AI Magazine Volume 6 Number 3 (1985) (© AAAI)

Introduction to the Comtex Microfiche Edition of the Rutgers University Artificial Intelligence Research Reports:

The History of Artificial Intelligence at Rutgers

Saul Amarel

Department of Computer Science and Laboratory for Computer Science Research, Rutgers University, New Brunswick, New Jersey 08903

Background and Overview

The founding of a new College at Rutgers in 1969 became the occasion for building a strong Computer Science presence in the University. Livingston College thus provided the home for the newly organized Department of Computer Science (DCS) and for the beginning of Computer Science research at Rutgers. I came to chair the department after ten years at RCA Labs in Princeton, where I headed the Computer Theory Research group. My own work in the Labs concentrated mainly in AI. In the early 1960s, I became interested in problems of representation in problem solving, and in computer methods for building models and solving formation problems. As I continued working on these problems at Rutgers, their central significance for AI became increasingly clear to me; so was their difficulty. A colleague from RCA Labs, Chitoor V. Srinivasan, came with me to DCS in 1969 with an interest in description languages and frameworks for knowledge representation. The early stages of AI research at Rutgers were necessarily influenced by interests that the two of us brought to the new department.

As the faculty of DCS grew, the diversity of AI interests and projects expanded, and the overall size of the effort increased significantly. By the late seventies, Rutgers was already recognized as one of the major AI centers in the US, with a broad range of AI activities—from the pioneering development of expert systems in a variety of domains to the study of basic issues of representation and inference. At present, our AI work is organized within the Laboratory for Computer Science Research (LCSR) which was formed in 1977 as the research arm of DCS under my direction. Work in LCSR covers all areas of Computer Science—with AI continuing to be the dominant area of activity. About fifty people with Rutgers appointments—faculty, graduate students and technical staff—and twenty five collaborators in other institutions are currently involved in various aspects of AI research in LCSR.

The historical highlights of the development of AI research at Rutgers, with emphasis on the major strands and themes that brought us to the current state of our research, naturally reflects my personal perspective on this development.

The present collection of reports covers most of our work in AI from 1970 through 1983. As in the case of CMU (Newell, 1984), there is no single numbering scheme for these reports. Three sequences of reports are identified by major (sponsored) projects within which they were generated, and two sequences are identified as DCS or LCSR reports. The largest sequence of TR's (technical reports) and TM's (technical memoranda) has the CBM prefix, which stands for our large "Computers in Biomedicine" Resource project. In all the sequences, there are gaps in numbering, as no sequence is exclusively devoted to AI.

While it would be unrealistic to attempt to place each author and report in this collection in historical context, I will outline the development of our major research activities in AI over the past fifteen years, with the aim of providing a framework within which the relationships between reports can be seen.

Let me start with some "pre-history," to provide pointers to the intellectual links with other AI centers in the country, and to explain some of the goals and approaches that shaped parts of the Rutgers AI program, especially the early work. In 1966, I spent a sabbatical at CMU, where I joined in a seminar with Al Newell and

I thank Kaz Kulikowski, Thorne McCarty, Tom Mitchell, Natesa S. Sridharan, and Sholom Weiss for comments on an earlier draft. I would also like to thank Priscilla Rasmussen for her efforts in getting the collection together

Herb Simon on representational phenomena in problem solving. This seminar focused on the role that representational choices play in problem solving, and it helped us identify some of the difficulties in mechanizing shifts of representation (Newell, 1966). I became convinced that we must develop a better understanding of theory formation processes in order to mechanize certain important types of representational shifts. I also became interested in interpretation processes, *e.g.,* medical diagnosis and their relationship to theory formation. In the years that followed, the links with CMU continued; we interacted on conceptual frameworks for problem solving, on issues of problem representation, and also on problems of expertise acquisition and theory formation.

Many problems of theory formation, as well as some problems of interpretation and design, are handled by solution methods where the main thrust of reasoning is from candidate solutions (hypotheses) to the conditions that a solution to the problem must satisfy. I call problems of this type "formation problems" (Amarel, 1971). These problems have received less attention than problems of derivation type where the solution construction process proceeds mainly by reasoning from problem conditions to (pieces of) candidate solutions. Since 1960 I have been working on formation problems in the context of a program design task, where the program is specified in terms of I/O pairs. This task has similarities with theory formation processes, and it provides a relatively simple model for studying problems of representation and inference in them. In the late 1960s, I became aware of the DENDRAL project at Stanford (Buchanan *et al.,* 1969), (Lindsay *et al.,* 1980). I saw DENDRAL as an excellent vehicle for studying problems of scientific theory formation, and for experimenting with different ways of using domain-specific knowledge in formation processes. I also saw similarities between DENDRAL (and MetaDENDRAL) and my work in program formation. In the course of following the DENDRAL project and other knowledge-based systems that Ed Feigenbaum and his HPP colleagues were building, I developed close links with the Stanford HPP group. We had many fruitful interactions in the decade of the seventies, especially in areas of AI applications, expert systems, and networks for national research collaborations.

Research Activities

The Rutgers Resource on Computers in Biomedicine

In the 1960s, my AI work at RCA Labs was supported by AFOSR. This support continued for about a year after my move to Rutgers. In 1970, with the encouragement of Bill Raub, who was heading the Biotechnology Resources Program of NIH at that time, we developed a proposal for the establishment of a Research Resource on Computers in Biomedicine at Rutgers. NIH approved this multidisciplinary effort, involving faculty from the fledgling Department of Computer Science as well as faculty from Biological Sciences, Psychology, Zoology, Ecology and the Medical School, and has been supporting it since 1971. The Resource has had a major impact on the stimulation and growth of AI research at Rutgers. I was principal investigator of the Rutgers Resource for its first six years, and in the following six years, Kulikowski and I were co-principal investigators. Currently, Kulikowski and Weiss are directing the Resource.

The broad objective of the Rutgers Research Resource was to bring advances in Computer Science—particularly in AI—into problems of biomedical inquiry. The new concept, relative to other Computer Resources at NIH, was to build a "people resource," where people with the help of computers would develop a computer-based framework for doing and applying research to important problems in biomedicine. The computer was to be used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in processes of experimentation and theory formation. An essential part of the effort was oriented to methodological problems of modeling and to the development of systems for the construction and use of models in a variety of tasks. Much of our work was concerned with processes of interpretation in medicine and psychology (Amarel, 1974). The philosophy of the project from the start has been to combine collaborative research on AI applications with core research on AI methodologies and with general system development activities that enhance the environment for AI research. The Resource started with six major areas of study. Work in three of the areas was in applications projects: Medical Modeling and Decision Making, Models of Belief Systems, and Modeling of Ecological Systems. The three remaining areas were core areas: Representations and Reasoning, Learning and Knowledge Acquisition, and AI System Development.

Medical Modeling and Decision Making: Expert Systems

From the very beginning, Kaz Kulikowski was the leader of research on Medical Modeling and Decision Making. He joined our faculty in 1970, with a Ph.D from the University of Hawaii, where his main interests were pattern recognition and medical decision making. Soon thereafter he began a close collaboration with Dr. Aran Safir, an ophthalmologist from the Mount Sinai School of Medicine. Joining them almost from the start was our first doctoral student in AI, Sholom Weiss. Their collaboration led to the development of the CASNET/Glaucoma system in the Resource (Weiss *et al.,* 1978). The system's task was to provide consultation for diagnosis and treatment of glaucoma—a chronic disease which is a leading cause of blindness in this country. At the heart of the system was a causal model of the disease in the form of a semantic net with weighted links (Kulikowski & Weiss, 1971). Another set of weighted links established associations between inputs to the system (*i.e.*, patient's findings) and parts of the causal model. Diagnostic conclusions were based on an interpretation of the patient's findings in light of the disease model. Treatment recommendations were guided by the diagnostic conclusions and by the individual pattern of findings in a patient. The term CASNET stands for "Causal Associational Network"; it was intended to characterize the types of links that enter in the network model. CASNET/Glaucoma was among the first few high performance expert systems in medicine whose development began in the early 1970s. Others were MYCIN at Stanford, PIP at MIT, and INTERNIST/CADUCEUS at the University of Pittsburgh (Szolovits & Pauker, 1978). All the (currently) well-known problems in large expert system development were encountered in building the Glaucoma System- the knowledge acquisition bottleneck, the need for high level programming frameworks/environments, the problems of validation and explanation, and the difficulties of collaboration with domain experts. Work on these problems has shaped much of the research on expert systems at Rutgers following the Glaucoma project.

The basic structure of CASNET/Glaucoma was already in place in late 1972. A strong research prototype was developed and tested, and found to perform impressively at the Academy of Ophthalmology Symposium on Glaucoma in October 1976 (Lichter & Anderson, 1977). By this time, a national network of ophthalmology collaborators (called ONET) from five medical schools in the U.S. was organized by Safir, Kulikowski and Weiss; its members worked closely with Resource researchers in system development and testing. The Arpanet and Tymnet networks, through which computer facilities at Rutgers were accessible, provided the technical basis for a new mode of effective "collaboration at a distance." Subsequently, medical researchers from the University of Tokyo joined ONET. In 1974, Sholom Weiss completed his Ph.D thesis (Weiss, 1974) with Kulikowski and after a brief period at Mount Sinai with Dr. Safir, joined Kulikowski on a research appointment in the Medical area of the Resource. Their fruitful collaboration continues. Work on CASNET/Glaucoma stimulated two other doctoral dissertations: One by Santosh Chokhani, in 1975 on coordination of mathematical models and inferential models (of the type used in CASNET) in diagnostic reasoning (Chokhani, 1975), and another by Michael Trigoboff in 1978 on a general semantic-net framework, called IRIS, for building diagnostic consultation systems (Trigoboff, 1978). Chokhani's work reflected a continuing strand of interest at Rutgers on finding ways to combine quantitative and qualitative reasoning in handling real-world problems (in medicine, physics, engineering). This work clarified many of the difficulties involved in effective coordination of quantitative and qualitative reasoning.

The Glaucoma effort was followed by several projects to build medical consultation systems in other areas—neuro-ophthalmology, infectious diseases of the eye, rheumatology and clinical pathology. In parallel with these efforts, work in the Medical area focused on the development of general frameworks and methodologies for building expert systems. This led to the EXPERT system, a generalized knowledge representation scheme for building expert systems, exercising them with individual consultation problems, and testing and analyzing their performance on large numbers of problem cases (Weiss & Kulikowski, 1979). EXPERT has been operational on DEC 10/20 computers since 1978; subsequently, it became available on VAX, IBM computers and on MC68000-based systems. Its representational and user-interface capabilities have been under constant improvement in recent years. The types of consultation problems best suited for EXPERT are classification and interpretation problems (Weiss & Kulikowski, 1981a). The PROSPECTOR system developed at SRI and the EMYCIN system developed at Stanford are also knowledge-based systems that specialize in problems of the same type. However, there are conceptual as well as technical differences in the design of EXPERT, where emphasis was placed on efficient performance and on portability. These features enabled important technology transfer results to occur in recent years. In 1981, an EXPERT program for the interpretation of serum protein electrophoretic patterns was translated from the DEC-20 computer on which it was developed onto a ROM, programmed in an assembly language version, and incorporated into a commercial instrument (Weiss, *et al.*, 1981b). This was the first commercially available stand-alone instrument to perform a direct interpretation of lab test results using expert system methods. As part of an intense effort to develop and test a comprehensive rheumatology consultation system, in close collaboration with Drs. Donald Lindberg, Gordon Sharp and their colleagues at the University of Missouri at Columbus (Lindberg *et al.*, 1980), the system, called AI/RHEUM, was transferred in 1982 from our DEC-20 to a stand-alone microprocessor and placed in a realistic clinical environment for routine tests. Similar transfers of EXPERT-based programs from a research computer to micros were carried out, starting in 1982, for various ophthalmological consultation systems in collaboration with Dr. Chandler Dawson of the UCSF (Kastner *et al.*, 1983).

The work on methodologies for designing expert systems led to two Ph.D dissertations. One by Peter Politakis in 1983 on the SEEK system for giving interactive advice about rule refinement during the development of an expert system (Politakis, 1983), and another by John Kastner in 1984 on a scheme for therapy planning (Kastner, 1984). The SEEK system makes an important contribution to the crucial problem of knowledge acquisition. While developed and tested in the context of the consultation system in rheumatology, its significance goes beyond medical applications of AI.

The experience obtained in the Medical Area and the development of the EXPERT-based programming environment for building expert systems led to a broadening of the activities of the group led by Kulikowski and Weiss, both in terms of exploration of new non-medical application domains and of research on general methodologies/techniques for expert system design. In 1981 a project started on expert reasoning in oil exploration, concentrating on problems of well-log interpretation (Weiss *et al.,* 1982). Supported by Amoco, this work proceeded in collaboration with domain experts at the Amoco Research Labs in Tulsa, Oklahoma. An important part of this project involves the development of a system, called ELAS, for managing a many-model consultation system in this domain. The knowledge base of the system includes not only the specifics of interpretations and actions in the domain, but also information about methods for solving subproblems, and the sequencing of subproblems that constitutes expert utilization of existing pieces of software for conventional modeling/simulation and signal analysis. This project, which resulted in a recent Ph.D dissertation by Chidanand Apte, has provided a rich task environment for the study of problems of coordinating qualitative-quantitative reasoning (Apte, 1984). Another non-medical project in the expert systems area, started in 1982 with IBM support, is concerned with the diagnosis of disk drive failures In 1983, a Center for Expert Systems Research (CESR) was created within LCSR, under the direction of Kulikowski and Weiss, to be a focus for the variety of projects growing out of the expert systems work in the Resource and from recent collaborative efforts with industry. The experience of over ten years of research by the expert systems group was recently compiled in *A Practical Guide to Designing Expert Systems*(Weiss & Kulikowski, 1984).

Modeling Belief Systems
Plan Recognition and Generation
Other Modeling Activities

A second major area of study in the Resource concerned Models of Belief Systems and Commonsense Reasoning in Planning. The key researcher in this area since 1971 was Chuck Schmidt, a social psychologist who subsequently received a joint appointment in Computer Science. The initial goal was to use the computer for developing a theory of how people arrive at an interpretation of the social actions of other people and to test the theory in certain communication situations. In 1974, Natesa S. Sridharan joined our department and began a close collaboration with Schmidt. Sridharan's previous research was on AI applications to chemical synthesis (at Stony Brook) and to analytical chemistry (at the DENDRAL project in Stanford). His interests included issues of representation and reasoning in problem solving.

The initial models developed in the area of Belief Systems were based on extensions of existing psychological ideas on social perception and were tested over interpretations of episodes involving three to four people in sequences of about a dozen actions. This work uncovered the complexity of some of the methodological problems in this area and led to a focus on models of Plan Recognition involving a single actor performing physical actions. The goal was to develop a model of an observer arriving at an hypothesis about an actor's plan from descriptions of the actions performed by the actor; furthermore, the model was to be consistent with data obtained from psychological experimentation that proceeded in parallel with the model building effort The Plan Recognition model developed was based on a "hypothesize and revise" process (Schmidt & Sridharan, 1977). The process of forming an initial hypothesis and its subsequent revisions bears a close relationship to the process of plan generation, execution monitoring, and plan repair. This model was formulated and tested on the AIMDS system, which was being developed in parallel by Sridharan as part of the core work in the Resource (Sridharan & Hawrusik, 1977). AIMDS is a high-level language and programming environment for building AI systems whose development was strongly motivated by the Plan Recognition effort.

The work on Plan Recognition led to a broader research effort by Schmidt and Sridharan in the area of planning problems. Starting in about 1980, they concentrated on human planning behavior and on the development of systems for generating and manipulating large and complex plans in situations where knowledge is incomplete. The goal was to obtain systems with qualities of robustness that characterize humans when they plan in realistic task domains (Sridharan & Bresina, 1982), (Sridharan, Bresina, & Schmidt, 1982). From 1981 to 1983 several interactive and flexible plan generating systems were developed.

The research on Belief Systems and on psychological models of plan recognition and generation was supported by the Rutgers Resource for ten years, until 1981. In 1982, Schmidt started a new project, under NSF support, oriented to theoretical extensions and to tests of his cognitive models in specific tasks of plan monitoring and plan recognition. In these tasks, a subject typically observes on a CRT screen a trace of the problem-solving performance of an intelligent agent; responses collected during this monitoring process allow for the analysis of the subject's interpretation processes. This work continues to provide an important linkage between our AI concerns and problems in cognitive psychology. It is also stimulating new work in the general area of man-machine communications.

The initial set of applications projects included Models of Ecological Systems, in addition to Medical Modeling and Models of Belief Systems. The goal was to mechanize part of the model construction process of an ecologist. The particular task considered was as follows: given data taken from a lake ecosystem, develop a compartmental model describing the circulation of a mineral (*e.g.,* phosphorus) in the ecosystem. By 1974, progress became slow due to difficulties in maintaining effective links with domain experts, and the project was terminated. This was one of several experiences that convinced us of the crucial role of good collaboration with domain experts on applications-oriented AI research. Work on AI methods in support of mathematical modeling resumed in the late 1970s focusing on enzyme kinetics and on respiration; it is being done in the context of biomedical applications in the Resource.

Core AI Work in the Resource
Representational Frameworks and AI Languages
Legal Reasoning

The core areas of the Resource have provided the environment for a great variety of AI studies and system developments that were not directly tied to specific application projects. Even though the individual core activities cannot be neatly classified, it is possible to discern a few stable orientations along which projects clustered over more than a decade. These are: Representational Frameworks and related Reasoning Approaches, AI Languages, Learning and Expertise Acquisition, Theory Formation, and Problem Solving Methods. I will describe projects in each of these areas, jumping whenever appropriate between areas to show linkages between projects, and to trace the evolution of activities that grew from seeds planted in the Resource.

The research of Srinivasan on Meta-Description Systems (MDS) was one of the early core activities. It received support from the Resource grant for about six years. MDS is a representational framework, with associated facilities for reasoning, that embodies many of the characteristics of AI knowledge representation systems based on structured objects. Systems of this type received widespread attention in the field after Marvin Minsky's "frames" paper (Minsky, 1975). The main ideas of MDS designs preceded Minsky's paper by several years. The framework includes templates for structuring knowledge in a domain, processes for instantiating templates, and procedures for managing consistency conditions in models within the domain (Srinivasan, 1976). In connection with the MDS framework, Srinivasan developed notions of "coherent information systems" and of architectures for general problem solving. Key parts of MDS were designed in the early 1970s, and considerable effort was devoted to their implementation. It became clear early on that high levels of programming support and computing power were needed to implement and test MDS designs. The grant support we received from DARPA, starting in 1973, provided some of the badly needed computer resources.

In the 1970s, our AI work grew without any substantial support from DARPA. In this respect, we were an anomaly relative to other major concentrations of AI work in the country (in particular, CMU, Stanford, MIT, and SRI) where AI received massive support from DARPA. Several contacts with DARPA in the early 1970s resulted in a project on Secure Operating Systems and Automatic Programming. Srinivasan and I were Co-PIs of this project, which was supported for four years starting in 1973. Parts of the Automatic Programming component of this research had a certain amount of AI content, consisting mainly of some aspects of Srinivasan's work in MDS and of parts of my work on program formation from I/O pairs. A report series with a prefix SOSAP covered work in the project.

The MDS-based activity in the DARPA grant centered on the concept of knowledge-based programming. In connection with this effort, new theoretical work on theorem-proving techniques, based on Gentzen's systems of logic, was started. The goal was to use a theorem prover in MDS in the model updating and template instantiation processes. Several technical reports, issued in the period from 1975 to 1980 (as parts of the CBM series produced by the Resource, of the SOSAP series produced by the DARPA project, and of the DCS series), describe research in this area. These reports are included in the present collection. David Sandford started work in 1977 with Srinivasan on Theorem proving in MDS; he completed his Ph.D in 1979 on a semantic refinement of theorem proving by resolution and on a theory of Model Specification (Sandford, 1979).

As part of the DARPA grant, a TIP to the Arpanet was installed at Rutgers in 1973. This permitted us to access the ISI PDP-10 machines (and later the SUMEX-AIM machine), thus obtaining the needed cycles for MDS implementation and testing. The link to the Arpanet also had a strong impact on all our other AI projects. Even though direct support by DARPA for our AI research in the 1970s was not significant, the DARPA grant brought us into the Arpanet community and facilitated our contacts with ongoing development in the field. In 1974 the University installed a PDP-10 for support of our research. A year later it was linked to the Arpanet. Our research computer has been linked to the Arpanet continuously since that time, except for a one year period (in 1978) when we were off the net.

Research on MDS has continued in recent years both on conceptual issues (Srinivasan, 1981), and on a system implementation in the context of a naval mission planning task. This work is being done in collaboration with (and support from) the AI Center of the Naval Research Laboratories in Washington, D.C.

In 1971, Bertram (Chip) Bruce joined our department after completing his Ph.D with Bob Simmons at the University of Texas in natural language processing. His interests included language understanding processes and knowledge representation frameworks with capabilities for handling time. His work on the CHRONOS knowledge representation system, on Case Systems for natural language, and on belief-guided language understanding (Bruce, 1975) was done as part of the core research in the Resource. Bruce made an important contribution by introducing the LISP culture to Rutgers. CHRONOS was the first relatively large AI system implemented in LISP in the Resource (Bruce, 1973). Another researcher who contributed to the early buildup of our LISP programming environment and of our research computer facilities was Gil Falk, who came to the department in 1970 with a Ph.D from Stanford in the area of Vision. He was part of a small group that brought time sharing to Rutgers. Subsequently, he became interested in computer networking and moved to BBN in 1974. Chip Bruce also moved to BBN in 1976, where he concentrated further on problems in the analysis and synthesis of discourse.

In 1974, Rick LeFaivre joined our department with a Ph.D from the University of Wisconsin in high level AI languages and fuzzy reasoning. He transferred his FUZZY language (a PLANNER-like language with capabilities for manipulating degrees of uncertainty) to Rutgers, continuing to work on it within the core of the Resource. Rick LeFaivre brought UCI LISP to Rutgers, introducing useful modifications to it (thus creating RUTGERS/UCI LISP), documenting the new dialect, and providing the support needed to make it widely used in both research and graduate work (LeFaivre, 1977). In 1978, LeFaivre moved to the Tektronix Research Labs on the West Coast.

As previously mentioned, the system framework for model development and testing in Belief Systems and Planning was provided by Sridharan's AIMDS (Sridharan *et al.,* 1983). The development of AIMDS was a major core activity in the Resource. It lasted for about five years, starting in 1976, and its early stages were closely guided by the Belief Systems application, and in particular by the Plan Recognition task. AIMDS was viewed initially as a high-level programming language which is especially suitable for building models of Action Interpretation; hence the first two letters in the system's name. Many of the initial concepts that entered in the design of AIMDS are based on Srinivasan's MDS ideas (accounting for the remaining letters). LeFaivre's RUTGERS/UCI LISP and parts of FUZZY provided the initial programming environment for implementing AIMDS. As experience with AIMDS grew, the system was increasingly seen as a general, high-level, programming environment for building AI systems (Sridharan, 1980). Starting in 1981, new applications of AIMDS were explored. The most important among them was Legal Reasoning. Another application, carried out in collaboration with researchers from the Centre National d'Études des Télécommunications in France, was in the area of engineering design concerning the specification of the behavior of switching systems.

Experience with the use of AIMDS showed that memory limitations reduced its usefulness in applications involving moderately large knowledge bases. In 1982, AIMDS was converted to run in ELISP, an extended-addressing version of RUTGERS/UCI LISP, which was developed by Chuck Hedrick for the DEC-20. This conversion removed a good part of the constraints on effective use of AIMDS. However, Sridharan continued to look into the possibilities of further radical performance improvements by exploring hardware and operating system designs that would permit parallel execution of AIMDS (Roach & Sridharan, 1982). Recently, Sridharan has moved to BBN where his work on parallelism in AI systems is continuing.

In the late 1970s, Sridharan established a collaboration with Thorne McCarty, a Professor of Law at SUNY Buffalo, on approaches to the tudy of Legal Reasoning using AIMDS. In the first stage of this collaboration, McCarty's TAXMAN I system was implemented in AIMDS. In this system, facts of corporate tax cases and concepts of the US Internal Revenue Code were represented, so that analysis of the tax consequences of a given corporate transaction could be produced by machine. The next stage involved the design and implementation of TAXMAN II, with the objective to develop a more realistic model of "open textured" legal concepts in the form of a prototype and a sequence of deformations that are dynamically built in the course of constructing a legal argument (McCarty & Sridharan, 1981). While the AI system framework used was based on core work in the Resource, this project opened a completely new area of application for us with representational problems that were especially challenging. Thorne McCarty came to Rutgers as an Distinguished Visiting Professor in 1981, joining us on a permanent basis as Professor of Law and Computer Science in 1983. Since 1982, research on a Computational Theory of Legal Argument received support from NSF. A special sequence of technical reports, with a LRP prefix, has been issued covering work in this area since 1979; which is part of the present collection. Several graduate students are now working on problems of knowledge representation. Recently, Susan Epstein completed her Ph.D in this area with Sridharan (Epstein, 1983).

Ray Reiter, from UBC, spent the year with us in 1982 and planted seeds of activity in logical foundations for representation of knowledge in AI and in database theory. He initiated a project on a logical reformulation of conventional relational database theory, with a focus on issues of incomplete knowledge, which is now continuing (with NSF support) in collaboration with Naftaly Minsky from our faculty (Reiter, 1983). Perhaps the most important result of this collaboration is a strengthening of the conceptual links between AI and work on databases at Rutgers.

Learning and Expertise Acquisition
Theory Formation
Problem Solving Methods

Let us return now to other core activities of the Rutgers Resource. Related to Bruce's research on natural language processing, we had several projects concentrating on language learning by machine. Bill Fabens joined us in 1970 with a Ph.D from the University of Wisconsin, and an interest in language analysis. Within the Resource, he developed a relaxation parser, called PEDAGLOT, used to learn grammar rules for completing or refining an "almost correct" grammar of a language on basis of presentations of strings in the language. A variant of PEDAGLOT provided the basis for interpreting descriptions of episodes in Schmidt's early psychological models of plan recognition. Bob Smith arrived in 1977 with a background in philosophy and computer-aided teaching/learning and language processing. The research of two of his students, Vic Ciesielski and Don Smith, was part of the core activities. Don Smith's Ph.D, completed in 1982, was on a method for language learning from examples which is guided by the learner's assumptions about the strategy of presentations used by the teacher. He built the FOCUSER system to test and demonstrate ideas in this area (Smith, 1982). Ciesielski's doctoral research, completed in 1980, concentrated on machine understanding of natural language text in a specialized domain (Ciesielski, 1980). He developed his methodology in the context of a natural language "front end' that accepts a description of a glaucoma case and transforms it into input that the CASNET/glaucoma expert system can understand. In 1974, Adrian Walker joined our department with an interest in formal languages and methods for grammatical inference. His research, which included methods for inferring stochastic grammars, was done within the core area. Fabens, Bob Smith and Walker are no longer at Rutgers; Don Smith, now a member of our faculty, is concentrating on AI approaches to debugging VLSI circuits.

Tom Mitchell joined us in 1978 after completing his Ph.D at Stanford with Bruce Buchanan in the area of concept learning by machine. His main interests were in machine learning and in applications of AI to problems of engineering design. At Stanford, he developed the concept of Version Spaces and applied it to the design of MetaDENDRAL (Mitchell, 1978). Shortly after coming to Rutgers, he initiated the LEX project, whose goal is to develop machine learning methods by which heuristic problem solving programs may improve their performance. Starting as a core activity, this project was subsequently supported jointly by the Resource and by a two-year NSF grant received in 1981. The task domain of the (forward reasoning) problem solver in LEX is symbolic integration, and the improvement in performance of the problem solver is obtained via learning good applicability conditions for the domain operators that are used during the search for solution. One method of learning that was studied, where the Version Space approach is used, is based on empirical generalizations from solution traces (Mitchell *et al.*, 1981). A second method is based on a detailed analysis of a single solution trace (Mitchell, 1984). In his "Computers and Thought" lecture in IJCAI-8 at Karlsruhe, Mitchell presented both methods within a unifying problem-solving paradigm (Mitchell 1983). The LEX project, along with related studies on machine learning, is covered in technical reports (starting in 1979) from the CBM, DCS, and LCSR sequences.

In 1984, a proposal was prepared for a major effort in learning and expertise acquisition, based on extensions of methods incorporated in LEX and their application in a new planning and design task. We intend to develop what we call a Learning Apprentice: An interactive knowledge-based consultant which directly assimilates knowledge in the course of its use, by observing and analyzing the design choices made by an expert user. This will be a learning component for a VLSI Design Consultant system that is currently being developed. This research is currently supported by a DARPA grant, with Tom Mitchell, Lou Steinberg and myself as Co-PIs.

From the beginning of the Rutgers Resource, my work on problem representations and theory formation was a part of the core activity. In recent years, after exploring characteristics of expert problem solving behavior (Amarel, 1980), I concentrated on processes of expertise acquisition that can be obtained via shifts in problem representation, *i.e.*, via problem reformulation (Amarel, 1981). I analyzed representational shifts in Tower of Hanoi problems that result in increased problem solving power in restricted classes of the domain and looked into the nature of knowledge and reasoning processes needed to bring about "appropriate" problem reformulations. This analysis showed that a wide range of theory formation capabilities are needed for extracting knowledge from problem solving experience so that it can be used to obtain problem formulations of increased power. In view of the centrality of theory formation problems in issues of reformulation and expertise acquisition, I recently returned to work on theory formation in the context of program design from I/O pairs. This work goes back to the early 1960s, but the current emphasis has shifted from a previous top-down approach (Amarel, 1971), to one that is primarily bottom-up and gives major weight to a detailed analysis of individual I/O pairs (Amarel, 1983).

Based on a proposal prepared by Sridharan, a related project started in 1984 on an experimental study of methods for problem reformulation. The goal is to create an environment for exploring processes of problem reformulation that involve the construction and modification of reduction operators. An important technique to be used is Plan Recognition, which is based on our earlier work on Planning. Research in this area is now supported by NSF, with Chuck Schmidt and myself as Co-PIs. As part of this work, a flexible framework for problem solving, which can embody a variety of reasoning methods, is being implemented. Several of us (especially Sridharan and I) have directed attention to this area in the past. One of the issues that comes up here is how to organize a knowledge base of operators and of other relevant information for controlling search processes in problem solving. In 1979, Larry Welsch completed his Ph.D with Sridharan on a method for "automatic synthesis of questions" in the context of a knowledge base from which relevant information was to be extracted for a problem-solving task (Welsch, 1979). Other recent research in problem solving methods includes work on constraint satisfaction problems and related analyses of computational complexity of search algorithms (Nudel, 1983).

Workshops and National Collaborations in the Rutgers Resource

From the very beginning, the activities of the Rutgers Resource included national collaborative efforts and information dissemination in AIM (Artificial Intelligence in Medicine). In 1973, the national SUMEX-AIM project was established at Stanford under NIH sponsorship. Its main goal was to provide hardware and software shared resources to AIM researchers at Stanford and (via computer networks) to AIM projects in other parts of the country. We worked closely with Stanford in the early stages of development of the SUMEX-AIM facility, and we participated in its governing committees from the start. The Rutgers Resource was one of the first national projects that accessed the Stanford computer facility, first via Arpanet and later via Tymnet. In 1974, our Resource was assigned by NIH the responsibility of organizing a series of annual national AIM Workshops. Their purpose was to bring researchers in AIM from all parts of the country together to share experiences on questions of content, methodology, direction, and organization problems of AIM community building. The first Workshop was held at Rutgers in June 1975. Its Proceedings appear as a technical report in the CBM sequence of this collection (Kulikowski & Sridharan, 1975). The Proceedings of two subsequent workshops are also parts of the CBM sequence. These Proceedings provide a good overall picture of developments in AI applications in the mid-1970s, showing the pioneering role of medical expert systems in these developments, and they record opinions/views of key investigators from virtually all the active AI applications projects in the country during that period. Reports presented at the Fourth AIM Workshop in 1978 on basic themes in AI research (Amarel, 1978) and on the relationship between AI and Psychology (Schmidt, 1978) are also part of the CBM sequence. In 1977, the research computing facilities at Rutgers were augmented with the help of NIH. A DEC-20 system was installed, providing the basis for an eastern node in the AIM network, called RUTGERS-AIM. National projects could now use the resources of both the Stanford and Rutgers nodes. In recent years, the national collaborative activities that are associated with the RUTGERS-AIM node include projects from the University of Pittsburgh, Ohio State, Harvard, Yale, and Penn. In the late 1970s, several members of the Rutgers Resource (mainly graduate students) collaborated with Stanford researchers in the preparation of the *Handbook of Artificial Intelligence*, an encyclopedic compilation of AI topics. A report covering the Rutgers contributions appears as part of the CBM sequence in the present collection (Ciesielski, 1979).

AI in Design

AI in Design is perhaps the single major area of current research at Rutgers whose origins cannot be traced to previous work in the Resource. It grew out of Tom Mitchell's interest in problems of digital circuit design. In 1980, we went back to DARPA to explore possible support of our AI research in several areas, including design. In 1981, we received a 2-year grant for research on methods for Automating Expertise in Digital System Design, with Tom Mitchell and myself as Co-PIs. Lou Steinberg, who came to our department in 1978 with a Stanford Ph.D, joined this project from the beginning. He had worked at Stanford with Cordell Green, and his interests were in AI approaches to software design and man-machine interaction. The focus of the project was to develop a system, which was later called REDESIGN, to provide interactive aid in the functional redesign of digital circuits. Given the description of a circuit and its functional specifications (*e.g.*, a computer terminal circuit and its function), REDESIGN would aid in the redesign of the circuit to meet altered functional specifications (*e.g.*, to alter the font used to display characters on the terminal screen). The work on REDESIGN required developing methods for representing and reasoning about digital circuits in ways that take into account the functional role of circuit submodules, and the decision that went into the original design of the circuit (Mitchell *et al.*, 1981). Two modes of reasoning were found essential to guiding redesign. One involves reasoning about propagation of signals and signal constraints through the circuit. A system called CRITTER was developed to implement this reasoning (Kelly & Steinberg, 1982). Work in this area led to a recent Ph.D dissertation by Van Kelly (1984). A second mode of reasoning is based on examining the design plan for the original circuit to determine the impact of replacing one component by another. The design plan summarizes the steps of progressive decomposition and implementation of the original specifications by the given circuit, the conflicts that arise from these design steps, and the purpose of the various submodules within the circuit. This work has shown us some of the elements that must go into a general framework for design. Extensions of this research into the VLSI area led to a project on Intelligent Aids for VLSI design. This work is supported by a three-year DARPA grant that started in 1983, with Tom Mitchell and Lou Steinberg as Co-PIs. One thrust of the project is to develop an interactive consultant, called VEXED, to aid in designing cells and arrays of cells of VLSI circuits. VEXED begins with the functional specifications of the cells and constraints on their interconnections and is intended to produce a design at the sticks level. As I indicated previously, we started work on a learning component for VEXED that will aid in acquiring and assimilating design rules in the course of the system's use by observing and analyzing choices made by an expert designer. Research on this learning component, which we call the Learning Apprentice, is being supported by a separate DARPA grant. A second thrust of the AI in VLSI project involves the development of an aid to assist in debugging VLSI circuits. One thesis of this research is that debugging is best approached by considering design and debugging as interrelated problems.

The research on Design has been an important factor in a major new development at Rutgers in the past two years: The establishment of the Center for Computer Aided Industrial Productivity (CAIP). This is part of a comprehensive effort by the State of New Jersey and the University to strengthen collaborative research between academia and industry in "high tech" areas. In addition to Computer Science, several Engineering disciplines as well as Operations Research and Mathematics faculty are participating. AI approaches to engineering, manufacturing, and management will be important parts of the research at CAIP. I believe that CAIP will strengthen our basic and applied work in Design and will stimulate work in components of AI that are new to us, *i.e.,* perception, manipulation, and their integration into complete robotic systems.

Computing Environment

By way of concluding this historical sketch, I would like to reiterate key points in the evolution of our research computing environment. Building a vigorous program of AI research at Rutgers was contingent on the availability of adequate computing resources. A number of people at Rutgers, most prominently, Saul Levy, Gil Falk, Bob Smith, and Chuck Hedrick, expended considerable and sustained effort in realizing our current level of computing power. Saul Levy is a member of our faculty who has been the computing coordinator for Computer Science for over eight years. He has been deeply involved in planning for our computing facilities.

We started the 1970s with no local research computing, and we are now at a point where the computing resources available to us compare favorably with those of other major AI research centers. In the early 1970s our computing was done on commercial time sharing services. In 1973, the installation of a TIP at Rutgers, courtesy of our first DARPA grant, permitted us to access PDP-10's at ISI and later at SUMEX-AIM. An important breakthrough was made in 1974 with the installation of a local PDP-10 by the University in support of our research. The Rutgers PDP-10 was linked to the Arpanet in January 1975. In 1977 our research mainframe was augmented to a DEC-20 with joint funding from the University and NIH. This system was linked to Tymnet as well as to Arpanet, and it received several augmentations (mainly memory) in recent years. The support of NIH—for the installation and augmentation of our DEC-20, for its link to Tymnet, and for sponsoring our link to Arpanet in the period between our DARPA grants—was a crucial element in the building of the computing environment for our research. Bill Baker, from the Biotechnology Resources Program of NIH, played a key role in these developments. At present, our DEC-20 (called locally the LCSR Red machine) is networked via Ethernet to several smaller computers and personal workstations. In the future, I believe that the main mode of augmentation of our computing resources will be in the form of additional personal workstations, such as LISP machines. Several years back, in 1978, the University formed a group of technical staff for operating our LCSR computer facilities. The head of this group, Chuck Hedrick, played a key role in shaping the computer environment for our research. Having joined our department in 1977, with a Ph.D from CMU where he worked with Herb Simon in machine learning, he decided a year later to devote full time to the development and management of our computing facilities. In addition, he is directing several DEC supported software development projects, including the implementation of extended addressing LISP and PASCAL on the DEC-20 and the design of interfaces between intelligent terminals and a DEC-20 mainframe.

Concluding Comments

As a result of a reorganization of the Rutgers New Brunswick campus in 1980, the Department of Computer Science became a unit of a unified Faculty of Arts and Sciences. LCSR continues to provide the supportive environment needed for ongoing research and the energy required to initiate work in new areas. The Rutgers Resource can claim several unique accomplishments in AI applications and has stimulated work in a number of new projects. The Resource is now concentrating on expert systems and their application to medicine and biology, and it has become the principal component of a recently organized Center for Expert Systems Research. Our growing activities on AI in Design have provided an important stimulus for the establishment of a new multi-disciplinary Center for Computer Aided Industrial Productivity at Rutgers, where AI is expected to play a key role.

Recently, I stepped down as Chair of Computer Science, after a fifteen-year period during which I saw an enormous growth in our educational and research activities. AI has gained substantial strength during this period, and it continues to be the most active area of research in the Rutgers Computer Science community. The major orientations of our current research are Expert Systems, Design and Planning, Learning and Theory Formation, and Issues of Representation. Our approaches range from the analysis of fundamental AI problems to the exploration of new applications and technology transfer. The multidisciplinary nature of our application projects will continue to induce rich connections with a variety of professional and scientific areas outside Computer Science, especially with medicine, engineering design, law, and management. A promising new development is the growth of links with other areas within Computer Science, in particular, with VLSI design, database theory, and design and analysis of algorithms. I believe that AI has much to contribute towards increasing the conceptual coherence of Computer Science. Our graduate program has a rich AI component. In addition to the regular courses, projects, seminars and colloquia, there are an increasing number of informal seminars that take place on short notice in many corners of the Hill Center where our Computer Science community lives. One of the things that I am especially pleased about is that graduate students are very much a part of this community; they make a major contribution to the intellectual climate in which our work is being done.

There have been many achievements in the discipline of AI over the past fifteen years, some of which resulted from work at Rutgers. There are still many basic problems that need serious attention—to provide a better scientific basis for the discipline, and to maintain the momentum of current developments. Several of these problems have been with us for over twenty years, *e.g.*, representation problems, while others emerged this last decade, *e.g.*, methodologies for expertise acquisition. I intend to spend more time working on some of these basic problems.

References

Amarel, S (1983) Program Synthesis as a Theory Formation Task—Problem Representations and Solution Methods Report CBM-TR-135, LCSR, Rutgers University To appear in Machine Learning: An artificial Intelligence Approach Michalski, Carbonell (Eds.), Mitchell Morgan Kaufmann, Los Altos, CA.

Amarel, S (1981) Problems of Representation in Heuristic Problem Solving: Related Issues in the Development of Expert Systems. Report CBM-TR-118, LCSR, Rutgers University Also appears in *Methods of Heuristics.* Groner, Groner and Bishof (Eds.), Hillsdale, NJ: Lawrence Erlbaum.

Amarel, S (1980) Initial Thoughts on Characterization of Expert Systems. Report CBM-TM-88, LCSR, Rutgers University

Amarel, S (1978) Basic Themes and Problems in Current AI Research Report CBM-TR-91, LCSR, Rutgers University

Amarel, S. (1975) Computer-based Modeling and Interpretation in Medicine and Psychology The Rutgers Research Resource in Computers in Life Science Research. Siler and Lindberg (Eds), Faseb and Planum Also available as Report CBM-TR-29, LCSR, Rutgers University, 1974

Amarel, S (1971) Representations and Modeling in Problems of Program Formation *Machine Intelligence 6*, Meltzer and Michie (Eds.), University of Edinburgh Press Also available as Report DCS-TR-5, LCSR, Rutgers University, 1970

Apte, C. (1984) A Framework for Expert Control of Interactive Software Systems Ph D Dissertation, Rutgers University

Bruce, B. (1975) Belief Systems and Language Understanding Report CBM-TR-41, LCSR, Rutgers University

Bruce, B. (1973) CHRONOS Users Manual Report CBM-TM-26, LCSR, Rutgers University

Buchanan, B G., Sutherland, G. L., & Feigenbaum, E A (1969) Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry *Machine Intelligence 4*, Meltzer and Michie (Eds), Edinburgh University Press

Chokhani, S (1975) On Interpretation of Bio-Mathematical Models Within a Class of Decision Making Procedures Ph D Dissertation, Rutgers University. Also available as Report CBM-TR-43, LCSR, Rutgers University.

Ciesielski, V. (1980) A Methodology for the Construction of Natural Language Front Ends for Medical Consultation Ph D Dissertation, Rutgers University Also available as Report CBM-TR-112, LCSR, Rutgers University.

Ciesielski, V. (1979) Artificial Intelligence in Medicine: Contributions of the Rutgers Research Resource to *The Handbook of Artificial Intelligence.* Report CBM-TR-105, LCSR, Rutgers University.

Epstein, S (1983) Knowledge Representation in Mathematics: A Case in Graph Theory. Ph D Dissertation, Rutgers University

Kastner, J (1984) Strategies for Expert System Consultation in Therapy Selection. Ph D Dissertation, Rutgers University

Kastner, J K. Weiss, S M , Kulikowski, C A , Dawson, C. R (1983) Therapy Selection in an Expert Medical Consultation System for Ocular Herpes Simplex. *Computers in Biology and Medicine.*

Kelly, V. (1984) The CRITTER System–automated critiquing of digital circuit designs Ph D Dissertation, Rutgers University

Kelly, V , & Steinberg, L (1982) *The CRITTER system: analyzing digital circuits by propagating behaviors and specifications* AAAI-82 Also appears as Report LCSR-TR-30, LCSR, Rutgers University.

Kulikowski, C A ,& Sridharan, N. S. (1975) Report of the First Annual AIM Workshop Report CBM-TR-56, LCSR, Rutgers University.

Kulikowski, C. A ,& Weiss, S M. (1971) Computer-Based Models for Glaucoma Report CBM-TR-3, LCSR, Rutgers University.

LeFaivre, R (1977) RUTGERS/UCI LISP Manual Report CBM-TR-79, LCSR, Rutgers University.

Lichter, P. R Anderson, D R (1977) Discussions on Glaucoma New York: Grune & Stratton.

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

Lindsay, K. R , Buchanan, B. G., Feigenbaum, E A., & Lederberg, J. (1980) *Applications of Artificial Intelligence for Organic Chemistry the Dendral project.* New York: McGraw-Hill

McCarty, L T & Sridharan, N S (1981) *The Representation of an Evolving System of Legal Concepts II: Prototypes and Deformations.* IJCAI 7 Also available as Report LRP-TR-11, LCSR, Rutgers University.

Minsky, M. (1975) A Framework for Representing Knowledge. In Winston (Ed), *The Psychology of Computer Vision* New York: McGraw-Hill.

Mitchell, T (1983) *Learning and problem solving Computers and thought lecture* IJCAI 8

Mitchell, T (1982) Toward Combining Empirical and Analytical Methods for Inferring Heuristics Report LCSR-TR-27, LCSR, Rutgers University. to appear in Elithorn and Banerji (Eds), *Human and Artificial Intelligence* North Holland

Mitchell, T., Steinberg, L , Smith, R G., Schooley, P , & Kelly, V (1981) *Representations for Reasoning About Digital Circuits.* IJCAI 7 Also available as Report LCSR-TR-11, LCSR, Rutgers University

Mitchell, T., Utgoff, P Nudel, B , & Banerji, R. (1981) *Learning problem solving heuristics by practice* IJCAI 7 Also available as Report LCSR-TR-15, LCSR, Rutgers University

Mitchell, T (1978) Version Spaces: An Approach to Concept Learning. Ph D Dissertation, Stanford University Also available as Report STAN-CS-78-711, Computer Science Department, Stanford University

Newell, A (1984) Introduction to the COMTEX Microfiche Edition of Reports on Artificial Intelligence from Carnegie-Mellon University *AI Magazine* 5:3

Newell, A (1966) On the Representation of Problems. *Computer Science Research Review.* Department of Computer Science, CMU.

Nudel, B. (1983) Consistent-Labeling Problems and Their Algorithms: Expected-Complexities and Theory-Based Heuristics. *AI Journal* 21:1,2. (Special Issue on Search and Heuristics, edited by Judea Pearl)

Politakis, P. (1983) Using Empirical Analysis to Refine Expert System Knowledge Bases Ph D Dissertation, Rutgers University

Reiter, R. (1983) Towards a Logical Reconstruction of Relational Database Theory in Data Modelling in AI In Brodie, M L (Ed), Goodson, J L (1983) AIMDS Users Manual (Version 3) LCSR. Rutgers University.

Sridharan, N S , & Bresina, J L (1982) *Plan formation in large realistic domains.* 4th Natl. Conf. of CSCSI/SCEIO Canadian Society for Computational Studies of Intelligence. Also available as Report CBM-TR-127, LCSR, Rutgers University

Sridharan, N S , Bresina, J L & Schmidt, C F (1982) Evolution of a Plan Generation System. Report CBM-TR-128, LCSR, Rutgers University.

Sridharan, N S (1980) Representational Facilities of AIMDS: A Sampling. Report CBM-TM-86, LCSR, Rutgers University

Sridharan, N S , Hawrusik, F (1977) Representation of Actions and Transformations in AIMDS. Report CBM-TR-76, LCSR, Rutgers University

Srinivasan, C V (1981) Knowledge Representation and Problem Solving in MDS Report DCS-TR-106, LCSR, Rutgers University

Srinivasan, C. V. (1976) Introduction to Meta Description System Report SOSAP-TR-18, LCSR, Rutgers University

Szolovits, P , Pauker, S G (1978) Categorical and Probabilistic Reasoning in Medical Diagnosis. *AI Journal* 11:1,2. (Special Issue on Applications to the Sciences and Medicine—edited by N. S. Sridharan).

Trigoboff, M (1978) IRIS: A Framework for the Construction of Clinical Consultation Systems Ph D Dissertation, Rutgers University.

Weiss, S M , Kulikowski, C. A. (1984) *A Practical Guide to Designing Expert Systems* Rowman and Allanheld.

Weiss, S , Kulikowski, C , Apte, C , Uschold, M , Patchett, J., Brigham, & R Spitzer, B (1982) *Building expert systems for controlling complex programs.* AAAI-82

Weiss, S M , & Kulikowski, C A (1981a) Expert Consultation Systems: The EXPERT and CASNET projects *Machine Intelligence,* (Infotech State of the Art Report) 9:3

Weiss, S. M , Kulikowski, C A , & Galen, R S. (1981b) *Developing microprocessor based expert models for instrument interpretation.* IJCAI 7. Also available as Report CBM-TR-121, LCSR, Rutgers University.

Weiss, S M., Kulikowski, C. A. (1979) *EXPERT: A system for developing consultation models* IJCAI 6. Also available as Report CBM-TR-97, LCSR, Rutgers University

Weiss, S. M., Kulikowski, C A., Amarel, S , & Safir, A (1978) A Model-Based Method for Computer-Aided Medical Decision-Making *AI Journal* 11:1,2 (Special Issue on Applications to the Sciences and Medicine—edited by N S Sridharan)

Weiss, S M (1974) A System for Model-Based Computer Aided Diagnosis and Therapy Ph D Dissertation, Rutgers University Also available as Report CBM-TR-27, LCSR, Rutgers University

Welsch, L. (1979) ASQ: Automatic Synthesis of Questions Ph D Dissertation, Rutgers University Also available as Report LCSR-TR-6, LCSR, Rutgers University