SUMEX

STANFORD UNIVERSITY

MEDICAL EXPERIMENTAL COMPUTER RESOURCE

RR - 00785

ANNUAL REPORT - YEAR 08

Submitted to

BIOTECHNOLOGY RESOURCES PROGRAM

NATIONAL INSTITUTES OF HEALTH

June 1, 1981

STANFORD UNIVERSITY SCHOOL OF MEDICINE

Edward A. Feigenbaum, Principal Investigator

DEPARTMENT OF HEALTH AND HUMAN SERVICES
PUBLIC HEALTH SERVICE
NATIONAL INSTITUTES OF HEALTH

DIVISION OF RESEARCH RESOURCES
BIOTECHNOLOGY RESOURCES PROGRAM

ANNUAL PROGRESS REPORT

1. PHS GRANT NUMBER: P41RR-00785-08

2. TITLE OF GRANT: S U MEDICAL EXPERIMENTAL COMPUTER RESOURCE (SUMEX)

3. NAME OF RECIPIENT INSTITUTION: Stanford University

4. HEALTH PROFESSIONAL SCHOOL: School of Medicine

5. REPORTING PERIOD:

   5a. FROM: August 1, 1980

   5b. TO: July 31, 1981

6. PRINCIPAL INVESTIGATOR:

   6a. NAME: Edward A. Feigenbaum, Ph.D.

   6b. TITLE: Professor and Chairman
   Department of Computer Science

   6c. SIGNATURE: Edward A Feigenbaum

7. DATE SIGNED: 6/17/81

8. TELEPHONE: (415) 497-4079

Table of Contents

List of Figures

Figure Page

I Narrative Description

This is an annual report for the Stanford University Medical EXperimental computer resource for applications of Artificial Intelligence in Medicine (SUMEX-AIM). It covers the period between May 1, 1980 and April 30, 1981.

We are about to begin a 5-year renewal of the SUMEX resource grant which will launch an important and exciting new phase for SUMEX-AIM community research. Recent successes in developing expert systems, many of them stemming from projects in the SUMEX-AIM community, have stimulated increasing interest in AI research from many fronts. At the same time, the on-going revolution in computational tools, made possible by larger and larger scale microelectronic integration, is making routine applications of AI systems more practical and effective. Our approved renewal goals focus principally on a merging of state-of-the-art community research in biomedical AI applications with these new computing tools and on the challenges they will bring to the SUMEX-AIM community and resource. We expect that the integration and exploitation of these emerging computer technologies that will have a profound effect on the development and export of practical biomedical AI programs.

This report on the last year in our current 3-year grant is thus, in a sense, a culmination of the early phase of the SUMEX resource. This phase has been characterized by the building of a national community of biomedical AI collaborators around a central resource located at Stanford University. Beginning with 5 projects in 1973, the AIM community grew to 11 major projects at our renewal in 1978 and currently numbers 16 fully authorized projects plus a group of 7 pilot efforts. Many of the computer programs under development by these groups are maturing into tools increasingly useful to the respective research communities. The demand for production-level use of these programs has surpassed the capacity of the present SUMEX facility and has raised important issues of how such software systems can be optimized for production environments, exported, and maintained.

To be sure, we will continue to seek interesting new AI applications in an expanding community of biomedical and computer scientists interacting through electronic media. However, we expect the SUMEX-AIM community to develop a somewhat different character in the coming years. It will become more decentralized in terms of computing resources, more diverse in scope, and even more heavily dependent on network communication facilities for interactions, collaborations, and sharing.

The following sections report on the activities of the SUMEX-AIM resource this past year including brief summaries of the objectives of SUMEX-AIM, a characterization of biomedical AI research, resource organization and operating procedures, recent core progress in system development and basic AI research, and progress in the collaborative projects.

I.A Summary of Research Progress

I.A.1 Overview of Objectives and Rationale

SUMEX-AIM ("SUMEX") is a national computer resource with a dual mission: a) promoting applications of computer science research in artificial intelligence (AI) to biological and medical problems and b) demonstrating computer resource sharing within a national community of health research projects. The central SUMEX-AIM facility is located physically in the Stanford University Medical School and serves as a nucleus for a community of medical AI projects at universities around the country. SUMEX provides computing facilities tuned to the needs of AI research and communication tools to facilitate remote access, inter- and intra-group contacts, and the demonstration of developing computer programs to biomedical research collaborators.

I.A.1.1 What is Artificial Intelligence

Artificial Intelligence research is that part of Computer Science concerned with symbol manipulation processes that produce intelligent action [1 - 7]. By "intelligent action" is meant an act or decision that is goal-oriented, is arrived at by an understandable chain of symbolic analysis and reasoning steps, and utilizes knowledge of the world to inform and guide the reasoning.

Placing AI in Computer Science

A simplified view relates AI research with the rest of computer science. The ways in which people use computers to accomplish tasks can be "one-dimensionalized" into a spectrum representing the nature of the instructions that must be given the computer to do its job; call it the What-to-How spectrum. At the How extreme of the spectrum, the user supplies his intelligence to instruct the machine precisely how to do his job, step-by-step. Progress in computer science may be seen as steps away from that extreme How point on the spectrum: the familiar panoply of assembly languages, subroutine libraries, compilers, extensible languages, etc. illustrate this trend.

At the other extreme of the spectrum, the user describes What he wishes the computer to do for him to solve a problem. He wants to communicate what is to be done without having to lay out in detail all necessary subgoals for adequate performance. Still, he demands a reasonable assurance that he is addressing an intelligent agent that is using knowledge of his world to understand his intent, complain or fill in his vagueness, make specific his abstractions, correct his errors, discover appropriate subgoals, and ultimately translate What he wants done into detailed processing steps that define How it shall be done by a real computer. The user wants to provide this specification of What to do in a language that is comfortable to him and the problem domain (perhaps

English) and via communication modes that are convenient for him (including perhaps speech or pictures).

The research activity aimed at creating computer programs that act as "intelligent agents" near the What end of the What-to-How spectrum can be viewed as a long-range goal of AI research.

Expert Systems and Applications

The national SUMEX-AIM resource is an outgrowth of a long, interdisciplinary line of artificial intelligence research at Stanford concerned with the development of concepts and techniques for building "expert systems" [1]. An "expert system" is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. For some fields of work, the knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the expert practitioners of that field.

The knowledge of an expert system consists of facts and heuristics. The "facts" constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The "heuristics" are the mostly-private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterize expert-level decision making in the field. The performance level of an expert system is primarily a function of the size and quality of the knowledge base that it possesses.

Currently authorized projects in the SUMEX community are concerned in some way with the application of AI to biomedical research (*). The tangible objective of this approach is the development of computer programs that will be more general and effective consultative tools for the clinician and medical scientist. There have already been promising results in areas such as chemical structure elucidation and synthesis, diagnostic consultation, and modeling of psychological processes.

Needless to say, much is yet to be learned in the process of fashioning a coherent scientific discipline out of the assemblage of personal intuitions, mathematical procedures, and emerging theoretical structure comprising artificial intelligence research. State-of-the-art programs are far more narrowly specialized and inflexible than the corresponding aspects of human intelligence they emulate; however, in special domains they may be of comparable or greater power, e.g., in the solution of formal problems in organic chemistry.

(*) Brief abstracts of the various projects can be found in Appendix A on page 278 and more detailed progress summaries in Section II on page 89.

I.A.1.2 Resource Sharing

Besides the biomedical AI research theme of SUMEX-AIM, another central goal is an exploration of the use of computer-based communications as a means for interactions and sharing between geographically remote research groups engaged in biomedical computer science research. This facet of scientific interaction is becoming increasingly important with the explosion of complex information sources and the regional specialization of groups and facilities that might be shared by remote researchers [8]. We expect an even greater decentralization of computing resources in the coming years with the emerging VLSI (*) technology in microelectronics and a correspondingly greater role for digital communications.

Our community building effort is based upon the current state of computer communications technology. While far from perfected, these developing capabilities offer highly desirable latitude for collaborative linkages, both within a given research project and among them. A number of the active projects on SUMEX are based upon the collaboration of computer and medical scientists at geographically separate institutions; separate both from each other and from the computer resource. The network experiment also enables diverse projects to interact more directly and to facilitate selective demonstrations of available programs to physicians, scientists, and students.

We have actively encouraged the development of additional affiliated computing resources within the AIM community and expect such decentralization to become the "way of the 80's". Since 1977, the facility at Rutgers University has allocated a portion of its capacity for national AIM projects and our network connections to Rutgers and common facilities for user terminals have been indispensable for effective interchanges between community members, workshop coordinations, and software sharing. In addition, the "Caduceus" project (**) (page 187) is expecting delivery of their own machine momentarily, the "Simulation of Cognitive Processes" project (page 226) already is doing most of their work on their own VAX computer, and several more projects have proposed machines dedicated to their own use.

The proliferation of distributed machines will serve to increase the importance of electronic communications to facilitate interactions and sharing. Even in their current developing state, communication facilities enable effective access to the SUMEX community resources from a great many areas of the United States and to a more limited extent from Canada, Europe, Japan, Australia, and other international locations.

(*) Very Large Scale Integration

(**) Previously called "Internist".

I.A.1.3 Impact of AI in Biomedicine

Artificial Intelligence is the computer science of symbolic representations of knowledge and symbolic inference. There is a certain inevitability to this branch of computer science and its applications, in particular, to medicine and biosciences. The cost of computers will continue to fall drastically during the coming two decades. As it does, many more of the practitioners of the world's professions will be persuaded to turn to economical automatic information processing for assistance in managing the increasing complexity of their daily tasks. They will find, from most of computer science, help only for those of their problems that have a mathematical or statistical core, or are of a routine data-processing nature. But such problems will be relatively rare, except in engineering and physical science. In medicine, biology, management -- indeed in most of the world's work -- the daily tasks are those requiring symbolic reasoning with detailed professional knowledge. The computers that will act as "intelligent assistants" for these professionals must be endowed with symbolic reasoning capabilities and knowledge.

The growth in medical knowledge has far surpassed the ability of a single practitioner to master it all, and the computer's superior information processing capacity thereby offers a natural appeal. Furthermore, the reasoning processes of medical experts are poorly understood; attempts to model expert decision making necessarily require a degree of introspection and a structured experimentation that may in turn improve the quality of the physician's own clinical decisions, making them more reproducible and defensible. New insights that result may also allow us more adequately to teach medical students and house staff the techniques for reaching good decisions, rather than merely to offer a collection of facts which they must independently learn to utilize coherently.

The knowledge that must be used is a combination of factual knowledge and heuristic knowledge. The latter is especially hard to obtain and represent since the experts providing it are mostly unaware of the heuristic knowledge they are using. Medical and scientific communities currently face many widely recognized problems relating to the rapid cumulation of knowledge, for example:

- codification of theoretical and heuristic knowledge
- effective use of the wealth of information implicitly available in textbooks, journal articles and from practitioners
- dissemination of that knowledge beyond the intellectual centers where it is collected
- customizing the presentation of that knowledge to individual practitioners as well as customizing the application of the information to individual cases

We believe that computers are the most hopeful technology to help overcome these problems. While recognizing the value of mathematical modeling, statistical classification, decision theory and other techniques, we believe that effective use of such methods depends on using them in conjunction with less formal knowledge, including contextual and strategic knowledge.

Artificial intelligence offers advantages for representing information and using it that will allow physicians and scientists to use computers as intelligent assistants. In this way we envision a significant extension to the decision making powers of individual practitioners without reducing the significance of the individuals.

Knowledge is power, in the profession and in the intelligent agent. As we proceed to model expertise in medicine and its related sciences, we find that the power of our programs derives mainly from the knowledge that we are able to obtain from our collaborating practitioners, not from the sophistication of the inference processes we observe them using. Crucially, the knowledge that gives power is not merely the knowledge of the textbook, the lecture and the journal but the knowledge of "good practice" -- the experiential knowledge of "good judgment" and "good guessing", the knowledge of the practitioner's art that is often used in lieu of facts and rigor. This heuristic knowledge is mostly private, even in the very public practice of science. It is almost never taught explicitly; almost never discussed and critiqued among peers; and most often is not even in the moment-by-moment awareness of the practitioner.

Perhaps the the most expansive view of the significance of the work of the SUMEX-AIM community is that a methodology is emerging therefrom for the systematic explication, testing, dissemination, and teaching of the heuristic knowledge of medical practice and scientific performance. Perhaps it is less important that computer programs can be organized to use this knowledge than that the knowledge itself can be organized for the use of the human practitioners of today and tomorrow.

The researchers of the SUMEX-AIM community currently constitute a large fraction of all the computer scientists whose work is aimed at the development of symbolic computational methods and tools. SUMEX-AIM is laying the scientific base so that medicine will be able to take advantage of these technological opportunities for inexpensive computer power. Medical diagnostic aids and tools for the medical scientist that operate in a environment of a network of "professional workstation" computers have the practical possibility of large-scale and low-cost use because of anticipated near-term developments in the computing industry.

I.A.2 Synopsis of Recent Progress

As we complete year 08, we can report substantial further progress in the overall mission of the SUMEX-AIM resource. We have continued the refinement of an effective set of hardware and software tools to support the development of large, complex AI programs for medical research and to facilitate communications and interactions between user groups. We have worked to maintain high scientific standards and AI relevance for projects using the SUMEX-AIM resource and have actively sought new applications areas and projects for the community. Many projects are built around the communications network facilities we have assembled; bringing together medical and computer science collaborators from remote institutions and making their research programs available to still other remote users. As discussed in the sections describing the individual projects, a number of the computer programs under development by these groups have matured into tools increasingly useful to the respective research communities. The demand for production-level use of these programs has surpassed the capacity of the present SUMEX facility and in preparation for our renewal goals, we have been investigating the general issues of how such software systems can be moved from SUMEX and supported in production environments.

A number of significant events and accomplishments affecting the SUMEX-AIM resource occurred during the past year:

1) In August 1980, under the chairmanship of Prof. Ted Shortliffe and with the assistance of Drs. L. Fagan and R. Blum, Stanford hosted the sixth AIM workshop. This workshop was innovative in that the presentations were fully "demo-based" using a live video projection of program typescripts and actual running sessions. The purpose of this approach was to allow participants to see more deeply into the inner workings of the various systems under development.

2) In conjunction with the 1980 workshop, Drs. Clancey and Shortliffe organized a continuing education tutorial for practicing physicians. The tutorial session was attended by over 135 doctors and included an introduction to computing, background information on decision theory and database applications in medicine, and presentations on a number of AI systems by 15 members and affiliates of the SUMEX-AIM community.

3) In November 1980, we defended our pending renewal application before a peer review site visit team. The SUMEX-AIM community was represented by several members of the AIM Executive Committee. A strong endorsement for future SUMEX goals and a recommendation for a 5-year renewal period resulted. These were confirmed by study section and council action. The technical substance of our future goals are outlined beginning on page 47.

4) The SUMEX-AIM collaborator project community has continued vigorous development of their respective programs. Details are reported by the individual investigators in Section II. The VM and ONCOCIN projects have begun preliminary clinical testing/evaluation this past year using SUMEX network and computing resources. The CADUCEUS

(INTERNIST) and SIMULATION OF COGNITIVE PROCESSES projects have been funded for and are setting up their own local VAX computing resources which should help reduce the load on SUMEX for newer pilot efforts. We have continued to work hard to meet the needs of collaborating projects and are grateful for their expressed appreciation.

5) We supported a highly successful, experimental dissemination of the MOLGEN programs into the molecular biology community. "Advertised" through presentations and demonstrations by MOLGEN investigators at several professional conferences, over 200 molecular biologists have used the system and most have found it easy to learn and highly effective as a research tool for their investigations.

6) We have continued development of the SUMEX facility hardware, software, and network systems to enhance throughput and to assist user access to existing and planned resources. A good range of internetwork software is available now including telnet, file transfer, and mail handling. Following the council recommendation for approval of our renewal application, our request to augment the AMPEX memory was funded by BRP. We have installed the new memory and are in the process of tuning the monitor to optimize use of the increased user memory.

7) We have actively explored options for professional workstation and VAX LISP systems in preparation for our renewal research. The current state of available systems is encouraging. However, delays in an operational version of Interlisp-VAX and an earlier than expected availability of Interlisp-Dolphin workstations has led us to recommend beginning the workstation phase of our research first.

I.A.3 Details of Technical Progress

The following material covers SUMEX-AIM resource activities over the past year in greater detail. These sections outline accomplishments in the context of the resource staff and the resource management. Details of the progress and plans for our external collaborator projects are presented in Section II beginning on page 89.

I.A.3.1 Facility Hardware

Over the past year, the SUMEX facility hardware configuration, including the main KI-10 machine (Figure 1), the 2020 satellite machine (Figure 2), and system network interconnections (Figure 3), have continued to develop according to plan and to operate effectively within capacity limitations. The primary facility hardware development efforts this year have been directed at:

1) Augmentation of the 256K word AMPEX memory to 512K words.

2) Implementation of Ethernet interface equipment for the KI-10 and other network server facilities.

3) Investigation and planning of hardware alternatives for the system development goals of our renewal grant.

4) Support of local project hardware needs.

Memory Augmentation

The SUMEX-AIM facility has been operating at capacity in terms of prime-time computing load for the past several years as documented in our previous reports. In spite of implementing a number of strategic facility augmentations over the years, we have not been able to satisfy the computing demands of our community. This condition has constrained the growth of the AIM community and our ability to bring AI programs nearing operational status in contact with potential external user communities while continuing to support on-going program development efforts. We have taken active steps to transfer prime time interactive loading to evening and night hours as much as possible including shifting personnel schedules (particularly for Stanford-based projects). We have implemented tools to control the fair allocation of CPU resources between various user communities and projects and have encouraged jobs not requiring intimate user interaction to run during off hours using batch job facilities. And we have acquired a 2020 system to offload program demonstrations and evaluations from the main research machine. Despite these efforts, our prime time loading has remained at saturation. Perhaps the most significant effect of the resulting poor response time is the deterrence of interactions with medical and other professional collaborators experimenting with available AI programs, whose schedules cannot be adjusted to meet computer loading patterns.

From the SUMEX viewpoint, we have attempted to do everything feasible and economically justified within available budgets to maximize the use of the existing hardware for productive work. One remaining step has been the expansion of our AMPEX memory from its current 256K word complement to its full 512K word capacity. The effect of this upgrade is to make more physical memory available to user programs thereby reducing swapping overhead (page faults and interrupt handling) and smoothing out system responsiveness under heavy load by keeping more working sets in core.

We requested approval for this expansion in May 1980. Following council approval of our renewal grant application, we received funding for the upgrade. The added memory was received and installed May 14, 1981, checked out during the following week, and a new 786K monitor brought up on May 21. This addition has increased user memory by about 60%. It is still too early to draw detailed conclusions about the effect of this enhancement and further tuning of monitor parameters controlling process scheduling and working set management needs to be done. We will report detailed results next year.

Local Network Interfaces

The initial design of the SUMEX system was that of a "star" topology centered on the KI-10 processors. In this configuration, all peripheral equipment and terminal ports were connected directly to the KI-10 busses. With the addition of new satellite machines, a unique focus no longer exists and some pieces of equipment need to be able to "connect" to more than one host. For example, a user coming into SUMEX over TYMNET will want to be able to make a selection of which machine he connects to. Another TYMNET user may want to make another choice of machine and so the TYMNET interface needs to be able to connect to any of the hosts. This could be accomplished by creating separate interfaces for each of the hosts to the TYMNET, each with a different address. Besides being expensive to duplicate such interface connections, it would be inconvenient for a user to reconnect his terminal from one host to another.

Over the past year and a half we have been developing a local, high-speed Ethernet to provide a flexible basis for our planned facility developments. The KI-10's and the 2020 were connected in time to support the AIM workshop last summer. Our development of Ethernet facilities has been guided by the goals of providing the most effective range of services for SUMEX community needs while remaining compatible with and able to contribute to and draw upon network developments by other groups. Since the early 3 Mbit/sec Ethernet was given to Stanford and several other universities by Xerox, an agreement has been reached between DEC, INTEL, and Xerox on the standards for an even higher performance network [13]. The new network runs at 10 Mbits/sec and supports a significantly larger packet address space. Xerox has started to market products for the new network but debugged interfaces, software, etc. for general use are not routinely available yet. Furthermore, even though three companies have agreed on a set of low level protocols and interface conventions, the rest of the world may not go along. There is already an alternative (but closely related) IEEE specification in preparation. Even among the three parties in the Ethernet specification, there is no agreement on higher level protocols.

All of this suggests that it is not time to jump to the newer and faster networks yet. We feel the 3 Mbit/sec network is adequate for our bandwidth needs in the near future and there is already a significant investment in 3 Mbit/sec network equipment at Stanford related to SUMEX community interests. In the longer term, we will want to upgrade to whatever hardware and protocol standard is broadly adopted.

In the meantime we are continuing to develop our 3 Mbit/sec PUP network services. This places a heavier burden on us to develop and maintain our own equipment for Ethernet support. We have tried to minimize the "home-brew" nature of this work by sharing common hardware and software designs with other groups in the same situation

The initial KI-10 interface was made via a PDP-11 connected to the I/O bus which is inefficient under heavy traffic. In anticipation of increased Ethernet demands on the KI-10's for high-speed terminals, file transfers, and other server functions, we have been designing and implementing a more efficient direct memory access interface. This interface uses a phase decoder (design borrowed from the SUN terminal project at Stanford) to detect the incoming serial Ethernet signal, an internal packet buffer to prevent overruns to and from the TENEX time-sharing system, and a memory bus interface to transfer data. The KI-10 DMA interface is partially debugged while highest priority work is proceeding on a gateway to the computer science building across campus.

In our initial connections of the KI-10 and 2020, we used a UNIBUS interface board designed by E. Markowski at Xerox. Because of the limited availability of these boards for our future work (an immediate need being for a gateway between various campus Ethernets), we began work an a PDP-11 interface board. This design is similar to that of the KI-10 interface and shares the serial phase decoder network front end. It provides several features not available on the Xerox board including more explicit error information and a more sophisticated filter on source addresses for incoming packets.

Planning for the Renewal Period

Over the past year we have spent considerable effort evaluating strategies and alternatives for planned system development in our renewal grant. Pending funding, council has approved our plan to acquire two VAX machines, five professional workstations, and a file server for the SUMEX resource starting in August 1981. We have debated at length the appropriate timing for purchases of this equipment within budget constraints. The Initial Review Group and Council enthusiastically endorsed the importance of optimizing the timing of our planned hardware acquisitions to coincide with the availability of desired technological developments and community needs. They recommended in their report that we be allowed considerable flexibility as to phasing of equipment purchases within the 5-year renewal period.

The rapidly changing technical and commercial situation vis a vis the research computing equipment we plan to buy if funded, indicates that there would be significant advantage to the SUMEX-AIM community in exercising this flexibility by delaying the purchase of our first VAX until the second renewal year (grant year 10) and advancing the purchase of the Professional Workstations to the first year (grant year 09). The rationale for this switch is as follows:

1) The INTERLISP language has been the basis for most SUMEX-AIM community AI research. Development of the VAX INTERLISP system at USC-ISI is substantially behind schedule. The most current estimate for completing a usable system is mid-1982 and no viable alternative version of LISP, with a fully developed programming support environment, will be available any sooner than that. Thus, if we purchased a VAX in year 09, we could not offer effective VAX LISP services before year 10. Strong pressure does exist within the ARPANET community to get VAX INTERLISP completed as expeditiously as possible so we believe that VAX will be a good machine choice by year 10 once INTERLISP is running. We are undertaking a separate study of this situation to assess the likelihood of VAX/INTERLISP being completed in a timely fashion and to estimate its performance characteristics on the VAX 11/780.

2) If the interchange in timing of the purchases of the first VAX and the Professional Workstations is approved, we have agreed that SUMEX-AIM will have shared access next year to the VAX 11/780 funded by ARPA to support Stanford Heuristic Programming Project research. This will minimize any delays in SUMEX-AIM work involving VAX that is not dependent on INTERLISP and will enable necessary systems development work and preliminary experimentation by SUMEX-AIM users to proceed without having to commit NIH grant funds. Because of long term commitments for the ARPA VAX and expected growth in SUMEX community needs, however, it can only substitute on a temporary basis during the first renewal grant year. After that a VAX dedicated solely to SUMEX-AIM community use will be needed.

3) The DEC VAX product line is continuously changing and there are some indications that new products may be offered on the 1982-1983 time scale that would be advantageous to SUMEX-AIM research. These may include features that enhance technical performance and/or cost effectiveness for our purposes. By year 10, we should be able to make a more judicious choice of the best configuration for our needs.

4) While VAX/INTERLISP is delayed, a suitable model of the professional workstation we need for our experimentation is available earlier than expected. The Xerox Dolphin is a system that has been in use as a research machine within Xerox for some time. It meets our technological needs including a high-bandwidth bit-mapped display terminal, full TENEX INTERLISP software compatibility, increased address space over the PDP-10 (but not as large as will be available on the VAX), acceptable capacity (roughly twice the single-user KA-10 speed), and existing Ethernet hardware/software support. Dolphins will be produced shortly in limited quantity by Xerox EOS, primarily for the ARPANET computer science community, and will be available for delivery beginning in August 1981. Their cost is currently higher than that expected for comparable systems several years from now. However, the

immediate purchase of the limited quantity planned will be cost-effective in allowing research to proceed in the SUMEX-AIM community on software that will be needed to exploit these later systems.

5) If we purchase the five INTERLISP-Dolphins in year 09, a significant increase in LISP processing capacity will be added to the SUMEX-AIM resource earlier than would be possible with VAX/INTERLISP. Even though these are intended primarily for stand-alone use, they nevertheless will afford badly needed relief for the overloaded central machines since the people using the Dolphins will not be running INTERLISP simultaneously on the KI-10's or 2020. In quantitative terms, taking a Dolphin to be about equal in speed to two KA-10's, the five Dolphins will roughly treble our current dual KI-10/2020 computing capacity.

Based on this rationale, it seems clearly to be in the best interest of the SUMEX-AIM community to delay the acquisition of the first VAX system and to accelerate the purchase of the five Professional Workstations.

Other Hardware Development

We have undertaken other hardware efforts as appropriate during the past year. Most significant of these was the development of a controller for a printer in the Stanford Oncology Clinic to support the ONCOCIN evaluation getting underway. This printer is part of an existing internal information system in use by the clinic. In order to integrate the printout from ONCOCIN sessions, we needed to provide a flexible connection to the SUMEX facility spoolers. We built a Z80-based microprocessor controller that senses status of the printer and performs buffering, flow control, and data rate conversion so it can act as a remote printer to the SUMEX machines when needed for ONCOCIN sessions.

In addition we have provided broad support to users for terminal and communications connections and repairs.

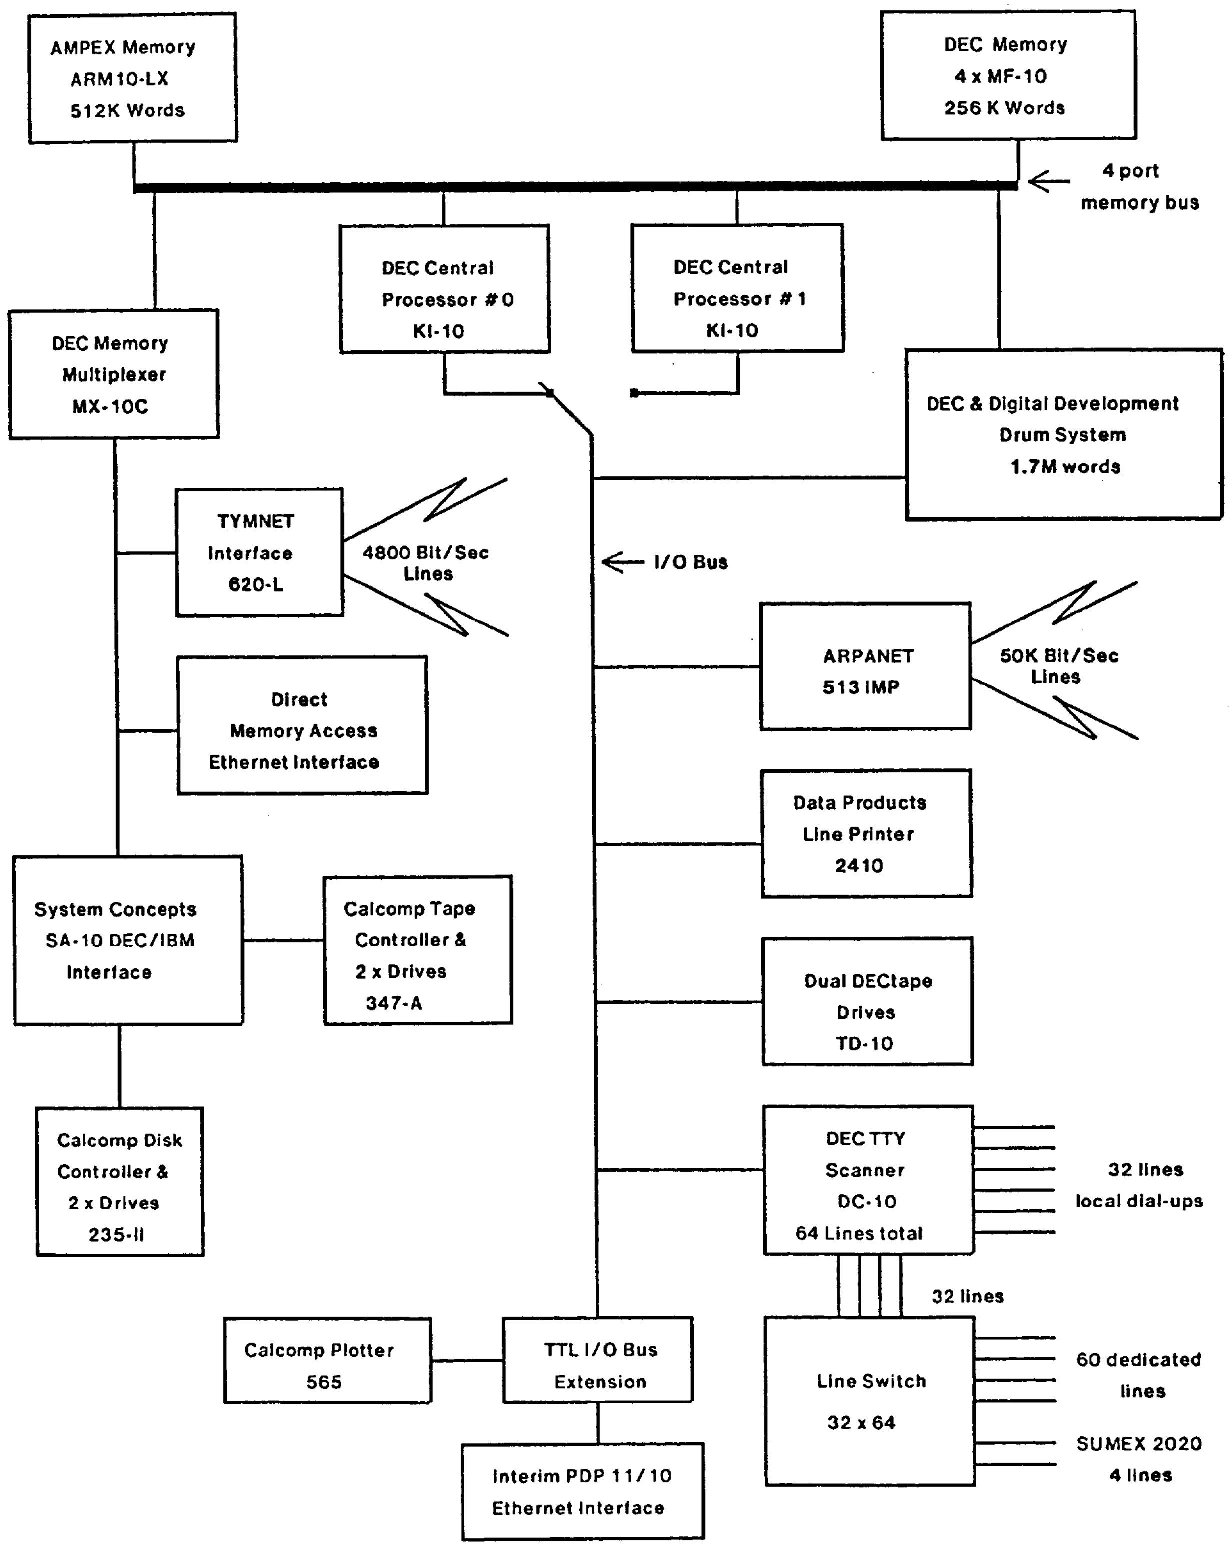

Figure 1. Current SUMEX-AIM KI-10 Computer Configuration

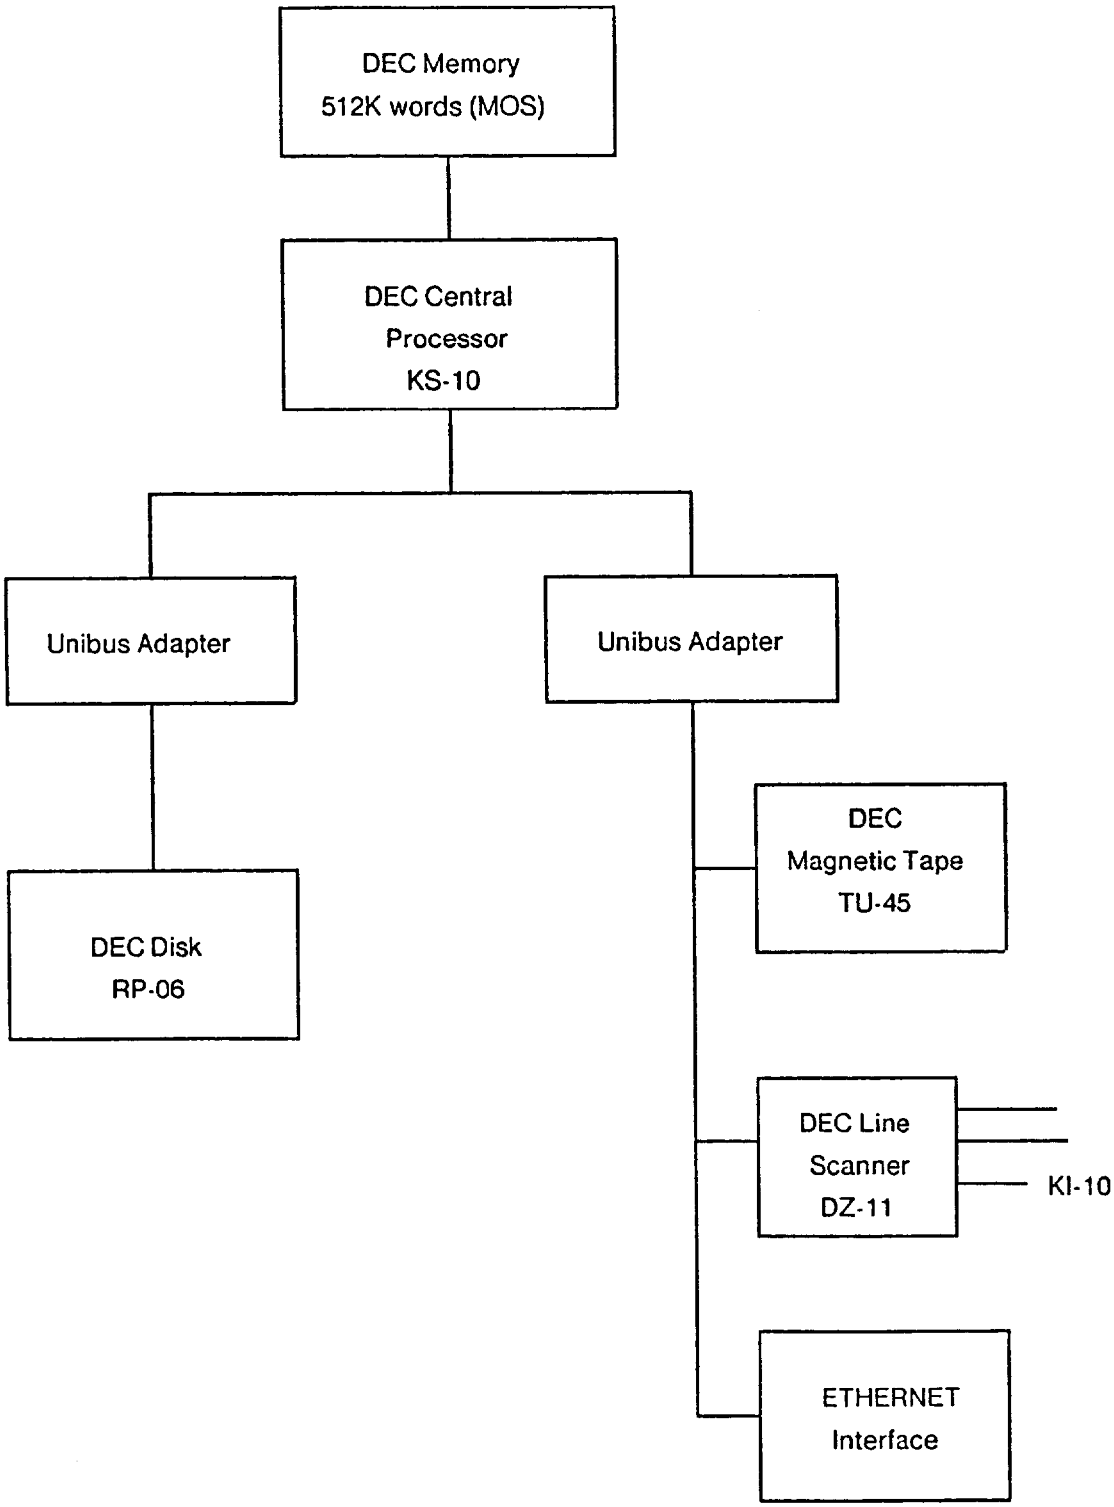

Figure 2. Current SUMEX-AIM 2020 Computer Configuration

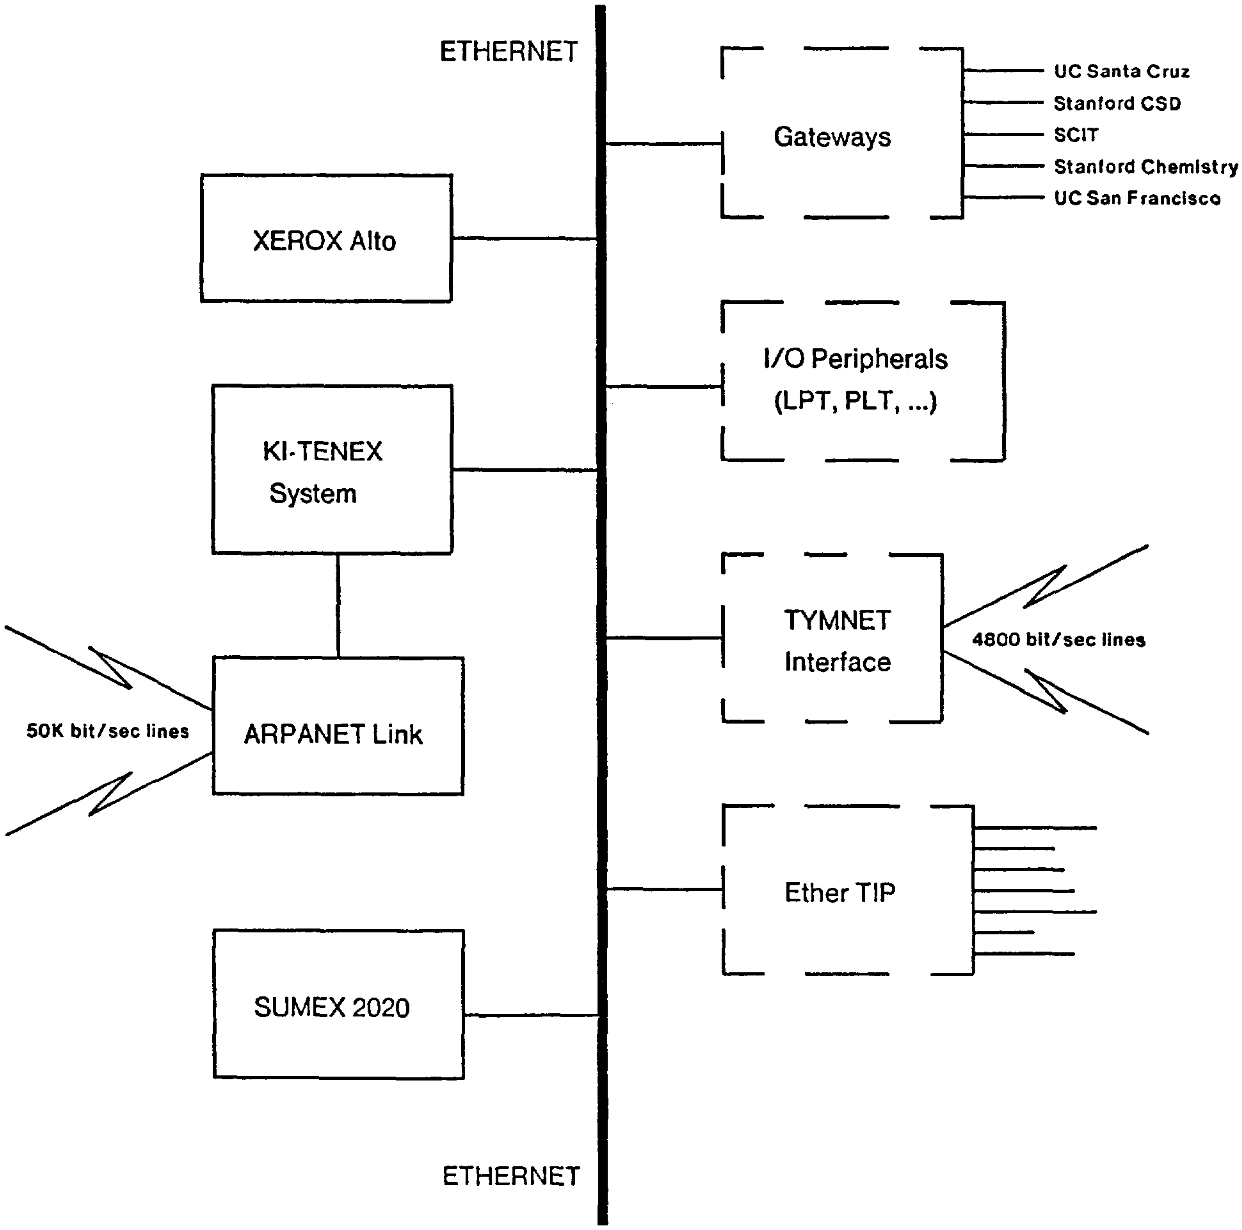

Figure 3. Intermachine Connections via ETHERNET

I.A.3.2 System Software

Our monitor software work this past year has concentrated on several areas including changes to support hardware development projects, upgrading and enhancing network interface service, correcting encountered system bugs, and implementing new features for better user community support. In addition we have invested substantial effort in becoming familiar with the VAX/UNIX system which will play a key role in our future research efforts.

Hardware Implementation

In parallel with our principal hardware efforts this past year to extend and improve local Ethernet connections, the necessary monitor changes to support the new hardware are being made. The largest effort, for which debugging is still on-going, has been the direct memory access Ethernet interface for the KI-10 duplex. This interface has been partially completed, including developing new interrupt service routines and facilities for doing hardware debugging during time-sharing so as not to disrupt availability of the system. Completion of this work has been delayed by placing higher priority on building a gateway connection to the Department of Computer Science building across campus. Since the KI-10's have a working, albeit inefficient, PDP-11 interface already, we decided our limited development resources were better used in establishing this badly needed new capability.

Additional work has gone into upgrading software support for the terminal hardline switch (SLM) developed last year. These improvements were to correct several problems in assigning terminals to available lines and to improve user feedback on system status while negotiating for a connection.

Network Interface Service

Effective January 1, 1981, the ARPANET formally changed the standard for packet "leaders" to allow addressing more hosts on an Interface Message Processor (IMP) and more IMP's on the network. This change required substantial upgrades to the monitor ARPANET service routines including the internal handling of data packets and two new JSYS's that communicate with user programs about network information. We imported much of the new code from ARPANET sites working on the development of network software (especially USC-ISI and SRI) but considerable work was required to adapt it to our dual-processor monitor and operating environment. The changeover went extremely smoothly with most users unaware that a change was taking place. We expect further changes to be required by early 1983 when the higher level communication protocols will move from NCP (network control protocol) to TCP (transmission control protocol).

This past year we have also continued to develop the Ethernet PUP software including improved hardware interfaces discussed above and numerous bug fixes. Many of the bug fixes relate to interactions between the PUP management software and other parts of the monitor such as the teletype handler. The Ethernet software is running very reliably now.

Monitor Bug Fixes and Improvements

We have continued to repair important bugs in our TENEX monitor. In general the system runs extremely reliably with most problems coming from explicit hardware malfunctions or periods of instability following significant monitor changes.

We found an additional number of subtle bugs in the system this past year that had been causing various problems. By now, all of the "obvious" bugs have been located and so those remaining are much more elusive, occurring infrequently or only after a long chain of rare events that is difficult to reconstruct. Examples of fixes include:

1) After an extended period of uptime, TYMNET users found all the ports to SUMEX in apparent use. This only happened after about 6.5 days of continuous uptime, itself a rather rare event. After a long search, we found an invalid index into one of the TYMNET connection database arrays which instead of testing the appropriate state bit, was testing a high order bit of a timer field. Thus, when that bit came on after being up for a long period, the test erroneously detected the port in use.

2) With the installation of the operational Ethernet, the overall timing of system functions changed, including the management of the drum service. Commands for page transfers were sent to the drum controller asynchronously as the requests were placed on the queue. This was done in the drum interrupt service and timing was such that new transfers could be posted "on the fly". As the Ethernet became operational, the timing of interrupt handling changed so that attempts to post these new transfer requests came at the wrong time for the drum controller and caused command sequence errors. A temporary fix was made to avoid this conflict but we still want to rework the drum management software to optimize performance.

3) Finally, users are invariably able to design system call arguments that present special cases to the monitor routines which don't work. We have repaired a number of such problems in various string handling JSYS's and in the floating point output JSYS.

System Loading Controls

We previously reported on the system load controls we have implemented on the KI-10 duplex to allocate available system capacity effectively among projects and users according to Executive Committee guidelines. These continue to operate effectively and we have not made any substantial changes in this area. All communities (National, Stanford, and Staff) are under load controls now. We have adjusted relative priorities for projects in the national community in accordance with Executive Committee reviews of the community in August 1980.

We have instituted a mechanism for reserving the 2020 for demonstrations and developmental testing of various expert systems (e.g., DENDRAL, ONCOCIN, etc.). Because of the unpredictability of usage during these reserved times, we feel that too much of the 2020 capacity is lost by simply dedicating the machine to such users. We are now reevaluating the reservation system, probably in favor of a "pie-slice" system that will guarantee dedicated users a large fraction of the machine but which allows other useful work to go on when their demand is low.

Executive Program

We have made several changes and improvements in the SUMEX EXECutive program this past year:

1) Many of the features of the EXEC that enhance its "friendliness" require access to auxiliary files. When we come up after a crash and there is file system damage to repair, these files may be compromised and in general extraneous file access at such times is undesirable. Thus we carefully reviewed the internals of the EXEC and made changes so that in debugging mode it operates "bare bones". All unneeded file accesses and interactions are eliminated.

2) Because of the difficulty in collecting definitive data about user experiences with network connections, we implemented a log of involuntary disconnects from network terminals. This log allows us to better correlate disconnects, looking for instances when all TYMNET users are disconnected or all users from a given node are disconnected as opposed to drops by individual users which may be caused by hanging up the telephone. We have now collected a database of these disconnects covering several months and indications are that TYMNET users are being dropped occasionally through some sort of network glitch. We are developing programs to better analyze these data so we can distinguish problems at the SUMEX end from those in the network so appropriate solutions can be worked out.

3) We implemented several layers of access constraints for GENET users (see page 84) including a limit to the number of simultaneous login's and a requirement for a user password to restrict access for commercial users. These developments have in fact limited the growth of the GENET community as recommended by AIM Executive Committee policy.

I.A.3.3 Network Communications

A highly important aspect of the SUMEX system is effective communication with remote users and between the growing number of machines available within the SUMEX resource. In addition to the economic arguments for terminal access, networking offers other advantages for shared computing. These include improved inter-user communications, more effective software sharing, uniform user access to multiple machines and special purpose resources, convenient file transfers, more effective backup, and co-processing between remote machines.

We continue to base our remote communication services on two networks - TYMNET and ARPANET for reasons detailed in previous annual reports. Users asked to accept a remote computer as if it were next door will use a local telephone call to the computer as a standard of comparison. Current network terminal facilities do not quite accomplish the illusion of a local call. Data loss is not a problem in most network communications - in fact with the more extensive error checking schemes, data integrity is higher than for a long distance phone link. On the other hand, networking relies upon shared community use of communication lines to procure widespread geographical coverage at substantially reduced cost. However, unless enough total line capacity is provided to meet peak loads, substantial queueing and traffic jams result in the loss of terminal responsiveness. Limited responsiveness for character-oriented TENEX interactions continues to be a problem for network users.

TYMNET

TYMNET provides broad geographic coverage for terminal access to SUMEX, spanning the country and also increasingly accessible from foreign countries. Technical aspects of our connection to TYMNET have remained unchanged this past year and have continued to operate reasonably reliably. As noted earlier, however, users have complained periodically about having their connections dropped and we have implemented a data collection facility in the EXEC program to help document and classify these failures. There are definitely episodes in which all connections are lost and the jobs are detached. These occur about once every few days but we are still analyzing these data to try to separate out local from network causes.

TYMNET has made few technical changes to their network that affect us other than to broaden geographical coverage. The previous network delay problems are still apparent although better cross-country trunks into New York and New England are available improving service there. TYMNET is still primarily a terminal network designed to route users to an appropriate host and more general services such as outbound connections originated from a host or interhost connections are only done on an experimental basis. This presumably reflects the lack of current economic justification for these services among the predominantly commercial users of the network. Whereas TYMNET is developing interfaces meeting X.25 protocol standards, the internal workings of the network will likely remain the same, namely, constructing fixed logical circuits for the duration of a connection and multiplexing characters in packets over each link between network nodes from any users sharing that link as part of their logical circuit.

We have continued to purchase TYMNET services through the NLM contract with TYMNET, Inc. Because of current tariff provisions, there is no longer an economic advantage to this based on usage volume. SUMEX charges are computed on its usage volume alone and not the aggregate volume with NLM's contribution to achieve a lower rate. We have implemented the "dedicated port" charging system for SUMEX use and have realized a substantial reduction in monthly usage costs. We will continue to work closely with NIH-BRP and NLM to achieve the most cost-effective purchase of these services.

ARPANET

We continue our advantageous connection to the Department of Defense's ARPANET, now managed by the Defense Communications Agency (DCA). Current ARPANET geographical and logical maps are shown in Figure 4 and Figure 5 on page 23. This connection has facilitated close collaboration with the Rutgers-AIM facility which is also on the net. Consistent with our long-standing agreements with ARPA and DCA we are enforcing a policy that restricts the use of ARPANET to users who have affiliations with DoD-supported contractors and system/software interchange with cooperating network sites. We are somewhat unique in this policy among other network sites since NIH has not become a member of the "sponsor's group" for the network. We would strongly encourage this step so that biomedical users could have more uniform access to the superior facilities of the ARPANET. This will become increasingly important as more NIH-sponsored sites desire access to the net and each other.

We have maintained good working relationships with other sites on the ARPANET for system backup and software interchange. Such day-to-day working interactions with remote facilities would not be possible without the integrated file transfer, communication, and terminal handling capabilities unique to the ARPANET. The ARPANET is also key to maintaining on-going intellectual contacts between SUMEX projects such as the Stanford Heuristic Programming Project authorized to use the net and other active AI research groups in the ARPANET community.

As indicated in the discussion of monitor software development, we implemented a significant change in ARPANET software support this past January 1, 1981. This change added support for the extended (96-bit) leader for packets that allow more Interface Message Processors (IMP) on the network and more ports per IMP. Substantial changes to the monitor network control program were necessary as well as to various user-level programs (TELNET, FTP, NETSER, RSSER, NETSTAT, etc.). The changeover went extremely smoothly with most users unaware of any effect.

ETHERNET

A substantial portion of our system effort this past year went into continued development of local network facilities to refine the connection between the KI-10 duplex and the 2020, to extend our network ties to other parts of campus (especially to the Computer Science Department building where the Heuristic Programming Project sits), and to prepare for the addition of new hardware in the renewal grant. As indicated in the earlier sections on monitor software and hardware, much has been done to implement more effective and efficient low-level system network connection facilities for our host systems.

We have also developed a number of software tools as a basis for implementing various kinds of Ethernet servers. These have been done in the language C, primarily because it is the language on which UNIX is based, has an active support community, and is being used for other network software that may be useful for our work. Specific areas of development include:

1) Server operating system: We have developed a simple operating system for use in servers that provides low-level interface to the Ethernet, hardware dependent interrupt service, process scheduling capabilities, and a series of defined monitor calls for invoking communication functions. This system is written initially for the PDP-11 but will also be portable to the MC-68000.

2) Higher Level Protocols: We have written routines that provide datagram, rendezvous/termination, and byte sequential protocol facilities on which other services such as EFTP, TELNET, etc. can be based.

3) We have written software for an Ethernet-to-Ethernet gateway that will establish connections between the SUMEX machine room and the Computer Science Department across campus. This system runs currently on a PDP-11/10 and supports dynamic assimilation of a routing table, periodic broadcast of this information to other hosts on connected networks, routing of addressed packets between connected networks, forwarding of key broadcast packets to allow distribution of network directories, and recording of gateway event status reports.

4) We have developed a wide range of diagnostic programs to assist in Ethernet software development including hardware diagnostics and downloading and debugging software.

5) We are actively working on the design of an Ethernet TIP to provide more terminal ports for the SUMEX system.

INTERNET SOFTWARE

One of the issues confronting the development of complex network-based systems, interconnected by gateway machines, is the support of internet communication of various kinds. For example, when a user at one of the Stanford Ethernet hosts wants to send a message to someone at MIT on a Chaosnet host, the mail handling programs have to know how to do the routing and the mail server programs have to be prepared to receive such mail for forwarding. Similarly, when establishing terminal telnet connections between such sites, the path of the link should be established automatically with the intervening sites merely acting as relay stations. In conjunction with groups at MIT and Stanford CSD, we have been developing prototypical systems for internet mail handling and telnet connections. The mail system is most highly developed and currently knows how to route messages between hosts on the Stanford Ethernet, the ARPANET, the MIT Chaosnet, the MIT LCSnet, and the Dialnet. This system has been operating since February.

We are also running a version of TELNET developed by Mark Crispin at SU-SCORE that allows a user to establish a connection across network boundaries without having to log into each intervening gateway and telnet further to the next station of a path to the desired destination host.

Figure 4. ARPANET Geographical Network Map

ARPANET GEOGRAPHIC MAP, MARCH 1981

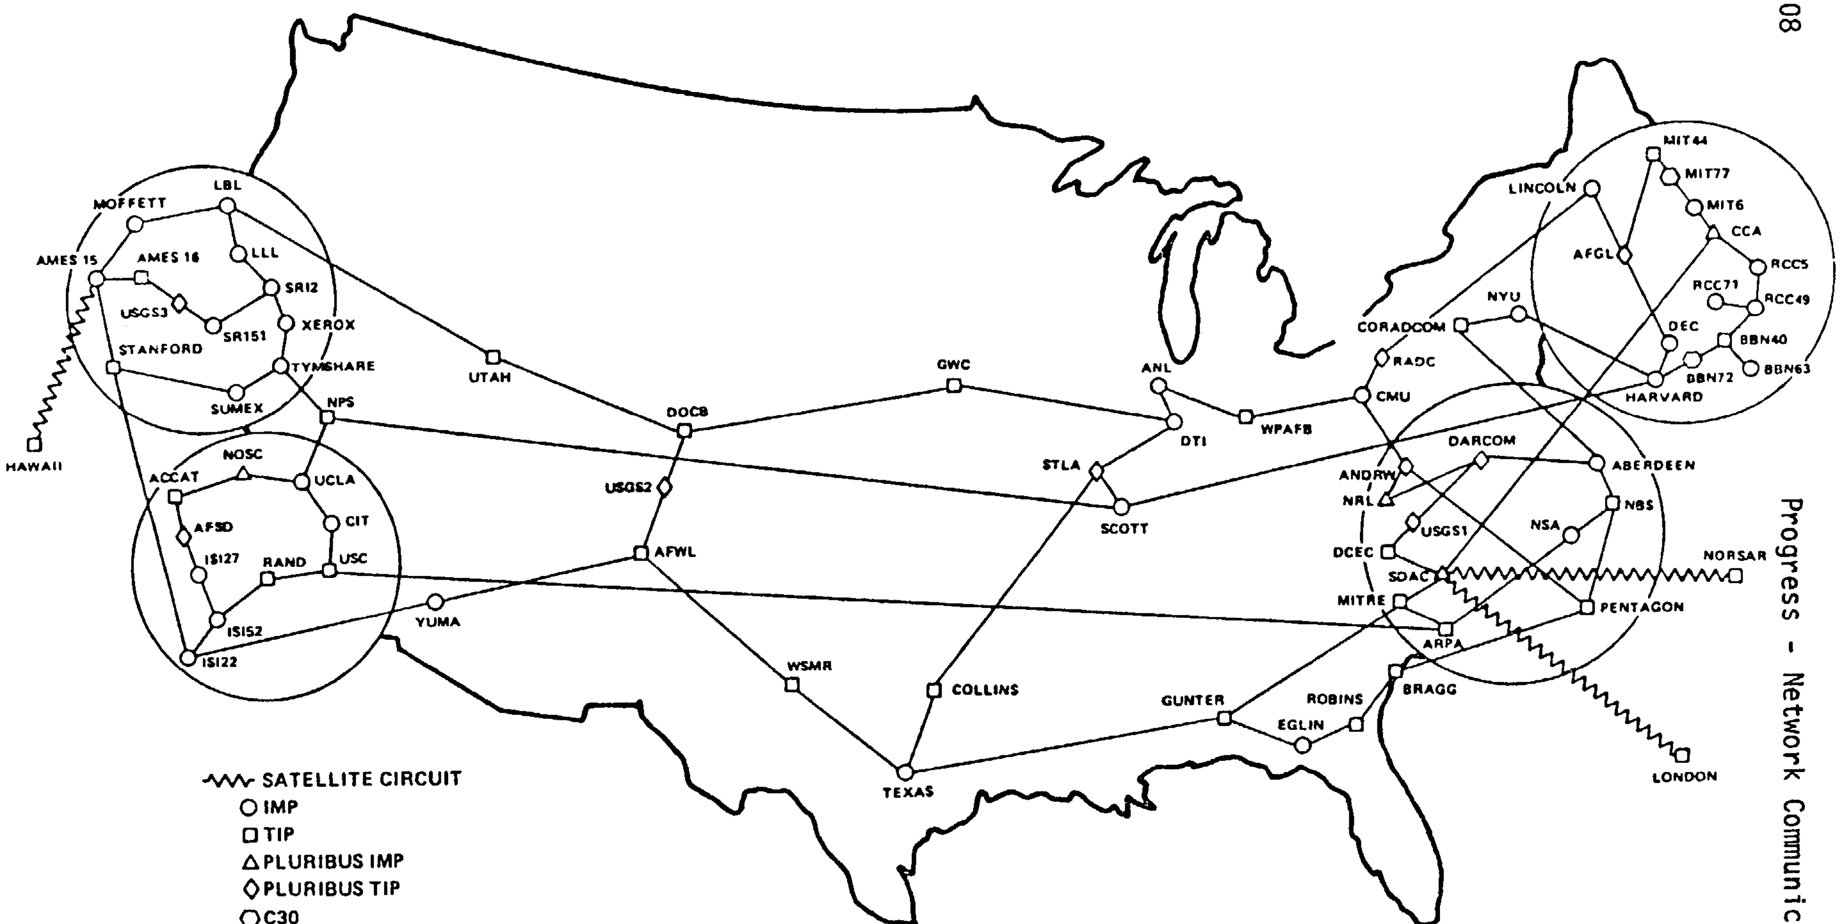

(NOTE: THIS MAP DOES NOT SHOW ARPA'S EXPERIMENTAL SATELLITE CONNECTIONS)
NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

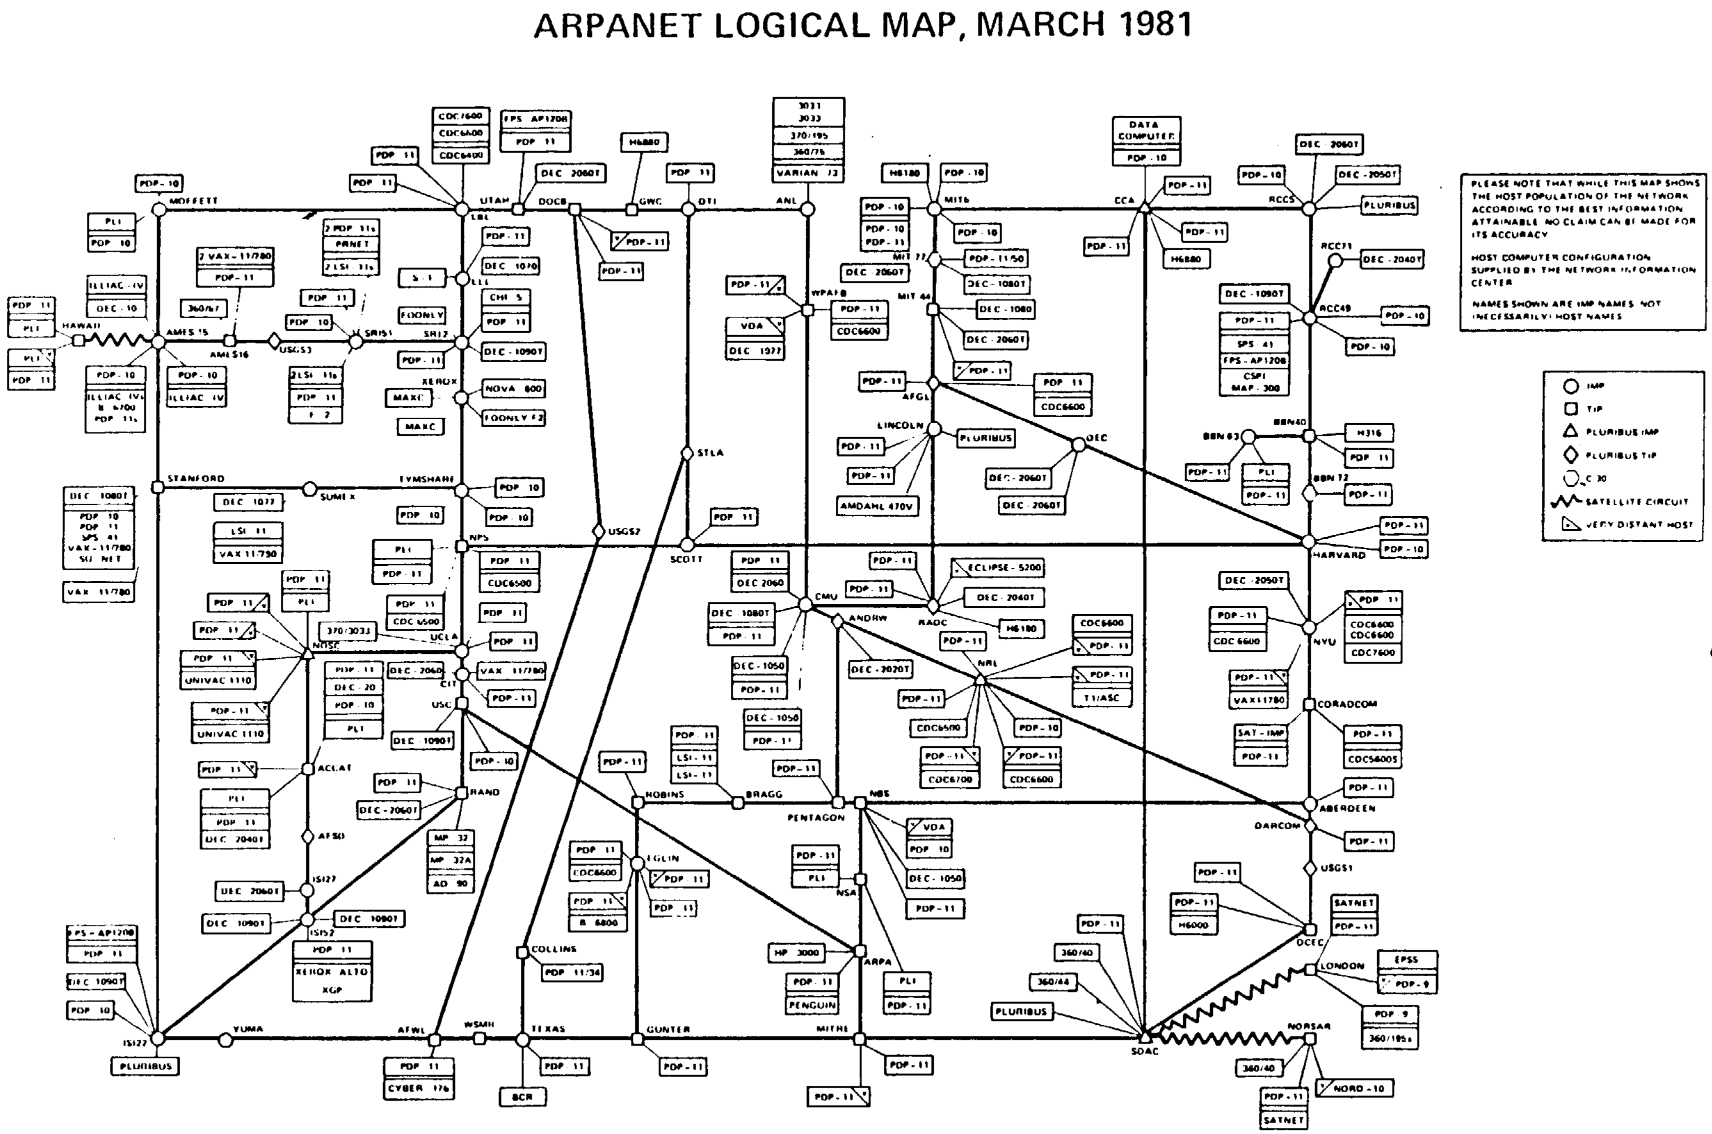

Figure 5. ARPANET Logical Network Map

I.A.3.4 User Software

We have continued to assemble and maintain a broad range of utilities and user support software. These include operational aids, statistics packages, DEC-supplied programs, improvements to the TOPS-10 emulator, text editors, text search programs, file space management programs, graphics support, a batch program execution monitor, text formatting and justification assistance, magnetic tape conversion aids, and many more. Over the past year we have undertaken several significant development efforts to provide needed new programs to the SUMEX-AIM community. These include:

1) TTYFTP - Many groups have had the need to move files between computers and do not have the sophisticated facilities of ARPANET or other local networks to help. These include for example the transfer of data between the PUFF project at Pacific Medical Center in San Francisco and SUMEX and the movement of instrument data in support of the Ultrasound Imaging (Ob-Gyn) project. We reported last year on the development of a file transfer program usable over any teletype line (hardline, dial-up, TYMNET, etc.) which incorporates appropriate control protocols and error checking. The design is based on the DIALNET protocols developed by Crispin at the Stanford AI Laboratory and extended by our group to achieve machine and data source independence. This past year we have had a number of requests from outside groups in similar situations for copies of this software. We have distributed copies to Rutgers, Stanford Research Institute, and the University of Texas.

2) C Compiler - We spent considerable effort bringing up a workable C compiler at SUMEX that would generate code for our KI-TENEX system and also cross compile to generate code for PDP-11's and other machines. We imported an early version of a TOPS-20 C compiler from MIT and adapted it for our system. The linker, code generator, and runtime package for this system were suitably modified to work under TENEX and code generators for other machines developed. The PDP-10 version still generates quite inefficient code in that bytes occupy full 36-bit words rather than being packed. This is satisfactory for debugging purposes but would have to be fixed if C were to become a system programming language for future TENEX work (we do not anticipate this).

3) TV Editor and Display Terminal Support - Much work was done to extend the TV editor which is widely used at SUMEX. This was done by importing the work done by Hedrick at Rutgers, adapting it to our needs, and extending it for additional features. Important improvements include multiple string searches, string replacement facilities, large block text relocation, and support for additional display terminals (Infoton, Zenith H-19, Concept-100, ADDS Regent 60, and Hewlett-Packard 2600 series terminals). We have also agreed to unify the sources for TV so that closer compatibility with other groups will exist (Rutgers, SUMEX, USC-ISI, SRI, and Stanford CONTEXT).

4) TYPER - For the AIM workshop held at Stanford last summer, we developed a program to assist with on-line typescript manipulation. The goal of the workshop (see Section I.E.1 on page 79) was to present a deeper insight into the workings of developing AI programs by interactively tracing sessions using them. In order to assure a reasonably organized presentation that could be prepared beforehand, we developed TYPER to allow a presenter to display a typescript of a typical session. TYPER provides facilities to randomly move between parts of the typescript, to display a table of contents, and to manage a hierarchical presentation of various parts of the session. At the highest level, only an overview of program operation need be given. By interactive commands, successive layers of detail can be flashed on the display screen as the discussion proceeds or as questions arise.

We have also implemented extensions and maintenance updates to many other existing programs including, for example, EDIR (a directory editing program), DUMPER (the file system backup program), BSYS (the user file archiving program), PA-1050 (the TOPS-10 compatibility package), BBD (the bulletin board program), and PUB (a text formatting program).

I.A.3.5 Documentation and Education

We have spent considerable effort to develop, maintain, and facilitate access to our documentation so as to accurately reflect available software. The HELP and Bulletin Board subsystems have been important in this effort. As subsystems are updated, we generally publish a bulletin or small document describing the changes. As more and more changes occur, it becomes harder and harder for users to track down all of the change pointers. Within manpower limits, we are in a continuous process of reviewing the existing documentation system for compatibility with the programs now on line and to integrate changes into the main documents. This will also be done with a view toward developing better tools for maintaining up-to-date documentation.

I.A.3.6 Software Compatibility and Sharing

At SUMEX-AIM we firmly believe in importing rather than reinventing software where possible. As noted above, a number of the packages we have brought up are from outside groups. Many avenues exist for sharing between the system staff, various user projects, other facilities, and vendors. The advent of fast and convenient communication facilities coupling communities of computer facilities has made possible effective intergroup cooperation and decentralized maintenance of software packages. The TENEX sites on the ARPANET have been a good model for this kind of exchange based on a functional division of labor and expertise. The other major advantage is that as a by-product of the constant communication about particular software, personal connections between staff members of the various sites develop. These connections serve to pass general information about software tools and to encourage the exchange of ideas among the sites.

Certain common problems are now regularly discussed on a multi-site level. We continue to draw significant amounts of system software from other ARPANET sites, reciprocating with our own local developments. Interactions have included mutual backup support, experience with various hardware configurations, experience with new types of computers and operating systems, designs for local networks, operating system enhancements, utility or language software, and user project collaborations. We have been able to import many new pieces of software and improvements to existing ones in this way. Examples of imported software include the message manipulation program MSG, TENEX SAIL, PASCAL, TENEX SOS, INTERLISP, the RECORD program, ARPANET host tables, and many others. Reciprocally, we have exported our contributions such as the crash analysis program, drum page migration system, KI-10 page table efficiency improvements, GTJFN enhancements, PUB macro files, the bulletin board system,. MAINSAIL, SPELL, SNDMSG enhancements, our BATCH monitor, and improved SA-10 software.

There are also several important examples of joint development efforts such as the internet mailer program (XMAILR). Because this program incorporates facilities for routing mail through many networks, it is important that the various sections of the program dealing with these specialized protocols be developed by the groups with expertise in the appropriate technology. Network connections have made a joint effort possible involving MIT, Stanford SCORE, and SUMEX.

We spent considerable effort developing a preliminary version of a TENEX/TOPS-20 compatibility package. The issue here is that as DEC develops TOPS-20, even though it is TENEX-like, it is not TENEX compatible and vice versa. Thus, a hope was to write a program that would resolve these compatibilities automatically rather than to force special adaptations for the two operating systems. The kinds of incompatibilities that exist include PDP-20 machine instructions that do not exist on earlier machines, new JSYS calls, incompatible changes to old JSYS calls, different syntax and facilities for device/file names, and different handling of error returns (types of return and error codes). It has proven unworkable to effectively handle all of these problems at the user level. Monitor changes are required to implement the widely used error return features (ERJMP/ERCAL) and make handling of other incompatibilities easier. We have not accomplished a complete compatibility package in any sense but have implemented requisite monitor changes and have developed several user packages that help emulate TOPS-20 JSYS calls for programs running on TENEX. We do not foresee being able to completely and effectively solve these problems within the expected lifetime of existing TENEX machines.

Finally, we have also assisted groups that have interacted with SUMEX user projects get access to software available in our community. For example, Prof. Dreiding's group in Switzerland became interested in some of the system software available here after attending the DENDRAL CONGEN workshops (see Section II.A.1.3 on page 103). We have provided him with the non-licensed programs requested. We have also provided software to Professor Bodmer's group at the Imperial Cancer Research Group in England in collaboration with the MOLGEN project (see Section II.A.1.5 on page 136).

I.A.3.7 Core Research

Over the past year we have supported several core research activities aimed at developing information resources, basic AI research, and tools of general interest to the SUMEX-AIM community. Principal areas of current effort include:

1) The AI Handbook which is a compendium of knowledge about the field of Artificial Intelligence being compiled by Professor Feigenbaum and collaborators. The handbook is broad in scope, covering all of the important ideas, techniques, and systems developed during 20 years of research in AI in a series of articles. Each is about four pages long and is a description written for non-AI specialists and students of AI. The handbook will be published in three volumes, the first of which is now on the market published by William Kaufmann, Inc. The AI Handbook effort is described in more detail in Section II.A.1.2 on page 99 and an outline of the current contents of the handbook can be found in Appendix B.

2) The AGE project which is an attempt to isolate inference, control, and representation techniques from previously developed knowledge-based programs; reprogram them for domain independence; write a rule-based interface that will help a user understand what the package offers and how to use the modules; and make the package available to other members of the AIM community. A more detailed description of progress on the AGE package can be found in Section II.A.1.1 on page 91.

It should be noted that SUMEX is providing only partial support for these projects with complementary support coming from an ARPA contract to the Heuristic Programming Project.

I.A.3.8 Resource Operations Statistics

The following data give an overview of various aspects of SUMEX-AIM resource usage. There are five sub-sections containing data respectively for:

1) Overall resource loading data

2) Relative system loading by community

3) Individual project and community usage

4) Network usage data

5) System reliability data

1. Overall resource loading data

The following plots display several different aspects of system loading over the life of the project. These include total CPU time delivered per month, the peak number of jobs logged in, and the peak load average. The monthly "peak" value of a given variable is the average of the daily peak values for that variable during the month. Thus, these "peak" values are representative of average monthly loading maxima and do not reflect the largest excursions seen on individual days, which are much higher.

These data show well the continued growth of SUMEX use and the self-limiting saturation effect of system load average, especially after installation of our overload controls early in 1978. Since late 1976, when the dual processor capacity became fully used, the peak daily load average has remained between about 5.5 and 6. This is a measure of the user capacity of our current hardware configuration and the mix of AI programs.

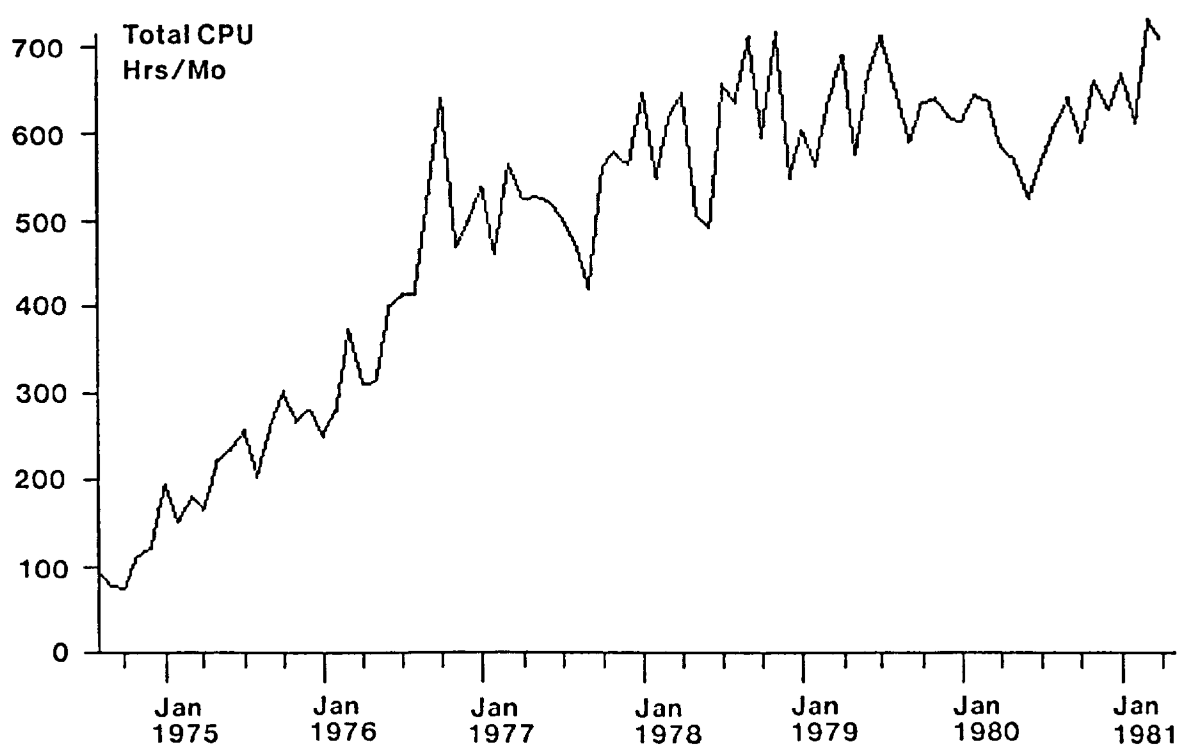

Figure 6. Total CPU Time Consumed by Month

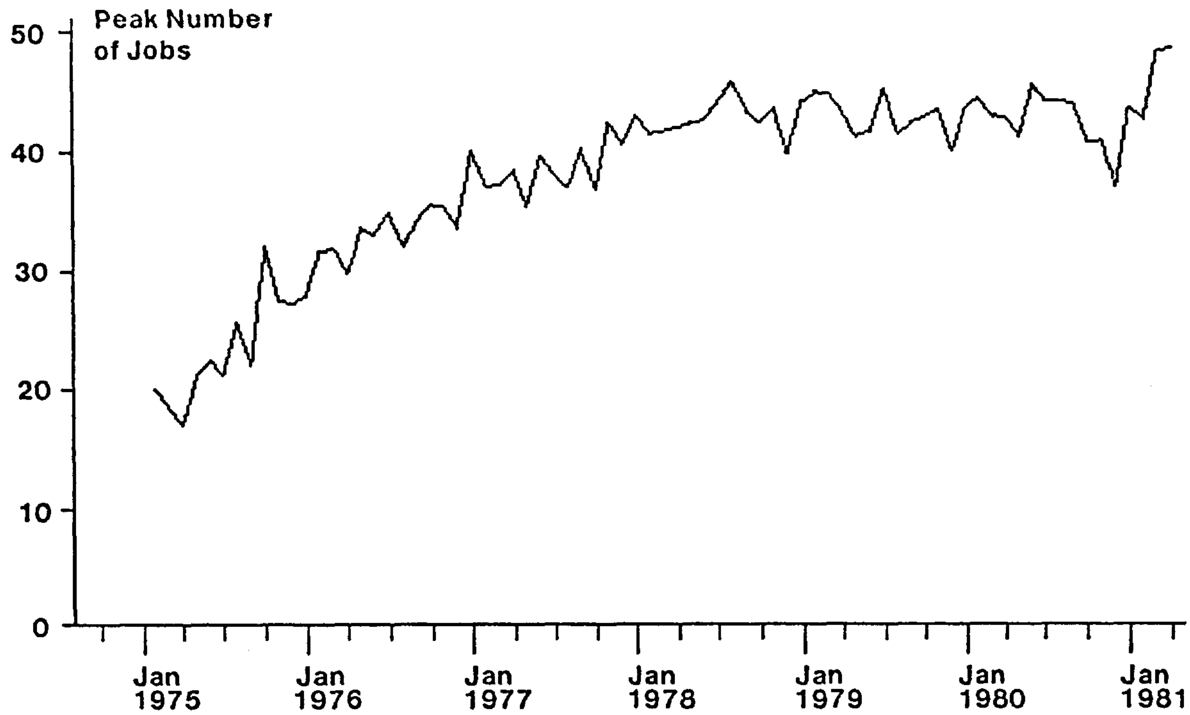

Figure 7. Peak Number of Jobs by Month

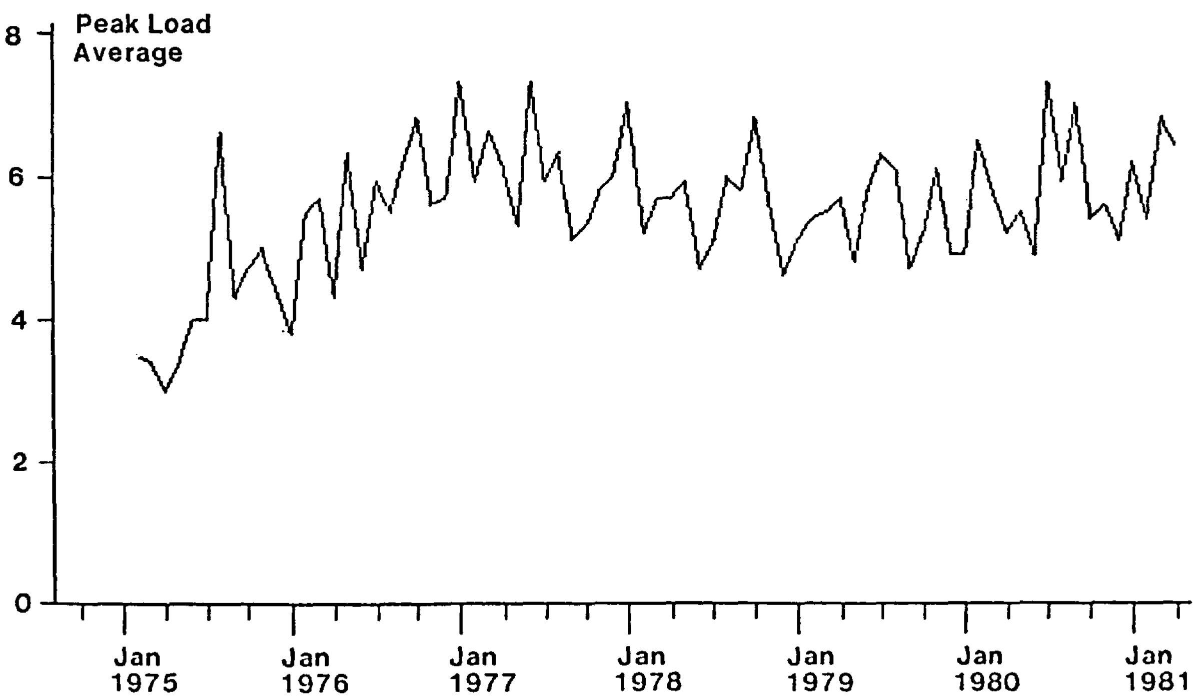

Figure 8. Peak Load Average by Month

2. Relative System Loading by Community

The SUMEX resource is divided, for administrative purposes, into 3 major communities: user projects based at the Stanford Medical School, user projects based outside of Stanford (national AIM projects), and common system development efforts. As defined in the resource management plan approved by BRP at the start of the project, the available system CPU capacity and file space resources are divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems, documentation, etc.

The monthly usage of CPU and file space resources for each of these three communities relative to their respective aliquots is shown in the plots in Figure 9 and Figure 10. Terminal connect time is shown in Figure 11. It is clear that the Stanford projects have held an edge in system usage despite our efforts at resource allocation and the substantial voluntary efforts by the Stanford community to utilize non-prime hours. This reflects the maturity of the Stanford group of projects relative to those getting started on the national side and has correspondingly accounted for much of the progress in AI program development to date.

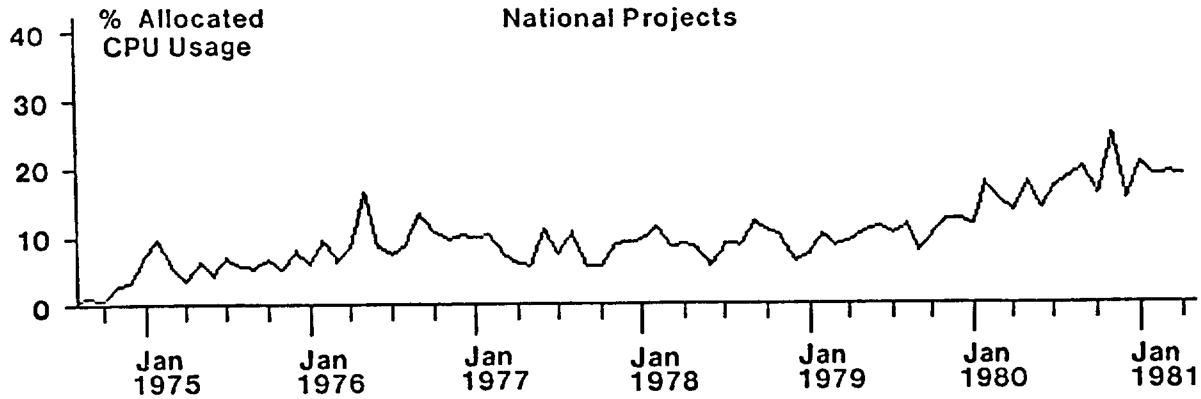

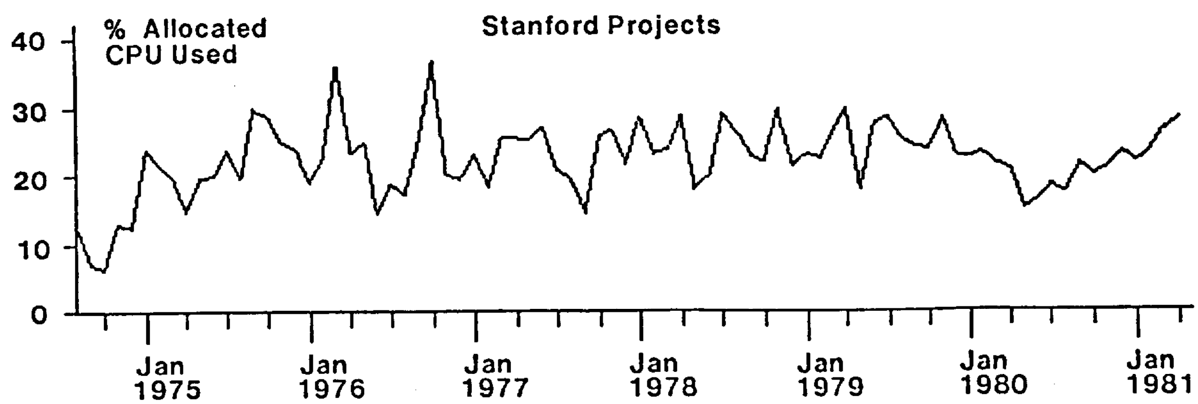

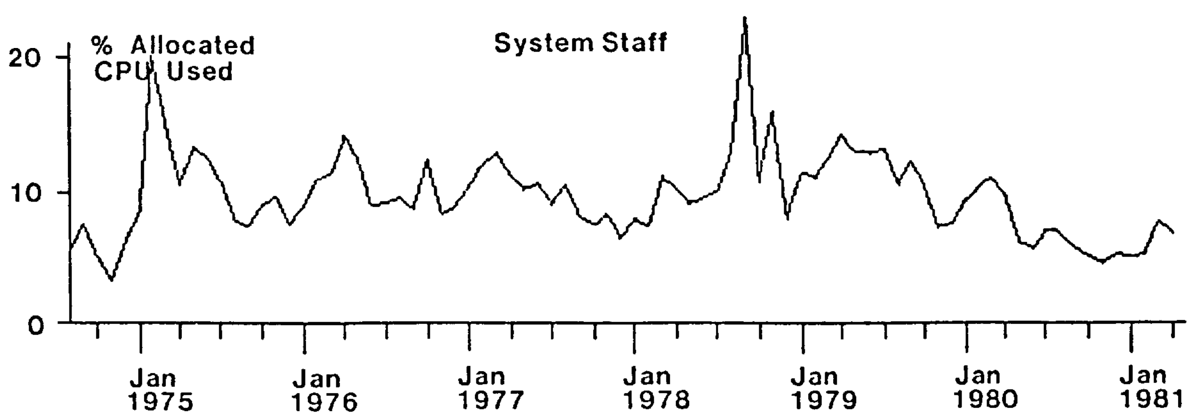

Figure 9. Monthly CPU Usage by Community

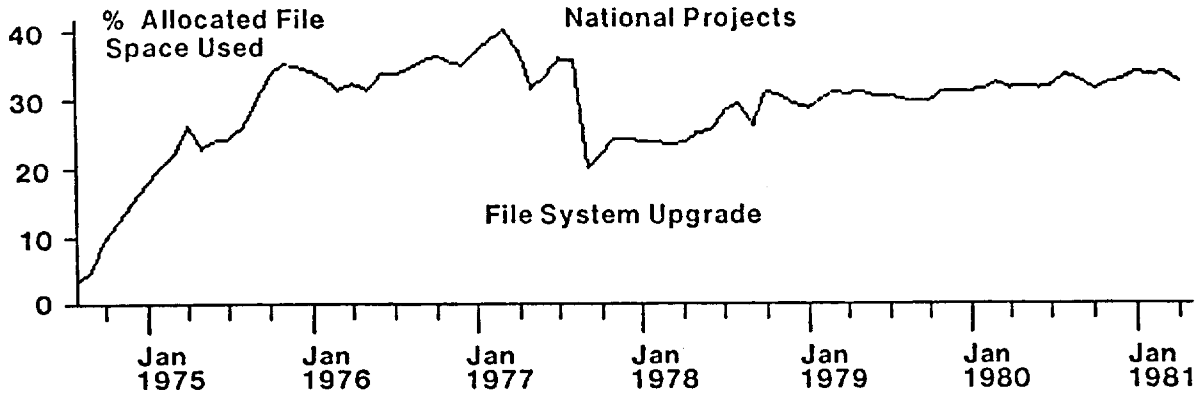

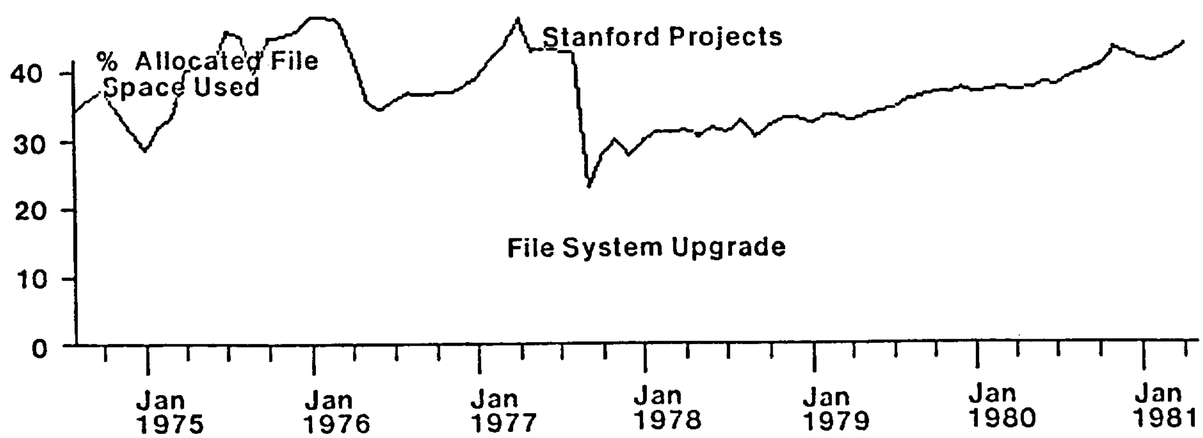

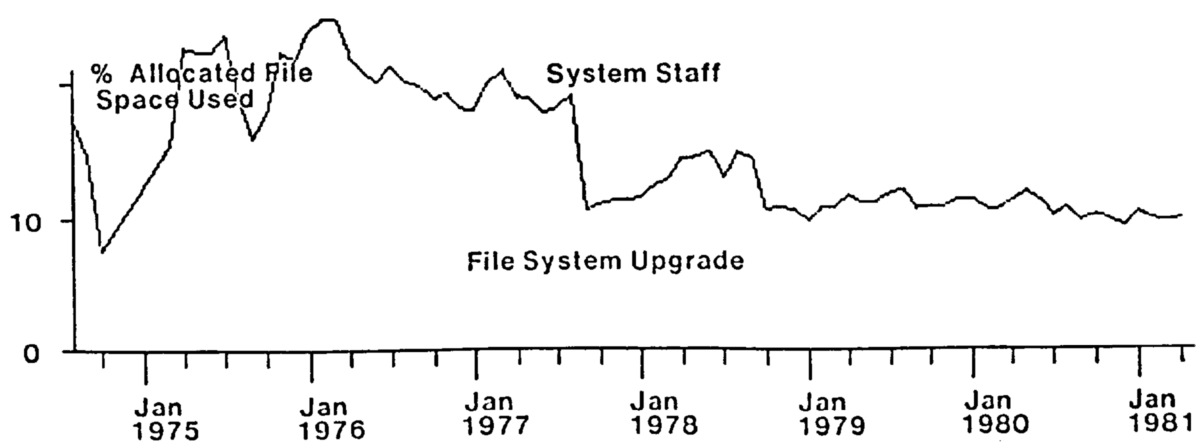

Figure 10. Monthly File Space Usage by Community

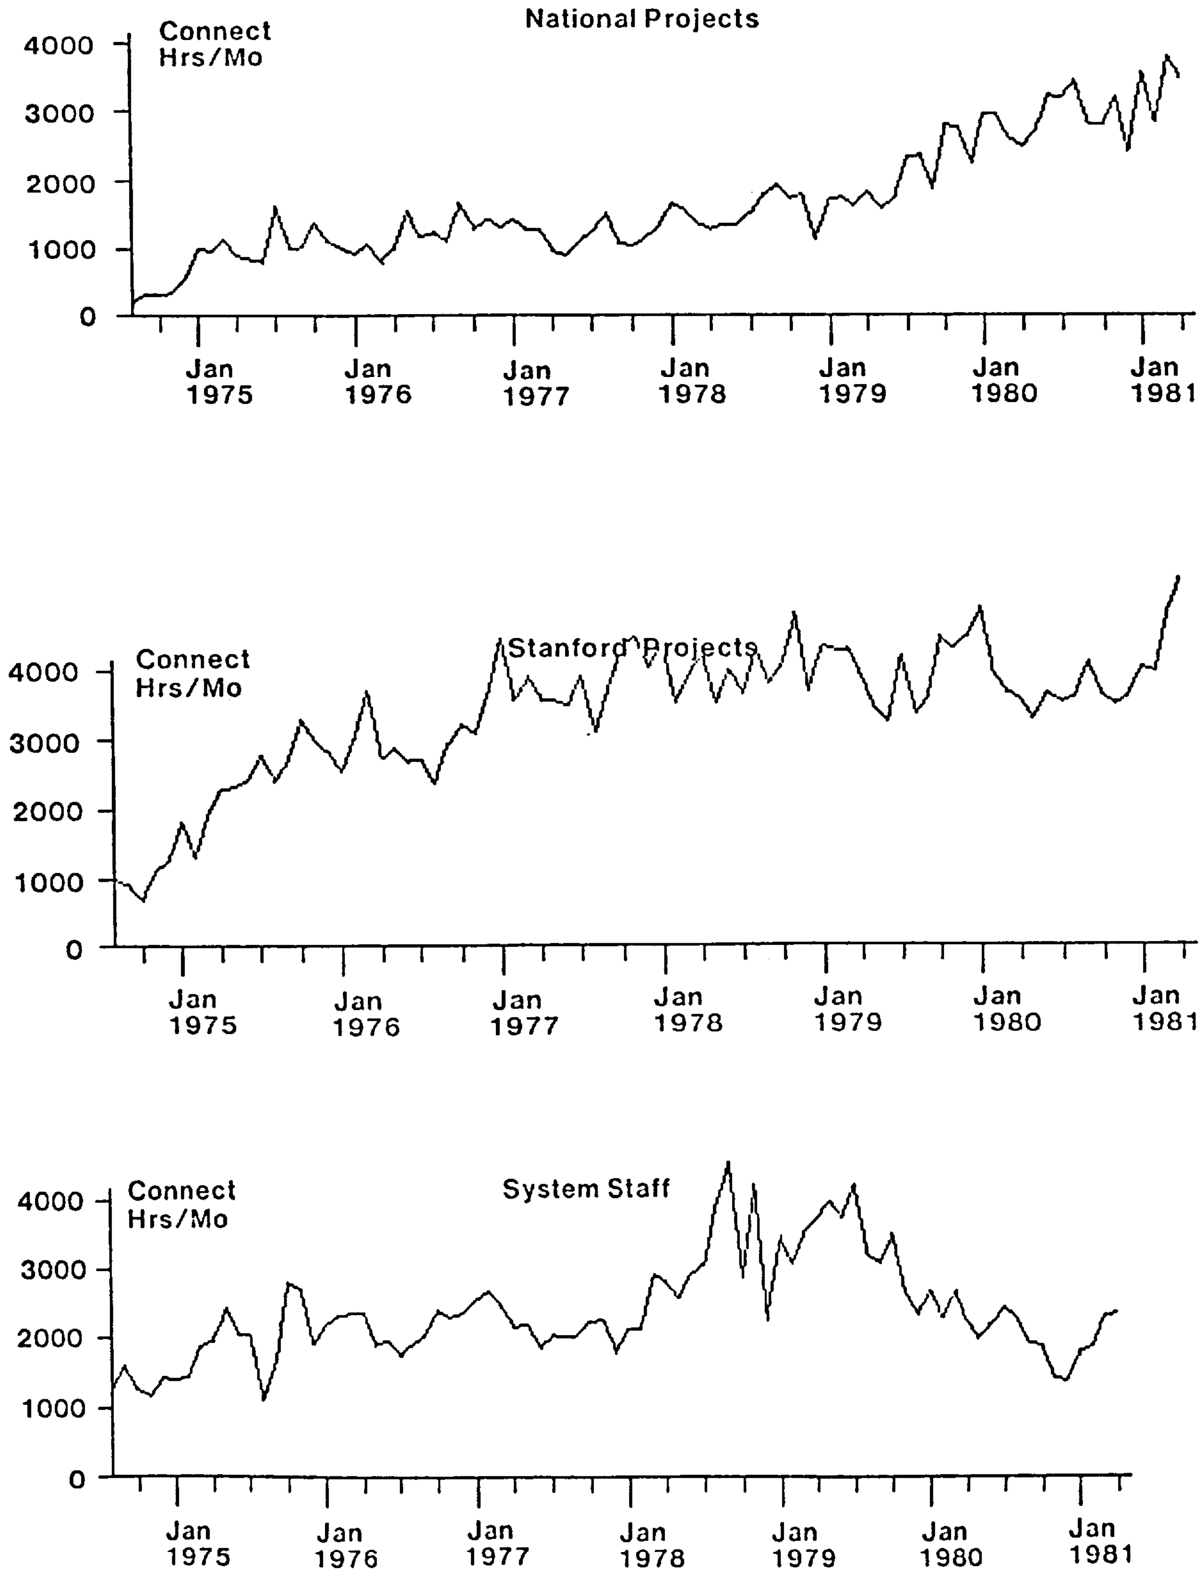

Figure 11. Monthly Terminal Connect Time by Community

3. Individual Project and Community Usage

The table following shows cumulative resource usage by project during the past grant year. The entries include a summary of the operational funding sources (outside of SUMEX-supplied computing resources) for currently active projects, total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May 1980 and April 1981.

Again the well developed use of the SUMEX resource by the Stanford community can be seen. It should be noted that the Stanford projects have voluntarily shifted a substantial part of their development work to non-prime time hours which is not explicitly shown in these cumulative data. It should also be noted that a significant part of the DENDRAL, MYCIN, AGE, AI Handbook, and MOLGEN efforts, here charged to the Stanford aliquot, support development efforts dedicated to national community access to these systems. The actual demonstration and use of these programs by extramural users (e.g., the GENET community) is charged to the national community in the "AIM USERS" category, however.

Several of the projects admitted to the National AIM community use the Rutgers-AIM resource as their home base. We do not explicitly list these projects in this annual report covering the Stanford SUMEX-AIM resource. We do record information about the Rutgers resource itself, however, and note its separate resource status with the flag "[Rutgers-AIM]".

Resource Use by Individual Project - 5/80 through 4/81

National AIM Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) ACT Project "Acquisition of Cognitive Procedures" John R. Anderson, Ph.D. Carnegie-Mellon Univ. ONR N00014-77-C-0242 9/78-9/80 $175,000 NSF IST-80-15357 2/81-2/84 $186,000	95.81	1214.90	2362.
2) SECS Project "Simulation & Evaluation of Chemical Synthesis" W. Todd Wipke, Ph.D. U. California, Santa Cruz NIH RR-01059-03S1 7/80-12/81 $36,949 NIH/NCI N01-CP-75816 1/80-7/81 $74,394	837.86	11968.57	10239
3) Mod Human Cogn Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado NIE-G-78-0172 9/80-8/81 $46,537 NIMH MH-15872-9-13 6/80-5/81 $32,880 ONR N00014-78-C-0433 6/80-5/81 $60,000 ONR N00014-78-C-0165 1/80-6/81 $85,000	150.08	2084.62	898
4) CADUCEUS Project "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. University of Pittsburgh NIH RR-01101-04 7/80-6/81 $465,199 NLM LM03710-01 7/80-6/81 $148,458 NLM LM03589-01 7/80-6/81 $32,750	344.92	5975.17	8365

5)	SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota NSF SE079-13036 NICHD T36-HD-17151 NICHD HD-01136 NSF/BNS-77-22075 NLM/NSF proposals pending	.55	13.94	9
6)	PUFF-VM Project "Biomedical Knowledge Engineering in Clinical Medicine" John J. Osborn, M.D. Inst. Medical Sciences, San Francisco Edward A. Feigenbaum, Ph.D. Stanford University NIH GM-24669 9/78-8/81 $164,000 (*) Renewal pending	97.01	5651.45	3785
7)	SCP Project "Simulation of Cognitive Processes" James G. Greeno, Ph.D. Alan M. Lesgold, Ph.D. University of Pittsburgh NIE-G-80-0014 12/80-11/81 $2,627,067 ONR N00014-79-C-0215 10/80-9/81 $247,053 NSF/NIE SED78-22289 12/78-8/81 $149,967	5.13	236.97	931
8)	*** [Rutgers-AIM] *** Rutgers Project "Computers in Biomedicine" Saul Amarel, D.Sc. NIH RR-00643 12/80-11/81 $495,079	12.87	339.45	8653

9) AIM Pilot Projects			
AI-COAG	4.02	96.89	672
EXCHANGE	7.64	100.89	60
HEADMED	.66	29.03	770
KRL	.57	6.99	268
MDX	.53	30.96	28
MELANOMA	.99	20.11	9
MISL	.73	14.42	813
SPA	1.42	30.10	251
SPEECH	11.10	224.32	537
UI	8.51	379.92	46
	-----	-------	----
AIM Pilot Totals	36.17	933.63	3459
10) AIM Administration	8.98	449.06	4141
11) AIM Users on Stanford Projects			
AGE	2.44	51.90	27
DENDRAL	62.59	905.71	1056
HMF	35.06	1026.88	2735
HPP	5.30	118.23	115
MOLGEN	411.85	5757.19	1095
MYCIN	7.44	620.40	105
AIM-Associates	1.65	37.00	153
Guest (all projects)	12.88	135.08	239
	------	-------	----
AIM User Totals	569.21	8652.39	5529
	------	--------	-----
Community Totals	2158.59	37520.15	48371

Stanford Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) AGE Project (Core) "Generalization of AI Tools" Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA-903-80-C-0107 (**) (partial support)	388.99	4625.17	3575
2) AI Handbook Project (Core) Edward A. Feigenbaum, Ph.D. Dept. Computer Science ARPA MDA-903-80-C-0107 (**) (partial support)	31.56	1680.34	2554
3) DENDRAL Project "Resource Related Research: Computers in Chemistry" Carl Djerassi, Ph.D. Dept. Chemistry NIH RR-00612-12 5/81-4/82 $237,387	554.47	8170.34	16366
4) EXPEX Project "Expert Explanation" Edward H. Shortliffe, M.D., Ph.D. Depts. Medicine/Computer Science ONR NR 049-479 1/81-12/81 $140,825	26.99	1163.37	513
5) MOLGEN Project "Experiment Planning System for Molecular Genetics" Edward A. Feigenbaum, Ph.D. Bruce G. Buchanan, Ph.D. Laurence H. Kedes, M.D. Douglas L. Brutlag, Ph.D. Depts. Computer Science/ Medicine/Biochemistry NSF ECS-8016247 10/80-9/81 $146,582 (*)	325.62	6146.69	6734

6)	MYCIN Projects "Computer-based Consult. in Clin. Therapeutics" Bruce G. Buchanan, Ph.D. Edward H. Shortliffe, M.D., Ph.D. Depts. Medicine/Computer Science NSF MCS-79-03753 7/79-3/81 $146,152 ONR/ARPA N00014-79-C-0302 3/79-3/82 $396,325 Kaiser Fdn. 7/79-12/80 $20,000 NLM LM-03395 7/80-6/81 $47,845 NLM LM-00048 7/80-6/81 $39,107	706.62	11373.94	13261
7)	Protein Struct Modeling "Heuristic Comp. Applied to Prot. Crystallog." Edward A. Feigenbaum, Ph.D. Dept. Computer Science NSF MCS-79-33666 12/79-11/81 $35,318	89.26	1194.98	3916
8)	RX Project Depts. Computer Science/Medicine Robert L. Blum, M.D. Gio C.M. Wiederhold, Ph.D. NLM New Invest. 7/79-6/82 $90,000 NCHSR 4/79-3/81 $35,000	64.15	1683.05	2222
9)	Stanford Pilot Projects			
	DECIDER (E. Johnson)	.27	3.70	0
	STRUCT (Abarbanel)	12.41	317.34	87
	SCANR (Brinkley)	36.81	758.69	771
		------	-------	----
	Stanford Pilot Totals	49.49	1079.73	858
10)	Stanford and HPP Assoc.	196.49	9538.63	11609
		-------	--------	-----
	Community Totals	2433.64	46656.24	61608

SUMEX Staff	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	602.53	20648.45	9571
2) System Associates	78.42	3081.22	4205
3) Misc. Usage	.16	1.84	770
	-------	---------	------
Community Totals	681.11	23731.51	14546

System Operations	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	2230.32	94975.93	69094
	=======	=========	======
Resource Totals	7503.66	202883.83	193619

(*) Award includes indirect costs.

(**) Supported by a larger ARPA contract MDA-903-80-C-0107 awarded to the Stanford Computer Science Department:

	Current Year (10/80-11/15/81)	Total Award (10/79-9/82)
Heuristic Programming Project	$ 538,262	$1,613,588
VLSI/CAD Network-based Graphics Development Resource	214,851	685,374
	-----------	----------
Total award	$ 753,113(*)	$2,298,962(*)

4. Network Usage Statistics

The plots in Figure 12 and Figure 13 show the monthly network terminal connect time for TYMNET and ARPANET. This forms the major billing component for SUMEX-AIM TYMNET usage. The terminal connect time does not reflect the time spent in file transfers and mail forwarding.

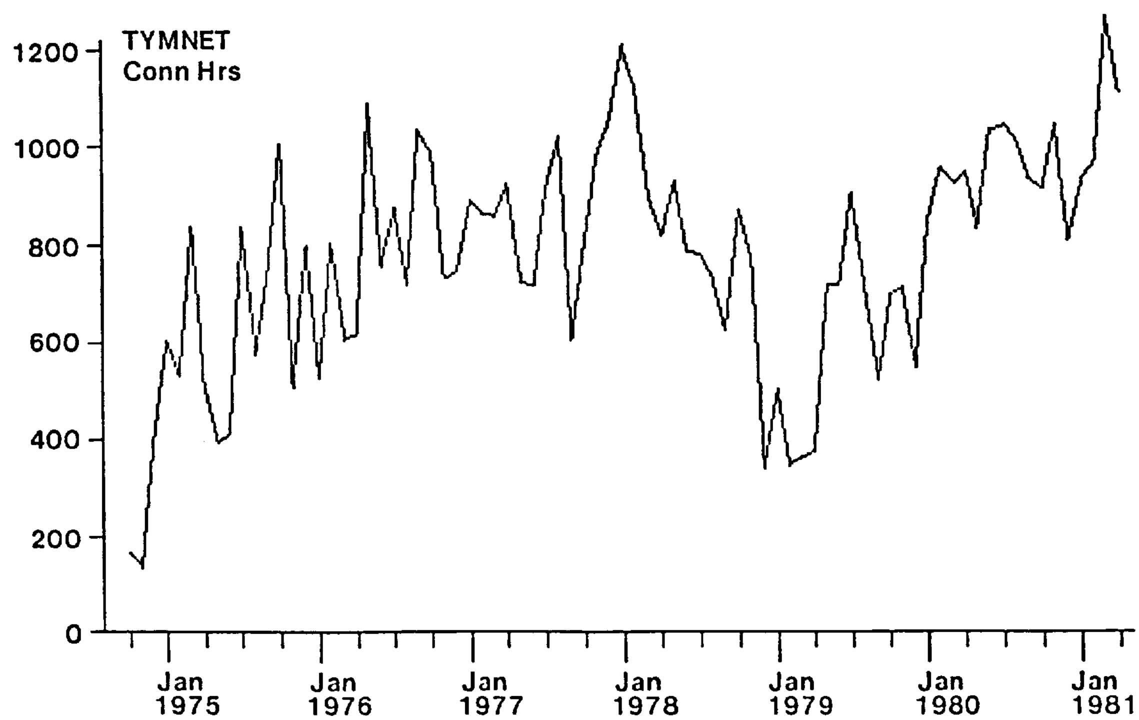

Figure 12. TYMNET Terminal Connect Time

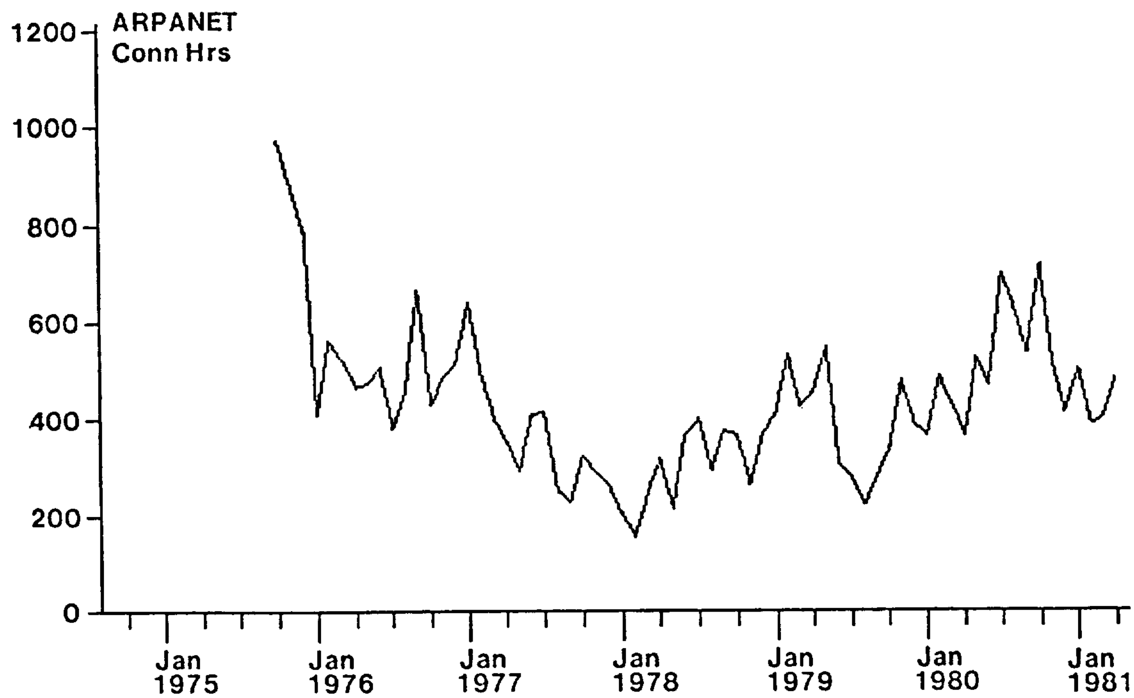

Figure 13. ARPANET Terminal Connect Time

5. System Reliability

System reliability has been very good on average with several periods of particular hardware or software problems. The table below shows monthly system reloads and downtime for the past year. It should be noted that the number of system reloads is greater than the actual number of system crashes since two or more reloads may have to be done within minutes of each other after a crash to repair file damage or to diagnose the cause of failure.

				1980						1981		
	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR	APR
RELOADS												
Hardware	3	2	2	8	8	16	2	0	6	2	4	0
Software	0	3	2	4	1	2	1	1	3	2	1	4
Environmental	0	1	0	3	1	0	0	1	0	0	0	1
Operator Error	0	0	1	0	0	0	1	1	0	0	2	1
Unknown Cause	2	0	0	0	0	0	0	0	0	0	0	0
	--	--	--	--	--	--	--	--	--	--	--	--
Totals	5	6	5	15	10	18	4	3	9	4	7	6
DOWNTIME (Hrs)												
Unscheduled	2	17	6	18	7	11	1	2	6	2	4	6
Scheduled	30	14	79	17	17	28	17	7	15	4	4	14
	--	--	--	--	--	--	--	--	--	--	--	--
Totals (Hrs)	32	31	95	35	24	39	18	9	21	6	8	20

TABLE 1. System Reliability by Month

I.A.3.9 SUMEX Staff Publications

The following are publications for the SUMEX staff and include papers describing the SUMEX-AIM resource and on-going research as well as documentation of system and program developments. Many of the publications documenting SUMEX-AIM community research are from the individual collaborating projects and are detailed in their respective reports (see Section II on page 89). Publications for the AGE and AI Handbook core research projects are given there.

[1] Carhart, R.E., Johnson, S.M., Smith, D.H., Buchanan, B.G., Dromey, R.G., and Lederberg, J, Networking and a Collaborative Research Community: A Case Study Using the DENDRAL Programs, ACS Symposium Series, Number 19, Computer Networking and Chemistry, Peter Lykos (Editor), 1975.

[2] Levinthal, E.C., Carhart, R.E., Johnson, S.M., and Lederberg, J., When Computers Talk to Computers, Industrial Research, November 1975

[3] Wilcox, C. R., MAINSAIL - A Machine-Independent Programming System, Proceedings of the DEC Users Society, Vol. 2, No. 4, Spring 1976.

[4] Wilcox, Clark R., The MAINSAIL Project: Developing Tools for Software Portability, Proceedings, Computer Application in Medical Care, October, 1977, pp. 76-83.

[5] Lederberg, J. L., Digital Communications and the Conduct of Science: The New Literacy, Proc. IEEE, Vol. 66, No. 11, Nov 1978.

[6] Wilcox, C. R., Jirak, G. A., and Dageforde, M. L., MAINSAIL - Language Manual, Stanford University Computer Science Report STAN-CS-80-791 (1980).

[7] Wilcox, C. R., Jirak, G. A., and Dageforde, M. L., MAINSAIL - Implementation Overview, Stanford University Computer Science Report STAN-CS-80-792 (1980).

Mr. Clark Wilcox also chaired the session on "Languages for Portability" at the DECUS DECsystem10 Spring '76 Symposium.

In addition, a substantial continuing effort has gone into developing, upgrading, and extending documentation about the SUMEX-AIM resource, the SUMEX-TENEX system, and the many subsystems available to users. These efforts include a number of major documents (such as SOS, PUB, TENEX-SAIL, and MAINSAIL manuals) as well as a much larger number of document upgrades, user information and introductory notes, an ARPANET Resource Handbook entry, and policy guidelines.

I.A.3.10 Future Plans

Our plans for the next grant year are based on those approved by the council review of our recent five-year renewal application scheduled to begin in August 1980. In addition to the specific plans for next grant year (discussed in some earlier sections too), we present a summary below of our overall objectives for the next five-year period to serve as a foundation for future reports. Near and long term objectives and plans for individual collaborating projects are discussed in Section II beginning on page 89.

The goals of the SUMEX-AIM resource are long term in supporting basic research in artificial intelligence, applying these techniques to a broad range of biomedical problems, experimenting with communication technologies to promote scientific interchange, and developing better tools and facilities to carry on this research.

Just as the tone of our renewal proposal derives from the continuing long-term research objectives of the SUMEX-AIM community, our approach derives from the methods and philosophy already established for the resource. We will continue to develop useful knowledge-based software tools for biomedical research based on innovative, yet accessible computing technologies.

For us it is important to make systems that work and are exportable. Hence, our approach is to integrate available state-of-the art hardware technology as a basis for the underlying software research and development necessary to support the AI work.

SUMEX-AIM will retain its broad community orientation in choosing and implementing its resources. We will draw upon the expertise of on-going research efforts where possible and build on these where extensions or innovations are necessary. This orientation has proved to be an effective way to build the current facility and community.

We have built ties to a broad computer science community; have brought the results of their work to the AIM users; and have exported results of our own work. This broader community is particularly active in developing technological tools in the form of new machine architectures, language support, and interactive modalities.

Toward a More Distributed Resource

The initial model for SUMEX as a centralized resource was based on the high cost of powerful computing facilities and not being able to duplicate them readily. This role is evolving with the introduction of more compact and inexpensive computing technology. Our future goals are guided by community needs for more computing capacity and improved tools to build more effective expert systems and to test operational versions of AI programs in real-world settings. In order to meet these needs, we must take advantage of a range of newly developing machine architectures and systems. As a result, SUMEX-AIM will become more a distributed community resource with heterogeneous computing facilities tethered to each other through communications media. Many of these machines will be located physically near the projects or biomedical scientists using them.

We have actively supported proposals from the more mature AIM projects for additional computing facilities tailored to their particular needs and designed to free the main SUMEX resource for new, developing applications projects. To date, the Rutgers resource has acquired a DEC 2050 facility, part of which is allocated for AIM usage; the "Simulation of Cognitive Processes" project has acquired a VAX which supports their needs; and the "Caduceus" (INTERNIST) project is acquiring a VAX to support experimental clinical testing of their program. Our future plans anticipate an even broader diversification of computing resources to meet the need of the AIM community.

The Continuing Role of SUMEX-Central

Even with more distributed computing resources, the central resource will continue to play an important role as a communication crossroads, as a research group devoted to integrating the new software and hardware technologies to meet the needs of medical AI applications, as a spawning ground for new application projects, and as a base for local AI projects. A key challenge will be to maintain the scientific community ties that grew naturally out of the previous colocation within a central facility.

Summary of Five-year Objectives

The following outlines the specific objectives of the SUMEX-AIM resource during the follow-on five year period. Note that these objectives cover only the resource nucleus; near and long-term objectives for individual collaborating projects are discussed in their respective reports in Section II. Specific aims are broken into three categories; 1) resource operations, 2) training and education, and 3) core research.

Resource Operations

1) Maintain the vitality of the AIM community. We will continue to encourage and explore new applications of AI to biomedical research and improve mechanisms for inter- and intra-group collaborations and communications. While AI is our defining theme, we may entertain exceptional applications justified by some other unique feature of SUMEX-AIM essential for important biomedical research. To minimize administrative barriers to the community-oriented goals of SUMEX-AIM and to direct our resources toward purely scientific goals, we plan to retain the current user funding arrangements for projects working on SUMEX facilities. User projects will fund their own manpower and local needs; will actively contribute their special expertise to the SUMEX-AIM community; and will receive an allocation of computing resources under the control of the AIM management committees. There will be no "fee for service" charges for community members. We will also continue to exploit community expertise and sharing in software development; and to facilitate more effective information sharing among projects.

2) Continue to provide effective computational support for AIM community goals. Our efforts will be to extend the support for artificial intelligence research and new applications work; to develop new computational tools to support more mature projects; and to facilitate testing and research dissemination of nearly operational programs. We will continue to operate and develop the existing KI-10/2020 facility as the nucleus of the resource. We will acquire additional equipment to meet developing community needs for more capacity, larger program address spaces, and improved interactive facilities. New computing hardware technologies becoming available now and in the next few years will play a key role in these developments and we expect to take the lead in this community for adapting these new tools to biomedical AI needs. We plan the phased purchase of two VAX computers to provide increased computing capacity and to support large address space LISP development, a 2000M byte file server to meet file storage needs, and a number of single-user "professional workstations" to experiment with improved human interfaces and AI program dissemination.

3) Provide effective and geographically accessible communication facilities to the SUMEX-AIM community for effective remote collaborations, communications among distributed computing nodes,

and experimental testing of AI programs. We will retain the current ARPANET and TYMNET connections for at least the near term and will actively explore other advantageous connections to new communications networks and to dedicated links.

Training and Education

Our goals during the follow-on period for assisting new and established users of the SUMEX-AIM resource are a continuation of those adopted for the previous grant term. Collaborating projects are responsible for the development and dissemination of their own AI programs. The SUMEX resource will provide community-wide support and will work to make resource goals and AI programs known and available to appropriate medical scientists. Specific aims include:

1) Provide documentation and assistance to interface users to resource facilities and programs. We will continue to exploit particular areas of expertise within the community for developing pilot efforts in new application areas.

2) Continue to allocate "collaborative linkage" funds to qualifying new and pilot projects to provide for communications and terminal support pending formal approval and funding of their projects. These funds are allocated in cooperation with the AIM Executive Committee reviews of prospective user projects.

3) Continue to support workshop activities including collaboration with the Rutgers Computers in Biomedicine resource on the AIM community workshop and with individual projects for more specialized workshops covering specific application areas or program dissemination.

Core Research

Our core research efforts will continue to emphasize basic research on AI techniques applicable to biomedical problems and the generalization and documentation of tools to facilitate and broaden application areas.

SUMEX core research funding is complementary to similar funding from other agencies and contributes to the long-standing interdisciplinary effort at Stanford in basic AI research and expert system design. We expect this work to provide the underpinnings for increasingly effective consultative programs in medicine and for more practical adaptations of this work within emerging microelectronic technologies. Specific aims include:

1) Continue to explore basic artificial intelligence issues for knowledge acquisition, representation, and utilization; reasoning in the presence of uncertainty; strategy planning; and explanations of reasoning pathways with particular emphasis on biomedical applications.

2) Support community efforts to organize and generalize AI tools that have been developed in the context of individual application projects. This will include work to organize the present state-of-the-art in AI techniques through the AI Handbook effort and the development of practical software packages (e.g., AGE, EMYCIN, UNITS, and EXPERT) for the acquisition, representation, and utilization of knowledge in AI programs. The objective is to evolve a body of software tools that can be used to more efficaciously build future knowledge-based systems and explore other biomedical AI applications.

Hardware Acquisition Rationale

As discussed in our progress report and supported by collaborating project reports, we have implemented an effective set of computing resources to support AI applications to biomedical research. At the resource core is the KI-TENEX/2020 facility, augmented by portions of the Rutgers 2050 and Stanford SCORE 2060 machines. These have provided an unsurpassed set of tools for the initial phases of SUMEX-AIM development in terms of operating system facilities, human engineering, language support for artificial intelligence program development, and community communications tools. As the size of our community and the complexity of knowledge-based programs have increased, several issues have become important for the continued development and practical dissemination of AI programs:

1) The community has a continuing need for more computing capacity. This arises from the growth of new applications projects, new core research ideas, and the need to disseminate mature systems within and outside of the AIM community. Nowhere is this felt more strongly than among the Stanford community where system access constraints have seriously impeded development progress. A picture of system congestion can be found in the summary of loading statistics beginning on page 29 and in the statements from many of our user projects.

2) Many programs require a larger virtual address space. As AI systems become more expert and encompass larger and more complex domains, they require ever larger knowledge bases and data structures that must be traversed in the course of solving problems. The 256K word address limit of the PDP-10 has constrained program development as discussed in our renewal proposal. Increasing effort has gone into "overlays" resulting in higher machine overhead, more difficulty in making program changes, and lost programmer time. Simpler hardware solutions are needed.

3) AI programs are being tested and disseminated increasingly beyond their development communities. We cannot continue to provide all of the computing resources this implies through central systems like SUMEX. The capacity does not exist. Network communications facilities are not able to support facile human interactions (high speed, improved displays, graphics, and speech/touch modalities). And a grant-supported research environment cannot meet the technical and administrative needs of a "production" community. Thus, we need to explore better ways to package complex AI software and distribute the necessary computing tools cost effectively into the user communities.

No single solution to these requirements for future development is available and we proposed and got peer approval to investigate a variety of machine architectures and support functions over the next grant period including:

1) experimentation with new shared centralized systems

2) distributed single-user "professional workstations"

3) improved communications tools to integrate them together effectively.

In addition to continuing operation of the existing resources, we plan to direct SUMEX research efforts to explore the potential of such newly available systems as solutions to AIM community needs. Our approach will be to integrate a heterogeneous set of network-connected hardware tools, some of which will be distributed through the user community. We will emphasize the development of system and application level software tools to allow effective use of these resources and continue to provide community leadership to encourage scientific communications.

Specific Hardware Plans for Year 09

In our proposal as approved by council, we described a carefully detailed plan for hardware acquisition. One of the approved purchases, the augmentation to the AMPEX core memory for the KI-10 duplex, was approved for the current year 08 and has already been implemented. In addition, for the technical reasons discussed on page 11, we have obtained BRP approval to accelerate the purchase of the five approved professional workstations to year 09 and to delay the first VAX purchase to year 10. The following then is a summary of planned hardware purchase for year 09:

- Buy five Interlisp Dolphin professional workstations for use in developing and experimenting with this means for AI program export and human interface enhancements.

- Develop a file server coupled to SUMEX host machines via the high speed Ethernet. This will minimize the need for redundant large file systems on each host and alleviate the file storage limitations of the AIM community.

- Acquire examples of state-of-the art display equipment including a bit-mapped display station and a hardcopy laser printing device.

- Buy additional required communications, interface, and test equipment to support the above acquisitions and community needs.

Continued Operation of Existing Hardware

The current SUMEX-AIM facilities represent a large existing investment. We do not propose any substantial changes to the existing KI-10 and 2020 hardware systems and we expect them to continue to provide effective community support and serve as a communication nucleus for more distributed resources. The proposed augmentation of the existing KI-10 AMPEX memory box in order to reduce page swapping overhead is underway.

It should be recognized that the KI-10 processors are now 6 years old and will be 12 years old at the end of the proposed grant term. We have already begun to feel maintenance problems from age such as poor electrical contacts from oxidization and dirt, backplane insulation flowing on "tight wraps", and brittle cables. These problems are quite manageable still and we expect to be able to continue reliable operation over the next grant term.

We plan no upgrades to the 2020 configuration. The current file shortage will be remedied in conjunction with that of the rest of the facility by implementing a community file server sharable and accessible via the Ethernet.

For both systems, we are actively working to complete efficient interfaces to the Ethernet to allow flexible, high speed terminal connections, file transfers, and effective sharing of network, printing, plotting, remote links, and other resources. This system will form the backbone for smooth integration of future hardware additions to the resource.

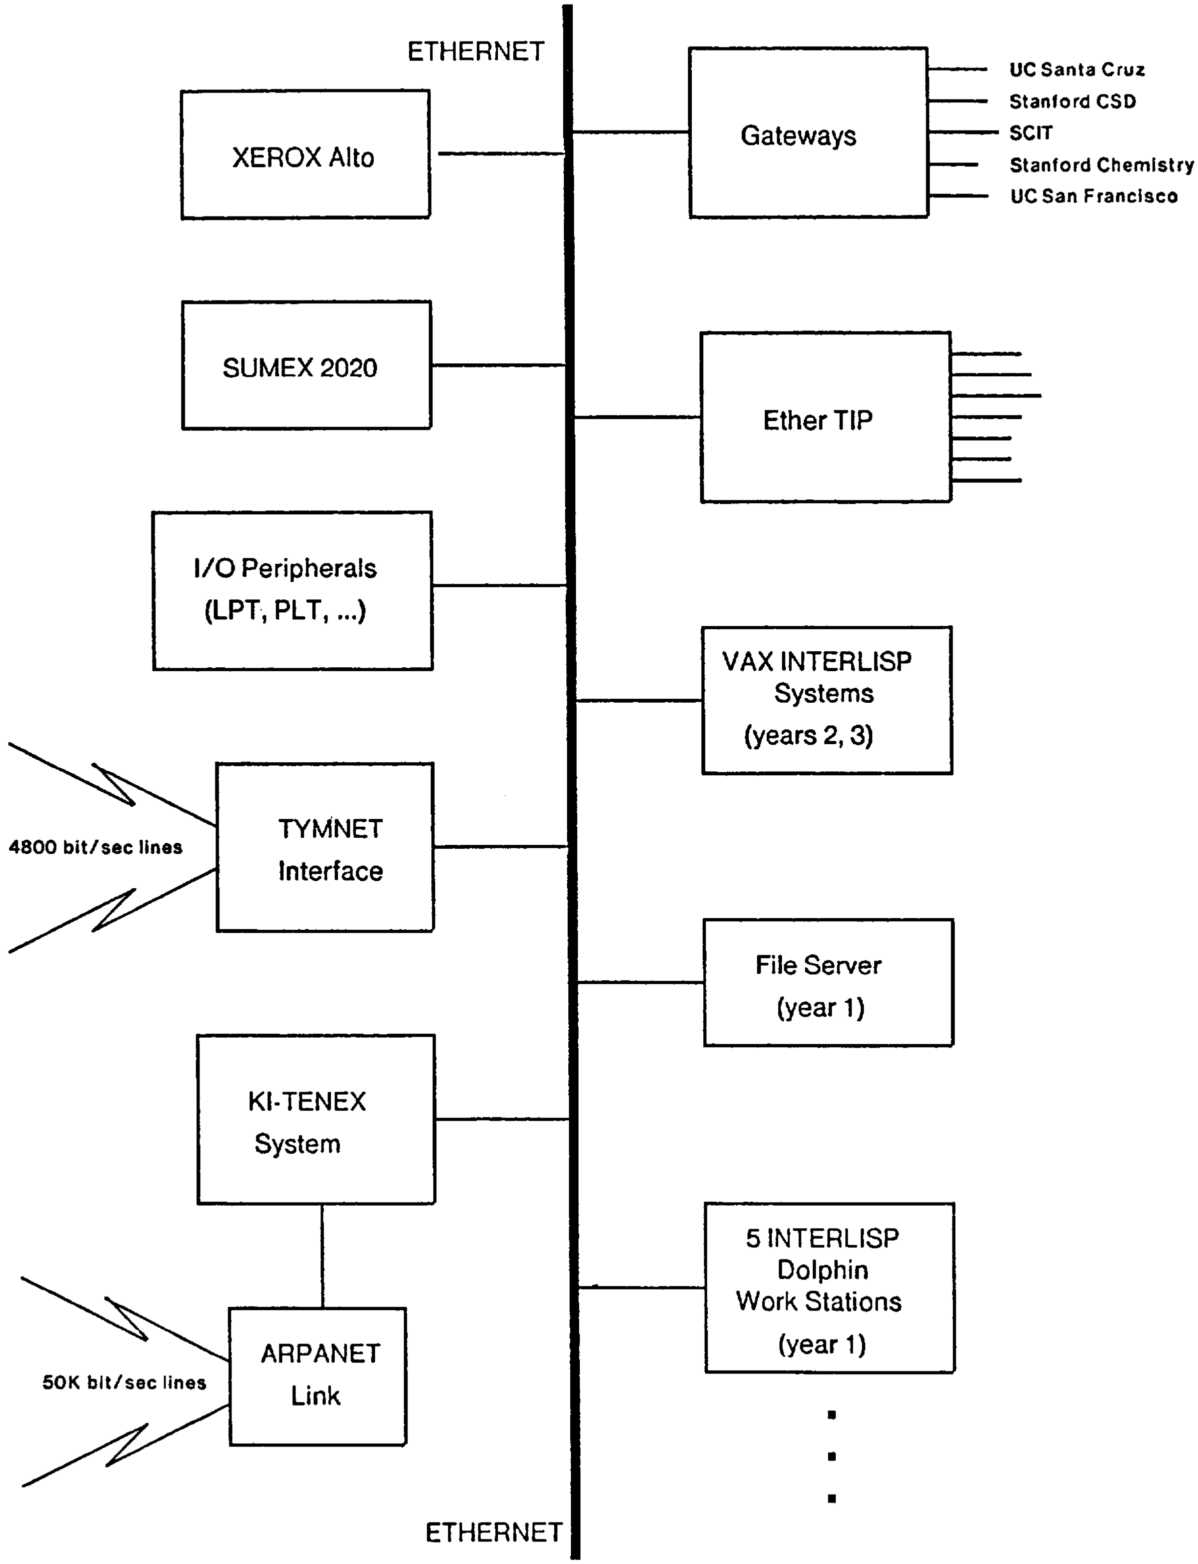

Figure 14. Planned Ethernet System to Integrate System Hardware

Communication Networks

Networks have been centrally important to the research goals of SUMEX-AIM and will become more so in the context of increasingly distributed computing. Communication will be crucial to maintain community scientific contacts, to facilitate shared system and software maintenance based on regional expertise, to allow necessary information flow and access at all levels, and to meet the technical requirements of shared equipment.

Long-Distance Connections

We have had reasonable success at meeting the geographical needs of the community during the early phases of SUMEX-AIM through our ARPANET and TYMNET connections. These have allowed users from many locations within the United States and abroad to gain terminal access to the AIM resources (SUMEX, Rutgers, and SCORE) and through ARPANET links to communicate much more voluminous file information. Since many of our users do not have ARPANET access privileges for technical or administrative reasons, a key problem impeding remote use has been the limited communications facilities (speed, file transfer, and terminal handling) offered currently by commercial networks. Commercial improvements are slow in coming but may be expected to solve the file transfer problem in the next few years. A number of vendors (AT&T, IBM, Xerox, etc.) have yet to announce commercially available facilities but TELENET is actively working in this direction. We plan to continue experimenting with improved facilities as offered by commercial or government sources in the next grant term. We have budgeted for continued TYMNET service and an additional amount annually for experimental network connections.

High-speed interactive terminal support will continue to be a problem since one cannot expect to serve 1200-9600 baud terminals effectively over shared long-distance trunk lines with gross capacities of only 9600-19200 baud. We feel this is a problem that is best solved by distributed machines able to effectively support terminal interactions locally and coupled to other AIM machines and facilities through network or telephonic links. As new machine resources are introduced into the community, we will allocate budgeted funds with Executive Committee advice to assure effective communication links.

Local Intermachine Connections

A key feature of our plans for future computing facilities is the support of a heterogeneous processing environment that takes advantage of newly available technology and shared equipment resources between these machines. The "glue" that links these systems together is a high speed local network. We have chosen Ethernet and the Xerox PUP [9, 12] protocols for these interconnections. This choice was based on the availability of that technology now and the economics of using already developed TENEX and other server software. We expect the Ethernet system to continue to meet our technical needs for the coming grant term and we plan to continue to use it. We are working closely with other groups here at Stanford and elsewhere to share hardware interface and software designs wherever possible.

Our goals are to complete integration of the SUMEX-AIM system, including making selected KI-10 peripherals available as Ethernet nodes, creating links to nearby campus resources, and establishing needed remote links to other groups not on the ARPANET such as Wipke at the University of California at Santa Cruz. A diagram of our Ethernet system is shown in Figure 14 on page 55 and includes the following major elements:

1) KI-10 direct memory access interface. We currently have an inefficient I/O bus connection.

2) Stanford campus gateway. Establish links to other Ethernets on campus to allow access to special resources (Dover printer, plotters, typesetting equipment, etc.) and to allow users to easily access various computing resources.

3) Ethertip. We need additional terminal ports into the system and the Ethernet provides a natural mechanism to do this supporting high speed terminals and connections to various resources (KI-10, 2020, VAX's, etc.).

4) TYMNET connection. This connection currently comes through the KI-10's and will be moved to a separate Ethernet node. This will free the KI-10's from handling the special TYMNET protocol and will allow TYMNET users to access any of the SUMEX-AIM resources. Similar facilities for the ARPANET may also be implemented depending on administrative constraints.

5) Printer/plotter service. We plan to make these local resources accessible from any of the SUMEX-AIM machines instead of being centered on the KI-10's. This will also free up the KI-10's from routine spooler tasks.

7) Connections for other machines (VAX's, Professional Workstations, file server, etc.)

Resource Software

We will continue to maintain the existing system, language, and utility support software on our systems at the most current release levels, including up-to-date documentation. We will also be extending the facilities available to users where appropriate, drawing upon other community developments where possible. We rely heavily on the needs of the user community to direct system software development efforts. Specific development areas for existing systems include:

1) completion of the Ethernet connections and necessary host software. This will include basic packet handling, PUP protocols at all levels, and relocation of shared existing resources to become Ethernet nodes.

2) bug fixes in the current monitors. We still have a number of bugs that cause infrequent crashes and that are hard to isolate because they cause system problems long after the fact. We will continue to work to repair these problems as time permits.

3) continued evaluation of system efficiency to improve performance.

4) compatibility issues. Our current compatibility package for TOPS-20 requires additional work to extend its features. We will also keep it up-to-date as DEC make new changes to their system.

5) continued work to create similar working and programming environments between our TENEX and TOPS-20 systems. This will include moving TENEX features like the SUMEX GTJFN enhancements and scheduling controls as needed to TOPS-20 and vice versa

6) continued work to improve system information and help facilities for users.

Our plans for augmenting the SUMEX-AIM resources will entail substantial new system and subsystem programming. Our goals will be to derive as much software as possible from the user communities of the new VAX and Professional Workstation machines but we expect to have to do considerable work to adapt them to our biomedical AI needs. Many features of these systems are designed for a computer science environment and lack some of the human engineering and "friendliness" capabilities we have found needed to allow non-computer scientists to effectively use them. We are beginning to experiment with physician needs for interfaces to our AI programs to be better able to adapt the new machines as professional aids. Also many of the utility tools that we take for granted in the well-developed TENEX and TOPS-20 environment (communications, text manipulation, file management, accounting, etc.) will have to be reproduced. We expect to set up many of the common information services as network nodes.

Within the AIM community we expect to serve as a center for software sharing between various distributed computing nodes. This will include contributing locally developed programs, distributing those derived from elsewhere in the community, maintaining up-to-date information on subsystems available, and assisting in software maintenance.

Community Management

We plan to retain the current management structure that has worked so well. We will continue to work closely with the management committees to recruit the additional high quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We expect the Executive and Advisory Committees to play an increasingly important role in advising on priorities for facility evolution and on-going community development planning in addition to their recruitment efforts. The composition of the Executive committee will grow as needed to assure representation of major user groups and medical and computer science applications areas. The Advisory Group membership rotates regularly and spans both medical and computer science research expertise. We expect to maintain this policy.

We will continue to make information available about the various projects both inside and outside of the community and thereby promote the kinds of exchanges exemplified earlier and made possible by network facilities.

The AIM workshops under the Rutgers resource have served a valuable function in bringing community members and prospective users together. We will continue to support this effort. This summer the AIM workshop will be held in Vancouver, British Columbia in conjunction with the International Joint Conference on Artificial Intelligence. We are actively helping to organize the meeting. We will continue to assist community participation and provide a computing base for workshop demonstrations and communications. We will also assist individual projects in organizing more specialized workshops as we have done for the DENDRAL and AGE projects.

We plan to continue indefinitely our present policy of non-monetary allocation control. We recognize, of course, that this accentuates our responsibility for the careful selection of projects with high scientific and community merit.

Training and Education Plans

We have an on-going commitment, within the constraints of our staff size, to provide effective user assistance, to maintain high quality documentation of the evolving software support on the SUMEX-AIM system, and to provide software help facilities such as the HELP and Bulletin Board systems. These latter aids are an effective way to assist resource users in staying informed about system and community developments and solving access problems. We plan to take an active role in encouraging the development and dissemination of community databases such as the AI Handbook, up-to-date bibliographic sources, and developing knowledge bases. Since much of our community is geographically remote from our machine, these on-line aids are indispensable for self help. We will continue to provide on-line personal assistance to users within the capacity of available staff through the SNDMSG and LINK facilities.

We budget funds to continue the "collaborative linkage" support initiated during the first term of the SUMEX-AIM grant. These funds are allocated under Executive Committee authorization for terminal and communications support to help get new users and pilot projects started.

Finally, we will continue to actively support the AIM workshop series in terms of planning assistance, participation in program presentations and discussions, and providing a computing base for AI program demonstrations and experimentation.

Core Research Plans

SUMEX core research includes both basic AI research and development of community tools useful for building expert systems. Expert systems are symbolic problem solving programs capable of expert-level performance, in which domain-specific knowledge is represented and used in an understandable line of reasoning. The programs can be used as problem solving assistants or tutors, but also serve as excellent vehicles for research on representation and control of diverse forms of knowledge. MYCIN is one of the best examples.

Because the main issues of building expert systems are coincident with general issues in AI, we appreciate the difficulty of proposing to "solve" basic problems. However, we do propose to build working programs that demonstrate the feasibility of our ideas within well defined limits. By investigating the nature of expert reasoning within computer programs, the process is "demystified". Ultimately, the construction of such programs becomes itself a well-understood technical craft.

The foundation of all of our core research work is expert knowledge: its acquisition from practitioners, its accommodation into the existing knowledge bases, its explanation, and its use to solve problems. Continued work on these topics provides new techniques and mechanisms for the design and construction of knowledge-based programs; experience gained from the actual construction of these systems then feeds back both (a) evaluative information on the ideas' utility and (b) reports of quite specific problems and the ways in which they have been overcome, which may suggest some more general method to be tried in other programs.

One of our long-range goals is to isolate AI techniques that are general, to determine the conditions for their use and to build up a knowledge base about AI techniques themselves. SUMEX resources are coordinated for this purpose with the multidisciplinary efforts of the Stanford Heuristic Programming Project (HPP). Under support from ARPA, NIH/NLM, ONR, NSF, and private funding, the HPP conducts research on five key scientific problem areas, as well as a host of subsidiary issues [1]:

1) Knowledge Representation - How shall the knowledge necessary for expert-level performance be represented for computer use? How can one achieve flexibility in adding and changing knowledge in the continuous development of a knowledge base? Are there uniform representations for the diverse kinds of specialized knowledge needed in all domains?

2) Knowledge Utilization - What designs are available for the inference procedure to be used by an expert system? How can the control structure be simple enough to be understandable and yet sophisticated enough for high performance? How can strategy knowledge be used effectively?

3) Knowledge Acquisition - How can the model of expertise in a field of work be systematically acquired for computer use? If it is true that the power of an expert system is primarily a function of the

quality and completeness of the knowledge base, then this is the critical "bottleneck" problem of expert systems research.

4) Explanation - How can the knowledge base and the line of reasoning used in solving a particular problem be explained to users? What constitutes an acceptable explanation for each class of users?

5) Tool Construction - What kinds of software packages can be constructed that will facilitate the implementation of expert systems, not only by the research community but also by various user communities?

Artificial Intelligence is largely an empirical science. We explore questions such as these by designing and building programs that incorporate plausible answers. Then we try to determine the strengths and weaknesses of the answers by experimenting with perturbations of the systems and extrapolations of them into new problem areas. The test of success in this endeavor is whether the next generation of system builders finds the questions relevant and the answers applicable to reduce the effort of building complex reasoning programs.

I.B Highlights

I.B.1 Handbook of Artificial Intelligence

The AI Handbook is a compendium of knowledge about the field of Artificial Intelligence being assembled under Professor Edward Feigenbaum and Messrs Avron Barr and Paul Cohen. It is being compiled by students and investigators at several research facilities across the nation. The AI Handbook Project is a good example of community collaboration using the SUMEX-AIM communication facilities to prepare, review, and disseminate this reference work on AI techniques. The Handbook articles exist as computer files at the SUMEX facility. All of our authors and reviewers have access to these files via the network facilities and use the document-editing and formatting programs available at SUMEX. This relatively small investment of resources has resulted in what we feel will be a seminal publication in the field of AI, of particular value to researchers, like those in the AIM community, who want quick access to AI ideas and techniques for application in other areas.

The AI Handbook Project was undertaken as a core activity by SUMEX in the spirit of community building that is the fundamental concern of the facility. We feel that the organization and propagation of this kind of information to the AIM community, as well as to other fields where AI is being applied, is a valuable service that we are uniquely qualified to support.

The scope of the work is broad: Two hundred articles cover all of the important ideas, techniques, and systems developed during 20 years of research in AI. Each article, roughly four pages long, is a description written for non-AI specialists and students of AI. Additional articles serve as overviews, which discuss the various approaches within a subfield, the issues, and the problems.

We expect the Handbook to reach a size of approximately 1000 pages. Roughly two-thirds of this material will constitute Volumes I and II of the Handbook. The material in Volumes I and II will cover AI research in Heuristic Search, Representation of Knowledge, AI Programming Languages, Natural Language Understanding, Speech Understanding, Automatic Programming, and Applications-oriented AI Research in Science, Mathematics, Medicine, and Education. Researchers at Stanford University, Rutgers University, SRI International, Xerox PARC, RAND Corporation, MIT, USC-ISI, Yale, and Carnegie-Mellon University have contributed material to the project. The current schedule for publication of the several volumes is as follows. It should be noted that Volume I has been selected by the Library of Computer Science as their August 1981 book club selection.

May, 1981: Publication of Volume 1 by publisher (Wm. Kaufmann Inc., Los Altos, Ca.)

August, 1981: Submission of final copy to publisher for Volume II (publication by end of 1981).

August-September, 1981: Completion of Technical Reports containing chapters of Handbook

October, 1981: Submission of final copy of Volume III to publisher (for publication first quarter 1982)

I.B.2 Tutorial on AI in Clinical Medicine

In conjunction with the 1980 AIM Workshop, a continuing education tutorial designed for physicians was held at Stanford on August 17-18, 1980. The tutorial was entitled "Computers in Medicine -- Applications of Artificial Intelligence Techniques" and was organized by Drs. W. Clancey and E. Shortliffe. The tutorial was well-attended by 135 physicians, 18 students, 10 members of the press, and several non-physician researchers. It was accredited for postgraduate medical education through Stanford University School of Medicine. Enrollees came from as far away as Mexico and the East Coast.

The course included an optional introduction to computers for those who had no prior experience with the technology, an overview of SUMEX-AIM research, and an introduction to background materials regarding decision theory and data base applications in medicine. Speakers also provided detailed presentations on MYCIN, CASNET/EXPERT, INTERNIST and GUIDON. The course closed with a panel discussion on the problems and promise of AI in Medicine. A syllabus was distributed including a comprehensive survey of medical AI research and is comprised of recent articles written by the tutorial faculty, mostly for a clinical audience. The faculty consisted of 15 distinguished researchers from the national AI community, including 7 physicians and 9 speakers from centers other than Stanford. Coordination and planning for the tutorial was was facilitated by sending electronic messages; almost all speakers regularly use SUMEX or another ARPANET machine.

The course was exceedingly well received. Attendees were fascinated by the content, generally felt it was well presented, and indicated they would recommend the course to others if it were made available again. Many physicians requested a follow-up course that would introduce them to more technical detail than had been possible in the introductory tutorial.

To evaluate the impact of the tutorial on the participants, and to assess baseline opinions regarding the field, we undertook a survey of the physicians' knowledge about computers as well as their attitudes towards medical consultation systems. The statistical analysis of these questionnaires has now been completed, and a paper summarizing the results submitted for publication (*). In brief, the survey showed that physicians were willing to accept the possibility of computer-based clinical decision aids but placed severe demands on the capabilities of such systems if they were to be acceptable for routine use.

(*) Teach, R.L. and Shortliffe, E.H. "An Analysis of Physician Attitudes Regarding Computer-based Clinical Consultation Systems." Submitted for publication, March 1981.

I.B.3 GENET - Dissemination of AI Tools for Molecular Genetics

The MOLGEN project at Stanford has focused on applications of artificial intelligence and symbolic computation in the field of molecular biology. The research began in 1975 and by early 1980, through many collaborative contacts, it was realized that some of the systems developed by MOLGEN were already of direct utility to many scientists in the domain.

In order to broaden MOLGEN's base of scientist collaborators to molecular biologists at institutions other than Stanford and to experiment with the use of a SUMEX-like resource to disseminate sophisticated AI software tools to a generally computer-naive community, we initiated an experimental user group called GENET. The response to our very limited announcement of this facility has been most enthusiastic.

We have offered three main programs to assist molecular genetics users: SEQ, a DNA-RNA sequence analysis program; MAP, a program that assists in the construction of restriction maps from restriction enzyme digest data; and MAPPER (written and maintained by William Pearson from Johns Hopkins University), a simplified version of the MOLGEN MAP program that is somewhat more efficient than the MOLGEN version. Some of the other more sophisticated programs being developed by MOLGEN research efforts have not been offered because they are not ready for novice users. In addition, the GENET users have had access to the SUMEX-AIM programs for electronic messaging, text editing, file searching, etc.

The GENET community, begun in spring of 1980, started to grow exponentially until they were consuming SUMEX resources on a scale equal to the largest AI research project. We were obliged to place restrictions on the number of simultaneous GENET users and to otherwise limit the growth of the community. Even with these restrictions, the community currently consists of approximately 200 users from 63 research institutions. Of these 200 users, approximately 35 are consistently active users. That is, they log in, run programs, and interact with the MOLGEN members on an almost daily basis. Many of these users have made valuable contributions to our work. About 100 others are frequent, but not regular users. They log in only when they have a major analysis task to perform, which seems to be on the order of once a month.

The remaining users rarely use the system. They have logged in a few times, but for one reason or another they never become regular users of the system. Quite often this is because a lab group will settle on having one or two graduate students or post-doctoral associates become the "computer experts" of the group, and as a result, the computer use by the other people in the lab drops to a very low level. An equally prevalent reason for users to stop using the GENET account is a lack of SUMEX resources. The major complaint that we get from GENET users concerns the lack of compute time and availability of the system. One account just is not enough for 200 people to share.

We have succeeded in the goals set out for GENET. Many of our GENET guests have become active collaborators in core MOLGEN research. We are also pleased by the numerous comments SUMEX has received from GENET users praising the user-sensitive nature of the resource, especially in comparison to typical university computer centers. It is clear we have only had the resources to whet the appetite of this large, active, international community.

I.B.4 AGE - A Tool for Knowledge-Based System Development

One of the most difficult, time-consuming, and expensive aspects of building knowledge-based systems (indeed any kind of software system) is the human effort involved in designing and coding them from the ground up. A major goal of SUMEX core research has been to demystify and make explicit the art of knowledge engineering. More concretely, we have attempted to isolate inference, control, and representation techniques from previous knowledge-based programs; reprogram them for domain independence; write an interface that will help a user understand what the package offers and how to use the modules; and make the package available to other members of the AIM community and the general scientific community to assist in knowledge-based program development. The AGE (Attempt to Generalize) package, developed by H. P. Nii and E. Feigenbaum is one of the earliest experimental examples of such a system and has reached a level of practical utility.

The design and implementation of the AGE program is based primarily on the experience gained with knowledge-based programs at the Stanford Heuristic Programming Project in the last decade. The programs that have been, or are being, built include: DENDRAL, meta-DENDRAL, MYCIN, HASP, AM, MOLGEN, CRYSALIS [Feigenbaum 1977], and SACON [Bennett 1978]. Initially, the AGE program embodies the AI techniques used in these programs but longer range goals are to integrate those developed at other AI laboratories as well. It is hoped that AGE will speed up the process of building knowledge-based programs and facilitate the dissemination of AI techniques by (1) packaging common AI software tools so that they need not be reprogrammed for every problem; and (2) helping people who are not knowledge engineering specialists write knowledge-based programs.

AGE is being developed along two separate fronts: the "kit" of tools for implementing knowledge-based systems and the "intelligent" interface to assist users make use of them. The current AGE system provides a set of preprogrammed "components" or "building blocks". A "component" is a collection of functions and variables that support conceptual entities in program form. For example, the production rule component, consists of a rule interpreter and various strategies for rule selection and execution.

The components in AGE have been carefully selected and modularly programmed to be useable in combinations. For those users not familiar enough to experiment on their own, AGE provides two predefined configurations of components--each configuration is called a "framework". One framework, called the Blackboard framework, is for building programs that use a globally accessible data structure called a "blackboard" [Lesser 77], and independent sources of knowledge which cooperate to form hypotheses. The other framework, called the Backchain framework, is for building programs that use backward-chained production rules as their primary mechanism of generating inferences (e.g., MYCIN).

Currently AGE-1 is available on a limited basis on the SUMEX-AIM resource and on the Stanford SCORE 2060 computer in the Computer Science Department. We held a three-day workshop in March 1980 to familiarize invitees with the use of AGE and to allow each participant to implement a running program related to his application area. For the 1980 AIM Workshop we reimplemented a major portion of the VM program using AGE. In addition to demonstrating a variety of features of AGE, we were able to illustrate the relatively short implementation time required once the goals of the application and the necessary knowledge were delineated -- a first-year graduate student had the program running in three weeks. We are still working to broaden the user community for AGE and to learn from their experiences what directions our future research efforts should take.

I.B.5 ONCOCIN- An Oncology Chemotherapy Advisor

Work on the oncology chemotherapy consultation system, named ONCOCIN, was begun in July 1979. It is one of the newest application areas being investigated in the Stanford SUMEX-AIM community and is designed to be an interactive system for assigning and managing patients on chemotherapy protocols. This spring, it was installed for initial experimental use by faculty and fellows in the Debbie Probst Oncology Day Care Center at Stanford University Medical Center. Overall goals for ONCOCIN are (1) to demonstrate that a rule-based consultation system with explanation capabilities can be usefully applied and accepted in a busy clinical environment; (2) to improve the tools available for building knowledge-based expert systems for medical consultation; and (3) to establish an effective scientific relationship with a group of physicians that will facilitate future research and implementation of knowledge-based tools for clinical decision making.

In addition to ONCOCIN's basic AI research goals, it is directed toward the development of a clinically useful oncology consultation tool that will: (1) assist with the identification of protocols that may apply to a given patient and to help determine the patient's eligibility for a given protocol; (2) provide detailed information on protocols in response to questions from clinic personnel; (3) assist with chemotherapy dose selection and attenuation for a given patient; (5) provide reminders, at appropriate intervals, of follow-up tests and films required by the protocol in which a given patient is enrolled; and (6) reason about managing current patients in light of stored data from previous visits of the individual patients or aggregate data about groups of "similar" patients.

We are pursuing a five-year plan for accomplishing these goals. We spent the first year working out a prototype ONCOCIN system, drawing from programs and capabilities developed for the EMYCIN system-building project. We also undertook a detailed analysis of the day-to-day activities of the Stanford oncology clinic in order to determine how to introduce ONCOCIN with minimal disruption of an operation which is already running smoothly. Much of this early effort was spent giving careful consideration to the most appropriate mode of interaction with physicians in order to optimize the chances for ONCOCIN to become a useful and accepted tool in this specialized clinical environment.

More recently we have detailed the design and have implemented an actual experimental system. This system is based on multiple processes that manage the physician interface, reasoning and problem-solving, and patient database management. All of the system work has been completed to allow installation of ONCOCIN in the clinic. Following the initial prototype development based on lymphoma protocols, we checked to verify that the representation method we are using will be adequate for arbitrary protocol knowledge that may be encountered in the future. So we decided to encode and briefly test the knowledge of a non-lymphoma protocol. We chose the complicated protocol for oat cell (small cell) carcinoma of the lung because it involves a large number of possible therapies and complex interweaving of chemotherapy and radiotherapy. After approximately one month's effort, the oat cell protocol was encoded and run successfully on a number of test cases. In addition, the lymphoma protocol specifications used in the clinic were changed and we spent a few weeks entering the necessary corrections. In all cases the ONCOCIN representation scheme was adequate to accommodate the protocol knowledge with only minor changes, and we are confident that the system will be able to adapt to other protocols that need to be encoded in the coming years.

ONCOCIN has been extensively debugged through runs on several hundred sample patient cases with the results reviewed in detail by the collaborating oncologists. We have just begun to offer the ONCOCIN system for use by the oncology faculty and fellows in the morning chemotherapy clinics in which most of the lymphoma patients receive their treatment. We have taken care in introducing ONCOCIN to provide needed baseline information so we can formally evaluate its impact and effectiveness in the oncology clinic.

I.C Administrative Changes

The SUMEX-AIM resource has undergone several administrative changes this past year that serve to enhance its position within the Stanford Medical School as a resource for AI research:

1) Professor Edward Shortliffe was appointed as co-Principal Investigator of SUMEX-AIM. Professor Shortliffe has been central in the development of the MYCIN group of projects and has long worked closely with Professor Feigenbaum in planning the future development of SUMEX. This appointment takes formal recognition of this role for Professor Shortliffe and strengthens SUMEX-AIM through his close scientific and administrative ties to the Stanford medical community.

2) In parallel with Professor Shortliffe's appointment as co-Principal Investigator, SUMEX moved administratively from the Department of Genetics to the Department of Medicine. It is now administered jointly between the Departments of Medicine and Computer Science. As part of the largest clinical medicine department at Stanford, SUMEX now has increased visibility and opportunity to broaden its local scientific collaborations.

3) Professor Elliott Levinthal began a two-year leave of absence to take a position as head of the Defense Sciences Office at DARPA. Professor Roy Maffly has replaced him as AIM liaison in charge of coordinating the reviews of new project applications and serving as the interface to collaborative projects.

I.D Resource Management and Allocation

The mission of SUMEX-AIM, locally and nationally, entails both the recruitment of appropriate research projects interested in medical AI applications and the catalysis of interactions among these groups and the broader medical community. User projects are separately funded and autonomous in their management. They are selected for access to SUMEX on the basis of their scientific and medical merits as well as their commitment to the community goals of SUMEX. Currently active projects span a broad range of application areas such as clinical diagnostic consultation, molecular biochemistry, belief systems modeling, mental function modeling, and instrument data interpretation (descriptions of the individual collaborative projects are in Section II beginning on page 89).

I.D.1 Management Committees

Since the SUMEX-AIM project is a multilateral undertaking by its very nature, we have created several management committees to assist in administering the various portions of the SUMEX resource. As defined in the SUMEX-AIM management plan adopted at the time the initial resource grant was awarded, the available facility capacity is allocated 40% to Stanford Medical School projects, 40% to national projects, and 20% to common system development and related functions. Within the Stanford aliquot, Prof. Feigenbaum and BRP have established an advisory committee to assist in selecting and allocating resources among projects appropriate to the SUMEX mission. The current membership of this committee is listed in Appendix C.

For the national community, two committees serve complementary functions. An Executive Committee oversees the operations of the resource as related to national users and makes the final decisions on authorizing admission for new projects and revalidating continued access for existing projects. It also establishes policies for resource allocation and approves plans for resource development and augmentation within the national portion of SUMEX (e.g., hardware upgrades, significant new development projects, etc.). The Executive Committee oversees the planning and implementation of the AIM Workshop series currently implemented under Prof. S. Amarel of Rutgers University and assures coordination with other AIM activities as well. The committee will play a key role in assessing the possible need for additional future AIM community computing resources and in deciding the optimal placement and management of such facilities. The current membership of the Executive committee is listed in Appendix C.

Reporting to the Executive Committee, an Advisory Group represents the interests of medical and computer science research relevant to AIM goals. The Advisory Group serves several functions in advising the Executive Committee; 1) recruiting appropriate medical/computer science projects, 2) reviewing and recommending priorities for allocation of resource capacity to specific projects based on scientific quality and medical relevance, and 3) recommending policies and development goals for the resource. The current Advisory Group membership is given in Appendix C.

These committees have actively functioned in support of the resource. Except for the meetings held during the AIM workshops, the committees have "met" by messages, net-mail, and telephone conference owing to the size of the groups and to save the time and expense of personal travel to meet face to face. The telephone meetings, in conjunction with terminal access to related text materials, have served quite well in accomplishing the agenda business and facilitate greatly the arrangement of meetings. Other solicitations of advice requiring review of sizable written proposals are done by mail.

We will continue to work with the management committees to recruit the additional high quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We will continue to make information available about the various projects both inside and outside of the community and thereby promote the kinds of exchanges exemplified earlier and made possible by network facilities.

I.D.2 New Project Recruiting

The SUMEX-AIM resource has been announced through a variety of media as well as by correspondence, contacts of NIH-BRP with a variety of prospective grantees who use computers, and contacts by our own staff and committee members. The number of formal projects that have been admitted to SUMEX has more than trebled since the start of the project to a current total of 8 national AIM projects and 8 Stanford projects. Others are working tentatively as pilot projects or are under review.

We have prepared a variety of materials for the new user ranging from general information such as is contained in a SUMEX-AIM overview brochure to more detailed information and guidelines for determining whether a user project is appropriate for the SUMEX-AIM resource. A questionnaire is available to assist users seriously considering applying for access to SUMEX-AIM. Pilot project categories have been established both within the Stanford and national aliquots of the facility capacity to assist and encourage new projects in formulating possible AIM proposals and pending their application for funding support. Pilot projects are approved for access for limited periods of time after preliminary review by the Stanford or AIM Advisory Group as appropriate to the origin of the project.

These contacts have sometimes done much more than provide support for already formulated programs. For example, Prof. Feigenbaum's group at Stanford previously initiated a major collaborative effort with Dr. Osborn's group at the Institutes of Medical Sciences in San Francisco. This project in "Pulmonary Function Monitoring and Ventilator Management - PUFF/VM" (see Section II.A.2.4 on page 201) originated as a pilot request to use MLAB in a small way for modeling. Subsequently the AI potentialities of this domain were recognized by Feigenbaum, Nii, and Osborn and a joint proposal was submitted to and funded by NIH. This past summer John Kunz from Dr. Osborn's laboratory spent approximately half time at Stanford to learn more about AI research and to participate more closely in the development of the PUFF/VM program.

Similarly, Prof. Feigenbaum and Ms. Nii recently spent two days with Profs. Kintsch and Polson at the University of Colorado, introducing them to the newly developed AGE package for use in formulating their program on modeling aspects of human cognition.

A list of the fully authorized projects currently comprising the SUMEX-AIM community can be found with brief abstracts in Appendix A on page 278. More detailed descriptions of collaborative project activities can be found in Section II.

As an additional aid to new projects or collaborators with existing projects, we provide a limited amount of funds for use to support terminals and communications needs of users without access to such equipment. We are currently providing support for 6 terminals and 4 modems for users as well as a leased line between Stanford and the University of California at Santa Cruz for the Chemical Synthesis project.

I.D.3 Stanford Community Building

The Stanford community has undertaken several internal efforts to encourage interactions and sharing between the projects centered here. Professor Feigenbaum organized a project with the goal of assembling a handbook of AI concepts, techniques, and current state-of-the-art. This project has had enthusiastic support from the students and substantial progress made in preparing many sections of the handbook (see Section II.A.1.2 on page 99 for more details).

Weekly informal lunch meetings (SIGLUNCH) are also held between community members to discuss general AI topics, concerns and progress of individual projects, or system problems as appropriate. In addition, presentations from a substantial number of outside speakers are invited.

I.D.4 Existing Project Reviews

We have conducted a continuing careful review of on-going SUMEX-AIM projects to maintain a high scientific quality and relevance to our medical AI goals and to maximize the resources available for newly developing applications projects. At meetings of the AIM Advisory Group and Executive Committee this past year, all the national AIM projects were reviewed.

These groups recommended continued access for most formal projects on the system. However, they recommended that the Higher Mental Functions project could better meet it current goals through computer support at UCLA and we have therefore reduced this project to "associate" status.

I.D.5 Resource Allocation Policies

As the SUMEX facility has become increasingly loaded, a number of diverse and conflicting demands have arisen which require controlled allocation of critical facility resources (file space and central processor time). We have already spelled out a policy for file space management; an allocation of file storage is defined for each authorized project in conjunction with the management committees. This allocation is divided among project members in any way desired by the individual principal investigators. System allocation enforcement is implemented by project each week. As the weekly file dump is done, if the aggregate space in use by a project is over its allocation, files are archived from user directories over allocation until the project is within its allocation.

We have implemented effective system scheduling controls to attempt to maintain the 40:40:20 balance in terms of CPU utilization and to avoid system and user inefficiencies during overload conditions. The initial complement of user projects justifying the SUMEX resource was centered to a large extent at Stanford. Over the past five years of the SUMEX grant, a substantial growth in the number of national projects was realized. During the same time the Stanford group of projects has matured as well and in practice the 40:40 split between Stanford and non-Stanford projects is not ideally realized although the demand from the national community has increased substantially (see Figure 9 on page 33 and the tables of recent project usage on page 36).

Our job scheduling controls bias the allocation of CPU time based on percent time consumed relative to the time allocated over the 40:40:20 community split. The controls are "soft" however in that they do not waste computer cycles if users below their allocated percentages are not on the system to consume the cycles. The operating disparity in CPU use to date reflects a substantial difference in demand between the Stanford community and the developing national projects, rather than inequity of access. For example, the Stanford utilization is spread over a large part of the 24-hour cycle, while national-AIM users tend to be more sensitive to local prime-time constraints. (The 3-hour time zone phase shift across the continent is of substantial help in load balancing.) During peak times under the overload control system reported previously, the Stanford community still experiences mutual contentions and delays while the AIM group has relatively open access to the system. We did enable overload controls for the national community this past year, however, because of their substantial increase in demand. For the present, we propose to continue our policy of "soft" allocation enforcement for the fair split of resource capacity.

Our system also categorizes users in terms of access privileges. These comprise fully authorized users, pilot projects, guests, and network visitors in descending order of system capabilities. We want to encourage bona fide medical and health research people to experiment with the various programs available with a minimum of red tape while not allowing unauthenticated users to bypass the advisory group screening procedures by coming on as guests. So far we have had relatively little abuse compared to what other network sites have experienced, perhaps on account of the personal attention that senior staff gives to the logon records, and to other security measures. However, the experience of most other computer managers behooves us to be cautious about being as wide open as might be preferred for informal service to pilot efforts and demonstrations. We will continue developing this mechanism in conjunction with management committee policy decisions.

We have also encouraged mature projects to apply for their own machine resources in order to preserve the SUMEX-AIM resource for research and development efforts and to support projects unable to justify their own machines. The INTERNIST project has received approval for a VAX machine to support their planned development and program testing work. Also Profs. Lesgold and Greeno's "Simulation of Cognitive Processes" project has moved the bulk of their work to their own local VAX.

I.E Dissemination Efforts

Throughout its existence, SUMEX-AIM has devoted substantial efforts toward disseminating information about its activities as a resource and about the work of individual collaborative projects. We continue to make many presentations at professional meetings, to provide services to demonstrate developed AI programs for interested groups and individuals, and to work in organizing workshops within the SUMEX-AIM community to introduce our work to collaborating professional communities. We have also spent substantial efforts in the past working with the Research Resources Information Center to produce the "Seeds of Artificial Intelligence" monograph to address a broader community of technical and lay people.

The following sections summarize some of the activities undertaken this past year:

I.E.1 Sixth AIM Workshop

The Sixth Annual AIM (Artificial Intelligence in Medicine) Workshop was held at Stanford University on 13-16 August 1980. The program chairman was Dr. E. Shortliffe, the chairman for demo-based sessions was Dr. L. Fagan, and the short report chairman was Dr. R. Blum. This was the first Workshop to be held in California, and was be held in conjunction with the first annual meeting of the AAAI Society (American Association for Artificial Intelligence).

Among the goals of this year's conference was the development of a format for scientific exchange that would help clarify the technical details of the programs that are under development throughout the AIM community. Many individuals have observed that it can be difficult at meetings such as this to obtain detailed understanding of one another's work. Formal presentations with slides and a description of data structures typically are divorced from a sense of the program's operation as seen to the user. As a result, many of us have had to complement our annual Workshop participation with visits to other sites so that we can learn about others' work in depth. In 1980 we experimented with a format that tried to simulate the kind of detailed interactions that have previously occurred only in individual sessions after hours or at times other than the Workshop.

Demo-Based Sessions

This year the major portion of the conference was devoted to detailed discussions of AIM systems through the vehicle of specially prepared demonstrations. Each of the established AIM systems was represented with a two hour presentation. Each speaker had a display terminal, special projection system and high-speed connections to the SUMEX 2020 computer. Rather then rely on an impromptu live demonstration, each project was asked to prepare a typescript of an interactive session, subject to the following guidelines:

(1) The typescript was to represent the interaction exactly as it occurred on the screen to the user (i.e., the presenters were asked not to delete mistakes, problems, garbage collect messages, etc.), and was to be augmented only with the following:

(a) annotations to clarify specific points.

(b) "break" interruptions as described below.

(2) The typescript was to be presented in short segments with enough discussion to identify the current point in the program's reasoning process.

(3) At pertinent points the researchers were asked to break into the program's operation during typescript preparation and display pertinent data structures to illustrate the system's internal representation and organization.

A computer program was written by the SUMEX staff to facilitate the display of annotated and formatted typescripts. The input to the program is a typescript file that has special control characters inserted into text to mark off pages of information. Other control characters are used to highlight (brighten) important points in the typescript, to turn pages or move to a specific page, and to provide for different levels of detail. The provision for different levels of detail was designed to show selectively information in response to questions, or to adjust presentations for different audiences (e.g., physicians vs. computer scientists). Because the program's output is treated as a text file by the system, slide-line material or diagrams can be inserted into the running transcript. A more complete description of the program is available on-line on the SUMEX computer.

No detailed evaluation of the demonstration techniques was undertaken, but our general impression was that the extended speaking time and concentration on program typescripts did orient the talks towards the details of how the programs operate. The major limitations were adequate but less-than-optimal imaging quality from the projection system (particularly in the largest auditorium), and the limited experience of AIM users with the equipment and software used. The SUMEX 2020 with the KI-10's as backup provided excellent computer support for the display technology. One group, the BELIEVER project from Rutgers, augmented their typescript with a "live" demonstration running on the Rutgers AIM resource. The SUMEX staff provided excellent support in the development of programs, equipment setup, and computer support.

A series of 20-minute parallel sessions was also provided for newer AIM systems under development. These talks used standard visual aids. If AIM conferences use demonstration sessions in the future, the featured programs should probably be chosen from among these developing systems.

Since the Workshop, several projects (including MOLGEN, GUIDON, and VM) have used the stored typescript for demonstrations. They have been useful when visitors wish to see a particular program but resources are not available to run the program in a real-time setting. Each of the demonstration files is available on the SUMEX system, and is available for access by all SUMEX users.

I.E.2 Tutorial on AI in Medicine

In conjunction with the AIM Workshop, a continuing education tutorial designed for physicians was held at Stanford on August 17-18, 1980. The tutorial was entitled "Computers in Medicine -- Applications of Artificial Intelligence Techniques" and was organized by Drs. W. Clancey and E. Shortliffe. The tutorial had a remarkably good attendance by physicians as well as several other individuals with an interest in the field. The course included an optional introduction to computers for those who had no prior experience with the technology, an overview of SUMEX-AIM research, and an introduction to background materials regarding decision theory and data base applications in medicine. Speakers also provided detailed presentations on MYCIN, CASNET/EXPERT, INTERNIST and GUIDON. The course closed with a panel discussion on the problems and promise of AI in Medicine. It was accredited for postgraduate medical education through Stanford University School of Medicine; the 135 physicians in attendance earned 11.5 continuing education credits. In addition, 18 students, several non-physician researchers, and 10 members of the press attended. Enrollees came from as far away as Mexico and the East Coast. For the reasonable fee of $40 covering the two days of lectures, the attendees also received a syllabus of readings and two lunches.

The syllabus is a comprehensive survey of medical AI research and is comprised of recent articles written by the tutorial faculty, mostly for a clinical audience. The faculty consisted of 15 distinguished researchers from the AIM community, including 7 physicians and 9 speakers from centers other than Stanford. By holding the tutorial immediately after the AIM Workshop and before the first Annual Meeting of the American Association for Artificial Intelligence (AAAI), we were able to secure the participation of expert physicians in the field who were already at Stanford (Drs. Greenes, Lindberg, Myers, and Pauker), as well as computer scientists from the East Coast (Drs. Davis, Kulikowski, Pople, Szolovits, and Swartout). Stanford speakers included Drs. Blum, Buchanan, Clancey, Feigenbaum, Fries, and Shortliffe. Coordination and planning for the tutorial was was facilitated by sending electronic messages; almost all speakers regularly use SUMEX or another ARPANET machine.

To evaluate the impact of the tutorial on the participants, and to assess baseline opinions regarding the field, we undertook a survey of the physicians' knowledge about computers as well as their attitudes towards medical consultation systems. The statistical analysis of these questionnaires has now been completed, and a paper summarizing the results submitted for publication (Teach, R.L. and Shortliffe, E.H. "An analysis of physician attitudes regarding computer-based clinical consultation systems." Submitted for publication, March 1981). In brief, the survey showed that physicians were willing to accept the possibility of computer-based clinical decision aids but placed severe demands on the capabilities of such systems if they were to be acceptable for routine use.

In addition, attendees were asked to evaluate the course itself, as well the the talks by individual speakers. These forms showed that the course was exceedingly well received. Attendees were fascinated by the content, generally felt it was well presented, and indicated they would recommend the course to others if it were made available again. Many physicians requested a follow-up course that would introduce them to more technical detail than had been possible in the introductory tutorial.

In conclusion, we believe that the tutorial was an encouraging success, and demonstrated the effectiveness of this kind of forum for introducing physicians to the research efforts within the AIM community. The faculty is enthusiastic about repeating the course, possibly on the East Coast in conjunction with a future AIM Workshop. Several members of the audience expressed interest in detailed, small group discussions of particular AIM programs. We believe these discussions could be a valuable way of exporting our methods and approach beyond the immediate AIM community.

I.E.3 GENET - An Experiment in AI System Dissemination

Background

The MOLGEN project at Stanford (see Section II.A.1.5 on page 136) has focused on applications of artificial intelligence and symbolic computation to the field of molecular biology. The research began in 1975 and is currently in the first year of a three year grant renewal. In early 1980 it was realized that some of the systems developed by MOLGEN were of direct utility to many scientists in the domain. Accordingly, with the cooperation of the SUMEX-AIM staff and close coordination with the AIM Executive Committee, it was decided in February 1980 to provide a carefully limited guest service for the community use of such systems.

There were two major reasons for the establishment of this guest service, which took the form of the GENET account on SUMEX. The first was to broaden MOLGEN's base of scientist collaborators, to find molecular biologists at institutions other than Stanford who could contribute actively to our knowledge-based approach to problem solving. The second was to introduce a generally computer-naive community to the benefits of resource sharing provided by a system like SUMEX, with the hope of serving as a model for the dissemination of other AI software and possibly for an eventual resource for molecular biology.

We believe that we have succeeded in these two goals. Many of our GENET guests have become active collaborators in core MOLGEN research. These collaborators include Professor Allan Maxam at Harvard Medical School, Dr. Walter Goad at Los Alamos, Dr. Richard Roberts at Cold Spring Harbor, Dr. William Pearson at Johns Hopkins, Drs. Walter Bodmer, Julia Bodmer, and Robert Kamen at the Imperial Cancer Research Fund, Professor Fred Blattner at Wisconsin, Dr. Andrew Taylor at University of Oregon, and Dr. Dan Davison of SUNY-Stonybrook. We are also pleased by the numerous comments SUMEX has received from GENET users praising the user-sensitive nature of the resource, especially in comparison to typical university computer centers.

GENET has been important both for MOLGEN and for the national community of molecular biology. It has ensured a steady flow of ideas for the artificial intelligence research that is core to both the MOLGEN grant and the SUMEX-AIM mission. It has also provided a useful service to an international community that is not readily available elsewhere.

GENET Community Management

Our decision to support the GENET guest experiment and our approach to doing so within the SUMEX-AIM resource has been reviewed and approved both by the AIM Executive Committee and by the Initial Review Group/National Advisory Research Resources Council in the course of the peer review of our pending SUMEX renewal application. We have tried to manage the GENET guest experiment in such a way that we maintain the "friendly" interface of the SUMEX-AIM resource for molecular biologists unfamiliar with computers while taking appropriate steps so that GENET usage does not detract from on-going AI research and so that we assure prudent administration SUMEX as an NIH-BRP resource. The key elements in our management approach include:

1) Controlled announcement of the GENET opportunity -- Beginning in February 1980, the availability of GENET services was announced, primarily by talks at professional conferences with accompanying program demonstrations. We decided against publishing "blanket" announcements in professional journals in order to maintain a very high standard of collaborator interest and scientific expertise within the limited group we could serve with available SUMEX resources.

2) Close coordination with the AIM Executive Committee -- We kept the AIM Executive Committee apprised of plans for the GENET experiment and of progress and growth of the community. At the August 1980 AIM Workshop meeting of the Executive Committee, Professor L. Kedes of the MOLGEN project made a presentation on the status of GENET. The Executive Committee approved continuation of the GENET service but because of the significant growth in the number of GENET users and their consumption of CPU resources, a limit of two simultaneous GENET jobs was placed on the community. The Executive Committee also approved the concept of a proposed Molecular Biology Computing Resource related to but separate from the existing SUMEX resource.

3) Careful control of GENET usage -- We have closely monitored the very rapid growth in GENET usage of SUMEX (see data below). With Executive Committee advice and in cooperation with the MOLGEN project personnel managing the GENET community, we have instituted several successively stringent controls on GENET users:

   a) All GENET users run out of the same directory so scheduler control limits are enforced to hold GENET usage as a whole down relative to that of AI research projects during heavy loads.

   b) The GENET directory has been intentionally limited in disk space allocation so that large numbers of files cannot be retained.

   c) Starting in October 1980, a limit of two simultaneous logged-in GENET jobs was placed on the community.

   d) Starting in December 1980, a policy statement was issued restricting GENET use to academic collaborators. MOLGEN project management informed industrial collaborators that they could no longer use the GENET facility and actively monitored adherence to this policy. Previously, valuable feedback had been obtained from a small group of industrial collaborators for MOLGEN AI program development. However, with the rapid growth of the highly competitive molecular genetics industry, there was no way we could adequately control industrial users consistent with SUMEX's status as a federally funded national resource. Thus, we decided to exclude them. In April 1981, we instituted a GENET user password checking system to further control community access, particularly in regard to industrial users.

4) Limited commitment of SUMEX staff resources -- The day to day management of the GENET community has been the responsibility of MOLGEN project personnel. SUMEX personnel have only contributed to developing system facilities to help manage GENET (guest and GENET password capabilities), assisted with technical communications problems, and advised in establishing GENET management policies consistent with AIM Executive Committee and SUMEX Principal Investigator resource policies. The total commitment of staff time has been on the order of 1-2 man-months.

Scope of the GENET User Community

The GENET community consists of approximately 200 users from 63 research institutions. Of these 200 users, approximately 35 are consistently active users. That is, they log in, run programs, and interact with the MOLGEN members on an almost daily basis. Many of these users have made valuable contributions to our work. About 100 others are frequent, but not regular users. They log in only when they have a major analysis task to perform, which seems to be on the order of once a month.

The remaining users rarely use the system. They have logged in a few times, but for one reason or another they never become regular users of the system. Quite often this is because a lab group will settle on having one or two graduate students or post-doctoral associates become the "computer experts" of the group, and as a result, the computer use by the other people in the lab drops to an almost non-existent level. Unfortunately, an equally prevalent reason for users to stop using the GENET account is a lack of resource time. Probably the major complaint that we get from GENET users is concerning the lack of compute time and availability of the system. One account just is not enough for 200 people to share, especially when it is restricted to 2 jobs at one time. We constantly remind the GENET users to use there resources wisely. We encourage them to use the BATCH system to run job in the wee hours of the morning, and we remind them to be prepared to do their work quickly when they log in to the system, but their efforts do not seem to help the problem very much.

Most GENET users use only a small set of programs. These consists of text editors, which are used to set up the data files that for the MOLGEN analysis programs; XSEARCH, which GENET users use to effectively search through our database for sequences that can assist them in their research; and the electronic mail facilities. Very few of our GENET users actually feel comfortable using programs other than the ones that we maintain, not because the other programs would not be useful, but instead because the users do not have the computer time to experiment with what is available.

There are three note-worthy programs that we provide for GENET users that are used extensively. SEQ, a DNA-RNA sequence analysis program, which is continually being improved, is the most widely used. MAP, a program that assists in the construction of restriction maps from restriction enzyme digest data, is also used a great deal. Finally, a new program, MAPPER (written and maintained by William Pearson from Johns Hopkins University), is a simplified version of the MOLGEN MAP program that is somewhat more efficient than the MOLGEN version. The MOLGEN UE program and special molecular genetics knowledge bases are not available to the general GENET user at this time for two reasons. First of all, the UE program is quite costly to use (in terms of computer cycles), and secondly, we feel that the knowledge base is not quite ready for the computer novice to learn and use without a significant amount of initial assistance. A few GENET users (mostly Stanford associates) that have had a significant interest in the knowledge base have become EXO-MOLGEN users and are developing knowledge bases on their own which we hope will eventually be added to the ones that MOLGEN is developing and maintaining.

GENET Usage Statistics

Following is a table of monthly statistics for GENET usage of SUMEX. Note "TOTAL CONNECT HOURS" includes connect time for local dialups, hardlines, ARPANET, and TYMNET. "TYMNET CONNECT HOURS" includes that part of the total connect time which is via TYMNET and for which SUMEX pays a separate usage charge. Recent GENET TYMNET usage has been about 20-25% of the total SUMEX TYMNET connect time. Our monthly TYMNET bills are about $5,000, so monthly GENET TYMNET usage is about $1,125. Most GENET users come from other parts of the country and no additional local dial-up lines have been installed to support GENET usage.

Month/ Year	CPU Hours	Total Connect Hours	TYMNET Connect Hours	GENET % of Sumex TYMNET Use	File Pages
Feb/80	3.23	32.72	18.88	2.0%	57
Mar/80	1.28	51.57	12.80	1.4	95
Apr/80	8.37	117.87	51.73	5.4	209
May/80	9.20	104.46	66.65	8.0	166
Jun/80	11.08	188.35	118.03	11.7	253
Jul/80	19.21	342.87	189.00	18.2	231
Aug/80	18.71	257.23	188.53	18.2	367
Sep/80	57.32	409.83	254.53	28.5	626
Oct/80	36.47	348.66	211.95	23.3	920
Nov/80	82.90	648.56	308.40	31.1	1133
Dec/80	19.86	295.85	188.67	22.8	1110
Jan/81	48.00	747.91	277.30	27.2	996
Feb/81	22.58	265.39	163.55	16.1	962
Mar/81	29.73	613.74	313.57	25.0	982
Apr/81	43.04	662.57	unavail	unavail	1633

Plots of the CPU usage, connect time, and file usage data can be found in Figures 16-18.

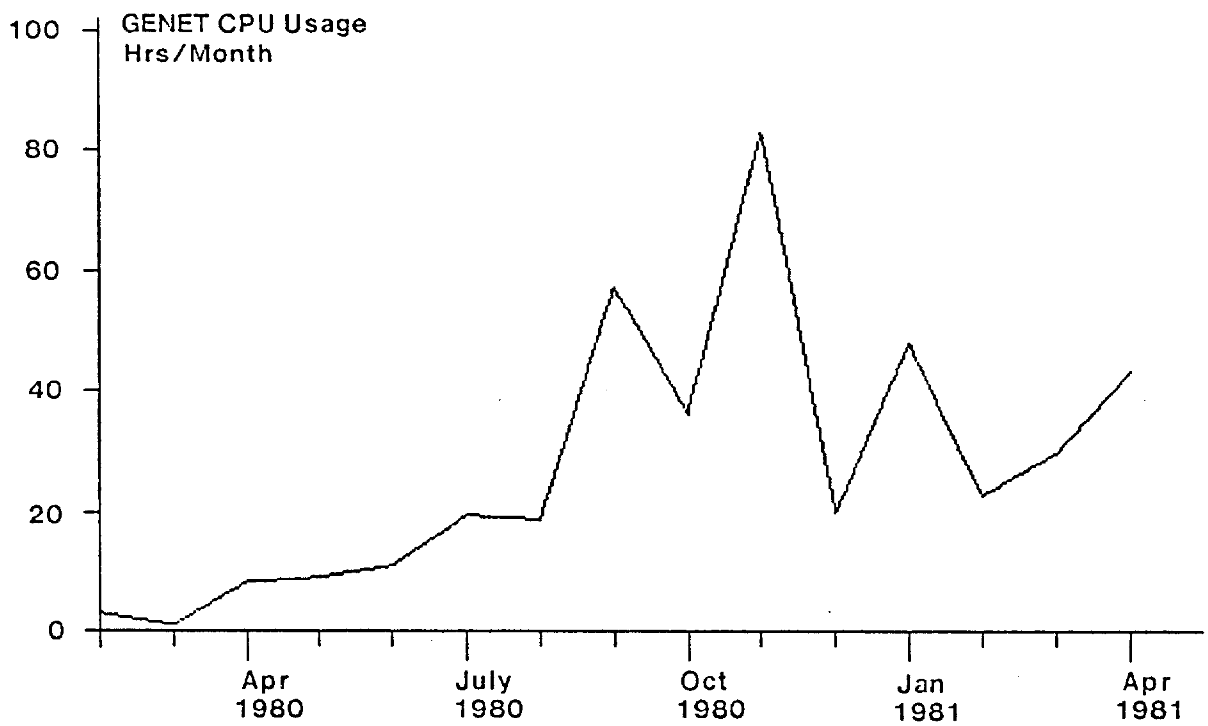

Figure 16. GENET CPU Usage by Month

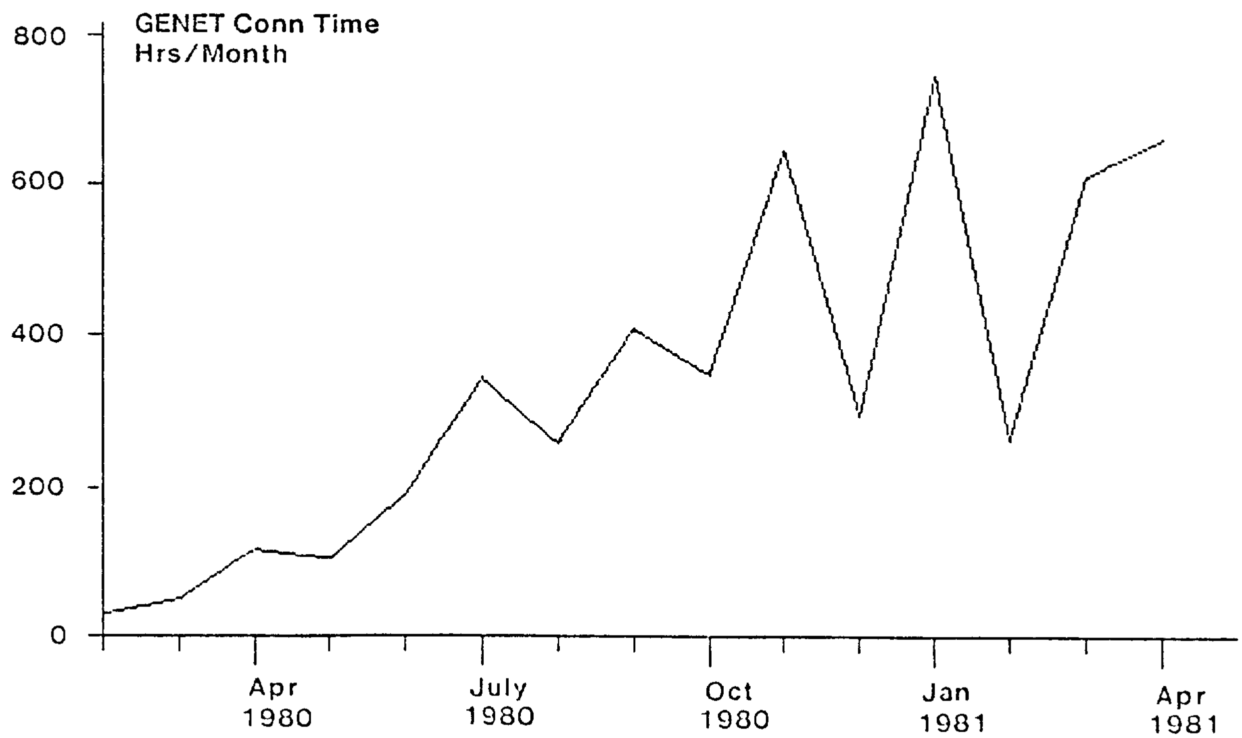

Figure 17. GENET Connect Time by Month

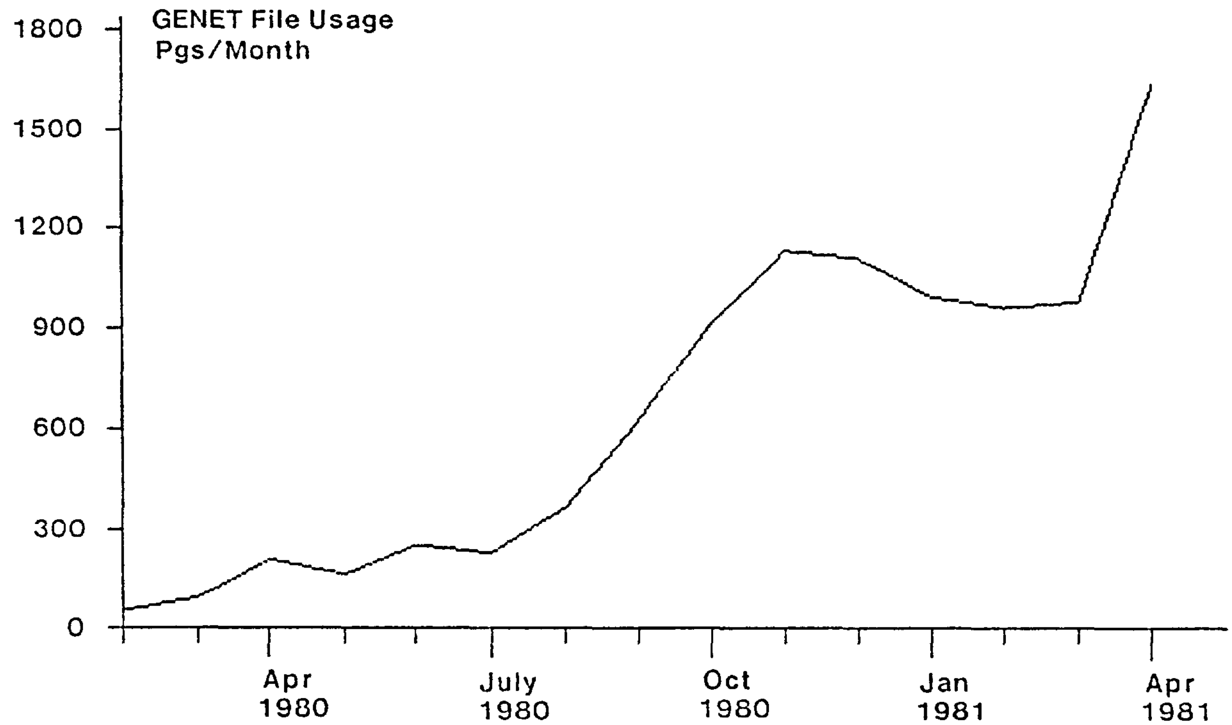

Figure 18. GENET File Space by Month

I.F Comments on the Biotechnology Resources Program

Resource Organization

We firmly believe that the Biotechnology Resources Program is one of the most effective vehicles for developing and disseminating technological tools for biomedical research. The goals and methods of the program are well-designed to encourage building of the necessary multi-disciplinary groups, merging appropriate technological and medical disciplines. In our experience with the SUMEX-AIM resource, several elements of this approach seem to emerge as key to the development and management of an effective resource:

1) Effective Management Framework - there needs to be an explicit agreement between the BRP and the resource principal investigator that sets out a clear mandate for the resource and its allocation, provides worthwhile incentives for the host institution and investigator to invest the necessary substantial professional career time to develop and manage the resource, and ensures equitable distribution of resource services to its target community.

2) Close Working Relationship with NIH - a resource is a major and often long-term investment of money and human energy. A close and mutually supportive working relationship between resource management, its advisory committees, and the NIH administration is essential to assure healthy development of the resource and its relationship to its user community. We at SUMEX-AIM have benefited immensely from such a relationship with Dr. William R. Baker, Jr. in the evolution of the SUMEX-AIM community.

3) Freedom to Explore Resource Potential - a resource, by its nature, operates at the "cutting edge" in developing its characteristic technology and learning how to effectively disseminate it to the biomedical community at large. BRP should not impose artificial constraints on the resource for commercializing its efforts (fees for service) or developing its potential (budget ceilings). Such artificial policy impositions can serve to undermine the very goals central to BRP's reason for existence. Satisfactory policies in this regard have been worked out recently and should be retained.

Electronic Communications

SUMEX-AIM has pioneered in developing more effective methods for facilitating scientific communication. Whereas face to face contacts continue to have their place, in the longer term we feel that computer-based communications will become increasingly important to NIH and the biomedical community. We would like to see BRP take a more active role in promoting these tools within NIH and its grantee community. A concrete step would be to become a sponsoring agency for the ARPANET which remains the most effective means for a very broad spectrum of services to promote good communications. This could serve as a base for interconnecting sponsored machines and offering a broader range of services and promoting broader collaboration among the biomedical community at large.

II Description of Scientific Subprojects

II.A Scientific Subprojects

The following subsections report on the AIM community of projects and "pilot" efforts including local and national users of the SUMEX-AIM facility at Stanford. Those using the Rutgers-AIM facility are annotated with "[Rutgers-AIM]". In addition to these detailed progress reports, we have included briefer summary abstracts of the fully authorized projects in Appendix A on page 278.

The collaborative project reports and comments are the result of a solicitation for contributions sent to each of the project Principal Investigators requesting the following information:

I. SUMMARY OF RESEARCH PROGRAM
- A. Project rationale
- B. Medical relevance and collaboration
- C. Highlights of research progress
  - --Accomplishments this past year
  - --Research in progress
- D. List of relevant publications
- E. Funding support (see details below)

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE
- A. Medical collaborations and program dissemination via SUMEX
- B. Sharing and interactions with other SUMEX-AIM projects (via computing facilities, workshops, personal contacts, etc.)
- C. Critique of resource management (community facilitation, computer services, communications services, capacity, etc.)

III. RESEARCH PLANS (8/80-7/86)
- A. Project goals and plans
  - --Near-term
  - --Long-range
- B. Justification and requirements for continued SUMEX use
- C. Needs and plans for other computing resources beyond SUMEX-AIM
- D. Recommendations for future community and resource development

We believe that the reports of the individual projects speak for themselves as rationales for participation; in any case the reports are recorded as submitted and are the responsibility of the indicated project leaders.

II.A.1 Stanford Projects

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Feigenbaum as Principal Investigator.

II.A.1.1 AGE - Attempt to Generalize

AGE - Attempt to Generalize

H. Penny Nii and Edward A. Feigenbaum
Computer Science Department
Stanford University

ABSTRACT: Isolate inference, control, and representation techniques from previous knowledge-based programs; reprogram them for domain independence; write an interface that will help a user understand what the package offers and how to use the modules; and make the package available to other members of the AIM community and labs doing knowledge-based programs development, and the general scientific community.

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The general goal of the AGE project is to demystify and make explicit the art of knowledge engineering. It is an attempt to formulate the knowledge that knowledge engineers use in constructing knowledge-based programs and put it at the disposal of others in the form of a software laboratory.

The design and implementation of the AGE program is based primarily on the experience gained in building knowledge-based programs at the Stanford Heuristic Programming Project in the last decade. The programs that have been, or are being, built are: DENDRAL, meta-DENDRAL, MYCIN, HASP, AM, MOLGEN, CRYSALIS [Feigenbaum 1977], and SACON [Bennett 1978]. Initially, the AGE program will embody artificial intelligence methods and techniques used in these programs. However, the long-range aspiration is to integrate those developed at other AI laboratories. The final product is to be a collection of building-block programs combined with an "intelligent front-end" that will assist the user in constructing knowledge-based programs. It is hoped that AGE will speed up the process of building knowledge-based programs and facilitate the dissemination of AI techniques by: (1) packaging common AI software tools so that they need not be reprogrammed for every problem; and (2) helping people who are not knowledge engineering specialists write knowledge-based programs.

B. Medical Relevance and Collaboration

AGE is relevant to the SUMEX-AIM Community in two ways: as a vehicle for disseminating cumulated knowledge about the methodologies of knowledge engineering and as a tool for reducing the amount of time needed to develop knowledge-based programs.

(1). Dissemination of Knowledge: The primary strategy for conducting AI research at the Stanford Heuristic Programming Project is to build complex programs to solve carefully chosen problems and to allow the problems to condition the choice of scientific paths to be explored. The historical context in which this methodology arose and summaries of the programs that have been built over the last decade at HPP are discussed in [Feigenbaum 1977]. While the programs serve as case studies in building a field of "knowledge engineering," they also contribute to a cumulation of theory in representation and control paradigms and of methods in the construction of knowledge-based programs.

The cumulation and concomitant dissemination of theory occur through scientific papers. Over the past decade we have also cumulated and disseminated methodological knowledge. In Computer Science, one effective method of disseminating knowledge is in the form of software packages. Statistical packages, though not related to AI, are one such example of software packages containing cumulated knowledge. AGE is an attempt to make yesterday's "experimental technique" into tomorrow's "tool" in the field of knowledge engineering.

(2). Speeding up the Process of Building Knowledge-based Programs: Many of the programs built at HPP are intelligent agents to assist human problem solving in tasks of significance to medicine and biology (see separate sections for discussions of work and relevance). Without exception the programs were handcrafted. This process often takes many years, both for the AI scientists and for the experts in the field of collaboration.

AGE will reduce this time by providing a set of preprogrammed inference mechanisms and representational forms that can be used for a variety of tasks. Close collaboration is still necessary to provide the knowledge base, but the system design and programming time of the AI scientists can be significantly reduced. Since knowledge engineering is an empirical science, in which many programming experiments are conducted before programs suitable for a task are produced, reducing the programming and experimenting time would significantly reduce the time required to build knowledge-based programs.

C. Highlights of Research Summary

Last year we reported the addition of Backchaining framework (the chaining of production rules in the manner similar to that used in MYCIN) and an interface to the Units package (for additional representational form and its use from AGE rules). In the past year we placed our research emphasis on (1) improving the existing component parts and the user interface, (2) developing debugging facilities, and (3) producing additional documents.

We completed the implementation of Trace and Break packages, as well as a facility for trace-back explanation. Using the trace-back facility users can inquire about the program's actions; AGE answers the questions by using the execution history list. Some example questions are: "What was the hypothesis before the execution of rule 2 in KS X?", "What Event led to the activation of KS X?". Since AGE has no knowledge of the application domain, it cannot "explain" the program actions in the language of the domain, but it produces "explanations" that are useful to the implementers.

We found that the specification and editing protocols for the various components were awkward and difficult for the users to learn. We redesigned this particular portion of the interface and have completed about 75% of the re-implementation.

In addition to the standard documents (a user's guide and a reference manual), we began a documented series of examples. These examples are actually implemented and running programs; each document consists of a description of the example problem, its formulation in terms of AGE, reasons for the particular formulation, and a complete program listing. In addition, the programs are available for the users to run. We observed that our documents, like most other program documentations, are useful only to those people who are already familiar with AGE. The Example Series is an experiment to see if a combination of standard documents and examples would be of any significant help to new users.

D. Publications

Nii, H. Penny and Aiello, Nelleke, "AGE: a knowledge-based program for building knowledge-based programs," Proc. of IJCAI-6, pp. 645-655, vol. 2, 1979.

Nii, H. Penny, "An Introduction of Knowledge Engineering, Blackboard Model, and AGE," HPP Working Paper, HPP-80-29.

Aiello, N. and Nii, H.P., "The Joy of AGE-ing: A User's Guide to AGE-1."

Aiello, N., Bock, C., Nii, H.P., White, W., "AGE Reference Manual."

AGE Example Series 1: "BOWL: A Beginner's Program."

AGE Example Series 2: "AGEPUFF: A Simple Event-Driven Program."

II. INTERACTION WITH THE SUMEX-AIM RESOURCES

AGE Availability:

Currently AGE-1 is available to a limited number of groups on the PDP-10 at the SUMEX-AIM Computing Facility and on the PDP-20/60 at the SCORE Facility of the Computer Science Department. The current implementation is described briefly in a later section.

Dissemination:

We previously reported a three-day workshop that we conducted in March 1980. The aims of the workshop were to familiarize the attendees with the use of AGE, and for each participant to implement a running program related to his application area. Of the attendees, the group from the Institute of Medical Electronics, University of Tokyo, has continued to use AGE to develop a medical diagnosis program.

In addition, many of the activities of the past year described earlier were direct results of what we learnt at the workshop.

For the 1980 AIM Workshop we reimplemented in AGE a major portion of the VM program (described elsewhere). In addition to demonstrating a variety of features of AGE, we were able to demonstrate the relatively short implementation time required once the goals of the application and the necessary knowledge were delineated -- a first-year graduate student had the program running in three weeks.

Profile of the Current AGE System:

To correspond to the two general technical goals described earlier, AGE is being developed along two separate fronts: the development of tools and the development of "intelligent" user interface.

Currently Implemented Tools:

The current AGE system provides the user with a set of preprogrammed modules called "components" or "building blocks". Using different combinations of these components, the user can build a variety of programs that display different problem-solving behavior. AGE also provides user interface modules that help the user in constructing and specifying the details of the components. A component is a collection of functions and variables that support conceptual entities in program form. For example, production rule, as a component, consists of: (1) a rule interpreter that support the syntactic and semantic description of production-rule representation as defined in AGE, and (2) various strategies for rule selection and execution.

The components in AGE have been carefully selected and modularly programmed to be useable in combinations. For those users not familiar enough to experiment with combining the components, AGE currently provides the user two predefined configuration of components--each configuration is called a "framework". One framework, called the Blackboard framework, is for building programs that are based on the Blackboard model [Lesser 77]. Blackboard model uses the concepts of a globally accessible data structure called a "blackboard", and independent sources of knowledge which cooperate to form hypotheses. The Blackboard model has been modified to allow flexibility in representation, selection, and utilization of knowledge. The other framework, called the Backchain framework, is for building programs that use backward-chained production rules as its primary mechanism of generating inferences.

The Front-End:

To support the user in the selection, specification, and use of the components, AGE is currently organized around four major subsystems that interact in various ways. Around it is a system executive that allows the user access to the subsystems through menu selection. Figure 1. shows the general interrelationship among these subsystems.

The Browse and Design subsystems help to familiarize the user with AGE and to guide the user in the construction of his programs through the use of predefined frameworks. The third subsystem is a collection of interface modules that help the user specify the various components of the framework. The last subsystem is designed for testing and refining the user program. Each of the subsystem is described in more detail below:

BROWSE: The function of Browse subsystem is to guide the user in browsing through its textual knowledge base, called the MANUAL. The MANUAL contains (a) a general description of the building-block components on the conceptual level; (b) a description of the implementation of these concepts within AGE; (c) a description of how these components are used within the object program; (d) how they can be constructed by the user; and (e) various examples. The information in the MANUAL is organized to represent the conceptual hierarchy of the components and to represent the functional relationship among them.

DESIGN: The function of the DESIGN subsystem is to guide the user in the design and construction of his program through the use of predefined configuration of components, or framework. Each framework is defined in DESIGN-SCHEMA, a data structure in the form of AND/OR tree, that, on one hand, represents all the possible configuration of components within the framework; and, on the other hand, represents the decisions the user must make in order to design the details of the user program. Using this schema, the DESIGN subsystem guides the user from one design decision point to another. At each decision point, the user has access to the MANUAL and also to advice regarding design decisions at that point. An appropriate ACQUISITION module can be invoked from the DESIGN subsystem so that general design and implementation specifications can be accomplished simultaneously.

ACQUISITION: For each component that the user must specify, there is a corresponding acquisition/editor module that queries the user for task-specific information. The calling sequence of the acquisition module is guided by DESIGN-SCHEMA when the user is using the DESIGN subsystem. They can also be accessed directly from the system menu or Interlisp.

INTERPRETER: This subsystem contains several modules that help the user run and debug his program. The Check module checks for the completeness and correctness of the specification for an entire framework. The Interpreter executes the user program. The Trace and Break modules are run-time debugging aids. The Editor, Check, Trace, Break, and the Explanation (described below) modules are designed to complement each other, and to help the user observe the workings of his program and to make corrections as necessary.

EXPLANATION: AGE has enough information to replay its execution steps, and it has reasonable justifications for the actions within the various framework. AGE provides a back-trace explanation facility whereby questions related to the execution history can be answered by the system interactively. However, AGE is totally ignorant of the user's task domain and has no means of conducting a dialogue about the specifics of the domain. A detailed history of the execution steps is available to the user to build his own domain specific explanation, if necessary.

```
    SYSTEM KNOWLEDGE        SUBSYSTEM            RESULT
                                |
     +-----------+       +-----V------+
     | MANUAL    |....>| BROWSE      |
     +-----------+...   +-----+------+
                  .           |
     +-----------+ .    +-----V------+      +------------+
     | DESIGN    |....>| DESIGN      |....>|USER SYSTEM |
     | SCHEMA    |...   |             |     |   DESIGN   |
     +-----------+ .    +-----+------+      +------+-----+
                   .          |<.................|
     +-----------+ .    +-----V------+      +------------+
     |COMPONENTS |....>| ACQUISITION|....>|    USER    |
     |           |      |   EDITOR   |     |   SYSTEM   |
     +-----------+      +------------+      +------+-----+
                              |<.................|
                        +------V-----+
                        |INTERPRETER |.....> EXECUTION
                        +------|-----+       HISTORY LIST
```

Figure 1. AGE System Organization
(... = data flow; --- = control flow)

III. RESEARCH PLAN

Research Topics:

The task of building a software laboratory for knowledge engineers is divided into two main sub-tasks:

1. The isolation of techniques used in knowledge-based programs: It has always been difficult to determine if a particular problem solving method used in a knowledge-based program is "special" to a particular domain or whether it generalizes easily to other domains. In existing knowledge-based programs, the domain specific knowledge and the manipulation of such knowledge using AI techniques are often so closely coupled that it is difficult to make use of the programs for other domains. One of our goals is to isolate the AI techniques that are general and determine precisely the conditions for their use.

2. Guiding the user in the initial application of these techniques: Once the various techniques are isolated and programmed for use, an intelligent agent is needed to guide the user in the application of these techniques. In AGE-1, we assume that the user understands AI techniques, knows what she wants to do, but does not understand how to use the AGE system to accomplish his task. A longer range interest involves helping the user determine what techniques are applicable to his task, i.e. it will assume that the user does not understand the necessary techniques of writing knowledge-based programs.

Research Plan:

In our judgement the the first research task has progressed enough to a point where can continue on to the second task. The system that embodies the results to date is called AGE-1. The structure of AGE-1 is now frozen and only minor modifications are being made. We will continue to support it by correcting bugs and adding requested features that are easily implementable.

AGE-2

AGE-2 will try to address the second of the research tasks described above.

Although the current Design subsystem provides specification functions that allow the user to interactively specify the knowledge of the domain and the control structure, it does not (aside from simple advise) provide the user any help in the actual design process. For example, AGE should be able to provide some aids to the user on what kinds of inference mechanisms and representations are appropriate for his application problem. We have stated this problem in our previous reports without any promising ideas on how we might attack this problem. With the variety of feedbacks we received from our experimental users, we now understand a few of the problems the inexperienced users are faced with. With these in mind, we have begun, and will continue, to explore ways in which we can redesign and add facilities that will help users who are not familiar with knowledge engineering techniques and methodologies.

One of the major obstacles in the way of AGE-2 development is the way in which AGE-1 is implemented. Although the syntax of AGE-1 is clearly defined (see the Reference Manual), the semantics are not well-defined. They are defined in ad hoc fashion in the Editor, the Interpreter, and the Check modules. In order for AGE-2 to be able to conduct a dialogue about itself with the user, its semantics, as well as its syntax, must be uniformly represented. Since very little research results are available in the area of representing the semantics of systems (one exception is in the automatic programming research), we need to experiment with a variety of approaches. We have already begun to look into some alternative representations. In changing the representation of the AGE system, no new components will be added, and minimum amount of changes will be made to the definition of the existing components.

Concurrent with re-representing the AGE system, we will identify a dozen or so framework, in addition of the existing two, that have simpler constructs and are easier for the novice users to understand. The simplicity will be achieved by providing less options for the user -- options which, because of their nature, are confusing to new users. Limiting the degrees of freedom for the user has the side benefit of allowing AGE to provide more specific description and aids. For example, in a very constrained framework we can provide a library of "standard" predicates for the users, which can have associated with them English translations; with such texts available the rules and the back-trace explanation can be printed in English-like form. Once the user is comfortable with the more simple frameworks, he can add complexity simply by replacing the predefined options selected for the frameworks.

Computing Resources and Management:

We believe the computing and communication resources provided by the SUMEX Facility make it one of the best in the country. The management is responsive to the needs of the research community and provides superb services. However, the system is getting to a point where no serious research and development is possible, because of the lack of computing cycles due to overcrowding. It is a compliment to the facility that there are so many users. On the other hand, our productivity has gone down in recent months, because of the heavy load on the system. It would appear that the situation will not improve on its own, since many of the projects that were small a few years ago are maturing into larger, more complex systems. Which is the way it should be. The environment in which the work is done also needs to grow. In short, without augmentation to the current computing power and storage space (which had never been generous), our ability to make research progress at SUMEX will be drastically curtailed.

II.A.1.2 AI Handbook Project

Handbook of Artificial Intelligence

E.A. Feigenbaum, A. Barr, and P. Cohen
Stanford Computer Science Department

I. SUMMARY OF RESEARCH PROGRAM

A. Technical Goals

The AI Handbook is a compendium of knowledge about the field of Artificial Intelligence. It is being compiled by students and investigators at several research facilities across the nation. The scope of the work is broad: Two hundred articles cover all of the important ideas, techniques, and systems developed during 20 years of research in AI. Each article, roughly four pages long, is a description written for non-AI specialists and students of AI. Additional articles serve as Overviews, which discuss the various approaches within a subfield, the issues, and the problems.

There is no comparable resource for AI researchers and other scientists who need access to descriptions of AI techniques like problem solving or parsing. The research literature in AI is not generally accessible to outsiders. And the elementary textbooks are not nearly broad enough in scope to be useful to a scientist working primarily in another discipline who wants to do something requiring knowledge of AI. Furthermore, we feel that some of the Overview articles are the best critical discussions available anywhere of activity in the field.

To indicate the scope of the Handbook, we have included an outline of the articles as an appendix to this report (see page 303).

B. Medical Relevance and Collaboration

The AI Handbook Project was undertaken as a core activity by SUMEX in the spirit of community building that is the fundamental concern of the facility. We feel that the organization and propagation of this kind of information to the AIM community, as well as to other fields where AI is being applied, is a valuable service that we are uniquely qualified to support.

C. Progress Summary

Because our objective is to develop a comprehensive and up-to-date survey of the field, our article-writing procedure is suitably involved. First drafts of Articles are reviewed by the staff and returned to the author (either an AI scientist or a student in the area). His final draft is then incorporated into a Chapter, which when completed is sent out for review to one or two experts in that particular area, to check for mistakes and omissions. After corrections and comments from our reviewers are incorporated by the staff, the manuscript is edited, and a final computer-prepared, photo-ready copy of the Chapter is generated.

We expect the Handbook to reach a size of approximately 1000 pages. Roughly two-thirds of this material will constitute Volumes I and II of the Handbook. The material in Volumes I and II will cover AI research in Heuristic Search, Representation of Knowledge, AI Programming Languages, Natural Language Understanding, Speech Understanding, Automatic Programming, and Applications-oriented AI Research in Science, Mathematics, Medicine, and Education. Researchers at Stanford University, Rutgers University, SRI International, Xerox PARC, RAND Corporation, MIT, USC-ISI, Yale, and Carnegie-Mellon University have contributed material to the project.

D. List of Relevant Publications

Many of the chapters of Volumes I and II of the AI Handbook have already appeared in preliminary form as Stanford Computer Science Technical Reports, authored by the respective chapter-editors. References follow. Other chapters of Volumes II and III will appear as Technical Reports in the summer and fall of 1981.

HPP-79-12 (STAN-CS-79-726)
Ann Gardner. Search.

HPP-79-17 (STAN-CS-79-749)
William Clancey, James Bennett, and Paul Cohen.
Applications-oriented AI Research: Education.

HPP-79-21 (STAN-CS-79-754)
Anne Gardner, James Davidson, and Terry Winograd.
Natural Language Understanding.

HPP-79-22 (STAN-CS-79-756)
James S. Bennett, Bruce G. Buchanan, and Paul R. Cohen.
Applications-oriented AI Research: Science and Mathematics.

HPP-79-23 (STAN-CS-79-757)
Victor Ciesielski, James S. Bennett, and Paul R. Cohen.
Applications-oriented AI Research: Medicine.

HPP-79-24 (STAN-CS-79-758)
Robert Elschlager and Jorge Phillips. Automatic Programming.

HPP-80-3 (STAN-CS-80-793)
Avron Barr and James Davidson. Representation of Knowledge.

E. Funding Support Status

The Handbook Project is partially supported under the Heuristic Programming Project contract with the Advance Research Projects Agency of the DOD, contract number MDA 903-77-C-0322, E. A. Feigenbaum, Principal Investigator and under the core research activities of the SUMEX-AIM resource.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Collaborations and Medical Use of Programs via SUMEX

We have had a modest level of collaboration with a group of students and staff at the Rutgers resource, as well as occasional collaboration with individuals at other ARPA net sites.

B. Sharing and Interactions with Other SUMEX-AIM Projects

As described above, we have had moderate levels of interaction with other members of the SUMEX-AIM community, in the form of writing and reviewing Handbook material. During the development of this material, limited arrangements have been made for sharing the emerging text. As final manuscripts are produced, they will be made available to the SUMEX-AIM community both as on-line files and in the hardcopy, published edition.

C. Critique of Resource Management

Our requests of the SUMEX management and systems staff, requests for additional file space, directories, systems support, or program changes, have been answered promptly, courteously and competently, on every occasion.

III. RESEARCH PLANS (8/80 - 7/83)

A. Long Range Project Goals

The following is the schedule for completion and publication of the AI Handbook:

May, 1981: Publication of Volume 1 by publisher (Wm. Kaufmann Inc., Los Altos, Ca.)

August, 1981: Submission of final copy to publisher for Volume II (publication by end of 1981).

August-September, 1981: Completion of Technical Reports containing chapters of Handbook

October, 1981: Submission of final copy of Volume III to publisher (for publication first quarter 1982)

(note: Volume I has been selected by the Library of Computer Science as their August,1981 book club selection)

B. Justifications and Requirements for Continued SUMEX Use

The AI Handbook Project is a good example of community collaboration using the SUMEX-AIM communication facilities to prepare, review, and disseminate this reference work on AI techniques. The Handbook articles currently exist as computer files at the SUMEX facility. All of our authors and reviewers have access to these files via the network facilities and use the document-editing and formatting programs available at SUMEX. This relatively small investment of resources will result in what we feel will be a seminal publication in the field of AI, of particular value to researchers, like those in the AIM community, who want quick access to AI ideas and techniques for application in other areas.

C. Needs and Plans for Other Computational Resources

We use document preparation programs at SUMEX and the Computer Science Department's SCORE machine. We have used and will continue to use a Computer Science Department phototypesetting machine, the Alphatype, to produce the final copy of the AI Handbook. The phototypesetting software called TEX, developed at Stanford, is the vehicle for this production.

D. Recommendations for Future Community and Resource Development

None.

II.A.1.3 DENDRAL Project

The DENDRAL Project
Resource-Related Research: Computers in Chemistry

Prof. Carl Djerassi
Department of Chemistry
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

The DENDRAL Project is a resource-related research project. The resource to which it is related is SUMEX-AIM, which provides DENDRAL its sole computational resource for program development and dissemination to the biomedical community.

A. Project Rationale

The DENDRAL project is concerned with the application of state-of-the-art computational techniques to several aspects of structural chemistry. The overall goals of our research are to develop and apply computational techniques to the procedures of structural analysis of known and unknown organic compounds based on structural information obtained from physical and chemical methods and to place these techniques in the hands of a wide community of collaborators to help them solve questions of structure of important biomolecules. These techniques are embodied in interactive computer programs which place structural analysis under the complete control of the scientist working on his or her own structural problem. Thus, we stress the word assisted when we characterize our research effort as computer-assisted structure elucidation or analysis.

Our principal objective is to extend our existing techniques for computer assistance in the representation and manipulation of chemical structures along two complementary, interdigitated lines. We are developing a comprehensive, interactive system to assist scientists in all phases of structural analysis (SASES, or Semi-Automated Structure Elucidation System) from data interpretation through structure generation to data prediction. This system will act as a computer-based laboratory in which complex structural questions can be posed and answered quickly, thereby conserving time and sample. In a complementary effort we are extending our techniques from the current emphasis on topological, or constitutional, representations of structure to detailed treatment of conformational and configurational stereochemical aspects of structure.

By meeting our objectives we will fill in the "missing link" in computer assistance in structural analysis. Our capabilities for structural analysis based on the three-dimensional nature of molecules is an absolute necessity for relating structural characteristics of molecules to their observed biological, chemical or spectroscopic behavior. These capabilities will represent a quantum leap beyond our current techniques and open new vistas in applications of our programs, both of which will attract new applications among a broad community of structural chemists and biochemists who will have access to our techniques. This access depends entirely on our access to and the continued availability of SUMEX-AIM. These issues are discussed in detail in the subsequent section, Interactions with the SUMEX-AIM Resource.

The primary rationale for our research effort is that structure determination of unknown structures and the relationship of known structures to observed spectroscopic or biological activity are complex and time-consuming tasks. We know from past experience that computer programs can complement the biochemist's knowledge and reasoning power, thereby acting as valuable assistants in solving important biomedical problems. By meeting our objectives we feel strongly that our programs will become essential tools in the repertoire of techniques available to the structural biochemist.

We are currently beginning the second year of our three year grant. This period represents a transition in the sense that we have pushed our research efforts in techniques for spectral interpretation, structure generation (e.g., CONGEN) and spectral prediction to their limits within the confines of topological representations of molecular structure. At this time, these techniques are perceived to be of significant utility in the scientific community as evidenced by our workshops, the demand for the exportable version of CONGEN and the number of persons requesting collaborative or guest access to our programs at Stanford (see Interactions with the SUMEX-AIM Resource). These existing techniques will, for some years to come, remain as important first steps in solving structural problems. However, in order to anticipate the future needs of the community for programs which are more generally applicable to biological structure problems and more easily accessible we must address squarely the limitations inherent in existing approaches and search for ways to solve them. Our major objectives are based on the following rationale.

None of our techniques (or the techniques of any other investigators) for computer-assisted structure elucidation of unknown molecular structures make full use of stereochemical information. As existing programs were being developed this limitation was less important. The first step in many structure determinations is to establish the constitution of the structure, or the topological structure, and that is what CONGEN, for example, was designed to accomplish. However, most spectroscopic behavior and certainly most biological activities of molecules are due to their three-dimensional nature. For example, some programs for prediction of the number of resonances observed in 13CMR spectra use the topological symmetry group of a molecule for prediction. However, in reality it is the symmetry group of the stereoisomer that must be used. This group reflects the usually lower symmetry of molecules possessing chiral centers and which generally exist in fewer than the total possible number of conformations. This will increase the number of carbon resonances observed over that predicted by the topological symmetry group alone. More generally, few of the techniques in the area of computer-assisted structure elucidation can be used in accurate prediction of structure/property relationships, whether the properties be spectral resonances or biological activities.

A structure is not, in fact, considered to be established until its configuration, at least, has been determined. Its conformational behavior may then be important to determine its spectroscopic or biological behavior. For these reasons we are emphasizing in our current grant period development of stereochemical extensions to CONGEN, our newly-developed structure generator, GENOA (see References 17, 18), and related programs such as the C-13 Nuclear Magnetic Resonance (NMR) programs (see References 15, 16), including machine representations and manipulations of configuration (see References 1, 10) and conformation (see Reference 19) and constrained generators for both aspects of stereochemistry (see References 6, 9, 11, 12).

None of the existing techniques for computer-assisted structure elucidation of unknown molecules, excepting very recent developments in our own laboratory, are capable of structure generation based on inferred partial structures which may overlap to any extent. Such a capability is a critical element in a computer-based system, such as we propose, for automated inference of substructures and subsequent structure generation based on what is frequently highly redundant structural information including many overlapping part structures. Important elements of our research are concerned with further developments of such a capability for structure generation (the GENOA program, (see Reference 17)).

Given the above tools for structure representation and generation, we can consider new interpretive and predictive techniques for relating spectroscopic data (or other properties) to molecular structure (see References 2, 3, 7, 8, 14, 15, 16). The capability for representation of stereochemistry is required for any comprehensive treatment of: 1) interpretation of spectroscopic data (see References 15, 16); 2) prediction of spectroscopic data (see References 15, 16); 3) induction of rules relating known molecular structures to observed chemical or biological properties (see Reference 19). These elements, taken together, will yield a general system for computer-aided structural analysis (the SASES system) with potential for applications far beyond the specific task of structure elucidation.

Parallel to our program development we have embarked on a concerted effort to extend to the scientific community access to our programs, and critical parts of our research effort are devoted to methods for promoting this resource sharing. Our rationale for this effort is that the techniques must be readily accessible in order to be used, and that development of useful programs can only be accomplished by an extended period of testing and refinement based on results obtained in analysis of a variety of structural problems, analyzed by those scientists actively involved in solutions to those problems. Our efforts in this area are summarized in Section II.A, Scientific Collaboration and Program Dissemination).

B. Medical Relevance and Collaboration

The medical relevance of our research lies in the direct relationship between molecular structure and biological activity. The sciences of chemistry and biochemistry rest on a firm foundation of the past history of well-characterized chemical structures. Indeed, structure elucidation of unknown compounds and the detailed investigation of stereochemical configurations and conformations of known compounds are absolutely essential steps in understanding the physiological role played by structures of demonstrated biological activity. Our research is focussed on providing computational assistance in several areas of structural chemistry and biochemistry, with primary attention directed to those aspects of the problem which are most difficult to solve by strictly manual methods. These aspects include exhaustive and irredundant generation of constitutional isomers, and configurational and conformational stereoisomers under chemical, biological and spectroscopic constraints with a guarantee that no plausible stereoisomer has been overlooked.

Although our programs can be applied to a variety of structural problems, in fact most applications by our group and by our collaborators are in the area of natural products, antibiotics, pheremones and other biomolecules which play important biochemical roles. In discussions of collaborative investigations involved with actual applications of our programs we have always stressed the importance of strong links between the structures under investigation and the importance of such structures to health-related research. This emphasis can be seen by examination of the affiliations of current DENDRAL-related investigators and the brief description of current collaborative efforts in Interactions with the SUMEX-AIM Resource.

C. Highlights of Research Progress

In this section we discuss briefly some major highlights of the past year and research currently in progress.

1. Past Year

1.1 _Exportable version of the CONGEN program for computer-assisted structure elucidation_. CONGEN is an interactive computer program whose task is to provide to the structural biochemist _all_ chemical structures which are possible candidates for the structure of an unknown chemical compound. Based on this information, experiments can be designed to pinpoint the correct structure, thereby facilitating rapid and unambiguous identification of novel, bioactive chemicals. During the past two years we have completed an exportable version of the CONGEN program and have exported it to a variety of structural analysis laboratories in academic, private and industrial research organizations. CONGEN is being utilized at Stanford and at export sites in the hands of investigators who use it as a tool in solving their own structural problems. We have been exporting versions of CONGEN for about 18 months. The program has been used as an aid in the solution of many new structures and recent results have formed the basis for at least eight formal lectures by users of CONGEN at remote sites.

1.2 _Version I of the GENOA program for structure generation with overlapping atoms_. GENOA (see Reference 17) is an outgrowth of CONGEN whose purpose is to suggest candidate structures for an unknown based on redundant and ambiguous structural inferences. This program, which utilizes CONGEN as an integral part of the computational procedures, is far simpler to use by the practicing biochemist. This results from GENOA's capability to construct structures based on substructural information obtained from a variety of spectroscopic, chemical and biochemical techniques. The program itself considers the structural implications of each new piece of structural data and automatically ensures that all overlaps are considered, thereby freeing the investigator from concerns about the potential for overlapping, or redundant substructural information. In addition, GENOA is the ideal tool for interfacing to automated procedures for spectral interpretation (see References 14, 15), because the necessity for manual intervention in the assignment of substructures is no longer required as it was for CONGEN.

1.3 Programs for Interpretation and Prediction of Spectral Data. We are actively pursuing several novel approaches to the automated interpretation of spectral data, concentrating on carbon-13 magnetic resonance (CMR), proton magnetic resonance (PMR) and mass spectral (MS) data. These approaches utilize large data bases of correlations between substructural features of a molecule and spectral signatures of such features. Our approaches are unique in that: 1) we can incorporate stereochemical features of substructures into the data bases; and 2) we can use the same data bases for both interpretation and prediction of data.

We have recently reported several new developments in the area of analysis of mass spectral data, including methods for mass spectral data interpretation (see Reference 14) and mass spectrum prediction (see References 3, 7, 8).

For either interpretation or prediction of magnetic resonance data, stereochemical substructure descriptors are absolutely essential. Resonance positions are a strong function of the local environment of a resonating atom, including position in space relative to other neighboring atoms. Descriptors which include the three dimensional relationships among atoms in a substructure are required in order to obtain meaningful correlations. We have recently completed the first phases of development of the data base and associated interpretation and prediction programs for C-13 NMR data (see References 15, 16). This approach uses a structure and substructure representation which incorporates configurational stereochemistry (see Reference 16).

Such data bases can be used to interpret spectral data to obtain substructures to be used in CONGEN and GENOA, the structure generating programs (see References 15, 17). Continued automation of this aspect of structure elucidation will significantly ease the burden on the structural biochemist because the computer-based files are much more comprehensive and easier to use than correlation tables or diffuse literature sources. The same data bases can be used to predict spectral signatures in the context of a set of complete molecular structures. Comparison of predicted and observed spectra allows a rank-ordering of candidates and will be very useful in directing the attention of the investigator to the most plausible alternatives (see References 7, 8, 15).

1.4 Constrained generation of configurational stereoisomers. During the previous grant period we solved the problem of computer generation of configurational stereoisomers. These are isomeric chemical structures that differ from one another in the arrangement of atoms in three-dimensional space. We have developed this method further, including now the capability for construction of all possible stereoisomers under stereochemical constraints (see Reference 9). Previously, CONGEN and GENOA were capable only of generation of constitutional isomers which convey no information about the structure in three dimensions. The interaction of biomolecules with biochemical systems is based on their three dimensional nature, not simply their constitution. Therefore, these new developments are crucial to use of computational techniques in structural studies.

Now, for the first time, a computer program can be used to begin with the molecular formula of an unknown compound and using constraints on both molecular connectivity and configuration arrive at a set of structural alternatives which include potential stereochemical variability. This capability allows use of spectral data whose interpretation depends strongly on stereochemical features of molecules. Most importantly, it gives us a structural representation and methods for structure generation and manipulation which represent the foundations for future developments of the one important remaining aspect of structural analysis, treatment of molecular conformations.

2. Research in Progress

The following are some highlights of research in progress. The common theme of these studies is representation of stereochemistry and use of stereochemical information in answering questions concerning the nature of known or unknown molecular structures.

2.1 Development of GENOA and STRCHK. GENOA can now deal with representations of configurational stereochemistry, although it does not make active use of such representations in generating constitutional isomers. The STRCHK (for Structure Checking) program represents the next stage in development of the post-generation analysis programs. STRCHK provides the entry point into the STEREO program for constrained generation of stereoisomers (see Reference 9) from constitutional isomers generated by either CONGEN or GENOA. In the case of GENOA, stereochemical information is passed to STRCHK where it can be used in STEREO. In addition, the mass and C-13 spectrum prediction and ranking programs (see References 7, 8, 14, 15) are available from STRCHK, together with several other utility programs for examining structural candidates. Both programs will be developed further, to the point where export to other computer facilities, as was done with CONGEN, will be possible.

2.2 Development of the C-13 data base and interpretive program. We plan further expansion of the C-13 NMR data base, using data obtained by us from the literature and supplied by others in collaborative efforts. Eventually we would like to pass this work on to an organization better equipped to build and maintain data bases. For the time being, however, our work is sufficiently experimental that we will maintain responsibility for the data base. The C-13 spectrum interpretation program will continue to be developed, as we attempt to make the program more "intelligent" chemically.

2.3 Representation and manipulation of conformational stereochemistry. The next year will see intensive efforts to develop programs for representation of molecular conformations. Preliminary work has led to an algorithm for representation and enumeration of conformations, and to a method for searching for common three-dimensional substructures in a set of structures (see Reference 19). The former study will first be directed to ward a program for representation of substructures with conformation designations. This will lead directly to a method for development of a data base and prediction program for any spectroscopic technique, such as proton NMR, where the spectral signatures are strongly influenced by molecular configurations. Subsequently, a program for generation and, eventually, constrained generation of molecular conformations, will be developed. Parallel to this work, the program for searching for three-dimensional common substructures, a problem important in structure/biological activity correlations, will be developed and tested extensively on previous studies presented in the literature. This work (see Reference 19) is based to an extent on similar work carried out for constitutional representations of structure (see Reference 5).

D. List of Recent Publications

(1) J.G. Nourse, R.E. Carhart, D.H. Smith, and C. Djerassi, "Exhaustive Generation of Stereoisomers for Structure Elucidation," J. Am. Chem. Soc., 101, 1216 (1979).

(2) C. Djerassi, D.H. Smith, and T.H. Varkony, "A Novel Role of Computers in the Natural Products Field," Naturwiss., 66, 9 (1979).

(3) N.A.B. Gray, D.H. Smith, T.H. Varkony, R.E. Carhart and B.G. Buchanan, "Use of a Computer to Identify Unknown Compounds. The Automation of Scientific Inference," Chapter 7 in "Biomedical Applications of Mass Spectrometry, First Supplementary Volume," G.R. Waller and O.C. Dermer, Eds., John Wiley and Sons, Inc., New York, 1980, p. 125.

(4) T.C. Rindfleisch, D.H. Smith, W.J. Yeager, M.W. Achenbach, and A. Wegmann, "Mass Spectrometer Data Acquisition and Processing Systems," in Chapter 3 of "Biomedical Applications of Mass Spectrometry, First Supplementary Volume," G.R. Waller and O.C. Dermer, Eds., John Wiley and Sons, Inc., New York, 1980, p. 55.

(5) T.H. Varkony, Y. Shiloach, and D.H. Smith, "Computer-Assisted Examination of Chemical Compounds for Structural Similarities," J. Chem. Inf. Comp. Sci., 19, 104 (1979).

(6) J.G. Nourse and D.H. Smith, "Nonnumerical Mathematical Methods in the Problem of Stereoisomer Generation," Match, (No. 6), 259 (1979).

(7) N.A.B. Gray, R.E. Carhart, A. Lavanchy, D.H. Smith, T. Varkony, B.G. Buchanan, W.C. White, and L. Creary, "Computerized Mass Spectrum Prediction and Ranking," Anal. Chem., 52 1095 (1980).

(8) A. Lavanchy, T. Varkony, D.H. Smith, N.A.B. Gray, W.C. White, R.E. Carhart, B.G. Buchanan, and C. Djerassi, "Rule-Based Mass Spectrum Prediction and Ranking: Applications to Structure Elucidation of Novel Marine Sterols," Org. Mass Spectrom., 15 355 (1980).

(9) J.G. Nourse, D.H. Smith, and C. Djerassi, "Computer-Assisted Elucidation of Molecular Structure with Stereochemistry," J. Am. Chem. Soc., 102, 6289 (1980).

(10) J.G. Nourse, "Applications of Artificial Intelligence for Chemical Inference. 28. The Configuration Symmetry Group and Its Application to Stereoisomer Generation, Specification, and Enumeration.", J. Amer. Chem. Soc., 101, 1210, (1979).

(11) J.G. Nourse, "Application of the Permutation Group to Stereoisomer Generation for Computer Assisted Structure Elucidation.", in "The Permutation Group in Physics and Chemistry", Lecture Notes in Chemistry, Vol. 12, Springer-Verlag, New York, (1979), p. 19.

(12) J.G. Nourse, "Applications of the Permutation Group in Dynamic Stereochemistry" in "The Permutation Group in Physics and Chemistry", Lecture Notes in Chemistry, Vol. 12, Springer-Verlag, New York, (1979), p. 28.

(13) J.G. Nourse, "Selfinverse and Nonselfinverse Degenerate Isomerizations," J. Am. Chem. Soc., in press (1980).

(14) N.A.B. Gray, A. Buchs, D.H. Smith, and C. Djerassi, "Computer-Assisted Structural Interpretation of Mass Spectral Data," Helv. Chim. Acta, in press (1981).

(15) N.A.B. Gray, C.W. Crandell, J.G. Nourse, D.H. Smith, and C. Djerassi, "Computer-Assisted Interpretation of C-13 Spectral Data," J. Org. Chem., 46 703 (1981).

(16) N.A.B. Gray, J.G. Nourse, C.W. Crandell, D.H. Smith, and C. Djerassi, "Stereochemical Substructure Codes for C-13 Spectral Analysis," Org. Magn. Res., 15, 375 (1981).

(17) R.E. Carhart, D.H. Smith, N.A.B. Gray, J.G. Nourse, and C. Djerassi, "GENOA: A Computer Program for Structure Elucidation Based on Overlapping and Alternative Substructures," J. Org. Chem, 46, 1708 (1981).

(18) D.H. Smith, N.A.B. Gray, J.G. Nourse, and C.W. Crandell, "The DENDRAL PROJECT: Recent Advances in Computer Assisted Structure Elucidation," Anal. Chim. Acta, Computer Techniques and Optimization, in press, (1981).

(19) D.H. Smith, J.G. Nourse, and C.W. Crandell, "Computer Techniques for Representation of Three-Dimensional Substructures and Exploration of Potential Pharmacophores," Proceedings of a Chemical Industries Institute of Technology Symposium on "Structure Activity Correlation as a Predictive Tool in Toxicology," Feb. 10-12, 1981, Raleigh, NC, in press.

E. Funding Support

Title:

RESOURCE RELATED RESEARCH: COMPUTERS IN CHEMISTRY (grant)

Principal Investigator:

Carl Djerassi, Professor of Chemistry, Department of Chemistry, Stanford University

Dennis H. Smith (Associate Investigator), Senior Research Associate, Department of Chemistry, Stanford University

Funding Agency:

Biotechnology Resources Program, Division of Research Resources, National Institutes of Health

Grant Identification Number:

RR-00612-12

Total Award and Period:

Total - 5/1/80 - 4/30/83 --------- $641,419

Current Award and Period:

Current - 5/1/81 - 4/30/82 -------- $237,387

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

In the coming period of our research, our computational approaches to structural biochemistry will become much more general and we plan wide dissemination of the programs resulting from our work. These more general approaches to aids for the structural biochemist will yield computer programs with much wider applicability than, for example, the existing CONGEN, GENOA, STEREO and STRCHK programs. We expect that this will create a significant increase in requests for access to our programs, placing heavy emphasis on our relationship with SUMEX to provide this access (see Justification and Requirements for Continued SUMEX Use for additional details).

For these reasons, in our current grant period the SUMEX-AIM resource is identified as the resource to which our research is related. The SUMEX-AIM resource has provided the computational basis for our past program developments and for initial exposure of the scientific community to these programs. The resource is, however, funded completely separately from our own research; we are only one of a nationwide community of users of the SUMEX-AIM facility. Our relationship to SUMEX is one which goes far beyond mere consumption of cycles on the SUMEX machine. It has been the goal of the SUMEX project to provide a computational resource for research in symbolic computational procedures applied to health-related problems. As such research matures, it produces results, among which are computer programs, of potential utility to a broad community of scientists. A second goal of SUMEX has been to promote dissemination of useful results to that community, in part by providing network access to programs running on the SUMEX-AIM facility during their development phases. SUMEX does not, however, have the capacity to support extensive operational use of such programs. It was expected from the beginning that user projects would develop alternative computing resources as operational demands for their programs grew. Such a state has been reached for the CONGEN, GENOA, STEREO and STRCHK programs and future developments in the DENDRAL Project to yield more generally useful programs will simply magnify the problem.

We will, therefore, under our relationship with SUMEX-AIM, participate as before in the SUMEX-AIM community in sharing methods and results with other groups during development of new programs. In addition, we plan to utilize the small machines requested as part of the SUMEX renewal. Our project will benefit by being able to provide more extensive operational access to our existing and developing programs using these machines, and to provide a test environment for adapting our programs to a more realistic laboratory computing environment than the special-purpose SUMEX resource (see Justification and Requirements for Continued SUMEX Use for additional information). SUMEX will benefit by moving a substantial part of the DENDRAL production load to more cost-effective systems, thereby freeing the SUMEX resource for new program development. Collaborators who wish to use existing programs for specific problems would access SUMEX via the network as before, but now would be routed to new machines. New program developments will be carried out on SUMEX itself, taking advantage of the much more extensive repertoire of peripheral devices, languages, debugging tools and text editors, i.e., precisely the tasks for which that system was designed.

Our proposed relationship to SUMEX-AIM has important implications beyond the practical considerations mentioned above. There is a significant research component to our proposal to make small machines as integral part of the resource sharing aspects of our relationship to SUMEX. The DENDRAL project is one of the first of the SUMEX-AIM projects to have developed sufficient maturity to require additional computer facilities to support production use and to facilitate export of its programs to be applied to real-world, biomedical structural problems. In a sense, then, we will be acting in a pathfinding role for the rest of the SUMEX-AIM community as other projects reach maturity and seek realistic mechanisms for dissemination of their software to meet the computational needs of their collaborators. Cooperating with SUMEX in the use of small machines, implementing new software, regulating access to divert development and applications to the appropriate machine are all experiments which we are willing to undertake together with SUMEX, knowing that we will be providing direction to future efforts along similar lines.

We will also be in a pathfinding role for a large segment of the biochemical community involved in computing, as we explore the utility of machines which will be much more widely available in Department and laboratory environments than DEC-10's and -20's. There are currently very few widely available computing resources which provide access to symbolic, problem solving programs operating in an interactive environment, We would be able to fulfill that need to the extent that applications have direct biomedical relevance, to the limits of our share of the SUMEX-AIM computing resource.

A. Scientific Collaboration and Program Dissemination

Scientific Collaborations:

The following is a brief description of collaborative efforts that have been taking place or will soon commence in the use of DENDRAL programs for various aspects of structural analysis.

1) Drs. Larry Anderson and Elliott Organick, Depts. of Fuels Engineering and Computer Science, University of Utah.

Dr. Anderson's research is in establishing the structure of coal and related polymers via various thermal and chemical degradation schemes. The degradation products are of interest to both energy and environmental studies. Professor Organick is responsible in part for the computer and graphics facility on which CONGEN and related programs can be run. We are exploring with them structure representations based on the Superatom concept in CONGEN as a means of representing families of structures. Access to our programs is primarily via the computer facility at Utah.

2) Dr. Raymond Carhart, Lederle Laboratories.

Dr. Carhart (a former member of our group) is engaged in research concerned with computer applications to structure/activity relationships. Program development is done jointly between Lederle and Stanford with free exchange of software. Lederle applications are carried out on their own computer facility.

3) Dr. Janet Finer-Moore, University of Georgia.

Dr. Finer-Moore is engaged in structure analysis of alkaloids in Dr. Peletier's group at Georgia. This research makes extensive use of 13C NMR. Our collaboration involves the development and application of our 13C interpretive and predictive programs in structure elucidation of new compounds based on an extensive set of 13C data available on closely related compounds. Access is via network to our programs at Stanford. We have just completed the draft of a manuscript as a result of this collaboration. (Dr. Finer-Moore has recently moved to the University of California, San Francisco.)

4) Dr. Brenda Kimble, University of California, Davis.

Dr. Kimble's research is in structural analysis of compounds which are present in trace amounts in environmental milieus and which show mutagenic activity. Many of these compounds are largely aromatic. We are developing the capabilities of our programs to deal efficiently with large, polynuclear aromatic compounds. Access to our programs is via network to Stanford.

5) Dr. Fred McLafferty, Cornell University.

Dr. McLafferty's research is involved with instrumental and analytical aspects of mass spectrometry. We are working with him on the development and application of an interface between his STIRS system and CONGEN/GENOA for structure determination based on mass spectral data. Part of this collaboration is development of IBM versions of some of our programs. Access is in part to Stanford, shifting primarily to Cornell as development proceeds.

6) Dr. David Cowburn, The Rockefeller University

Dr. Cowburn's research is in the area of conformational analysis, primarily of peptides. We are working with him on the development and application of our programs for generation of molecular conformations. Dr. Cowburn's works with large ring peptides which represent a significant challenge for a conformation generator. His participation will help assure an eventual program of practical use rather than just theoretical interest. Collaboration will be via network access to our programs at Stanford.

7) Dr. Gilda Loew, SRI International and The Rockefeller University.

Dr. Loew's research is in the area of quantitative structure/activity relationships, using primarily the methods of quantum mechanical calculations. We are working with her to interface our conformational generator to her coordinate-based calculation methods. Collaboration is carried out via accounts at Stanford with concurrent development of her programs on a VAX facility (NASA Ames Research Center).

8) Dr. D.C. Rohrer, Medical Foundation of Buffalo Research Laboratories, Buffalo, New York.

We have initiated a collaboration with Dr. Rohrer on the problem of finding the common 3-dimensional substructural features of a set of chemical structures. The use of such a program would be to postulate substructural features which are responsible for similar biological or spectral properties. The initial approach is similar to that used successfully to find the greatest common subgraph of a set of constitutional structures. Collaboration will be via network access to Stanford.

9) Dr. J.N. Shoolery, Varian Associates, Palo Alto

We are collaborating with Dr. Shoolery and others at Varian to obtain high quality C-13 spectra of several marine sterols available only in very small quantities. This is being done as part of our ongoing project to develop programs which are capable of spectral interpretation and prediction. The Varian people access our programs directly or via network.

Program Dissemination:

We have provided access to our programs to a community of collaborators via 1) distribution of the CONGEN program to other laboratories, and 2) guest or individual accounts on the SUMEX computer facility here at Stanford. These methods to promote the dissemination and use of our programs are elaborated below, followed by a brief description of some of our collaborations.

a) Program Export

The past two years we have distributed CONGEN to a number of laboratories owning computers on which the exportable version can now execute. These currently include DEC PDP-10 and -20 systems operating under the TENEX, TOPS-10 and TOPS-20 operating systems, and more recently, the beginnings of a version for IBM systems. The following persons are currently running CONGEN on their own laboratory computers:

Dr. Larry Anderson - University of Utah
(work described in section on collaborations)

Dr. Hartmut Braun - Organische-Chemisches Institut der Universitat Zurich, Switzerland
A former member of Prof. Wipke's group at UC Santa Cruz. He has only recently installed the program at ETH, Zurich.

Dr. Raymond Carhart - Lederle Laboratories
(work described in section on collaborations)

Dr. Roy Carrington - Shell Biosciences Laboratory, England
Dr. Carrington has used the program both as a guest user and recently in export. He has given presentations on the use of CONGEN and has applied the program to the structure determination of a new acidic amino acid, 2,4-methanoglutamic acid, and other compounds from plant seeds. This work was done in collaboration with Prof. Jon Clardy at Cornell who is also a guest user.

Dr. Robert Carter - University of Lund, Sweden
Dr. Carter obtained a version of the program for use of several groups at Universities in Sweden.

Dr. Daniel Chodosh - Smith, Kline & French Laboratories
He has installed CONGEN and written an extensive users' manual for the use of SKF chemists.

Dr. Henry Dayringer - Monsanto Agricultural Products Co.
He and Dr. Schwenzer (now at Gulf) were responsible for obtaining and installing CONGEN. Primary use is as an aid to structure elucidation of photoproducts and metabolites of agricultural chemicals.

Dr. Douglas Dorman - Lilly Research Labs
Dr. Dorman has been one of our best users. He attended our 1978 workshop and has given several presentations on the use of CONGEN. He has used the program as an aid in solving a number of structures including some beta-lactam antibiotic derivatives.

Dr Philip Ihrig - Amoco Standard Oil (Indiana)

Dr. Martin Huber - Ciba-Geigy, Switzerland
Dr. Huber is a former member of Prof. Wipke's group at UC Santa Cruz. He has recently received the program and is currently working to interest his coworkers at Ciba in computer assisted structure elucidation.

Dr. Carroll Johnson - Oak Ridge National Laboratory
Dr. Johnson is a long time colleague who spent a year at Stanford in 1976. He is involved with the analytical group at Oak Ridge and is using the program as an analytical aid and as a model for programs he is developing.

Dr. G. Jones - ICI Pharmaceuticals, England
He has installed CONGEN and is currently evaluating its utility for use by analytical chemists at ICI.

Dr. Fred W. McLafferty - Cornell University
(work summarized under collaborations)

Dr. Peter W. Milne - CSIRO Division of Computing Research, Australia
He contacted us through his association with the Heuristic Programming Project at Stanford. He has acted as the Australian contact for distribution of CONGEN in that country.

Dr. James Morrison - Latrobe University, Australia
(see Milne, above)

Dr. David Pensak - E.I. duPont de Nemours and Company
(see EXODENDRAL account DUPONT, and workshop)

Dr. Joseph SanFilippo - Rutgers University
Dr. SanFillippo is using CONGEN in conjunction with his work on superoxide chemistry and in the evaluation of mass spectral data for environmental samples.

Dr. William Sieber - Sandoz, Ltd., Switzerland
He has installed CONGEN for use by structural chemists at Sandoz. Currently they are evaluating its utility.

Dr. M.D. Sutherland - University of Queensland, Australia
(see Milne, above)

Dr. R.O. Watts - Australian National University
(see Milne, above)

b) EXODENDRAL Account

We reserve a special account on SUMEX for persons interested in access to our programs. Initially, this account was used for anyone desiring access, independent of expected level of use or eventual interest. As the SUMEX system became more heavily loaded a mechanism for guest access was provided and at that point we began to differentiate our users by level of interest. For those desiring merely to try programs we provide guest access (see page 119). If there is interest in continuing collaboration, EXODENDRAL status is given, which provides access to more system facilities and good file management capabilities. The persons who have been active under EXODENDRAL status this year are the following (with the account name followed by the contact person and association):

<BRAEKMAN>
Dr. Jean-Claude Braekman - Universite Libre de Bruxelles, Belgium
He is a former post doctoral fellow in our group, and accesses CONGEN from Belgium for natural products structure elucidation.

<BRAUN>
Dr. Hartmut Braun - Organische-Chemisches Institut der Universitat Zurich, Switzerland
(see section on export)

<CARRINGTON>
Dr. Roy Carrington - Shell Biosciences Laboratory, England
(see section on export)

<COWBURN>
Dr. David Cowburn - The Rockefeller University
(see section on collaborations)

<DORMAN>
Dr. Douglas Dorman - Lilly Research Laboratories
(see section on export)

<DREIDING>
Dr. Andre Dreiding - Organische-Chemisches Institut der Universitat Zurich, Switzerland
He has used CONGEN and STEREO extensively in structural studies. He has also worked closely with Braun (see section on export under Braun).

<DUPONT>
Dr. Earl Abrahamson - E.I. duPont de Nemours and Company
Dr. Abrahamson and 4 colleagues attended our 1980 workshop. They are attempting to integrate our program into their overall computer software system which includes a wide variety of programs for applications to chemical problems.

<FINER-MOORE>
Dr. Janet Finer-Moore - University of Georgia
(see section on collaborations)

<GASH>
Dr. Kenneth Gash - California State College at Dominguez Hills

<HELLER>
Dr. Steven Heller - Environmental Protection Agency
We are continuing our work with the NIH/EPA Chemical Information System, through Heller, to attempt to find mechanisms for making CONGEN accessible through that system.

<HUBER>
Dr. Martin Huber - Ciba-Geigy, Switzerland
(see section on export)
<MILNE>
Dr. Peter W. Milne - CSIRO Division of Computing Research, Australia
(see section on export)

<MONSANTO>
Dr. Henry Dayringer - Monsanto Company
(see section on export)

<MWOOD>
Dr. Mark Wood - Rutgers University

<RCARHART>
Dr. Raymond Carhart - Lederle Laboratories
(see section on collaborations)

<ROHRER>
Dr. Douglas C. Rohrer - Medical Foundation of Buffalo
(see section on collaborations)

<ROUSSEL>
Dr. Jean Mathieu - Roussel UCLAF
(see section on guest access under Delaroff)

<SIEBER>
Dr. William Sieber - Sandoz Ltd., Switzerland
(see section on export)

<VARIAN>
Dr. James Shoolery - Varian Associates
(see section on collaborations)

c) GUEST Access

We have provided GUEST access to our programs for those persons desiring occasional access to study a structural problem and for those who wish a "hands-on" introduction to the programs. Persons who have received information about this method of access are listed below (and the names of those who have actually logged in as guests are preceded with an asterisk):

*Dr. Robert Adamski - Alcon Labs

*Dr. A. Bothner-by - Carnegie Mellon University
Dr. Bothner-by has requested access to aid others in the Chemistry Department with structure elucidation work.

*Dr. Reimar Bruening - Institut fur Pharmazeutische Arzneimittellehre der Universitat, West Germany
Dr. Bruening has used the program to aid in his solution of the structure of the alkaloid Cassine. He was a participant in our 1978 workshop and has maintained interest since then. He has given at least one presentation in Germany on our programs.

*Dr. William Brugger - International Flavors and Fragrances
Dr. Brugger is interested in eventually obtaining CONGEN for use at IFF in natural products structure elucidation.

*Dr. Robert Carter - University of Lund, Sweden
(see section on export)

*Dr. Francois Choplin - Institut Le Bel, France

*Dr. Jon Clardy - Cornell University
He has used CONGEN on occasion to determine the potential structural variety for an unknown prior to obtaining the X-ray crystal structure.

Dr. Brian Coleman - Koninklijke/Shell-Laboratorium, Holland

*Dr. Mike Crocco - American Hoechst Corp.

*Dr. V. Delaroff - Roussel UCLAF, France
Dr. Delaroff attended our 1980 workshop. He is in charge of a spectroscopic team which checks structures and suggests structures for unknown compounds with important biological activities. They have been using our programs by remote access to aid these investigations.

*Dr. Dan Dolata - University of California at Santa Cruz
He is one of our contacts with Prof. Wipke's group at UC Santa Cruz.

*Dr. Bruno Frei - Laboratorium f. Organische Chemie, Switzerland

*Dr. Y. Gopichand - University of Oklahoma
He has worked with Prof. Schmitz on the solution of several structures of various marine natural products.

*Dr. John Gordon - Kent State University
Dr. Gordon has been using CONGEN while working at Chemical Abstracts in Columbus, Ohio. He has been using CONGEN to investigate general issues of structure representation.

Dr. Peter Gund - Merck, Sharpe and Dolme Research Labs

*Ms. Wendy Harrison - University of Hawaii at Manoa
Ms. Harrison is a student with Prof. Scheuer at Hawaii. She attended our 1978 workshop and has used the programs occasionally as an aid to structure determination in marine chemistry.

Dr. J. Hartenstein - Goedecke Co., Germany

*Dr. Richard Hogue - University of California at Santa Cruz
He is another contact with Prof. Wipke's group.

Dr. H. Honig - Institut fur Organische-Chemie u. Organisch-Chemische Technologie, Austria

Dr. Kenneth Houk - Louisiana State University
Dr. Houk has recently moved to Pittsburgh where he hopes to develop closer contact with our group.

Dr. H. Kating - Institut fur Pharmazeutische Biologie, Germany

Dr. Brenda Kimble - University of California at Davis
(see section on collaborations)

Dr. Sydell Lewis - University of California at Berkeley

*Dr. David Lynn - University of Virginia
Dr. Lynn attended our 1978 workshop when he was working with Prof. Nakanishi at Columbia.

*Dr. In Ki Mun - Cornell University
(see section on collaborations under McLafferty)

*Dr. Koji Nakanishi - Columbia University
We have worked with him and his students (see Lynn) on structures of several synthetic and natural products.

*Dr. Suba Neir - Washington University, St. Louis
Dr. Neir used the program to aid in determination of the structure of a mutagen.

Dr. A. Neszmelyi - Central Research Institute for Chemistry of the Hungarian Academy of Sciences

Dr. A.C. Oehlschlager - Simon Fraser University, Canada

*Ms. Connie Oshirio - Lawrence Berkeley Labs

Dr. J.R. Jocelyn Pare - The J.R.J. Pare Establishment for Chemistry Ltd., Canada

Dr. James M. Perry - Worcester Polytechnic Institute, Massachusetts

*Dr. Philip Pfeffer - USDA (Philadelphia)

*Dr. Ned Phillips - University of Florida

*Dr. J.D. Roberts - California Institute of Technology

Dr. Robert Santini - Purdue University

Dr. Norm Stemple - Alcon Labs

Dr. Richard Teeter - Chevron Chemical Co.
We have used CONGEN and the mass spectrum analysis programs to verify the structural assignment of an unknown compound.

*Dr. Babu Venkataraghavan - Lederle Laboratories
(see section on collaborations)

Dr. Stephen Wilson - Indiana University

*Dr. W.T. Wipke - University of California at Santa Cruz
We have worked closely with Prof. Wipke's group for several years on problems of structure representation and manipulation in our complementary areas of computer applications in chemistry.

*Dr. Michael Zippel - Institut fur Biochemie Zentrale Arbeitsgruppe Spectroskopie, Germany
Dr. Zippel used CONGEN to investigate the possible connection with their spectral search system.

d) Industrial Affiliates Program

The high level of interest shown by industrial research laboratories in our programs has always presented us with delicate questions about access to SUMEX-AIM. In the past we have granted access for trials of our programs under the conditions that access is necessarily limited and that the recording mechanisms of our programs be used to ensure that all such trial use be in the public domain. As of April, 1980, we began solicitation of interested industrial organizations to participate in a DENDRAL Project Industrial Affiliates Program. As of May 1, 1981, we have six members. We intend to use this program as a means by which we can offer collaborations with our on-going research to industrial organizations separate from SUMEX-AIM. Although EXODENDRAL accounts to such organizations are used to facilitate communication and sharing of new programs and concepts of interest with the community as a whole, all significant and certainly all proprietary use of our programs will be carried out on their own computational facilities.

e) Program License

We are currently exploring the mechanism of program license to commercial firms as a method for dissemination of well-developed programs, for example CONGEN. This mechanism involves a negotiated agreement between a company and Stanford University for rights to access to and dissemination of identified computer programs. Currently, two companies are negotiating with Stanford. We see this mechanism as serving the function of technology transfer in a very realistic way. We do not, as a research project, have the charter or the resources to do what is essentially final engineering of a program and integration of the program into an existing, larger system. Such "value added" effort is crucial to broad acceptance of a computer-based method. In addition, a participating company would take on the burden of maintenance, documentation and training, freeing our personnel to pursue our research objectives and to bring experimental programs to the level of performance where they, too, can be disseminated by licenses.

B. Interactions with Other SUMEX-AIM Projects

We routinely collaborate with other projects on SUMEX most closely related to our own research. In particular, these collaborations have taken place with the CRYSALIS project, MOLGEN, SECS and have begun with Dr. Carroll Johnson at Oak Ridge.

CRYSALIS is concerned with new approaches to the interpretation of X-ray crystallographic data. X-ray crystallography is another approach to molecular structure elucidation. One of our long-term interests is exploring ways in which CONGEN or GENOA generated structures might be used to guide the search of electron density maps. We are also communicating with Prof. Jon Clardy at Cornell on this problem. It is hoped that having narrowed down the structural possibilities for an unknown using physical and chemical data, the few remaining candidates can be used to guide interpretation of such maps.

Most of the structural problems investigated by MOLGEN involve much larger molecules than the size normally investigated in DENDRAL research. Thus, structural representations involving higher levels of abstraction are of utility in MOLGEN, making our structure manipulation tasks quite different. However, many of the ways in which MOLGEN manipulates its structural representations drew on past experience in DENDRAL in developing algorithms to perform these manipulations.

We collaborate frequently with the SECS project in a number of ways. Although our research efforts are in one sense directed toward opposite ends of work on chemical structures, SECS being devoted to synthesis, DENDRAL being devoted to analysis, the underlying problems of structural manipulation share many common aspects. We have exchanged software where possible, particularly in the area of chemical structure display. We have held several discussions in joint group meetings and at several symposia including the AIM Workshops on common problems, including substructure searching, canonical representations and representation and manipulation of stereochemistry. Persons visiting one laboratory often take the opportunity to visit the other. For example, recent visitors to both laboratories have included Prof. Andre Dreiding, Zurich, Dr. Martin Huber, Basel, and Prof. Robert Carter, Lund.

Dr. Carroll Johnson has collaborated on the CRYSALIS project in the past. More recently he has taken an interest in the use of knowledge-based programs for certain problems in spectral data interpretation. For this reason he is exploring the AGE and EMYCIN systems as frameworks for his program structure, and is involved in discussions with DENDRAL to see where common areas of data interpretation can be identified so that he can draw on our experience and programs. This effort is just beginning at this time; we plan to meet early in May at Stanford to continue discussions.

C. Critique of Resource Management

The SUMEX-AIM environment, including hardware, system software and staff, has proven absolutely ideal for the development and dissemination of DENDRAL programs. The virtual memory operating system has greatly facilitated development of large programs. The emphasis on time-sharing and interactive programs has been essential to us in our development of interactive programs. Our experience with other computer facilities has only emphasized the importance of the SUMEX environment for real-world applications of our programs. To run CONGEN, for example, in a batch computing environment would make no sense whatever because the program (and our other, related programs) is successful in large part because an investigator can closely monitor and control the program as it works toward solution. We have no complaints whatsoever about the computing environment.

We do have, however, significant problems with SUMEX-AIM capacity, both in available computer cycles and on-line file storage. In a sense DENDRAL suffers from its success. The rapid progress made during the last grant period and now continuing into the next period has led to development of many new programs as adjuncts to CONGEN and GENOA and at the same time has inspired many persons in the scientific community to request some form of access to our programs. The net result is that it is often very difficult to carry on at the same time development and collaborations involving applications of our programs to structural problems due to high load average on the system.

The current overcrowding we see on SUMEX creates two major problems for us in the conduct of our research. First, it diminishes productivity as many people compete for the resource; the "time-sharing syndrome" leads to idle, wasted time at the terminal waiting for trivial computations to be completed. Second, the slow response time of the system is an aggravation to an outside investigator who is anxiously trying to solve a structural problem. At some point even the most interested persons will give up, log off the computer and resort to manual methods where possible.

We have taken many steps within our project to try to work around heavy use periods on SUMEX. Our group works a staggered schedule, both in terms of the actual hours worked each day and in terms of what days each week are worked. This results in some problems in intra-group communication, but fortunately the message and other communication systems of SUMEX help alleviate that situation. We try to run all demonstrations on the DEC-2020 to help ease the burden on the dual KI-10 system. We encourage our collaborators to avoid prime-time use of the system when possible.

For these reasons, we strongly support the planned augmentation of the SUMEX-AIM hardware. Any part of our computations which can be shifted to another machine will not only facilitate export of our software but will ease the load on the DEC-10s and make it easier to continue our research. Both will serve to make SUMEX more responsive and our productivity higher.

III. RESEARCH PLANS

A. Project Goals and Plans

Current research efforts were described in highlight form in the first section, Summary of Research Program. In this section we discuss in outline form the major goals of our current grant period (5/1/80 - 4/30/83), with an indication of the progress made to date.

Our goals include the following:

1) Develop SASES (Semi-Automated Structure Elucidation System) as a general system for computer aided structural analysis, utilizing stereochemical structural representations as the fundamental structural description. SASES will represent a computer-based "laboratory" for detailed exploration of structural questions on the computer. It will have as key components the following:

A) Capabilities for interpretation of spectral data which, together with inferences from chemical or other data, would be used for determination of (possibly overlapping) substructures. We have made considerable progress in the areas of mass spectrometry (see References 3, 14) and C-13 NMR spectroscopy (see References 15, 18);

B) The GENOA (structure Generation with Overlapping Atoms) program which will have the capability of exhaustive generation of (topological and stereochemical) structural candidates and include as an essential component the existing CONGEN program. We have developed Version I of GENOA for use by our collaborators (see Reference 17);

C) Capabilities for prediction of spectral (and biological) properties to rank-order candidates on the basis of agreement between predicted and observed properties. Again, we have made considerable progress in mass (see References 3, 7, 8) and C-13 NMR (see References 15, 16, 18) spectroscopy;

2) Develop the GENOA program and integrate it with CONGEN. GENOA will represent the heart of SASES for exploration of structures of unknown compounds, or configurations or conformations of known compounds. GENOA will be a completely general method for construction of structural candidates for an unknown based on redundant, overlapping substructural information, and it will include capabilities for generation of topological and stereochemical (see References 1, 6, 9, 10, 11) isomers;

3) Develop automated approaches to both interpretation and prediction of spectroscopic data, including but not limited to the following spectroscopic techniques:

A) carbon-13 magnetic resonance (13CMR) (see References 15, 16, 18);

B) proton magnetic resonance (1HMR);

C) infrared spectroscopy (IR);

D) mass spectrometry (MS) (see References 3, 7, 8);

E) chiroptical methods including circular dichroism (CD), magnetic circular dichroism (MCD).

The interpretive procedures will yield substructural information, including stereochemical features, which can be used to construct structural candidates using GENOA. We have illustrated this method in recent publications (see References 14, 18). The predictive procedures will be designed to provide approximate but rapid predictions of expected spectroscopic behavior of large numbers of structural candidates, including various conformers of particular structures. Such procedures can be used to rank-order candidates and/or conformers. The predictive procedures will also be designed to provide more detailed predictions of structure/property relationships for known or candidate structures in specific biological applications. These procedures have been illustrated in recent publications (see References 3, 7, 8, 15, 18).

4) Develop a constrained generator of stereoisomers, (see Reference 9) including:

A) design and implement a complete and irredundant generator of possible conformations for a given known, or a candidate for an unknown, structure;

B) provide constraints for the conformation generator so that proposed structures for a known or unknown compound possess only those features allowed by: i) intrinsic structural features such as ring closure and dynamics of the chemical structure; and ii) data sensitive to molecular conformations (e.g., MCD, NMR);

C) integrate the stereochemical developments with the GENOA program as a final, comprehensive solution to the structure generation problem and allow for interface of the program with other methods dependent on atomic coordinates.

5) Promote applications of these new techniques to structural problems of a community of collaborators, including improved methods for structure elucidation and potential new biomedical applications, through resource sharing involving the following methods of access to our facilities and personnel;

A) nationwide computer network access, via the SUMEX-AIM computer resource;

B) exportable versions of programs to specific sites;

C) workshops at Stanford to provide collaborators with access to existing and new developments in computer-assisted structure elucidation in an environment where complex questions of utility and application can be answered directly by our own scientific staff;

D) interface to a commercially available graphics terminal for structural input and output, at as low a cost as possible, so that chemists can draw or visualize structures more simply and intuitively than with our current, teletype-oriented interfaces.

B. Justification and Requirements for Continued SUMEX Use

In previous sections we discussed the relationship between the DENDRAL Project and SUMEX-AIM, methods for using SUMEX-AIM for dissemination of our programs to a broad community of structural chemists and biochemists and a critique of resource management. In this section we wish to emphasize certain factors which were not discussed earlier and to show how our future directions and interests are closely related to the proposed continuation and augmentation of the SUMEX-AIM resource.

As resource-related research, DENDRAL is intimately tied to the SUMEX resource. Our involvement with SUMEX goes far beyond simple use of the facility. We use SUMEX as the focal point for a number of collaborative efforts, for export of our software and for the communication facilities essential to maintaining close contact with remote research groups working with us. SUMEX provides computational facilities for our workshops, where we bring outside investigators to Stanford to use new programs applied to real structural problems. We have already discussed in our critique the difficulties we have, in view of heavy SUMEX load, of maintaining both our research effort and the resource-sharing aspects of our project.

In view of these factors and because <u>SUMEX is our sole source of computational facilities</u>, we took certain steps in our renewal proposal to attempt to alleviate our situation. Specifically, we requested a computer for our own project, a DEC VAX 11/780, to be linked to SUMEX via ETHERNET. This computer was meant to help offload some of the computational burden DENDRAL places on SUMEX, to provide a facility for production use of our programs by our collaborators and to represent a model for the type of low-cost, scientific computer available in the future to many investigators who could then run our programs in their own laboratories.

Our request for the VAX was turned down with specific comments made that SUMEX facilities should be used to support development of new programs and to the extent possible, encourage <u>preliminary</u> production use of our programs by outside persons. In our opinion this view is somewhat shortsighted, because SUMEX is currently overloaded to the extent that even development is impeded. In addition, our current situation leaves no room for the computational burden created by some of our collaborators who need considerably more than "preliminary" access because they have no access to a computer suitable for running our programs.

For these reasons, we strongly support the effort of SUMEX to acquire a VAX and other small machines in future years. Although we realize that such machines will have to be shared among the SUMEX-AIM community as a whole, the augmentation of the resource would go a significant way to meeting the computational requirements of our project and provide a variety of systems of potential use for future export of our programs.

C. Needs and Plans for Other Computing Resources

For several years now we have directed some attention toward alternative computing resources which could be used to support all "production" use of our programs, i.e., all applications designed to use the programs to solve real problems. Although this would have the severe disadvantage of separating our research effort from many of the applications, it has been our hope that emerging technology in networking would enable us to keep in reasonably close contact with another resource. Two resources have emerged as candidates for systems where our programs can be accessed and used in problem-solving. Unfortunately, neither has so far proven feasible for several reasons (mentioned below). At this time we cannot determine if the problems will be resolved. Until such time, we will remain completely dependent on SUMEX for <u>all</u> our computational needs.

One alternative resource is the NIH/EPA Chemical Information System. For more than three years we have been working with them to obtain sufficient contract money to provide a version of CONGEN integrated into that system. The concept and the funds were approved but a contract has never been issued due to administrative problems at the EPA. Although there have been some developments recently, we still have no firm idea on when such a contract will be issued. If this effort is successful, then we can encourage persons who desire access to our programs to consider using the NIH/EPA system.

A second alternative is the National Resource for Computation in Chemistry (NRCC). This Resource has recently had its funding terminated.

We are now pursuing an alternative discussed previously, that of arranging license agreements with private industry for dissemination of our software. This will likely be the focus of our future efforts to disseminate programs to those researchers who merely wish to use them rather than work together with us in collaborative arrangements to develop more powerful programs.

D. Recommendations for Future Resource and Community Development

We have discussed previously our recommendation for the hardware augmentation, particularly with regards to purchase of small machines to facilitate future export. We also have increasing need for more file storage on-line. This is a result of building large data bases as part of our research in spectral interpretation. For the time being we are working with experimental programs and small data bases. As time progresses, however, these data bases will grow rapidly as our group and a number of our collaborators add additional structures and associated spectral data.

Another capability which is of increasing importance to our own work is access to low-cost graphics systems. Our programs will develop increasing dependence on graphics for visualization of three-dimensional molecular structures. Scientists desiring access to our programs will need a graphics terminal for optimum use of our systems. Currently available vector displays are simply too expensive for the average investigator. The emerging technology of low-cost raster display systems offers a more promising possibility. However, no currently available machine has the required capabilities for under $10,000, and this is an area where machines like the Alto hold more promise. SUMEX could perhaps initiate an effort to obtain a system which has the hardware necessary for frame-based display. Such a system allows rotation of three-dimensional objects in a way which permits visualization of the actual shape of the object.

II.A.1.4 EXPEX Project

EXPEX - Expert Explanation Project

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University

Michael R. Genesereth, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

EXPEX is a new Stanford project that joined the AIM community only a few months before this report was prepared. We therefore have little to report in the way of progress other than our background work that led to a recently funded proposal and the initiation of this new research. The major thrust of the work is the development of powerful representational schemes to facilitate knowledge acquisition and explanation. This includes not only the study of fundamental representational formalisms but also the encoding of various types of knowledge, such as causal information and user models.

We believe that the productivity of basic computer science research tends to be heightened by experiments that deal with significant real world problem domains. Challenges drawn from chemistry, medicine, and molecular biology have introduced additional complexity to expert systems work at Stanford, but have simultaneously forced system developers to respond to pragmatic constraints and user demands that have had a significant impact on the basic AI techniques selected or developed. Thus, we believe that creative investigation into symbolic reasoning techniques is facilitated by working in real world settings where the application forces us to avoid oversimplification. The explanation portion of the research effort will therefore deal with a medical domain (endocrinology) and be undertaken on SUMEX, whereas the knowledge acquisition portion will deal with nonmedical topics and use other computing resources at Stanford. Our report here will only describe EXPEX, the research on expert medical explanation, but it should be understood that this is actually only one part of a coordinated effort to tie together research in knowledge acquisition and explanation through common representation techniques.

B. Medical Relevance and Collaboration

Our interest in explanation derives from the insights we gained in developing explanatory capabilities for the MYCIN system. In the case of MYCIN and its descendents, we have been able to generate intelligible explanations by taking advantage of its rule-based representation scheme. Rules can be translated into English for display to a user, and their interactions can also be explicitly demonstrated. By adding mechanisms for understanding questions expressed in simple English, we were able to create an interactive system that allowed physicians to convince themselves that they agreed with the basis for the program's recommendations. The limitations of the explanations generated in this way have become increasingly obvious, however, and have led to improved characterization of the kinds of explanation capabilities that must be developed if clinical consultation systems are to be accepted by physicians.

C. Highlights of Research Progress

MYCIN's explanation capabilities were generalized in EMYCIN and thus became available for any EMYCIN consultation system. They were further modified and utilized in both TEIRESIAS and GUIDON. Although we had experienced problems using MYCIN's rules for certain kinds of explanations (e.g., control mechanisms that were sometimes encoded in rules, or algorithmic knowledge such as the mechanisms for drug selection), it was in the setting of GUIDON that the inadequacies of MYCIN's approach became most apparent. Consider, for example, a simple MYCIN rule such as:

If: the patient is less than 8 years old
Then: don't give tetracycline

This rule is adequate for MYCIN's decision making task, and would be understood by most physicians if it were used in an explanation, but it is obvious to a casual observer that it contains a giant leap in logic. It is accordingly difficult for GUIDON to teach this rule to a novice medical student because the underlying pathophysiologic knowledge (i.e., that tetracycline is deposited in the developing bone and teeth of youngsters, weakening the former and disfiguring the latter) is not explicitly represented in MYCIN. Examples such as this one emphasize that a variety of knowledge forms are necessary if an intelligent system is to customize its explanations to the individual who is using the program. Underlying structural and causal relationships are generally required in addition to the high level judgmental rules that had contained almost all of the domain knowledge in MYCIN and the other EMYCIN systems.

We therefore began to study in more detail the nature of the explanatory process, and were surprised to find that there are very few writers who have addressed the issues which now interest us. Perhaps the most relevant studies of explanation are in the education literature; several educators have tried to identify the characteristics of explanations which make individuals good teachers. These analyses are accordingly relevant to computer-aided instruction work, such as GUIDON, although issues of automation are not addressed explicitly. On the other hand, they seem less pertinent when applied to the "persuasiveness" of a justification offered by a scientist to a colleague.

A weekly seminar group has been formed to discuss knowledge representation and to analyze the characteristics of good explanations. We have often kept our discussions separate from computer science issues, concentrating instead on the psychology of explanation and planning to return eventually to consider ways in which our developing theory might be implemented in knowledge-based consultation systems. Although there are several subproblems, it was agreed that the problems of explanation can generally be divided into four categories: (1) modeling the knowledge of the system user; (2) selecting a response strategy; (3) modeling contextual information regarding the interaction; and (4) understanding the question. One goal of our new work, then, is to build an explanation system which explicitly addresses these topics.

Modeling the User's Knowledge

GUIDON and other ICAI systems have recognized the need to keep an internal model of the student, i.e., what he has shown he knows, what you have already told him, and perhaps a record of where his greatest weaknesses lie. Similarly, it is clear than an expert human consultant customizes his explanations so that they can be understood by the person requesting the consultation (and are thereby maximally convincing). The expert starts with certain suppositions about his client's knowledge (e.g., a teacher may presume his student is starting from scratch, but a cardiologist will assume that another physician requesting advice probably already knows a fair amount of cardiology). The default presumption is modulated, however, as the interaction proceeds and the client demonstrates his strengths or weaknesses.

We have recently begun some experiments to investigate methods for encoding, along with the domain knowledge, the complexity and importance of that knowledge. These two parameters seem to be independently important in deciding whether to include a given reasoning step in an explanation. "Key" points (i.e., those that are highly important) probably should be mentioned even if they are not complex and are likely to be known to the user. On the other hand, less important but complex items probably need not be mentioned unless an expert user is really pressing for details of a decision pathway. Thus, static measures of complexity and importance can be compared with user descriptors that are initially assigned by default (depending upon the status of the user, e.g., expert vs. student), but are later altered dynamically in response to the course of the dialog and what it has revealed about the user's background knowledge.

These ideas have been encoded in a small computer program which uses a limited knowledge base of rules and associations from the domains of pharyngitis (sore throats) and calcium metabolism. We have experimented with a semantic network representation in which the nodes are values of attributes and rules are only one form of link between nodes. All nodes and rules have complexity and importance measures associated with them. An "opinion" regarding a specific patient can be represented as a subset of the nodes in the network, plus the links between them that account for how it has been determined which nodes are active. In this setting, a question tends to ask how it has been determined that a given node is active for a given patient. The appropriate explanation could be very complex if an effort were made to explain every link leading from data observations to the node descriptor in question. A customized explanation is therefore generated based on three variables which can be dynamically manipulated by the program: (1) the focus of the dialog (e.g., broad-based vs. localized),

(2) the expertise of the user, and (3) the degree of generality which is appropriate. These three variables are clearly not independent, and we are experimenting with ways to have their values manipulated in a reasonable fashion as the dialog proceeds.

This early effort has provided the basis for further discussions in our seminar group as we have attempted to arrive at an optimal representation for the research to follow. We have been fortunate to enlist the collaboration of an endocrinologist at Stanford, Dr. Larry Crapo, who is eager to work with us on building an endocrinology knowledge base. It is likely that we will select the pathophysiology of calcium disorders as a small focused area to study. This domain is appealing for computer-based representation because the relationships are well-understood and there are some challenging problems.of feedback homeostasis that will need to be represented. In the years ahead, we will encode this knowledge base in detail and begin experiments on the generation of explanations using the kinds of techniques outlined above.

Selecting A Response Strategy

Our explanation efforts to date have tended to be simple reiterations of individual reasoning steps, but it is clear that experts and teachers use several alternate strategies for conveying their ideas or key facts. Many of these techniques draw upon common sense world knowledge (e.g., analogies with familiar concepts outside the domain), but we have thus far failed to capitalize on these teaching strategies in our work. Thus another goal of the work that lies ahead will be to develop structures for drawing parallels or otherwise representing the strategies used by good "explainers."

Modeling Contextual Information Regarding the Interaction

We have already mentioned some of the ways in which contextual information may be useful in determining the best way to answer a question. For example, a more accurate model of the user's knowledge can be developed over time, and the extent to which a given conversation is focused on a particular local topic can be assessed. Note that we are emphasizing here issues other than those related to natural language understanding; computational linguists also often cite the need to record contextual dialog information in order to handle problems such as anaphora. An understanding of the "flow" of a dialog is also important in understanding the meaning of subsequent questions, as we discuss below.

Understanding The Question

This issue interfaces with the problem of natural language understanding, but we view it in a somewhat different light. We emphasize instead the ways in which the model of the user and contextual information may allow us to disambiguate questions. To draw from a medical example that we have frequently discussed, consider the following scenario. A reasoning program for pharyngitis diagnosis and management has just diagnosed strep throat and recommended penicillin and the user asks the question "Why would you give penicillin?" In the most obvious case, one might imagine a response that itemizes the risks of streptococcal infections and the reasons for treating early with penicillin. Similarly, one might expect a more detailed response for a student and a quick summary for a physician using the system.

However, an alternate interpretation is that EVERY physician knows the theoretical reasons for giving penicillin in strep pharyngitis, and that if the user is a physician and is asking the question then he must be asking something different than the simple informational question. In this case the query might be interpreted as a challenge (one that might have been conveyed by tone of voice if it had been asked of a human consultant). Apparently the user has reason to doubt that penicillin was the appropriate agent in this case, or thinks that no drug was required. Other background information and contextual knowledge should also help, and an intelligent program might thereby answer the question in a given case in any of the following ways:

"Because the patient has pre-existing rheumatic heart disease."

"Because I doubt that he is allergic to penicillin, even though he reported that he is."

"Because he is unreliable and I am afraid I will not be able to reach him to call him back if his strep culture comes back positive."

"Because I tend to treat conservatively and give penicillin for strep throat even though I know there hasn't been a case of rheumatic heart disease in California in over 10 years."

Note how different these kinds of explanations are from the simple justification that a program such as MYCIN might have given:

"Because streptococcal pharyngitis may be followed by rheumatic myocarditis or glomerulonephritis, mediated by immune complexes, and I can prevent this complication by giving penicillin (to which streptococci are uniformly sensitive)."

The ideal intelligent assistant should be able to determine from knowledge of the user, the domain, the individual case, and the context of the dialog, which of the preceding responses is most appropriate. We will attempt to identify methods for giving our program this kind of capability.

D. Publications Since January 1980

Wallis, J.W. and Shortliffe, E.H. Explanatory power for expert systems: studies in the representation of causal relationships for medical consultations. Internal working memo, Heuristic Programming Project, Stanford University, May 1981.

E. Funding Support

Grant Title: "The Development of Representation Methods to Facilitate Knowledge Acquisition and Exposition in Expert Systems"
Principal Investigator: Edward H. Shortliffe
Agency: Office of Naval Research
ID Number: NR 049-479
Term: January 1981 to December 1983
Total award: $456,622
Current award (1981): $140,825

II. INTERACTION WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

We are only beginning program development at this time, and have therefore had no opportunity to share our results with others as of yet.

B. Sharing and Interaction with Other SUMEX-AIM Projects

We anticipate frequent ongoing interactions with other SUMEX-AIM research efforts because the development of explanation techniques is a pertinent research issue for all expert systems work in medicine. Bill Clancey's work on GUIDON is addressing many of the same issues and we expect frequent opportunities for interchange.

C. Critique of Resource Management

Although we have not yet placed significant demands on SUMEX management, our previous experience working with Tom Rindfleisch and his staff would suggest that this new project will receive the same kind of laudatory service for which SUMEX has become known.

III. RESEARCH PLANS (6/81-12/83)

A. Project Goals and Plans

We intend to investigate optimal techniques for the computer-based representation of expert knowledge. Because we have come to recognize the limitations of any single representation technique taken alone, a principal objective will be to merge alternate approaches, augmented with new capabilities. We will, in turn, evaluate the effectiveness of the new representation scheme by focusing on issues of both knowledge acquisition and explanation. Furthermore, we will perform these experiments in two expert domains, medical reasoning (EXPEX) and computer circuitry debugging (DART). These areas were selected because we have local expertise in each, but also because they are sufficiently different from one another that they will force us to ensure the generality of the techniques we are developing. Utilizing a single representation scheme for all aspects of the work will also encourage generality of the developed techniques because this will force us to avoid concentrating on either the input (knowledge acquisition) or output (explanation) functions alone.

Initially we shall concentrate on defining the knowledge representation scheme to be used. Although modifications will of course be necessary in response to additional lessons learned thereafter, we expect to reach an early consensus on the major components of the internal representation we will be using.

Subsequently our efforts will divide into two components, each of which will utilize the representation scheme devised in the initial period. EXPEX will concentrate on manually constructing a knowledge base regarding calcium metabolism and pathophysiology, whereas the DART effort will be concentrating on knowledge acquisition for their non-medical domain. In the EXPEX work, the construction of the clinical knowledge base will have created an environment for the development of the explanation capabilities which are the second thrust of our work. Drawing on the early work of Jerry Wallis, described in the memo referenced in the publications section of this report, we will next construct an expository system that uses the endocrinology knowledge base in order to generate interactive explanations.

Ultimately we hope to perform experiments using both the knowledge acquisition and explanation tools that will have been developed in the two separate domains. One task will be to see if we can develop from scratch the endocrinology knowledge base that will have been hand-coded for the EXPEX effort. Because this knowledge will have previously been encoded manually, we will have a well-defined model of the form the knowledge should take as it is acquired interactively using the new system building tools developed in the electronics environment. At the same time, some of us will be developing experiments to test the validity and effectiveness of the explanation tools that were developed for EXPEX. An excellent test of the generality we have been seeking will be to build the circuit debugging knowledge base using the new knowledge acquisition tools, and then to demonstrate the utility of the explanation routines for exposition of the knowledge in this new domain.

Our ultimate goal, then, is to have developed a unified system of knowledge representation that facilitates both system building, through interactive knowledge acquisition, and explanation, through interactive responses to knowledge base queries. It should be emphasized that throughout our work the focus will be on the underlying representation issues and not on polished text generation nor natural language understanding.

II.A.1.5 MOLGEN Project

MOLGEN - A Computer Science Application to Molecular Biology

Prof. E. Feigenbaum and Dr. P. Friedland
Department of Computer Science
Stanford University

Assoc. Prof. D. Brutlag
Department of Biochemistry
Stanford University

Assoc. Prof. L. Kedes
Department of Medicine
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. We plan to expand and improve these systems and build new ones to meet the rapidly growing needs of the domain of recombinant DNA technology. We do this with the view of including the widest possible national user community through the facilities available on the SUMEX-AIM computer resource.

It is only within the last few years that the domain of molecular biology has needed automated methods for experimental assistance. The advent of rapid DNA cloning and sequencing methods has had an explosive effect on the amount of data that can be most readily represented and analyzed by computer. Moreover we have already reached a point where progress in the analysis of the information in DNA sequences is being limited by the combinatorics of the various types of analytical comparison methods available. The application of judicious rules for the detection of profitable directions of analysis and for pruning those which obviously lack merit will have an autocatalytic effect on this field in the immediate future.

The MOLGEN project has continuing computer science goals of exploring issues of knowledge representation, problem-solving, and planning within a real and complex domain. The project operates in a framework of collaboration between the Heuristic Programming Project (HPP) in the Computer Science Department and various domain experts in the departments of Biochemistry, Medicine, and Genetics. It draws from the experience of several other projects in the HPP which deal with applications of artificial intelligence to medicine, organic chemistry, and engineering.

During the next three years of MOLGEN research we intend to begin a transition from being primarily a computer science research project to being an interdisciplinary project with a strong applications focus. The tools that we have already developed will be improved to the point where they make a significant contribution to both research and engineering in the domain of molecular biology.

B. Medical Relevance and Collaboration

The field of molecular biology is nearing the point where the results of current research will have immediate and important application to the pharmaceutical and chemical industries. Recombinant DNA technology has already demonstrated the possibility of harnessing bacteria to produce nearly limitless amounts of such drugs as insulin and somatostatin. Governmental reports estimate that there are more than 200 new and established industrial firms already undertaking product development using these new genetic tools.

The programs being developed in the MOLGEN project have already proven useful and important to a considerable number of molecular biologists. Currently several dozen researchers in various laboratories at Stanford (Prof. Paul Berg's, Prof. Stanley Cohen's, Prof. Laurence Kedes', Prof. Douglas Brutlag's, Prof. Henry Kaplan's, and Prof. Douglas Wallace's) and over 300 others throughout the country are using MOLGEN programs over the SUMEX-AIM facility. We have exported some of our programs to users outside the range of our computer network (University of Geneva [Switzerland], Imperial Cancer Research Fund [England], and European Molecular Biology Institute [Heidelberg] are examples).

C. Highlights of Research Progress

<u>Accomplishments</u>:

The current year has seen the completion of what might be considered the first phase of the MOLGEN project. This section will summarize the major accomplishments of that first phase.

1. Representation Research

The domain of molecular biology has proven a fruitful testbed in the development of a flexible software package, the Unit System, for symbolic representation of knowledge. The package is already in use by a variety of research projects both within the Heuristic Programming Project at Stanford and at other institutions. It provides for acquisition and storage of many different types of knowledge, ranging from simple declarative types like integers and strings to complex declarative types like nucleic acid restriction maps to procedural types like a rule language in a subset of English.

A major effort has been made in the past year to take the Unit System and, observing the experience of its many scientific users, improve and enhance those features which are most important. This has resulted in a speed improvement of at last two orders of magnitude for the most used functions in the representation system. The MOLGEN project has provided a unique laboratory for the conversion of a theoretically-based knowledge base system into a practical package for knowledge acquisition and manipulation. This is because of the active daily use of the system for real laboratory problems. The Unit System has become what may be considered the first "second-generation" knowledge representation package.

We have concentrated on representation methods that are unique to molecular biology, particularly, convenient methods for storing information about nucleic acid sequences and maps of those sequences, as well as an English- like language for manipulating that information. This language has allowed and encouraged the molecular biologist members of the MOLGEN project to become their own "programmers" without having to worry about the underlying representation structures of their knowledge bases. For example, the phrase:

JOIN SV40 FROM FIRST ECOR1 SITE TO FIRST BAMH1 SITE TO PBR322 FROM 250 TO 3000 INTO NEWSEQUENCE

will perform the clearly indicated operation.

2. Planning Research

The problem of designing laboratory experiments in molecular biology has been fundamental to MOLGEN research. The work has been split into two major subparts, each resulting in a doctoral thesis in computer science. The two systems, developed by Peter Friedland and Mark Stefik, produce reasonable experiment designs on test problems suggested by laboratory scientists.

The majority of MOLGEN planning work has awaited the successful completion of the latest phase of representation improvement described above. Dr. Rene Bach, a post-doctoral fellow, has recently completed a DNA sequencing experiment adviser, built entirely with the Unit System and associated procedural description language. This system provides guidance in developing a sequencing protocol given a partial initial restriction map for a new nucleic acid sequence. Other members of the MOLGEN group have begun exploring the representations needed with the Unit System to provide transparent descriptions (i.e., invisible to the non-computer scientist user) of growing experimental plans, so that plans may be treated in a manner identical with all other types of knowledge.

3. Knowledge Base Construction

Over six man-years have now been spent in constructing knowledge bases for various fields of molecular biology. Professors Kedes and Brutlag, Dr. Bach, and several students have cooperated on one knowledge base which is expert in restriction enzyme methodology and another which is competent for a wide range of general laboratory techniques. Professor Kedes has worked on a knowledge base for his own interests in gene structure. Professor Brutlag is concentrating on a knowledge base for satellite DNAs. Dr. Bach has built a knowledge base expert in sequencing methods for his sequencing advisor (see above). Professor Sninsky and Dr.

Bach have collaborated in a knowledge base for expression vectors. Professor Sninsky and Dr. Abarbanel have collaborated on the beginnings of a protein knowledge base to explore methods for predicting secondary protein structure from primary amino acid sequence. Several researchers in Professor Hogness's laboratory in the Department of Biochemistry have begun to build a knowledge base for storing information about many different lambda vector clones. Finally, several scientists in Professor Kaplan's Cancer Biology Research Laboratory have started to explore a monoclonal antibody knowledge base.

The knowledge bases developed under the MOLGEN grant have begun to find their way into the daily laboratory practice of many of the scientists associated with the project. They have a provided a mechanism for managing the explosive growth of data and strategies in many areas of molecular biology without the necessity of building special purpose systems for each area. Also, the expert scientists themselves have been able to design and build their own systems, avoiding the time and reliability problem of a knowledge base passing through the filter of a computer scientist intermediary. The knowledge bases have served as "intelligent encyclopedias," as simulation systems, and as training vehicles.

It should also be noted that the Unit System allows for the easy transfer of knowledge from one knowledge base to another, and indeed the various expert molecular biologists have freely shared information as they work on related knowledge bases.

4. Other Applications of Symbolic Computation to Molecular Biology

MOLGEN programmers have spent full time enhancing existing MOLGEN applications programs and developing new systems. The SEQ program is a general purpose nucleic acid sequence analysis system. It provides a range of functions including translation, lexicography, regions of richness, restriction mapping, general string search, and intra- and inter-sequence homologies, symmetries, and dyad symmetries. The program calculates statistical probabilities for homologies, symmetries, and dyad symmetries, and also determines approximate free energy contribution of dyad symmetry structures in RNA. SEQ is highly interactive and provides many built-in explanation and help facilities.

MAP is a program which determines restriction sites from enzymatic digest data. A recent collaboration with Dr. Pearson has been exploring ways to combine the ideas from MAP and Dr. Pearson's system for solving similar problems.

Work has begun on two major new applications programs: GEL, a a system which provides bookkeeping and overlap determining assistance for "shotgun" sequencing experiments; and AA a program which provides most of the functions of SEQ for amino acid sequences.

During the first year of the current MOLGEN grant, we have provided guest access to the SEQ and MAP programs to the national academic community through the facilities of the SUMEX-AIM computer system. This has meant free, dial-up access from almost anywhere in the United States. Over 300 researchers at over 80 institutions have used the service. It has been so popular that the SUMEX-AIM Executive Committee has found it necessary to limit the service to at most two simultaneous users at any one time. The facilities provided to MOLGEN guests has been very limited, with a single directory and 250 disk pages serving the entire national community. Despite this, a wide variety of interesting research has been done, and MOLGEN is most grateful to the SUMEX-AIM staff for their kind assistance.

<u>Research in Progress</u>:

The remainder of the current grant period will be spent on the further development of the tools that have been constructed for experiment design and sequence analysis and on expansion and improvement of the knowledge base. This section details those research plans.

1. Representation

The Unit System is now at a stage of general utility. The MOLGEN group will continue to enhance and improve both the underlying representation methods and the user interface. A further order of magnitude improvement in the speed of common operations will be achieved during the next year. There will be at least a doubling of the genetic-specific vocabulary of the procedural description language. Mechanisms will be developed to allow the Unit System to communicate effectively with all of the major MOLGEN applications programs, particularly SEQ, MAP, and GEL.

We also anticipate a major effort, beginning in the summer of 1981, to adapt the Unit System to run on one of the newly available personal scientific work-station computers--most likely the Xerox Dolphin. This will provide a qualitative improvement in user interaction because of the large bit-map display and graphics capabilities, and the ability to free individual knowledge base builders from time-sharing system load.

2. Knowledge Base Development and Planning

Planning work will proceed by moving into at least one new sub-domain and by synthesizing the research of previous years. Dr. Bach will construct a knowledge-based system for the planning of cloning experiments using the experience he has gained from the simpler task of designing sequencing experiments. The previous MOLGEN planning research has produced two major ideas. One, the "skeletal-plan" approach of Dr. Friedland, involves the selection and refinement of planning strategies, which may range from abstract to specific, provided as part of an expert knowledge base. This idea resulted from a study of the way in which molecular biologists design experiments. The other idea was the "constraint-posting" method of Dr. Mark Stefik, which concentrates on the extensive evaluation of the constraints introduced by each step of a growing plan in order to guide the selection of the next step of the plan. The MOLGEN group will undertake the project of synthesizing these two ideas, combining the practical efficiency of skeletal-plan refinement with the general power of constraint posting.

We will also begin work on plan verification and optimization systems within the same general framework; i.e. systems to check proposed plans for suitability and to improve such plans. We hope to begin to extend this to a plan debugging system; one which interacts with an experimenter to determine where an experiment failed and how to correct the problem.

Knowledge base construction will proceed, with each MOLGEN collaborator exploring the particular sub-domain most interesting to him. Professor Maxam will begin building a knowledge base, most likely in the area of detailed nucleic acid structure, this summer. This will lead to joint work with Dr. Friedland on the application of knowledge-based methods to the study of the mechanism of gene regulation. Several researchers at the Imperial Cancer Research Fund in London, England will join the effort starting in the early autumn of 1981. .

In particular, we expect the area of cloning methodology to occupy a large portion of Dr. Bach's time in knowledge base construction. We will draw from the wide and varied expertise within the Stanford community during this effort.

3. Applications Systems

Work will continue on improving and enhancing SEQ and MAP. The GEL and AA programs will begin operation. The currently available tools within the Unit System allow the molecular biologists themselves to design and construct special-purpose applications systems within their knowledge bases. Professors Brutlag and Kedes have already built over a dozen such "programs" for construction of specific restriction maps and fragment tables and the simulation of recombinant DNA operations. As the domain experts decide that these systems are generally useful, they will be optimized and packaged into stand-alone programs.

The MOLGEN group will continue to cooperate with any national efforts to develop a sequence analysis data bank and facility for the academic community. We hope our current collaborative service activities on SUMEX-AIM will serve as the prototype for a larger and more comprehensive national facility.

D. Publications

Feitelson J., Stefik M.J., A Case Study of the Reasoning in a Genetics Experiment, Heuristic Programming Project Report HPP-77-18 (Working Paper) (May 1977)

Friedland P., Knowledge-Based Experiment Design in Molecular Genetics, Proceedings Sixth International Joint Conference on Artificial Intelligence, 285-287 (August 1979)

Friedland P., Knowledge-Based Experiment Design in Molecular Genetics, Ph.D. Thesis, Stanford CS Report CS79-760 (December 1979)

Martin N., Friedland P., King J., Stefik M.J., Knowledge Base Management for Experiment Planning in Molecular Genetics, Fifth International Joint Conference on Artificial Intelligence. 882-887 (August 1977)

Stefik M., Friedland P., Machine Inference for Molecular Genetics: Methods and Applications, Proceedings of the National Computer Conference, (June 1978)

Stefik M.J., Martin N., A Review of Knowledge Based Problem Solving As a Basis for a Genetics Experiment Designing System, Stanford Computer Science Department Report STAN-CS-77-596. (March 1977)

Stefik M., Inferring DNA Structures From Segmentation Data: A Case Study, Artificial Intelligence 11, 85-114 (December 1977)

Stefik, M., An Examination of a Frame-Structured Representation System, Proceedings Sixth International Joint Conference on Artificial Intelligence, 844-852 (August 1979)

Stefik, M., Planning with Constraints, Ph.D. Thesis, Stanford CS Report CS80-784 (March 1980)

E. Funding Support

The MOLGEN grant is titled: MOLGEN: A Computer Science Application to Molecular Biology. It is NSF Grant ECS-8016247. Current Principal Investigators are Edward A. Feigenbaum and Bruce G. Buchanan, Professors of Computer Science, Laurence H. Kedes, Investigator, Howard Hughes Medical Institute and Associate Professor of Medicine, and Douglas L. Brutlag, Associate Professor of Biochemistry. MOLGEN is currently funded from 10/80 to 9/81 at $146,582 including indirect costs as the first year of a three year renewal.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

Until this year, all system development had taken place on the SUMEX-AIM facility. The facility has not only provided excellent support for our programming efforts but has served as a major communication link among members of the project. Systems available on SUMEX-AIM such as INTERLISP, TV-EDIT, and BULLETIN BOARD have made possible the project's programming, documentation and communication efforts. The interactive environment of the facility is especially important in this type of project development.

Unfortunately, the computing environment at SUMEX has suffered in the recent past from heavy demands on cycle time creating serious real-time delays for programmers and knowledge-base building especially. The Units Editor is especially sensitive because of its relatively large demands on cpu-resources. Accordingly, a significant fraction of the MOLGEN group activity has been transferred to the SCORE computer in the Department of Computer Science at Stanford. When SUMEX hardware is updated, we anticipate that its response time will improve and the MOLGEN computing will return full time to SUMEX. It is clear, however, that the MOLGEN project continues to thrive and prosper because of the computing environment only available at SUMEX: the interactive environment including instantaneous communications among collaborators who are physically distant (even on the Stanford campus), and especially the unique telecommunications facilities that have allowed the development of the GENET community with its access to MOLGEN applications tools are two clear examples.

We have taken advantage of the collective expertise on medically-oriented knowledge-based systems of the other SUMEX-AIM projects. In addition to especially close ties with other projects at Stanford, we have greatly benefitted by interaction with other projects at yearly meetings and through exchange of working papers and ideas over the system.

The ability for instant communication with a large number of experts in this field has been a determining factor in the success of the MOLGEN project. It has made possible the near instantaneous dissemination of MOLGEN systems to a host of experimental users in laboratories across the country. The wide-ranging input from these users has greatly improved the general utility of our project.

We find it very difficult to find fault with any aspect of the SUMEX resource management. It has made it easy for us to expand our user group, to give demonstrations (through the 20/20 adjunct system), and to disseminate software to non-SUMEX users overseas.

III. RESEARCH PLANS

A. Justification and Requirements for Continued SUMEX Use

The MOLGEN project depends heavily on the SUMEX facility. We have already developed several useful tools on the facility and are continuing research toward applying the methods of artificial intelligence to the field of molecular biology. The community of potential users is growing nearly exponentially as researchers from most of the bio-medical fields become interested in the technology of recombinant DNA. We believe the MOLGEN work is already important to this growing community and will continue to be important. The evidence for this is an already large list of pilot exo-MOLGEN users on SUMEX.

SUMEX is currently having difficulty meeting the research needs of the MOLGEN project adequately. We expect to need more file space as our knowledge bases grow; perhaps an additional 5000 disk blocks in the next few years for that work. Our real difficulties will come in the applications testing of MOLGEN tools. We support with great enthusiasm the acquisition of satellite computers for technology transfer and hope that the SUMEX staff continue to develop and support these systems. One of the oft-mentioned problems of artificial intelligence research is exactly the problem of taking prototypical systems and applying them to real problems. SUMEX gives the MOLGEN project a chance to conquer that problem and potentially supply scientific computing resources to a national audience of bio-medical research scientists.

II.A.1.6 MYCIN Projects Group

MYCIN Projects

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University

Bruce G. Buchanan, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The MYCIN Projects are a set of related research programs, each devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The name was derived from our first system, the MYCIN program. That research has now given way to three active sub-projects (EMYCIN, GUIDON, and ONCOCIN), each of which is discussed in the sections that appear below. The key issue for all sub-projects has been to develop programs that can provide advice similar in quality to that given by human experts, and to develop systems that are easy to use and acceptable to physicians and medical students.

The success of the original MYCIN infectious disease consultation program has led us to try to generalize and expand the methods employed in that program to a number of ends:

(1) to develop consultation systems for other domains (our generalized system-building tool is known as "Essential MYCIN", or EMYCIN, and has been applied in several new areas; ONCOCIN is our newest consultation system and was inspired by EMYCIN although it is actually an entirely new program);

(2) to explore other uses of the MYCIN knowledge base (our tutoring system, GUIDON, uses the infectious disease knowledge in MYCIN to teach medical students about diagnosis and management of infections);

(3) to continue to improve the interactive process, both for the developer of a knowledge-based system, and for the user of such a system (both EMYCIN and ONCOCIN have stressed simplified techniques for interacting with a knowledge base and entering data); and

(4) to experiment with alternate techniques for knowledge representation, recognizing that the pure production rule method used in MYCIN was inadequate at times and frequently led to confusion regarding the separation of strategy or control

processes from domain knowledge (ONCOCIN uses production rules as only one of several knowledge representation techniques, and the work on GUIDON has led to a more robust revised version of MYCIN known as NEOMYCIN).

B. Medical Relevance and Collaboration

By utilizing our EMYCIN system to collaborate on building the PUFF program, we learned that it is possible in a short period of time to develop a clinically useful consultation system using the domain-independent parts of MYCIN. EMYCIN has since been applied in a number of additional medical domains. With each successive application we learn more about the representation of medical knowledge and the scope and limitations of the production rule formalism used in EMYCIN. For example, it has become clear that "shallow" rules relating signs and symptoms to diagnoses through a few intermediate concepts can be sufficient for high performance in medical diagnosis. On the other hand, such shallow rules are not always sufficient for teaching medical students because they lack the deeper causal links needed for justifying and remembering the more shallow associations (see GUIDON discussion below).

Although EMYCIN was not used to build our new ONCOCIN program, the lessons learned in building prior production rule systems have allowed us to create a large oncology protocol management system much more rapidly than was the case when we started to build MYCIN. We are introducing ONCOCIN for use by Stanford oncologists in the Spring of 1981. This would not have been possible, of course, without the active collaboration of Stanford oncologists who helped with the construction of the knowledge base and also kept project computer scientists aware of the psychological and logistical issues related to the operation of a busy outpatient clinic.

In addition, there is a growing realization that medical knowledge, originally codified for the purpose of computer-based consultations, may be utilized in additional ways that are medically relevant. Using the knowledge to teach medical students is perhaps foremost among these, and GUIDON continues to focus on methods for augmenting clinical knowledge in order to facilitate its use in a tutorial setting. A particularly exciting aspect of this work is the insight that has been gained regarding the need to structure knowledge differently, and in more detail, when it is being used for different purposes (e.g., teaching as opposed to clinical decision making). This aspect of the GUIDON research has led to the development of a modified version of MYCIN, NEOMYCIN, which is an evolving computational model of medical diagnostic reasoning that we hope will enable us to better understand and teach diagnosis to students.

C. Highlights of Research Progress

1. Accomplishments This Past Year

EMYCIN

In the last year, substantial research efforts were completed and described in publications. First, a complete EMYCIN system, including a rule compiler and interactive rule editor, was packaged and documented in Bill van Melle's thesis. Second, an investigation into mixing production rules and frames as a representation of medical knowledge was completed and documented by Jan Aikins. Third, a redesign of the EMYCIN system to include more of the structural and strategic knowledge needed for tutoring was completed and documented (see description of NEOMYCIN in the GUIDON sections below). In addition, the LISP code and the EMYCIN Manual were both improved considerably in response to suggestions from outside users.

The complete EMYCIN system includes: a) a rule interpreter, b) an explanation facility, c) an abbreviated rule language and editor for rule input, d) a debugging package, and e) a rule compiler. Some new developments in these parts are described briefly below and details are given in van Melle's thesis.

The rule interpreter has remained much the same since its conceptualization as a procedure that traces backward through chains of rules asking questions of the user only when values of parameters cannot be deduced. One important new development was a change in the certainty factor model used to propagate degrees of certainty from multiple pieces of positive and negative evidence to a conclusion. The new model is commutative, which means that it is no longer necessary to accumulate separate measures for positive and negative evidence, but a single measure reflecting both. The new model gives the same result as the old one in combining certainty factors of like sign, but is more gentle when combining CF's of opposite sign. The previous scheme had the problem of compressing too much information into the region near 1, and, as a result, ten conclusions of CF= 0.9 could be substantially overthrown by a single conclusion of CF= -0.8.

The explanation capability was improved by implementing the system's dictionary in hash tables. This gives the reasoning program more working space, since the dictionary can be kept on secondary storage until needed. Also, access time for parts of the dictionary is small because of the hash coding. The mechanisms used in EMYCIN have been used in other systems as well.

The input language for new rules has been simplified and stylized in an abbreviated rule language (ARL). ARL exploits the fact that associative triples are almost as easy for a person to read and manipulate as English text -- "the X of Y is Z", although stylized, is understandable. ARL resembles a shorthand form we have seen several domain experts use to sketch out sets of rules. The parameter names used in ARL are simply the labels the expert uses in defining the parameters of the domain. The conciseness of ARL makes it much easier to input than English, which is an important consideration when entering a large body of rules. Its conciseness is also a benefit when EMYCIN prints large numbers of rules for the expert to examine. Because knowledge acquisition is a critical problem for building expert systems [cf. Buchanan, 1981], we look on ARL as a pragmatic solution to a large problem we are addressing more fully in a knowledge acquisition system called ROGET.

The debugging package brings together pieces that have been part of the system for some time. These include: 1) the EMYCIN explanation facility; 2) a program that automatically explains how the system arrived at the results of a consultation; 3) a program that reviews each result of a consultation, allowing the user to judge whether the result is correct, and assisting the user in refining the knowledge base in order to correct any errors noted in the result or in intermediate conclusions; and 4) a program that automatically compares the results of a consultation to stored "correct" results for the same case, and explains any errors in the conclusions.

The rule compiler, described in last year's progress report, has been integrated with the whole system. Production rules, while convenient in their modularity, are not the best representation for speedy execution. The rule compiler transforms a program's production rules into a decision tree, eliminating the redundant computation inherent in a rule interpreter, and compiles the resulting tree into machine code. The program can thereby use an efficient deductive mechanism for running the actual consultation, while the flexible rule format remains available for acquisition, explanation, and debugging.

Finally, the EMYCIN user's manual has been improved. This manual is designed to be used by system builders who are creating a consultation system, not by the eventual users of the consultation system itself.

The second major research effort completed this year was an investigation of the efficacy of mixing production rules and frames in a medical reasoning program. This system, called CENTAUR, was described in Jan Aikins' thesis. The medical problem is identical to the topic for the PUFF system, namely the diagnosis of pulmonary function disorders. The medical knowledge, too, was fixed for some of the experiments. The ways the knowledge was represented and used in CENTAUR, however, were changed in order to determine whether improvements over PUFF could be found.

The representation of prototypical cases in CENTAUR improves the understandability of the program's line of reasoning and gives it more focus than PUFF has. Frames (called "prototypes" by Aikins) encode the knowledge of typical cases. These are linked in a hierarchy of ever more specialized descriptions of the subtypes, and are linked to production rules that associate evidence with hypotheses.

The control of reasoning in CENTAUR is also represented in frames, which allows explicit changes in reasoning strategies and experimentations with alternatives. Because of this explicitness, the program can explain its control strategy to a user, thus making it more understandable than in PUFF.

EMYCIN Applications

EMYCIN is intended for use by system builders who wish to construct a consultation program around a production rule representation of knowledge (with individual facts represented as associative triples) and a backward chaining control structure. It is not a universal programming language, but because the form of the final consultation program is fixed in advance, EMYCIN can save the system builder considerable effort. We have demonstrated this fact with several applications. Four were described in last year's report (PUFF, SACON, HEADMED, and CLOT). Three more were undertaken in this year, and a fourth was recently begun in the field of dermatology by Dr. Blois and his associates at UCSF.

GRAVIDA

A medical consultant called GRAVIDA, was developed to track an expectant mother through her pregnancy. Constructed by Dr. Val Catanzarite, currently a resident at Santa Clara Valley Medical Center, the system acquires information about current and past medical problems of the mother, any previous pregnancies, and general historical data about the patient. GRAVIDA then keeps track of the patient on a per-visit basis, recommending tests, detecting potentially dangerous medical conditions, and estimating the current age of gestation. The construction of this consultant required the extension of the rule language (by adding several new predicate functions) to look for simple trends and events over a series of previous visits.

DART

The other consultants are applied to non-medical domains. In conjunction with the IBM Corporation we have developed a consultant, called DART, that identifies probable causes of failures in teleprocessing subsystems of IBM 370-class computer systems. The system accepts stylized descriptions of the observed failure (e.g., lost data, machine went into a loop, terminal doesn't respond, etc.) and then directs the acquisition of data which are collected from traces available to field service personnel. Finally, DART uses these data to indict specific components, both hardware and software, which might be broken.

LITHO

The other major consultant now under development seeks to identify rock formations found at various depths of an oil-well bore hole. The consultant, called LITHO, examines geological and physical data of individual zones of interest to identify various aspects of the geological formations. This consultant is being constructed in conjunction with the Schlumberger Corporation and is similar to the GEO consultant developed with the AGE system.

With the publication of van Melle's EMYCIN thesis, which deals with the design of improved knowledge acquisition facilities for EMYCIN, and with the availability of an EMYCIN manual, each of the three consultants described above was constructed largely by the experts themselves. After initial discussions concerning the design of the system's goals, the identification of the data to be gathered, and the basic flow of the consultant's dialogue, the process of writing and inputting the hundreds of rules and parameters per system has been done primarily by the expert. All of them have remarked on the ease with which the current facilities allow this interaction to occur. As a result of these experiments, numerous improvements and modifications, both to the EMYCIN system and to the manual, are being incorporated into the package.

GUIDON

The original version of GUIDON, described in Clancey's 1979 thesis, was developed as an experiment to test the educational potential of MYCIN's rules and the ability to use the rule base outside of the consultation setting. Experiments with medical students indicated that GUIDON's framework for teaching knowledge was reasonably satisfactory, but the teaching points were not always clear in the rules. We then conceived a two step plan: first, to analyze the rule set and change it as necessary for teaching purposes (retaining the consultative capability), and second, to test this revised rule set in a new version of GUIDON. In 1980, the first step of this plan was achieved.

Analysis of MYCIN's Rules

From the period February 1 through December 1, 1980 we met regularly with a physician consultant for the purpose of revising MYCIN's rules so that teaching points were clear. Protocol analysis (presenting cases MYCIN had solved) was the chief method. We also experimented with sorting of medical findings, direct lectures, and undirected recall ("tell me everything you know about..."). We attended a series of courses taught by the same physician and compared them to another physician's handling of the same course.

Key Findings:

a. Our framework of structural, support, and strategic knowledge for organizing, justifying and controlling the use of heuristic rules served well in knowledge acquisition dialogues. We would always ask ourselves, "What kind of rationale is he giving me? A data/hypothesis rule? Why does he believe a rule? Why did he think to consider that association (the indexing, the approach)?" We put our analysis on this psychological footing from the start, because we learned in GUIDON1 that a tutorial program must incorporate knowledge that people use to access and control their heuristics.

b. It is not sufficient to revise MYCIN's rules; the decomposition of knowledge into subgoals is itself sometimes imprecise or non-standard.

c. Knowledge has to be added, namely the expertise for when to use MYCIN: when should one think about meningitis? what might it be confused with? MYCIN was not designed to be the "primary care" physician, but teaching diagnosis, our goal, involves expanding the knowledge base to include initial problem formulation.

d. The physician's approach was logical and easy for us to emulate. He was consistent from case to case, and moreover did what he told students to do. This is not necessarily typical. Other teachers we observed were not able to articulate their approach as clearly and seemed to be less sure of what students were thinking. There were common strategical concepts, however, that our three experts all used to explain their reasoning.

Development of NEOMYCIN

We implemented a prototype consultation system that constitutes a psychological model of diagnostic problem solving. This system is upward compatible with EMYCIN systems, and thus could replace the EMYCIN language and interpreter. Key theoretical features of the design:

a. Forward-directed reasoning from data to hypotheses and state-categories, emulating expert problem solving:

   1) Trigger rules place hypotheses in the differential diagnosis directly as data are received. The differential is maintained so that more specific causes replace general hypotheses.

   2) Data are abstracted immediately, e.g., "diplopia" is thought of as an "abnormal neurological finding"

   3) Process-oriented questions are immediately asked if they are relevant to the domain, even if not directed to any particular hypothesis, e.g., asking when a symptom began and how it has changed over time.

   4) Data suggest causal state-categories, possibly jumping over a chain of causal links to conjecture some generic problem.

   5) Data/hypothesis associations are applied in the context of the current differential diagnosis (working memory of hypotheses).

b. Explicit, separate representation of:

   1) a problem-space hierarchy to which data/hypothesis rules are attached ("etiological taxonomy") (previously implicit as the "context clauses" of rules);

   2) causal rules that ultimately tie into this hierarchy;

   3) world relations that constrain the relevance of data (previously implemented as "screening clauses");

   4) disease process knowledge that cuts across the etiological distinctions, useful for initial problem formulation.

c. A hierarchical set of domain-independent meta-rules constitute a diagnostic meta-strategy. These rules examine the knowledge sources listed above and the current differential to select an hypothesis to focus on and the next datum to collect.

The key strategical idea to teach students is that collecting circumstantial evidence is preparation for making physical measurements. Its purpose is to "establish the hypothesis space," to determine the range of possibilities that might be causing the problem. Strategies for achieving this involve considering common and unusual causes, looking for evidence that will broaden the space of possibilities.

There are two orientations when establishing the hypothesis space: 1) "group and differentiate" -- upward-looking, initial problem formulation in which one tries to cluster the data under some generic process (cause); and 2) "explore and refine" -- attempting to confirm successively more specific causes.

The diagnostic meta-rules are generally applied as a pure-production system for each subtask (e.g., "find a new focus" is a subtask). Abort conditions are inherited to simulate shifting of focus (and return to higher goals) as data broaden the differential or exploration suggests that a conjecture is unlikely.

2. Research in Progress

Short-term Plans for NEOMYCIN and GUIDON

We are shifting development of GUIDON to the Dolphin computer, now on loan from Xerox and located in the Computer Science Department building at Stanford. GUIDON must be revised to be compatible with the NEOMYCIN system. These revisions take two forms: a) simplifications to the code (NEOMYCIN is designed to make it easier to index rules as they are used in the tutorial), and b) extensions to take advantage of knowledge now represented explicitly in NEOMYCIN (taxonomy of problems, world facts, diagnostic strategy).

NEOMYCIN is essentially a psychological model of diagnosis that enables us to monitor the student's problem solving and provide assistance in ways that were not possible before. For example, we will be teaching forward-directed inferences -- leaps from data to hypotheses -- that we represent in NEOMYCIN's trigger rules. With this additional knowledge of how experts think, GUIDON version 2 will have leverage for interrupting the student to test his knowledge, as well as having a better basis for understanding a student's partial solutions.

We expect that complete revisions of GUIDON so that it can take advantage of what is now in NEOMYCIN will require 6 months. This includes an entirely redesigned student model, plus the new capabilities for interruption, assistance, and evaluation of student hypotheses. In parallel, we will be refining NEOMYCIN by testing it on the 100 meningitis cases in our library. Two students will be revising GUIDON; a third student will continue development of NEOMYCIN (on the SUMEX-AIM computer facility). Dr. Clancey will direct and participate in both aspects of the project.

Formalization of Teaching Principles

One of the students who will be revising GUIDON is a doctoral candidate in the Education Department at Stanford. For his thesis research, this student will be parameterizing GUIDON's tutorial rules so they are controlled by a higher order model of teaching methods. Design of this model is complete on paper now. It will be implemented after GUIDON2 is working.

Formal Experimentation

Through our contacts with the medical school, we have arranged to test GUIDON with medical students during the period September '81 - March '82. This aspect of the project will be managed by the doctoral student in Education. Plans are to do exploratory experimentation 1) to test the usefulness of the diagnostic model for interpreting student behavior 2) determine whether theoretical differences in tutoring behavior are detectable by the students. Analysis of results should provide a basis for extending the diagnostic model.

Development of a Mechanical/Electronic Diagnostic Program

We have begun collaboration with researchers from IBM to develop a system similar to NEOMYCIN in the domain of computer failure diagnosis. The purpose of this project will be to determine to what extent the domain-independent strategies we formalized from experience in the medical domain are applicable to electronic troubleshooting. In the past year, Prof. Buchanan supervised development of an EMYCIN consultation program, named DART, for diagnosing teleprocessing problems. Through this experience, IBM personnel learned about our techniques, and we were introduced to the hardware and software problems they need to solve. We will be drawing upon this experience in the next year.

ONCOCIN

The oncology chemotherapy consultation system, named ONCOCIN, has achieved many of its goals since work on the project began in July 1979. We are developing an interactive system to be used by oncology faculty and fellows in the Debbie Probst Oncology Day Care Center at Stanford University Medical Center. Our overall goals are:

(1) to demonstrate that a rule-based consultation system with explanation capabilities can be usefully applied and gain acceptance in a busy clinical environment;

(2) to improve the tools currently available, and to develop new tools, for building knowledge-based expert systems for medical consultation; and

(3) to establish both an effective relationship with a specific group of physicians, and a scientific foundation, that will together facilitate future research and implementation of computer-based tools for clinical decision making.

Specific Objectives:

The ONCOCIN research goals are directed both towards the basic science of artificial intelligence and towards the development of clinically useful oncology consultation tools.

Artificial Intelligence Objectives

We have undertaken AI research with the following aims:

(1) to implement and evaluate recently developed techniques designed to make computer technology more natural and acceptable to physicians;

(2) to extend the methods of rule-based consultation systems to interact with a large database of clinical information; and

(3) to continue basic research into the following problem areas: mechanisms for handling time relationships, techniques for quantifying uncertainty and interfacing such measures with a production rule methodology, approaches to acquiring knowledge interactively from clinical experts, assessment of knowledge base completeness and consistency.

Oncology Clinic Objectives

We have begun to develop and implement a protocol management system, for use in the oncology day care center with the following capabilities:

(1) to assist with identification of current protocols that may apply to a given patient;

(2) to assist with determining a patient's eligibility for a given protocol;

(3) to provide detailed information on protocols in response to questions from clinic personnel;

(4) to assist with chemotherapy dose selection and attenuation for a given patient;

(5) to provide reminders, at appropriate intervals, of follow-up tests and films required by the protocol in which a given patient is enrolled;

(6) to reason about managing current patients in light of stored data from previous visits of (a) the individual patients, or (b) the aggregate of all "similar" patients.

Overview of Goals for 1980:

We have described a five-year plan for accomplishing the above goals. As discussed at this time last year, we spent our first year developing a prototype ONCOCIN consultation system, drawing from programs and capabilities developed for the EMYCIN system-building project. During that year, we also undertook a detailed analysis of the day-to-day activities of the Stanford oncology clinic in order to determine how to introduce ONCOCIN with minimal disruption of an operation which is already running smoothly. We also spent much of our time in the first year giving careful consideration to the most appropriate mode of interaction with physicians in order to optimize the chances for ONCOCIN to become a useful and accepted tool in this specialized clinical environment.

During our second year of the project, we have accomplished all the goals we identified for 1980:

(1) We have completed a special interface program that responds to commands from the customized keypad described last year;

(2) We encoded the rules for one more chemotherapy protocol (oat cell carcinoma of the lung) and updated the Hodgkin's Disease protocols when new versions were released late in 1980; these exercises demonstrated the generality and flexibility of the representation scheme we have devised;

(3) We developed the software protocols for achieving communication between the interface program and the reasoning program;

(4) We have coordinated the printing routines needed to produce hardcopy flowsheets, patient summaries, and encounter sheets;

(5) Lines have been installed between the SUMEX machine room and the oncology clinic, and the new terminal and a hard copy device have been installed in the Oncology Day Care Center for final testing and debugging; and

(6) We have just begun to offer the ONCOCIN system for use by oncology faculty and fellows in the morning chemotherapy clinics in which most of the lymphoma patients receive their treatment.

We had two additional goals, not explicitly stated in last year's report. One was to design formal evaluation studies that would allow us to assess the impact of ONCOCIN and its acceptance by the physicians for whom it is designed. Second, we wanted to experiment with computational techniques for verifying the completeness and consistency of a developing knowledge base.

PROGRESS - 1980/81:

Further Development and Testing of the Reasoning Program

The early prototype of the Reasoning system was described in last year's report in some detail. A more recent summary has been submitted for presentation at the 7th International Joint Conference on Artificial Intelligence. The Reasoner is coded in Interlisp, and is running on the SUMEX computers (both the PDP-10, and the 20/20 on which we have been running when the system is used in the oncology clinic).

The Reasoner has been extensively debugged this year. Several hundred sample patient cases have been run, and the results have been reviewed in detail by the collaborating oncologists. When problems have been uncovered by this process, changes in the Reasoner program (or in the encoding of the lymphoma protocol knowledge) have been undertaken.

Verification of the Adequacy of the Knowledge Representation Scheme

In an effort to verify that the representation scheme we are using will be adequate for arbitrary protocol knowledge that may be encountered in the future, we decided to encode and briefly test the knowledge of a non-lymphoma protocol. We chose the complicated protocol for oat cell (small cell) carcinoma of the lung because it involves a large number of possible therapies and complex interweaving of chemotherapy and radiotherapy. After approximately one month's effort by an experienced programmer, the oat cell protocol had been encoded and ran successfully on a few test cases. In addition, the lymphoma protocols themselves were changed in late 1980, and we spent a few weeks in early 1981 entering the changes implicit in these new versions. In all cases the ONCOCIN representation scheme was adequate to accommodate the protocol knowledge with only minor modifications, if any, and for this reason we are confident that our system will be able to adapt to any other protocols that need to be encoded in the coming years.

Physician/Computer Interaction

The actual mechanics of computer terminal interaction is as important to a clinical system's acceptance as the quality of the program's advice. If a system is slow or cumbersome, physicians will tend to reject it. With this in mind, we have sought to develop an optimal interactive mechanism that will not unreasonably tax the budget of the project.

In last year's report we indicated that this interactive system was to be written in PASCAL. After some initial experiments, however, we decided to use SAIL instead. The system is referred to as the "Interviewer", and it has now been fully implemented and debugged.

As we emphasized when outlining our research goals, we have wanted ONCOCIN to maintain the explanation and justification capabilities that we have argued are crucial to the acceptance of clinical consultation systems. The Interviewer uses a specialized split-screen display that enables the physician to enter patient data entries in one region while pertinent explanations are displayed in another.

Development of Mechanisms for Interprocess Communication

Because the Reasoner (the Interlisp reasoning program) and the Interviewer (the SAIL program with which the physician interacts) must run in parallel in two different processes on the same machine, we needed to devise mechanisms for allowing these two programs to communicate with one another. This has been a major systems programming task, but we are pleased with the effectiveness of the generalized interprocess communication mechanism that we devised.

Designing an Evaluation of the ONCOCIN System

Because we wish to evaluate formally the impact of ONCOCIN and its effectiveness in the oncology clinic, we have devised a set of three experiments, two of which are already underway. The study designs are outlined in detail in an evaluation document that we have prepared.

Verifying the Completeness and Consistency of the Knowledge Base

An important question for AI researchers involved with the development of expert systems is how to ascertain that a knowledge base for a consultation program is complete and consistent. Dr. Motoi Suwa, a visitor to Stanford from Japan, became fascinated with this question and collaborated with us on a formal analysis of the developing ONCOCIN knowledge base. His paper describing that work was submitted for presentation at the 7th International Joint Conference on Artificial Intelligence.

D. Publications Since January 1980

Aikins, Janice S. Prototypes and Production Rules: A Knowledge Representation for Computer Consultations. PhD Dissertation, Stanford University. Memo HPP-80-17, August, 1980.

Bennett, S.W., and Scott, A.C. Computer-assisted customized antimicrobial dosages. Amer. J. Hosp. Pharm. 37:523-9 (1980).

Buchanan, B.G. Research on Expert Systems. Memo HPP-81-1, January, 1981. To appear in D. Michie (ed.) Machine Intelligence 10.

Campbell, A.B., Chang, P., Cho, J., Hickam, D., Shortliffe, E.H., and Teach, R. Preliminary proposal for the evaluation of the ONCOCIN System. Internal memo, the ONCOCIN Project, April 1981.

Clancey, W.J. The Epistemology of a Rule-based Expert System: the effect of proceduralization of knowledge on explanation. To appear in the Journal of Artificial Intelligence.

Clancey, W.J. and Letsinger R. NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. Submitted to IJCAI7.

Clancey, W.J. Methodology for building an intelligent tutoring system. To appear in a book on Cognitive Science Methodology, edited by Kintsch, Miller, and Polson.

Fagan, L.M., Shortliffe, E.H., and Buchanan, B.G. Computer-based medical decision making: from MYCIN to VM. Automedica 3:97-106 (1980).

Gerring, Phil. System documentation: interprocess communication system (TopDog and Interactor). Internal memo, the ONCOCIN Project, November 1980.

Shortliffe, E.H., Scott, A.C., Bischoff, M.B., Campbell, A.B., van Melle, W., and Jacobs, C.D. ONCOCIN: An expert system for oncology protocol management. Submitted to Proceedings of the 7th IJCAI, Vancouver, B.C., August 1981.

Shortliffe, E.H. Consultation systems for physicians. Proceedings of the CSCSI/SCEIO Conference, 14-16 May 1980, University of Victoria, British Columbia, pp. 1-11.

Shortliffe, E.H. Medical Cybernetics: The Challenges of Clinical Computing. To appear in Cybernetics, Technology, and Growth, S. Basheer Ahmed, editor; Lexington Books, 1981.

Shortliffe, E.H. Medical Computing: Another Basic Science? Proceedings of the 4th Symposium on Computer Applications in Medical Care, Washington, D.C., November 1980.

Shortliffe, E.H. The computer as clinical consultant (editorial). Arch. Int. Med., March 1980.

Suwa, M., Scott, A.C., and Shortliffe, E.H. An approach to verifying completeness and consistency in a rule-based expert system. Submitted to Proceedings of 7th IJCAI, Vancouver, B.C., August 1981.

Teach, R.L. and Shortliffe, E.H. An analysis of physician attitudes regarding computer-based clinical consultation systems. Submitted for publication in Comput. Biomed. Res., February 1981.

van Melle, W., Shortliffe, E.H., Buchanan, B.G. EMYCIN: A domain-independent system that aids in constructing knowledge-based consultation programs. To appear in Pergamon-Infotech State of the Art Report on Machine Intelligence, 1981.

van Melle, W. A domain-independent system that aids in constructing knowledge-based consultation programs. PhD thesis, Computer Science Department, Stanford University, June, 1980.

E. Funding Support

Grant Title: "Knowledge-Based Consultation Systems"
Principal Investigator: Bruce G. Buchanan
Agency: National Science Foundation
ID Number: MCS-7903753
Term: July 1979 to March 1981
Total award: $146,152
Current award (1980): $72,493

[No continuation proposal was submitted to the NSF since the current version of the system successfully completes our proposed work. We intend to use EMYCIN as a vehicle for experimental research under other funding, including SUMEX core research, but we are not proposing further research or development on EMYCIN itself.]

Contract Title: "Exploration of Tutoring and Problem-Solving Strategies"
Principal Investigator: Bruce G. Buchanan
Agency: Office of Naval Research and
Advanced Research Projects Agency (joint)
ID number: N00014-79-C-0302
Term: March 1979 to March 1982
Total award: $396,325

Grant Title: "Explanatory Patterns In Clinical Medicine"
Principal Investigator: Edward H. Shortliffe
Agency: Kaiser Family Foundation
Term: July 1979 to December 1980
Total award: $20,000
SKIP 2
Grant Title:"Research Program: Biomedical Knowledge Representation"
Principal Investigator: Edward A. Feigenbaum
Co-Principal Investigator (ONCOCIN Project): Edward H. Shortliffe
Agency: National Library of Medicine
ID Number: LM-03395
Term: July 1979 to June 1984
Total award: $497,420
Current award (1980-1981): $99,400
Administered through Medicine: ONCOCIN suballocation ($47,845)

Grant Title: "Symbolic Computation Methods For Clinical Reasoning"
Principal Investigator: Edward H. Shortliffe
Agency: National Library of Medicine
ID Number: LM-00048
Term: July 1979 to June 1984
Total award: $196,425
Current award (1980-1981): $39,107

II. INTERACTION WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

A great deal of interest in both MYCIN and EMYCIN has been shown by the medical and academic communities. For two years in succession we were invited by the American College of Physicians to demonstrate MYCIN at the organization's annual meeting (San Francisco, March 1979, and New Orleans, April 1980). The physicians have uniformly been enthusiastic about the program's potential and what it reveals about one current approach to computer-based medical decision making. In both cases, the demonstrations were performed on-line using network access to the SUMEX computer.

We have demonstrated our programs to both physicians and computer scientists on numerous additional occasions. At the AIM tutorial in August 1980, both MYCIN and GUIDON were presented to introduce physicians to the field of AI in medicine. GUIDON was also demonstrated on the Dolphin machine at the Xerox-PARC open house during the AAAI in August. In addition, both EMYCIN and GUIDON were featured demonstrations at the annual AIM Workshop, held the same week at Stanford. The TYPER program, developed by SUMEX staff in collaboration with Dr. Larry Fagan of Stanford, was used to good effect at this workshop as well as for informal demonstrations throughout the year.

Several project members contributed to the Expert Systems Workshop, sponsored by RAND and ARPA and held in San Diego in August, where EMYCIN was one of the "system building tools" that was studied in detail. The Workshop has led to the preparation of a book, "Building Expert Systems," and many of our research group have written portions of that volume (Buchanan, Clancey, Scott, Aikins, Shortliffe, van Melle).

NEOMYCIN was presented to the contractors of the ONR "Instructional systems and advanced training" division, held in Pittsburgh, in January 1981. Presentations of this kind carry SUMEX-AIM results out to cognitive psychologists from around the country. Dr. Clancey also presented a talk on GUIDON research at the annual conference of the Association for Development of Computer Instructional Systems in Atlanta, in March 1981, and at the annual conference of AERA in Los Angeles, in April 1981.

Several medical school and computer science teachers have also asked to use MYCIN in their computer science or medical computing courses, and we continue to make the programs available frequently to researchers around the world who access SUMEX using the GUEST account.

EMYCIN has generated considerable interest in the academic and business communities. We have been in frequent contact with Bud Frawley and Alain Bonnet, of Schlumberger, Chuck Brodnax and Milt Waxman of the Hughes Aircraft Corporation, and Harry Reinstein and Cliff Hollander from IBM Scientific Research Center. EMYCIN, on SUMEX, has been used at the University of Illinois and Michigan State University to explore the construction of expert systems.

B. Sharing and Interaction with Other SUMEX-AIM Projects

We have continued collaboration with the RX, VM, and PUFF projects. Our development of a domain-independent system is facilitated by having a number of very different working systems on which to test our additions and modifications to EMYCIN. All the projects have provided us with useful comments and suggestions.

The community created on the SUMEX resource has other benefits that go beyond actual shared computing. Because we are able to experiment with other developing systems, such as INTERNIST, and because we frequently interact with other workers (at the AIM Workshop or at other meetings around the country), many of us have found the scientific exchange and stimulation to be heightened. Several of us have visited workers at other sites, sometimes for extended periods, in order to pursue further issues which have arisen through SUMEX- or Workshop-based interactions. In this regard, the ability to exchange messages with other workers, both on SUMEX and at other sites, has been crucial to rapid and efficient exchange of ideas. For example, most of the invitations and planning for the 6th AIM Workshop, held at Stanford in August 1980 and described in detail elsewhere in this report, were accomplished via SUMEX or ARPANET mail. Certainly it is unusual for a small community of researchers with similar scholarly interests to have at their disposal such powerful and efficient communication mechanisms, even among those on opposite coasts of the country.

C. Critique of Resource Management

The SUMEX facility has maintained the high standards that we have praised in the past. The staff members are always helpful and friendly, and work as hard to please the SUMEX community as to please themselves. As a result, the computer is as accessible and easy to use as they can make it. More importantly, it is a reliable and convenient research tool. We extend special thanks to Tom Rindfleisch for maintaining high professional standards for all aspects of the facility.

Due to the introduction of our ONCOCIN work with its special hardware and communication needs, we continue to be aware that we are taxing the limited resources of SUMEX with regards to technical hardware support. It has been next to impossible for one technical specialist (Nick Veizades) to balance the numerous diverse demands on his time. This is not a problem with management of the Resource but a reflection of the need for additional technical personnel associated with SUMEX. We perceive this to be a particularly important requirement in the future as the Resource undertakes an expanded role in the implementation and testing of new hardware.

Special mention should be made of the remarkable role played by Tom Rindfleisch and his staff in helping to organize remote demonstrations of SUMEX-AIM programs. In October, 1980, when the NIH Council on Research Resources met in Atlanta, demonstrations of MYCIN and INTERNIST on the DEC 2020 at Stanford were so carefully arranged as to make them seem commonplace. We salute Tom and the staff for their uncomplaining assistance, and are grateful for the efforts they have made to provide a mechanism for facilitating future demonstrations at remote locations.

Finally, we continue to feel the need for more computing power. Much of our research and development continues to take place in the hours from 7 p.m. to 10 a.m., but it is unreasonable to expect all our programming staff to adjust their own schedules around a computer. The existence of the 20/20 has been helpful in permitting demonstrations with good response time, and it has allowed us to introduce ONCOCIN in a real clinical environment, but ongoing R&D on the main machine remains difficult much of the time. Even the evening hours are now seeing higher load averages than was once the case. We anticipate considerable improvement in this regard as the recently approved additional computing hardware becomes available. In the meantime, much of the work on EMYCIN has been moved to the SCORE computer in the Computer Science Department. Response time aside, we have shifted our development of GUIDON to the Xerox Dolphin in order to take advantage of the larger address space. This also frees up disk space so that we can comfortably develop NEOMYCIN on SUMEX.

We also strongly support the creation of the new position assumed by Anne Fadenrecht; her excellent early efforts should be especially helpful in taking some of the load off of Carole Miller and Tom Rindfleisch.

III. RESEARCH PLANS

A. Project Goals and Plans

EMYCIN

Now that the design and capabilities of EMYCIN are essentially fixed, we are planning to develop new applications and to use the system as an experimental tool. The applications to electronic fault diagnosis and geology will continue and we expect to find additional medical applications as well. Many of these we expect will be undertaken by other research groups. Because we view artificial intelligence as an experimental science [Buchanan, 1981], we wish to collect data on the nature of problems EMYCIN can help solve and the limitations of the problem solving method embodied in EMYCIN.

Our research on knowledge acquisition depends on the existence of a working EMYCIN system. In the ROGET program, currently under development, hierarchical knowledge about consultation systems and their knowledge bases is used to help an expert define a new knowledge base to be used by EMYCIN. For example, the meningitis and pulmonary function knowledge bases both contain rules associating diagnoses with laboratory tests and with clinical findings. ROGET will be able to use this fact to help an expert divide a new rule set into rules using test results and measurements as evidence and another rule set using more subjective evidence.

GUIDON

We have now established a good framework for organizing knowledge in an expert system to be used for tutoring. We characterize knowledge by its use for: structuring knowledge sources, supporting (justifying) knowledge sources, or controlling their invocation. In the most general terms, our plans are to do research in acquiring, representing and presenting structural, support, and strategic knowledge.

We used this framework to design NEOMYCIN. Experiments during the coming year will provide a basis for developing our model of diagnosis. In particular, we propose:

a) to extend NEOMYCIN's model of diagnostic strategy to include common, non-expert approaches. Besides improving the program's ability to model the student, this enumeration of the space of strategies will allow us to follow a plan of research similar to Brown's and Burton's, but in the domain of diagnostic strategy as opposed to subtraction procedures. Eventually, we want to develop a principled psychological model that will relate strategies to knowledge and processing abilities.

b) further studies of expert reasoning in domains that require "forming a picture" of a malfunctioning process. Experience with NEOMYCIN showed that expert diagnosticians attempt to order the data they collect causally, on a time line. Interpretation of observations can be partially understood as an attempt to match this description of onset, course, severity (intensity, frequency), and causal relations of findings onto known malfunctions that are recalled (indexed) by these process variables. This work will build upon recent advances in understanding causality (e.g., deKleer and Brown).

c) exploitation of new technology for experimentation with teaching methods. How can we take advantage of the Dolphin's graphic capabilities in a GUIDON tutorial? Besides graphically presenting rule relationships, we might show the student the same kind of diagrams that we use when describing our knowledge bases to our AI colleagues (hierarchies, diagrams relating compiled associations to underlying causal chains). Other than presentation strategies, we would like to experiment with different interfaces, perhaps to break away from a continuous dialogue to use the screen more as a work space for annotating and examining the knowledge base, and organizing data and hypotheses in a diagnostic problem.

d) incorporate GUIDON as an integral part of the curriculum in in medical diagnosis at Stanford. We propose to make GUIDON available at the Fleischmann Learning Center at Stanford Medical School, just as the traditional programs built at Massachusetts General and Ohio State were made available. In addition, we will work with one or more teaching fellows at the medical school to include GUIDON as part of the "clinical diagnosis" course which is taught regularly at Stanford. This will continue our commitment to empirical research to develop our model of diagnosis and the teaching procedures.

ONCOCIN

During the coming year, there are four principal areas in which we expect to expend our efforts on the ONCOCIN System:

(1) The system will be implemented for ongoing use in the Stanford Oncology Clinic, with an experimental evaluation period to begin July 1, 1981.

(2) The system will be formally evaluated with regard to its impact on (a) the attitudes of the oncologists, (b) the accuracy and completeness of data collection, and (c) the adequacy of the management decisions made in the clinic.

(3) We will begin to encode additional protocols as the lymphoma system comes into regular use and physicians begin to demand the inclusion of a greater percentage of the protocols used in the management of cancer patients at Stanford.

(4) We will begin to devote a greater percentage of our time to experiments in encoding complex judgmental reasoning of the sort that is usually performed by expert oncologists and is not formally specified in the protocol documents themselves.

Throughout the year we shall continue to relate the requirements of the system we are developing to the underlying artificial intelligence methodologies. We are convinced that the basic science frontiers of AI are best explored in the context of systems for real world use; thus ONCOCIN serves as a vehicle for developing an improved understanding of the issues that underlie all forms of knowledge engineering.

B. Requirements for Continued SUMEX Use

All the work we are doing (EMYCIN, GUIDON, ONCOCIN, plus continued use of the original MYCIN program) is totally dependent on continued use of the SUMEX resource. The programs all make assumptions regarding the computing environment in which they operate, and the ONCOCIN design in particular depends upon proximity to the DEC 2020 which enables us to use a 9600 baud interface.

In addition, we have long appreciated the benefits of GUEST and network access to the programs we are developing. SUMEX greatly enhances our ability to obtain feedback from interested physicians and computer scientists around the country. Network access has also permitted high quality formal demonstrations of our work both from around the United States and from sites abroad (e.g., Japan, Sweden, Switzerland).

We plan to continue development of NEOMYCIN on SUMEX during the next year, whereas the GUIDON/Dolphin effort will continue on the crowded Computer Science Department Dolphin only until the SUMEX individual workstations become available. Using the main SUMEX machine, we intend to make NEOMYCIN fully usable as a consultation program so that it can be compared with MYCIN. In particular, we will be comparing cases run through both MYCIN and NEOMYCIN to see whether simplification and clarification of the rules for purposes of teaching will in turn change the program's accuracy.

C. Requirements for Additional Computing Resources

The acquisition of the DEC 2020 by SUMEX has been crucial to the growth of our research work, both to insure high quality demonstrations and to enable us to develop a system such as ONCOCIN for real-world use in a clinical setting. As we continue to develop systems that are potentially useful as stand-alone packages (e.g., an exportable EMYCIN), the additional small computers that are planned will be particularly valuable resources. It is not yet clear which machines are optimal for the LISP-based applications we are developing, and an opportunity to test our systems on several small-to-medium machines will be invaluable and in keeping with our desire to move some of the AIM products into a community of service users.

As we have mentioned, the response time on the main machine continues to be a major problem, both during the daytime hours and frequently in the evenings as well. The proposed SUMEX acquisitions that will provide additional cycles and permit off-loading of some users from the PDP-10 will significantly benefit the SUMEX research community.

In addition, we believe that our GUIDON experience using the Dolphin personal computer is a significant part of our research. First, the Dolphin's large address space will permit development of the large knowledge base that an intelligent tutoring system requires; we have overgrown the facilities available at SUMEX. Second, the Dolphin's graphics will enable us to develop new methods for presenting material from the knowledge base. Third, the Dolphin will provide a reliable, constant "load-average" machine, for running experiments with students. Finally, the development of GUIDON on the Dolphin demonstrates the feasibility of running intelligent tutoring systems on small, affordable machines in schools and remote sites.

We seem to have an insatiable appetite for disk storage space, even though ONCOCIN received an additional substantial allocation since our report last year. ONCOCIN, in particular, has become an extremely large system, and the data files for a clinic full of patients will require substantial additional space. We hope that the planned SUMEX file-server will allow the allocation of several thousand more pages. It should also help alleviate the need to keep copies of all patient files on both the 10 and the 20/20.

D. Recommendations for Future Community and Resource Development

In last year's report we made two recommendations for new SUMEX developments: (1) the acquisition of several small machines, linked to the main processor through the Ethernet and able to run INTERLISP, and (2) the formal establishment of a mechanism for providing hardware and communications equipment for SUMEX demonstrations at a distance. Both of these have been acted upon by SUMEX, and we are delighted by this kind of responsiveness.

The AIM community is small and close-knit, but there remain communication problems within it. The AIM Workshops are excellent means of transferring information annually, but between Workshops all of us are remiss in not communicating new technical reports and articles. It would be very desirable to maintain a list of current publications from all the AIM research groups, for distribution by ARPANET or U.S. Mail to all others. No group will add to the list, however, unless the benefit of the information gained from such a list exceeds the cost of adding to it. SUMEX may be able to function as a catalyst to this kind of community communication.

II.A.1.7 Protein Structure Project

Protein Structure Modeling Project

Prof. E. Feigenbaum and Mr. Allan J. Terry
Department of Computer Science
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Technical Goals

The goals of the protein structure modeling project are to 1) identify critical tasks in protein structure elucidation which may benefit by the application of AI problem-solving techniques, and 2) design and implement programs to perform those tasks. We have identified two principal areas which are of practical and theoretical interest to both protein crystallographers and computer scientists working in AI. The first is the problem of interpreting a three-dimensional electron density map. The second is the problem of determining a plausible structure in the absence of phase information normally inferred from experimental isomorphous replacement data. Current emphasis is on the implementation of a program for interpreting electron density maps (EDM's).

B. Medical Relevance and Collaboration

The biomedical relevance of protein crystallography has been well stated in an excellent textbook on the subject (Blundell & Johnson, Protein Crystallography, Academic Press, 1976):

> "Protein Crystallography is the application of the techniques of X-ray diffraction ... to crystals of one of the most important classes of biological molecules, the proteins. ... It is known that the diverse biological functions of these complex molecules are determined by and are dependent upon their three-dimensional structure and upon the ability of these structures to respond to other molecules by changes in shape. At the present time X-ray analysis of protein crystals forms the only method by which detailed structural information (in terms of the spatial coordinates of the atoms) may be obtained. The results of these analyses have provided firm structural evidence which, together with biochemical and chemical studies, immediately suggests proposals concerning the molecular basis of biological activity."

The project involves a collaboration between computer scientists at Stanford University and crystallographers at Oak Ridge National Laboratories (Dr. Carroll Johnson), the University of California at San

Francisco (Dr. Robert Langridge), and the University of California at San Diego (under the direction of Prof. Joseph Kraut). Our principal collaborator at UCSD is Dr. Stephan Freer. We also collaborate with Dr. Eric Grosse at Bell Laboratories, whose field is numerical analysis.

C. Progress Summary

We have completed a major cycle of design review and program reorganization, resulting in the system described in publications four and nine below. The system now has a completely hierarchical, rule-based control structure proceeding from strategy rules, to a set of task rules, ending with individual knowledge sources. This new design seems powerful and flexible enough to provide the basis of a useful EDM interpretation system for protein structure determination.

After building the control structure we wanted, we have worked on building up the knowledge base. Large chunks of knowledge are called "tasks"; we have implemented five out of a projected set of nine. To date, we have implemented the Initialization task, two tracing tasks, a task to split "group toeholds", and a version of a task that finds "second generation" toeholds. Further details of these tasks and their content can be found in publication number four.

We have also continued our efforts to improve the power of our data representations. Towards this end we have implemented a new preprocessor based on Dr. Grosse's thesis research. This program is an improved method for finding the critical points of a function. In our case, the peaks of the electron density map are useful guides to atom locations and the full set of critical points are used in the ridge-line analysis discussed in publication one.

Finally, we are compiling documentation on the system and the knowledge it embodies. These documents should be sufficiently complete so that we, or other groups, will have little difficulty picking up where we leave off. We also feel that explicit documentation of our model-building heuristics will be useful to the crystallographic community as it provides a new viewpoint, complementary to traditional crystallographic methods.

The work currently in progress can be characterized as additions to the knowledge base and work on new data representations. The five tasks currently implemented form the core of the system and suffice to solve about a half of a small protein. The remaining tasks will embody knowledge about finding new toeholds (to restart the trace when it is blocked) and about tracing in areas of the data too complex to interpret with present heuristics. One of the main areas of work along these lines is the incorporation of some notion of stereochemistry and the constraints on three dimensional structure it provides. This will be useful in the matching of features and in the prediction of secondary structure. The last item of work in progress is an attempt to design a data representation that captures volume information. Current representations such as the skeleton preserve topology but do not preserve shape. With the inclusion of volume information, we should be able to capture much of the expert's knowledge of shape and form that presently goes unused.

D. List of Publications

(1) Carroll Johnson and Eric Grosse, "Interpolation Polynomials, Minimal Spanning Trees, and Ridge-Line Analysis in Density Map Interpretation", American Crystallographic Association Program and Abstracts, 4:2, Evanston, Ill. Aug. 1976

(2) Robert S. Engelmore and H. Penny Nii, "A Knowledge-Based System for the Interpretation of Protein X-Ray Crystallographic Data," Heuristic Programming Project Memo HPP-77-2, January, 1977. (Alternate identification: STAN-CS-77-589)

(3) E.A. Feigenbaum, R.S. Engelmore, C.K. Johnson, "A Correlation Between Crystallographic Computing and Artificial Intelligence," in Acta Crystallographica, A33:13, (1977). (Alternate identification: HPP-77-15)

(4) Robert Engelmore and Allan Terry, "Structure and Function of the CRYSALIS System", Proc. 6IJCAI, 1979. pp250-256 (Alternative identification: HPP-79-16)

(5) R.S. Engelmore, A. Terry, S.T. Freer, and C.K. Johnson, "A Knowledge-Based System for Interpreting Protein Electron Density Maps", Abstracts of Amer. Crystallographic Ass. 7,1 (1979) p38

(6) E.H. Grosse, "Approximation and Optimization of Electron Density Maps", Stanford University Ph.D. Thesis, Dec. 1980 (Alternative identification: STAN-CS-80-835)

(7) R. Engelmore and A. Terry, Article VII.C3 (Crysalis) in Barr, A., and Feigenbaum, E. A. (eds.), The Handbook of Artificial Intelligence, Vol. II, Stanford Ca., HeurisTech Press, Los Altos, Ca.: Kaufman, 1981

(8) A. Terry and R. Engelmore, "A Knowledge-Based Approach to the Interpretation of Protein Electron Density Maps", to appear in a forthcoming book on expert systems by Pergamon Infotech International, Maidenhead, England

(9) A. Terry, "Hierarchical Control of Production Systems", paper submitted to 7IJCAI

E. Funding status

Grant title: The Automation of Scientific Inference: Heuristic Computing Applied to Protein Crystallography

Principal Investigator: Prof. Edward A. Feigenbaum

Funding Agency: National Science Foundation

Grant identification number: MCS 79-33666

Term of award: December 1, 1979 through November 31, 1981

Amount of award: $35,318 (direct costs only)

II. INTERACTION WITH THE SUMEX-AIM RESOURCE

A. Collaborations

The protein structure modeling project has been a collaborative effort since its inception, involving co-workers at Stanford and UCSD (and, more recently, at Oak Ridge, UCSF, and Bell Laboratories). The SUMEX facility has provided a focus for the communication of knowledge, programs and data. Without the special facilities provided by SUMEX the research would be seriously impeded. Computer networking has been especially effective in facilitating the transfer of information. For example, the more traditional computational analyses of the UCSD crystallographic data are made at the CDC 7600 facility at Berkeley. As the processed data, specifically the EDM's and their Fourier transforms, become available, they are transferred to SUMEX via the FTP facility of the ARPA net, with a minimum of fuss. (Unfortunately, other methods of data transfer are often necessary as well -- see below.) Programs developed at SUMEX, or transferred to SUMEX from other laboratories, are shared directly among the collaborators. Indeed, with some of the programs which have originated at UCSD and elsewhere, our off-campus collaborators frequently find it easier to use the SUMEX versions because of the interactive computing environment and ease of access. Advice, progress reports, new ideas, general information, etc. are communicated via the message and/or bulletin board facilities.

B. Interaction with Other SUMEX-AIM Projects

Our interactions with other SUMEX-AIM projects have been mostly in the form of personal contacts. We have strong ties to the MYCIN, AGE and MOLGEN projects and keep abreast of research in those areas on a regular basis through informal discussions. The SUMEX-AIM workshops provide an excellent opportunity to survey all the projects in the community. Common research themes, e.g. knowledge-based systems, as well as alternate problem-solving methodologies were particularly valuable to share.

C. Critique of Resource Services

The SUMEX facility provides a wide spectrum of computing services which are genuinely useful to our project -- message handling, file management, Interlisp, Fortran and text editors come immediately to mind. Moreover, the staff, particularly the operators, are to be commended for their willingness to help solve special problems (e.g., reading tapes) or providing extra service (e.g. immediate retrieval of an archived file). We would also like to commend the staff for its extensive help in setting up a link between SUMEX and Dr. Langridge's group at UCSF. Such cooperative behavior is rare in computer centers.

There are several facilities we wish to single out as particularly useful in furthering our research goals. Since the members of the project are physically distant, the MSG program is very useful. Similarly, the file system, the ARCHIVE facility, and the general ease of getting backup files from the operator greatly aid our efforts at coordinating the efforts of collaborators using many large data sets and programs. The crystallographers in the project find SUMEX to be a friendly environment which allows them to do their work with a minimum of dealing with operating system details.

It has become increasingly evident, however, that as CRYSALIS expands, the facility cannot provide enough machine cycles during prime time to support the implementation and debugging of new features. For example, our segment-labeling preprocessor requires about an hour of machine time per 100 residues of protein (this is typically five to eight hours of terminal time during working hours) even when the Lisp code is compiled.

III. <u>USE OF SUMEX DURING THE REMAINING GRANT PERIOD</u> (8/79 - 7/81)

A. Long-Range Goals

Our short term goals are to build up the knowledge base to the point where it can solve a small, known protein from "live" data. This will probably entail the implementation of at least seven tasks. By this point we should also have a package of data-reduction programs suitable for export to interested crystallographers.

Our long range goals are the exploitation of the rule-based control structure for investigating alternative problem-solving strategies, the investigation of modes of explanation of the program's reasoning steps, and the expansion and generalization of the system to cover a wider range of input data.

B. Justification for Continued Use of SUMEX

We feel that SUMEX is the ideal vehicle for further research on CRYSALIS. While some of our work is numerical in nature and uses such facilities as FORTRAN, our main interest is in artificial intelligence. Besides being an expert system of use to the crystallographic community, CRYSALIS is an exploration of the general signal processing problem. We are vitally concerned with issues such as proper architecture for using a wide variety of heuristics effectively and hypothesis formation when both data and model are poor. The utility of our work to the AI community is partially demonstrated by the development of the AGE project, an extension of Ms. Nii's early work on CRYSALIS.

This project progresses by the collaboration of several physically-separated groups. SUMEX provides a unique resource, an electronic community of researchers in our field, through the many systems such as net mail, country-wide access, and community workshops. We feel that CRYSALIS would not be possible outside of such a community.

C. Needs and Plans for Other Computing Resources

Our major need for other computing resources is for graphical display of our data and results. This need will be met by use of Dr. Langridge's Evans and Sutherland Picture System at UCSF and Dr. Johnson's raster-based graphics system at ORNL. The major impediment is SUMEX's current inability to support data transfer to other machines at more than 1200 baud. We are attempting to link SUMEX to UCSF by using FTP over the ARPAnet to the LBL machine and then use an existing link from LBL to UCSF.

We will make minor use of the Stanford Computer Science Department's SCORE machine, mostly to run the SCRIBE text formatting program until such time as it is available on SUMEX.

D. Recommendations for Future Community and Resource Development

There are two recommendations we wish to make, the first and most important is to expand the computing power available to SUMEX users. CRYSALIS is an inherently-large problem. Proteins contain hundreds, to thousands of atoms which means large hypothesis structures, large quantities of data, and a compute-bound inference program. As the system grows to maturity, we expect increasingly serious problems with address space limitations and with machine cycle availability.

The second recommendation is that SUMEX develop some relatively inexpensive file transfer facility for machines not on the ARPAnet. Software for this already exists in the form of the TTYFTP program (or possible future programs like it, but in a more portable language), the development needed is in hardware and in the TENEX operating system so that transfer rates greater than 1200 baud can be achieved. We are motivated to recommend this not only by our own need for such a facility, but also by the belief that it would aid other collaborations involving SUMEX and outside computers (the SECS project for example), and aid in the dissemination of useful programs from the research setting of SUMEX to user laboratories.

II.A.1.8 RX Project

The RX Project: Deriving Medical Knowledge from Time-Oriented Clinical Databases

Robert L. Blum, M.D.
Department of Computer Science
Stanford University

Gio C. M. Wiederhold, Ph.D.
Departments of Computer Science and Electrical Engineering
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Technical Goals

Introduction:

Medical and Computer Science Goals

The objective of the RX Project is to develop a medical information system capable of accurately deriving knowledge of the course and consequences of treatment of chronic diseases from a large collection of stored patient records.

Computerized clinical databases and automated medical records systems have been under development throughout the world for at least a decade. Among the earliest of these endeavors was the ARAMIS Project, (American Rheumatism Association Medical Information System) under development at Stanford by Dr. James Fries and his colleagues since 1969. A prototype ambulatory records system was generalized in the early 1970's by Prof. Gio Wiederhold and Stephen Weyl in the form of a Time-Oriented Database (TOD) System. The TOD System, run on the IBM 370/3033 at the Stanford Center for Information Processing (SCIP), now supports the ARAMIS Project as well as a host of other chronic disease databases which store patient data gathered at many institutions nation-wide. At the present time ARAMIS contains records of over 14,000 patients with a variety of rheumatologic diagnoses. Over 62,000 patient visits have been recorded, accounting for 50,000 patient-years of observation.

The fundamental objective of ARAMIS, the other TOD research groups, and all other clinical data bank researchers is to use the raw data which has been gathered by clinical observation in order to study the evolution and medical management of chronic diseases. Unfortunately, the process of reliably deriving knowledge from raw data has proven to be refractory to existing techniques because of problems stemming from the complexity of disease, therapy, and outcome definitions; the complexity of time relationships; complex causal relationships creating strong sources of bias; and problems of missing and outlying data.

A major objective of the RX Project is to explore the utility of symbolic computational methods and knowledge-based techniques at solving this problem of accurate knowledge inference from non-randomized, non-protocol patient records. A central component of RX is a knowledge base of medicine and statistics, organized as a hierarchy or taxonomic tree consisting of nodes with attached data and procedures. Nodes representing diseases and therapeutic regimens contain procedures which use a variety of time-dependent predicates to label patient records in the database, facilitating the retrieval of time-intervals of interest in the records. The database is then inverted so that each node or object in the knowledge base contains pointers to all time-intervals during which its definition is satisfied.

Nodes in the knowledge base also contain lists of other nodes which are causally related. These functional dependencies are used to infer causal pathways among nodes for purposes of selecting confounding variables which need to be controlled for in the study of a specific hypothesis. Causal pathways may also be used in an exploratory mode to assist in the discovery of new hypotheses.

To study a particular causal hypothesis the knowledge base also contains information on the applicability of various statistical procedures and procedures for applying them.

B. Medical Relevance and Collaboration

As a test bed for system development our focus of attention has been on the records of patients with systemic lupus erythematosus (SLE) contained in the Stanford portion of the ARAMIS Data Bank. SLE is a chronic rheumatologic disease with a broad spectrum of manifestations which can lead to death in the third decade of life. With many perplexing diagnostic and therapeutic dilemmas, it is a disease of considerable medical interest.

In the future we anticipate possible collaborations with other project users of the TOD System such as the National Stroke Data Bank, the Northern California Oncology Group, and the Stanford Divisions of Oncology and of Radiation Therapy.

The RX Project is a new research effort only in existence for about two years, and, hence the project is still in a developmental stage. The primary issues being addressed at this stage are those concerned with the specifics of knowledge representation.

We believe that this research project is broadly applicable to the entire gamut of chronic diseases which constitute the bulk of morbidity and mortality in the United States. Consider five major diagnostic categories which are responsible for approximately two thirds of the two million deaths per year in the United States: myocardial infarction, stroke, cancer, hypertension, and diabetes. Therapy for each of these diagnoses is fraught with controversy concerning the balance of benefits versus costs.

1) Myocardial Infarction: Indications for and efficacy of coronary artery bypass graft vs. medical management alone. Indications for long-term antiarrhythmics ... long-term anticoagulants. Benefits of cholesterol-lowering diets, exercise, etc.

2) Stroke: Efficacy of long-term anti-platelet agents, long-term anticoagulation. Indications for revascularization.

3) Cancer: Relative efficacy of radiation therapy, chemotherapy, surgical excision - singly or in combination. Optimal frequency of screening procedures. Prophylactic therapy.

4) Hypertension: Indications for therapy. Efficacy versus adverse effects of chronic antihypertensive drugs. Role of various diagnostic tests such as renal arteriography in work-up.

5) Diabetes: Influence of insulin administration on microvascular complications. Role of oral hypoglycemics.

Despite the expenditure of billions of dollars over recent years for randomized controlled trials (RCT's) designed to answer these and other questions, answers have been slow in coming. RCT's are expensive of funds and personnel. The therapeutic questions in clinical medicine are too numerous for each to be addressed by its own series of RCT's.

On the other hand, the data regularly gathered in patient records in the course of the normal performance of health care delivery is a rich and largely underutilized resource. The ease of accessibility and manipulation of these data afforded by computerized clinical data banks holds out the possibility of a major new resource for acquiring knowledge on the evolution and therapy of chronic diseases.

The goal of the research which we are pursuing on SUMEX is to increase the reliability of knowledge derived from clinical data banks with the hope of providing a new tool for augmenting knowledge of diseases and therapies as a supplement to knowledge derived from formal prospective clinical trials. Furthermore, the incorporation of knowledge from both clinical data banks and other sources into a uniform knowledge base should increase the ease of access by individual clinicians to this knowledge and thereby facilitate both the practice of medicine as well as the investigation of human disease processes.

C. Highlights of Research Progress

1. 1 July 1980 to 1 May 1981

Our predominant objective was to detail the overall conceptual framework for the knowledge base and to develop the extensive computational machinery necessary for retrieving, analyzing, and displaying defined time-intervals within patient records.

The RX Knowledge Base (KB):

The central component of RX is a knowledge base of medicine and statistics, organized as a frame-based, taxonomic tree consisting of units with attached data and procedures. Units representing diseases and therapies contain procedures which use a variety of time-dependent predicates to label the patient records, facilitating the retrieval of time-intervals of interest in the records. Other units representing statistical techniques are used to map hypotheses onto study designs and event definitions. Implementing the algorithms and data structures of this KB was one of the major tasks of the current year.

At the current time the RX KB contains about 200 units of which 75 contain definitions and other relevant .information pertaining to disease courses, effects of drugs, lab values, etc. This information compromises a small subset of medical knowledge dealing with some of the signs and symptoms of systemic lupus erythematosus (SLE) as well as the effects and indications of some drugs used for this disease. Other units contain machine-readable knowledge of statistical techniques needed for testing entered hypotheses. There are approximately 40 time-dependent functions used to map from the database values onto defined units.

The entire RX system currently contains approximately 400 INTERLISP functions accounting for 150 disk pages of code. The KB is about 60 disk pages. One disk page = 512 words * 36 bits per word. Also one disk page = approx. 1.5 typed pages on 8.5 by 11.5 inch paper.

Statistical Interfaces:

Once the relevant episodes have been defined and retrieved from the database they must be analyzed statistically. To do this we have recently adopted the IDL or Interactive Data-Analysis Language package developed at the Xerox Palo Alto Research Corp. IDL is a matrix manipulation language similar to APL and is built upon INTERLISP as is RX itself. The use of IDL for statistical analysis confers a tremendous advantage in that analyses are now highly interactive. IDL has completely supplanted our use of SPSS.

Time-Oriented Graphics Package:

This package enables data on an individual patient to be graphed over time, either linearly by visit or by calendar time with a "telescoping" capability. The program overlays graphs of both point data and data represented as episodes.

Study Editor:

Dr. Jerrold Kaplan, a research associate affiliated with the project, has implemented an additional package of programs which display to the clinician user those decisions which have been made by the knowledge base concerning which statistical techniques are to be employed, which variables are to be controlled for, and which time intervals are to be excluded. This affords the user with a means for seeing a sketch of the study plan before it is executed, and enables him to modify that plan.

Clinical Study: The Effect of Prednisone on Cholesterol

As a testbed for the prototype system we have been investigating the hypothesis that the steroid, prednisone, produces a significant elevation of plasma cholesterol. To test this hypothesis, the records of 50 patients with systemic lupus erythematosus (SLE) were transferred from the ARAMIS Database to SUMEX. Of these patients, 18 were found to have five or more cholesterol determinations and to have had sufficient variance in their prednisone regimens to be testable. The KB is used to elaborate a complex causal model for the prednisone/cholesterol hypothesis which is tested using a hierarchical multiple regression method with time-lagged values. The KB is used to determine sources of possible bias and to control for those variables in the regression or to eliminate corresponding time-intervals from records. An empirical Bayes method is used to average the estimated effects in patients with varying amounts of data.

The result, a highly statistically significant elevation of cholesterol by prednisone, will be submitted for publication during the coming year.

2. Research In Progress

Much work remains to be done in expanding the system software and in expanding the knowledge base. Current work is addressed to increasing the flexibility of the time-segmentation functions and enriching the data structures which encode relationships among objects.

We are trying to make increasingly general the class of medical hypotheses which the system can analyze automatically. This requires incorporating knowledge of additional statistical methods into the KB. We are also attempting to generalize our algorithms for selecting the set variables which may potentially confound a given hypothesis.. As a means for testing and expanding the system's capabilities we intend to perform several specific studies of importance in the management of the rheumatic diseases. Our study of the effect of prednisone on cholesterol was mentioned above. Other studies now being planned include the effect of chronic aspirin ingestion on liver function in rheumatoid arthritis, the specific incidence of infectious complications of steroids as a function of dose and duration, and the utility of various autoantibodies in the prediction of flares of SLE as compared to the utility of other indicators.

Finally, we are developing a methodology for discovering hypotheses of interest in the database using a heuristically guided search of non-parametric lagged cross correlations.

D. Publications

Blum, Robert L.: Discovery and Representation of Causal Relationships from a Time-Oriented Clinical Database: The RX Project. Stanford University Doctoral Dissertation (in press), 1981

Blum, Robert L.: Displaying Clinical Data from a Time-Oriented Database. (in press) Computers in Biology and Medicine, 1981

Wiederhold, Gio: Databases in Healthcare. Compendium Series on Technology in Healthcare, sponsored by the Healthcare Technology Center, Univ. of Missouri, Columbia, Mo., also available as Stanford CS Report 80-790

Blum, Robert L.: Automating the Study of Clinical Hypotheses on a Time-Oriented Data Base: The RX Project. Submitted for publication to MEDINFO80, Tokyo, Japan, Oct. 1980

Blum, Robert L., Wiederhold, Gio: Inferring Knowledge from Clinical Data Banks Utilizing Techniques from Artificial Intelligence. Proc. of The 2nd Annual Symp. on Computer Applications in Medical Care, pp. 303 to 307, IEEE, Washington, D.C., November 5-9, 1978

E. Funding Support Status

1) A Computer-Based System for Advising Physicians on Clinical Therapeutics
Robert L. Blum, M.D.: Awardee
Post-Doctoral Research Fellowship in Clinical Pharmacology
Pharmaceutical Manufacturers' Association Foundation
Total award: $32,500 (direct)
Term: July 1, 1978 to June 30, 1980

2) Integrating Medical Knowledge and Clinical Data Banks
Robert L. Blum, M.D.: Principal Investigator
National Library of Medicine, New Investigator Award
Total award: $90,000 (direct)
Term: July 1, 1979 to June 30, 1982

3) Integrating Medical Knowledge and Clinical Data Banks
Gio C. M. Wiederhold, Ph.D.: Principal Investigator
National Center for Health Services Research, Small Grants
Total award: $35,000 (direct)
Term: April 1, 1979 to March 31, 1981

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Collaborations

Since our project is relatively new, we do not yet have public versions of the programs. There is, however, a large sphere of collaboration which we expect in the future. Once the RX program is developed, we would anticipate collaboration with all of the ARAMIS project sites in the further development of a knowledge base pertaining to the chronic arthritides. The ARAMIS Project at SCIP is used by a number of institutions around the country via commercial leased lines to store and process their data. These institutions include the University of California School of Medicine, San Francisco and Los Angeles; The Phoenix Arthritis Center, Phoenix; The University of Cincinnati School of Medicine; The University of Pittsburgh School of Medicine; Kansas University; and The

University of Saskatchewan. All of the rheumatologists at these sites have closely collaborated with the development of ARAMIS, and their interest in and use of the RX project is anticipated. We hasten to mention that we do not expect SUMEX to support the active use of RX as an on-going service to this extensive network of arthritis centers, but we would like to be able to allow the national centers to participate in the development of the arthritis knowledge base and to test that knowledge base on their own clinical data banks.

B. Interactions with Other SUMEX-AIM Projects

Several of the concepts incorporated into the design of the RX Project have been inspired by other SUMEX-AIM Projects. The RX knowledge base is similar to the Units Package of the MOLGEN PROJECT. The production rule inference mechanism used by us is similar to that in the MYCIN Project.

Several programs developed by the MYCIN group are regularly used by RX. These include disk hash file facilities, text editing facilities, and miscellaneous LISP functions. Regular communication on programming details is facilitated by the on-line mail system.

C. Critique of Resource Management

The SUMEX KI-10 has been severely overloaded for at least a year. Working in LISP is impossible during the day and is even difficult at times which were formerly low utilization times. This has forced us to rely increasingly on other local computation facilities.

The SUMEX resource management, per se, has always been accessible and cooperative in trying to provide our project with adequate resources subject to prevailing constraints.

III. RESEARCH PLANS

A. Project Goals and Plans

The overall goal of the RX Project is to develop a computerized medical information system capable of accurately extracting medical knowledge pertaining to the therapy and evolution of chronic diseases from a database consisting of a collection of stored patient records.

1. Short-Term Goals

Goals for the year August, 1980 to July, 1981 have been detailed in section IC. above on research in progress. To summarize that section, our main short-term goal is to generalize and refine our methods for labeling and retrieving time-intervals or episodes from individual patient records and to generalize the class of hypotheses which the system is capable of analyzing. This requires further refinements in RX's algorithms for choosing and controlling for variables which may potentially confound an hypothesis of interest.

2. Long-Range Goals: August, 1981 to July, 1986

There are two inter-related long-range goals of the RX Project: 1) automatic discovery of knowledge in a large time-oriented database and 2) provision of assistance to a clinician who is interested in testing a specific hypothesis. These tasks overlap to the extent that some of the algorithms used for discovery are also used in the process of testing an hypothesis.

We hope to make these algorithms sufficiently robust that they will work over a broad range of hypotheses and over a broad spectrum of data distributions in the patient records.

B. Justification for Continued Use of SUMEX

Computerized clinical data banks possess great potential as tools for assessing the efficacy of new diagnostic and therapeutic modalities, for monitoring the quality of health care delivery, and for support of basic medical research. Because of this potential, many clinical data banks have recently been developed throughout the United States. However, once the initial problems of data acquisition, storage, and retrieval have been dealt with, there remains a set of complex problems inherent in the task of accurately inferring medical knowledge from a collection of observations in patient records. These problems concern the complexity of disease and outcome definitions, the complexity of time relationships, potential biases in compared subsets, and missing and outlying data. The major problem of medical data banking is in the reliable inference of medical knowledge from primary observational data.

We see in the RX Project a method of solution to this problem through the utilization of knowledge engineering techniques from artificial intelligence. The RX Project, in providing this solution, will provide an important conceptual and technologic link to a large community of medical research groups involved in the treatment and study of the chronic arthritides throughout the United States and Canada, who are presently using the ARAMIS Data Bank through the SCIP facility via TELENET.

Beyond the arthritis centers which we have mentioned in this report, the TOD (Time-Oriented Data Base) User Group involves a broad range of university and community medical institutions involved in the treatment of cancer, stroke, cardiovascular disease, nephrologic disease, and others. Through the RX Project, the opportunity will be provided to foster national collaborations with these research groups and to provide a major arena in which to demonstrate the utility of artificial intelligence to clinical medicine.

<u>SUMEX as a Resource</u>:

To discuss SUMEX as a resource for program development, one need only compare it to the environment provided by our other resource, the IBM 370/3033 installation at SCIP - the major computing resource at Stanford. Of the programs which we use daily on SUMEX -INTERLISP, MSG, TVEDIT, BBD, LINK - there is nothing even approaching equivalence on the 370, despite its huge user community. These programs greatly facilitate communication with other researchers in the SUMEX community, documentation of our programs, and the rapid interactive development of the programs themselves. The development of a program involving extensive symbolic processing and as large and complex as RX at the SCIP facility, would require a staff many times as large as ours. The SUMEX environment greatly increases the productive potential of a research group such as ours to the point where a large project like RX becomes feasible.

Computation Resources Required by RX:

Disk Allocation:

RX requires the use of two large data files which need to be kept on-line: the patient database (DB) and the knowledge base (KB). In the course of testing a hypothesis several other files are used: inverted files, source files for statistical processing, LISP SYSOUT files, etc. Our current total disk allocation of 1500 pages for all RX group members has been just adequate. In the future, with anticipated expansions in numbers of patients and size of the KB, we intend to request an increase of our total allocation to 2000 pages.

C. Other Computational Resources

It is clear that the scope of potential application of the RX Project is large. Within the term of the SUMEX-AIM grant projected through July, 1986, we anticipate the involvement of several of the national ARAMIS collaborating institutions in developing and testing arthritis knowledge bases which reflect their own patient populations and therapeutic biases. The current SUMEX machine configuration will not be able to support this national interaction because the central processors of the KI-10 are already taxed to the limit. Ours is among the SUMEX groups which would greatly benefit by the addition of one or more PDP-10 compatible machines, which could provide support to our anticipated national user community. Another resource which would be highly desirable is a faster and more reliable means for transferring data interactively between SUMEX and the SCIP IBM 370. Our current method utilizes a 2400 baud line with transmission from SCIP to SUMEX only, and is fraught with a high error rate. The addition of a reliable local network facility would greatly facilitate our ability to transfer patient files from SCIP to SUMEX.

D. Recommendations for Resource Development

SUMEX is heavily loaded everyday and almost every evening. Program research is next to impossible during those periods. Program development would be greatly facilitated by the addition of any resources which lessened this loading: upgrading the current machine to a KL or adding core to decrease page swapping.

II.A.2 National AIM Projects

The following group of projects is formally approved for access to the AIM aliquot of the SUMEX-AIM resource or the Rutgers-AIM resource. Their access is based on review by the AIM Advisory Group and approval by the AIM Executive Committee.

II.A.2.1 Acquisition of Cognitive Procedures (ACT)

Acquisition of Cognitive Procedures (ACT)

Dr. John Anderson
Carnegie-Mellon University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

To develop a production system that will serve as an interpreter of the active portion of an associative network. To model a range of cognitive tasks including memory tasks, inferential reasoning, language processing, and problem solving. To develop an induction system capable of acquiring cognitive procedures with a special emphasis on language acquisition and problem-solving skills.

B. Medical Relevance and Collaboration

1. The ACT model is a general model of cognition. It provides a useful model of the development of and performance of the sorts of decision making that occur in medicine.

2. The ACT model also represents basic work in AI. It is in part an attempt to develop a self-organizing intelligent system. As such it is relevant to the goal of development of intelligent artificial aids in medicine.

We have been evolving a collaborative relationship with James Greeno and Allan Lesgold at the University of Pittsburgh. They are applying ACT to modeling the acquisition of reading and problem solving skills. We have made ACT a guest system within SUMEX. ACT is currently at the state where it can be shipped to other INTERLISP facilities. We have received a number of inquiries about the ACT system. ACT is a system in a continual state of development but we periodically freeze versions of ACT which we maintain and make available to the national AI community.

C. Highlights of Research Progress

Our ACTF system is a production system that operates in a semantic network data base. Our learning work has been focused on ways of increasing the power of production systems for performing various tasks. One class of learning mechanisms concern what we call knowledge compilation. This involves automatic mechanisms for creating productions that directly perform behavior that formerly required interpretative processing of knowledge in the semantic network. These compilation mechanisms also model the process by which human experts develop special purpose procedures to deal with the different types of problems that occur in their domain of expertise.

Another class of learning mechanisms are concerned with tuning existing procedures so that they apply more appropriately. There are various mechanisms concerned with extending or generalizing the range of application of a procedure. In the past year we have been working at reducing these different generalization processes to a common partial matching process. In addition to generalization, tuning occurs in the ACTF system by means of discrimination and composition. Discrimination is a process for restricting the range of applicability of a production. Composition attempts to build macro-operators out of a series of productions.

The third direction of our learning work has been concerned with developing a flexible strength-based set of conflict resolution rules. Here we are concerned with modelling the gradual improvement seen in human cognitive skills and also providing the system with the resilience so that it can recover from noise and changes in environmental contingencies.

We have been applying this theory in detail to a simulation of how students acquire proof skills in geometry. We have a more or less thorough analysis of how students learn new postulates of geometry; we initially use these postulates in an interpretative fashion, integrating them with prior knowledge; how they compile special purpose procedures that directly apply this knowledge to proof generation; and how these procedures become tuned with practice. This application has provided strong evidence for most of the learning developments in the ACT system. It has also forced us to develop formalisms for how planning and problem-solving should be structured within a production-system framework.

D. List of Project Publications

[1] Anderson, J.R. Language, Memory, and Thought. Hillsdale, N.J.: L. Erlbaum, Assoc., 1976.

[2] Kline, P.J. & Anderson, J.R. The ACTE User's Manual, 1976.

[3] Anderson, J.R., Kline, P. & Lewis, C. Language processing by production systems. In P. Carpenter and M. Just (Eds.). Cognitive Processes in Comprehension. L. Erlbaum Assoc., 1977.

[4] Anderson, J.R. Induction of augmented transition networks. Cognitive Science, 1977, 125-157.

[5] Anderson, J.R. & Kline, P. Design of a production system. Paper presented at the Workshop on Pattern-Directed Inference Systems, Hawaii, May 23-27, 1977.

[6] Anderson, J.R. Computer simulation of a language acquisition system: A second report. In D. LaBerge and S.J. Samuels (Eds.). Perception and Comprehension. Hillsdale, N.J.: L. Erlbaum Assoc., 1978.

[7] Anderson, J.R., Kline, P.J., & Beasley, C.M. A theory of the acquisition of cognitive skills. In G.H. Bower (Ed.). Learning and Motivation, Vol. 13. New York: Academic Press, 1979.

[8] Anderson, J.R., Kline, P.J., & Beasley, C.M. Complex Learning. In R. Snow, P.A. Frederico, & W. Montague (Eds.). Aptitude, Learning, and Instruction: Cognitive Processes Analyses. Hillsdale, N.J.: Lawrence Erlbaum Assoc., 1980.

[9] Anderson, J.R. & Kline, P.J. A learning system and its psychological implications. In the Proceedings of the Sixth International Joint Conference on Artificial Intelligence, 1979.

[10] Reder, L.M. & Anderson, J.R. Use of thematic information to speed search of semantic nets. Proceedings of the Sixth International Joint Conference on Artificial Intelligence, 1979, 708-710.

[11] Neves, D.M. & Anderson, J.R. Knowledge compilation: Mechanisms for the automatization of cognitive skills. In J.R. Anderson (Ed.), Cognitive Skills and their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

[12] Anderson, J.R., Greeno, J.G., Kline, P.J., & Neves, D.M. Acquisition of Problem-solving skill. In J. R. Anderson (Ed.), Cognitive Skills and their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

E. Funding Support

An information-processing analysis of learning in geometry
John R. Anderson, Principal Investigator
National Science Foundation (IST-80-15357)
$186,000 Feb 15, 1981 - Feb 15, 1984

II. INTERACTION WITH THE SUMEX-AIM RESOURCE

A. Collaborations, Interactions, and Sharing of Programs via SUMEX.

We have received and answered many inquiries about the ACT system over the ARPANET. This involves sending documentations, papers, and copies of programs. The most extensive collaboration has been with Greeno and Lesgold who are also on SUMEX (see the report of the Simulation of Comprehension Processes project). There is an ongoing effort to assist them in their research. Feedback from their work is helping us with system design.

We find the SUMEX-AIM workshops (those that we could manage to attend) ideal vehicles for updating ourselves on the field and for getting to talk to colleagues about aspects of their work of importance to us.

Due to memory space problems encountered by ACT we expect that soon we will need to make use of the smaller version of INTERLISP developed at SUMEX for use in the CONGEN program.

B. Critique of Resource Management

The SUMEX-AIM resource has been well suited for the needs of our project. We have made the most extensive use of the INTERLISP facilities and the facilities for communication on the ARPANET. We have found the SUMEX personnel extremely helpful both in terms of responding to our immediate emergencies and in providing advice helpful to the long-range progress of the project. Despite the fact that we are not located at Stanford, we have not encountered any serious difficulties in using the SUMEX system; in fact, there are real advantages in being in the Eastern time zone where we can take advantage of the low load on the system during the morning hours. We have been able to get a great deal of work done during these hours and try to save our computer-intensive work for this time.

Two location changes by the ACT project (from Michigan to Yale in the summer of 1976 and from Yale to Carnegie-Mellon in the summer of 1978) have demonstrated another advantage of working on SUMEX: In both cases we were back to work on SUMEX the day after our arrival.

III. RESEARCH PLANS (8/80-7/86)

A. Project Goals and Plans

Our long-range goals are: (1) Continued development of the ACT system; (2) Application of the system to modeling of various cognitive processes; (3) Dissemination of the ACT system to the national AI community.

This is a period of major evolution for the ACT theory. We have been developing three special versions of the ACTF learning that allow us to more efficiently simulate learning in three domains: proving theorems in geometry, speaking a new language, and writing programs in LISP. We are also performing special purpose simulations of the processes of spreading activation in memory retrieval and of pattern-matching processes in reading. We will be assimilating our experiences with these special purpose simulations in putting forth a major revision of the ACT theory. A research monograph is being written setting forth this theory and is scheduled for completion in late 1982. Subsequent to the writing of this monograph we intend to create an ACTG successor to ACTF that will embody the new conceptions.

B. Justification for Continued Use of SUMEX:

Our goal for the ACT system is that it should serve as a ready-made "programming language" available to members of the cognitive science community for assembling psychologically-accurate simulations of a wide range of cognitive processes. Our intention and ability to provide such a resource justifies our use of the SUMEX facility. This facility is designed expressly for the purpose of developing and supporting such national AI resources and is, in this regard, clearly superior to the facilities we have available locally from the Carnegie-Mellon computer science department. Among the most important SUMEX advantages are the availability of INTERLISP on a machine accessible by either the ARPANET or TYMNET and the existence of a GUEST login. It appears that, at least for the time being, ACT has no hope of being a national resource unless it resides at SUMEX and, given the local unavailability of a network-accessible INTERLISP, it would even be very difficult to shift any significant portion of our development work from SUMEX to CMU.

C. Needs and Plans for Other Computational Resources

Carnegie-Mellon's plans to begin upgrading its PDP-10 hardware to emerging state-of-the-art machines (VAX, LISP machines, etc.) promises to provide a excellent resource eventually, and we hope to have access to that resource as it develops. However, given that a considerable amount of software development will be required, a sophisticated LISP system such as INTERLISP is not likely to be available on this hardware in the near future.

D. Comments and Suggestions for Future Resource Goals

We are beginning to feel squeezed by various limitations of the SUMEX facility. The problem of peak load is quite serious. We have also been struggling with the address limitations of the current INTERLISP which is made more grievous by the amount of space INTERLISP requires. The computation time and address space limitations have meant that we have not been able to pursue certain projects that we would have otherwise. We applaud any efforts to increase computational power, to increase the address space of INTERLISP (e.g. VAXes), or to create significantly more space efficient versions of INTERLISP.

II.A.2.2 CADUCEUS Project (INTERNIST)

CADUCEUS Project (*)

J. D. Myers, M.D. and H. Pople, Ph.D.
University of Pittsburgh
Pittsburgh, Pennsylvania

I. SUMMARY OF RESEARCH PROGRAM

A. Medical Rationale

The principal objective of this project is the development of a high-level computer diagnostic program in the broad field of internal medicine as an aid in the solution of complex and complicated diagnostic problems. To be effective, the program must be capable of multiple diagnoses (related or independent) in a given patient.

A major achievement of this research undertaking has been the design of a program called INTERNIST-I, along with an extensive medical data base now encompassing over 500 diseases and some 3450 individual manifestations of disease.

Although this consultative program is designed primarily to aid skilled internists in complicated medical problems, the program may have spin-off as a diagnostic and triage aid to physicians assistants, rural health clinics, military medicine and space travel.

Development of the INTERNIST-I system was begun about ten years ago. The system was successfully demonstrated for the first time in 1974 and has been used since that time in the analysis of hundreds of clinical problems.

A major point of departure for the design of the original INTERNIST program was the realization that the task of clinical decision making in internal medicine is an ill-structured problem. In other domains, the task of diagnosis is often viewed as one of pattern recognition or discrimination: there is available a predefined collection of possible classifications (characterizing disease entities or clinical states), one and only one of which is considered possible in the case being studied. A diagnostic problem solver dealing with such a well structured domain has the fairly straightforward task of selecting that one of this fixed set of alternatives which best fits the facts of the case. Many statistical,

---

(*) For a variety of reasons, including a request from an agency alleging a prior claim on the name, future generations of the diagnostic program originally called INTERNIST will subsequently be referred to as CADUCEUS. This universal symbol of the medical profession seems appropriate to the expanded role we see for this type of program in the years to come. To avoid confusion in this report, the original program will continue to be called INTERNIST-I while references to the successor system, originally called INTERNIST-II, will now employ the new name.

pattern recognition, and algorithmic techniques have been employed successfully in performing computer aided diagnosis in these well structured clinical problem domains.

Primarily because complex cases often involve two or more concurrently active disease processes, no set of exhaustive and mutually exclusive classifications can be developed to structure the diagnostic problem in internal medicine. In principle, it might be argued that this more complex problem domain could be reduced to a simple discrimination task if, in addition to the individual disease entities, one includes appropriate multiple disease complexes in the set of allowable patient descriptors. However, since our experience indicates that as many as ten or twelve individual descriptors may apply in a complex clinical problem, and considering that there are a thousand or more individual descriptors of interest in Internal Medicine, the prospect of recording explicitly all possible multiple disease classifications is clearly infeasible.

Our thesis is that, in the absence of explicit structure derived from the problem domain, the successful clinician engages in heuristic imposition of structure so that effective problem solving strategies might be selected and employed for decision making relative to the postulated problem structure.

In INTERNIST-I, this concept of heuristic imposition of structure is expressed primarily by means of a novel "problem-formation" heuristic. In effect, the program composes dynamically, on the basis of evidence provided, what in context constitutes a presumed exhaustive and mutually exclusive subset of disease entities that can explain, more or less equally well, some significant subset of the observed findings in a clinical case. This heuristic problem structuring procedure is invoked repeatedly during the course of a diagnostic consultation in order to deal sequentially with the component parts of a complex clinical problem.

Because this program is intended to serve a consulting role in medical diagnosis, it has been challenged with a wide variety of difficult clinical problems: cases published in the medical journals, cpc's, and other interesting and unusual problems arising in the local teaching hospitals. In the great majority of these test cases, the problem-formation strategy of INTERNIST-I has proved to be effective in sorting out the pieces of the puzzle and coming to a correct diagnosis, involving in some cases as many as a dozen disease entities.

On the basis of this extensive test of the initial INTERNIST-I system, it has become clear that many aspects of the system's performance could be significantly enhanced if it would be possible to deal with the various component problems and their interrelationships simultaneously. This has led to the design of CADUCEUS, a system embodying strategies of concurrent problem-formation which we expect will yield more rapid convergence to the correct diagnosis in many cases, and in at least some cases provide more acceptable diagnostic behavior.

B. Medical Relevance and Collaboration

The program inherently has direct and substantial medical relevance.

The institution of collaborative studies with other institutions has been deferred pending completion of the programs and knowledge base enhancements required for CADUCEUS. The installation of our own, dedicated VAX computer expected this summer will considerably aid future collaboration.

C. Highlights of Research Progress

Accomplishments This Past Year:

a) Prototypic computer programs have been written to operate CADUCEUS in the new diagnostic mode. The entire medical data base for the liver and biliary tract diseases has been reorganized into a form compatible with and utilizable by the CADUCEUS programs. Implementation of this work is pending the installation of the VAX computer when all of the programs must be written or rewritten using the FRANZ-LISP language.

b) The medical knowledge base comprising now just over 500 individual diseases and some 3450 manifestations of disease and hundreds of thousands of individual medical "facts," has been cumulative for the past eight years. Much effort has been spent during the past year in updating several dozens of diseases, most of which had been profiled years ago, and in establishing uniformity and consistency in this vast knowledge base. In addition, 17 new diseases have been profiled. The pediatric knowledge base has been expanded and now includes 78 diseases.

c) INTERNIST up to this time has been deficient in anatomic knowledge, particularly in topographical anatomy and anatomic laterality. An anatomic knowledge base beginning with neuroanatomy (the most complex) is being built for later incorporation into CADUCEUS. The knowledge base for the peripheral nervous system and the spinal cord is largely completed. The topographical anatomy of the abdomen and thorax are partially completed.

Research in Progress:

There are five major components to the continuation of this research project:

1) The completion, continued updating, refinement and testing of the extensive medical knowledge base required for the operation of INTERNIST-I.

2) The completion and implementation of the improved diagnostic consulting program, CADUCEUS, which has been designed to overcome certain performance problems identified during the past five years' experience with the original INTERNIST-I program.

3) Institution of field trials of CADUCEUS on the clinical services in internal medicine at the Health Center of the University of Pittsburgh.

4) Expansion of the clinical field trials to other university health centers which have expressed interest in working with the system.

5) Adaptation of the diagnostic program and data base of CADUCEUS to subserve educational purposes and the evaluation of clinical performance and competence.

Current activity is devoted mainly to the first two of these, namely, the continued development of the medical knowledge base, and the implementation of the improved diagnostic consulting program (CADUCEUS). The development of the anatomic knowledge base is mentioned above.

Doctor Gordon Banks, a skilled neurologist who also has a Ph.D. in physics and considerable experience in computing, will be joining the team as of July 1, 1981 and will provide manpower and expertise for the further development of the sizeable and important neurological component of the medical knowledge base and its manipulation by the CADUCEUS programs.

D. List of Relevant Publications

1. Pople, H.E. "The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning", Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Boston, August 1977.

2. Pople, H.E. "On the Knowledge Acquisition Process in Applied A.I. Systems", Report of Panel on Applications of A.I., Proceedings of Fifth International Joint Conference on Artificial Intelligence, 1977.

3. Pople, H.E., Myers, J. D. & Miller, R.A. "The DIALOG Model of Diagnostic Logic and its Use in Internal Medicine, Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, September 1975.

4. Pople, H.E. "Artificial Intelligence Approaches to Computer-Based Medical Consultation, Proceedings IEEE Intercon, New York, 1975.

5. Myers, J. D. "The Process of Clinical Diagnosis and Its Adaptation to the Computer," in The Logic of Discovery and Diagnosis in Medicine, University of Pittsburgh Series in the Philosophy and History of Science, University of California Press (in press).

6. Myers, J. D., Pople, H. E. & Miller, R. A. "INTERNIST: Can Artificial Intelligence Help?" in Clinical Decision and Laboratory Use, University of Minnesota Press (in press).

7. Pople, H. E. "Coming to Grips with the Multiple Diagnosis Problem," in Computer-Assisted Decision Making Using Clinical and Paraclinical (Laboratory) Data. B. Statland & S. Bauer (eds.) Mediad Inc., Tarrytown, N. Y., 1980, pp. 81-88. Reprinted in The Logic of Discovery and Diagnosis in Medicine, University of Pittsburgh Series in the Philosophy and History of Science, University of California Press (in press).

8. Pople, H. E. "Heuristic Methods for Imposing Structure on Ill-Structured Problems: The Structuring of Medical Diagnostics," in Artificial Intelligence in Medicine, AAAS Symposium Series, Westview Press (forthcoming 1981).

E. Funding support

1. Clinical Decision Systems Research Resource

   Harry E. Pople, Jr., Ph.D.
   Associate Professor Business

   Jack D. Myers, M.D.
   University Professor (Medicine)
   University of Pittsburgh

   Division of Research Resources
   National Institutes of Health

   2 R24 RR01101-04

   07/01/80 - 06/30/85
   $1,607,717

   07/01/80 - 06/30/81
   $465,199

2. INTERNIST: A Computer-Based Diagnostic Consultant

   Harry E. Pople, Jr., Ph.D.
   Associate Professor of Business

   Jack D. Myers, M.D.
   University Professor (Medicine)
   University of Pittsburgh

   National Library of Medicine
   National Institutes of Health

   1 RO1 LM03710-01

   07/01/80 - 06/30/85
   $817,884

   07/01/80 - 06/30/81
   $148,458

3. New Computer-Based Patient Case Simulator

Randolph A. Miller, M.D.
Associate Professor of Medicine
University of Pittsburgh

National Library of Medicine - New Investigator
National Institutes of Health

1 R23 LM03589-01

07/01/80 - 06/30/83
$89,350

07/01/80 - 06/30/81
$32,750

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A, B. Collaborations and Medical Use of Program Via SUMEX

CADUCEUS remains in a stage of research and development. As noted above, we are continuing to develop better computer programs to operate the diagnostic system, and the knowledge base cannot be used very effectively for collaborative purposes until it has reached a critical stage of completion. These factors have stifled collaboration via SUMEX up to this point and will continue to do so for the next year or two. In the meanwhile, through the SUMEX community there continues to be an exchange of information and states of progress. Such interactions particularly take place at the annual AIM Workshop.

C. Critique of Resource Management

SUMEX has been an excellent resource for the development of CADUCEUS. Our large program is handled efficiently, effectively and accurately. The staff at SUMEX have been uniformly supportive, cooperative, and innovative in connection with our project's needs.

III. RESEARCH PLANS (7/81-6/86)

A. Project Goals and Plans

The prototype CADUCEUS programs and the trial reorganization of the liver and biliary tract diseases will be installed in the VAX over the summer and fall of this year. As rapidly as possible and pending further refinement and reorganization from experience with the new system, the remainder of the medical knowledge base will be entered. Local and later collaborative field trials must necessarily be postponed until this development has been accomplished.

At least 200 important medical diseases remain to be programmed. Renewed effort in this direction is now being expanded now that other tasks have been surmounted. Expanded efforts in the fields of neurology and pediatrics are included as described above.

B. Justification and Requirements for Continued SUMEX Use

Our use of SUMEX will obviously decline upon the installation of our VAX. Nevertheless, the excellent facilities of SUMEX are expected to be used for certain developmental work. It is intended, further, to keep INTERNIST-I at SUMEX for comparative use as CADUCEUS is developed here. Our team hopes to remain as a component of the SUMEX community and to share experiences and developments.

C. Needs and Plans for Other Computing Resources beyond SUMEX-AIM

Our predictable needs in this area will be met by the dedicated VAX computer soon to be installed.

D. Recommendations for Future Community and Resource Development

Whether a program like CADUCEUS, when mature, will be better operated from centralized, larger computers or from the developing self-contained personal computer is difficult to predict. For the foreseeable future it would seem that centralized, advanced facilities like SUMEX will be important in further program development and refinement.

II.A.2.3 Hierarchical Models of Human Cognition

Hierarchical Models of Human Cognition (CLIPR Project)

Walter Kintsch and Peter G. Polson
University of Colorado
Boulder, Colorado

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The two CLIPR projects have made substantial progress in their research in this past year. This progress is almost completely due to our access to the SUMEX facility. The prose comprehension group has completed one major project, and is currently interacting with other SUMEX projects with the goal of building a prose comprehension model that reflects state-of-the-art knowledge from psychology and artificial intelligence.

The main activity of the planning group during the last year has been the detailed analysis of thinking-out-loud protocols collected from both expert and novice software designers. SUMEX facilities have been used to store, edit, and reformat the raw protocols to facilitate later analysis. Results of successive analyses are then input to SUMEX, and SUMEX facilities are used to collate the various results.

Technical Goals:

The CLIPR project consists of two subprojects. The first, the text comprehension project, is headed by Walter Kintsch and is a continuation of work on understanding of connected discourse that has been underway in Kintsch's laboratory for over seven years. The second, the planning project, is headed by Peter Polson of the University of Colorado and Michael Atwood of Science Applications Incorporated, Denver, and is studying the processes of planning using software design tasks.

The goal of the prose comprehension project is to develop a computer system capable of the meaningful processing of prose. This work has been generally guided by the prose comprehension model discussed by Kintsch and van Dijk (1978), although our programming efforts have identified necessary clarifications and modifications in that model (Miller & Kintsch, 1980a). Our more recent research (Miller & Kintsch, 1980b) has emphasized the importance of knowledge and knowledge-based processes in comprehension, and we are accordingly working with the AGE and UNITS groups at SUMEX toward the development of a knowledge-based, blackboard model of prose comprehension. We hope to be able to merge the substantial artificial intelligence research on these systems with psychological interpretations of prose comprehension, resulting in a computational model that is also psychologically respectable.

The primary goal of the planning project is the development of a model of human performance on software design tasks. We intend to begin by modeling protocols of experts on solving a particular problem, eventually extending the model to other levels of experience and problems. We propose a two-pronged attack on the process of developing a model.

The first is to develop a deeper understanding of our protocol data, to increase our knowledge of the details of the planning processes and the knowledge structures that experts use in the process of planning. We have developed a method of protocol analysis that essentially involves the transforming of the protocol into a low level theoretical description of the processes used to solve the design problem. We have assumed a very simplified version of a blackboard model that is described in Atwood and Jeffries (1980). We currently carry out our analysis by hand, developing a form of this low level model for each protocol. However, much of the activities involved in developing this model are clerical in nature and involve the categorization of segments of a verbal protocol and then the reorganization of the categorized information. Much of this work can be automated, and we propose to develop a program that will facilitate our protocol analysis and the development of the low level models that we use to describe the behavior of individual subjects.

Our second and much longer term objective is the development of a substantive model in AGE that can simulate the design processes. We feel that the software tools that are being developed at SUMEX -- in particular AGE and the UNITS package -- will dramatically facilitate our ability to develop this substantive model. Furthermore, current theoretical ideas about both the process of design and the representation of knowledge involved in developing a design have been strongly influenced by the MOLGEN project at SUMEX (Stefik, 1980).

B. Medical Relevance and Collaboration

The text comprehension project impacts indirectly on medicine, as the medical profession is no stranger to the problems of the information glut. By adding to the research on how computer systems might understand and summarize texts, and determining ways by which the readability of texts can be improved, medicine can only be helped by research on how people understand prose. Development of a more thorough understanding of the various processes responsible for different types of learning problems in children and the corresponding development of a successful remediation strategy would also be facilitated by an explicit theory of the normal comprehension process.

Note that our goal of a blackboard model is particularly relevant to the understanding of learning difficulties. One important aspect of a blackboard model is the separation of cognitive processes into a set of interacting subprocesses. Once such subprocesses have been identified and constructed, it would be instructive to observe the model's performance when certain of these processes are facilitated or inhibited. Many researchers have shown that there are a variety of cognitive deficits (insufficient short-term memory capacity, poor long-term memory retrieval, and such) that can lead to reading problems. Having a blackboard model in which the power of individual components could be manipulated would be a significant step in determining the nature of such reading problems.

The planning project is attempting to gain understanding of the cognitive mechanisms involved in design and planning tasks. The knowledge gained in such research should be directly relevant to a better understanding of the processes involved in medical policy making and in the design of complex experiments. We are currently using the task of software design to describe the processes underlying more general planning mechanisms that are also used in a large number of task oriented environments like policy making.

Both the text comprehension project and the planning project involve the development of explicit models of complex cognitive processes; cognitive modelling is a stated goal of both SUMEX and research supported by NIMH.

The on-going development of the prose comprehension model would not be possible without our collaboration with the AGE and UNITS research groups. We look forward to a continued collaboration, with, we hope, mutually beneficial results. Several other psychologists have either used or shown an interest in using an early version of the prose comprehension model, including Alan Lesgold of SUMEX's SCP project, who is exporting the system to the LRDC vax. Needless to say, all of this interaction has been greatly facilitated by the local and network-wide communication systems supported by SUMEX. There has been considerable communication between members of the prose comprehension and AGE/UNITS groups as program bugs have been discovered and corrected; the presence of a mail system has made this process infinitely easier than if telephone or surface mail messages were required. The mail system, of course, has also enabled us to maintain professional contacts established at conferences and other meetings, and to share and discuss ideas with these contacts.

C. Progress Summary

The prose comprehension project has completed an initial version of a model of prose comprehension (Miller & Kintsch, 1980a). This model has been applied to a large number of texts, and has yielded quite reasonable predictions of recall and readability. Psychologists from other universities have used this system to derive reading time and recall predictions for their own experimental materials; publication of this work is pending. We are currently using the AGE and UNITS packages to extend this model toward one that can make use of world knowledge in its analyses.

The planning group has completed the detailed analysis of several long thinking-out-loud protocols collected from both expert and novice software designers. These analyses involved the development of a lower level model for each of the protocols. See Atwood and Jeffries (1980) for details and examples.

D. List of Relevant Publications

Atwood, M. E., & Jeffries, R. Studies in plan construction I: Analysis of an extended protocol. Technical Report SAI-80-028-DEN, Science Applications, Incorporated, Denver, Co. March, 1980.

Atwood, M. E., & Jeffries, R. Studies in plan construction II: Novice design behavior. Technical Report SAI-80-154-DEN, Science Applications, Incorporated, Denver, Colorado, December, 1980.

Polson, P. G., Jeffries, R., Turner, A., & Atwood, M. E. The process of designing software. To appear in J. R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, N.J.: Erlbaum.

Atwood, M. E., Polson, P. G., Jeffries, R., and Ramsey, H. R. Planning as a process of synthesis. Technical Report SAI-78-144-DEN, Science Applications, Incorporated, Denver, Co. December, 1978.

Kintsch, W. On modelling comprehension. Invited address at the American Educational Research Association convention. San Francisco, April 10, 1979.

Kintsch, W. and van Dijk, T. A. Toward a model of text comprehension and production. Psychological Review, 1978, 85, 363-394.

Miller, J. R., & Kintsch, W. Readability and recall of short prose passages: A theoretical analysis. Journal of Experimental Psychology: Human Learning and Memory, 1980, in press.

Miller, J. R., & Kintsch, W. Readability and recall of short prose passages. Text, 1981, in press.

Miller, J. R. A knowledge-based mode of prose comprehension: Applications to expository text. Paper presented at the American Educational Research Association meeting, April, 1981.

E. Funding Support Status

1. Readability and Comprehension.
   Walter Kintsch, Professor, University of Colorado
   National Institute of Education
   NIE-G-78-0172
   9/1/78 - 8/31/81: $96,627
   9/1/80 - 8/31/81: $46,537

2. Text Comprehension and Memory
   Walter Kintsch, Professor, University of Colorado
   National Institute of Mental Health
   5 Rol MH15872-9-13
   6/1/76 - 5/31/81: $159,060
   6/1/80 - 5/31/81: $32,880

3. Comprehension and Analysis of Information in Text
Walter Kintsch, Professor, University of Colorado, and
Lyle E. Bourne, Jr., Professor, University of Colorado
Office of Naval Research, Personnel and Training Programs
ONR N00014-78-C-0433
6/1/78 - 5/31/80: $68,315
6/1/80 - 5/31/81: $60,000

4. Procedural Net Theories of Human Planning and Problem Solving
Michael Atwood, Research Psychologist, Science Applications, Incorporated, Denver, Colorado
Office of Naval Research, Personnel and Training Programs
ONR N0014-78-C-0165
1/25/78 - 12/31/80: $230,000
1/1/80 - 6/30/81: $85,000

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Sharing and Interactions with Other SUMEX-AIM Projects

Our primary interaction with the SUMEX community has been the work of the prose comprehension group with the AGE and UNITS projects at SUMEX. Feigenbaum and Nii have visited Colorado, and one of us (Miller) recently attended the AGE workshop at SUMEX. Both of these meetings have been very valuable in increasing our understanding of how our problems might best be solved by the various systems available at SUMEX. We also hope that our experiments with the AGE and UNITS packages have been helpful to the development of those projects.

We should also mention theoretical and experimental insights that we have received from Alan Lesgold and other members of the SUMEX SCP project. The initial comprehension model (Miller & Kintsch, 1980) has been used by Dr. Lesgold and other researchers at the University of Pittsburgh, as well as researchers at Carnegie-Mellon University, the University of Manitoba, Rockefeller University, and the University of Victoria.

B. Critique of Resource Management

The SUMEX-AIM resource is clearly suitable for the current and future needs of our project. We have found the staff of SUMEX to be cooperative and effective in dealing with special requirements and in responding to our questions. The facilities for communication on the ARPANET have also facilitated collaborative work with investigators throughout the country.

III. RESEARCH PLANS (8/79 - 7/81)

A. Long Range Projects Goals and Plans

The primary long-term goal of the prose comprehension group is the development of a blackboard-based model of prose comprehension. Correspondingly, we anticipate continued use of the AGE and UNITS packages.

These packages allow us to model the knowledge structures possessed by people and the inferential processes that operate upon those structures, and are essential to our work.

The primary goal of the planning project is the development of a model, or a series of models, of human performance on the software design task. We intend to begin by modeling the protocols of experts on a particular task, eventually extending the model to other levels of experience and other tasks. To do this we will have to become more familiar with AGE and work on articulating our theory in a way that is compatible with the AGE framework. This will involve two parallel lines of effort. One is a deeper analysis of our protocol data, to increase our knowledge of the detailed planning processes and knowledge structures experts are using to solve these problems. The second is the development of a model in AGE that can simulate these processes. We have to date been using SUMEX only for the latter activity, but we are beginning discover that both objectives are so intertwined that it is counter-productive for us to be using separate computer systems. We have transferred much of our protocol analyses activities to SUMEX, making it easier for us to share this very rich data source with other investigators.

B. Justification and Requirements for Continued SUMEX Use

The research of the prose comprehension project is clearly tied to continued access to the AGE and UNITS packages, which are simply not available elsewhere. We hope that our continued use of these systems will be offset by the input we have been and will continue to provide to those projects: our relationship has been symbiotic, and we look forward to its continuation.

C. Needs and Plans for Other Computational Resources

We currently use two other computing systems located at the University of Colorado. One is the Department of Psychology's VAX 11/780, which is used primarily to run real-time experiments to be modeled on SUMEX. The second is the University of Colorado's CDC 6400, which is used for various types of statistical analysis.

When the ARPA-sponsored Vax/Interlisp project is completed, we would be most interested in experimenting with becoming a remote AGE/UNITS site. It would seem that this sort of development is the ultimate goal of the package projects, and this type of interaction, once it becomes feasible, would be a logical extension of our association with the SUMEX facility.

D. Recommendations for Future Community and Resource Development

Our primary recommendation for future development within SUMEX involves (a) the continued support of INTERLISP, which is needed for AGE and for other work we have underway on SUMEX and (b) the continued development of the AGE and UNITS projects. In particular, we would like to see an extension of AGE to include a wider variety of control structures so that our psychological models would not be confined to one particular view of knowledge-based processing. The limited physical capacity of SUMEX, both in terms of address space and overloading, is, as before, a major problem. The prose comprehension group can no longer use the publicly released AGE/UNITS system due to its severely limited address space, and has had to build a personal AGE system from a stripped-down version of Interlisp and a selected subset of AGE and UNITS. We heartily endorse the plans underway to obtain more computing capacity for the SUMEX project.

Given our acquisition of a VAX, we particularly support the ongoing and continued development of INTERLISP for the VAX, so that local use of AGE and UNITS would be possible. Since we, as well as other psychologists, need the real-time capability of VAX/VMS to run on-line experiments, we hope that the INTERLISP system to be developed will be compatible with VMS. Note that this need for real-time work coincides with real-world applications of SUMEX programs, in which a VAX might be devoted to both real-time patient monitoring and diagnostic systems such as PUFF or MYCIN.

II.A.2.4 PUFF-VM Project

PUFF-VM: Biomedical Knowledge Engineering in Clinical Medicine

John J. Osborn, M.D.
The Institutes of Medical Sciences (San Francisco)
Pacific Medical Center

and

Edward H. Shortliffe, M.D., Ph.D.
Department of General Internal Medicine
Stanford University Medical Center
Stanford University

The immediate goal of this project is the development of knowledge-based programs to interpret physiological measurements made in clinical medicine. The interpretations are intended to be used to aid in diagnostic decision making and in therapeutic actions. The programs will operate within medical domains which have well developed measurement technologies and reasonably well understood procedures for interpretation of measured results. The programs are:

(1) PUFF: the interpretation of standard pulmonary function laboratory data which include measured flows, lung volumes, pulmonary diffusion capacity and pulmonary mechanics, and

(2) VM: management of respiratory insufficiency in the intensive care unit.

The second, but equally important, goal of this project is the dissemination of Artificial Intelligence techniques and methodologies to medical communities that are involved in computer aided medical diagnosis and interpretation of patient data.

I. SUMMARY OF RESEARCH PROGRAM

PUFF:

A. Technical Goals

The task of PUFF program is to interpret standard measures of pulmonary function. It is intended that PUFF produce a report for the patient record, explaining the clinical significance of measured test results. PUFF also must provide a diagnosis of the presence and severity of pulmonary disease in terms of measured data, referral diagnosis, and patient characteristics. The program must operate effectively over a wide range of pathological conditions with a broad clinical perspective about the possible complexity of the pathology.

B. Medical Relevance and Collaboration

Interpretation of standard pulmonary function tests involves attempting to identify the presence of obstructive airways disease (OAD: indicated by reduced flow rates during forced exhalation), restrictive lung disease (RLD: indicated by reduced lung volumes), and alveolar-capillary diffusion defect (DD: indicated by reduced diffusivity of inhaled CO into the blood). Obstruction and restriction may exist concurrently, and the presence of one mediates the severity of the other. Obstruction of several types can exist. In the laboratory at the Pacific Medical Center (PMC), about 50 parameters are calculated from measurement of lung volumes, flow rates, and diffusion capacity. In addition to these measurements, the physician may also consider patient history and referral diagnosis in interpreting the test results and diagnosing the presence and severity of pulmonary disease.

Currently PUFF contains a set of about 250 physiologically based interpretation "rules". Each rule is of the form "IF <condition> THEN <conclusion>". Each rule relates physiological measurements or states to a conclusion about the physiological significance of the measurement or state.

The interpretation system operates in a batch mode, accepting input data and printing a report for each patient. The report includes: (1) Interpretation of the physiological meaning of the test results, the limitation on the interpretation because of bad or missing data; the response to bronchodilators if used; and the consistency of the findings and referral diagnosis. (2) clinical findings, including the applicability of the use of bronchodilators, the consistency of multiple indications for airway obstruction, the relation between test results, patient characteristics and referral diagnosis. (3) Interpretation Summary, which consists of the diagnosis of presence and severity of abnormality of pulmonary function.

C. Progress Summary

<u>Knowledge base</u>:

PUFF is implemented on the PDP-10 in a EMYCIN system which is designed to accept rules from new task domains. A typical rule is:

```
If (FVC>=80) and  (FEV1/FVC<predicted-5) then PEAK FLOW RATES ARE
REDUCED, SUGGESTING AIRWAY OBSTRUCTION OF DEGREE
        if (predicted-15<= FEV1/FVC <predicted-5) MILD
        if (predicted-25<=FEV1/FVC <predicted-15) MODERATE
        if (predicted-35<=FEV1/FVC <predicted-25) MODERATE TO SEVERE
        if (FEV1/FVC <predicted-35)               SEVERE
```

This rule compares the ratio of FEV1, the amount of air that can be forced out in the first second of exhalation with the total "forced vital capacity" (FVC) or total amount of lung volume that can be exhaled. The inability to force out a large percentage of air in the critical first second implies the presence of an obstruction in the airway.

Results:

The results of the PUFF system are reviewed in more detail in the 1978 SUMEX annual report and [Kunz 78]. A version of the PUFF system is now in routine daily use at Pacific Medical Center. Reports are reviewed by a physician pulmonary physiologist. Over 85 % of the reports are accepted by the physician without change; they are signed and entered into the patient record. Most of the remaining reports are edited on-line to modify a small point in the test interpretation.

Table 1 reviews a study of the agreement in severity of diagnoses made by two MD's and by PUFF rules. This study was made with a less complete rule base than what is currently available in the pulmonary lab. In 94% of 144 cases analyzed in a prospective study, the degree of severity (0=none; 1=mild; 2=moderate; 3=moderately-severe; 4=severe) of OAD diagnosed by the first MD was within a single degree of severity of OAD diagnosed by the second MD. In 96% of the 79 cases for which the first MD diagnosed OAD, the second MD diagnosed the severity of OAD within one level of the severity diagnosed by the first MD. Agreement within one degree of severity of the diagnoses by the first and second MD's was substantially lower in RLD and DD cases. These discrepancies occurred because the second MD consistently called RLD more severe than did the first MD, and he consistently did not diagnose diffusion defects when the first MD diagnosed DD of moderate or greater degree.

	Percent Agreement with 1st MD			
	All 144 cases		1st MD made Dx	
Diagnosis	Second M.D.	PUFF Rules	Second M.D.	PUFF Rules
Normal				
OAD	0.94	0.99	0.96	0.97
RLD	0.92	0.97	0.77	1.00
DD	0.87	0.87	0.60	0.80
Total	0.91	0.94	0.86	0.94

Table 1. Percent agreement within one degree of severity of diagnoses Approximately 1500 patients have been interpreted by the system, by two MD's and by the first MD and rules.

In addition to the use of PUFF as a working clinical tool, it has been very useful for evaluation of knowledge representation methods. The original PUFF knowledge base (around 60 rules) represents realistic medical knowledge but is small enough to use for experiments. The PUFF knowledge has been used in the AGE system, the CENTAUR system using a combination of rules and prototypes, and the WHEEZE system, a UNIT-based approach to knowledge-representation.

D. Relevant Publications

[1] "A Physiological Rule-Based System for Interpreting Pulmonary Function Test Results", J.C. Kunz, R.J. Fallat, D.H. McClung, B.A. Votteri, J.S. Aikins, H.P. Nii, L.M. Fagan, E.A. Feigenbaum, HPP 78-154, Stanford Heuristic Programming Project, 1978.

[2] "Prototypes: An Approach to Knowledge Representation for Hypothesis Formation", Aikins, J.S., HPP-79-10 (working paper), 1979. Also Int. Joint Conf. on Artif. Intell., Tokyo, Japan, August, 1979.

[3] "A Physiological Rule-Based System for Interpreting Pulmonary Function Test Results", J.C. Kunz, R.J. Fallat, D.H. McClung, B.A. Votteri, J.S. Aikins, H.P. Nii, L.M. Fagan, E.A. Feigenbaum, Proceedings of Computers in Critical Care and Pulmonary Medicine, IEEE Press, 1979.

[4] "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering", E.A. Feigenbaum, Proceedings of the IJCAI, (1977). (Also Stanford Computer Science Department Memo STAN-CS-77-612).

E. Funding Support

PUFF-VM is supported by NIH grant GM24669 for $164,000 from 1 September 1978 - 30 August 1981. Some indirect costs are included in this total. A renewal application is pending.

VM:

A. Technical Goals

The Ventilator Manager program (VM) interprets the clinical significance of time varying quantitative physiological data from patients in the ICU. This data is used to manage patients receiving ventilatory assistance. An extension of a physiological monitoring system, VM (1) provides a summary of the patient's physiological status appropriate for the clinician; (2) recognizes untoward events in the patient/machine system and provides suggestions for corrective action; (3) suggests adjustments to ventilatory therapy based on a long-term assessment of the patient status and therapeutic goals; (4) detects possible measurement errors; and, (5) maintains a set of patient-specific expectations and goals for future evaluation. The program produces interpretations of the physiological measurements over time, using a model of the therapeutic procedures in the ICU and clinical knowledge about the diagnostic implications of the data. These therapeutic guidelines are represented by a knowledge base of rules created by clinicians with extensive ICU experience. Physiological measurements are collected at the bedside every 2-10 minutes and processed by the PMC computer system. The PMC and SUMEX computers are currently linked by TYMNET and special data communication programs written by SUMEX staff member, William Yeager. the program's interpretation is sent back to terminals in the ICU for evaluation by research clinicians.

B. Medical Relevance and Collaboration

To assist in the interpretation process, VM must be able to recognize unusual or unexpected clinical events (including machine malfunction) in a manner specifically tailored to the patient in question. The interpretation task is viewed as an ongoing process in the ICU, so that the physiological measurements must be continually reevaluated producing a current clinical picture.

This picture can then be compared with previous summary of patient status to recognize changes in patient condition upon which therapy selection and modifications can be made. The program must also determine when the measurements are most likely to be sensitive to error or when external measurements would be of diagnostic significance.

VM offers a new approach toward more accurate recognition of alarm conditions by utilizing the history and situation of the patient in the analysis. This is in contrast to the use of static limits applied to measurements generated to fit the "typical patient" under normal conditions. Our program uses a model of interpretation process, including the types and levels of conclusions drawn manually from the measurements to provide a summary of patient condition and trends. The program generated conclusions are stated at levels more abstract than the raw data; for example, the presence of hemodynamic stability/instability rather than in terms of heart rate and mean arterial pressure. When the data is not reliable enough to make these conclusions, additional tests may be suggested. The recognition of important conclusion for which external verification is sought, will also elicit the suggestion for confirming tests from the program.

C. Progress Summary

VM is provided with the values of 30 physiological measurements on a 2- or 10-minute bases by an automatic monitoring system. The output is in the form of suggestions to clinicians and periodic summaries (see example case below).

<u>Example Case</u>:

The following case demonstrates typical output of the system. The data used in this example were obtained from a post-cardiac surgery patient from the ICU at Pacific Medical Center. The terms VOLUME, ASSIST, CONTROLLED MANDATORY VENTILATION (CMV), and T-PIECE refer to specific types of ventilatory assistance. The output format is: (a) ..time of day.., (b) generated comments for clinicians, starting with "**", and (c) commentary in {}.

```
 ..1350.. ..1351..
 ** SYSTEM ASSUMES PATIENT STARTING VOLUME VENTILATION.
                                        {monitoring started}
 ** HYPERVENTILATION                    {diagnostic conclusions
 ** TACHYCARDIA                          based on monitored data}
 ** PATIENT HYPERVENTILATING.           {suggested therapy based on
 ** SUGGEST REDUCING MINUTE VOLUME       diagnosis}
 ..1400..
 . . .
 ..1450..
** HYPERVENTILATION
** TACHYCARDIA
** PATIENT HYPERVENTILATING.
** SUGGEST REDUCING MINUTE VOLUME
 ..1500..
** HYPERVENTILATION
** PATIENT HYPERVENTILATING.
** SUGGEST REDUCING MINUTE VOLUME


Current conclusions:                {summary information}
HYPOTENSION PRESENT for 41 MINUTES
HYPERVENTILATION PRESENT for 33 MINUTES
SYSTOLIC B.P. LOW for 46 MINUTES
{etc.}

Conclusions:       {time of day}   |.....|.....|.....|.
                                   13    14    15    16
HEMODYNAMICS -- STABLE                        ====
HYPERVENTILATION -- PRESENT            =     == == ====
HYPOTENSION -- PRESENT                   ====  ========
TACHYCARDIA -- PRESENT                 ==== ==

patient is on ASSIST                         =====   ==
patient is on CMV                        ====     ===
patient is on VOLUME                   ==
patient is on NOT-MONITORED        ====

Goal is CMV                            ================
Goal is VOLUME                     ====
                                   |.....|.....|.....|.
                                   13    14    15    16
```

The availability of new measurements requires updated interpretations based on the changing values and trends. As the patient setting changes--e.g., as a patient starts to breathe on his own during removal (weaning) from the ventilator--the same measurement values lead to different interpretations. In order to properly interpret data collected during changing therapeutic contexts, the knowledge base includes a model of the stages that a patient follows from admission to the unit through the end of the critical monitoring phase. Recognition of the appropriate patient context is an essential step in determining the meaning of most physiological measurements.

The majority of the knowledge of the VM program is concerned with the relations between the various concepts known by the program. These concepts include: measurement values, typical therapeutic decisions, diagnostic labels, and physiological states. The connections between concepts are represented by a form of production rules using the structure "IF premise THEN action."

The rules in VM are of the form:

IF facts about measurements or previous conclusions are true

THEN

1) Make a conclusion based on these facts;
2) Print out suggestions for the clinician;
3) Establish expectations about the future values of measurements.

A sample VM rule is shown below.

```
STATUS RULE: STABLE-HEMODYNAMICS
DEFINITION: Defines stable hemodynamics for most settings
APPLIES to patients on VOLUME, CMV, ASSIST, T-PIECE
COMMENT: Look at mean arterial pressure for changes in
    blood pressure and systolic blood pressure for maximum
    pressures.
IF
  HEART RATE is ACCEPTABLE
  PULSE RATE does NOT CHANGE by 20 beats/minute in 15 minutes
  MEAN ARTERIAL PRESSURE is ACCEPTABLE
  MEAN ARTERIAL PRESSURE does NOT CHANGE by 15 torr in 15 minutes
  SYSTOLIC BLOOD PRESSURE is ACCEPTABLE
THEN
  The HEMODYNAMICS are STABLE
```

Figure 1. Sample VM Interpretation Rule. The meaning of 'ACCEPTABLE' varies with the clinical context--i.e., whether the patient is receiving VOLUME or CMV ventilation, etc. This rule makes a conclusion for internal system use. Similar rules also make suggestions to the user.

During this last year, we started clinical validation of the first medically oriented artificial intelligence program to be used in a continuous (real-time) setting. Our experiences with this project for data interpretation in the intensive care unit, has highlighted some basic research questions for time-oriented systems. We have concentrated on explanation, reasoning, and therapy planning in a changing clinical environment.

Our work in explanation has concentrated on showing the change in reasoning between one instant in time and the next. Expressed symbolically, a description of what has changed may often be more useful than a description of the current chain of reasoning. We are also working on methods to reason with data from different points in time, e.g., revising old hypotheses about the clinical setting when a delayed laboratory result becomes available.

Our third area of time-oriented research is the development of techniques for creating and modifying long-term therapeutic plans. These plans are created by combining management strategies provided by expert clinicians. Based on an analysis of previous response to therapy (e.g., following an oncology protocol), specialized treatment will be suggested for patients unable to follow the standard treatment plan. About twenty strategies, and fifteen appropriate patients have been identified for this task.

D. Relevant Publications

Fagan, L.M., Kunz, J.C., Feigenbaum, E.A. and Osborn, J.J.: A symbolic processing approach to measurement interpretation in the intensive care unit. Proc. Third Annual Symposium Computer Applications in Medical Care, Silver Spring, Maryland, October, 1979, pp. 30-33.

Fagan, L.M., Shortliffe, E.H. and Buchanan, B.G.: Computer-based medical decision making: From MYCIN to VM. Automedica 3(2), 1980.

Fagan, L.M.: VM: Representing Time-Dependent Relations in a Medical Setting, Ph.D. dissertation, Stanford University, 1980.

Osborn, J.J., Fagan, L.M., Fallat, R.J., et al: Managing the data from respiratory measurements. Med. Instrumentation, November-December, 1979. (Winner of the 'Best Article of the Year' Award for AAMI - 1979.)

E. Funding Support

PUFF-VM is supported by NIH grant GM24669 for $164,000 from 1 September 1978 - 30 August 1981. Some indirect costs are included in this total. A renewal application is pending.

II. <u>RESEARCH PLANS</u>

A. Long Range Goals and Plans

The main emphasis of this project has switched from the development of the PUFF system to the extension and evaluation of the VM system. This change is consistent with the goals of the NIH proposal, the current use of PUFF in a clinical setting and the research questions that remain in the VM portion of the project. Some long term interests, such as consensus building between experts, will be examined using both application areas.

The long range goal of the VM project is to develop and evaluate an interpretation system that will improve patient care in the ICU. Toward this goal, we plan to extend the rule set, provide better models of physiology and therapy, and start a formal evaluation of the program's therapeutic advice.

The rule set in VM will be extended to handle a greater number of patients. The current emphasis of the program has been on the management of post-surgical patients with normal pre-operative status. We will continue to concentrate on post-surgical patients, but the knowledge base will be augmented to handle patients with additional problems noted before surgery or those who have an unusual response to therapy after surgery. The majority of this knowledge will be used to create a more detailed classification of the patient population and the corresponding generation of expectations.

These rule set extensions will ultimately be limited by representation of the underlying cardiopulmonary physiology and the therapeutic plans used in the ICU. Still other improvements will come from a better model of the mechanical ventilator and other instrumentation. Each of these models will provide a structure upon which to build the rule base, and are motivated by the special problems of evaluating the patient's status in a dynamic clinical setting. These problems include the evaluation of the relationship between actual and anticipated response to therapy and the the recognition of a particular therapy step in the context of a larger therapeutic plan (e.g., the process of removing a patient from the ventilator when the patient has an underlying lung disease).

We will develop integrated causal models of the physiological response of the patient to changes in patient physiological state. In the clinical case, we assume that the model will consist of abstract characterizations of the active parts of the patient and life-support system. The relation between these active parts will characterize the physiology of the patient organ system or the operation of the mechanical life-support system. The general approach to development of models will be to consider a simple, well-understood but potentially useful case: a static model of ventilation; second, we will consider oxygenation as a more complex physiological process involving a larger number of physiological parameters. The oxygenation model will also include some representation of time dependency using both heuristic notions of time change and mathematical representation of time constant of change.

In order to determine the appropriate areas for these model building activities and to insure acceptance by physicians, a careful prospective validation will be carried out to identify the accuracy of the advice of the program.

III. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Collaborations and Medical Use of Programs via SUMEX

The PUFF-VM project requires very close collaboration between investigators at two institutions separated by fifty miles. This kind of collaboration, in which program development and testing proceeds concurrently on the same application system, requires a computer network facility for sharing of code, data and ideas. SUMEX has been used at PMC for running programs developed concurrently by Stanford and PMC staff, and data has been taken from the PMC computer system and transferred to SUMEX on magnetic tape for program development and testing. The SUMEX staff has developed a cooperating set of computer programs to allow the PMC computer and the SUMEX/2020 systems to actively exchange files and program data and output.

We also use the SUMEX system for purposes other than program development. A joint PMC-Stanford report of VM was prepared entirely through the the word communications and processing capabilities of SUMEX. Investigators from the two institutions have collaborated in writing reports together; the separate contributions are prepared on SUMEX, edited and merged with an exchange of messages but without ever requiring actual meetings. We have also used the system for trading bibliographic information with other AIM users. We have also experimentally run the Internist program using SUMEX.

B. Sharing and Interactions with Other SUMEX-AIM Projects

We have participated in the AIM workshop and had very fruitful interaction with a number of other SUMEX users, directly influencing our perception of important problems and potentially appropriate solutions. Personal contacts at other conferences, at Stanford AI weekly meetings, and at PMC with visiting members of the AIM community, have also been very helpful in keeping abreast of the current thinking of other members of the AI community and with members of the medical community interested in computer based physiological analysis and diagnosis. We believe that the use of a common machine and the existence of the AIM conference encourages increased recognition and better communication with other AIM workers. Within AIM we most closely collaborate with the MYCIN, MOLGEN and DENDRAL projects, who share common space, common techniques, and common attitudes.

C. Critique of Resource Management

The SUMEX community continues to be an extremely supportive environment in which to do research on uses of artificial intelligence in clinical medicine. The community has two equally vital resources -- the people with knowledge and interest in AI and the facility on which AI system development can proceed. They are equally excellent as resources, helping hands when faced with problems, and friendly support for continued productive research. The availability of INTERLISP; of a facility on which routine data processing functions (eg. manipulating magnetic tapes and making long listings) can take place; and of message-sending among remote users are all vital functions for our project. SUMEX provides them in an environment which is friendly and reliable. Management of the SUMEX facility is consistent and excellent.

D. Needs and Plans for Other Computational Resources

The future goals of the project (as described above) will require considerable computational requirements in the near future. These requirements will come in the form of active development of a large INTERLISP program, and extensive testing of the program in a clinical environment. We hope to perform as much of the evaluation work as possible on the 2020. System development of the program will probably continue on

SUMEX during off-hours or be off-loaded to the spare time on the 2020. Most subsidiary text processing tasks have been off-loaded from SUMEX to avoid the high load average situation during the day. The storage of usable versions of the program and the test files used in the evaluation of the program will require about 1000 additional pages on the SUMEX computer.

The on-line evaluation of the program is currently using the 2020 on a stand alone basis one afternoon a week. Batch runs after modifications of the rule set also use computer time during off hours. This level of computer usage should remain during 1981-82.

E. Recommendations for Future Community and Resource Development

We perceive the evolution of our.AI capability as moving from a highly speculative development state, for which the interactive development capabilities of SUMEX are vital, to a more stable but still changing validation-and-evaluation state. Ultimately we foresee rather stable specification of a program for routine clinical use. Thus, we see the need to transfer our AI techniques from the SUMEX PDP-10 to a local host. For this transfer, a principal long-range need is for software systems that will allow us to run AI systems on a mini-computer after they have been developed on the more powerful SUMEX facility. If the validation of PUFF-VM in the PMC clinical setting shows the programs to be effective in health care, then we hope and expect to be able to provide the capability on a routine basis.

We would also like to encourage SUMEX's role as a facilitator of information transfer between AIM users. This can happen by scheduling on-line demonstrations that any other user can "connect to," or by providing a common depository for AI and medicine information. This might take the form of on-line bibliographies, collecting common user packages, or connecting common research interests together. This communication service would complement the technical service facilities currently provided by the SUMEX staff.

II.A.2.5 Rutgers Computers in Biomedicine Project [Rutgers-AIM]

Rutgers Computers in Biomedicine

Rutgers Research Resource--Computers in Biomedicine

Principal Investigator: Saul Amarel
Rutgers University
New Brunswick, New Jersey

I. SUMMARY OF RESEARCH PROGRAM

A. Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for significant research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The focal concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is 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 problem solving and in processes of experimentation and theory formation.

At present, the total number of investigators who participate in scientific activities of the Resource is 81; of these, 34 have Rutgers appointments, 26 are outside investigators who participate in collaborative research projects that are mainly located at Rutgers, and 21 are investigators from collaborative national AIM projects that are located in different parts of the country. In addition, the Resource has 11 other members in Administrative, Computer Systems/Operations and general programming and secretarial functions. Thus, the Rutgers Resource community numbers at present a total of 92 participants.

Resource activities include research projects (collaborative research and core research) training/dissemination projects, and computing services in support of user projects. The research projects are organized in three main AREAS OF STUDY. These areas of study and the senior investigators in each of these are:

1) Medical Modeling and Decision Making (C. Kulikowski, S. Weiss)

2) Modeling Belief Systems and Commonsense Reasoning (C. Schmidt and N.S. Sridharan)

3) Artificial Intelligence: Representations, Reasoning, and System Development (S. Amarel)

The training/dissemination activities of the Rutgers Resource include sponsorship of the Annual AIM Workshop - whose main objective is to strengthen interactions between AIM investigators, to disseminate research methodologies and results, and to stimulate collaborations and imaginative resource sharing within the framework of AIM. Starting in 1979, the Workshop is being organized and hosted on a rotational basis by the members of the AIM community, in coordination with the Rutgers Research Resource. The Sixth AIM Workshop was organized by the SUMEX-AIM Resource and it was held at Stanford University in August, 1980.

B. Medical Relevance; Collaborations

During 1980-81 we continued the development of a versatile system for building consultation models, called EXPERT. We reached a new milestone in technology transfer by implementing a scheme for the automatic translation of EXPERT models (with some restrictions) into microprocessor compatible algorithmic form. 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. They also represent an active area of biomedical and clinical research, in which a group of our medical collaborators at the University of Missouri under Dr. Gordon Sharp has been noted for its contributions. The application of A.I. approaches to problems of medical decision making in this domain was facilitated by our collaboration with Dr. Donald Lindberg, Director of the Health Care Technology Center at the University of Missouri.

Our experience with the design of the rheumatology model has shown us that the knowledge engineering tools and know-how that we developed so far in the Resource 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 every 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.

During the current research period the initial prototype model in connective tissue disease diagnosis developed by Dr. Sharp and his colleagues was subjected to careful review and evaluation by an outside panel of experts. As a result, the model was significantly broadened in scope (of both findings and diagnoses), and a new criterion-based scheme for consistently organizing knowledge in the specialty was devised. The Resource investigators at Rutgers formalized this approach within the framework of EXPERT, and developed a method for updating criterion-based models using the cases stored in their data base. Another important development motivated by the rheumatology project has been the introduction of a special logical feature for reporting the diagnostic interpretations in a sub-language that more closely follows clinical terminology.

In ophthalmology work continues with the development of the neuro-ophthalmological model by Dr. William Hart at Washington University, the glaucoma model by Dr. Y. Kitazawa at Tokyo University and the 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 (Columbia University and Overlook Hospital) in the development of a model for the interpretation of serum protein electrophoretic patterns. The prototype model was tested by specialists in the domain and improved to the point that it could be used to produce an automatic interpretation for the output of a scanning densitometer. In order to facilitate the transfer of the EXPERT advice scheme to the instrument environment, we developed an automatic translating program which produces an algorithmic form of the EXPERT model (which is restricted slightly in its scope of primitives for this purpose). The algorithmic version was in turn automatically translated into a microprocessor assembler code. Thus we have completed a full cycle of transfer of technology from the large scale developmental system of the Resource to a clinically useable instrument.

In biomedical modeling applications, the EXPERT scheme has been used to build preliminary prototype models that will interact with simulation programs, giving advice on the modeling process itself. In collaboration with Dr. William Yamamoto of George Washington University, we are working on a model of respiration, while enzyme kinetics modeling continues to be investigated with Dr. David Garfinkel of the University of Pennsylvania.

C. Highlights of Research Progress

1. Medical Modeling and Decision-Making

Research during the past year has concentrated on problems of representation, knowledge acquisition and treatment planning strategies, resulting in significant extensions of the EXPERT formalism, and its application in specific medical domains (see I.B above). From the developmental perspective, we have achieved a great step forward by producing automatic translation programs for EXPERT models into microprocessors compatible form.

The major achievements of our work can be summarized as follows:

1.1 Technology Transfer

A program for automatically translating an EXPERT model into algorithmic form was first developed. This required a restriction on the logical primitives used in the EXPERT model, which did not prove to be a significant limitation.

A program for converting the algorithmic version of the EXPERT model into a microprocessor assembly language (for the Motorola 6000) was then developed.

The translation programs were tested with prototypes and an actual expert-derived model for serum protein electrophoresis interpretation. The utility and dissemination of expert models will be greatly enhanced once they can be used with clinically useful and routinely applied instruments.

1.2 EXPERT System Changes

a) Procedure-Invocation Capability/Distributed Control

We have added a command within EXPERT that permits suspension of execution of the reasoning strategies while (user defined) procedure is invoked. This can accept and pass back parameters from the model, so that specialized procedural knowledge (in the form of mathematical simulation models, search procedures, pathway tracing procedures, etc.) in a domain can be explicitly coded and its results used by the consultation model.

This capability also allows one EXPERT model to call another, thereby simulating the passage of information and control from one specialist to another when dealing with a complex medical case. This provides an initial focus for studies in <u>distributed</u> <u>knowledge</u> <u>base</u> design and reasoning in consultation.

b) Criterion Based Representation

A particular representation formalism has been developed to encode clinical criteria in a natural manner within the EXPERT scheme. This involves the specification of an intermediate hypothesis structure for possible, probable and definite truth-value modifiers of a disease, in the form of major, minor and exclusionary findings. It is implemented through a special knowledge acquisition front-end which then translates the criteria from tabular form to EXPERT rules.

c) Output Language Capabilities

These have been implemented as special logical functions that can be used to override the conventional EXPERT output with special clinical language templates. The templates are filled with the specific interpretations and supporting findings for each individual case, thus making for more specific (or personalized) reporting of conclusions and advice.

1.3 Knowledge Acquisition and Model Building Strategies

We have developed several heuristic methods for giving advice on the updating of a knowledge base specified in the Criterion-based form, as well as a front end for helping build such a model. The learning capabilities are interactive so that a meta-EXPERT system suggest rules that could be generalized or restricted depending on the performance of the current model on a data base of cases. The user can select a particular modification of the model based on these suggestions (or he can try any other of his choice), and the meta system will test and comment on the performance of the modified model. If the results are not as satisfactory as expected, further changes will be suggested and the process repeated until performance is improved.

1.4 Treatment Planning: Representation and Strategies

An investigation of treatment planning representations has resulted in a new system that emphasizes facilities for encoding sequences of actions (plans of management) and producing explanations of reasoning based on them. Strategies of treatment selection involve a generalized scheme for positive (therapeutic) effects and negative (contraindication, allergy, complication) effects, that are used as the basis of sorting procedures. A system has been implemented and is being tested in the area of infectious eye disease management.

1.5 Other Research and Development

a) Data Base Compatibility of EXPERT

A method for transferring case data from an existing data base into EXPERT form was developed and implemented for our thyroid data base.

b) Model-Based Knowledge Acquisition

A scheme for prompting the acquisition of EXPERT-type rules by a model of domain function has been investigated.

2. Modeling of Belief Systems and Commonsense Reasoning

2.1 Recognition of Commonsense Plans: Plan Generation and Revision

The human information processing system possesses a richly structured long-term memory which contains an enormous store of general knowledge, concepts and facts about the domains of commonsense discourse. During this period our research has focused on ways in which this type of information can be represented and used in plan generation and revision.

In plan generation we have developed a rule-based representation for normative and normalcy knowledge. A plan generator has been implemented which uses this knowledge as a basis for default reasoning so that planning may proceed despite the unavailability of knowledge about the specific planning context.

In both planning and plan revision we have explored the ways in which default assumptions and the hierarchy of commonsense concepts may be used. By using the hierarchy of concepts, we have been able to develop strategies for dealing with uncertain and incomplete knowledge which at the same time provides the basis for guiding future refinement and revision of a plan hypothesis.

Further, we have reimplemented our plan generator so that (a) it is organized as a pair of cooperative processes that propose plan alternatives and criticize the proposals, (b) both processes can be interactively guided, and (c) user may call for replanning by functioning as a critic.

2.2 Reasoning with Commonsense Concepts

In the area of goal-directed inference about individuals we have studied ways of representing both the intension and extension associated with hierarchically structured commonsense concepts. Using this representation we have investigated strategies of plausible inference and the attendant problem of evaluating the consistency of the resulting inferences and proposed strategies for dealing with incomplete and/or inconsistent inferences.

In a recent report, we have explored a formal construct called PERMISSIVENESS of a default rule as an aid to resolving inconsistencies among default rules. Permissiveness requires no prior probabilities or additional domain knowledge.

3. Artificial Intelligence: Representations and System Development

3.1 AIMDS: System for Studying Knowledge Representations and Commonsense Reasoning

The main objective of this research during the current period were: to continue experimental work on the present major applications of AIMDS; to enhance the description method, so that it becomes possible to construct complex descriptions, including the ability to describe system behavior to itself; to develop new processing frameworks for extended power and usability of the system; and to improve the AIMDS implementation for efficiency, reliability and quality of user interface. Substantial progress was made in each of these areas. The two main applications of AIMDS remain in the area of common-sense reasoning studied in the Resource, and in studies of legal reasoning that are funded by the NSF. The AIMDS system provides a solid foundation for our collaborative research on commonsense reasoning; it continued to be used this period in the representation of psychological concepts of plan, expectation, interpretation, as well as of common sense concepts of objects, places, persons, goals and of world knowledge governing the use of such concepts. In the past, the two AIMDS functions MAKE and FIND accepted partial descriptions of individuals in a domain. There were no functions that constructed descriptions. One of the significant new capabilities that was introduced this period is functions for building and manipulating descriptions within AIMDS. Also, improvements in AIMDS implementation which were developed this period include powerful methods for optimizing memory use. The development of interesting and convenient user interfaces has occupied a subordinate role in our project. But the recent introduction of a general purpose display front-end for AIMDS has begun to fill a gap. The display, called RELDISP, displays a fixed 7x7 matrix on the top half of the screen; using the bottom half as the traditional scroll mode user interaction panel. The rows and columns of the matrix display the seven most recently referenced instances in the associative memory, and the entries of the matrix display the name of the relation asserted or denied between any pair of instances. The display is quite effective in focusing the viewers' attention during a demonstration of any of the AIMDS application programs. We have yet to gather experience in its use for the programmer who develops application programs. The RELDISP package was also used to conduct a demo of the plan recognition project during the Sixth Annual AIM Workshop held at Stanford. Presently, a revised user manual is in preparation which has expanded sections that describe important new additions to the AIMDS system. The new manual and the improved user interface facilities are expected to contribute to further dissemination of the AIMDS system in the Resource and in the national AIM community.

3.2 Expertise Acquisition and Theory Formation

There are two main research activities in this subproject. One, which was initiated in the previous period, is focusing on the task of <u>acquiring good control strategies</u> for heuristic search via experimentation -- within the framework of a given problem representation. The second, which we have been pursuing for several years, is concerned with improvements in a problem solving system via <u>shifts in problem representation</u>. Related to the second approach, is the problem of how to acquire and use models at different levels of resolution in a given domain in order to increase problem solving performance (expert performance) in the domain. Problems of theory formation are at the heart of the processes of expertise improvement and acquisition that we have been studying here. During this period, we made substantial progress in these studies.

3.3 Knowledge Acquisition and Man-Machine Interfaces

In this subproject we have two research activities. The first, in which most of our effort was concentrated in the current period, is concerned with Natural Language Interfaces for Consultation Systems. This work has resulted this year in a doctoral dissertation by Vic Ciesielski, a research assistant in the Resource. The second activity is focusing on a problem of language acquisition in the context of a student/teacher dialogue.

3.4 Programming Environment for AI Research

Our work on the development of a supportive programming environment for AI research has continued this period with emphasis on the development of an enhanced new version of the Rutgers/UCI LISP.

D. Up-to-Date List of Publications

This following is an update of publications in the Rutgers Resource for the period 1980 and 1981 (only publications not listed in previous SUMEX annual reports are presented here).

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

Amarel, S., (1981) "Review of Characteristics of Current Expert Systems" Rutgers Technical Report CBM-TM-89.

Amarel, S. (1981) "Problems of Representation in Heuristic Problem Solving; Related Issues in the Development of Expert Systems" Rutgers Technical Report CBM-TR-118.

Banerji, R.B. and Mitchell, T., (1980) "Description Languages and Learning Algorithms: A Paradigm for Comparison" Rutgers Technical Report CBM-TR-107.

Ciesielski, V. (1980) "A Methodology for the Construction of Natural Language Front Ends for Medical Consultation Systems" Rutgers Technical Report CBM-TR-112.

Galen, R., Weiss, S., Kulikowski, C.A., (1981) "The Interpretation of Serum Protein Electrophoresis by Computer" Medical Laboratory Observer, (in press).

Goldberg, R.N. and Weiss, S.M., (1981) "An Experimental Transformation of a Large Expert Knowledge-Base" Proceedings of the Fourteenth Hawaii International Conference on Systems Sciences, pp. 828-836, (1981). (to appear in Journal of Medical Systems.) Rutgers Technical Report CBM-TR-110.

Goldberg, R. and Weiss, S.M. (1980) "Constructing an Expert Knowledge Base for Thyroid Consultation Using Generalized A.I. Techniques" Proc. Hawaii International Conference on Systems Science, Honolulu, Jan. 1980.

Goodson, J.L. (1980) "A Process for Evaluating Tree-Consistency" Rutgers Technical Report CBM-TR-111.

Kastner, J. and Weiss, S.M., (1981) "A Precedence Scheme for Selection and Explanation of Therapies" Rutgers Technical Report CBM-TM-90.

Kulikowski, C.A., (1980) "Artificial Intelligence Methods and Systems for Medical Consultation" IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-2, No. 5, Sept. 1980, pp. 464-476.

Kulikowski, C.A., Weiss, S.M. and Galen, R.S., (1980) "On the Diagnostic Frontier: Expert Consultation by Computer" Diagnostic Medicine, Nov/Dec. 1980, pp. 89-90.

Kulikowski, C.A. and Ostroff, J.H. (1980) "Constructing an Expert Knowledge Base for Thyroid Consultation Using Generalized A.I. Techniques" Proc. Annual Symposium on Computers in Health Care, Washington, Nov. 1980.

Lindberg, D. and Weiss, S.M. (1980) "On the Diagnostic Frontier: Expert Consultation by Computer" MEDINFO '80.

Lindberg, D., Sharp, G., Kingsland, L., Weiss, S., Hayes, S., Ueno, H., Hazelwood, S. (1980) Proceedings of the Third World Conference on Medical Informatics, pp. 1311-1315 (1980).

Mitchell, T., Utgoff, P. and Banerji, R. (1980) "Learning Problem-Solving Heuristics by Experimentation" Rutgers Technical Report CBM-TR-114.

Mitchell, T.M. (1980) "The Need for Biases in Learning Generalizations" Rutgers Technical Report CBM-TR-117.

Nagel, D., (1980) "An Experiment in Extracting Some Properties of Binary Relations" Rutgers Technical Report CBM-TM-84.

Nagel, D., (1980) "Some Considerations on Extracting Definitional Information about Relations" Rutgers Technical Report CBM-TM-85.

Nudel, B. and Utgoff, P. (1981) "A Bibliography on Machine Learning" Rutgers Technical Report CBM-TR-120.

Politakis, P. and Weiss, S.M., (1981) "A System for Empirical Experimentation with Expert Knowledge" Rutgers Technical Report CBM-TM-91.

Politakis, P. and Weiss, S. (1980) "Designing Consistent Knowledge Bases: An Approach to Expert Knowledge Acquisition" Rutgers Technical Report CBM-TR-113.

Politakis, P., Weiss, S. and Kulikowski, C.A. (1980) "Designing Consistent Knowledge Bases for Expert Consultation Systems," 13th Annual Hawaii International Conference on Systems Science, January, 1980, pp. 675-683.

Politakis, P., Weiss, S., (1980) "Toward a Systematic Methodology for the Design of Expert Consultation Systems", Proceedings ACM Computer Science Conference, p. 22, (1980).

Schmidt, C.F., (1980) "The Role of Object Knowledge in Human Planning" Rutgers Technical Report CBM-TM-87.

Schmidt, C.F. (1980) "Plan Recognition and Revision: Understanding the Observed Actions of Another Actor" Rutgers Technical Report CBM-TR-115.

Smith, D. (1980) "FOCUSER: A Strategic Interaction Paradigm for Language Acquisition" Rutgers Technical Report CBM-TR-116.

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

Sridharan, N.S., Schmidt, C.F., and Goodson, J. L. (1980) "The Role of World Knowledge in Planning" Rutgers Technical Report CBM-TR-109.

Sridharan, N.S., Schmidt, C.F., and Goodson, J.L. (1981) "Reasoning by Default" Rutgers Technical Report CBM-TR-119.

Weiss, S.M. and Kulikowski, C.A. (1980) "On the Diagnostic Frontier: Expert Consultation by Computer" MEDINFO '80.

Weiss, S.M., Kern, K., Kulikowski, C.A. and Uschold, M.F. (1981) "A Guide to the Use of the EXPERT Consultation System. revised 1981.

Weiss, S., Kulikowski, C.A., (1980) "Data Base Aids for Designing Consultation Systems", Proceedings of the Third World Conference on Medical Informatics, pp. 794-797 (1980).

Weiss, S., Kulikowski, C.A., (1981) "Expert Consultation Systems: The EXPERT and CASNET Projects, Pergamon INFOTECH State of the Art Report on Machine Intelligence, (in press).

Weiss, S., (1980) "Tools for Knowledge Engineering", Proceedings Workshop on Expert Systems, pp. 59-64 (1980).

Weiss, S. (1981) Article on The EXPERT System, to appear in the Knowledge Engineering Chapter in "Expert Systems," editors: Hayes-Roth, F., Waterman, D., (in press).

Weiss, S., Kulikowski, C.A., and Galen, R. (1981) "Developing Microprocessor Based Expert Models for Instrument Interpretation" Rutgers University Department of Computer Science Tech. Memo, (1981).

E. Funding Support

The Rutgers Resource is funded through an NIH grant entitled "Rutgers Research Resource on Computers in Biomedicine" -- grant number P41RR643. The Principal Investigator is Dr. Saul Amarel, Professor, Chairman Department of Computer Science, and Director of the Laboratory for Computer Science Research, Rutgers, the State University of New Jersey. This grant is in the first year of its fourth 3-year renewal extending from Dec. 1, 1980 through Nov. 30, 1983. The total direct cost awarded for the 3-year period is $1,395,580 with first year funding of $495,079 in direct costs from Dec. 1, 1980 through Nov. 30, 1981.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Dissemination; Interactions

The SUMEX-AIM facility provides one of the nodes where some of our medical collaborators can access programs developed at Rutgers.

An important responsibility of the Rutgers Resource within the national AIM community is to sponsor dissemination and training activities.

The focus of our efforts in this area continues to be centered around the AIM Workshops and sessions on AIM research at national and international conferences.

In order to increase the dissemination of AIM work within specialty fields of medicine, we have also presented tutorial papers at relevant conferences.

1) Sixth AIM Workshop (1980):

This past year the AIM Workshop was hosted by Stanford University, according to the new scheme of rotation. Dr. E. H. Shortliffe organized the Workshop, coordinating with Dr. C. A. Kulikowski for the Rutgers Resource, which continues as the main sponsor: the budget was made available from the Rutgers Resource to cover most local arrangements at Stanford and travel funds for the participants.

2) AAAI Conference:

Resource investigators presented papers and were panelists at the first AAAI Conference that followed the workshop. This gave an opportunity for wider dissemination of the AIM research in the artificial intelligence community.

3) Expert Systems Workshop:

A strong representation of Resource investigators at this workshop (which was organized by Rand) enabled the contribution of ideas and experience from the biomedical applications of the Resource to be documented in the upcoming book that will cover this new general area of research.

4) MEDINFO:

The medical modelling and decision making group of the Resource presented a paper on the EXPERT system and its updating capabilities at this important international conference. The conference was co-chaired by one of the groups close medical collaborators, Dr. Donald Lindberg, who also presented the joint work on rheumatology model development using the Rutgers EXPERT system.

5) Hawaii International Conference on Systems Sciences:

Investigators from the Resource's medical group received the best paper award in the medical decision session at this Conference for the contribution entitled: "An Experimental Transformation of a Large Expert Knowledge Base".

6) Fourth Annual Symposium on Computers in Health Care:

Research on large data base testing via consultation models applied in the area of thyroid diseases was presented by members of the medical group of the Resource.

B. National AIM Projects at Rutgers

The national AIM projects, approved by the AIM Executive Committee for the 1980 grant year, who use the Resource system are listed here. BRIGHT is continuing under the direction of Dr. W. Gordon Walker of Johns Hopkins University, mostly for research in clinical medicine. The CONGEN and DENDRAL projects, under Dr. Djerassi and Dr. D. Smith, are predominantly at SUMEX with the Resource used when the system is overloaded, as a backup for demonstrations, and to allow for ease in collaborations with Dr. San Felippo of Rutgers Chemistry Department. The INTERNIST project has been using the Resource for backup during demos, etc, all text-editing functions and when SUMEX is down for extended periods. Most of the actual research on the Ohio State project is continuing to be done at Rutgers under Dr. Chandrasekaran and Artificial Intelligence models of clinical reasoning are being developed by Dr. Greenes of Harvard University on the Rutgers system. Dr. Garfinkel, heading the Penn project, has continued his work on modeling of metabolic pathways on the Rutgers system, in collaboration with Resource investigators. During this period, the AIM Executive Committee has authorized a group of about 25 researchers in biomedical modeling and simulation (from various parts of the country) to access the Rutgers/AIM system on a pilot basis -- to help them formulate collaborative research programs.

C. Critique of Sumex-AIM Resource Management

Rutgers is currently using the Sumex facility for three purposes:

- communication with other researchers working at Sumex, and with Sumex staff.

- building collaborative ophthalmology models. The researchers working on this project are from ophthalmology groups at a number of institutions around the country.

- backup computing for demonstrations, conferences and site visits.

In the first two cases Sumex is being used because it provides a convenient center for communication among people spread around the country. This is true because of its comparatively good Tymnet facilities, as well as because of the good support that its staff provide for remote users, particularly those who are not experienced computer users.

Our usage is currently running at less than 200 connect hours per year, with an overall connect/CPU ratio of about 42. This usage has gone down since last year because of the fact that Rutgers is now again on the ARPAnet. Much of our cooperation with users and staff at Sumex can be done using mail and file transfer through the ARPAnet. This does not require Rutgers users to log into Sumex. Such cooperation continues to be important for us, even though it is no longer reflected in our usage statistics.

The quality of the support given by Sumex staff continues to be very high. Our users report that the load at Sumex continues to be very high, despite attempts by staff to control it. However the problem has not been enough to cause serious difficulties for most of our activities.

III. RESEARCH PLANS

A. Project Goals and Plans

We are planning to continue along the main lines of research that we have established in the Resource to date. Our medical collaborations will continue with emphasis on development of consultation systems in rheumatology, ophthalmology and endocrinology. The psychological component of work on belief systems and common sense reasoning will be phased out, while the basic AI issues of plan recognition/generation and default reasoning will receive increased attention. Our core work will continue with emphasis on further development of the EXPERT framework and also on AI studies in representations and problems of knowledge and expertise acquisition. We also plan to continue our participation in AIM dissemination and training activities as well as our contribution -- via the RUTGERS/LCSR computer -- to the shared computing facilities of the national AIM network.

B. Justification and Requirements for Continued SUMEX Use

Continued access to SUMEX is needed for:

1) Backup for demos, etc.

2) Programs developed to serve the National AIM Community should be runnable on both facilities.

3) There should be joint development activities between the staffs at Rutgers and Sumex in order to ensure portability, share the load, and provide a wider variety of inputs for developments.

C. Needs and Plans for Other Computing Resources beyond SUMEX-AIM

Beyond the current SUMEX-AIM facility there is need for access to a more 'personal' type of computing facilities (e.g., Dolphins). In addition SUMEX might provide a high quality output device (e.g., line printer or XGP) for the community.

D. Recommendations for Future Community and Resource Development

We strongly support the efforts of Sumex staff to develop personal computing resources. Our own local planning efforts have been helped considerably by the assistance of Sumex staff. As physical facilities are spread out in more and more places, we believe that centers such as Sumex can play a growing role in keeping track of technology, planning, and technical support of users. It may also prove convenient for Sumex to serve as a center for mail and other communications activities carried out by personal computers at remote sites.

We believe that computing is going to move towards smaller machines spread around more centers. As this happens, there is some danger that the scientific community which has been built up around the Sumex-AIM computing resource will tend to become fragmented. In this context it becomes increasingly important for the AIM project to sponsor meetings and other kinds of visits involving members of the community. It also becomes important to separate Sumex's role in communications, coordination, and technical consultation from their role as suppliers of computing time. The latter role is likely to decrease outside the Stanford community as time goes on. It is therefore important to make sure that the former continues to be supported, and in fact to grow.

II.A.2.6 Simulation of Cognitive Processes

Simulation of Cognitive Processes

James G. Greeno
Alan M. Lesgold
Learning Research and Development Center
University of Pittsburgh

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Our goal is to improve understanding of the basic cognitive processes involved in the use and acquisition of cognitive skills. Concern is spread between basic competence in reading, mathematics and general problem solving on the one hand and complex expertise in domains such as radiology on the other. A basic form for our theorizing is computer simulation of human performance. Such simulation modeling allows richer understanding in the course of theory building and stronger empirical validation of theory than would otherwise be possible in domains where several separate cognitive systems or components interact during skilled performance.

B. Medical Relevance

Increased understanding of basic cognitive processes is relevant to medical needs in two ways. One form of relevance involves performance of professional medical tasks. Lesgold has been working with other local colleagues to understand the nature of the skill of reading and diagnosing chest radiographs. The work is leading toward both insights into techniques for improving the training of radiologists and improved understanding of the aspects of radiology that seem to be learned only with many years of experience. This latter improved specification of what is hard to learn and why this is the case will hopefully inform decision making on automatic diagnostic facilities and diagnostic "prostheses" whose cost-effectiveness is otherwise difficult to establish. The second form of relevance of basic research in cognition to medical needs is in development of understanding of the cognitive requirements of elementary skills such as reading and arithmetic computation, in which cognitive deficits can constitute severe disability. Improved understanding of these basic skills should provide principles useful in developing a more valid approach to diagnosis and remediation of learning disabilities.

C. Highlights of Research Progress

1. Accomplishments This Past Year

In the study of the acquisition of reading processes, firm data began to emerge documenting the causal role of slow and nonfacile word recognition in hampering progress in the acquisition of higher-order reading skills (Lesgold, Resnick, Roth, & Hammond, 1981; Lesgold & Curtis, in press; Resnick, Lesgold & Roth, in preparation). Recent results of longitudinal research in our laboratory have shown curriculum differences in the primary grades that significantly affect the specific patterns of learning word recognition that different children exhibit. In addition, there has been an improved understanding of the nature of component process interactions in reading that will drive a simulation of the specific problems in learning comprehension skills that come from inadequately practiced word recognition.

The work on radiological skill has been progressing (Lesgold, Feltovich, Glaser, and Wang, 1981). After extensive review of protocol data from both expert radiologists and residents, we have a good sense of the ways in which radiology skill is a blend of opportunistic problem solving skills from medical diagnosis and more specific spatial representation skills for human anatomy. Modeling work is beginning in earnest, now. We will be experimenting with both the sorts of blackboard approaches subsumed under the AGE system and with parallel activation models of the sort being used in perception work (e.g., CAPS).

A modeling effort that was carried out in its initial phase with the SUMEX facility, enabling use of Anderson's ACT system, involved study of the use of spatial information in solving problems in elementary probability theory. In this project we are simulating the processing of information in Venn diagrams in the calculation of probabilities of composite events. The study extends earlier analyses of spatial information in problem solving in the domain of geometry proof exercises, which also included collaboration with Anderson and use of SUMEX in early phases of simulation programming (Anderson, Greeno, Kline, & Neves, in press).

We have developed a simulation of learning by children who have defective procedures for arithmetic calculation. This work has depended on concepts and techniques with which we became familiar initially in collaborative work with Anderson's group. The model simulates learning from special instruction designed to communicate understanding of concepts of place value as well as a correct computational procedure, and thus provides a preliminary hypothesis about the nature of learning in which conceptual understanding and procedural knowledge interact meaningfully.

2. Research in Progress

Research is continuing on all these topics. In each case, we are building, through porting of code from other centers and through some efforts of our own, an appropriate simulation environment on our own VAX-11-780 system. We are clearly building on what we've learned within the SUMEX community in doing this.

The work in reading and radiology has many aspects in common, and we hope to exploit this commonality in simulation modeling. However, the emphases will be slightly different. The radiology effort will concentrate on a highly refined level of skill, even in the learner, while the reading effort deals with the earliest stages of skill acquisition and is particularly concerned with coordinations in time (in active memory) of representations for concepts that are interrelated in a text.

Studies of mathematical cognition are continuing, with increasing emphasis on individual differences and on learning. We are pursuing empirical and theoretical studies of individuals' differences in the use of composite forms in Venn diagrams THAT correspond to composite events that are involved in problems. We are extending the work on learning in arithmetic, relating our model of acquiring understanding to some general principles of heuristic procedure modification, and we will apply these results to analysis of other learning phenomena in the early arithmetic domain, such as those studied by Riley, Greeno, and Heller (forthcoming).

D. List of Relevant Publications

Anderson, J.R., Greeno, J.G., Kline, P.J., & Neves, D.M. Acquisition of problem-solving skill. In J.R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Lawrence Erlbaum Associates, in press.

Lesgold, A.M., Resnick, L.B., Roth, S.F., & Hammond, K.L. Patterns of learning to read. Paper presented a the biennial meeting of the Society for Research in Child Development, 1981. Lesgold, A.M., Feltovich, P., Glaser, R., & Wang, Y. Radiological expertise. Symposium presentation at the annual meeting of the American Educational Research Association, 1981.

Lesgold, A.M., & Curtis, M.E. Learning to read words efficiently. To appear in A.M. Lesgold & C.A. Perfetti (Eds.), Interactive processes in reading, Hillsdale, NJ: Erlbaum, in press.

Lesgold, A.M., & Perfetti, C.A. Interactive processes in reading: Where do we stand? In A.M. Lesgold & C.A. Perfetti (Eds.), Interactive processes in reading, Hillsdale, NJ: Erlbaum, in press.

Lesgold, A.M., & Perfetti, C.A. (Eds.). Interactive processes in reading, Hillsdale, Nj: Erlbaum, in press.

Resnick, L.B., & Lesgold, A.M. A longitudinal study of difficulties in learning to read. To appear in J.P. Das, R. Mulcahy, & A.E. Wall (Eds.), Theory and Research in Learning Disability, forthcoming.

Riley, M.S., Greeno, J.G., & Heller, J.I. Development of children's problem-solving ability in arithmetic. In H.P. Ginsburg (Ed.), Development of mathematical thinking. New York: Academic Press, in press.

Lesgold, A.M. Cognitive simulation modeling on the VAX-11. Behavior Research Methods and Instrumentation, in press.

E. Funding Support

1) National Institute of Education

1. Title: Research on learning and schooling.
2. Principal Investigators: Robert Glaser, University Professor and Co-Director, Learning Research and Development Center (LRDC), University of Pittsburgh; and Lauren Resnick, University Professor of Psychology and Co-Director, LRDC.
3. Funding Agency: National Institute of Education
4. Grant Number: NIE-G-80-0014
5. Total Award: 1 December 79 to 30 November 82, $7,880,729.
6. Current Period: 1 December 80 to 30 November 81, $2,627,067.

2) Office of Naval Research

1. Title: Cognitive and instructional factors in the acquisition and maintenance of skill.
2. Principal Investigators: Robert Glaser, University Professor and Co-Director, LRDC; James Greeno, University Professor; Alan Lesgold, Research Associate Professor of Psychology; Michelene Chi, Research Assistant Professor of Psychology.
3. Funding Agency: Office of Naval Research
4. Contract Number: N00014-79-C-0215
5. Total Award: 1 January 79 to 30 September 83, $1,676,950.
6. Current Period: 1 October 80 to 30 September 81, $247,053.

3) National Science Foundation/National Institute of Education

1. Title: Invention and understanding in the acquisition of computation.
2. Principal Investigators: Lauren B. Resnick, Professor of Psychology and Co-Director, LRDC; and James Greeno, University Professor of Psychology.
3. Funding Agency: National Science Foundation and National Institute of Education (Joint Program)
4. Grant Number: SED78-22289
5. Total Award: $149,967.
6. Current Period: 1 December 78 to 31 August 81

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Distribution via SUMEX

The work on development of radiology skills is being done in collaboration with Dr. Yen Wang, Clinical Professor of Medicine, University of Pittsburgh.

B. Sharing and Interaction with Other SUMEX-AIM Projects

We have shared software from the ACT and AGE projects via SUMEX. In addition, the ties to the ACT project have led to interactions and joint software and programming environment development with psychologists at Carnegie-Mellon University. We expect to maintain a common simulation development environment for psychologists in both institutions, largely because SUMEX proved the value of such standardization. Similarly, especially if the COGNET proposal for a cognitive Science Network is funded, we expect to support a standard demo and guest interface in both places.

C. Critique of Resource Management

None.

D. Future Involvement

This is a final report for this project. Because of the successful acquisition of on-site facilities for artificial intelligence work, we are moving from national project to associate status within the SUMEX-AIM community. We are particularly grateful for the many opportunities SUMEX involvement provided when we most needed them. We have learned from them in building our own resource and expect to continue to share in the community-building aspects of AIM in the future.

II.A.2.7 SECS - Simulation and Evaluation of Chemical Synthesis

SECS - Simulation and Evaluation of Chemical Synthesis

PI: W. Todd Wipke
Board of Studies in Chemistry
University of California
Santa Cruz, California 95064

Coworkers:

D. Dolata	(Grad student)
I. Kim	(Grad student)
D. Rogers	(Grad Student)
J. Chou	(Postdoctoral)
P. Condran	(Postdoctoral)
T. Moock	(Postdoctoral)
H. Kuehmstedt	(Visiting Professor, Univ. of Greifswald, East Germany)

I. SUMMARY OF RESEARCH PROGRAM

A. Technical Goals. The long range goal of this project is to develop the logical principles of molecular construction and to use these in developing practical computer programs to assist investigators in designing stereospecific syntheses of complex bio-organic molecules. Our specific goals this past year included adding a starting material library, initial exploration of strategies based on starting materials, developing the representation for a reasoning module, and completion of the user-defined transform capabilities. The objectives for the XENO project were to establish extensive collaborations with metabolism experimentalists to test XENO predictions and to begin development on methods for assessing the potential biological activity of each metabolite. The world-wide SECS User's Group for sharing chemical transforms was initiated.

B. Medical Relevance and Collaboration

The development of new drugs and the study of how drug structure is related to biological activity depends upon the chemist's ability to synthesize new molecules as well as his ability to modify existing structures, e.g., incorporating isotopic labels or other substituents into biomolecular substrates. The Simulation and Evaluation of Chemical Synthesis (SECS) project aims at assisting the synthetic chemist in designing stereospecific syntheses of biologically important molecules. The advantages of this computer approach over normal manual approaches are many: 1) greater speed in designing a synthesis; 2) freedom from bias of past experience and past solutions; 3) thorough consideration of all possible syntheses using a more extensive library of chemical reactions than any individual person can remember; 4) greater capability of the computer to deal with the many structures which result; and 5) capability of computer to see molecules in graph theoretical sense, free from bias of 2-D projection.

The objective of using XENO (a spinoff of SECS) in metabolism is to predict the plausible metabolites of a given xenobiotic in order that they may be analyzed for possible carcinogenicity. Metabolism research may also find this useful in the identification of metabolites in that it suggests what to look for. Finally, it seems there may even be application of this technique in problem domains where one wishes to alter molecules so certain types of metabolism will be blocked.

C. Highlights of Research Progress

1. Progress and Accomplishments

RESEARCH ENVIRONMENT: At the University of California, Santa Cruz, we have a GT40 and a GT46 graphics terminal connected to the SUMEX-AIM resource by 1200 and 2400 baud leased lines (one leased line supported by SUMEX). We also have a TI725, TI745, CDI-1030, DIABLO 1620, and an ADM-3A terminal used over 300 baud leased lines to SUMEX. UCSC has only a small IBM 370/145, a PDP-11/45, 11/70 and a VAX 11/780, (the 11's are restricted to running small jobs for student time-sharing) all of which are unsuitable for this research. The SECS laboratory is located in a newly renovated room with raised floor in 125 Thimann Laboratories, adjacent to the synthetic organic laboratories at Santa Cruz so the environment is excellent.

2. SECS Program Developments

The Simulation and Evaluation of Chemical Synthesis (SECS) program has undergone many additions to improve its capabilities and usefulness to synthetic chemists.

The ALCHEM language has been made extensible to facilitate addition of new groups. New functional groups and ring systems can be defined and referred to by name within a transform and individual atoms and bonds within the named entity can also be referenced. Thus, a transform may now inquire if an indole ring system exists elsewhere in the molecule, and may then inquire about substituents at various positions on that ring system. Such a query was previously impossible. This kind of query arises frequently in heterocyclic chemistry.

The User Defined Transform module (UDT) allows the chemist to define a chemical transform on-line in the middle of a synthetic analysis. We have now enabled the program to "learn" the transform, i.e., remember it and be able to apply it again where applicable. This is the first work in machine learning of chemical reactions. A graphical front-end to the UDT module was added and the whole module incorporated into the production SECS version.

A META-SECS top-level plan generator is being implemented to reason using synthetic principles and conclude plans which will then be used to guide the existing SECS program in synthetic analysis. The First Order

Predicate Calculus is being used to represent the synthetic strategies and principles. The statement parser and data structure manipulation routines have been implemented, the inference engine is in progress at this time. This explicit reasoning in synthetic strategies should prove very interesting as a way to control development of the synthesis tree.

Synthesis tree output is now possible on the PRINTRONIX printer/ plotter using a new software vector to raster conversion program developed by our group. The resulting raster file is compressed, transmitted from SUMEX to Santa Cruz, then expanded and fed to the printer. This approach avoids our previous problems resulting from the long transmission time required to plot the big trees over a slow communication line and the mechanical problems associated with pen plotters.

The Aldrich catalog of commercially available chemicals has been converted from Wisswesser line notation to connection table and to the SEMA unique representation. This allows SECS to determine if a precursor is commercially available. Initial studies of using this library to develop synthetic strategies has begun. Several algorithms have been developed to identify potential starting materials for a target molecule The effectiveness of these algorithms is being evaluated against literature syntheses.

3. XENO - A Program to Predict Plausible Metabolites

The XENO program was developed to assist metabolism researchers in predicting plausible metabolites of compounds foreign to an organism, and in evaluating the potential biological activity of the resulting metabolites. The knowledge base of XENO has been revised completely and now includes 110 types of metabolic processes. We have specialized on rat and mouse systems to date. The XENO program takes graphical input of a compound to be metabolized and stepwise generates a tree of metabolite structures which might result.

The second phase of XENO which evaluates potential biological activity is underway. We are developing a series of rules for each class of compounds to relate structure to biological activity. Work to date has concentrated on aromatic amines and polycyclic aromatic hydrocarbons.

Collaborations with metabolism experimentalists have begun in order that XENO can make predictions for compounds actively being studied in the laboratory. Initial feedback from ICI Pharmaceuticals on one particular drug indicates that the major metabolites were correctly generated and that XENO proposed structures that may explain some polar compounds that ICI has not yet been able to identify. XENO generated several possible pathways to explain an unusual metabolite. ICI is seeking to differentiate between the pathways. We are similarly collaborating with Mead Johnson. Initial results there indicated several transforms were missing from the XENO library or were improperly represented. Those have been corrected and new analyses sent for evaluation. Other collaborations are in progress with investigators at NIH and at the Australian National University.

We are also comparing XENO analyses to results reported in the literature for compounds such as cyclophosphamide. This work is sponsored by the National Cancer Institute.

D. List of Current Project Publications

M.L. Spann, K.C. Chu, W.T. Wipke, and G. Ouchi, "Use of Computerized Methods to Predict Metabolic Pathways and Metabolites," J. of Env. Pathology and Toxicology, 2, 123 (1978); also reprinted in "Hazards from Toxic Chemicals," ed. M.A. Mehlman, R.E. Shapiro, M.F. Cranmer and M.J. Norvell, Pathotox Publishers, Inc., Park Forest South, Ill., 1978, pp. 123-121.

P. Gund, E.J.J. Grabowski, D.R. Hoff, G.H. Smith, J.D. Andose, J.B. Rhodes, and W.T. Wipke, "Computer-Assisted Synthetic Analysis at Merck," J. Chem. Info. and Comput. Sci., 20, 288 (1980).

S.A. Godleski, P.v.R. Schleyer, E. Osawa, and W. T. Wipke, "The Systematic Prediction of the Most Stable Neutral Hydrocarbon Isomer," Progress in Physical Organic Chemistry, Vol. 13, 1981, pp. 63-118.

R.E. Carter and W.T. Wipke, "SECS--EH Hjalpmedel Vid Organisk Syntesplanering," Kemisk Tidskrift, in press.

W.T. Wipke, G. Ouchi, and J. Chou, "Computer-Assisted Prediction of Metabolites," Proceedings of Conference Structure Activity Relations, Research Triangle Park, N.C., Feb., 1980.

E. Funding Status

1) Resource-Related Research: Biomolecular Synthesis
   PI: W. Todd Wipke, Associate Professor, UCSC
   Agency: NIH, Research Resources
   No: RR01059-03S1
   7/1/80-12/31/81 $ 36,949 TDC

2) Computer-Aided Prediction of Metabolites for Carcinogenicity Studies
   PI: W. Todd Wipke
   Agency: NIH, National Cancer Institute
   No: NO1-CP-75816
   1/1/80-7/31/81 $74,394 TDC

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX. SECS is available in the GUEST area of SUMEX for casual users, and in the SECS DEMO area for serious collaborators who plan to use a significant amount of time and need to save the synthesis tree generated. Much of the access by others has been through the terminal equipment at Santa Cruz because graphic terminals make it so much more convenient for structure input and output.

Prof. H. Kuehmstedt of the University of Greifswald, E. Germany used SECS to generate some interesting and novel synthetic routes to progesterone. Demonstrations and sample synthetic analyses were generated for Drs. Terry Brunck and Steve Roman of Shell Development, John Harper of Amoco Chemicals, Prof. Fujiwara, University of Tsukuba, Japan, Dr. Peder Berntsson, Hassle, Sweden. Other visitors included Dr. M. Onozuka, A. Tomonaga and H. Itoh, Kureha Chemical Co., Tokyo, Japan, Dr. Rhyner, Director of Research, Ciba-Geigy, Basel, Switzerland. Demonstrations of SECS in Sweden were performed by Dr. Carter at many universities and companies. Ned Phillips of the College of Pharmacy, Univ. of Florida is accessing SECS via SUMEX GUEST access. A synthesis of vellerolactone, a substance found to be toxic and teratogenic was generated for Prof. R. E. Carter, Univ. Lund, Sweden.

Dr. Wipke has also used several SUMEX programs such as CONGEN in his course on Computers and Information Processing in Chemistry. Communication between SECS collaborators is facilitated by using SUMEX message drops, especially since the time difference between the U.S. and Europe and Australia makes normal telephone communication practically impossible. Testing and collaboration on the XENO project with researchers at the NCI depend on having access through SUMEX and TYMNET.

B. Examples of Sharing, Contacts and Cross-Fertilization with Other SUMEX-AIM Projects. This year the SECS and XENO project have made use of the teletype plot program which Ray Carhart of the CONGEN project wrote at Stanford. We modified the program to fit the needs of our projects. This was facilitated by being able to transfer the programs within areas on the same computer system at SUMEX. We continue to have intellectual interactions with the DENDRAL and MOLGEN project in areas where we have common interests and have had people from those projects speak at our group seminars. SUMEX also is used for discussions with others in the area of artificial intelligence on the ARPANET.

We developed a local print capability through SUMEX with the help of the SUMEX staff which has facilitated our work greatly. We have also communicated with SUMEX staff regarding selection of terminals and other computer equipment.

C. Critique of Resource Services. We find the SUMEX-AIM network very well human engineered and the staff very friendly and helpful. The SECS project is probably one of the few on the AIM network which must depend exclusively on remote computers, and we have been able to work rather effectively via SUMEX. Basically we have found that SUMEX-AIM provides a productive and scientifically stimulating environment and we are thankful that we are able to access the resource and participate in its activities.

SUMEX-AIM gives us at UCSC, a small university, the advantages of a larger group of colleagues, and interaction with people all over the country. We especially thank SUMEX for support of the leased line for our GT40, and for helping develop our remote print capability.

SUMEX however has fallen short of our goals and desires: the load average on SUMEX has increased and reduced my group's efficiency greatly--the system is too overloaded. We are installing SECS on the 2020 to begin to make use of that additional capability. We also have not been able to utilize the 4800 baud high speed line we purchased because SUMEX limitations forced running at 2400 baud. We had hoped to be able to write tapes locally with the 4800 baud line, but at 2400 baud it is too slow to be practical. We would like to see some of their local lines slowed down so those remote people doing graphics can run at a higher speed. We have found that when a FORTRAN program is overlayed, the symbol table is lost, making symbolic debugging with DDT impossible, we wish that could be corrected.

D. Collaborations and Medical Use of Programs via Computers other than SUMEX. SECS 2.9 now resides on the CompuServe computer networks so anyone can access it without having to convert code for their machine. This has proved very useful as a method of getting people to try this new technology. Dr. George Purvis of Battelle is accessing SECS via CompuServe, as are Gene Dougherty of Rohm and Haas and many others. SECS also resides on the Medicindat machine at the University of Gothenborg, Sweden, and is available all over that country by phone. Similarly in Australia, SECS resides at the University of Western Australia and is available throughout the country over CSIRONET. Plans are underway for a similar situation in Japan.

III. RESEARCH PLANS (6/81-6/82)

A. Long Range Project Goals and Plans. The SECS project now consists of two major efforts, computer synthesis and metabolism, the latter being a very young project. Our plans for SECS for the next year include completing the high level reasoning module for proposing strategies and goals, and providing control which continues over several steps. This reasoning module also will be able to trace the derivation of goals and thus explain some of its reasoning. We also plan to focus on bringing the transform library up in sophistication to improve the performance and capabilities of SECS. In particular we plan to allow a transform to have access to the precursors generated as well as to the product, this will allow much greater control and more natural transform writing, but it requires extensive changes in the SECS control structure to permit this.

We will continue to explore starting material oriented strategies based on the Aldrich Chemical file we now have implemented. We especially are interested in chirality based strategies which we feel are very strong.

We plan to explore running SECS on a virtual memory 32-bit computer like a VAX-11/780 or a PRIME since many chemistry departments now have these machines available and thus could run SECS.

The XENO metabolism project will be expanding the data base to cover more metabolic transforms, including species differences, sequences of transforms, and stereochemical specificities of enzymatic systems. Development of the second phase which assesses the biological activity of the metabolites will continue as will efforts to simulate excretion and incorporation, the endpoints of metabolism. Finally, application of the current program to the molecules actively being investigated by metabolism researchers will occur concurrently to test and verify the work done to date on XENO and provide examples for publication.

In the next five years we foresee the SECS and XENO projects reaching a stage of maturity where they will find much application in other research groups. Our research will continue in these areas, but turn to some new programs that approach the problems from different viewpoints and allow us an opportunity to begin fresh taking advantage of what we have learned from the building of SECS and XENO.

B. Justification and Requirements for Continued Use of SUMEX. The SECS and XENO projects require a large interactive time-sharing capability with high level languages and support programs. I am on the campus computing advisory committee and am the campus representative to the UC systemwide computing advisory committee and know that the UCSC campus is not likely in the future to be able to provide this kind of resource. Further there does not appear to be in the offing anywhere in the UC system a computer which would be able to offer the capabilities we need. Thus from a practical standpoint, the SECS and XENO projects still need access to SUMEX for survival.

Scientifically, interaction with the SUMEX community is still extremely important to my research, and will continue to be so because of the direction and orientation of our projects. Collaborations on the metabolism project and the synthesis project need the networking capability of SUMEX-AIM, for we are and will continue to be interacting with synthetic chemists at distant sites and metabolism experts at the National Cancer Institute. Our requirements are for good support of FORTRAN.

Our needs for SUMEX include fixing the overlay loader so that an overlaid program can retain its symbol 1 table and permit symbolic use of DDT. This is a serious problem we hope can be fixed by SUMEX staff because without symbols, debugging is very difficult and time-consuming, since we must run SECS and XENO overlaid.

C. Needs beyond SUMEX-AIM. We do plan to acquire a virtual memory minicomputer like a VAX or PRIME in the future to offload some of our processing from SUMEX. Such a machine would enable us to do some production and development work locally and would explore the feasibility of those types of machines as hosts for SECS and XENO. A local machine would also free us from the problems we have experienced in the winter when the telephone lines to Stanford get wet and are too noisy to use. Even if we had such a machine we still need to use SUMEX because we plan to continue to develop and maintain the PDP-10 version of SECS and we need SUMEX for its networking capabilities. In the future if we had a mini at UCSC, we would lighten our load on SUMEX, but currently we see our load increasing as our group grows and as we start new projects yet must maintain existing large programs.

We especially need the local capabilities to read and write magnetic tape because we receive and send many tapes between our collaborators. Driving to SUMEX to write a tape is not efficient for our personnel and hinders communication with collaborators via tape. The problem will worsen because the SECS Users Group will be sending UCSC tapes of chemical transforms on a regular basis.

D. Recommendations for Community and Resource Development. The AIM Workshops have been excellent in the past and should be continued. We feel the SUMEX resource is too heavily utilized at times to get any productive work done. SUMEX staff could lighten the load on the machine by reducing the speed of text terminals at Stanford from 2400 baud and above down to 1200 baud which is plenty fast for humans to read, and giving remote users faster capabilities, say 4800 baud. We feel the community would benefit if remote users such as we had a virtual minicomputer so the load could be distributed more and not have everything go through Stanford which is highly congested and quite expensive for multiple leased lines. SUMEX can not currently handle an increase in the outside community using SECS or XENO for testing. The response time guests and outside collaborators see is not a good reflection on the actual efficiency of the programs.

II.A.2.8 SOLVER Project

SOLVER: Problem Solving Expertise

Dr. P. E. Johnson
Center For Research In Human Learning
University of Minnesota

Dr. W. B. Thompson
Department of Computer Science
University of Minnesota

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

This project focuses upon the development of strategies for discovering and documenting the knowledge and skill of expert problem solvers. In the last fifteen years, great progress has been made in synthesizing the expertise required for solving extremely complex problems. Computer programs exist with competency comparable to human experts in diverse areas ranging from the analysis of mass spectrograms and nuclear magnetic resonance [DENDRAL] to the diagnosis of certain infectious diseases [MYCIN].

Design of an expert system for a particular task domain usually involves the interaction of two distinct groups of individuals, "knowledge engineers," who are primarily concerned with the specification and implementation of formal problem solving techniques, and "experts" (in the relevant problem area) who provide factual and heuristic information of use for the problem solving task under consideration. Typically, the knowledge engineer, after consulting with one or more experts, decided on a particular knowledge representational structure and inference strategy. Next, "units" of factual information are specified. That is, properties of the problem domain are decomposed into a set of manageable elements suitable for processing by the inference operations. Once this organization has been established, major efforts are required to refine representations and acquire factual knowledge organized in an appropriate form. Major research problems exist in developing more effective representations, improving the inference process, and in finding better means of acquiring information from either experts or the problem area itself.

Programs currently exist for empirical investigation of some of these questions for a particular problem domain [AGE, UNITS, RLL]. These tools allow the investigation of alternate organizations, inference strategies, and rules bases in an efficient manner. What is still lacking, however, is a theoretical framework capable of reducing dependence on the expert's intuition or on near exhaustive testing of possible organizations. Despite their successes, there seems to be a consensus that expert systems could be better than they are. Most expert systems embody only the limited amount of expertise that individuals are able to report in a particular, constrained language (e.g. production rules). If current systems are approximately as good as human experts, given that they represent only a portion of what individual human experts know, then improvement in the "knowledge capturing" process should lead to systems with considerably better performance.

B. Medical Relevance and Collaboration

Collaboration with Dr. James Moller, Department of Pediatrics, and Dr. Donald Connelly in the Department of Laboratory Medicine and Pathology, both at the University of Minnesota Medical School.

C. Highlights of Research Progress

Accomplishments of this past year. Prior research at Minnesota on expertise in diagnosis of congenital heart disease has resulted in a theory of diagnosis and an embodiment of that theory in the form of a computer simulation model which diagnoses cases of congenital heart disease.

At a macroscopic level, the simulation model contains four categories of knowledge used by the expert physician (pediatric cardiologist) in diagnosis. First, the model has clinical knowledge of disease. This knowledge is hierarchically structured, including categories of disease, specific diseases, and variants of the same disease that differ in presentation. From each element in the hierarchy, there is knowledge of the associated anatomy and physiology and the expected clinical manifestations. Second, the model has deductive knowledge of disease - knowledge of principles of cardiovascular pathophysiology and the clinical manifestations useful in detecting underlying pathology. The causal knowledge deductively relates cardiovascular defects to hemodynamics and to expected patient data. This category of knowledge is not typically used in diagnosis, since the expert physician can simply recall expected clinical manifestations through clinical knowledge of disease, rather than deduce them. Third, the model has heuristic knowledge of disease and of clinical findings useful in its diagnosis. One aspect of this knowledge provides indices from clinical manifestations to diseases and from disease to disease, useful in choosing diagnostic hypotheses to consider. Another aspect of heuristic knowledge is related to evaluation of diagnostic alternatives -- rules of thumb for ruling in and ruling out alternatives as correct or incorrect. Another form of heuristic knowledge screens abnormal from normal clinical findings and identifies subsets of findings which are likely to have the same underlying cause. All heuristics are aimed at basically the same end -- reducing the cognitive demands of diagnosis without loss of diagnostic accuracy. The fourth and last category is knowledge of data acquisition techniques: interviewing methods, physical exam maneuvers, special procedures, and laboratory utilization.

Diagnosis is characterized as a heuristic search process, with four sources of knowledge involved. Clinical knowledge of disease is the hierarchical structure to be searched. Deductive knowledge of disease is useful in construction of missing pieces of that hierarchy and in justifying the clinical information it contains. Heuristic knowledge aids in limiting the section of clinical knowledge to be searched and in providing simple evaluation functions for use in search. Data acquisition for knowledge is essential in obtaining patient information to be used in the search process.

The simulation model is currently programmed in UT Lisp on the University of Minnesota CDC Cyber 74; it contains just over one million characters. One of the first goals of our project is to transfer this program over to the Sumex-AIM system - see Project goals and Plans (below).

<u>Research</u> <u>in</u> <u>progress</u>. The methodology of our research derives from the discipline of cognitive science, and from our study of expert problem solvers. This methodology consists of: (1) extensive use of verbal thinking aloud protocols as well as other experimental data as a source of information from which to make inferences about underlying cognitive structures and processes; (2) development of computer models as a means of testing the adequacy of inferences derived from the protocol studies; (3) testing and refinement of the cognitive models based upon the study of human and model performance in experimental settings.

This past year we have been investigating expertise in solving certain classes of physics problems. We have chosen this area because reasonable data on human competency is available from past work at Minnesota and elsewhere, some work on formal modeling has already been done, and both the problems and required knowledge are well specified. Effort will be concentrated on investigating control structures and search heuristics. Those portions of the system unique to physics will be isolated so that generality can be easily studied by extending the program to other domains.

Once a computational model has been implemented, three classes of experiments will provide useful information. First, a series of validation tests will be used to estimate the effectiveness of the model. A series of problems will be given to the model and a solution trace produced. The degree to which the program's behavior corresponded to the ways in which the humans solved the same problems will be determined. In addition, the effectiveness of the heuristics in producing efficiencies relative to alternative approaches to the problems will be estimated. In the second set of experiments, the model will be optimized for different types of problems within the same domain to determine which heuristics are likely to be problem specific. We are particularly interested in the degree to which the problems are well specified and the effects of problem space topology.

Expert/novice differences will be investigated by operating the model with different levels of sophistication in its knowledge base. The last set of experiments will investigate errors in problem solving. Classes of problems likely to produce an incorrect answer will be identified. The degree to which these errors are a necessary consequence of the search heuristics will be investigated. Finally, the model will be extended so that when garden path errors are recognized, a generalization process is invoked so that similar situations can be avoided in the future.

D. List of Relevant Publications

Connelly, D., & Johnson, P.E. The medical problem solving process. Human Pathology, 1980, 11, 412-419

Elstein, A., Gorry, A., Johnson, P., & Kassirer, J. New research direction. In Clinical Decision Making and Laboratory Use. D.C. Connelly, E. Benson, & D. Burke (Eds.), University of Minnesota Press (in press).

Feltovich, P.J., Knowledge based components of expertise in medical diagnosis. Unpublished doctoral dissertation, University of Minnesota, 1981.

Johnson, P.E. Cognitive models of medical problem solvers. In D.C. Connelly, E. Benson, & D. Burke (Eds.) Clinical Decision Making and Laboratory Use. University of Minnesota Press (in press).

Johnson, P.E., Severance, D.G., & Feltovich, P.J. Design of decision support systems in medicine: Rationale and principles from the analysis of physician expertise. Proceedings of the Twelfth Hawaii International Conference on System Sciences, Western Periodicals Co., 1979, 3, 105-118.

Johnson, P.E., Barreto, A., Hassebrock, F., Moller, J., & Prietula, M. Expertise and error in diagnostic reasoning. Cognitive Science (in press).

Johnson, P.E., & Thompson, W.B. Strolling down the garden path: Detection and recovery from error in expert problem solving. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, B.C., August 1981.

Moller, J.H., Bass, G.M., Jr., & Johnson, P.E. New techniques in the construction of patient management problems. Medical Education (in press).

Swanson, D.B., Computer simulation of expert problem solving in medical diagnosis. Unpublished doctoral dissertation, University of Minnesota, 1978.

Swanson, D.B., Feltovich, P.J., & Johnson, P.E. Analysis of physician expertise: Implications for the design of decision support systems. In D.B. Shires & H. Wold (Eds.), Medinfo77. Amsterdam: North-Holland Publishing Co., 1977.

E. Funding and Support

Work being done in scientific reasoning is sponsored under a current NSF (SE079-13036) grant to Paul Johnson. The work in law has been supported by by the Minnesota Center for Research in Human Learning and is described in a proposal which is currently under review by the NSF Law and Social Science Program. The work in medicine has been supported by NICHD

(T36-HD-17151 and HD-01136) and NSF (NSF/BNS-77-22075) grants to the Minnesota Center for Research in Human Learning of which Paul Johnson is a principal investigator. Additional support is being sought from the National Library of Medicine and the Information Sciences Program at NSF.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

Collaboration with Tufts University - New England Medical Center (See III.A.)

B. Sharing and Interactions with Other SUMEX-AIM Projects

Our work is complementary to many of the current projects supported on Sumex-AIM. We will be investigating certain aspects of expert problem solving in order to develop better organizational and knowledge acquisition strategies. Such work requires that we be able to build upon the extensive experience in knowledge engineering within the Sumex-AIM communities. Specifically, we first need to investigate the number of existing programs in order to determine the degree to which they satisfy the design goals which we will be establishing. We then hope to use the program construction tools that are available in order to build prototype systems to illustrate our ideas.

C. Critique of Resource Management

(None)

III. RESEARCH PLANS

A. Project Goals and Plans

Near term. The research for which we wish to use the Sumex-AIM resource has two subcomponents, one directed at selective reimplementation of the diagnosis program (described above) and the other directed at extensions of the research more broadly. Reimplementation of the diagnosis program will be carried out first, using AGE and UNITS knowledge engineering tools, and Interlisp. The present program is a blend of a semantic network representation (for clinical knowledge of disease) in a production system/blackboard control structure (heuristic activation of portions of clinical knowledge). This structure is nicely congruent with the facilities provided through UNITS and AGE linked together. This relatively well specified reimplementation task will provide an appropriate environment for learning about Sumex-AIM and the resources it provides. It will also provide an interesting comparison of performance of similar diagnosis programs implemented in different ways. The second subcomponent of the proposed research will focus upon data collection and patient management skills of the expert physician. Work thus far has focused largely on diagnosis because of our lack of familiarity with the medical domain and because of the paramount importance of methodological development, as discussed in the next section on knowledge capturing. In pediatric cardiology, we now have the substantive knowledge and methodological tools required to approach medical problem solving more broadly, including knowledge and procedures related to both collection of patient data for diagnostic purposes and patient management. In collaboration with experimental psychologists and physicians at Minnesota, and computer scientists and physicians at Tufts University in Boston we propose to investigate the stimulus information utilized by physicians in various patient data sources (X-ray, EKG, and heart sounds).

Long range. We propose to investigate the "knowledge capturing" process that occurs in the early stages of the development of expert systems when problem decomposition and solution strategies are being specified. Several related questions will be addressed: What are the performance consequences of different organization approaches, how can these consequences be evaluated, and what tools can assist in making the best choice? How can organizations be determined which not only perform well, but are structured so as to facilitate knowledge acquisition from human experts?

B. Justification and Requirements for Continued SUMEX Use

We are currently using a CDC Cyber 74 and Cyber 172 for most of our work in problem solving. The Cyber computers are not well suited for interactive computing and have serious limitations with respect to address space and available support software.

C. Needs and Plans for Other Computing Resources beyond SUMEX-AIM

We expect delivery on a VAX-11/780 in June 1981 at the University of Minnesota. Hopefully, much of our work will eventually be implementable on the VAX. However, this depends both on our ability to acquire additional peripheral hardware and on increased availability of AI software for the VAX. Until that time, access to a TENEX facility is extremely desirable.

D. Recommendations for Future Community and Resource Development

(None)

II.A.3 Pilot Stanford Projects

The following are descriptions of the informal pilot projects currently using the Stanford portion of the SUMEX-AIM resource pending funding, and full review and authorization.

II.A.3.1 Protein Secondary Structure Project

Protein Secondary Structure Project

Robert M. Abarbanel, M.D.
University of California Medical Center
University of California at San Francisco

I. SUMMARY OF RESEARCH

A. Project Rationale

Development of a protein structure knowledge base and tools for manipulation of that knowledge to aid in the investigation of new structures. System to include cooperating knowledge sources that work under the guidance of other system drivers to find solutions to protein secondary structure problems. Evaluations of structure predictions using known proteins and other user feedbacks available to aid user in developing new methods of prediction.

B. Medical Relevance and Collaboration

Many important proteins have been sequenced but have not, as yet, had their secondary or tertiary structures revealed. The systems developed here would aid medical scientists in the search for particular configurations, for example, around the active sites in enzymes. Predictions of secondary structure will aid in the determination of the full "natural" configuration of important biological materials. Development of systems such as these will contribute to our knowledge of medical scientific data representation and retrieval.

C. Highlights of Research Progress

This is a relatively new project at SUMEX. During the last year, a representation of protein sequences and rules for their manipulation have been built using the UNITS system developed at SUMEX. At this time, various existing structure prediction algorithms are being implemented to explore the needs for tools in a fully developed system where a scientist may describe a new prediction scheme in nearly natural language, and see it run on protein sequences with aids to discovery of errors and problems in the new methods developed.

D. List of Relevant Publications

None.

E. Funding Support

New Investigator Award proposal pending with NLM.

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Medical Collaborations

None.

B. Sharing and Interactions with SUMEX Projects

This project is closely allied with the MOLGEN group, both in computer and scientific interests. Some pattern matching methodology created for the protein data base has been adopted and used in the various DNA knowledge bases. The principal persons in the MOLGEN group have contributed to this project's use and understanding of knowledge base software and resources.

C. Critique of Resource Management

System load has been a significant problem. Except during late night or early morning hours, use of complex packages like the UNITS editor or other system available text editors is a major burden. Often the schedule manager has moved the active processes to background. This is a problem shared with many others in the community using knowledge base tools. It is a common practice to transfer files to other computing resources for interaction during daylight hours. Communications via MSG have aided this project in establishing connections with other researchers and groups with parallel interests.

III. RESEARCH PLANS

A. Project Goals and Plans

Near-Term: Design of rules modules in UNITS editor that will allow an inexperienced user to express algorithms for structure prediction in near-natural language. These prediction schemata will then be translated into appropriate combinations of invocations of knowledge sources, editors, and feedback tools to aid the user in further refinement of his algorithms.

Long-Term: Systems to be developed for the discovery of rules and/or algorithms that can transform a protein sequence into a secondary structure, possibly with implications about tertiary structure. This will involve an effort along the lines of the meta-DENDRAL project.

B. Need for Resources

1. SUMEX Resources

Environment of knowledge base tools and people is the primary motive for doing this work using SUMEX. Access to both established and developing systems aids this project in setting down standards of excellence, forward thinking about computing tools and methodologies, and active exchange of techniques and ideas. The close collaboration with the MOLGEN researchers is particularly useful in this regard.

2. Other Computing Resources

A soon to be established network connection with the Computer Graphics Laboratory at UCSF will provide access to 1) the latest in protein structural information, and 2) color line drawing graphics facilities for evaluation and display of this projects product. A real time display using color graphics will become a possibility.

Connections with the HPP Unix system may also prove useful for program sharing.

II.A.3.2 Ultrasonic Imaging Project

ULTRASONIC IMAGING PROJECT

James F. Brinkley, M.D.
W.D. McCallum, M.D.
Depts. of Computer Science,
Obstetrics and Gynecology
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The long range goal of this project is the development of an ultrasonic imaging and display system for three-dimensional modeling of body organs. The models will be used for non-invasive study of anatomic structure and shape as well as for calculation of accurate organ volumes for use in clinical diagnosis. Initially, the system will be used to determine fetal volume as an indicator of fetal weight; later it will be adapted to measure left ventricular volume, or liver and kidney volume.

The general method we are using is the reconstruction of an organ from a series of ultrasonic cross-sections taken in an arbitrary fashion. A real-time ultrasonic scanner is coupled to a three-dimensional acoustic position locating system so that the three-dimensional orientation of the scan plane is known at all times. During the patient exam a dedicated microcomputer based data acquisition system is used to record a series of scans over the organ being modelled. The scans are recorded on a video tape recorder before being transferred to a video disk. 3D position information is stored on a floppy disk file. In the proposed system the microprocessor will then be connected to SUMEX where it will become a slave to an AI program running on SUMEX. The SUMEX program will use a model appropriate for the organ which will form the basis of an initial hypothesis about the shape of the organ. This hypothesis will be refined at first by asking the user relevant clinical questions such as (for the fetus) the gestational age, the lie of the fetus in the abdomen and complicating medical factors. This kind of information is the same as that used by the clinician before he even places the scan head on the patient. The model will then be used to request those scans from the video disk which have the best chance of giving useful information. Heuristics based on the protocols used by clinicians during an exam will be incorporated since clinicians tend to collect scans in a manner which gives the most information about the organ. For each requested scan a prototype outline derived from the model will be sent to the microcomputer. The requested scan will be retrieved from the video disk, digitized into a frame buffer, and the prototype used to direct a border recognition process that will determine the organ outline on the scan. The resulting outline will be sent to SUMEX where it will be used to update the model. The scan requesting process will be continued until it is judged that enough information has been collected. The final model will then be used to determine volume and other quantitative parameters, and will be displayed in three dimensions.

We believe that this hypothesize verify method is similar to that used by clinicians when they perform an ultrasound exam. An initial model, based on clinical evidence and past experience, is present in the clinician's mind even before he begins the exam. During the exam this model is updated by collecting scans in a very specific manner which is known to provide the maximum amount of information. By building an ultrasound imaging system which closely resembles the way a physician thinks we hope to not only provide a useful diagnostic tool but also to explore very fundamental questions about the way people see.

We are developing this system in phases, starting with an earlier version developed at the University of Washington. During the first phase the previous system has been adapted and extended to run in the SUMEX environment. Clinical studies have been initiated to determine its effectiveness in predicting fetal weight and left ventricular volume. At the same time computer vision techniques are being studied in order to develop the system further in the direction of increased applicability and ease of use. We thus hope to develop a limited system in order to demonstrate the feasibility of the technique, and then to gradually extend it with more complex computer processing techniques, to the point where it becomes a useful clinical tool.

B. Medical Relevance

This project is being developed in collaboration with the Ultrasound Division of the Department of Obstetrics at Stanford, of which W.D. McCallum is the head.

Fetal weight is known to be a strong indicator of fetal well-being: small babies generally do more poorly than larger ones. In addition, the rate of growth is an important indicator: fetuses which are "small-for-dates" tend to have higher morbidity and mortality. It is thought that these small-for-dates fetuses may be suffering from placental insufficiency, so that if the diagnosis could be made soon enough early delivery might prevent some of the complications. In addition such growth curves would aid in understanding the normal physiology of the fetus. Several attempts have been made to use ultrasound for predicting fetal weight since ultrasound is painless, noninvasive, and apparently risk-free. These techniques generally use one or two measurements such as abdominal circumference or biparietal diameter in a multiple regression against weight. We recently studied several of these methods and concluded that the most accurate were about +/-200 gms/kg, which is not accurate enough for adequate growth curves (the fetus grows about 200 gms/week). The method we are proposing is based on the assumption that fetal weight is directly related to volume since the density of fetal tissue is nearly constant. We are hoping that by utilizing three dimensional information more accurate volumes and hence weights can be obtained.

In addition to fetal weight, the current implementation of this system is now being evaluated for its ability to determine other organ volumes. In collaboration with Dr. Richard Popp of the Stanford Division of Cardiology we have started to evaluate the system on in vitro hearts. Left ventricular volumes are routinely obtained by means of cardiac catheterization in order to help characterize left ventricular function. Attempts to determine ventricular volume using one or two dimensional information from ultrasound has not demonstrated the accuracy of angiography. Therefore, three-dimensional information should provide a more accurate means of non-invasively assessing the state of the left ventricle.

C. Highlights of Research Progress

During the past year we have evaluated a large portion of the first phase of our implementation. As mentioned in previous reports the first phase involves the manual outlining of scans with a light pen and the interpolation of an ad hoc model to the data. The system as it stands now is entirely data directed and hence has no notion of what kind of organ it is looking at. However we expect that it will still give more accurate volumes than any now obtainable, and thus will provide the medical impetus for continued research into model directed approaches. The following is a list of accomplishments over the past year:

1. Elimination of most remaining bugs in the data acquisition and SUMEX modelling software.

2. Development of a user's guide, to the point where the system is now routinely used by a person other than the system developer.

3. An evaluation of the ability of the system to locate a point in three dimensions, the result of which shows that our accuracy is about 5-6mm. The sources of this uncertainty were evaluated separately and found to come about equally from the resolution of the ultrasound scanner (2mm), the resolution of the position locator (2mm) and the combination light pen and display monitor. An additional error source, namely the exact parameters describing the mounting of the 3D position locating device to the ultrasound scanner, was essentially eliminated by employing a relaxation procedure to optimize these parameters.

4. The system's ability to calculate balloon volumes was tested on 10 balloons ranging between 257 and 2092 cc. Each balloon was imaged 3 times. Most of the errors between measured and calculated volumes fell within 2%.

5. The system's ability to calculate kidney volumes was tested on two each of pork, sheep and human kidneys, again imaging each organ 3 times. Most of the errors fell within 5%.

6. A clinical study of the ability of the system to predict fetal weight has been underway for about 9 months. A clinical protocol, which directs how scans are to be acquired from the patients, has been developed and refined over that time. This protocol has evolved into a "bottom up" method of collecting on video tape and floppy disk essentially all the information that it is possible to collect from the fetus. Although all this information is not currently being used it will remain

available for use by later model directed versions of the system. In the present study we are comparing the ability of fetal head and trunk 3D reconstructions versus a series of 3D points and axes to predict fetal weight. Two patient populations are being tested: dead babies in a water tank (to test the accuracy in an "ideal" case), and live term babies in utero. In each case calculated volumes are compared against measured weights in a multiple regression using the program STATPACK, which is present on SUMEX. At present we have examined 26 dead babies and 31 live babies. The results so far indicate, somewhat surprisingly, that a combination of simple measurements such as head and abdominal diameters, does about as well as head and trunk 3D reconstructions at predicting weight. However, neither method does well enough to meet our initial goal of better than 200 gms/kg. The most likely explanation seems to be that although we are acquiring a large amount of information, we are not processing it in an "intelligent" manner. Computer plots of the fetal reconstructions show that because of our repeatability errors (see C3 above) and because of fundamental limitations of ultrasound, we are not able to obtain complete and accurate 3D fetal surface points. Although a human looking at the data can easily infer the fetal shape the computer is unable to do this because it has no notion of the shape of a fetus. Thus this study merely justifies more than ever the need for a model based approach.

The research that is currently underway includes the following:

1. Continuation of the clinical study on live term babies since 31 patients is not enough to make definite conclusions. Also since this data will remain available, the evaluation of a model based system should be much easier in the future.

2. In vitro testing of the accuracy of the system on hearts, in collaboration with the Division of Cardiology. Methods of measuring true left ventricular volume are being perfected along with the collection of data. We expect (and have already partially shown on balloons and kidneys) that the current system will do fairly well on hearts, which are much simpler than fetuses. However, even if accurate heart volumes are achieved a model directed approach will be necessary in order to reduce the extreme cumbersomeness of the present system. (It now takes about 2 hours to manually outline scans and to produce a volume).

3. A literature survey is being conducted in the areas of computer vision and 3D modelling in order to determine what methods have been used in the past to represent 3D shapes. The most promising method at present seems to be generalized cylinders, as proposed by Tom Binford at Stanford. The cylinders would probably be represented as uniform cubic splines, as suggested by Uri Shani at Rochester. An additional

requirement will be shape constraints so that the class of all hearts or fetuses can be parameterized such that the given data determines an instance of the generic model.

D. Publications

Brinkley, J.F., Moritz, W.E., Baker, D.W., "Ultrasonic Three-Dimensional Imaging and Volume From a Series of Arbitrary Sector Scans", Ultrasound in Medicine and Biology, vol 4, pp 317-327.

Brinkley, J.F., McCallum, W.D., Daigle, R.E., "A Distributed Computer System for Fetal Weight Determination", Proceedings of the 24th Annual Meeting of the American Institute of Ultrasound in Medicine, Montreal, August 27-31, 1979, p 113.

Brinkley, J.F., McCallum, W.D., Muramatsu, S., "Three-Dimensional Display of Fetus, Placenta and Uterus using an Ultrasonic Computer Modelling System, Proceedings of the 25th Annual Convention of the American Institute of Ultrasound in Medicine, New Orleans, Sept 15-19, 1980.

Brinkley, J.F., Muramatsu, S., McCallum, W.D., "Fetal Weight from Three-Dimensional Ultrasonic Data", submitted to the 26th Annual Convention of the American Institute of Ultrasound in Medicine, San Francisco, 1981.

McCallum, W.D., Brinkley, J.F., "Estimation of Fetal Weight from Ultrasonic Measurements", American Journal of Obstetrics and Gynecology, 133:2, pp.195-200, Jan. 1979.

E. Funding Status

1) "Ultrasonic Measurement of Fetal Volume and Weight"
Principal Investigator: W.D. McCallum, M.D.
Assistant Professor
Department of Obstetrics and Gynecology
Stanford University
Funding agency: National Institute of Child Health and Human Development
Number: 1-R01 HD12327-01
Total term and direct cost: 7/1/79-6/30/81, $111,823
Current funding period: 7/1/79-6/30/80, $60,423

2) "Ultrasonic Three-dimensional Organ Modelling", individual postdoctoral fellowship.
Fellow: James F. Brinkley
Sponsor: W.D. McCallum
Funding Agency: National Institute of General Medical Sciences
Number: 1 F32 GM08092-01
Total term and direct cost: 7/1/81-6/30/84 (3 years) $65,452 (stipend)
Current funding from this fellowship: none

II. INTERACTIONS WITH SUMEX-AIM RESOURCE

A. Collaborations

We are collaborating more with medical people than anyone else. The project is located in the Obstetrics Department at Stanford where W.D. McCallum manages the ultrasound patients. We have also been collaborating with Dr. Richard Popp in the Division of Cardiology at Stanford.

B. Sharing and Interactions with SUMEX Projects

Mostly personal contacts with the Heuristic Programming Project and MYCIN project at Stanford. The message facilities of SUMEX have been especially useful for maintaining these contacts. Since the first phase of the project is now essentially complete we have been interacting more with other SUMEX projects in order to develop the AI ideas.

C. Resource Management

In general SUMEX has been a very usable system, and the staff has been very helpful. The only complaint is that it is impossible to get anything done in the afternoons since we always get bumped.

III. Research Plans

A. Project Goals and Plans

As mentioned in Part I we are implementing this system in phases, each phase requiring use of more sophisticated artificial intelligence techniques. The major phases are as follows (in chronological order:

1. Set up prototype system. Perform engineering tests and clinical evaluations of the ability of the system to predict fetal weight, heart and kidney volume.

As mentioned under Highlights of Research Progress, this first implementation phase is complete, as are the initial engineering evaluations. The clinical study to predict fetal weight should be complete at the end of June when the current grant runs out. Heart studies are just beginning with the Cardiology Division and kidney studies are expected to begin soon with the Department of Radiology, in collaboration with Dr. Barbara Carroll. However, our initial patient studies have demonstrated the basic limitations of the system, which are inadequate models and difficulty of use. From a medical point of view the next phases will be attempts to remove these limitations.

2. Explore other methods for geometric modelling, AI techniques of goal directed problem solving.

In order to develop adequate models and control strategy it will be necessary to examine other AI methods of generating models and using them to guide problem solving programs. This phase of the research is now under way. For this aspect of our research the SUMEX-AIM community has been especially useful.

3. Develop program, as outlined in the introduction, with several limitations:

- Only a simple organ will be modelled at first, ie not the entire fetus including limbs.

- The computer will still request certain scans to be retrieved from the video disk but the operator will outline them with the light pen. Since ultrasound image quality is improving so rapidly it makes sense to wait as long as possible before attempting automated border recognition. The models and control strategies developed during this phase should be useful when actual border recognition is attempted however.

4. Extend the technique to more irregular objects structured models will be developed so that the fetal limbs can be included.

5. Add image processing hardware, develop automated border recognition software.

The models developed in the last two phases will be used to guide the border recognition process. As these phases are implemented they will continue to be tested against the clinical data acquired and stored on floppy disk by the data acquisition system. In this way we can develop new ideas while continually upgrading the clinical utility of the system.

B. Justification for Continued Use of SUMEX

The goals of this project seem to be compatible with the general goals of SUMEX, ie to develop the uses of artificial intelligence in medicine. The problem of three-dimensional modelling is a very general one which is probably at the very heart of our ability to see. By developing a medical imaging system that models the way clinicians approach a patient we should not only develop a useful clinical tool but also explore some very fundamental problems in AI.

C. Need for Resources

1. SUMEX Resources

Our current share of the SUMEX resources is adequate.

2. Other Resources

Judging from our present experience it appears that SUMEX could not handle the amount of data required for image processing on digitized ultrasound scans. This is one of the main reasons we are proposing a distributed system in which SUMEX only directs a smaller machine to do the actual number crunching. It is also one of the reasons we are postponing direct digitization until later. As microprocessors become more powerful they will be capable of acting as slaves to an intelligent SUMEX program. The AI program will direct the image processing functions of the micro so that the data is processed in an intelligent way, but SUMEX will only see the results of that processing, not the actual data. We will thus need to keep track of developments in microcomputers so that we can develop this kind of distributed system.

D. Recommendations

Since we are planning to develop a distributed system we would hope to see these kind of systems being developed by the SUMEX resource. Projects that would be of direct interest are networks (such as ETHERNET), personal computer stations, graphics displays, etc.

II.A.3.3 DECIDER Project: The Psychology of Expert Judgment

The Psychology of Expert Judgment

Eric J. Johnson
Department of Psychology
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

This project focuses on the psychology of expert judgment under uncertainty. There are two separate reasons for pursuing this topic. First, there seems to be a conflict in the view of expertise in two different disciplines. Experts have started to receive much attention in artificial intelligence and cognitive science. The reason for this interest, by and large, is their apparent success in various task domains. Behavioral Decision Theory, in contrast, has called the existence of expert judgment in question, and supplies a number of empirical studies supporting this conclusion. Further exploration of these two conflicting views seems warranted. The current project concentrates on further empirical research and the development of a conceptual analysis of the skills experts bring to judgment tasks.

Second, existing models of novice choice behavior may be successfully extended to account for expert behavior. These models must incorporate domain expertise in order to account for the observed differences between novice and expert subjects in empirical research. Such models can help provide the basis of user interfaces into expert systems, serving as user models aiding in the construction of queries and instructions to the user.

B. Medical Relevance

The initial empirical investigations have focussed on the selection of medical personnel, in particular, the selection of housestaff for post-graduate training. Further extensions are planned using medical diagnosis. Thus, while the study of expert judgment has obvious implications for medicine, this research also uses medicine as its substantive focus.

II. INTERACTION WITH SUMEX-AIM RESOURCES

My connection with the SUMEX resource has encouraged contact with a number of members of the SUMEX and Stanford community. Although my office is located in a building next to the Computer Science Dept., the use of SUMEX for mail has greatly facilitated my interaction with several members of that community.

This interaction has included discussions with members of the MYCIN group and resulted in a talk discussing this work to the Stanford Heuristic Programming Project luncheon meetings.

III. FUTURE PLANS

I am currently preparing a grant proposal based on the initial results of this work. I anticipate that empirical studies will be conducted in the next year. Currently, various medical diagnosis tasks are being considered. The emphasis will be on the relationship between the decision strategies used by the expert judge and that judge's accuracy. I also anticipate that SUMEX will allow access to expert performance systems. One possibility is the comparison of these systems to simple baseline and actuarial models. For these reasons, SUMEX provides a unique resource for this research.

II.A.4 Pilot AIM Projects

The following are descriptions of the informal pilot projects currently using the AIM portion of the SUMEX-AIM resource or the Rutgers-AIM resource pending funding, and full review and authorization.

II.A.4.1 AI-COAG: Coagulation Expert Project

COAGULATION EXPERT Project

Dr. D.A.B. Lindberg
School of Medicine
University of Missouri-Columbia

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Preliminary experiment to form a clinical consultant program based on a formal representation of medical knowledge of the Blood Coagulation Expert.

B. Medical Relevance and Collaboration

Experts in clotting are few and tend to be based at University hospitals or large tertiary care facilities. It would be extremely helpful if this knowledge could be made available to physicians via an automated system.

Relevance of such a proposed system would be with respect to diagnosis, management, and continuing medical education.

The team at UMC consists of the following individuals:

Lamont Gaston, M.D.
Lawrence C. Kingsland III, Ph.D.
Donald A. B. Lindberg, M.D.
Johannes Yesus, M.D.
Anthony Vanker, Ph.D.

Dr. Gaston is a consulting hematologist, Director of the Coagulation Laboratory and Dr. Yesus is Director of the Blood Bank. Dr. Lindberg is Director of the Information Science Group; Dr. Vanker is a Postdoctoral Fellow whose Ph.D. is in Physiology; and Dr. Kingsland is a Postdoctoral Fellow whose Ph.D. is in Electrical Engineering.

Expertise in the field as well as clinical laboratory and patient records are being provided by UMC to build and test the consultant. In addition, testing of the initial program and consultation has been obtained from Heinz Joist, M.D., Director, Coagulation Laboratory, St. Louis University Hospitals.

A formal research proposal was submitted to NIH on October 27, 1980, based on the studies performed on SUMEX and at UMC.

C. Highlights of Research Progress

Accomplishments:

Use of UNITS/AGE: an initial model has been created on SUMEX.

Experimental use of EMYCIN: a feasibility test with a textbook level consultant model has been created on SUMEX.

Use of local LSI-11: In addition the initial knowledge base has been assembled into a simpler (but operational) system on a DEC LSI-11 using RT-11 and BASIC.

We have selected a strategy for development. This is to begin with the interpretation of clinical laboratory tests: first the Coagulation Screen (of 6 tests), then the Partial Coagulation Screen (of 3 tests), then the individual determinations. Laboratory and clinical features will be taken into account, and other features will be added.

Research in Progress:

We have tested the initial model against actual clinical records for 315 patients. This is partly as a validation of the work done, and partly as a means to bring to our attention the unusual circumstances and unforeseen problems which we know will be present. That is, we have allowed for all feasible patterns of results but (probably) have not yet allowed for all the surrounding clinical circumstances. Initial results: 80% of model responses are appropriate as is. 20% are inappropriate as is, because the model does not yet allow for patient clinical history and certain special assays. We are now constructing a clinical history-taking program.

D. List of Relevant Publications

Lindberg, D.A.B., Gaston, L.W., Kingsland, L.C. III, Vanker, A.D., AI/COAG: A Knowledge Based System for Consultation about Human Hemostasis Disorders: Progress Report. Submitted to Fifth Annual Symposium on Computer Applications in Medical Care, November 1-4, 1981, Washington, D.C.

E. Funding Support

This preliminary research phase was supported from two sources.

1. USPHS Grant No. T15 LM 07006, "Training Program in Medical Information Science." Full funding is $161,410/year. About $25,000/year is being devoted to this project.

2. USPHS Grant No. HS 02569, "Health Care Technology Center." Current funding is $500,000/year. About $12,000/year was devoted to this project.

A project research proposal has been submitted jointly to the National Library of Medicine and the Blood Resources Program of the National Institutes of Health.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Disseminations via SUMEX

Dr. Vanker will give an oral presentation of our work at the Spring meeting of Trainees and Directors, NLM Training Programs, in May at Washington, D.C. We have also given individual demonstrations of our models to visiting scientists (including some from Japan) at UMC.

B. Sharing and Interactions with Other SUMEX-AIM Projects

In February, 1980, David Goldman, a medical student at UMC and former pre-doctoral fellow in the Information Science Group, spent a week at Stanford University becoming acquainted with the various artificial intelligence systems in development at SUMEX. In fact, with the help of members of the SUMEX-AIM community, he was able to implement a simple workable coagulation model in EMYCIN.

In March, 1980, Dr. Ueno attended a workshop on AGE at Stanford. Through this workshop he was able to learn a great deal that was directly applicable to our work. He also obtained a better understanding of the UNITS package and how it might be used to interface with AGE. All of us in the study group attended the AIM-tutorial at Stanford in August 1980.

Since the AI systems in which we are interested in are in some stage of implementation on the SUMEX computer and since partial documentation does exist for these systems, we have been able to learn a great deal on our own by an interactive, trial-and-error method. Of course we have had many questions and we have been able to obtain prompt and helpful information from various members of the SUMEX-AIM community via the network electronic message system.

Our decision to make the main line of our development based on the LSI-11 microprocessors is possible only because of the existence of SUMEX and the larger central system as back-up to us should the problem exceed our local capacity.

C. Critique of Resource Management

We have found the people at SUMEX to be uniformly helpful and more than willing to aid us in our attempts to understand the various aspects of artificial intelligence (AI) in medicine. Both Mr. Goldman and Dr. Ueno were delighted with their experiences at Stanford and commented on the willingness of otherwise very busy people to help them with their problems.

One of the drawbacks of SUMEX is that quite often the interaction is slow. There are days when we must wait up to several minutes between exchanges between our terminal and SUMEX. This is apparently due to a high average load on SUMEX at the time. We have had no other problems with the resources at SUMEX and we feel the management has done a good job.

It is clear that truly collaborative research with SUMEX (as opposed to research which merely draws on SUMEX computational resources) requires the designation of a collaborating computer scientist at Stanford. Collaboration of this sort can arise out of personal contact through events such as the AIM Workshop, through sabbaticals, and scientific meetings. In addition, collaboration typically involves capturing the interest of graduate students, and must provide these students with problems which are suitable for dissertation research. These things have not happened as yet in the case of the AI/COAG project but we look forward in the future to this possibility. In any event, this necessary personal component in remote collaboration teaches us that the training, education, and dissemination aspects of resources such as SUMEX are extremely important and cannot be separated from the research elements.

III. RESEARCH PLANS

A. Project Goals and Plans

Plans for the Coming Year:

1. Continue assembling the knowledge data base, with emphasis on documentation of the primary literature sources of the KS's.

2. Continue to study UNITS/AGE and EMYCIN.

3. Continue comparison of the two potentially complex models vs. the microprocessor version.

Long Range Plans:

Our long range plans are to develop the consultation system with Knowledge Sources in the following areas: interpretation of laboratory screening tests, hemostasis patient history, management of oral anticoagulants, diagnosis of hemophilia/von Willebrand's disease, platelet function studies, diagnosis and management of disseminated intravascular coagulation, and pre-operative hemostasis studies.

B. Justification and Requirements for Continued SUMEX Use

As our knowledge base grows, the capabilities of UNITS/AGE will become increasingly more important to us. The UNITS package has built-in means of dealing with large amounts of knowledge in a hierarchical fashion. AGE is a knowledge-based program design to build other knowledge-based programs.

An ancillary, but still important, objective of our work in AI in medicine is to learn about the strengths and weaknesses of the particular AI programming systems in use and development in order to understand better how knowledge can be stored and manipulated. This understanding, in itself important, may then be applied in the design of simpler, but perhaps more accessible programs which can be implemented on micro- or mini-computers.

II.A.4.2 DATA Project

DATA: Data Analysis Tutorial and Advisory System

Asst. Prof. Ruven E. Brooks
Department of Psychiatry and Behavioral Sciences
University of Texas Medical Branch

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

1. The Problem of Biomedical Data Analysis Skills

The lack of sophistication of most biomedical researchers in data analysis is a pervasive constant in the professional functioning of journal editors, funding study groups, and the like. The deficiencies range from the merely regrettable, such as the use of a few too many t-tests or failures to use all applicable analysis techniques, to ones which invalidate entire data analyses, such as failures to substantially meet the statistical assumptions underlying analysis techniques or erroneous interpretation of analysis results. While the source of all of these problems can be covered with the broad label, "inadequate knowledge of statistics," they can, in fact, be traced back to several distinct knowledge deficiencies:

(1) Insufficient knowledge of statistical and probability theory. The type of problems this gives rise to are ones in which valid data have been collected, but the analyses used significantly violate tenets of statistical theory. Typical examples might include performing multiple tests without correction of significance levels ("probability pyramiding") or the use of an improper procedure for estimating missing data.

(2) Faulty data analysis inference/faulty "statistical logic." This phrase is used to describe lack of knowledge about the relationship between questions being asked about the data and the information yielded by analysis techniques. The problems that this gives rise to are ones in which valid data have been collected and statistically acceptable analyses have been performed but in which results from the analysis technique don't match the given scientific question. Typically, these problems manifest themselves as misinterpretations of the results of analyses, for example, using the order of entry of variables into a stepwise regression as an indicator of how strongly the individual variables predict the criterion. Additionally, they can also reveal themselves in inappropriate uses of inferential analyses when exploratory approaches would be more appropriate.

(3) Faulty scientific inference. This is intended to cover situations in which the scientific or experimental questions do

not match the data that was collected, i.e., in which the variables being measured cannot answer the scientific question being asked, regardless of the statistical analysis used. While there are many humorous stories of researchers asking statisticians to make sense out of useless data, there is a fine line dividing errors of this type from legitimate scientific questions as to the correct interpretation to be given to particular measurements.

We assert that deficiencies of the first type can be eliminated by more thorough and wider use of existing tools and techniques, and that deficiencies of the third type must be resolved by better training within individual biomedical disciplines. We will argue, however, that remedying deficiencies of the second type require a new approach.

2. Deficiencies of Existing Approaches

2.1 Improved Graduate Education

Early initial and substantial training in the skills of statistics and data analysis must be a part of any worthwhile graduate program in the biomedical sciences. However, more, better, or longer courses in biomedical statistics do not entirely solve the problem of providing the working biomedical research with data analysis competence in useable form. For most researchers, the design of an experiment and the analysis of the final results together are but a small fraction of their research activities, occupying a relatively small portion of their work time. Hence, their data analysis skills, however good they may have been originally, are certain to have deteriorated significantly by the time they are needed. Moreover, data analysis techniques are continuously evolving, and while this evolution may take place relatively slowly from the statistician's viewpoint, it still may take place rapidly enough to rule out the feasibility of one-time education.

The previous paragraph suggests another possibility, continuing education for researchers, in analogy to continuing education programs for physicians. Indeed, there are short courses and seminars for researchers on both analytic techniques and particular computer packages. However, these courses are not of value to most researchers since they are rarely scheduled to coincide with the completion of a study and the deadline for a final report or journal submission.

2.2 Self-Educational Materials

Given the inadequacy of formal courses, most researchers rely on one of two strategies: The first is to use self-education to recover the skills acquired during graduate education. (This is, of course, presuming that adequate mastery of these skills was obtained in the first place.) Indeed, this can be perfectly adequate for brushing up on one's knowledge of statistical or probability theory, for example, for reviewing the assumptions that must be met to use normal analysis of variance. It is not, however, much use in learning or relearning the process of statistical logic or data analysis inference. This process involves first, selecting a statistical analysis technique that is appropriate both to the properties of the measurements that were made and to the scientific question being asked and, then, interpreting the results of the analysis in terms of this scientific question. Useful tutorial materials in the logic of data analysis must help users relearn or learn this process, rather than just a set of static facts.

Existing materials are not well suited to teaching this process. Consider, first of all, statistics textbooks or sourcebooks. Almost all of these texts are organized by analytic technique or model, not by the structure of the data or by the question to be answered by the data analysis. This organization is not well suited to teaching the process of data analysis inference. Suppose, for example, that the researcher had collected categorical data and wanted to review methods for determining sets of predictors for it; the relevant material would be so scattered through the chapters of most multivariate analysis sourcebooks that the researcher is likely to miss important parts of it.

A form of tutorial material that does not share the drawbacks of the textbook or sourcebook is the case study book which follows the analysis of a complete case from the formulation of the initial data analysis questions or goals through the selection of analytic techniques to the interpretation of final results. While such casebooks could be useful, they are rarely available in the researcher's area of expertise and it is usually difficult to locate the cases that are relevant to the researcher's problem.

Consultants:

Given the inadequacy of sourcebooks and the lack of sufficient case books, most biomedical researchers today use statistical consultants as their source of knowledge for performing data analyses. As convenient as this may be for the researcher who has a good consultant available, it is by no means desirable as a permanent solution to the problem. Such consultants are invariably overburdened and pressed for time. They, therefore, are forced to structure their task into one of learning, as quickly as possible, the minimum knowledge needed to understand the user's research problem and then telling the user what analyses to perform and how to interpret them. Teaching the researcher the rationale behind the analysis selection and interpretation frequently becomes an option to be indulged in if there's any time left over. Thus, interaction with a consultant rarely results in much of an improvement of the biomedical researcher's own skills.

From the consultant's perspective, current practice also leaves much to be desired. After finally learning enough about the research domain to understand the issues, many problems turn out to be, if not routine, at least ordinary enough to present little challenge while others turn out to be hopeless, no matter how much creativity the consultant supplies. Where, ideally, the consultant should play the role of expert confirming the applicability of standard practice or creating a new analysis for an unusual case, in current practice, the consultant's role is frequently just to compensate for deficiencies in the skills a competent researcher should be expected to display.

3. Proposed Research

The primary goal of this research is the construction of a computer system to teach and advise users on how to statistically analyze their data. The intended user population is biomedical researchers who have a rudimentary knowledge of statistical techniques as might be acquired through one or, even, several graduate courses, but, who, because they do data analysis or experimental design infrequently, require tutoring and guidance. It is not intended to provide complete initial training in statistics, though it could be a useful tool in such training.

The basic method of operation of the system would be to guide users through case studies of the logical process of selecting data analysis methods and interpreting their results, based on the nature of the measurements and on the experimental questions. The guidance would take the form of both feedback to specific user decisions and a general question answering capability. This latter capability could be used to generate explanations on topics ranging from the reasons for analysis selection and interpretation in the specific case at hand to background questions about how analyses are performed, what the validity criteria are for interpreting them, etc. The amount and level of detail in these explanations would be controlled by a user model based on the individual user's sequencing and on the details of the case.

The processing of a case will take place in four phases. In the variable description phase, the user will be shown a description of what was measured and asked to characterize it with terms such as metrical versus categorical, maximum-minimum values, etc. In the question clarification phase, the user will be given a scientific question to be answered in the study and asked to recast it in terms of questions about the variables that were actually measured. For the analysis selection phase, the user will be asked to select analyses and will be given feedback in terms of reasons for a selection to the quality of the selections. Finally, in the analysis interpretation phase, the user will be shown the results of an analysis for different sets of sample data, and asked to give a correct interpretation.

The system will have three main operating modes. In the stored case mode, the properties of the variables, the data analysis questions, and the data analysis technique will all be stored in fixed form. In the update mode, users will be allowed to alter the properties of the variables and the analysis techniques and to observe the effects on question formulation, analysis selection, and analysis interpretation. Finally, in new case mode, the user will enter all of the information, starting with the scientific question. While this mode will not permit checking whether the data analysis questions are consistent with the scientific questions, it will still be useful as a tool for permitting users to explore potential experimental designs.

4. Expected Benefits of this Research

The proposed tutorial system would have a number of significant advantages over current approaches. First, because it is organized around the case, rather than around the properties of individual analysis techniques, it will do a better job of conveying the relationships between scientific questions, measurements, analysis techniques, and interpretation. Further, since the presented cases can be tailored to fit the problem the researcher has at hand, there is less likelihood that significant and relevant information will be overlooked. Additionally, since the component that generates the example cases is separate from the one that generates the text to describe them, it is easy to present the cases in the language of the researcher's specialty.

A second advantage of the proposed system is that it could be used by researchers to explore potential experimental designs before actually running an experiment. Though researchers are frequently advised to talk to a statistician before beginning a study, the prospect of attempting to schedule several meetings with a busy consultant is frequently enough to discourage researchers from following this advice. A learning tool, such as the one proposed here should prove conducive to better designing of studies.

These two direct advantages have the potential to combine and produce an indirect benefit: greater productivity on the part of statistical consultants. The argument can be made that researchers should always go to a consultant, regardless of their own level of competence, if only to get a second opinion from someone equally competent. Given the validity of this argument, the more the researcher knows, the more useful the consultation is likely to be with enhanced opportunities to produce the creative or optional analysis, rather than the merely adequate one.

C. Highlights of Research Progress

During this past year, proposals were submitted to the National Science Foundation and to the National Institutes of Health; both were declined. The critiques from both sets of reviews shared a common theme, that the construction of a computer system to give advice on statistical analysis was not a desirable goal, whether or not it was technically feasible. The basis for this position was that any system which increased an investigator's reliance on a computer system, rather than on his or her own intellectual resources, would lead to a decline in research quality. To quote one of the more vituperative NSF reviewers:

> ...If successful, the system purportedly would reduce the need for research assistants who interface with the computer systems, as well as some portions of what statistical consultants do and, I greatly fear, further reduce the amount of thinking that goes into analysis.
>
> ...More importantly, however, the thrust of the proposal is all wrong. It suggests the use of computers in precisely the way they are most inappropriate, that is, trying to substitute computer generated rules for human thought. While a portion of the program may be of some use, the main thrust of the proposal is to substitute advice based on a set of computer generated questions and

> answers for the investigator's own knowledge. Just as the existence of the computer software packages have enabled people who don't understand statistical techniques to use them by reducing the cost of computation, so, I think, this program would give people a false sense of security and lead them to think they were using correctly techniques they do not understand. The effect of this program, if it were successful, would be to further reduce the thought that goes into designing and analyzing data, and make research even more mindless than it all too frequently is now.

Since the nature of this criticism left little hope of ever obtaining funding for the project as originally conceived, a major re-orientation of the project was undertaken, shifting from a purely advisory system towards a tutorial system. This shift has had a number of consequences, the most trivial of which being the development of a new acronym (DATA system = Data Analysis Tutorial and Advisory system). More important, it has caused a shift in the intellectual focus of the project from being primarily one of the application of existing software to a new domain, to one that touches on two issues of more fundamental A.I. interests, methods for providing explanation and methods for problem generation in intelligent CAI.

D. List of Relevant Publications

None.

E. Funding Support

None.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations via SUMEX

B. Sharing and Interaction with Other SUMEX-AIM Projects

In August, 1980, I attended the AIM workshop at Stanford. While in Palo Alto, I visited Tom Moran at Xerox PARC who gave me a manual for the Interactive Data Language, an Interlisp based statistical package developed by Ron Kaplan and Beau Shiel. Since this system has useful potential for both the DATA project and the RX project, I inquired whether it was available at SUMEX and learned that, somewhat surprisingly, SUMEX projects were unaware of its existence. It has since been transferred down the hill and is now being used by the RX project and will be available for use in the DATA project.

C. Critique of Resource Management

As this project has not involved significant utilization of machine resources over the past year, no comments in this area are appropriate.

III. Research Plans

A. Project Goals and Plans

The major short-term goal for the next year will be to complete the first-level software design for the entire system and to begin construction of the problem generation component. This component, which generates case descriptions of statistical analysis problems, is, by itself, an operational system since the case descriptions can be evaluated and used independently of the rest of the system.

B. Justification for Continued SUMEX Use

SUMEX continues to be the only nationally available resource that is oriented towards support of the remote (as versus on-site) artificial intelligence in medicine researcher. While I am located in the nation's sixth largest medical school, I am not located near any suitable research facilities; the artificial intelligence group at Austin is more than 200 miles away, and my local environment has neither the hardware nor the software support for any kind of artificial intelligence work nor is there any possibility of acquiring any before 1983.

C. Needs and Plans for Other Resources

The Medical Computer Science Dept. at the University of Texas Health Science Center in Dallas is exploring the possibility of establishing a facility for A.I. research. Should their efforts prove fruitful, it would provide a feasible alternative to SUMEX, though not for at least 2-3 years.

D. Recommendations for Future Resource Development

If the efforts with single-user systems prove fruitful at Stanford, it will be important to develop plans for software licensing and support at other sites. Given that the true annual cost of an on-site programmer to maintain a personal computer is a substantial share of the one-time cost of the machine, there will still be considerable need for SUMEX to provide support even if it becomes possible for each research group to have their own hardware.

II.A.4.3 EXCHANGE Project

Japan/U.S. Medical Consultation System Exchange

Drs. S. Kaihara, T. Koyama,
and F. Mizoguchi
Hospital Computer Center
University of Tokyo Hospital

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Our goal is two-fold, namely the international comparison of medical knowledge formalized in the consultation systems in some clinical areas, and secondly, the trial of cooperative development of computer programs in the international settings. Both objectives can be achieved only by using such facilities as SUMEX-AIM which can be accessed internationally by the collaborating researchers.

B. Medical Relevance and Collaboration

Due to rapid increase of medical knowledge required for daily practice of medicine, the interest in medical consultation systems is increasing. Our project involves two most prominent expert physicians, Prof. Shigenobu Nagataki, Professor of Medicine at Nagasaki University and Prof. Yoshiyuki Suzuki, Professor of Pediatrics at University of Tokyo, who are both enthusiastic for the development of consultation systems in their field.

C. Highlights of Research Progress

Three projects are in progress:

1) Comparison of medical expert knowledge formalized in the system at international settings. Two systems are now being developed using the EXPERT system developed at Rutgers University. The first one is the consultation system of thyroid diseases led by Prof. Nagataki. The core of the system has been completed and the knowledge will be compared with the American expert knowledge in the near future.

2) Development of practical consultation system using "AGE". Our experience showed that the diagnosis obtained by deduction covers only a small part of physicians' interests. They often express a need for such functions as data error checking, strategies for obtaining data, evaluation of missing data on the diagnosis, reference data related to the diagnosis, etc. A consultation system which has the above features is now being developed, using "AGE" which was developed at Stanford University. Preliminary results in the consultation of inborn errors of metabolism showed encouraging results.

3) Development of MECS-AI. A general purpose system named as MECS-AI is under development. This system is designed to assist in constructing knowledge-based inference systems with time-oriented data treatments.

D. List of Relevant Publications

None yet.

E. Funding Support

i) The Japan-US Cooperative Science Program (grant)
Shigekoto Kaihara, Hospital Computer Center, University of Tokyo
The Japan Society for the Promotion of Science
(no number)
March 1979 to February 28, 1981 5,280,000 Japanese yen

ii) Grant in Aid for Scientific Research (grant)
Shigekoto Kaihara, Hospital Computer Center, University of Tokyo
The Japan Society for the Promotion of Science
549009
April 1980 to March 1983 8,736,000 Japanese yen

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaboration and Program Dissemination via SUMEX

Dr. Lindberg's group at University of Missouri is collaborating with us for the evaluation of consultation programs.

Dr. Kulikowski's group at Rutgers University is also collaborating by supplying EXPERT and evaluating developed systems.

B. Sharing and Interactions with Other SUMEX-AIM Projects

We are collaborating with the "AGE" developing project at Stanford University, which supplies us with the newest version of AGE. Our group also attended the AIM workshop held at Stanford in August, 1980.

C. Critique of Resource Management

We have found the staff professional and helpful. We have had almost no troubles in using SUMEX facilities from Tokyo, Japan.

III. RESEARCH PLANS

A. Project Goals and Plans

Project goals are two, namely 1) international comparison of medical knowledge formalized in the consultation system in clinical fields; 2) international collaboration for development of more practical medical consultation systems using AGE or MECS-AI. The procedures described in Section I.C. will be continued.

B. Needs and Plans for Other Computing Resources beyond SUMEX-AIM

We are using computer facilities located in Japan whenever possible to reduce the cost of international data communication. This will be continued.

C. Recommendations for Future Community and Resource Development

Development of systems which can be transferred to smaller machines such as VAX is desirable in our environment, for there is only one DEC 20 in the large city of Tokyo.

II.A.4.4 MELANOMA Project

Decision Support for Clinical Melanoma

Drs. M.S. Blois and Hyo Kim
Section on Medical Information Science
University of California Medical Center
University of California at San Francisco

We have not made significant progress on our project of "Decision Support for Clinic Melanoma". Due to the low priority we have and the large size of EMYCIN program we had only limited access to the computer. However we have become familiarized with the system and we hope that we will be able to make slow but steady progress in future.

II.B Books, Papers, and Abstracts

The publications for the various collaborative projects are summarized in their respective progress reports.

II.C Resource Summary Table

Detailed resource usage information is summarized starting on page 36. Tabulation of this information on the NIH-provided form is yet to be completed.

Appendix A

Community Growth and Project Synopses

This appendix contains a graphical display of the development of the SUMEX-AIM community over the years and brief synopses of currently active projects. Figure 15 below illustrates the substantial growth in the cumulative number of projects in the Stanford, national SUMEX, and Rutgers-AIM communities since the resource began operation in 1974.

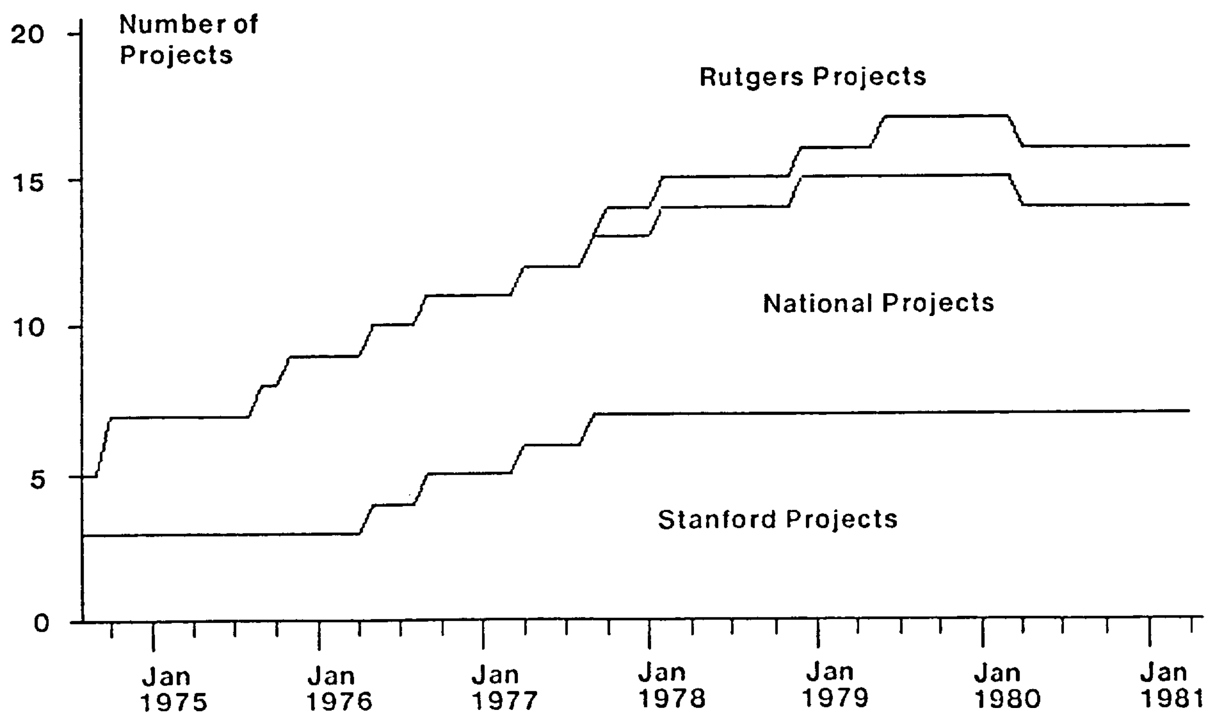

Figure 15. SUMEX-AIM Growth by Community

National AIM Project: ACQUISITION OF COGNITIVE PROCEDURES (ACT)

Principal Investigator: John R. Anderson, Ph.D.
Department of Psychology
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213
(412) 578-2788 (ANDERSON@SUMEX-AIM)

The ACT Project combines a semantic network data-base with a production system to simulate human cognition. Prominent among the reasons for using a production system architecture as a framework for developing such a program is the possibility of modeling learning as the acquisition of new productions. ACT possesses a number of learning mechanisms which have been used to model the learning of procedural skills such as language comprehension and geometry theorem proving. Some of these mechanisms have the effect of either extending or restricting the set of circumstances in which a particular behavior is performed so as to produce better performance. Others have the effect of speeding up cognitive operations by compressing the effects of a series of production applications into the application of a single production. Out of this set of productions ACT applies those that usually result in desirable outcomes. In this way it is able to model the human ability to learn even when given unreliable feedback. Another feature of ACT that reflects its psychological orientation is its willingness to model human limitations. Here the hope is that by being faithful to the human mind even in its failings, it eventually may be possible to emulate its successes.

SOFTWARE AVAILABLE ON SUMEX

The ACT production system is available to GUEST users of SUMEX.

REFERENCES

Anderson, J.R.: Language, Memory, and Thought. Lawrence Erlbaum Associates, Hillsdale, N.J., 1976.

Anderson, J.R., Kline, P.J. and Lewis, C.H.: A production system model of language processing. IN M.A. Just and P.A. Carpenter (eds.), Cognitive Processes in Comprehension. Lawrence Erlbaum Associates, Hillsdale, N.J., 1977.

Anderson, J.R. and Kline, P.J.: A learning system and its psychological implications. Proc. Sixth IJCAI, Tokyo, August, 1979.

National AIM Project: CADUCEUS (previously INTERNIST)

Principal Investigators: Jack D. Myers, M.D. (MYERS@SUMEX-AIM)
Harry E. Pople, Ph.D. (POPLE@SUMEX-AIM)
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
Dr. Pople: (412) 624-3490

The major goal of the CADUCEUS Project is to produce a reliable and adequately complete diagnostic consultative program in the field of internal medicine. Although this program is intended primarily to aid skilled internists in complicated medical problems, the program may have spin-off as a diagnostic and triage aid to physicians' assistants, rural health clinics, military medicine and space travel. In the design of CADUCEUS and its predecessor INTERNIST I, we have attempted to model the creative, problem-formulation aspect of the clinical reasoning process. The program employs a novel heuristic procedure that composes differential diagnoses, dynamically, on the basis of clinical evidence. During the course of a CADUCEUS or INTERNIST I consultation, it is not uncommon for a number of such conjectured problem foci to be proposed and investigated, with occasional major shifts taking place in the program's conceptualization of the task at hand.

SOFTWARE AVAILABLE ON SUMEX

Versions of INTERNIST are available for experimental use, but the project continues to be oriented primarily towards research and development; hence, a stable production version of the system is not yet available for general use.

REFERENCES

Pople, H.E., Myers, J.D. and Miller, R.A.: The DIALOG model of diagnostic logic and its use in internal medicine. Proc. Fourth IJCAI, Tbilisi, USSR, September, 1975.

Pople, H.E.: The formation of composite hypotheses in diagnostic problem solving: An exercise in synthetic reasoning. Proc. Fifth IJCAI, Boston, August, 1977.

National AIM Project: HIERARCHICAL MODELS OF HUMAN COGNITION

Principal Investigators: Walter Kintsch, Ph.D. (KINTSCH@SUMEX-AIM)
Peter G. Polson, Ph.D. (POLSON@SUMEX-AIM)
Computer Laboratory for Instruction
in Psychological Research (CLIPR)
Department of Psychology
University of Colorado
Boulder, Colorado 80302
(303) 492-6991

Contact: Dr. James Miller (JMILLER@SUMEX-AIM)

The CLIPR Project is concerned with the modeling of complex psychological processes. It is comprised of two research groups. The prose comprehension group has completed a project that carries out the microstructure text analysis described by Miller and Kintsch (1980), yielding predictions of the recall and readability of that text by human subjects. More recently, this group has been interacting with the Heuristic Programming Project at Stanford, using the AGE and UNITS packages to build a more complex model of the knowledge-based processes characteristic of prose comprehension. The planning group is working toward a model of the planning processes used by expert computer software designers. The initial development of this model requires the detailed analysis of expert software design protocols for subsequent simulation.

SOFTWARE AVAILABLE ON SUMEX

A set of programs has been developed to perform the microstructure text analysis described in Kintsch and van Dijk (Psychological Review, 1978) and Miller and Kintsch (1980). The program accepts a propositionalized text as input, and produces indices that can be used to estimate the text's recall and readability. A more complex model based in AGE and UNITS, which emphasizes the knowledge-based aspects of comprehension, is currently under development.

REFERENCES

Atwood, M.E., Polson, P.G., Jeffries, R. and Ramsey, H.R.: Planning as a process of synthesis. Technical Report SAI-78-144-DEN. Science Applications, Inc., Denver, Colorado, December, 1978.

Jeffries, R., Turner, A.A., Polson, P.G., and Atwood, M.A.: The processes involved in designing software. In J.R. Anderson (Ed.), Cognitive skills and their acquisition. Hillsdale, NJ: Erlbaum, 1981 (forthcoming).

Kintsch, W.: On modeling comprehension. Educ. Psychologist, 14:3-14, 1979.

Miller, J.R. and Kintsch, W.: Readability and recall of short prose passages: A theoretical analysis. J. Experimental Psychology: Human Learning and Memory, 1980. (In press)

National AIM Project: BIOMEDICAL KNOWLEDGE ENGINEERING IN CLINICAL MEDICINE (PUFF-VM)

Principal Investigators: John J. Osborn, M.D.
The Institutes of Medical Sciences
Pacific Medical Center
San Francisco, California 94115
(415) 567-0900 (OSBORN@SUMEX-AIM)

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

The PUFF-VM Project has produced two knowledge-based programs for the interpretation of physiologic measurements made in clinical medicine. The interpretations are intended to aid in diagnostic decision-making and in selecting therapeutic actions. The programs are: PUFF--the evaluation of pulmonary function laboratory data, and VM--the evaluation and management of respiratory status for patients in the intensive care unit.

The task of the PUFF PROGRAM is to interpret standard measures of pulmonary function. In the laboratory at the Pacific Medical Center (PMC), about 50 parameters are calculated from measurement of lung volumes, flow rates, and diffusion capacity. In addition to these measurements, patient history and referral diagnosis also are used to interpret the test results. PUFF produces a report for the patient record, explaining the clinical significance of measured test results. It also provides a diagnosis of the presence and severity of pulmonary disease. The interpretation process is accomplished by examination of expert knowledge represented by a set of production rules. Each rule relates physiologic measurements or states to a conclusion about the physiologic significance of the measurement or state. A version of the PUFF program is used daily at the PMC.

The VENTILATOR MANAGER (VM) PROGRAM is designed to interpret on-line physiologic data in the intensive care unit (ICU). These data are used to manage post-surgical patients receiving mechanical assistance in breathing. VM is an extension of a physiologic monitoring system, and is designed to perform 5 specialized tasks in the ICU: 1) to detect possible measurement errors; 2) to recognize untoward events in the patient/machine system and suggest corrective action; 3) to summarize the patient's physiologic status; 4) to suggest adjustments to therapy based on the patient's status over time, and long-term therapeutic goals; and 5) to maintain a set of patient-specific expectations and goals for future evaluation by the program. The program produces interpretations of the physiologic measurements over time, using a model of the therapeutic procedures in the ICU and clinical knowledge about the diagnostic implications of the data.

SOFTWARE AVAILABLE ON SUMEX

The PUFF and VM programs will be available to GUEST users for use on pre-existing (non-identifiable) cases. No packages currently exist for program development.

REFERENCES

Fagan, L.M., Kunz, J.C., Feigenbaum, E.A. and Osborn, J.J.: A symbolic processing approach to measurement interpretation in the intensive care unit. Proc. Third Annual Symposium Computer Applications in Medical Care, Silver Spring, Maryland, October, 1979, pp. 30-33.

Fagan, L.M., Shortliffe, E.H. and Buchanan, B.G.: Computer-based medical decision making: From MYCIN to VM. Automedica 3(2), 1980.

Kunz, J.C., Fallat, R.J., McClung, D.H., et al: A physiological rule based system for interpreting pulmonary function test results. Heuristic Programming Project Report HPP-78-164, Computer Science Dept., Stanford Univ., November, 1978.

Osborn, J.J., Fagan, L.M., Fallat, R.J., et al: Managing the data from respiratory measurements. Med. Instrumentation, November-December, 1979.

Rutgers AIM Project: RUTGERS RESEARCH RESOURCE - COMPUTERS IN BIOMEDICINE

Principal Investigator: Saul Amarel, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546 (AMAREL@RUTGERS)

The broad objective of the Resource is to apply advanced methods in computer science, particularly in artificial intelligence (AI), to biomedical problems. The Resource has three major areas of study: 1) Medical Modeling and Decision Making in several medical domains with emphasis on collaborative development of consultation systems in rheumatology and ophthalmology; 2) Modeling of Belief Systems and Commonsense Reasoning with emphasis on the psychology of plan recognition and handling of stereotypes; and 3) Artificial Intelligence studies with emphasis on Representations, Interpretation processes, and problems of knowledge and expertise acquisition. The studies in Medical Modeling and Decision Making are performed jointly by computer and medical scientists at Rutgers and elsewhere in the Country and abroad.

The Resource also sponsors national Artificial Intelligence in Medicine (AIM) Workshops for the AIM community.

SOFTWARE AVAILABLE ON SUMEX

CASNET--System for consultation in the diagnosis and treatment of glaucoma.

Documentation available:

"A Model-Based Method for Computer-Aided Medical Decision Making" Weiss, S., Kulikowski, C., Amarel, S., Safir, A., (1978) pp. 145-172, AI Journal, North Holland Press, 1978.

EXPERT--System for designing and applying consultation models using a relatively simple language to describe the models. It has been used in a variety of medical and non-medical applications (mainly rheumatology, ophthalmology and endocrinology).

Documentation available:

"A Guide to the Use of the EXPERT Consultation System" Weiss, S., Kern, K., and Kulikowski, C., CBM-TR-94, 11/78.

"EXPERT: A System for Developing Consultation Models," Proceedings IJCAI, Tokyo, August 1979, pp. 942-947, (also Rutgers Computer Science Report CBM-TR-97).

Anyone interested in using CASNET or EXPERT may contact either Kulikowski@Rutgers or Weiss@Rutgers.

REFERENCES

Amarel, S.: Computer-based interpretation and modeling in medicine and psychology: The Rutgers Research Resource. IN Siler and Lindberg (eds.), Computers in Life Science Research. Foseb and Planum, 1975.

Schmidt, C.F., Sridharan, N.S., and Goodson, J.L.: The plan recognition problem: An intersection of psychology and artificial intelligence. AI Journal 11(1,2), August, 1978 (special issue on applications to the sciences and medicine).

Weiss, S., Kulikowski, C.A., Amarel, S. and Safir, A.: A model-based method for computer-aided medical decision-making. AI Journal 11(1,2), August, 1978 (special issue on applications to the sciences and medicine).

National AIM Project: SIMULATION OF COGNITIVE PROCESSES (SCP)

Principal Investigators: James G. Greeno, Ph.D. (GREENO@SUMEX-AIM)
Alan M. Lesgold, Ph.D. (LESGOLD@SUMEX-AIM)
Learning Research and Development Center
University of Pittsburgh
Pittsburgh, Pennsylvania 15260
Dr. Lesgold: (412) 624-4901

The general purpose of the SCP Project is to develop increased understanding of normal and deficient cognitive functions, especially in reading and mathematics. Earlier work included simulations of interactive processes of grapheme-phoneme decoding and word recognition, and of semantic processes in comprehension of quantitative information in arithmetic word problems. The main emphasis at this time is on a collaboration with John Anderson, using the ACTF system to explore mechanisms of learning in the domain of geometry proofs. The SCP part of this work includes development of a system that learns by reading example proofs. The goal is to identify conceptual structures that are required for a learner to acquire planning strategies.

SOFTWARE AVAILABLE ON SUMEX

Programs are in a developmental stage and not yet available for use.

REFERENCES

Greeno, J.G., Magone, M.E. and Chaiklin, S.: Theory of constructions and set in problem solving. IN Memory and Cognition (In press). (Also available as Technical Report 1979/9, Learning Research and Development Center, Univ. Pittsburgh.)

Greeno, J.G.: Preliminary steps toward a cognitive model of learning primary mathematics. IN K. Fuson and W. Geeslin (eds.), Models of Children's Mathematical Learning, ERIC Information Center. (In press)

Lesgold, A.M. and Curtis, M.E.: Learning to read words efficiently. IN A. Lesgold and C. Perfetti (eds.), Interactive Processes in Reading, Erlbaum, Hillsdale, N.J. (In progress)

National AIM Project: SIMULATION AND EVALUATION OF CHEMICAL SYNTHESIS (SECS)

Principal Investigator: W. Todd Wipke, Ph.D.
Department of Chemistry
University of California at Santa Cruz
Santa Cruz, California 95064
(408) 429-2397 (WIPKE@SUMEX-AIM)

The SECS Project aims at developing practical computer programs to assist investigators in designing syntheses of complex organic molecules of biological interest. Key features of this research include the use of computer graphics to allow chemist and computer to work efficiently as a team, the development of knowledge bases of chemical reactions, and the formation of plans to reduce the search for solutions. SECS is being used by the pharmaceutical industry for designing syntheses of drugs.

A spin-off project, XENO, is aimed at predicting the plausible metabolites of foreign compounds for carcinogenicity studies. First, the metabolism is simulated; then the metabolites are evaluated for possible carcinogenicity.

SOFTWARE AVAILABLE ON SUMEX

SECS-- Available with a reaction library of over 400 reactions. The user needs a TTY or a DEC GT40 type graphics terminal.

XENO-- (for prediction of metabolites of xenobiotic compounds) is available for preliminary exploration since the project is still in the early development stages.

PRXBLD--(for building approximate molecular models from two-dimensional molecular models) is an energy minimization approach which is available both stand-alone and included within SECS.

REFERENCES

Spann, M.L., Chu, K.C., Wipke, W.T. and Ouchi, G.: Use of computerized methods to predict metabolic pathways and metabolites. J. Env. Pathology and Toxicology 2:123, 1978.

Wipke, W.T., Smith, G., Choplin, F. and Sieber, W.: SECS--Simulation and Evaluation of Chemical Synthesis: Strategy and planning. IN Computer-assisted Organic Synthesis Planning. ACS Symposium Series, 1977, pp. 97-127.

Wipke, W.T., Ouchi, G. and Krishnan, S.: Simulation and Evaluation of Chemical Synthesis--SECS. An application of artificial intelligence techniques. Artificial Intelligence 10:999, 1978.

National AIM Project: PROBLEM SOLVING EXPERTISE

Principal Investigators:

Paul E. Johnson, Ph. D.
Center for Research in Human Learning
205 Elliott Hall
University of Minnesota
Minneapolis, Minnesota 55455
(612) 373-5302 (PJOHNSON@SUMEX-AIM)

William B. Thompson, Ph. D.
Department of Computer Science
136 Lind Hall
University of Minnesota
Minneapolis, Minnesota, 55455
(612) 373-0132 (THOMPSON@SUMEX-AIM)

The Minnesota SOLVER project focuses upon the development of strategies for discovering and representing the knowledge and skill of expert problem solvers. Although in the last fifteen years great progress has been made in synthesizing the expertise required for solving complex problems, most expert systems embody only the limited amount of expertise that individuals are able to report in a particular constrained language (e.g. production rules). What is still lacking is a theoretical framework capable of reducing dependence upon the expert's intuition or on the near exhaustive testing of possible organizations. Our methodology consists of: (1) extensive use of verbal thinking aloud protocols as a source of information from which to make verbal inferences about underlying cognitive structure and process; (2) development of computer models as a means of testing the adequacy of inferences derived from protocol studies; (3) testing and refinement of the cognitive models based upon the study of human and model performance in experimental settings. We are currently investigating problem solving expertise in domains of medicine, science, and law.

SOFTWARE AVAILABLE ON SUMEX

A version of the DIAGNOSER simulation model is being implemented on SUMEX; it should be available sometime this year.

REFERENCES

Johnson, P.E., Barreto, A., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P., & Swanson, D. Expertise and error in diagnostic reasoning. Cognitive Science (in press). Johnson, P.E. Cognitive Models of Medical Problem Solvers. In D.C. Connelly, E. Benson, & D. Burke (eds.), Clinical decision making and laboratory use. University of Minnesota Press (in press). Johnson, P.E., & Thompson, W.B. Strolling down the garden path: Error prone tasks in expert problem solving. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, August, 1981.

Stanford Project: GENERALIZATION OF AI TOOLS (AGE)

Principal Investigators: H. Penny Nii (NII@SUMEX-AIM)
Edward A. Feigenbaum, Ph.D.
Department of Computer Science
Stanford University
Stanford, California 94305
H.P. Nii: (415) 497-2739

The long-range objective of AGE, a SUMEX CORE RESEARCH Project, is to build a software laboratory for building knowledge-based, application programs. It is an attempt to define and accumulate knowledge-engineering tools, with rules to guide in the use of these tools. The design and implementation of the AGE program will be based primarily on the experiences gained in building knowledge-based programs by the Stanford Heuristic Programming Project in the last decade (The programs that have been or are being built are: DENDRAL, META-DENDRAL, MYCIN, HASP, AM, MOLGEN, GUIDON, CRYSALIS, PUFF, VM and SACON.). The initial AGE program contains a collection of tools suitable for constructing user programs based on the Blackboard paradigm (used in HASP and CRYSALIS) and Backward-chained production rules (used in MYCIN). In addition, AGE has facilities to aid the user in the construction, debugging, and running of his program.

SOFTWARE AVAILABLE ON SUMEX

AGE-1 is available on an experimental basis to a limited number of users. A public version of the programs, together with reference manuals and user guides, is planned for July, 1980.

REFERENCES

Nii, H.P. and Feigenbaum, E.A.: Rule-based understanding of signals. IN D.A. Waterman and F. Hayes-Roth (eds.), Pattern-directed Inference Systems. Academic Press, 1978, pp. 483-501.

Nii, H.P. and Aiello, N.: AGE: A knowledge-based program for building knowledge-based programs. Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 645-655.

Aiello, N., Bock, C., Nii, H.P. and White, W.: Joy of AGE-ing: An Introduction to the Use of AGE-1 System, HPP Technical Memo, July, 1980.

Stanford Project: HANDBOOK OF ARTIFICIAL INTELLIGENCE

Principal Investigator: Edward A. Feigenbaum, Ph.D.
Avron Barr, Research Associate (BARR@SUMEX-AIM)
Department of Computer Science
Stanford University
Stanford, California 94305

The AI Handbook Project is a part of SUMEX CORE RESEARCH aimed at making the important results of AI research accessible to the large, multidisciplinary community of scientists who want to build AI systems in their own problem areas. Students and researchers at Stanford and other AI laboratories have prepared over 300 short articles describing the fundamental ideas, useful techniques, and exemplary programs developed in the field over the last 20 years. These articles have been written for computer-literate scientists and engineers in other fields who are unfamiliar with AI research and jargon. The Handbook will provide a scientist who, for instance, might want to know what a "heuristic" is or how to build a "natural language" front end, with information about all of the relevant AI techniques and existing systems, as well as abundant pointers into the field's literature.

The Handbook is being published in report and book form. It also will be made available to the SUMEX community via an on-line information retrieval system. Following is a TOPIC OUTLINE for Volumes I and II: HANDBOOK OF ARTIFICIAL INTELLIGENCE

INTRODUCTION: The Handbook of Artificial Intelligence; Overview of AI Research; History of AI; An Introduction to the AI Literature

SEARCH: Overview; Problem Representation; Search Methods for State Spaces, AND/OR Graphs, and Game Trees; Six Important Search Programs

REPRESENTATION OF KNOWLEDGE: Issues and Problems in Representation Theory; Survey of Representation Techniques; Seven Important Representation Schemes;

AI PROGRAMMING LANGUAGES: Historical Overview of AI Programming Languages; Comparison of Data Structures and Control Mechanisms in AI Languages; LISP

NATURAL LANGUAGE UNDERSTANDING: Overview - History and Issues; Grammars; Parsing Techniques; Text Generation Systems; Machine Translation; The Early NL Systems; Six Important Natural Language Processing Systems

SPEECH UNDERSTANDING SYSTEMS: Overview - History and Design Issues; Seven Major Speech Understanding Projects

APPLICATIONS-ORIENTED AI RESEARCH--SCIENCE AND MATHEMATICS: Overview; TEIRESIAS - Issues in Expert Systems Design; Research on AI Applications in Mathematics (MACSYMA and AM); Research on AI Applications in Chemistry (DENDRAL, CRYSALIS, etc.); Other Scientific Applications Research

APPLICATIONS-ORIENTED AI RESEARCH--MEDICINE: Overview of Medical Applications Research; Six Important Medical Systems

APPLICATIONS-ORIENTED AI RESEARCH--EDUCATION: Historical Overview of AI Research in Educational Applications; Issues and Components of Intelligent CAI Systems; Seven Important ICAI Systems

AUTOMATIC PROGRAMMING: Overview; Techniques for Program Specification; Approaches to AP; Eight Important AP Systems

The following sections of the Handbook are still in preparation and will appear in Volume III: Theorem Proving; Vision; Robotics; Information Processing Psychology; Learning and Inductive Inference; Planning and Related Problem-solving Techniques.

Stanford Project: DENDRAL--RESOURCE RELATED RESEARCH - COMPUTERS IN CHEMISTRY

Principal Investigator: Carl Djerassi, Ph.D.
Department of Chemistry
Stanford University
Stanford, California 94305

Contact: Dr. Dennis SMITH@SUMEX-AIM
(415) 497-3144

The DENDRAL Project involves research in computer-assisted structure elucidation of unknown organic compounds of biological importance. This research has three major components: 1) program development; 2) biochemical applications; and 3) resource-sharing.

Recent program developments have been directed toward building more powerful interactive programs to assist chemists in the three major areas of structure elucidation: analysis of data to yield substructural information about an unknown ("planning"), advanced methods for assembly of substructures into complete structures ("structure generation"), and the prediction of data for structural candidates to rank-order the candidates by comparison of predicted and observed data ("testing"). Important problems of structure representation have been solved which have enabled dealing with stereochemical (three-dimensional) aspects of structure throughout the procedures.

Major areas of application of the programs in the research of this group and other collaborative projects include: a) marine natural products, particularly marine steroids and halogenated compounds which display biological activity; b) antibiotics and other derivatives of known or potential drugs; c) terpene alkaloids; d) photoproducts related to vitamin A; and e) conformational studies of narcotic analogs and polypeptides.

These programs are shared among a community of collaborators and guest users at SUMEX, with communication via computer network from a variety of sites in the U.S., Europe and Australia. Exportable versions of some programs are maintained. These versions have been installed successfully in more than 15 research laboratories throughout the world.

SOFTWARE AVAILABLE ON SUMEX

CONGEN--An interactive program for structure generation to yield candidate structures for an unknown based on inferred substructural components (exportable).

GENOA--An advanced structure generator capable of handling overlapping substructural information; uses CONGEN as a core component (exportable).

Meta-DENDRAL--An INTSUM, RULEGEN and RULEMOD sequence for automatic rule formation to relate observed data to substructures in mass spectrometry and carbon magnetic resonance spectroscopy.

REACT--A program for carrying out a complex sequence of chemical reactions and exploration of the consequences of those reactions.

NMR--For substructural inference and spectrum prediction in carbon magnetic resonance spectroscopy (will be exportable).

REFERENCES

Carhart, R.E., Varkony, T.H. and Smith, D.H.: Computer assistance for the structural chemist, IN D.H. Smith (ed.), Computer-Assisted Structure Elucidation. American Chemical Society, Washington, D.C., 1977, p. 126.

Djerassi, C., Smith, D.H. and Varkony, T.: A novel role of computers in the natural products field. Naturwiss. 66:9, 1979.

Nourse, J.G., Carhart, R.E., Smith, D.H. and Djerassi, C.: Exhaustive generation of stereoisomers for structure elucidation. J. Am. Chem. Soc. 101:1216, 1979.

Stanford Project: EXPEX -- THE EXPERT EXPLANATION PROJECT

Principal Investigator: Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University
Stanford, California 94305
(415) 497-6970 (SHORTLIFFE@SUMEX-AIM)

EXPEX is a new Stanford project that joined the AIM community in the Spring of 1981. The major thrust of the research is the development of new representation schemes to facilitate knowledge acquisition and explanation. This includes not only the study of fundamental representational formalisms but also the encoding of various types of knowledge, such as causal information and user models. The explanation portion of the research effort is dealing with a medical domain (endocrinology) and is being undertaken on SUMEX, whereas the knowledge acquisition portion deals with nonmedical topics and uses other computing resources at Stanford.

Our interest in explanation derives from the insights we gained in developing explanatory capabilities for the MYCIN system. In that system and its descendents, we were able to generate intelligible explanations by taking advantage of a rule-based representation scheme. The limitations of the justifications generated using MYCIN's explanation capabilities have become increasingly obvious, however, and have led to improved characterization of the kinds of explanation capabilities that must be developed if clinical consultation systems are to be accepted by physicians. EXPEX is devoted to the development of new theoretical insights into this problem.

REFERENCE

Wallis, J.W. and Shortliffe, E.H. Explanatory power for expert systems: studies in the representation of causal relationships for medical consultations. Internal technical report, Heuristic Programming Project, May 1981.

Stanford Project: MOLGEN--AN EXPERIMENT PLANNING SYSTEM FOR MOLECULAR GENETICS

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

Laurence H. Kedes, M.D. (KEDES@SUMEX-AIM)
Department of Medicine
Stanford University
Stanford, California 94305
(415) 497-5897

Contact: Dr. Peter FRIEDLAND@SUMEX-AIM
(415) 497-1740

The goal of the MOLGEN Project is to apply the techniques of artificial intelligence to the domain of molecular biology with the aim of providing assistance to the experimental scientist. The most substantial problem under consideration is the task of experiment design. Two major approaches to this problem have been explored, one which instantiates abstracted experimental strategies with specific laboratory tools, and one which creates plans *in toto*, heavily influenced by the role played by interactions between plan steps. As part of the effort to build an experiment design system, a knowledge representation and acquisition package--the UNITS System, has been constructed. A large knowledge base, containing information about nucleic acid structures, laboratory techniques, and experiment-design strategies, has been developed using this tool. Smaller systems, such as programs which analyze primary sequence data for homologies and symmetries, have been built when needed.

SOFTWARE AVAILABLE ON SUMEX

Knowledge-based Experiment Design system (Friedland).

Meta-planning with Constraints experiment design system (Stefik).

UNITS system for knowledge representation and acquisition.

Interactive KORN Program for DNA sequence analysis.

GA1 program for restriction map construction.

SAFE program for gene excision.

REFERENCES

Friedland, P.E.: Knowledge-based experiment design in molecular genetics (Ph.D. thesis). Computer Science Dept. Report, CS-79-771, Stanford Univ.

Friedland, P.E.: Knowledge-based experiment design in molecular genetics. Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 285-287.

Stefik, M.J.: An examination of a frame-structured representation system. Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 845-852.

Stefik, M.J.: Planning with constraints (Ph.D. thesis). Computer Science Dept. Report, Stanford Univ. (In progress)

Stanford Project: The MYCIN PROJECTS -- KNOWLEDGE ENGINEERING FOR MEDICAL CONSULTATION

Principal Investigators: Bruce G. Buchanan, Ph.D.
Department of Computer Science
Stanford University
Stanford, California 94305
(415) 497-0935 (BUCHANAN@SUMEX-AIM)

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University
Stanford, California 94305
(415) 497-6970 (SHORTLIFFE@SUMEX-AIM)

Subproject Directors:

MYCIN:	Dr. Shortliffe, and A. Carlisle Scott
EMYCIN:	Dr. Buchanan, and William van Melle
GUIDON:	Dr. Buchanan, and William J. Clancey
ONCOCIN:	Dr. Shortliffe, and A. Carlisle Scott

The MYCIN Projects are a set of research efforts overseen by a collaborative group of physicians and computer scientists who are developing intelligent systems using the techniques of knowledge engineering. The research focus includes knowledge acquisition, inexact reasoning, explanation, education, and the representation of time and of expert thinking patterns. Project members currently are working in a variety of medical domains including infectious disease therapy selection, intelligent computer-aided instruction, and the management of cancer chemotherapy protocols. Recent emphasis in the research has included intensive work regarding human engineering in an effort to implement the cancer therapy system for routine use by physicians. There is also a heightened interest in gearing representation, knowledge acquisition, and explanation more to the way that an expert actually thinks.

MYCIN Project (continued)

SOFTWARE AVAILABLE ON SUMEX

MYCIN--A consultation system designed to assist physicians with the selection of antimicrobial therapy for severe infections. It has achieved expert level performance in formal evaluations of its ability to select therapy for bacteremia and meningitis. Although MYCIN is no longer the subject of an active research program, the system continues to be available on SUMEX for demonstration purposes and as a testing environment for other research projects.

EMYCIN--The "essential MYCIN" system is a generalization of the MYCIN knowledge representation and control structure. It is designed to facilitate the development of new expert consultation systems for both clinical and non-medical domains.

GUIDON--A system developed for intelligent computer-aided instruction. Although it is being developed in the context of MYCIN's infectious disease knowledge base, the techniques are generalizable to any EMYCIN domain. The current research emphasis has been on an improved understanding of how the expert thinks so as to optimize the learning experience for the student.

ONCOCIN--This newest subproject is a system designed to provide oncologists with consultations regarding the management of patients receiving chemotherapy for cancer. Much of the knowledge in this domain is already well-specified in protocol documents, but expert judgments also need to be understood and modeled. Other research topics include human engineering and the development of capabilities to help achieve clinical acceptance of the program.

REFERENCES

Clancey, W.J., Shortliffe, E.H. and Buchanan, B.G.: Intelligent computer-aided instruction for medical diagnosis. Proc. Third Symposium on Computer Applications in Medical Care, Silver Spring, Maryland, October, 1979, pp. 175-183.

Shortliffe, E.H. Consultation systems for physicians. Proceedings of the CSCSI/SCEIO Conference, 14-16 May 1980, University of Victoria, British Columbia, pp. 1-11.

Shortliffe, E.H. The computer as clinical consultant (editorial). Arch. Int. Med., March 1980.

van Melle, W. A domain-independent system that aids in constructing knowledge-based consultation programs. PhD thesis, Computer Science Department, Stanford University, June 1980.

Stanford Project: PROTEIN STRUCTURE MODELING (CRYSALIS)

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

Contact: Allan TERRY@SUMEX-AIM
(415) 497-1740

The CRYSALIS system is an application of artificial intelligence methodology to the task domain of protein crystallography. The focus is the structure determination problem: the derivation of an atomic model of the protein from an indistinct image of the electron density. The crystallographer interprets these data in light of the known chemical composition of the protein, general principles of protein chemistry, and his own experience. The goal of the CRYSALIS Project is to integrate these diverse sources of knowledge and data into a program that matches the crystallographer's level of performance in electron density map interpretation. A successful solution to this problem must deal with issues such as representation and management of a large knowledge base, opportunistic reasoning, and appropriate description of the emerging hypothesis, while keeping human engineering considerations in sight. Automation of this task would shorten the time for protein determination by several weeks to several months and would fill a major gap in the construction of a fully-automated system for protein crystallography.

SOFTWARE AVAILABLE ON SUMEX

CRYSTALLOGRAPHIC DATA REDUCTION PROGRAMS (in FORTRAN):

- A density map skeletonizer (SKEL37) based on an improved version of Greer's algorithm.
- A package for locating the critical points in a map.
- A general map-manipulation utility (INSPCT) that can find peaks, display regions, and compute various statistics.

TWO LISP SYSTEMS (with the caveat that both are under active development):

- A system (SEGLABELING) which heuristically parses the segmented map into labels similar to those a crystallographer would use.
- The inference system (CRYSALIS).

REFERENCES

Engelmore, R.S. and Nii, H.P.: A knowledge-based system for the interpretation of protein x-ray crystallographic data. Heuristic Programming Project Report HPP-77-2, Computer Science Dept., Stanford Univ., January, 1977.

Engelmore, R. and Terry, A.: Structure and function of the CRYSALIS system. Proc. Sixth IJCAI, Tokyo, August, 1979.

Nii, H.P. and Feigenbaum, E.A.: Rule-based understanding of signals. Heuristic Programming Project Report HPP-77-7, Computer Science Dept., Stanford Univ., April, 1977.

Stanford Project: RX--DERIVING KNOWLEDGE FROM TIME-ORIENTED CLINICAL DATABASES

Principal Investigators: Robert L. Blum, M.D.
Departments of Medicine
and Computer Science
Stanford University
Stanford, California 94305
(415) 497-3088 (BLUM@SUMEX-AIM)

Gio C.M. Wiederhold, Ph.D.
Departments of Computer Science
and Electrical Engineering
Stanford University
Stanford, California 94305
(415) 497-0635 (WIEDERHOLD@SUMEX-AIM)

The objective of clinical database (DB) systems is to derive medical knowledge from the stored patient observations. However, the process of reliably deriving causal relationships has proven to be quite difficult because of the complexity of disease states and time relationships, strong sources of bias, and problems of missing and outlying data.

The goal of the RX Project is to explore the usefulness of knowledge-based computational techniques in solving this problem of accurate knowledge inference from non-randomized, non-protocol patient records. Central to RX is a knowledge base (KB) of medicine and statistics, organized as a taxonomic tree consisting of frames with attached data and procedures. The KB is used to retrieve time-intervals of interest from the DB and to assist with the statistical analysis. Derived knowledge is incorporated automatically into the KB. The American Rheumatism Association DB containing 7,000 patient records is used.

SOFTWARE AVAILABLE ON SUMEX

RX--(excluding the knowledge base and clinical database) consists of approximately 200 INTERLISP functions. The following groups of functions may be of interest apart from the RX environment:

SPSS Interface Package: Functions which create SPSS source decks and read SPSS listings from within INTERLISP.
Statistical Tests in INTERLISP: Translations of the Piezer-Pratt approximations for the T,F, and Chi-square tests into LISP.
Time-Oriented Data Base and Graphics Package: Autonomous package for maintaining a time-oriented database and displaying labelled time-intervals.

REFERENCES

Blum, R.L. and Wiederhold, G.: Inferring knowledge from clinical data banks utilizing techniques from artificial intelligence. Proc. Second Annual Symposium Computer Applications in Medical Care, IEEE, Washington, D.C., November, 1978, pp. 303-307.

Blum, R.L.: Automating the study of clinical hypotheses on a time-oriented database: The RX project. Submitted to MEDINFO80, Third World Conference on Medical Informatics, Tokyo, 1980.

Weyl, S., Fries, J., Wiederhold, G. and Germano, F.: A modular self-describing clinical databank system. Comp. and Biomed. Res. 8(3):279-293, June, 1975.

Wiederhold, G., Fries, J.F.: Structured organization of clinical data bases. AFIPS Conference Proc. 44:479-485, 1975.

Appendix B

AI Handbook Outline

E. A. Feigenbaum and A. Barr
Computer Science Department
Stanford University

This is a list of the Chapters in the Handbook. Articles in the first five Chapters appear in Volume I. Articles in Chapters VI through X will appear in Volume II, and the remaining chapters are expected to appear in Volume III. A list of all of the articles in each Chapter (although tentative for Volume III) follows.

VOLUME I:

I. Introduction
II. Search
III. Knowledge Representation
IV. Understanding Natural Language
V. Understanding Spoken Language

VOLUME II:

VI. Programming Languages for AI Research
VII. Applications-oriented AI Research: Science
VIII. Applications-oriented AI Research: Medicine
IX. Applications-oriented AI Research: Education
X. Automatic Programming

VOLUMES III (Tentative):

XI. Models of Cognition
XII. Automatic Deduction
XIII. Vision
XIV. Robotics
XV. Learning and Inductive Inference
XVI. Planning and Problem Solving

VOLUME I

I. INTRODUCTION

A. What is Artificial Intelligence?
B. The AI Handbook
C. Accessing the AI literature

II. SEARCH

A. Overview
B. Problem representation
  1. State-space representation
  2. Problem-reduction representation
  3. Game trees
C. Search methods
  1. Blind state-space search
  2. Blind AND/OR graph search
  3. Heuristic state-space search
    a. Basic concepts in heuristic search
    b. A*--Optimal search for an optimal solution
    c. Relaxing the optimality requirement
    d. Bidirectional search
  4. Heuristic search of an AND/OR graph
  5. Game tree search
    a. Minimax procedure
    b. Alpha-beta pruning
    c. Heuristics in game tree search
D. Sample search programs
  1. Logic Theorist
  2. General Problem Solver
  3. Gelernter's geometry theorem-proving machine
  4. Symbolic integration programs
  5. STRIPS
  6. ABSTRIPS

III. KNOWLEDGE REPRESENTATION

A. Overview
B. Survey of representation techniques
C. Representation schemes
  1. Logic
  2. Procedural representations
  3. Semantic networks
  4. Production systems
  5. Direct (analogical) representations
  6. Semantic primitives
  7. Frames and scripts

IV. UNDERSTANDING NATURAL LANGUAGE

A. Overview
B. Machine translation
C. Grammars
   1. Review of formal grammars
   2. Transformational grammars
   3. Systemic grammar
   4. Case grammars
D. Parsing
   1. Overview of parsing techniques
   2. Augmented transition networks
   3. The General Syntactic Processor
E. Text generation
F. Natural language processing systems
   1. Early natural language systems
   2. Wilks's machine translation system
   3. LUNAR
   4. SHRDLU
   5. MARGIE
   6. SAM and PAM
   7. LIFER

V. UNDERSTANDING SPOKEN LANGUAGE

A. Overview
B. Speech systems architecture
C. The ARPA SUR projects
   1. HEARSAY
   2. HARPY
   3. HWIM
   4. The SRI/SDC speech systems

VOLUME II

VI. PROGRAMMING LANGUAGES FOR AI RESEARCH

A. Historical overview
B. Features of AI programming languages
   1. Overview and comparison
   2. Data structures
   3. Control structures
   4. Pattern matching
   5. Programming environment
   6. Truth maintenance
C. Major AI programming languages
   1. LISP
   2. PLANNER and CONNIVER
   3. QLISP
   4. SAIL
   5. POP-2
   6. FUZZY

VII. APPLICATIONS-ORIENTED AI RESEARCH: SCIENCE

A. Overview of applications-oriented AI research
B. TEIRESIAS--Issues in designing expert systems
C. Research on applications in chemistry
   1. Applications in chemical analysis
   2. The DENDRAL programs
      a. DENDRAL
      b. CONGEN and its extensions
      c. Meta-DENDRAL
   3. CRYSALIS
   4. Applications in organic synthesis
D. Other scientific applications
   1. MACSYMA
   2. The SRI computer-based consultant
   3. PROSPECTOR
   4. AI in database management

VIII. APPLICATIONS-ORIENTED AI RESEARCH: MEDICINE

A. Overview
B. Medical applications systems
   1. MYCIN
   2. CASNET
   3. INTERNIST
   4. Present Illness Program
   5. Digitalis Advisor
   6. IRIS
   7. EXPERT

IX. APPLICATIONS-ORIENTED AI RESEARCH: EDUCATION

A. Historical overview of AI applications in education
B. Issues in the design of tutoring systems
C. Computer-based tutoring systems
   1. SCHOLAR
   2. WHY
   3. SOPHIE
   4. WEST
   5. WUMPUS
   6. GUIDON
   7. BUGGY
   8. EXCHECK
D. Research on nontutorial uses of AI in education

X. AUTOMATIC PROGRAMMING

A. Overview--Methods of program specification

B. Basic approaches to automatic programming
C. Automatic programming systems
   1. PSI
   2. SAFE
   3. Programmer's Apprentice
   4. PECOS
   5. DAEDALUS
   6. PROTOSYSTEM-1
   7. NLPQ
   8. LIBRA--Automatic program optimization

VOLUME III (Tentative)

XI. MODELS OF COGNITION

A. Overview
B. General Problem Solver
C. Models of cognitive development
D. EPAM
E. Semantic-network models of memory
   1. Quillian's semantic memory system
   2. HAM
   3. ACT
   4. MEMOD
F. Belief systems

XII. AUTOMATIC DEDUCTION

A. Overview
B. Resolution-based theorem proving
C. Nonresolution theorem proving
D. Applications of theorem proving
E. Nonmonotonic logic

XIII. VISION

A. Overview
B. Blocks-world understanding
C. Processing of visual data
D. Shape understanding
E. Representation and control methods in vision
F. Sample applications in vision research

XIV. Robotics

A. Overview
B. Computation in a physical environment
C. Engineering and kinematics
D. Languages and simulation
E. Planning and representation

XV. Learning and Inductive Inference

- A. Overview
- B. Rote learning
- C. Advice taking
- D. Learning from examples
  - 1. Overview
  - 2. Adaptive learning
  - 3. Learning single concepts
  - 4. Learning multiple concepts
  - 5. Learning by doing

XVI. Planning and Problem Solving

- A. Overview
- B. Linear planners
- C. Hierarchical planners
  - 1. NOAH and extensions
  - 2. MOLGEN
- D. Opportunistic planning

Appendix C

AIM Management Committee Membership

The following are the membership lists of the various SUMEX-AIM management committees at the present time:

AIM Executive Committee:

LEDERBERG, Joshua, Ph.D. (Chairman)
President
The Rockefeller University
1230 York Avenue
New York, New York 10021
(212) 360-1234, 360-1235

AMAREL, Saul, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546

BAKER, William R., Jr., Ph.D. (Exec. Secretary)
Biotechnology Resources Program
National Institutes of Health
Building 31, Room 5B43
9000 Rockville Pike
Bethesda, Maryland 20205
(301) 496-5411

FEIGENBAUM, Edward, Ph.D.
Principal Investigator - SUMEX
Department of Computer Science
Margaret Jacks Hall, Room 216
Stanford University
Stanford, California 94305
(415) 497-4079

LINDBERG, Donald, M.D. (Adv Grp Member)
605 Lewis Hall
University of Missouri
Columbia, Missouri 65201
(314) 882-6966

MYERS, Jack D., M.D.
School of Medicine
Scaife Hall, 1291
University of Pittsburgh
Pittsburgh, Pennsylvania 15261

SHORTLIFFE, Edward H., M.D., Ph.D.
Co-Principal Investigator - SUMEX
Division of General Internal Medicine, TC117
Stanford University Medical Center
Stanford, California 94305
(415) 497-5821

AIM Advisory Group:

LINDBERG, Donald, M.D. (Chairman)
605 Lewis Hall
University of Missouri
Columbia, Missouri 65201
(314) 882-6966

AMAREL, Saul, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546

BAKER, William R., Jr., Ph.D. (Exec. Secretary)
Biotechnology Resources Program
National Institutes of Health
Building 31, Room 5B43
9000 Rockville Pike
Bethesda, Maryland 20205
(301) 496-5411

FEIGENBAUM, Edward, Ph.D. (Ex-officio)
Principal Investigator - SUMEX
Department of Computer Science
Margaret Jacks Hall, Room 216
Stanford University
Stanford, California 94305
(415) 497-4079

LEDERBERG, Joshua, Ph.D.
President
The Rockefeller University
1230 York Avenue
New York, New York 10021
(212) 360-1234, 360-1235

MINSKY, Marvin, Ph.D.
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
545 Technology Square
Cambridge, Massachusetts 02139
(617) 253-5864

MOHLER, William C., M.D.
Associate Director
Division of Computer Research and Technology
National Institutes of Health
Building 12A, Room 3033
9000 Rockville Pike
Bethesda, Maryland 20205
(301) 496-1168

MYERS, Jack D., M.D.
School of Medicine
Scaife Hall, 1291
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
(412) 624-2649

PAUKER, Stephen G., M.D.
Department of Medicine - Cardiology
Tufts New England Medical Center Hospital
171 Harrison Avenue
Boston, Massachusetts 02111
(617) 956-5910

SHORTLIFFE, Edward H., M.D., Ph.D. (Ex-officio)
Co-Principal Investigator - SUMEX
Division of General Internal Medicine, TC117
Stanford University Medical Center
Stanford, California 94305
(415) 497-5821

SIMON, Herbert A., Ph.D.
Department of Psychology
Baker Hall, 339
Carnegie-Mellon University
Schenley Park
Pittsburgh, Pennsylvania 15213
(412) 578-2787 or 578-2000

Stanford Community Advisory Committee:

FEIGENBAUM, Edward, Ph.D. (Chairman)
Department of Computer Science
Margaret Jacks Hall, Room 216
Stanford University
Stanford, California 94305
(415) 497-4079

SHORTLIFFE, Edward H., M.D., Ph.D.
Co-Principal Investigator - SUMEX
Division of General Internal Medicine, TC117
Stanford University Medical Center
Stanford, California 94305
(415) 497-5821

DJERASSI, Carl, Ph.D.
Department of Chemistry, Stauffer I-106
Stanford University
Stanford, California 94305
(415) 497-2783

MAFFLY, Roy H. Maffly, M.D.
Division of Nephrology
Veterans Administration Hospital
3801 Miranda Avenue
Palo Alto, California 94304
(415) 858-3971

References

1. Feigenbaum, E.A., The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering, Proceedings of the 1978 National Computer Conference, AFIPS Press, (1978).

2. Nilsson, N.J., Principles of Artificial Intelligence, Tioga Publishing Company, Palo Alto, California (1980).

3. Winston, P.H., Artificial Intelligence, Addison-Wesley Publishing Co., (1977).

4. Nilsson, N.J., Artificial Intelligence, Information Processing 74, North-Holland Pub. Co. (1975).

5. Barr A. and Feigenbaum, E.A. (Eds.), The Handbook of Artificial Intelligence Volume I, William Kaufmann, Inc. Los Altos, Calif. (1981)

6. Boden, M., Artificial Intelligence and Natural Man, Basic Books, New York, (1977).

7. McCorduck, P., Machines Who Think, W.H. Freeman and Co., San Francisco (1979).

8. Coulter, C. L., Research Instrument Sharing, Science, Vol. 201, No. 4354, August 4, 1978.

9. Metcalfe, R.M. and Boggs, D.R., Ethernet: Distributed Packet Switching for Local Computer Networks, Comm. ACM, Vol. 19, No. 7 (July 1976).

10. Shoch, J.F. and Hupp, J.A., Performance of an Ethernet Local Network -- A Preliminary Report, Proceedings of the Local Area Communications Network Symposium, Boston,May 1979.

11. Taft, E.A., Implementation of PUP in TENEX, Internal XEROX PARC memorandum, June 1978.

12. Boggs, D.R., Shoch, J.F., Taft, E.A., and Metcalfe, R.M., PUP: An Internetwork Architecture, XEROX PARC report CSL-79-10, July 1979.

13. Digital Equip. Corp., Intel Corp., and Xerox Corp., The Ethernet - Data Link and Physical Layer Specifications, Version 1.0, September 30, 1980.