construction" and "modification" of legal concepts that occurs during the analysis of a difficult case. In the system we are now developing, called TAXMAN II, we are attempting to remedy these deficiencies and attempting at the same time to develop a cognitive theory of the patterns of argument adopted by lawyers and judges in the early years of the corporate tax code. We are currently testing the TAXMAN II model on several stock dividend and corporate reorganization cases decided by the Supreme Court in the 1920s and 1930s and the results so far are encouraging, although still quite incomplete.

The TAXMAN system is implemented at present in the AIMDS representation language (see previous research summary). Given this representational framework, the TAXMAN system is constructed by encoding the various concepts of corporate tax law into the framework of the abstraction-expansion hierarchy. The basic "facts" of a corporate tax case can be represented in a relatively straightforward manner: corporations issue securities, transfer property, distribute dividends, etc. Below this level there is an expanded representation of the meaning of a security interest in terms of its component rights and obligations: the owners of the shares of a common stock, for example, have certain rights to the "earnings," the "assets," and the "control" of the corporation. Above this level there is the "law": the statutory rules which classify transactions as taxable or nontaxable, ordinary income or capital gains, dividend distributions or stock redemptions, etc. We have demonstrated that the TAXMAN I system is capable of representing the full set of facts of an actual corporate tax case, such as United States v. Phellis, 257 U.S. 156 (1921), and capable also of representing the statutory rules and concepts which classify such cases as tax-free reorganizations under Sections 368(a)(1)(B), (C), and (D) of the Internal Revenue Code. (See, e.g., McCarty, 1977, 1980 ) Furthermore, as long as we confine our efforts to the general areas of corporate and commercial law, we believe that there are numerous practical applications for a system of this sort. We have discussed these possibilities in McCarty (1980).

However, the main goal of our present research is to move beyond the limitations of the original TAXMAN paradigm and to develop a more realistic model of the structure and dynamics of legal concepts In the TAXMAN II system, as in the TAXMAN I system, the precise statutory rules are represented as logical templates, a term intended to suggest the way in which a "logical" pattern is "matched" to a lower-level factual network during the analysis of a corporate tax case. But the more amorphous concepts of corporate tax law, the concepts typically constructed and reconstructed in the process of a judicial decision, are represented in the TAXMAN II system by a prototype and a sequence of deformations of the prototype (McCarty and Sridharan, 1981). The prototype is a relatively concrete description selected from the lower-level factual network itself, and the deformations are selected from among the possible mappings of one concrete description into another. We have argued that these "prototype-plus-deformation" structures play a crucial role in the process of legal argument, and that they contribute a degree of stability and flexibility to an emerging system of concepts that would not exist with the template structure alone. We have illustrated these ideas with a detailed analysis of Eisner v. Macomber, 252 U.S. 189 (1920), the early stock dividend case, and our analysis is now undergoing a full-scale implementation (McCarty and Sridharan, 1982).

### Automating Expertise in Digital System Design

*Participants:* Tom Mitchell, Louis Steinberg, Van Kelly, Patricia Schooley, Jeffrey Shulman, Timothy Weinrich, Reid Smith, Saul Amarel

*Research Support:* DARPA

The Digital Design Project, which began at Rutgers about a year ago, is concerned with applying Artificial Intelligence methods to problems of computer-aided digital circuit design, and with studying the nature of design problems in general. The current focus of this project is on constructing a computer program to help redesign existing digital circuits to accommodate desired changes to their functional specifications.

Basic computer science research issues arise in the task of automated circuit redesign. In order to redesign a given circuit, the relationship between the function and structure of that circuit must be represented in some useful manner. General knowledge about electronics, circuit analysis, and circuit design strategies must also be represented and used effectively Methods for managing the large number of interdependent constraints involved in design must be developed, as well as methods for satisfying these constraints during design and redesign.

In the initial stages of this work we have focussed on representation issues, and on the basic modes for reasoning about an existing circuit design. We have developed initial representations (Mitchell, Steinberg, Smith,, Schooley, Jacobs, and Kelly, 1981) for circuit topology, functions of circuit components, circuit timing and the flow of data streams through the circuit, and have used this representation to describe portions of a computer terminal circuit. We are currently developing a system, called CRITTER (Kelly and Steinberg, in press), which considers the topology of the circuit and the known functions of circuit components to answer questions such as "Given that the signal at point $X$ must have a period of 100 microseconds, what can be said about signals at other points in the circuit?" More generally, CRITTER is intended to perform digital circuit analysis by propogation of signals and signal constraints through the circuit. CRITTER forms the basis for a program that generates plausible localized circuit redesigns.

In addition to being able to represent and reason about the operation of a circuit, it is also important to be able to represent and reason more globally about the design plan of the circuit. The design plan details the way in which the functional specifications are decomposed and implemented by the current circuit, and thus contains information about the roles, or intended uses, of various circuit components. Each implementation step within the hierarchical design plan is characterized as an instantiation of some general rule of circuit design, such as "If you want to convert a parallel signal to the equivalent serial signal, then you may use a shift register." We are currently experimenting with alternative schemes for representing design plans, and are considering their utility for directing the redesign process.

Our research thus far suggests that the redesign task can be directed and constrained by relying upon

1. the kind of circuit analysis that CRITTER is intended to provide, and
2. examination of the interdependencies of implementation choices recorded in the design plan of the circuit.

Our near term research will focus on combining these sources of knowledge and others to automate portions of the redesign task.

### LEX: Learning Problem-Solving Heuristics by Experimentation

*Participants:* Tom Mitchell, Paul Utgoff, Ranan Banerji, Adam Irgon

*Research Support:* NSF

The goal of this research is to develop computer learning methods by which heuristic problem-solving programs may improve their performance through practice. The research has two major thrusts:

1. to design and analyze computer methods for self-improvement of problem solving strategies; and
2. to construct a working program to test our ideas empirically.

We have developed a heuristic search program, called LEX, that learns heuristics to guide its search in a particular task area: solving symbolic integration problems. This program is given a set of operations that can be performed on integrals (e.g., integration by parts, removing a constant to outside the integral, etc.). It is able to learn heuristics that recommend in which cases the various operations are likely to lead to solutions. One heuristic which is typical of those that LEX has learned may be paraphrased as follows: "If the integrand is the product of $x$ and any transcendental function of $x$, then try Integration by Parts with $u$ bound to $x$, and $dv$ bound to the transcendental function times $dx$."

The design of LEX consists of four distinct program modules:

1. the Problem Solver, which uses available heuristics to guide its search for solutions to practice problems;
2. the Critic, which analyzes successful solutions to isolate training examples of particularly useful or wasteful search steps;
3. the Generalizer, which formulates and later refines general heuristics, given the specific training examples produced by the Critic; and
4. the Problem Generator, which generates new practice problems for the system to try to solve.

In Mitchell, Utgoff, Nudel, and Banerji (1981), we describe a system based upon the first three of these modules, and results obtained by presenting the system with hand-generated practice problems. A prototype Problem Generator module is currently being implemented.

Our current research on this project is in three major directions.

**Combining deductive with inductive methods for inferring heuristics.** The Generalizer forms heuristics in an empirical, data-driven fashion, by finding features common to several problems in which the same operator is found useful. The generalization of heuristics is based on the Version Space approach of maintaining multiple plausible generalizations of the heuristic and refining this set as additional data becomes available. However, it is also possible to learn a great deal about the appropriate generalization of the heuristic by conducting a detailed analysis of a single solution trace in which that heuristic should apply. In Mitchell, Utgoff, Nudel, and Banerji (1981), we sketch how this kind of deductive analysis of solution traces can guide generalization of heuristics and suggest how this deductive analysis might be combined with the current empirical, inductive methods to produce a much stronger Generalizer which relies upon both deductive and inductive methods to infer heuristics.

**Automatically extending the language of heuristics** One of the most troubling aspects of current approaches to concept learning and generalization is that the language for stating concepts strongly biases the range of concepts that can be learned In LEX, for example, the language for describing heuristics is too restricted to correctly describe all the heuristics for which training data is observed. To handle this problem, the system must be able to

1. detect situations in which its current language cannot describe the appropriate heuristic, and
2. consider extending its language in some way that will allow it to describe that heuristic

Paul Utgoff is considering this problem area in his Ph.D. dissertation. He is considering several bases for generating new terms for the language: the observed data, biases in the current language, and the same kind of deductive analysis of problem solutions noted in the above paragraph.

**Designing and experimenting with the Problem Generator.** We are currently implementing two tactics for suggesting new training problems. These tactics are as follows:

1. find a proposed heuristic for which alternative descriptions are plausible, then generate a training problem that resolves among these alternatives, and
2. find two of the given operators whose preconditions overlap, but for which there is no heuristic that chooses between them, then generate a problem to which they both apply.

Both of these are tactics for generating practice problems that remove specific ambiguities in the system's heuristic problem-solving strategy. We plan to experiment with these tactics, and to also consider more global strategies for planning sequences of several practice problems.

## Expertise Acquisition through Shifts in Problem Representations

*Participant:* Saul Amarel

*Research Support:* NIH (BRP)

The problem of representations in problem solving is concerned with the choice of formulation of a problem for a system which is organized in accordance with a given problem schema. Key issues are the following:

1. understanding the relationships between such a choice and the efficiency with which the system can be expected to find a solution to the problem;
2. finding ways of choosing an "appropriate" problem formulation—from the point of view of minimizing the computational effort needed to construct a solution; and
3. finding ways of shifting from one problem formulation to another in a manner that increases problem solving performance

These are fundamental and difficult issues in AI which we have been studying for about fifteen years (Amarel, 1967, 1968, 1971).

Progress in this area has been stimulated in recent years by efforts in our other projects at Rutgers to develop AI methodologies for building expert systems and, in particular, by problems of expertise acquisition Expert behavior requires the conceptualization/formulation of a given task within a highly "appropriate" representational framework; and major improvements in expertise are often attained via appropriate shifts in problem formulation. We are now studying such processes of improvement in problem solving expertise via representational shifts. We recently developed a conceptual framework for handling problem representations in which the grammatical specification of solutions for a problem class plays an important role (Amarel, 1981). A shift in problem representation amounts to a change in solution grammar or/and to a change in the way in which problem conditions control the process of solution generation from a given grammar Within this framework, we are exploring representational shifts in several problems of reasoning about actions. For Tower of Hanoi problems, we considered a sequence of about ten shifts in formulation that effect transitions between heuristic production schemas, goal-directed relaxed reduction schemas, and reduction schemas that are based on macromoves. Each transition between formulations involves the acquisition and transformation of a certain amount of new knowledge. Such knowledge may originate from experience of previous problem solving activity or from analysis of basic concepts in the domain. Our recent work suggests that a system for mechanizing shifts in problem representations must be able to flexibly configure itself in accordance with various problem solving schemas and *to manage/coordinate simultaneously more than one representation* of a problem (Amarel, 1982). These are also desired characteristics of robust expert systems, that is, systems that can perform at peak performance in "central" parts of some domain and still perform in an acceptable manner in the domain's "periphery." Systems of this type are now being studied To obtain shifts in problem representations, a system must be able to perform a variety of theory formations and program synthesis tasks. Our goal for the near future is to focus on some of these tasks in the context of certain types of transitions between problem formulations that were identified in the Tower of Hanoi domain.

### The Representation of Knowledge in AI and in Database Theory

*Participant:* Ray Reiter

The major goal of this research is to provide logical foundations for representation theory in both AI and database theory. Recent results in these areas are summarized below.

**Database theory.** A reformulation of conventional relational database theory has been given in terms of first order logic (Reiter, in press). The reformulation provides proof theoretic definitions for query evaluation and the satisfaction of integrity constraints. In addition, it proposes a representation reflecting the "closed world assumption" and in terms of this provides a semantics for null values and disjunctive information. Finally, it indicates how this logical approach accommodates a natural generalization of the relational model to represent more real world knowledge.

**Non-monotonic reasoning in AI.** Reiter (1980) provides a formalization of default reasoning. Reiter and Criscuolo (1981) discusses a variety of problems regarding default inheritance in hierarchies. Recent work by D. Etherington and R. Reiter focuses on providing algorithms for inheritance of defaults in hierarchies of the kind treated by NETL (Fahlman, 1979), with the objective of proving these algorithms correct with respect to an appropriate default logic.

Circumscription (McCarthy, 1980) and predicate completion (Clark, 1978) provide forms of non-monotonic reasoning.

**Result.** For a wide variety of first order theories, circumscription implies predicate completion.

### Consultation System for Naval Mission Planning

*Participants:* C. V Srinivasan, David Sandford

*Research Support:* Naval Research Laboratories

We are using a representational framework called MDS (Meta Description System) to model the mission planning process in the U.S. Navy as a constraint satisfaction problem. The planning process is typically viewed as consisting of several stages (like "commander's interpretation of mission," "own force status," "enemy force status," "commander's decision," etc.). Our objective is to create a consultation system that can assist a mission planner by providing reasons for either supporting or refuting given interpretations of a mission or courses of actions chosen for the mission.

In the first stage of the project we are now creating a pilot system to model typical kinds of reasoning that occur in the planning process. The reasoning process is based on "residues" and a rule based system for the residue selection and constraint propogation. Work in the project is being conducted in close collaboration with the AI center at the Naval Research Laboratories, Washington D.C.

### References

Amarel, S (1967) An approach to heuristic problem solving and theorem proving in the propositional calculus In J F. Hart & S. Takasu (Eds ), *Systems and computer science* U. of Toronto Press

Amarel, S. (1968) On representations of problems of reasoning about actions. In D Michie (Ed ), *Machine intelligence 3* Edinburgh University Press

Amarel, S. (1971) Representations and modeling in problems of program formation In B Meltzer & D. Michie (Eds.), *Machine intelligence 6.* Edinburgh University Press.

Amarel, S (1981) Problems of representation in heuristic problem solving: Related issues in the development of expert systems Tech. Rep CBM-TR-118, Laboratory for Computer Science Research, Rutgers University

Amarel, S. (1982) Expert behavior and problem representations. Tech. Rep. LCSR-TR-25, Laboratory for Computer Science Research, Rutgers University.

Clark, K. L. (1978) Negation as failure. In H. Gallaire & J Minker (Eds.), *Logic and data bases* New York: Plenum.

Fahlman, S E. (1979) *NETL: A system for representing and using real-world knowledge.* Cambridge, Mass : MIT Press

Kastner, J. K., & Weiss, S. M. (1981) A precedence scheme for selection and explanation of therapies. *IJCAI 7,* 908–909

Kelly, V., & Steinberg, L (in press) CRITTER: A constraint propagation system for digital circuit analysis Paper submitted to the AAAI

Kulikowski, C. A. (1980) Artificial intelligence methods and systems for medical consultation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-2

Kulikowski, C. A., Weiss, S. M , & Galen, R. S (1981) Computerized diagnosis in the lab *Medical Laboratory Observer,* 41–57.

Lindberg, D., Sharp, G , Kingsland, L , Weiss, S., Hayes, S., Ueno, H , & Hazelwood, S. (1980) Computer based rheumatology consultant. *Proceedings of the Third World Conference on Medical Informatics,* 1311–1315

McCarthy, J. (1980) Circumscription—A form of non-monotonic reasoning. *Artificial Intelligence* 13:27–39.

McCarty, L T (1977) Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning *Harvard Law Review* 90:837–893.

McCarty, L. T (1980) Some requirements for a computer-based legal consultant. *AAAI 1,* 298–300

McCarty, L. T., & Sridharan, N. S (1981) The representation of an evolving system of legal concepts: II Prototypes and deformations *IJCAI 7,* 246–253

McCarty, L T , & Sridharan, N. S (1982) A computational theory of legal argument. Tech Rep. LRP–TR–13, Laboratory for Computer Science Research, Rutgers University.

Mitchell, T M., Steinberg, L., Smith, R G , Schooley, P , Jacobs, H , & Kelly, V. (1981) Representations for reasoning about digital circuits. *IJCAI 7,* 343–344

Mitchell, T. M., Utgoff, P E., Nudel, B., & Banerji, R (1981) Learning problem-solving heuristics by practice. *IJCAI 7,* 127–134

Politakis, P , & Weiss, S. (1981) A system for empirical experimentation with expert knowledge. Tech Rep., Computer Science Dept , Rutgers University.

Reiter, R (1980) A logic for default reasoning. *Artificial Intelligence* 13:81–132

Reiter, R (in press) *Towards a logical reconstruction of relational data-base theory.*

Reiter, R., & Criscuolo, G (1981) On interacting defaults. *IJCAI 7,* 270–276

Schmidt, C F., Sridharan, N. S., & Goodson, J L (1976) Recognizing plans and summarizing actions. *Proceedings of AISB Conference,* 291–306.

Schmidt, C F., Sridharan, N S., & Goodson, J. L. (1978) The plan recognition problem: An intersection of psychology and artificial intelligence *Artificial Intelligence* 11:45-83.

Weiss, S M., & Kulikowski, C A. (1979) EXPERT: A system for developing consultation models *IJCAI 6,* 942–947.

Weiss, S. M., Kulikowski, C A., & Galen, R S. (1981) Developing microprocessor based expert models for instrument interpretation *IJCAI 7,* 853–855

Weiss, S., Kulikowski, C , & Safir, A. (1978) Glaucoma consultation by computer *Computers in Biology and Medicine* 1:25–40

# ANNOUNCEMENT

## The AAAI Tioga Prize

The Tioga Publishing Company has donated $1000 to the AAAI to be awarded as a prize for the best paper presented at the 1982 AAAI meeting in Pittsburgh

The prize will be awarded as follows:

1 The program committee will make the selection.

2. The prize will be awarded to that paper that makes an unusually significant technical contribution to AI and that is also a shining example of elegant and clear exposition. The program committee will have the option of deciding that no paper this year meets these criteria. In that case, the prize money will be held over and (possibly) awarded at the 1983 AAAI meeting

3. The program committee will select several papers as candidates for the prize at the time it selects papers for the conference These finalists will be selected on the basis of the short papers only. Technical significance will be a primary factor in making the selection The authors will be notified that they are finalists in the contest at the time of notification of acceptance of their papers The winner will be selected on the basis of quality of exposition of the final (and perhaps longer) paper.

4 Multiple authors of a winning paper will share the prize equally In case of an unresolvable tie between papers, the several authors of the tied papers will also share the prize equally

5. When the winning paper is given at the conference, it will be announced that it is the winning paper and the prize will be awarded by the session chairperson.

(*Editor's note:* Tioga Publishing Company is not committing itself to more than just this year's prize Nils Nilsson has suggested that this could be the start of a tradition of awarding a "publisher's prize " Each year, the name of the publisher or publishers donating the prize for that year would be mentioned in the AAAI program. If any other publishers are interested, contact Nils Nilsson.)