TC Rindfleisch

SUMEX-AIM Resource (1973 – 1992)

(Adapted from an Annual Report dated 1980. Another account by Gregory Freiherr and commissioned by the Division of Research Resources of the National Institutes of Health in March 1980 is available: *The Seeds of Artificial Intelligence: SUMEX-AIM*, **NIH Publication No. 80-2071**.)

Principal Investigator:	Joshua Lederberg (1973-1976) Edward Feigenbaum (1976-1983) Edward Shortliffe (1983-1992)
Resource Director:	Thomas Rindfleisch (1973-1992)

Resource Goals and Background

The SUMEX-AIM project (Stanford University Medical EXperimental computer for Artificial Intelligence in Medicine) was a national computer resource funded by the National Institutes of Health (NIH) between 1973 and 1992. It encompassed a dual mission: 1) the promotion of applications of artificial intelligence (AI) computer science research to biological and medical problems and 2) the demonstration of network-based collaboration and computer resource sharing within a national community of health research projects.

Artificial Intelligence is a branch of computer science which attempts to discern the underlying principles involved in the acquisition and utilization of knowledge in reasoning, deduction, and problem-solving activities. The idea of intelligent machines is very old in fiction, but present work dates from the time stored program electronic computers became available starting in 1949. It is perhaps reasonable to date AI from A.M. Turing's 1950 paper: "Computing machinery and intelligence" (*Mind*, 59, 433-460). Any behavior that can be carried out by a mechanical device can be represented in a computer, and getting a particular behavior is "just" a matter of writing a program (unless the behavior requires special input and output equipment). AI systems are typically characterized by complex computational processes that are primarily non-numeric, e.g., graph-searching and symbolic pattern analysis. They involve procedures whose execution is controlled by diverse types and forms of knowledge about a given task domain, such as models, fragments of "advice", and systems of constraints or heuristic rules. Unlike conventional algorithms, commonly based on a well-tailored method for a given task, AI procedures may use a multiplicity of methods in a highly conditional manner – depending on the specific data in the task and a variety of sources of relevant information.

Each authorized project in the SUMEX community was concerned in some way with the application of these principles to real-world biomedical problems. The tangible objective of this approach was the development of computer programs which, using formal and informal knowledge bases together with mechanized hypothesis formation and problem-solving procedures, would be more general and effective consultative tools for the clinician and biomedical scientist. The exhaustive search potential of computerized hypothesis formation and knowledge base utilization, constrained where appropriate by heuristic rules or interactions with the user, already had begun to produce promising results in areas such as chemical structure elucidation, diagnostic consultation, and mental function modeling.

One of the overarching goals of the SUMEX-AIM project was to facilitate the process of fashioning a coherent scientific discipline out of the assemblage of personal intuitions, mathematical procedures, and emerging theoretical structure of the "analysis of analysis" and of problem solving. State-of-the-art programs were far more narrowly specialized and inflexible than the corresponding aspects of human intelligence they emulated; however, in special domains they had comparable or greater power, e.g. in the solution of formal problems in organic chemistry or in the integral calculus.

The community-building role of SUMEX-AIM was based upon the primitive state of computer communications technologies of the time. While far from perfected, those new capabilities offered highly desirable latitude for collaborative linkages, both within a given research project and among them. SUMEX began in an era predating personal computers and initially offered access to centralized time-shared computing resources (a Digital Equipment

Corporation PDP-10 TENEX system) via the early communications networks of the 1970s – ARPANET (SUMEX was host 56 on the 63-node ARPANET and was the only non-Department of Defense node at the time); TYMNET, TELENET, and WATS lines).

As computing systems evolved, SUMEX moved step by step toward a fully distributed model with community links provided through the evolving Internet. Several of the active projects on SUMEX were based upon the collaboration of computer and medical scientists at geographically separate institutions; separate both from each other and from the computer resource. Another major goal of the network experiment was to enable diverse projects to interact more directly and to facilitate selective demonstrations of available programs to physicians and medical students. Even in their developing state, such communication facilities allowed access to the rather specialized SUMEX computing environment and programs from a great many areas of the United States (even to a limited extent from Europe) for potential new research projects and for research product dissemination and demonstration.

The SUMEX-AIM resource resided administratively within the Stanford University School of Medicine and served as a nucleus for a growing community of projects at Stanford and nationally, eventually numbering about 20. SUMEX provided computing facilities specifically tuned to the needs of AI research and developed many tools for encouraging and facilitating community relationships among collaborating projects and medical researchers. SUMEX also initiated various efforts to build inter- and intra-group interactions through AIM workshops, a local “mini-conference” on AI techniques for the Stanford community of projects, and a seminar project initiated by Professor Feigenbaum to assemble from the community a “handbook” of AI concepts, methods, and state-of-the-art applications. The AI research and systems development projects at SUMEX-AIM led to the founding of a number of spin-off start-up companies, including Teknowledge, Intelligenetics, Xidak, Kinetics Fastpath, and Cisco Systems.

User projects were separately funded and autonomous in their management. They were selected for access to SUMEX, with peer review consultation, on the basis of their scientific and medical merits as well as their commitment to the community goals of SUMEX. Active projects spanned a broad range of application areas, including molecular biochemistry; clinical diagnostic and therapy consultation; belief systems modeling; mental function modeling; and instrument data interpretation.

SUMEX-AIM Research Communities

National AIM Projects

- **Acquisition of Cognitive Procedures (ACT)** (11/1975)

John Anderson, Ph.D.
Department of Psychology
Carnegie-Mellon University
Pittsburgh, Pennsylvania

Abstract: ACT (Acquisition of Cognitive Procedures) is a program intended to represent the development and performance of decision-making, much as other programs model aspects of human cognition, such as memory, inferential reasoning, language-processing, and problem-solving. ACT is a basic research project in AI, containing a logic scheme that may be transferable to applications in specific areas in or outside of medicine. The goal is that future versions of ACT will resemble very closely the process by which people learn to make decisions. ACT’s knowledge base consists of two components. One contains facts and serves as the program’s memory – essentially a data base. The other is a set of rules used to make decisions based on what is contained in the memory. Machine learning functions have been built into ACT to conceive new decision-making rules and to modify old ones on a continuing basis. A stumbling block in all learning systems is that rules commonly used in human decision-making often defy description, even by those who use them. A fallback method has been built in; for relatively unstructured situations, ACT uses trial-and-error.

- **Chemical Synthesis Project (SECS)** (9/1975)

W. Todd Wipke, Ph.D.
Department of Chemistry
University of California
Santa Cruz, California 95064

Abstract: The SECS (Simulation and Evaluation of Chemical Synthesis) project is aimed at developing symbolic and analytic methods for describing the logical principles used when synthesizing or constructing molecules. SECS is intended to facilitate the development of new and modified drugs, as well as synthetic compounds modeled after those that occur naturally. In particular, the project is concentrated on assisting the chemist to design and select syntheses of biologically important molecules.

- **Hierarchical Models of Human Cognition (CLIPR Project)** (12/1978)

Peter G. Polson, Ph.D.
Walter Kintsch, Ph.D.
Computer Laboratory for Instruction in Psychological Research (CLlPR)
Department of Psychology
University of Colorado

Abstract: The HIERARCHICAL MODELS OF HUMAN COGNITION project seeks to explore the complex cognitive processes that underlie text comprehension and planning. One aspect is to study the means by which people understand and summarize texts in order to determine ways to improve the readability of texts. An explicit theory of normal comprehension might lead scientists to the factors that cause learning problems in children, as well as suggest ways to overcome these problems. The other focus of the project is to model how people create plans and design complex systems. Specifically, the study will compare how experts and novices use their knowledge to design computer software. Given a coherent formulation of these processes, aids could be developed that would help people perform this task.

- **Higher Mental Functions Project** (10/1973)

Kenneth M. Colby, M.D.
University of California (UCLA)
Department of Psychiatry
Los Angeles, California 90024

Abstract: The Higher Mental Functions project is devoted to investigating models for human personality problems, specifically paranoia and adult neuroses. A computer simulation of paranoid thought processes (PARRY) is used to test the consistency of a theory describing the pathology. PARRY also serves as a training device in teaching students or psychiatric residents about various aspects of paranoia. The program has proved its ability to do both. Recently PARRY was interviewed by five psychiatrists via teletype a number of times. The psychiatrists were advised at the start that they would be communicating with either a patient or a computer. In the test, PARRY's responses were so indistinguishable from those of the paranoid patient that the psychiatrists could not tell the difference between the two. Although the test does not prove that the theory on which PARRY relies is all-inclusive, it shows that the theory is good enough to serve as a tool in teaching students about the pathology. In using AI techniques to classify neuroses, the project hopes to sharpen the rules that identify patients with different neuroses. This program works opposite to the way PARRY does. Rather than interpreting questions presented by interviewers and returning paranoid answers, the AI program in neuroses must take neurotic answers and work backward to the underlying concepts that distinguish the patient's pathology from those of other patients. These key ideas would then be clustered to form the profile of a certain type of patient.

- **INTERNIST Project** (10/1974)

Jack D. Myers, M.D.
Harry E. Pople, Ph.D.
Decision Systems Laboratory
University of Pittsburgh

Abstract: The goal of the INTERNIST project is to solve difficult diagnostic problems in internal medicine by simulating the reasoning processes used by human physicians. To be effective, the program must be able to diagnose several diseases in parallel if they are present in a single patient. Some of the expert diagnostic criteria the program emulates include: observations fed into the computer must evoke appropriate hypotheses about the disease; hypotheses must generate a list of manifestations that would be present in the patient if the diagnosis were correct; the computer must be able to rank models of disease according to their probability of being correct and must be able to decide when the weight of evidence is sufficient to permit reasonably confident judgment; the program must be able to group hypotheses into mutually exclusive subsets corresponding to different diagnoses. Since beginning work in 1970, the group has developed a working system that has partially achieved these objectives. Spin-offs from the program might be used by physicians' assistants or in rural health care clinics, military outposts, and spacecraft.

- **MEDICAL INFORMATION SYSTEMS LABORATORY Project** (1/1975)

Bruce McCormick, Ph.D.
Department of Information Engineering
University of Illinois, Chicago Circle
Morton Goldberg, M.D.
Department of Ophthalmology
University of Illinois Medical Center

Abstract: The Medical Information Systems Laboratory (MISL) is exploring artificial intelligence techniques applied to automated clinical decision-making in ophthalmology. Ancillary goals include investigations into the construction of a data base in ophthalmology, and into distributed data base design. Core research concerns the exploration of inferential relationships between analytic data and the natural history of selected eye diseases, both in treated and untreated form.

- **PUFF/VM: Biomedical Knowledge Engineering in Clinical Medicine** (10/1977)

John J. Osborn, M.D.
The Institutes of Medical Sciences
San Francisco, California
Edward A. Feigenbaum, Ph.D.
Stanford University

Abstract: PUFF/VM (Pulmonary Function and Ventilator Management) is a collaborative effort between respiratory disease experts at the Pacific Medical Center (PMC) in San Francisco and computer scientists at Stanford. Research includes development of two systems: one for the diagnosis and assessment of therapy for pulmonary function (PUFF), and the other for more intelligent monitoring of the condition and progress of patients who must use ventilators in intensive care units (ICUs) to assist breathing (VM). PUFF depends on some 250 decision-making rules, which are similar in form to those used by MYCIN. These rules are used to interpret a variety of patient signs related to pulmonary function. In VM, the goal is to develop a program with wider use in intensive care medicine. For example, VM explicitly considers the effects of time variation in interpreting a patient's condition. Measurements taken immediately after cardiac surgery might be within a normal range in that context, but a day later these same measurements may indicate problems. VM must change its expectations for the patient and interpret the measurements in new contexts.

- **Rutgers Computers in Biomedicine Resource** (10/1973)

Saul Amarel, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey
a. EXPERT AND CASNET/Glaucoma
Casimir Kulikowski, Ph.D.
Sholom Weiss, Ph.D.
b. BELIEVER
Charles Schmidt, Ph.D.

Abstract: CASNET is a reasoning framework, originally developed for diagnosing and treating glaucoma, that is being generalized as an AI tool and will be applied as a proof of concept in rheumatology. When the CASNET project began in 1971, the goal was to explore different ways of representing and using, in a computer program, expert medical knowledge to arrive at diagnoses and treatments. Researchers picked the domain of glaucomas because these eye diseases and their treatments have mechanisms that are relatively well-understood. Consequently, they were good candidates for in-depth modeling of the causes and effects of disease processes. Following their initial success, investigators are now applying techniques learned while constructing CASNET/Glaucoma to the design of a new computer program for diagnosis and treatment of rheumatic disease. The system currently has preliminary models of seven rheumatic diseases and a data base of over 150 patient cases. One of the rheumatology program's most important features is that its medical knowledge can be rapidly organized and updated.

The BELIEVER project seeks to refine a theory of human information-processing. The goal is to define the way people assemble facts into coherent, understandable patterns, for example, in order to understand the observed actions of others.

- **Simulation of Cognitive Processes (SCP)** (2/1978)

James G. Greeno, Ph.D.
Alan M. Lesgold, Ph.D.
Learning Research & Development Center
University of Pittsburgh

Abstract: In the Simulation of Cognitive Processes project, a version of ACT is being applied to model the acquisition of reading and arithmetic problem-solving skills. These two skills were chosen because they are very basic cognitive requirements for getting along in the world. The project is founded on studies showing that various word-processing skills and arithmetic procedures are underdeveloped in children who do poorly in reading, mathematics, or both. One model simulates the process of solving simple word problems. It is intended to determine the degree to which semantic and linguistic factors, rather than arithmetic knowledge, are responsible for children's difficulty in solving these problems at early grade levels. By providing a framework in which the effects of different levels of skill acquisition can be understood, it is hoped that criteria can be developed that will separate subjects who have brain damage from those whose cognitive skills have developed poorly.

Stanford Projects

- **AI Handbook Project** (4/1977)

Edward A. Feigenbaum, Ph.D.
Department of Computer Science
Stanford University

Abstract: The AI HANDBOOK project is an effort aimed at speeding the dissemination of information about AI techniques. In its final form, the handbook will contain some 200 articles covering the most important ideas, techniques, and systems developed during the past 20 years of AI research. The articles, each about four pages long, will be written in language suited to the student of AI, as well as to professionals outside the field. Later this year, the first of two volumes is expected to be printed. It will cover techniques for heuristic search, knowledge representation, AI programming languages, natural language understanding, speech understanding, application-oriented research in AI, and automatic programming. The authors represent work at both academic and commercial research centers.

- **Attempt to Generalize (AGE)** (9/1977)

H. Penny Nii
Edward A. Feigenbaum, Ph.D.
Department of Computer Science
Stanford University

Abstract: The ATTEMPT TO GENERALIZE (AGE) project is intended to "despecialize" software, making knowledge engineering more generally available to the scientific community. Most projects in the SUMEX-AIM community, such as DENDRAL, MYCIN, and MOLGEN, develop "handcrafted" intelligent agents to assist human problem-solving in task domains of medicine and biology. This process is very time-consuming. AGE grew out of a speech-understanding program (HEARSAY) that is based on a "blackboard" model for developing, organizing, and solving problems using different types of knowledge sources. Much as EMYCIN abstracts the framework for rule-based systems. AGE is aimed at abstracting the knowledge acquisition and problem-solving framework of blackboard systems.

- **DENDRAL Project** (10/1973)

Carl Djerassi, Ph.D.
Department of Chemistry
Edward A. Feigenbaum, Ph.D.
Department of Computer Science
Stanford University

Abstract: The goal of the DENDRAL (Dendritic Algorithm) project uses computerized symbolic reasoning to determine biomolecular structures. The inputs to DENDRAL are instrument data such as from a high-resolution mass spectrometer (MS) coupled with a gas chromatograph; from a nuclear magnetic resonance (NMR) spectrometer; or from an infrared spectrometer; together with other constraints on structural features in the molecule. These constraints describe configurations of atoms and provide limits within which the structural candidates for an unknown compound must fit. Meta-DENDRAL is a program that learns mass spectrometer cleavage rules for specific types of compounds by examining data from a set of examples. Both DENDRAL and meta-DENDRAL use similar production system or rule-based logic.

- **MOLGEN Project** (9/1976)

Laurence H. Kedes, M.D.
Veterans Administration Hospital
Palo Alto, California 94304
Edward A. Feigenbaum, Ph.D.
Douglas Lenat, Ph.D.
Department of Computer Science
Stanford University

Abstract: The MOLGEN (Molecular Genetics) project is developing tools to help with experiment planning for the manipulation of DNA. Experimental steps include methods to analyze and modify nucleic acids using physical methods such as electron microscopy or enzymatic modification. MOLGEN is mainly focused on organizing experimental techniques and determining the order in which they should be applied to achieve specified goals.

- **MYCIN/ONCOCIN/GUIDON Project** (10/1973)

Bruce G. Buchanan, Ph.D.
Edward H. Shortliffe, M.D., Ph.D.
Department of Computer Science
Stanford University

Abstract: MYCIN specializes in the diagnosis and therapy selection for patients with specific infectious diseases. Its goal is to provide sound therapeutic advice, using available information to identify all the organisms likely to be causing the infection. The system builds on and extends techniques evolved through years of DENDRAL experience encoding domain-specific knowledge using “production rules” (essentially IF … THEN statements) which represent facts and their interrelationships. The MYCIN knowledge base currently consists of some 500 such rules. In its domain of expertise, MYCIN scores better than human prescribers. The EMYCIN project extracts the essential rule-based problem-solving strategy from MYCIN so it can be applied to other fields. This framework includes the logic behind diagnosis, therapy recommendations, explanation, and knowledge acquisition.

ONCOCIN is an oncology protocol management system that is designed to assist physicians in the treatment of cancer patients. The system consists of a set of programs, including a rule-based reasoner that has the necessary knowledge of cancer chemotherapy, and a graphical user interface that makes the system acceptable to oncologists. Work on ONCOCIN began in mid-1979 and the system was installed for preliminary use in May 1981.

The GUIDON project is an attempt to teach the diagnostic and therapy planning logic of domain experts (like MYCIN) to people. The program uses AI techniques to represent both subject material and teaching strategies. Unlike traditional computer-aided instruction (CAI) programs, GUIDON constructs models of the student’s knowledge, reflecting their weaknesses, strengths, and preferred style of learning.

- **Protein Structure Project** (10/1973)

Edward A. Feigenbaum, Ph.D.
Robert Engelmore, Ph.D.
Department of Computer Science
Stanford University

Abstract: The goal of the PROTEIN STRUCTURE project is to build computer models of protein structures from crystallographic data, particularly electron density maps. Electron density maps represent the structures in three dimensions, but these maps are usually crude and ambiguous. The program depends heavily on background information, such as the amino acid sequence in a protein, for guidance in forming hypotheses about the compound’s three-dimensional structure. The program attempts to capture the knowledge protein chemists use for this task in heuristic rules.

- **RX Project** (1/1979)

Robert L. Blum, M.D.
Gio Wiederhold, Ph.D.
Department of Computer Science

Stanford University

Abstract: The goal of the RX project is to develop a system for extracting knowledge about the evolution and treatment of chronic diseases from data in patient records stored in computerized clinical data banks. Chronic diseases comprise the majority of the diseases taking the greatest toll in terms of death and disability: arteriosclerosis; cancers; high blood pressure; arthritis; diabetes; and others. Expensive and difficult prospective trials have been the most common way to evaluate alternative treatments. RX seeks to combine statistical approaches with AI techniques utilizing large knowledge bases to better discern the complex relationships among the many variables that influence the progression of chronic disease. The knowledge base of RX will contain knowledge of the various diseases, symptoms, therapies, outcomes, laboratory tests, and the many interrelationships which exist among them. The knowledge base will be used primarily to abstract the key events occurring in the computerized patient charts. These abstracted records will then be used to assess the degree of correlation between various therapies and disease outcomes.

Pilot National AIM Projects

- **Communication Enhancement Project** (3/1977)

John B. Eulenberg, Ph.D.
Carl V. Page
Department of Computer Science
Michigan State University
Collaborator: Kenneth Colby (UCLA)

Abstract: The COMMUNICATION ENHANCEMENT pilot project seeks to design intelligent speech prostheses for persons with severe communication handicaps. Research includes the design of input devices that can be used by persons whose movement is greatly restricted, the development of software for text-to-speech production, and the production of a microcomputer-based portable speech prosthesis. Early work resulted in a portable communication system for a 10-year-old boy with cerebral palsy who cannot speak or use his hands to write. The partially successful device influenced the design of a lap-board communication aid, called SAL (Semantically Accessible Language). It translates Bliss symbols (developed by C. K. Bliss) into spoken language. When using the lap-board, patients choose symbols for various words, which are translated by a microcomputer into orthographic and phonetic strings which are turned into sounds by a voice synthesizer and into typed words by a visual display unit. The goal is to greatly extend the intelligence of the program to produce more personalized and grammatically correct output.

- **Computerized Psychopharmacology Advisor (HEADMED)** (5/1976)

Jon F. Heiser, M.D.
Ruven E. Brooks, Ph.D.
Department of Psychiatry and Behavioral Sciences
University of Texas Medical Branch

Abstract: HEADMED is a psychopharmacology advisor that addresses the problem of drug misuse in psychiatry, where often there are inadequate rationales for prescriptions and dosages may be wrong. In addition, many physicians may not know how such medications affect a patient's general health, and other physical disorders or treatments that the patient may be getting. The HEADMED project seeks to adapt MYCIN's rule-based architecture to fit a medical domain much different than infectious diseases. A side benefit from developing the program is that the rules for patient assessment and management will be objectively described and will be of value as educational tools.

Pilot Stanford Projects

- **Ultrasonic Imaging Project** (11/1978)

W. Desmond McCallum, M.D.
James Brinkley, M.D.
Department of Gynecology and Obstetrics
Stanford University Medical School

Abstract: The ULTRASONIC IMAGING project builds on past experience in the use of computers and ultrasound techniques to help model body organs in three dimensions. Models are to be used to study anatomic structures noninvasively and to determine the volume of organs. Data obtained will be applied to clinical diagnosis. Initially, the system will be designed to determine the volume of a fetus as an indicator of its weight. Later it will be extended to measure the volume of the liver, the kidney, or the left ventricle of the heart.