SUMEX

STANFORD UNIVERSITY
MEDICAL EXPERIMENTAL COMPUTER RESOURCE

RR - 00785

COMPETING RENEWAL APPLICATION

Submitted to

BIOTECHNOLOGY RESOURCES PROGRAM
NATIONAL INSTITUTES OF HEALTH

and Ann. Rept. 07.

June 1, 1980

STANFORD UNIVERSITY SCHOOL OF MEDICINE
Edward A. Feigenbaum, Principal Investigator

Prelude: An Overview and Personal Statement

by Edward A. Feigenbaum, Principal Investigator

This prelude is unabashedly a statement of advocacy. As we prepare this proposal, gathering up the threads of our past achievement and weaving them into a coherent picture of our future, there is in the SUMEX Project a sense of pride and accomplishment, and a feeling of exhilaration and momentum regarding the future.

SUMEX was established with three main goals:

1. to provide computing resources and human assistance to those scientists working on applications of artificial intelligence research in medicine and biology;
2. to test the idea that it was feasible to provide resources and assistance to the nation from a single site, with time-shared operating systems, national computer communication networks, and a staff oriented toward the special problems of remote users;
3. to grow, from seed to plant, the community of scientists interested in working on applications of AI to the biomedical sciences; facilitating the growth, health, and vigor of the community by the use of electronic communications linking its members. One question we were asking was, "Is there a new style of science that will emerge in a communications-enhanced setting of national, rather than institutional, scope?"

These goals were and are unique to SUMEX, and their pursuit has given rise to a "spirit of SUMEX"--a spirit that unfortunately does not come across well in the dry recitations of a proposal document; hence, this personal prelude.

SUMEX's success as a national research resource

The SUMEX Project has demonstrated that it is possible to operate a computing research resource with a national charter--that the services providable over networks were those that were facilitative of the growth of AI-in-Medicine. Previous NIH computer RR's were mostly institutional in scope, occasionally regional (like the UCLA resource).

Some of the most notable projects in the history of Artificial Intelligence were done with terminal-and-network, without a computer on site. In human terms, this means,of course, without the headaches and energy drains of proposing a machine, installing it, maintaining it and its software, hiring its system programmers and operators, dealing with communication vendors, etc. The famous INTERNIST program was developed from Pittsburgh in this way. And the ACT computer model was begun at

Michigan, continued at Yale, and later at Carnegie-Mellon, all without moving the program or losing a day's work because of machine transition problems.

The projects SUMEX supports have generally required substantial computing resources with excellent interaction. This is hard to obtain in all but a few universities. SUMEX is, in a sense, a "great equalizer". A scientist gains access by virtue of the quality of his/her research ideas, not by the accident of where s/he happens to be situated--in other words, the ethic of the scientific journal.

SUMEX has demonstrated that a computer resource is a useful "linking mechanism" for bringing together and holding together teams of experts from different disciplines who share a common problem focus. For example, computer scientists have been collaborating fruitfully with physical chemists, molecular biochemists, geneticists, crystallographers, internists, ophthalmologists, infectious disease specialists, intensive care specialists, oncologists, psychologists, biomedical engineers, and other expert practitioners. And in some of these cases, the interdisciplinary collaboration, usually so difficult to achieve in the best of circumstances, was achieved in spite of geographical distance between the participants, using the computer networks.

SUMEX has achieved successes as a community builder. AI concepts and software are among the most complex products of computer science. Historically it has not been easy for scientists in other fields to gain access to and mastery of them. Yet the collaborative outreach of SUMEX has been able to bridge the gap in a number of cases. For example, Dr John Osborn (Pacific Medical Center, San Francisco) and I found common scientific interests in the application of AI to intensive care, and initiated a SUMEX-based collaboration. That project resulted in a system of potential significance to intensive care medicine; in two Stanford computer science Ph.D. dissertations, hence two new doctoral-level recruits to the ranks of computers-in-medicine specialists; in one computer science/physiology Special Ph.D. Program for one of Dr. Osborn's biomedical engineers; and an award to Dr. Osborn's team in 1979 from the Association for the Advancement of Medical Instrumentation.

I wish to contrast this success story with the traditional difficulties I have encountered outside the health research field in trying to bridge the gap to engineering-oriented industrial firms. The human resource and motivation was present. The SUMEX base of easily available shared software technology was not. The resulting problems have generally raised too high a threshold to overcome.

The SUMEX mission has been able to capture the contributions of some of the finest computers-in-medicine specialists and computer scientists in the country. For example, Professor Joshua Lederberg (SUMEX's first PI, now President of The Rockefeller University) is Chairman of SUMEX's Executive Committee; and Professor Donald Lindberg, M.D., Director of the University of Missouri's Health Care Technology Center, is Chairman of the AIM Advisory Group. Professor Herbert Simon of Carnegie-Mellon University, Professor Marvin Minsky of MIT, and many other distinguished scientists serve on that peer review committee. These people are active participants in SUMEX. Lederberg and Lindberg are continuing collaborators in the research itself. And Simon, for example, was the person who prompted our collaboration with psychologists at the University of Colorado.

SUMEX now has the reputation of a model national resource, pulling together the best available interactive computing technology, software, and computer communications in the service of a national scientific community. Planning groups for national facilities in cognitive science, computer science, and biomathematical modeling have discussed and studied the SUMEX model.

SUMEX and Artificial Intelligence Research

The SUMEX Project is a relative latecomer to AI research. Yet its scope has given strong impetus to this historic development in computer application. AI research is that part of computer science that investigates symbolic reasoning processes, and the representation of symbolic knowledge for use in inference. It views heuristic knowledge to be of equal importance with "factual" knowledge, indeed to be the essence of what we call "expertise". In its "Expert Systems" work, it seeks to capture the expertise of a field, and translate it into programs that will offer intelligent assistance to a practitioner in that field.

For computer applications in medicine and biology, this research path is crucial, indeed ineluctable. Medicine and biology are not presently mathematically-based sciences; not like physics and engineering capable of exploiting the mathematical characteristics of computation. They are essentially inferential, not calculational, sciences. If the computer revolution is to affect biomedical scientists, computers will be used as inferential aids.

Perhaps the larger impact on medicine and biology will be the exposure and refinement of the hitherto largely private heuristic knowledge of the experts of the various fields studied. The ethic of science that calls for the public exposure and criticism of knowledge has traditionally been flawed for want of a methodology to evoke and give form to the heuristic knowledge of scientists. The AI methodology is beginning to fill that need. Heuristic knowledge can be elicited, studied, critiqued by peers, and taught to students.

The tide of AI research and application is rising. AI is one of the fronts along which university computer science groups are expanding. The NSF's program in Intelligent Systems is vigorous and growing. The pressure from student career-line choices is great: to cite an admittedly special case, approximately one-third of the students applying to Stanford's computer science Ph.D. program cite AI as a possible field of specialization. In industry, new groups have been forming regularly: Texas Instruments two years ago formed a substantial AI group; so did the oil-industry-service firm, Schlumberger, Inc.; IBM has reinitiated its AI work; and the new genetic engineering firms are becoming interested.

The tide is rising largely because of the development in the 1970's of methods and tools for the application of AI concepts to difficult professional-level problem solving; and the demonstration in various areas of medicine and other life sciences that these methods and tools really work. Here SUMEX has played a key role, so much so that it is regarded as "the home of applied AI."

SUMEX has been the nursery, as well as the home, of such well-known AI systems as DENDRAL (chemical structure elucidation), MYCIN (infectious disease diagnosis and therapy), INTERNIST (differential diagnosis), and ACT (human memory organization). These, and other programs developed at SUMEX, have played a seminal role in structuring modern AI paradigms and methodology. First among these has been a shift of AI's focus from inference procedures to knowledge representation and use. There is now a recognition that the power of problem solvers derives primarily from the knowledge that they contain--of the elements of the problem domain, of the strategies for solving problems in that domain, and of the forms in which the knowledge is to be acquired. In 1977, Goldstein and Papert of MIT, writing in the journal Cognitive Science, described the change of focus as a "paradigm shift" in AI. This shift was induced largely (though of course not exclusively) by the work at SUMEX, beginning with the DENDRAL development in 1965.

Toward the mid-'80s: the Future of SUMEX

Success breeds its problems. The revolution in computer technology and costs adds complexity to their solution.

At the beginning, the SUMEX community was small, and idea-limited. The SUMEX computer facility was an ideal vehicle for the research. Now the community is large, and the momentum of the science is such that its progress is now limited by computing power. The size and scientific maturity of the SUMEX community has fully consumed the resource in every critical dimension: CPU power, main memory size, and file space.

The limitation that AI researchers agree most critically limits their scientific imagination, and adds inordinately to program development time, is the 256K word main memory space, brought about by the 18 bit address of the PDP-10's and 20's. Economically, main memory size need not be much of a limitation any more, but it is essential to move to a machine with more addressing bits.

But which machine? In the turmoil of the computer developments of today, this is not easy to answer. Computers will come in many different sizes and prices and each will fit a particular class of needs. Our planning axiom for the period 1981-86 has been: the need to accommodate a HETEROGENEITY of computers and peripheral devices. We must maintain a flexible posture with respect to the introduction of new capabilities and changing costs during this continuing revolution. Yet we must choose.

Our plan, sketched below, is conservative in maintaining and extending SUMEX's current service level; yet is forward-looking enough to position SUMEX properly for mid-course corrections and for the computing world of the late 1980's. Here it is, briefly sketched.

The existing DEC KI-10 duplex, with its superb software, will be "filled out"--stretched to the point of diminishing returns from hardware addition; then frozen. It is an amiable workhorse. We can not (indeed dare not) do without it during this period of turbulence. But it has seen better days, and will be ineffective by the end of the grant period.

A DEC VAX 11/780 will be acquired in the first year. Based on more modern technology and a more competitive price, it has the extra address bits that are required. On VAX we get the same kind of low-cost ride on the software work of others that we got when we adopted TENEX and INTERLISP for the KI-10's. The UNIX operating system is available, and is being further developed under ARPA support. ARPA is also supporting the reprogramming of INTERLISP for VAX. For integrated circuit design research, ARPA has already placed two VAX computers at our Computer Science Department, so we are building experience rapidly in VAX use. And, de facto, the VAX has become the "computer science machine" of the early '80s, so that nationally its software development is moving rapidly. A family of VAX's, both more and less powerful, at (hopefully) appropriate prices, is in the wings.

The "technology transfer" machine to which we will move the heavy national use of SUMEX's mature AI applications (such as DENDRAL, SECS, MOLGEN, VM) will be another DEC VAX, acquired in the middle of the period. This machine's role is intended to be entirely analogous to the role currently played by the DEC 2020 at SUMEX vis a vis the KI-10 duplex. It will be the VAX-era prototype of the "spinoff" machine, loosely tethered to SUMEX by networks. In the last DENDRAL Project renewal, the NIH Study Section denied such a machine to DENDRAL, suggesting that the required resource would better be provided by SUMEX. We seek, and plan, to assume this obligation.

And what about the single-user professional scientific workstation--the powerful, small, cheap officemate that will serve most of the researcher's computing needs? Much of the present turbulence in the computing world swirls around this question. Yes, we believe it is coming, and will probably be an economically viable concept in the late '80s. No, we do not believe it will be powerful enough or cheap enough for most routine research needs in the planning period. Yet we must begin to explore the space of possibilities opened up by these machines, eschewing articles of faith for real experience. We must learn to build systems of these machines and to build and manage graceful software for these systems. If decentralization is in our future, we must learn its technical characteristics. Consequently, we have planned the acquisition of a number of such single-user workstations over the course of the coming period, some to be placed at Stanford, some in the national community, at the decision of the Executive Committee.

These machines will be tethered to the SUMEX central facility and staff by local digital network at Stanford and by national network to the non-Stanford community. With DEC 10's, 20, VAX's, and workstations coexisting to serve community needs, it is economical and convenient to continue the centralization of file storage, and the networks make it possible for most applications at Stanford and many applications nationally. Computer scientists are in general agreement that economies of scale will continue to dominate in secondary storage for some time. We have planned, therefore, to alleviate the present file space shortage not by adding discs to machines in an ad hoc fashion but by adding a common file server to the resource. To facilitate the transfer of software and access to valuable common facilities, the SUMEX complement of equipment will be linked by local digital networks to other major centers of computing at Stanford, most important of which is the Computer Science Department.

The success of SUMEX is the success of its dedicated and extraordinarily competent staff, headed by Tom Rindfleisch. This human resource of SUMEX should not, and will not, be decentralized. In the world of computer systems talent and user-assistance expertise, there are indeed continuing large "economies of scale".

The smoothly operating management structure of SUMEX is one of its joys and victories. We do not plan to fix something that is not broken. We plan that the Executive Committee and the AIM Advisory Committee will continue to function as they now do.

So this is it in a nutshell:

Run the present configuration with more main memory; acquire two VAX large-memory systems (years 1 and 3) for new research and for maturing project communities; cautiously add some single-user professional workstations; acquire a common file server; link everything in a transparent digital networking scheme; continue the central staff and management structure, essentially unchanged in size and function.

As we add up the budget (flinchingly, I hasten to say), we note that the cost will not be cheap, despite the much-touted fall in the cost of computing. But we believe we have been conservative; that the scientific community we serve needs these resources; and that by its science and its applications orientation, it has earned them.

I look at the widely acclaimed NSF report calling for the refurbishing of computer equipment for experimental computer science (the so-called "Feldman Report") and note that it calls for "refurbishing" expenditures for just a single department greater than that budgeted in this proposal, with a "refresh" cycle of five years to accommodate advancing technology. The scientific work of the SUMEX-AIM community is the quintessence of experimental computer science. It is advancing, and gaining acceptance, beyond expectations. SUMEX serves the nation, not one university or department. I believe that its budget accords well with the national interest and with the scientific interest.

Conclusion: the "Spirit of SUMEX"

I would like to conclude not with my own words but with the words of Professor Douglas Brutlag, a Stanford Biochemist who collaborates with my group on the MOLGEN project and who sent me, unsolicited, the letter quoted below in its entirety. Nothing I could say could more accurately portray the "spirit of SUMEX" mentioned earlier.

> "My original role in the Molgen project was that of a biochemist advisor to those developing a knowledge base of molecular biological information and techniques. I rapidly found that SUMEX could be very useful to my own work in ways that I had never expected. First, MOLGEN was a success very early and I now routinely use the artificial intelligence methods incorporated within the frame oriented knowledge base in my everyday work in the laboratory. I use the knowledge base not only to store our results from experiments and to analyze them, but I can readily interact with the knowledge base to examine the data from several different viewpoints and display it in different ways.
>
> In addition to the interactive nature of knowledge base work, I have found computer networks and file transfer protocols to be exceptionally useful. The nation wide commercial networks have permitted many of my colleagues across the country to try out the software we have developed at Stanford in pilot projects. This together with message sending capabilities has resulted in instantaneous feed back about the work we have done and allowed us to develop our program and to incorporate ideas from a much larger base of expertise. Several collaborative arrangements have been set up and some have even become involved in our programming efforts. Moreover, our software has had such general utility that subsequently many of the other workers have obtained accounts on their local computers and we have sent them the software by file transfer protocols. Electronic information transfers have saved both time and energy in preparing hard copy versions as well as facilitated the update programs at many distant locations.
>
> I think that one of the major reasons that SUMEX works so well is that it is designed with the naive user in mind. Because it is so interactive and user oriented, the activation energy to learn how to use the system is very low. Of all of the interactive systems with which I have worked (five in all), SUMEX was not only the easiest, but was indeed a real pleasure. I felt more like the system was working for me from the very beginning, rather than me fighting the system. Hence, my productivity on SUMEX has increased immeasurably. In addition, I have no hesitation encouraging others at remote sites to use SUMEX in the collaborative efforts mentioned above."

Table of Contents

Section Page

Prelude: An Overview and Personal Statement i

List of Figures xiii

1. Biographical Sketches 2

2. Budget 3

2.1 First Year Budget Detail (8/1/81 - 7/31/82) 3

2.1.1 Total First Year Budget 3

2.1.2 First Year Personnel Detail 4

2.2 5-year Budget Summary (8/81 - 7/86) 5

2.3 Budget Explanation and Justification 6

3. Introduction and Aims 13

3.1 Overview of Objectives and Rationale 14

3.1.1 Definitions of Artificial Intelligence 14

3.1.2 Resource Sharing 16

3.2 SUMEX-AIM Background 16

3.3 Specific Aims 18

3.3.1 Resource Operations 19

3.3.2 Training and Education 20

3.3.3 Core Research 20

4. Significance 22

5. Progress 30

5.1 Brief Statement of Prior Goals 30
5.1.1 Resource Operations 30
5.1.2 Training and Education 30
5.1.3 Core Research 31
5.2 Summary of Progress: 11/77 - 4/80 32
5.3 Detailed Progress Highlights 34
5.3.1 Resource Operations 34
5.3.1.1 System Hardware 34
5.3.1.2 System Software 39
5.3.1.3 Network Communication Facilities 46
5.3.1.4 Resource Management 47
5.3.2 Core Research 49
5.3.3 SUMEX Staff Publications 51

6. Methods of Procedure 52
6.1 Resource Operations Plans 52
6.1.1 Resource Hardware 52
6.1.1.1 Rationale for Future Plans 52
6.1.1.2 Summary of Proposed Hardware Acquisitions . . 55
6.1.1.3 Existing Hardware Operation 56
6.1.1.4 Large Address Space Machines 57
6.1.1.5 Single-User Professional Workstations . . . 57
6.1.1.6 File Server 59
6.1.2 Communication Networks 62
6.1.2.1 Long-Distance Connections 62
6.1.2.2 Local Intermachine Connections 62
6.1.3 Resource Software 64
6.1.4 Community Management 66

6.2 Training and Education Plans 68

6.3 Core Research Plans 69

6.3.1 Knowledge Representation 71

6.3.1.1 RLL -- The Representation Language Language 71

6.3.1.2 Research on Planning 73

6.3.1.3 Causal Models 77

6.3.2 Knowledge Utilization and Tools for Building Expert Systems 78

6.3.2.1 Attempt to Generalize (AGE) 78

6.3.2.2 AI Handbook 78

6.3.2.3 Research in Automated Consultation about Expert Systems 79

6.3.2.4 EMYCIN 80

6.3.3 Knowledge Acquisition 82

6.3.4 Explanation 85

7. Available Facilities 89

8. Literature Cited 90

9. Collaborative Project Reports 135

9.1 Stanford Projects 136

9.1.1 AGE - Attempt to Generalize 137

9.1.2 AI Handbook Project 145

9.1.3 DENDRAL Project 149

9.1.4 MOLGEN Project 171

9.1.5 MYCIN Project 186

9.1.6 Protein Structure Project 205

9.1.7 RX Project 211

9.2 National AIM Projects 220

9.2.1 Acquisition of Cognitive Procedures (ACT) 221

9.2.2 SECS - Simulation and Evaluation of Chemical Synthesis 226

9.2.3 Hierarchical Models of Human Cognition 234

9.2.4 HMF - Higher Mental Functions 240

9.2.5 INTERNIST Project 244

9.2.6 PUFF/VM Project 250

9.2.7 Simulation of Cognitive Processes 261

9.2.8 Rutgers Computers in Biomedicine Project [Rutgers-AIM] 267

9.2.9 Decision Models in Clinical Diagnosis [Rutgers-AIM] 282

9.2.10 Heuristic Decisions in Metabolic Modeling [Rutgers-AIM] 286

9.3 Pilot Stanford Projects 289

9.3.1 Ultrasonic Imaging Project 290

9.4 Pilot AIM Projects 297

9.4.1 Coagulation Expert Project 298

9.4.2 Communication Enhancement Project 302

9.4.3 A Computerized Psychopharmacology Advisor 309

9.4.4 Computer-Aided Refinement of Medical Knowledge . . 317

9.4.5 Interactive Statistical Package Advisor 321

9.4.6 Conceptual Structures for Medical Diagnosis [Rutgers-AIM] 323

Appendix A

Community Growth and Project Synopses 331

Appendix B

Resource Operations and Usage Statistics 355

Appendix C

Local Network Integration 374

Appendix D

Remote Network Communication Facilities 376

Appendix E

Resource Management Structure 383

Appendix F

LISP Address Space Limitations 390

Appendix G

AI Handbook Outline 392

Appendix H

MAINSAIL System Demonstration 398

Appendix I

AIM Management Committee Membership 399

List of Figures

Figure Page

1. Current SUMEX-AIM KI-10 Computer Configuration 36

2. Current SUMEX-AIM 2020 Computer Configuration 37

3. Intermachine Connections via ETHERNET 38

4. Proposed VAX configuration 60

5. Planned Ethernet System to Integrate System Hardware 61

6. SUMEX-AIM Growth by Community 331

7. Total CPU Time Consumed by Month 356

8. Peak Number of Jobs by Month 357

9. Peak Load Average by Month 357

10. Monthly CPU Usage by Community 359

11. Monthly File Space Usage by Community 360

12. Monthly Terminal Connect Time by Community 361

13. Average Diurnal Loading (4/80): Number of Jobs 369

14. Average Diurnal Loading (4/80): Load Average 370

15. Average Diurnal Loading (4/80): Percent Time Used 370

16. TYMNET Terminal Connect Time 371

17. ARPANET Terminal Connect Time 372

18. TYMNET Network Node List 379

19. ARPANET Geographical Network Map 380

20. ARPANET Logical Network Map 381

21. TELENET Geographical Network Map 382

Form Approved
O.M.B. 68-R0249

SECTION I

DEPARTMENT OF
HEALTH, EDUCATION, AND WELFARE
PUBLIC HEALTH SERVICE

GRANT APPLICATION

LEAVE BLANK

TYPE	PROGRAM	NUMBER
REVIEW GROUP		FORMERLY
COUNCIL *(Month, Year)*		DATE RECEIVED

TO BE COMPLETED BY PRINCIPAL INVESTIGATOR *(Items 1 through 7 and 15A)*

1. TITLE OF PROPOSAL *(Do not exceed 53 typewriter spaces)*
S U Medical EXperimental Computer Resource (SUMEX)

2. PRINCIPAL INVESTIGATOR

2A. NAME *(Last, First, Initial)*
Feigenbaum, Edward A.

2B. TITLE OF POSITION
Professor and Chairman
Department of Computer Science

2C. MAILING ADDRESS *(Street, City, State, Zip Code)*
SUMEX Computer Project - Room TB105
Stanford University Medical Center
Stanford, California 94305

2D. DEGREE
Ph.D.

2E. SOCIAL SECURITY NO.

2F. TELEPHONE DATA — Area Code: 415 — TELEPHONE NUMBER AND EXTENSION: 497-4079

2G. DEPARTMENT, SERVICE, LABORATORY OR EQUIVALENT *(See Instructions)*
Departments of Genetics/Medicine

2H. MAJOR SUBDIVISION *(See Instructions)*
School of Medicine

3. DATES OF ENTIRE PROPOSED PROJECT PERIOD *(This application)*
FROM: 08/01/81
THROUGH: 07/31/86

4. TOTAL DIRECT COSTS REQUESTED FOR PERIOD IN ITEM 3
$ 6,793,862

5. DIRECT COSTS REQUESTED FOR FIRST 12-MONTH PERIOD
$ 1,336,864

6. PERFORMANCE SITE(S) *(See Instructions)*
Stanford University

[illegible] Rul
ann. rep.

7. Research Involving Human Subjects *(See Instructions)*
A. ☒ NO B. ☐ YES Approved: ________ Date
C. ☐ YES – Pending Review

8. Inventions *(Renewal Applicants Only - See Instructions)*
A. ☒ NO B. ☐ YES – Not previously reported
C. ☐ YES – Previously reported

TO BE COMPLETED BY RESPONSIBLE ADMINISTRATIVE AUTHORITY *(Items 8 through 13 and 15B)*

9. APPLICANT ORGANIZATION(S) *(See Instructions)*
Stanford University
Stanford, California 94305
IRS No. 94-1156365
Congressional District No. 12

10. NAME, TITLE, AND TELEPHONE NUMBER OF OFFICIAL(S) SIGNING FOR APPLICANT ORGANIZATION(S)
Larry J. Lollar
Sponsored Projects Officer
Sponsored Projects Office
Telephone Number (s) (415) 497-2883

11. TYPE OF ORGANIZATION *(Check applicable item)*
☐ FEDERAL ☐ STATE ☐ LOCAL ☒ OTHER *(Specify)* Private Non-Profit University

12. NAME, TITLE, ADDRESS, AND TELEPHONE NUMBER OF OFFICIAL IN BUSINESS OFFICE WHO SHOULD ALSO BE NOTIFIED IF AN AWARD IS MADE
K.D. Creighton
Associate Vice President - Controller
Stanford University
Stanford, California 94305
Telephone Number (415) 497-2251

13. IDENTIFY ORGANIZATIONAL COMPONENT TO RECEIVE CREDIT FOR INSTITUTIONAL GRANT PURPOSES *(See Instructions)*
01 School of Medicine

14. ENTITY NUMBER (Formerly PHS Account Number)
IRS No. 94-1156365

15. **CERTIFICATION AND ACCEPTANCE.** We, the undersigned, certify that the statements herein are true and complete to the best of our knowledge and accept, as to any grant awarded, the obligation to comply with Public Health Service terms and conditions in effect at the time of the award.

SIGNATURES (Signatures required on original copy only. *Use ink, "Per" signatures not acceptable)*	~~A. SIGNATURE OF PERSON NAMED IN ITEM 2A~~	DATE
	B. SIGNATURE(S) OF PERSON(S) NAMED IN ITEM 10 Larry J. Lollar	DATE 5/27/80

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

The undersigned agrees to accept responsibility for the scientific and technical conduct of the project and for the provision of required progress reports if a grant is awarded as the result of this application.

5/21/80
Date

Edward A. Feigenbaum
Edward A. Feigenbaum
Principal Investigator

SECTION 1

DEPARTMENT OF HEALTH, EDUCATION, AND WELFARE
PUBLIC HEALTH SERVICE

RESEARCH OBJECTIVES

LEAVE BLANK

PROJECT NUMBER

NAME AND ADDRESS OF APPLICANT ORGANIZATION

Stanford University, Stanford, California 94305

NAME, SOCIAL SECURITY NUMBER, OFFICIAL TITLE, AND DEPARTMENT OF ALL PROFESSIONAL PERSONNEL ENGAGED ON PROJECT, BEGINNING WITH PRINCIPAL INVESTIGATOR

E. Feigenbaum	Principal Investigator	Computer Science
E. Shortliffe	Co-Principal Invest.	Medicine
T. Rindfleisch	Facility Manager	Genetics/Medicine
E. Levinthal	AIM Liaison	Genetics

(See continuation page for additional professional personnel engaged on project.)

TITLE OF PROJECT

Stanford University Medical EXperimental Computer Resource (SUMEX)

USE THIS SPACE TO ABSTRACT YOUR PROPOSED RESEARCH. OUTLINE OBJECTIVES AND METHODS. UNDERSCORE THE KEY WORDS (NOT TO EXCEED 10) IN YOUR ABSTRACT.

Stanford University is developing and operating a NATIONAL SHARED COMPUTING RESOURCE in partnership with the NIH Biotechnology Resources Program to explore advanced application of COMPUTER SCIENCE in health research. There are two main objectives of the facility: 1) the managerial, administrative and technical demonstration of a national shared technological resource for health research, and 2) the specific encouragement of applicatio of ARTIFICIAL INTELLIGENCE IN MEDICINE (AIM). Besides the economic advantages of resource sharing made possible by emerging DATA COMMUNICATION technologies, a closer interaction between diverse research efforts is expected to promote a more systematic exchange of research products and ideas. This may be particularly true in applications of computer science. Multilateral community building rather than unilateral service is the project's essential mandate.

The term "artificial intelligence" (AI) is applied to research aimed at increasing the computer's effectiveness as a tool through the emulation of aspects of human SYMBOLIC REASONING and PROBLEM-SOLVING. The field emphasizes the judgmental manipulation of symbolic (non-numeric) representations of knowledge of a task domain for model-building and decision-making. Current applications include programs which assist in inferring chemical structures from spectrographic data, suggesting diagnoses and treatments within various classes of diseases, and modeling aspects of human behavior patterns.

Additional users of the facility will be selected within available resource computer capacity with the help of an AIM Executive Committee and Advisory Group on the basis of reviews of the proposed research. Selection criteria will include general scientific interest and merit, relevance to the AI mission, and community orientation of the collaborator.

LEAVE BLANK

RESEARCH OBJECTIVES (continuation page)

Stanford University Medical EXperimental Computer Resource (SUMEX)
Stanford University, Stanford, California 94305

Additional Professional Personnel Engaged on Project:

A. Sweer	System Programmer	Genetics/Medicine
F. Gilmurray	System Programmer	Genetics/Medicine
M. Bizzarri	System Programmer	Computer Science
M. Achenbach	System Programmer	Genetics/Medicine
W. Yeager	System Programmer	Genetics/Medicine
R. Tucker	System Programmer	Genetics/Medicine
B. Buchanan	Adjunct Professor	Computer Science
H.P. Nii	Research Associate	Computer Science
W. van Melle	Research Associate	Computer Science
N. Aiello	Scientific Programmer	Computer Science
N. Veizades	Electronics Engineer	Genetics/Medicine

1 Biographical Sketches

In order to reduce the bulk at the beginning of this already lengthy proposal, we have placed the biographical sketches for all professional personnel contributing to the project in the section starting on page 94.

SECTION II – PRIVILEGED COMMUNICATION

DETAILED BUDGET FOR FIRST 12-MONTH PERIOD	FROM	THROUGH
	08/01/81	07/31/82

DESCRIPTION *(Itemize)* PERSONNEL NAME	TITLE OF POSITION	TIME OR EFFORT %/HRS.	AMOUNT REQUESTED *(Omit cents)* SALARY	FRINGE BENEFITS	TOTAL
(see next page)	PRINCIPAL INVESTIGATOR				
			462,319	99,645	561,964

CONSULTANT COSTS None		--
EQUIPMENT		465,000*
Communications, interfaces, test equipment, etc.	10,000	
KI-10 AMPEX core expansion	65,000	
VAX 11-780	250,000	
AIM file server	120,000	
Terminals/displays/printers	20,000	
SUPPLIES		32,000
Computer operations	12,000	
Office supplies	5,000	
Engineering parts	15,000	
TRAVEL DOMESTIC		6,000
TRAVEL FOREIGN None		--
PATIENT COSTS None		--
ALTERATIONS AND RENOVATIONS None		--
OTHER EXPENSES		271,900
Equipment maintenance	108,400	
DEC KI-10 (51,000), Calcomp disks/tapes (13,900), DEC 2020 (15,000), DEC VAX (10,000), File Server (10,000), DEC PDP-11/GT-40 (4,000), Local terminals (4,500)		
Equipment lease	3,000	
Office telephones	7,500	
Local dataphones	10,000	
Software lease and license	6,000	
Technical Services/Repro./Books	4,000	
System and program documentation	3,000	
Network communications	100,000	
SUMEX-AIM collaborative linkages	30,000	
TOTAL DIRECT COST *(Enter on Page 1, Item 5)* →		1,336,864

INDIRECT COST *(See Instructions)*

DATE OF DHEW AGREEMENT:

_____ % S&W*

58 % NTDC August 8, 1979

☐ WAIVED

☐ UNDER NEGOTIATION WITH:

*IF THIS IS A SPECIAL RATE *(e.g. off-site)*, SO INDICATE.

2.1.2 First Year Personnel Detail

		%	Salary	
Project Management				
E. Feigenbaum	Principal Investigator	10		
E. Shortliffe	Co-Princ Invest	10		
T. Rindfleisch	Facility Manager	100		
E. Levinthal	AIM Liaison	25		
C. Miller	Admin Assistant	100		
E. Henderson	Office Assistant	100		
D. Vian	Office Assistant	25		
System Staff				
A. Sweer	System Programmer	100		
F. Gilmurray	System Programmer	100		
M. Bizzarri	System Programmer	100		
M. Achenbach	Syst Prog/User Cons	100		
W. Yeager	Syst Prog/User Cons	100		
R. Tucker	Syst Prog/Opns Mgr	100		
E. Hedberg	Syst Prog - Stud R.A.	62		
J. Clayton	Syst Prog - Stud R.A.	62		
Core Research Staff				
B. Buchanan	Adj Professor	10		
H. Nii	Research Assoc	60		
W. Vanmelle	Research Assoc	50		
N. Aiello	Sci Prog	50		
P. Cohen	Sci Prog - Stud R.A.	62		
D. Smith	Sci Prog - Stud R.A.	62		
J. Kunz	Sci Prog - Stud R.A.	62		
Electrical Engineering Staff				
N. Veizades	Electronics Engineer	100		
E. Schoen	Stud. Electronics Aide	62		
Student Syst Prog/Opns Support				
W. Aviles	Syst Prog - Student	50		
G. Noga	Syst Prog - Student	50		
D. Powers	Syst Prog - Student	50		
C. Robinson	Syst Prog - Student	50		
	Total Salaries		462319	
	Staff Benefits		99645	
********** Total Personnel				561964

BUDGET ESTIMATES FOR ALL YEARS OF SUPPORT REQUESTED FROM PUBLIC HEALTH SERVICE
DIRECT COSTS ONLY (Omit Cents)

DESCRIPTION		1ST PERIOD (SAME AS DETAILED BUDGET)	ADDITIONAL YEARS SUPPORT REQUESTED *(This application only)*					
			2ND YEAR	3RD YEAR	4TH YEAR	5TH YEAR	6TH YEAR	7TH YEAR
PERSONNEL COSTS		561,964	621,220	686,694	767,130	848,623	--	--
CONSULTANT COSTS *(Include fees, travel, etc.)*		--	--	--	--	--	--	--
EQUIPMENT (*)		465,000	280,500	416,025	171,576	132,155	--	--
SUPPLIES		32,000	35,200	38,720	42,592	46,851	--	--
TRAVEL	DOMESTIC	6,000	6,600	7,260	7,986	8,785	--	--
	FOREIGN	--	--	--	--	--	--	--
PATIENT COSTS		--	--	--	--	--	--	--
ALTERATIONS AND RENOVATIONS		--	--	--	--	--	--	--
OTHER EXPENSES		271,900	299,995	326,747	346,433	365,906	--	--
TOTAL DIRECT COSTS		1,336,864	1,243,515	1,475,446	1,335,717	1,402,320	--	--

TOTAL FOR ENTIRE PROPOSED PROJECT PERIOD *(Enter on Page 1, Item 4)* ⟶ $ 6,793,862

REMARKS: *Justify all costs for the first year for which the need may not be obvious. For future years, justify equipment costs, as well as any significant increases in any other category. If a recurring annual increase in personnel costs is requested, give percentage. (Use continuation page if needed.)*

(*) Equipment Purchase items are not included in the Net Total Direct Cost base used to compute Indirect Costs.

(see continuation pages for budget justification)

2.3 Budget Explanation and Justification

The following paragraphs explain in detail our budget plan over the proposed 5-year grant term. Indirect costs are not shown in the budget and will be computed separately on the basis of Net Total Direct Costs (Total Direct Costs less funds for Equipment Purchase). In the most recent agreement between Stanford and the DHEW dated August 8, 1979, the indirect cost rate is 58%.

Personnel

The proposed personnel budget is based on the current staffing for resource management, development, and operations with the addition of a system programmer and an engineering aide to support planned new hardware and software development work. Individual salary figures are not included in the "first year budget detail" plan but have been submitted separately to NIH in confidence. The salary estimates reflect current actual rates and include anticipated increases averaging 10% annually based on recent experience with inflation. Staff benefits are computed using rates currently projected by Stanford University: 21.0% for 8/81, 21.6% for 9/81-8/82, 22.2% for 9/82-8/83, 22.8% for 9/83-8/84, 24.8% for 9/84-8/85, and 25.4% for 9/85-8/86.

Project Management and Technical Direction:

Prof. Feigenbaum is budgeted at 10% as project principal investigator, Prof. Shortliffe at 10% as co-principal investigator for medical liaison (*), Mr. Rindfleisch at 100% is responsible for facility implementation and management, Dr. Levinthal at 25% is responsible for liaison with the national AIM community and the AIM management committees, and Ms. Miller and Ms. Henderson at 100% each provide project administrative and office assistance for SUMEX and community affairs.

System programming:

The programming staff, while sharing a substantial joint responsibility for system development/maintenance, user assistance, subsystem and utility program development, and operational support, have specific areas of responsibility as follows. Messrs. Sweer and Gilmurray, and Bizzarri (100% each) share responsibility for monitor and system support. These duties include, for example, on-going development work for new machine integration into the facility, Ethernet implementation, performance analysis and improvement, system communications support, special device drivers and diagnostics, scheduler controls, and system maintenance. They also share responsibility for system software such as

---

(*) No salary is shown for Dr. Shortliffe for the first 3 years because he is supported by an NLM Research Career Development Award through 6/84. In order to assist his work on the project, we budget 25% support for D. Vian, his office assistant

EXECutive programs, languages, and other general utilities. Mr. Hedberg is a student system programmer who has been working with the project for several years and will continue to work on EXEC developments, network interface software, and software compatibility under supervision of the system staff.

System maintenance and operations:

Mr. Tucker (100%) is responsible for our network liaison, operations utility program development and maintenance, and overseeing system operations and backup. He is assisted in providing file system archive/restore service and backup dumps as well as system utility programming support by the four undergraduate students (currently Messrs. Aviles, Noga, Powers, and Robinson).

User support:

The user support staff includes Mr. Michael Achenbach (100%), Mr. William Yeager (100%), and a student research assistant, Ms. Jan Clayton. Messrs. Achenbach and Yeager will share responsibility for subsystem maintenance and user consulting as well as assisting with software to integrate planned new hardware. Mr. Achenbach also assists in interfacing user program packages into the system (e.g., DENDRAL, MYCIN), assuring appropriate documentation and assisting with initial user contacts. Mr. Yeager serves as the primary contact for user consultation, answering many questions himself and referring others to the appropriate staff members expert in particular areas. Mr. Yeager will also continue development of inter-user communication facilities. Ms. Clayton will be responsible for updating system documentation and developing more effective tools for users to access available documentation.

AI Core Research:

We budget partial support for specific members of the Heuristic Programming Project for core research work to explore basic AI issues relating to biomedical applications and to develop and generalize AI software tools important to the entire SUMEX-AIM community. Complementary support for related work within the HPP is received from other sources such as ARPA and NSF. Prof. Buchanan (10%) will provide technical direction for staff and students working on proposed core research efforts. Ms. Nii (60%) and Dr. Vanmelle (50%) will lead the AGE and EMYCIN efforts respectively. Ms. Aiello (50%) will provide programming support and the graduate research assistants, Messrs Cohen, Smith, and Kunz will work on thesis topics related to particular core research goals.

Electronics support:

Finally we budget Mr. Veizades (100%) and a student engineering aide for hardware engineering and maintenance. They are responsible for designing needed special purpose hardware (e.g., communications equipment, intermachine network hardware, and Ethernet interfaces), integrating new hardware into the facility, and maintaining facility equipment.

Consultant

We do not now plan any consulting support during the follow-on grant period.

Equipment

The "Equipment" budget covers only equipment purchases. Lease arrangements for collaborator terminal and communications support as well as maintenance contracts are discussed under "Other".

Minor Equipment:

$10,000 per year is allocated for minor equipment purchases including communications equipment, Ethernet interfaces, and test equipment. This budget is increased by 5% per year to accommodate inflation.

Major Equipment:

Following are budget estimates for the major equipment acquisitions planned. The prices quoted are best current estimates. Over the 5-year term of the grant prices will certainly change and alternate vendor options may become available for some subsystems. We will carefully review each purchase with BRP to achieve the most advantage in terms of technical and cost effectiveness.

yr 1 - Add 256K words of core to the existing KI-10 AMPEX memory to reduce page swapping overhead. This will cost $65,000 based on a quote from AMPEX for the memory modules and control logic to augment the existing ARM-10LX cabinet.

- Buy a VAX 11/780 with 2M bytes of memory, floating point accelerator, 1 RP-06 disk drive, 1 TE-16 tape drive, and 1 DZ-11 line group at $250,000 based on a current price quotation including tax. This machine will be used to provide large address space INTERLISP facilities, to experiment with AI program export, to support development of VAX system software for the community, and to alleviate congestion in the Stanford 40% of the SUMEX resource. This system has minimal memory for this initial integration work and will be expanded in year 2.

- Buy a bare PDP-11/34 processor with 64K of memory ($18,000), 2 Trident 300 Mbyte disk drives with controller ($49,000), and 2 STC 6250 BPI magnetic tape drives with controller ($53,000) to develop a community file server. This file server will be 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.

- $20,000 is allocated for a "Stanford University Network" bit-mapped display terminal station ($10,000) and a Canon laser printer for high quality hardcopy output ($10,000).

yr 2 - Add 2M bytes of memory to the VAX purchased in year 1 ($70,000).

- Add 630M bytes to the file server purchased in year 1 ($40,000). This will include 2 300 Mbyte drives which will fill the controller.

- Buy 5 single-user "professional workstations" (PWS) ($160,000 -- $30,000 each plus tax). This price is based on the projected cost of the Zenith-MIT NU system or its equivalent. These machines will be used to develop and experiment with user-dedicated machines for AI program development, export, and human interface enhancements. These machines will be distributed within the Stanford community initially to facilitate development and will be coupled by Ethernet with the main resource.

yr 3 - Add a second VAX 11/780 with 4 Mbytes memory, 1 RP-06 disk drive, 1 TE-16 tape drive, floating point accelerator, and 1 DZ-11 line group ($320,000) for general community support with large address space INTERLISP. This machine will be managed for program testing in a way similar to the existing 2020.

- Add 2 PWS systems ($65,000) to be distributed within the AIM community under Executive Committee control.

- $20,000 is allocated for an additional "Stanford University Network" bit-mapped display terminals ($10,000) and a Canon laser printer for high quality hardcopy output ($10,000) for the anticipated growing and distributed community of local users.

yr 4 - Add 3 PWS systems ($100,000) to be distributed within the AIM community under Executive Committee control.

- Add 630M bytes to the central file server to meet expected growth in community file storage needs. This will include a second controller with two drives ($60,000)

yr 5 - Add 3 PWS systems ($100,000) to be distributed within the AIM community under Executive Committee control.

- $20,000 is allocated for an additional "Stanford University Network" bit-mapped display terminals ($10,000) and a Canon laser printer for high quality hardcopy output ($10,000) for the anticipated growing and distributed community of local users.

Supplies

The computer supplies budget is an extension of our recent operating experience with the SUMEX-AIM facility and expected increases for the new machines. We estimate $12,000 for the first year covering paper, ribbons, tapes, disk packs, labels, and other supplies. We budget a 10% per year escalation of these costs. Office supplies are budgeted at $5,000 per year also based on past experience and are increased 10% per year. Engineering supplies cover needed parts and spares for interfacing and integrating new equipment and for maintaining in-house equipment. We budget $15,000 per year for this purpose with an annual inflation factor of 10%.

Travel

The travel budget covers travel to technical meetings, management committee meetings, and AIM workshop meetings as well as travel to assist user groups get started on SUMEX as needed. We budget for 4 east coast trips ($800 each), 3 midwest trips ($600 each), and 4 west coast trips ($250 each). Future years are inflated by 10% per year.

Other

Equipment Maintenance:

We budget for facility equipment maintenance based on our past experience with DEC and other vendors. We expect to retain our favorable cooperative maintenance arrangements with DEC for the KI-10 and 2020 systems and to add appropriate vendor contracts for the other equipment (VAX's, file server, Professional workstations, etc.) as acquired. We spend substantial staff effort in maintaining equipment to minimize costs in contracts and "time and materials" to outside vendors. We continue to investigate alternatives for maintenance: either in-house or from another vendor. So far we have not been able to project enough cost savings or improved service to justify a change. With costs continuously rising, we will periodically re-evaluate alternatives to achieve the most cost effective maintenance service for the resource. We have budgeted a 5% per year inflation for maintenance costs.

Equipment Lease:

We budget $3,000 per year for equipment lease related to on-going collaborative linkages to SUMEX. $2,000 per year is allocated for continued lease of a communication line between the SUMEX machine room and the SECS facilities at the University of California at Santa Cruz. $1,000 per year is for a line to Prof. Langridge's group at UC San Francisco. These lines were approved by the AIM Executive Committee.

Telephone Services:

We budget $7,500 per year for staff office and home terminal telephones and $10,000 per year to cover dataphone services for local Stanford community dialup ports on the SUMEX computer. These estimates are based on the current configuration of lines and expected growth for planned new equipment. We periodically review these arrangements to maintain satisfactory service at minimum cost.

Software Lease:

We budget $6,000 per year for software lease costs. These funds are used to maintain our license rights to and updates for such software as DEC monitors, language and utility products, SITBOL, STP, SPSS, SIMULA, etc. as well as additional packages the community may require.

Services and Documentation:

$4,000 per year is budgeted for books, publications, technical services, and reproduction based on previous experience. $3,000 per year is budgeted for providing to users up-to-date documentation for system and subsystem usage. Substantial efforts continue to upgrade documentation for the user community.

Communications support:

We budget a total of $100,000 per year for network services starting in year 1 and increased by 5% per year. Of this amount, $75,000 is allocated based on current experience for TYMNET services (including network interface, maintenance, and usage costs) projected to accommodate increased usage for the new equipment. In past years, these funds have been distributed directly from NIH/BRP through NLM contracts with TYMNET. This may still prove to be the most cost-effective approach and we will work closely with NIH/BRP to secure these critical services at the lowest cost.

The remaining $25,000 is budgeted as a contingency to experiment with other networks or communications media to support AIM work if justified by community needs and technological developments or to retain our highly beneficial ARPANET connection. A growing number of the AIM community members with local machines have expressed the need for a means to transfer files with SUMEX. This need will increase with more distributed AIM computing resources. Since TYMNET is not currently moving to provide this kind of service, further experimentation with TELENET or other vendors may be warranted.

At present SUMEX-AIM ARPANET costs are being borne by ARPA-IPTO as part of the Stanford Heuristic Programming Project contract. We have no information that this relationship will change (we do get frequent inquiries from ARPA about its status however). The $25,000 contingency may be needed to cover part of these costs should ARPA/DCA policies changes.

Collaborative Linkages:

We budget $30,000 per year for collaborative linkage needs. These funds will be available for terminals, lines, and other facilities to enable more effective inter-group collaborations and contacts with medical scientists. These funds have been very effective in the past in assisting new projects get connected to available computing resources within the AIM community pending grant support of their research. These funds are allocated in close cooperation with the AIM Executive Committee and BRP. We budget a 5% annual increase for this collaborative linkage support.

II. Research Plan

This is an application for renewal of a grant supporting the Stanford University Medical EXperimental computer research resource for applications of Artificial Intelligence in Medicine (SUMEX-AIM). We have attempted to keep this proposal as brief as possible and to place detailed background information in appendices. However, we felt obliged to exceed some of the page limitations stipulated in the NIH guidelines for a several reasons:

1) the computer science discipline of artificial intelligence is relatively new and its intersection with and significance to medicine requires more explanation than more traditional areas of biomedical research.

2) the SUMEX-AIM resource encompasses a national community of more than 20 research projects pursuing diverse applications areas. In order to illustrate the scope of the community and to provide the scientific basis for continued support of SUMEX as a resource, the objectives of these projects must be presented. We also include a brief description of the important operational base of the resource that may be unfamiliar to some reviewers.

3) this application is for a 5-year renewal term. Many of the core and collaborative research efforts are aimed at long term goals to assist biomedical researchers and clinicians in information management, analysis, and decision making. In order to provide a more efficient research environment, avoiding the overhead of additional proposal preparations and reviews on time scales shorter than expected result horizons, we hope to describe our goals in sufficient detail to justify the 5-year award period.

Specific Aims

3 Introduction and Aims

3.1 Overview of Objectives and Rationale

The SUMEX-AIM ("SUMEX") project is a national computer resource with a dual mission: a) the promotion of applications of computer science research in artificial intelligence (AI) to biological and medical problems and b) the demonstration of computer resource sharing within a national community of health research projects. The SUMEX-AIM resource 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.

In the body of this proposal, we offer definitions and explanations of these efforts at several levels of detail to meet the needs of reviewers from various perspectives. For this overview, we give only a brief definition of AI and a summary of the background, present status, and expectations of our research for the requested term of the renewal, the five years beginning August 1, 1981.

3.1.1 Definitions of 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 manner of use of computers by people 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 yet with 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

---

(*) Brief abstracts of the various projects can be found in Appendix A on page 331 and more detailed progress summaries in Section 9 on page 135.

special domains they may be of comparable or greater power, e.g., in the solution of formal problems in organic chemistry.

3.1.2 Resource Sharing

An equally important function of the SUMEX-AIM resource is an exploration of the use of computer 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. 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.

Even in their current developing state, communication facilities enable effective access to the specialized SUMEX computing environment from a great many areas of the United States and to a more limited extent from Canada, Europe, Australia, and other international locations.

3.2 SUMEX-AIM Background

Beginning in the mid-1960's with DENDRAL (**), a project focused on applications of artificial intelligence to problems of biomolecular

---

(*) Very Large Scale Integration

(**) Much of the early DENDRAL computation work was done on the ACME IBM 360/50 interactive computing resource at Stanford, which was funded by the NIH Biotechnology Resources Program between 1965 and 1973.

structure characterization, the Stanford Heuristic Programming Project has pioneered in expert systems research with funding support from NIH, ARPA, NSF, and NASA. Since 1973, SUMEX-AIM has developed as a national resource for applying these techniques to a broad range of biomedical research problems.

Funding of the SUMEX-AIM resource from the NIH Biotechnology Resources Program (BRP) began in December 1973 for a five year period. Prof. Joshua Lederberg was Principal Investigator and Prof. Edward A. Feigenbaum was co-Principal Investigator. The major hardware was delivered and accepted in April 1974, and the system became operational for users during the summer of 1974. In 1977, we applied for a five-year renewal grant to continue our national research effort. We received a recommendation for approval of the five year period from the study section but this was reduced to three years following Professor Lederberg's decision in early 1978 to accept the presidency of The Rockefeller University. The principal investigator role passed easily to Prof. Feigenbaum, Chairman of the Stanford Computer Science Department, based upon his long-time involvement with the project and close collaboration with Prof. Lederberg. The highly interdisciplinary spirit of SUMEX has been retained with very close ties to the Stanford Medical School through Drs. E. H. Shortliffe (current co-Principal Investigator of SUMEX) and S. N. Cohen.

Although six years is hardly long enough for a conclusive determination of the success of the SUMEX-AIM model, we can fairly take pride in the diligence and technical competence with which we have responded to the community responsibilities mandated by the terms of our grant. An important element in satisfying those responsibilities was the establishment of a mutually satisfactory management structure, on which we report in further detail later (see Appendix E on page 383). Good will and common purpose are of course the indispensable ingredients for an effective community resource, and we are grateful to have been able to offer this service in a congenial framework, and at the same time to be able to support our local computing research needs.

The present renewal application is therefore written from a perspective of having built a substantial community of active biomedical AI research projects and having just begun the new phase of our research to integrate and exploit emerging computer technologies that will have a profound effect on the development and export of practical medical AI programs. Beginning with 5 projects in 1973, the AIM community grew to 11 major projects at our renewal in 1978 and currently numbers 17 fully authorized projects plus a group of 8 pilot efforts. In addition to the Rutgers Computers in Biomedicine project, two of the formal projects and one of the pilots do their computing using the portion of the Rutgers University facility allocated to AIM community users. As discussed in the sections describing the individual projects (see Section 9 on page 135), 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.

3.3.1 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.

3.3.2 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.

3.3.3 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. The details of these are given in Section 6.3.

4 Significance

What is the significance of the artificial intelligence research and knowledge engineering work for which SUMEX is a resource? And what is the significance of SUMEX for achieving the goals of the enterprise?

In this section, we first sketch, in an abstract way, the significance of the scientific work. We then probe more deeply examining medicine, biochemistry, and psychology. Finally, we look at SUMEX's facilitative role, particularly in the light of the microelectronic revolution; and conclude with a discussion of the more general aspects of SUMEX's scientific role in enhancing scientific communication and knowledge.

A Brief Recapitulation

Artificial Intelligence research and its applications-oriented twin, Knowledge Engineering, are those parts of Computer Science that are concerned with the representation of symbolic knowledge for computer use; and the construction of programs for symbolic inference that can make use of the knowledge to achieve intelligent action. Examples of such actions include finding problem solutions, forming hypotheses, offering advice, inferring diagnoses, recommending therapeutic steps, and so on. 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.

Managing the Growth of Knowledge

Medical and scientific communities currently face many 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

These needs are widely recognized. In addition, computers are recognized as the most hopeful technology to overcome the problems. While recognizing the value of mathematical modeling, statistical classification, decision theory and other techniques, we believe that effective use of those 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.

More specifically...AI in the service of Medicine

Although computing technology is playing an increasingly important role in medicine, systems designed to advise physicians on diagnosis or therapy selection have received poor clinical acceptance. Despite diverse research efforts, and a literature on computer-aided diagnosis that has numbered at least 1000 references in the last 20 years, clinical consultation programs have seldom been used other than in experimental environments.

The reasons for attempting to develop such systems are self-evident. Growth in medical knowledge has far surpassed the ability of the 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.

In recent years observers have begun to analyze the reasons for poor acceptance of the systems that have sprung from such research, and some have argued that the problems have tended to lie not only with the decision-making performance of such programs but also with system design features that have failed to appreciate the physician's viewpoint or have made the interactive process unappealing. To correct,these deficiencies future systems must be fast, easy to use, and congenial. They must address important clinical problems with which physicians recognize they need assistance. But perhaps most important, in order to stress the primary physician's role as ultimate decision maker, they must be able to explain what they are doing, not through quotations of statistical theory but in terms of a line of reasoning that is familiar and similar to the kind of justification a clinician might expect from a human consultant. Explanation capabilities help the physician using the program decide whether to follow its advice; they thereby emphasize the computer's function as a helpful tool that is intended to complement rather than replace the primary physician's own decision-making powers.

Because of considerations such as these, the last decade has witnessed the development of new approaches to computer-based medical decision making. Of particular significance is research directed at the encoding and utilization of experts' judgmental knowledge -- the kind of practical experience which underlies the daily practice of medicine and is far-removed from the mathematical approaches of formal decision analysis. Artificial Intelligence is a particularly relevant computer science subfield because of its emphasis on symbolic reasoning capabilities rather than numeric computations. The AIM community's promising research into medical symbolic reasoning represents more than the application of well-established computing techniques. Although the approaches are young and experimental, significant accomplishments in codifying medical knowledge and modeling clinical reasoning have already been achieved. Additional investigation, in artificial intelligence and in related computer science subfields, will further facilitate the development of useful, congenial, high-performance consultation systems. These systems will improve when we know better how to manage such problems as (1) understanding the psychology of medical reasoning as practiced by specialists, (2) automated interpretation of written and spoken natural language, (3) acquisition and representation of knowledge obtained from collaborating experts, (4) encoding and utilization of time relationships central to many disease processes, and (5) mechanisms for representing and measuring inexact reasoning.

AI in the service of Biochemistry: why SUMEX?

Consider three major projects engaged in research in structural biochemistry:

1) DENDRAL, computer-assisted elucidation of molecular structure, including stereochemistry, with applications in the areas of natural products, bio-active compounds and conformational analysis

2) MOLGEN, investigations of experiment planning in molecular genetics, including structural studies of large biomolecules with emphasis on sequencing of nucleic acids

3) SECS, computer simulation and evaluation of chemical synthesis

In each case, a new type of computational assistance is being made available to a significant modern area of scientific research. Though in the past each field has made some use of the numeric and searching capabilities of computers, the use of advanced methods for symbolic manipulation, representation of knowledge, and inference is new, currently significant, and holds great promise in future development.

Over the past several years all three projects have matured to the point where specific programs are being disseminated to the scientific community via the mechanisms of outside access to SUMEX or direct program export to other laboratories. Each project is currently engaged in studies pointed toward both application of existing programs to real biochemical problems and research into new computer-based tools for future applications. The SUMEX resource provides a focal point for building a collaborative community with common interests in particular programs. The resource provides the computational capacity for new developments and a medium for communication for discussions of successes, and failures, aimed at improving application programs.

The rapid development of these programs, to the point of sharing the programs with a community of investigators, is due to several factors. These factors are important in understanding the special significance of the SUMEX resource and the role it plays in continued development and dissemination of the programs. All three projects share an important underlying thread, and that is the concept of a molecular structure. Even though the three projects deal with computer representations of molecular structures at varying levels of specificity, the fact that there are formal, precise descriptions of structure available greatly facilitates subsequent computer manipulation of the representations. A significant part of the structural manipulations which must take place can be treated algorithmically. Development of such algorithms has reached a highly sophisticated state; these developments represent a strong foundation on which to build subsequent procedures which rely on judgmental knowledge, or rules, to arrive at scientifically meaningful conclusions.

The "knowledge engineering" aspects represent a set of similar problems in system design shared by all three projects. Here the concept of community building and sharing of ideas, factors inherent in SUMEX as a resource, play an essential role in allowing the projects to learn from one another and from AI programs in other major areas.

The biochemistry projects have as a common goal the development of interactive programs which act as problem-solving assistants to an investigator. In order to be useful to a wide community, such programs must be capable of assisting in the solution of a variety of real scientific problems. Here SUMEX is indispensable. The resource provides many facilities for access to programs, for recording of terminal sessions, for rapid exchange of messages about problems and their solutions, and for development and export of versions of programs for use in other laboratories.

Using the DENDRAL project as a concrete example, SUMEX has been used for program development and application to many structural problems of the DENDRAL group and their collaborators throughout the country. Export of the CONGEN program began about eight months ago and already eighteen copies of the program have been distributed to other laboratories. SUMEX will continue to be used for development and for exposure of several new programs (adjuncts to or successors of CONGEN) to structural problems here at Stanford, with export taking place after developing confidence in the programs. In addition, new research projects have been undertaken with a small number of collaborators. These persons are interested in development of new techniques for structural analysis, especially in the area of stereochemistry. Network access to SUMEX has been provided so that development of the techniques themselves will take place at one central facility, with the message system providing the primary means of communication between DENDRAL project members and their collaborators. Specific structural problems, for example the conformational studies of Dr. Cowburn at Rockefeller University, come from the collaborators and exemplify the type of problem which the programs must be capable of solving in order to be useful to the community of persons engaged in related research.

Another example: AI methods in Psychology

The orientation of AI research toward the construction of intelligent agents -- known as "knowledge engineering" -- has always coexisted with an orientation toward the explication and understanding of human cognitive behavior viewed as information processing. Indeed the marriage of AI models and methods with the problems and techniques of Cognitive Psychology has been so fruitful that a field with its own name, society, and journal has been born thereof: Cognitive Science.

Since the health research community has long been a supporter of basic research in Cognitive Psychology through the NIMH, it has been appropriate that this branch of AI be supported by SUMEX. The gains thereby have been perceived to be so significant that the Cognitive Science field is itself now considering the establishment of a network-based community, for which SUMEX is one of the leading two models.

The significance of the AI methodology to the modeling of cognitive processes has always been seen as:

> precision of expression...computer programming languages are not only ideally suited for expressing the elementary information processes of the model and the postulated data structures, but admit no vagueness or incompleteness.
>
> complexity...the difficulty of managing the modeling process does not go up significantly as the model becomes richer (more complex); thus the methodology does justice to the complexity of human cognitive processes, does not force oversimplifications.
>
> testability...though the models are complex, the computer will generate in detail the remote consequences of the modeling assumptions for particular situations; thus the models are as testable and correctable, in principle, as any in the "hard" sciences.

In recent years, SUMEX-AIM has been one of the most significant forces impelling the forward motion of cognitive science. It has allowed the building of geographically dispersed communities around a single modeling effort; and it has reduced the "cost of entry" to this methodology.

The best example relates to the ACT model of human long-term associative memory, initially constructed by John Anderson. This elegant model has been explored, modified, and tested by a subcommunity of psychologists who gain access to it by the normal simple SUMEX-AIM procedures (bypassing the laborious process, sometimes impossible to achieve, of "bringing it up" at their own sites). As another example, Professor Kintsch and his group at the University of Colorado were able, on the second day of a visit by two Stanford researchers, to begin the process of using the Stanford-SUMEX-developed system, AGE, to model human story comprehension.

What is the GENERAL SIGNIFICANCE of SUMEX-AIM?

As a Research Resource...

SUMEX-AIM is widely viewed as a model national computing resource. Its service has been wide-ranging, in terms of user help and variety of software services provided; reliable; economical on a per-user or per-project basis; and effective in promoting the healthy growth of its research community. It is being studied by communities of scientists in molecular biology (both in the U.S. and Europe) and in cognitive science as a model of how to provide similar service to their sciences; and the term "SUMEX-like facility" was common in planning discussions for the National Center for Computation in Chemistry and for a proposed ARPA national computing resource for ARPA-sponsored DOD projects.

As an experiment in community building...

Lederberg's original vision extended far beyond the "resource" mandate. He said, in an earlier SUMEX renewal proposal,

> "We infer that many fields of scientific inquiry will have to use similar methods of exchange of critical commentary; that the electronic communications of computer programs is a prototype for the maintenance of other knowledge bases essential for the fabric of a complex and demanding society. The computer is at one time the node of a knowledge-sharing network, and the device for verifying the consistency and pertinence of the updates and criticisms that the users remit. Thus we can view our resource as exemplifying a technology that induces a new social organization of scientific effort."

SUMEX-AIM has been remarkably, though not uniquely, successful in pointing to this new direction for scientific integration and cumulation. The collection of computer science research centers on the ARPANET represents another example, but because the goals of SUMEX are more focused, its achievements at community building are more easily defined. The speed with which the relatively new MOLGEN programs are making their way into the relevant scientific community, by means of help from and access to SUMEX, is gratifying evidence of the community building spirit and technique of the resource. That this path cut by SUMEX in the '70s will become the highway of the 80's and '90s is very likely.

As a focus for the development of the inexpensive "intelligent assistant" in medicine and the biosciences...

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 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, in 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 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 such reasoning capabilities and knowledge. The researchers of the SUMEX-AIM community currently constitute a large fraction of all the computer scientists whose work is aimed at this inevitable development.

The day is not far off. There appeared in Business Week, April 14, 1980 an article on INTEL and their plans for the 1980's. INTEL is presently fourth in integrated circuit sales but is on a much faster growth curve than its competitors. Therefore its plans should be an important indicator of the technological environment to be expected in this coming decade.

INTEL's plans include a "minimainframe" more powerful than any chip computer so far announced, which includes the ability to be linked in networks for even higher performance. INTEL is investing about $100 million in software for a full-fledged operating system with capabilities in language understanding, mechanization of intellectual activity, pattern recognition etc..

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 VAX-like and $30,000 "professional workstation" computers have the practical possibility of large-scale and low-cost use because of these anticipated near-term industrial developments.

As a focus for the methodology that will explicate and disseminate the "private" -- heuristic -- knowledge of practice...

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.

Lederberg's statement from our previous proposal rounds out this larger view:

> "Although our substantive efforts are mostly concerned with the 'micro-problems' of scientific or clinical inference, there may be more important treasures in a macro-perspective on the integration of knowledge in medicine. I believe that it is reasonable to expect that the systematization of biomedical knowledge, to which computer AI will make an indispensable contribution, is an important side effect of these investigations in knowledge-engineering; and that this will lead in turn to the recognition of holes in the overall fabric that badly need patching. We have too little theory of the practice of science to offer more than case studies at this time."

5 Progress

This report covers only the resource nucleus; objectives and progress for individual collaborating projects are discussed in their respective reports in Section 9 beginning on page 135. These projects collectively provide much of the scientific basis for SUMEX as a resource and our role in assisting them has been a continuation of that adopted for the first grant term. Collaborating projects are autonomous in their management and provide their own manpower and expertise for the development and dissemination of their AI programs.

5.1 Brief Statement of Prior Goals

The following summarizes SUMEX objectives for the on-going three year grant, begun on August 1, 1978. It will be noted that the high-level goals for this work closely parallel those for the renewal period. These are the continuing basis for our long-term program in biomedical AI research and are resummarized here to comply with the requested NIH form for this proposal. Changes to previous detailed objectives because of explicit guidelines and funding limits in the council award are noted below.

5.1.1 Resource Operations

1) Continue the building of a community of projects applying AI techniques to medical problems including improving mechanisms for inter- and intra- group collaborations and communications.

2) Provide an effective computing resource to support the development and research dissemination of biomedical AI computer programs for a broad range of applications areas.

3) Provide effective and geographically accessible network communication facilities to the SUMEX-AIM community for remote collaborations, scientific communications, and experimentation with developing AI programs.

5.1.2 Training and Education

1) Provide documentation and assistance in interfacing users to resource facilities and programs.

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 technical workshop activities in collaboration with the Rutgers Computers in Biomedicine resource and individual application projects.

We had proposed support for a "visiting scientist" position to allow prospective qualified SUMEX-AIM project investigators or users to spend a term in close contact with on-going research work. Funding for this position was cut by the NIH review committees.

5.1.3 Core Research

1) Continue to encourage community efforts at organizing and developing AI techniques by supporting projects such as the AI Handbook, special language developments, and other projects community members may propose to contribute.

2) Explore generalizations of AI tools for knowledge acquisition, representation, and utilization.

3) Explore AI software implementation and export mechanisms such as machine-independent languages and special purpose computer systems. This includes the continued development of the MAINSAIL system and the investigation of satellite general purpose machines capable of running existing systems.

Because of guidelines and funding limits in the council-approved award, we removed several goals in the core research work as originally proposed including support for development of a general planning package, a heuristic knowledge acquisition system, and a general explanation system. We were also forced to limit the goals of the MAINSAIL effort to the completion of the language design and to a demonstration of implementations for five target systems. No export efforts for MAINSAIL or work on microprogrammed implementations were possible.

5.2 Summary of Progress: 11/77 - 4/80

1) We have continued to recruit a growing community of user projects and collaborators. The initial complement of 5 projects has grown to 17 fully authorized projects currently plus a group of 8 pilot efforts in various stages of formulation. Several of these projects use the AIM computing facility at Rutgers. 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.

2) SUMEX user projects have made good progress in developing and disseminating effective consultative computer programs for biomedical research. These performance programs provide expertise in analytical biochemical analyses and syntheses, medical diagnoses, and various kinds of cognitive and affective psychological modeling. We have worked hard to meet their needs and are grateful for their expressed appreciation. [see Section 9 beginning on page 135].

3) A first version of the AGE system has been completed. It uses the "blackboard model" control structure for coordinating multiple expert sources of knowledge for the solution of problems. The UNITS package [9] for a "frame-oriented" representation of knowledge is now being incorporated. AGE provides a general structure and an interactive facility for implementing knowledge-based systems. A workshop to introduce AGE to the AIM community was held at Stanford in February 1980. [see Section 9.1.1 on page 137].

4) We have completed the initial phases of a systematic effort to document AI concepts and techniques through the AI Handbook Project. It comprises a compendium of short articles about the projects, ideas, problems, and techniques that make up the field of AI. The first two volumes covering heuristic search, knowledge representation, natural language and speech understanding, AI languages, various applications domains, and automatic programming were completed in August 1979 and publication plans are in progress. All completed sections have been published as Stanford Computer Science Department technical reports. Work on a third volume is progressing well. [see Section 9.1.2 on page 145 and Appendix G on page 392]

5) We successfully completed the design and a demonstration of the MAINSAIL language system as a tool for software portability. A common compiler, code generators, and runtime support for TENEX, TOPS-10, TOPS-20, RT-11, and RSX-11 have been developed as part of this demonstration system and numerous applications programs written by collaborating research groups. Further work past this demonstration phase will be done independently of SUMEX through a private company, XIDAK, formed to continue the development, dissemination, and maintenance of MAINSAIL. Work is under way to develop MAINSAIL for the VAX and a number of other target machines. [see Appendix H on page 398].

6) We have continued refinement of the SUMEX facility hardware and software systems. We have worked to enhance throughput, to better control the allocation of resources among communities, to increase efficiency, to enhance human interfaces, to improve documentation, and to extend the range of software facilities available to user projects. We also completed installation and evaluation of a connection to TELENET as an alternate source of communications services for our community.

7) We completed planning and implementation of a satellite machine that supports more operational demonstrations of mature AI programs and helps alleviate system congestion for on-going program development. This acquisition of a DEC 2020 system was reviewed and approved by an ad hoc study section. We have installed the machine and are actively working on its integration into KI-10 facility by means of a local Ethernet [10]. Using an interim connection, it has been used extensively for workshops and program demonstrations.

8) We have smoothly completed the management transition. On July 1, 1978, Prof. Edward Feigenbaum assumed the role of SUMEX Principal Investigator following Prof. Joshua Lederberg's installation as president of The Rockefeller University. Prof. Lederberg continues to maintain close ties with SUMEX activities as chairman of the SUMEX-AIM Executive Committee. Close coordination of project activities with medical research is provided by Dr. E. H. Shortliffe, co-Principal Investigator of SUMEX. Dr. Shortliffe is Assistant Professor of General Internal Medicine and one of the key developers of the MYCIN system. Effective August 1, 1980, SUMEX will become part of the Department of Medicine where it will be centered in the largest clinical department of the Stanford Medical School. Previously, SUMEX had been in the Department of Genetics with Prof. Stanley Cohen, Dr. Lederberg's successor as chairman, assisting in project medical coordination.

5.3 Detailed Progress Highlights

The following material highlights in more detail SUMEX-AIM resource activities since the last review in the context of the resource staff and the resource management.

5.3.1 Resource Operations

Our core facility, initially installed in March 1974, is built around a Digital Equipment Corporation (DEC) KI-10 computer and the TENEX operating system. This facility has provided a superb base for the AI mission of SUMEX-AIM in terms of its interactive computing environment, its AI program development tools, and its network and interpersonal communication media. Biomedical scientists have found SUMEX easy to use in exploring applications of developing artificial intelligence programs for their own work and in stimulating more effective scientific exchanges with colleagues across the country.

These tools also give us access to a large computer science research community, including active artificial intelligence and system development research groups. Coupled through effective network facilities, these groups greatly enhance the SUMEX-AIM community environment through broader scientific interchange and software sharing.

Following are highlights for recent developments in various aspects of the facility. Detailed information about SUMEX loading can be found in Appendix B on page 355. Plots are given there for overall resource usage, diurnal loading, community/project usage, and network traffic.

5.3.1.1 System Hardware

1) Implemented a number of strategic facility augmentations over the years in response to growing community needs to increase system capacity and improve performance for interactive expert systems. These include: (3/74) - install KI-10 with 192K words of memory; (11/74) - add 64K words of memory; (5/76) - add second KI-10; (8/77) - add 256K words of memory and double on-line file space (see Figure 1 for a current configuration diagram).

2) Acquired a software-compatible satellite DEC 2020 computer as a dedicatable resource for improved interactive response for experimental testing of AI programs. This relatively inexpensive machine ($175,000) includes a KS-10 processor approximately half the speed of a KI-10, 512K words of memory, 1 disk and 1 tape drive, 16 terminal lines, and software license (see Figure 2 for a configuration diagram}). It runs TOPS-20 and is for the most part software-compatible with the KI-TENEX system. The 2020 was installed without problem in August 1979 and we have supported many program demonstrations on it for the DENDRAL, ONCOCIN, AGE, SECS,

INTERNIST, and MOLGEN projects. Major conferences for which the 2020 has been used include the Sixth International Joint Conference on AI from Tokyo, Japan in August 1979 and, most recently, the American College of Physicians meeting in New Orleans in April 1980.

3) Began implementation of a local Ethernet [10] as the basis for integrating the KI-10 facility with the 2020 and future planned hardware. Based on Xerox-developed protocols, this system will connect SUMEX resources through a 3.3 Mbit/sec network to allow uniform terminal access, file transfers, peripheral equipment sharing, and remote resource access through gateways. Figure 3 on page 38 shows current configuration plans for the SUMEX network. The KI-10's are fully operational on the Ethernet through an interim I/O bus PDP-11 interface. This uses a Xerox-designed PDP-11 interface board and an adaptation of their higher level software. The 2020 is connected electrically through its UNIBUS adapter. We are working to complete the 2020 connection software and to design a direct memory interface for the KI-10's to achieve higher performance and efficiency. [see Appendix C on page 374 for details].

4) We have designed and implemented communications control hardware to allow sensing of carrier drop on dial-up lines so that attached jobs can be detached to prevent users from inadvertently connecting to hanging jobs. We also implemented a software-controlled switch to allow more efficient use of available terminal scanner ports on the system. Hardwired and leased line connections no longer tie up scanner ports when not in use.

5) We have supported community hardware communication needs by installing and maintaining local terminals and connections; assisting in the acquisition and installation of terminals at remote user sites; assisting with dedicated links to remote user sites (e.g., UC Santa Cruz and UC San Francisco); and assisting with equipment installation for AI program demonstrations.

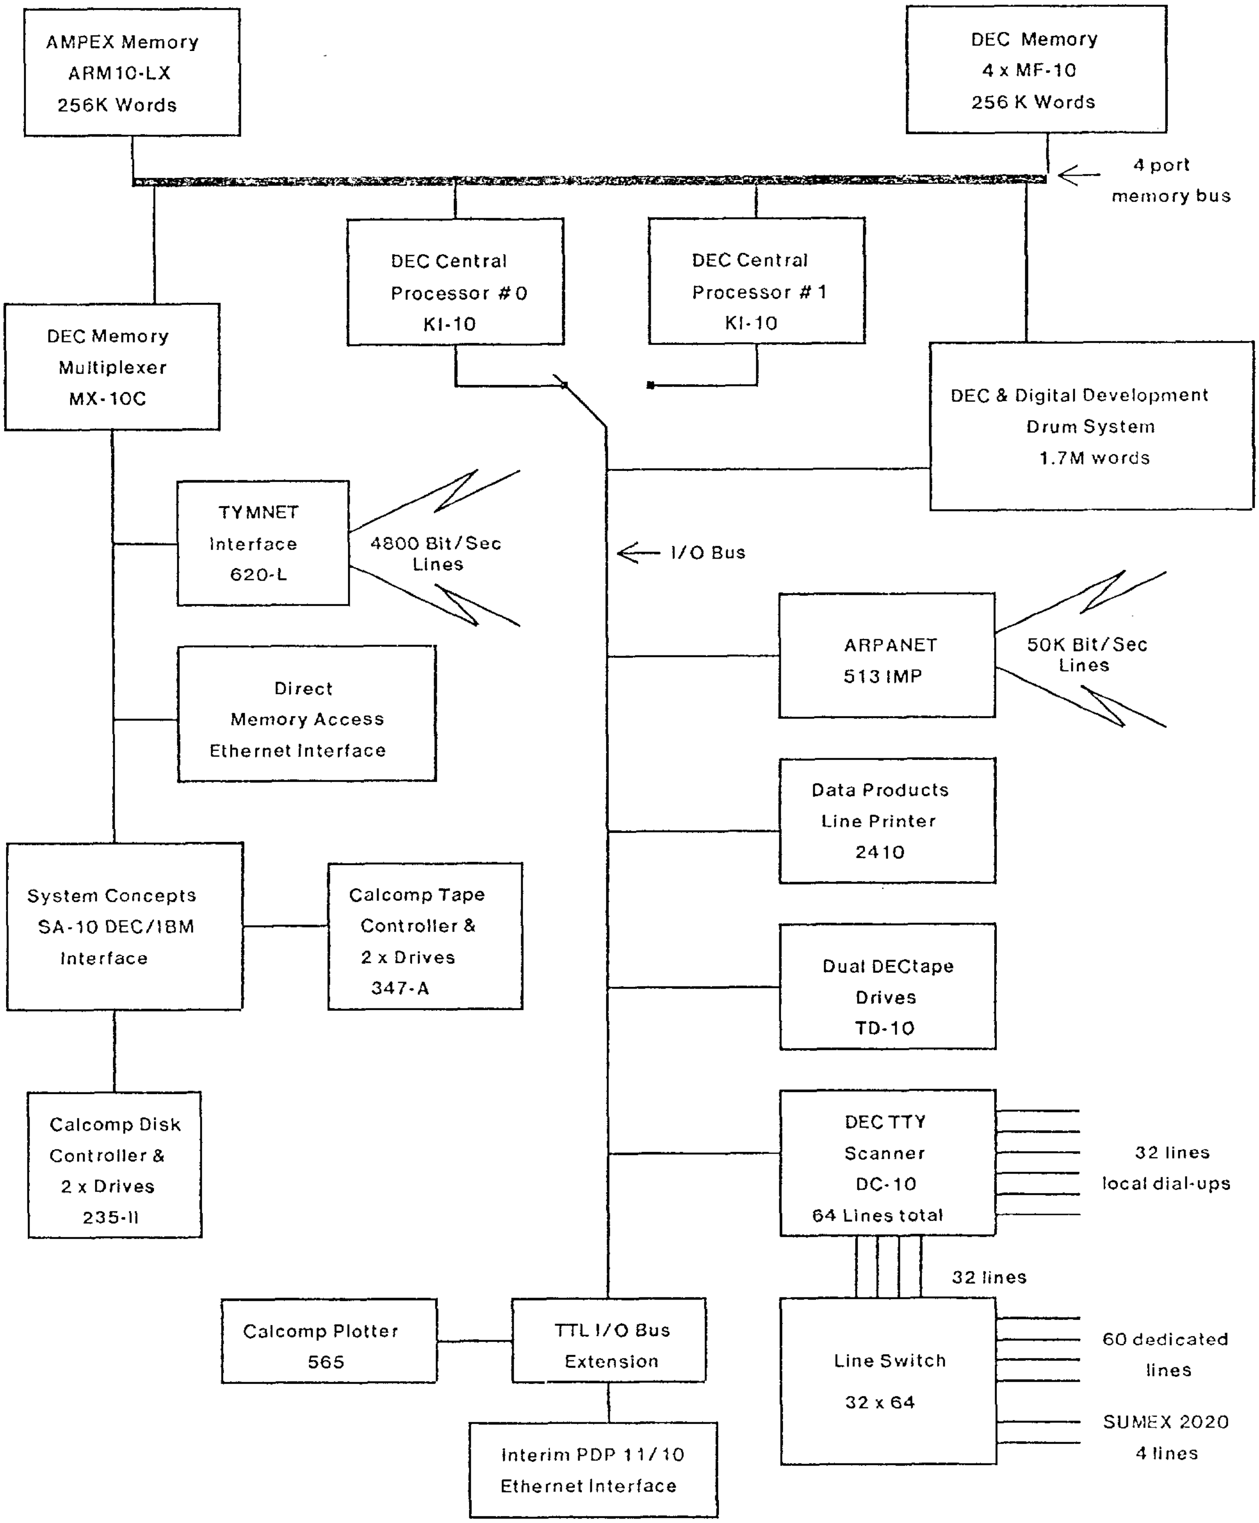
AMPEX Memory
ARM10-LX
256K Words
DEC Memory
4 x MF-10
256 K Words
4 port
memory bus
DEC Central
Processor # 0
KI-10
DEC Central
Processor # 1
KI-10
DEC Memory
Multiplexer
MX-10C
DEC & Digital Development
Drum System
1.7M words
TYMNET
Interface
620-L
4800 Bit/Sec
Lines
I/O Bus
ARPANET
513 IMP
50K Bit/Sec
Lines
Direct
Memory Access
Ethernet Interface
Data Products
Line Printer
2410
System Concepts
SA-10 DEC/IBM
Interface
Calcomp Tape
Controller &
2 x Drives
347-A
Dual DECtape
Drives
TD-10
Calcomp Disk
Controller &
2 x Drives
235-II
DEC TTY
Scanner
DC-10
64 Lines total
32 lines
local dial-ups
32 lines
Calcomp Plotter
565
TTL I/O Bus
Extension
Line Switch
32 x 64
60 dedicated
lines
SUMEX 2020
4 lines
Interim PDP 11/10
Ethernet Interface

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

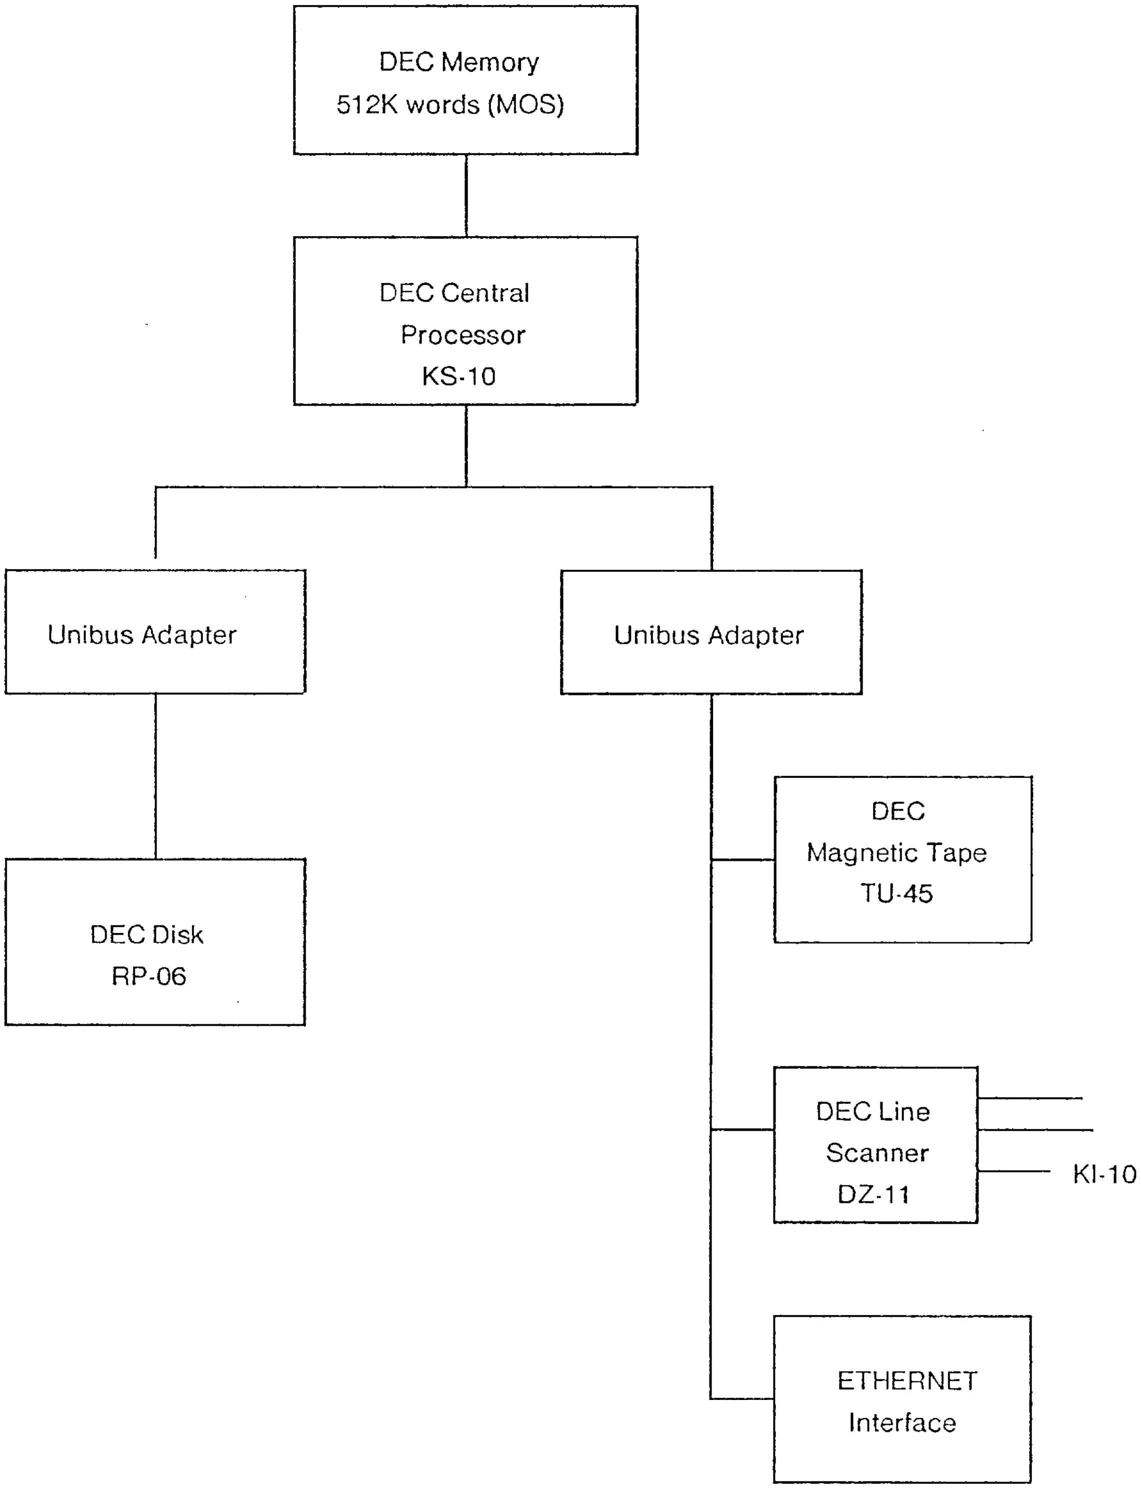
DEC Memory
512K words (MOS)
DEC Central
Processor
KS-10
Unibus Adapter
Unibus Adapter
DEC Disk
RP-06
DEC
Magnetic Tape
TU-45
DEC Line
Scanner
DZ-11
KI-10
ETHERNET
Interface

Figure 2. Current SUMEX-AIM 2020 Computer Configuration

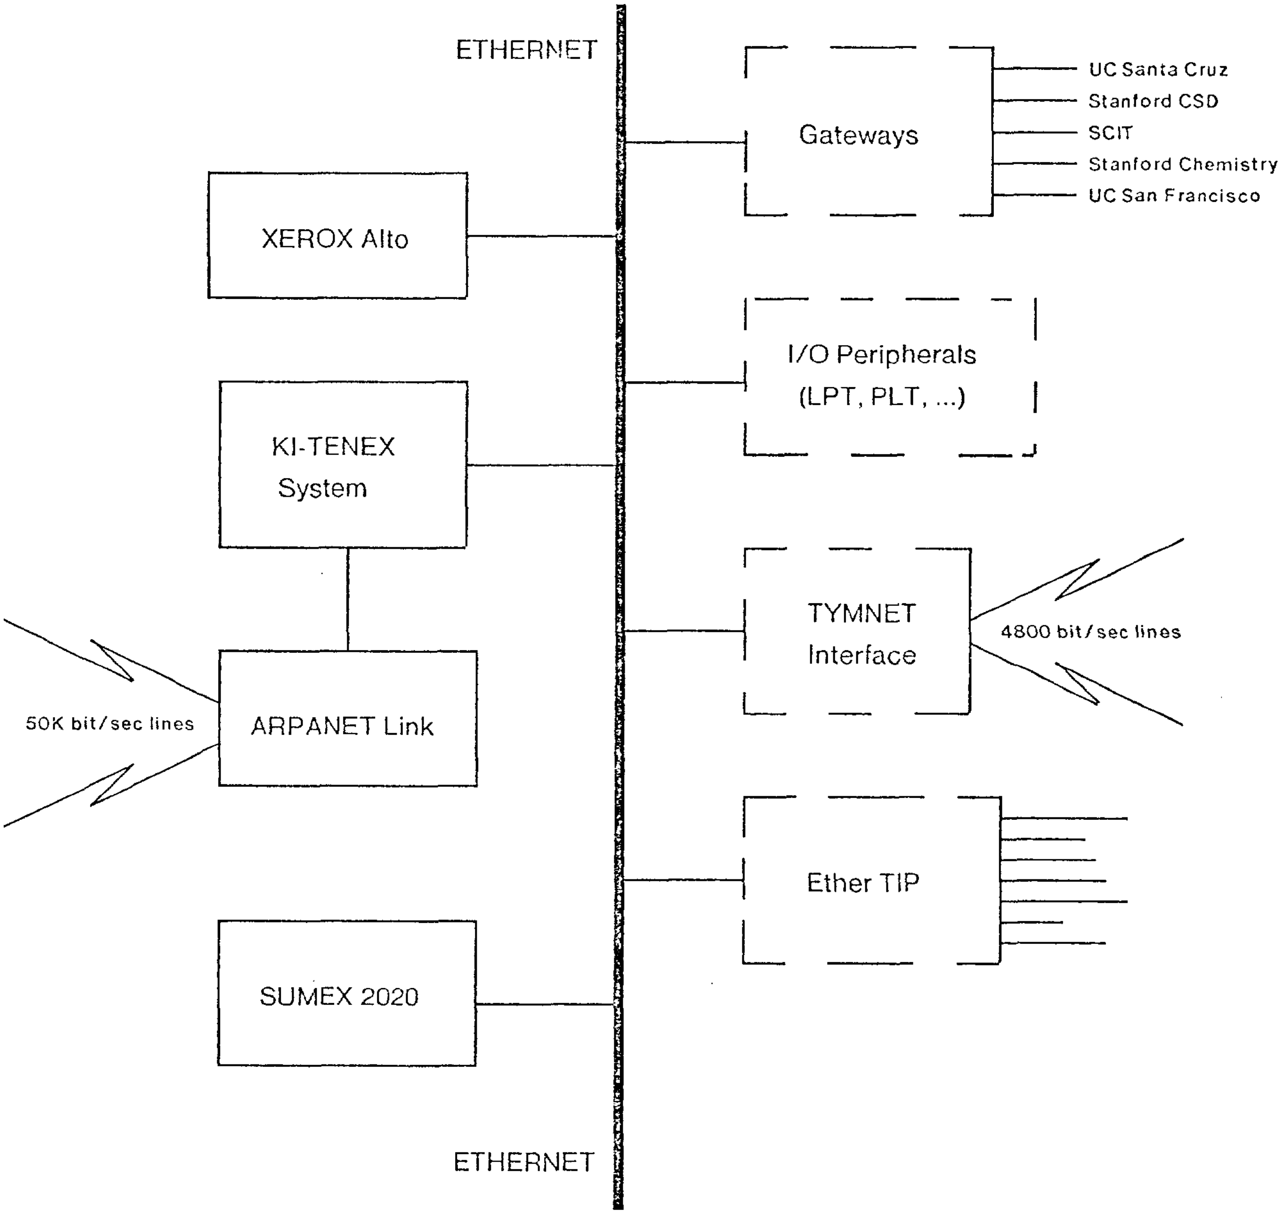
ETHERNET
Gateways
UC Santa Cruz
Stanford CSD
SCIT
Stanford Chemistry
UC San Francisco
XEROX Alto
I/O Peripherals
(LPT, PLT, ...)
KI-TENEX
System
TYMNET
Interface
4800 bit/sec lines
50K bit/sec lines
ARPANET Link
Ether TIP
SUMEX 2020
ETHERNET

Figure 3. Intermachine Connections via ETHERNET

5.3.1.2 System Software

In parallel with the choice of DEC PDP-10 hardware for the SUMEX-AIM facility, we selected the TENEX operating system developed by Bolt, Beranek, and Newman (BBN) as the most effective for our medical AI applications work. Together with the hardware, TENEX has provided a superb environment in which to pursue community biomedical AI applications work. Following are highlights of recent system software developments:

Monitor

1) we have made significant contributions to the KI-TENEX monitor that are now in use at other sites. These include efficiency improvements in the management of user page tables, implementation of a memory-shared TYMNET interface including outbound circuit facilities, design and implementation of the dual processor TENEX system, implementation of a page migration system to assure effective use of fixed-head swapping storage, and improvements in system routines for locating and recognizing file names.

2) developed overload control facilities that effectively limit the number of active processes on the system to those that can be supported with reasonable response time. These provide for "background" jobs, "demo priority" jobs, and mechanisms to temporarily suspend user jobs that have not cooperated with requests to reduce the system load. Active process slots are allocated on the basis of a priori resource percentages that communities and projects are entitled to.

3) implement monitor communication controls for the experimental TELENET network connection. These included special "Xon/Xoff" facilities to allow transmission of packets into the network at 1200 baud irrespective of terminal speed so that network transmission delays could be minimized. Network "backpressure" commands prevented overruns for slower terminals. [see Appendix D on page 376 for details].

4) implement monitor service routines for the "carrier detect" control and line switching hardware.

5) examined KI-TENEX page faulting behavior to measure the utility of block transferring pages in anticipation of faults. Data for a wide range of programs indicate that TENEX already does a good job of keeping needed pages in memory, limited by the amount of physical memory available. We propose to add another 256K of core memory to the system to reduce swapping overhead.

6) integrate the Ethernet and PUP monitor service routines adapted from Xerox PARC [10, 13]. This required redesigning the hardware interface code for our interim PDP-11 I/O bus interface (KI-10) and the 2020 18-bit UNIBUS adapter, changing executive "XCT" codes to conform to differences in hardware function between the Xerox microcoded PDP-10 and our KI-10's, and implementing needed

additional system calls (JSYS's). The KI-10 is fully working on our Ethernet. The extensive TOPS-20 monitor changes for the 2020 are still in progress.

7) adapt the TOPS-20 monitor from the Stanford DEC 2060 systems to the SUMEX 2020. We have made minimal changes to the monitor code except to accommodate the Ethernet interface and to provide needed controls for priority program demonstration and testing.

8) make numerous monitor bug repairs to provide for more reliable system operation and file integrity. Obvious bugs were removed long ago so those remaining are elusive and occur infrequently. We have found and fixed bugs in the management of multi-fork structures, the ARPANET control programs, the file page backup routines, the manipulation of special monitor pages mapped through the user page table, and the concatenation of drum I/O requests for latency reduction.

Utility Features

We have made a significant number of utility improvements to the monitor to add new features, improve compatibility with TENEX 1.34 and TOPS-20, or improve operational effectiveness. A brief list includes:

1) Printer device and spooler that manages a print queue for Prof. Wipke's group at UC Santa Cruz. This device allows interspersing use of the UCSC link as a terminal line and as a printer device.

2) Password error monitoring to log out jobs causing a high number of failures and to report the source and target directories to the operator. This is designed to catch occasional attempts at unauthorized entry into the system, generally from remote network connections.

3) Improved GTJFN features to partially recognize ambiguous file names up to the point of ambiguity and to recognize parts of the TOPS-20 name syntax for compatibility.

4) Upgrade routines and JSYS's to conform with TENEX 1.34 to provide desirable new features (selective expunge, group connect, improved file system physical format, and expanded directory hash table) and to retain compatibility with evolving ARPANET protocols.

5) Checksum monitor code as loaded to detect I/O device errors or memory problems.

6) Make the console teletype of the second processor available for use and improve operational procedures for taking crash dumps and reloading the system.

System Executive

One of the most important system programs is the EXECutive which is the basic user interface to manipulate files, directories, and devices; control job and terminal parameter settings; observe job and system status; and execute public and private programs. The SUMEX EXEC is quite well developed at this stage but we have made several recent improvements:

1) Implementation of LOGIN.CMD and COMAND.CMD files which are processed at login and upon starting any new EXEC. These files allow the user to give any available EXEC command automatically to set default parameters, print status information, etc.

2) Enhancement of the functions and improvement of the human interaction of the file archive/retrieval system. Users can now specify a list of files to be retrieved, edit their archive directories to remove old entries or collect groups of entries, annotate entries to better document contents, and interactively step forward and backward when searching for an entry.

3) Implementation of general wild card facilities for the COPY and RENAME commands. This allows users to copy/rename groups of files to new files with names derived by reorganizing selected substrings from the originals thereby reducing the manual typing required.

4) Implement the selective expunge command from TENEX 1.34 so that temporary files (e.g., MESSAGE.COPY) can be retained while expunging unneeded deleted files.

5) Improvement of the scheduling control information provided to users for planning their work around overloaded system conditions.

6) Implement demo controls for the 2020 EXEC to preserve its capacity during scheduled sessions for AI program tests or demonstrations.

System Utilities and Operations

We have made numerous improvements and bug fixes to the system utility and operations programs needed to assist smooth management of the system and to provide new facilities for users. A brief list of the most significant tasks includes:

1) Spooler improvements - allow users to retract requests to list files and implement a special spooler for printing files remotely at UC Santa Cruz for Prof. Wipke's group. This spooler communicates over a line also used for terminals and uses a specially designed protocol to coordinate line usage.

2) SYSJOB controls - several of the system utilities for TELNET connections, mail forwarding, statistics collection, TYMNET downtime msg updating, etc. were relocated to a separate system job to facilitate better resource allocation controls and to reduce

competition with other critical system functions (disk page backup and network control programs).

3) Overload controls - implement the user-level demo priority and uncooperative job controls for overloaded system conditions based on the monitor control functions described earlier.

4) File archive/retrieval - improvements to BSYS incorporating user status information on retrieval processing and the latest BBN system for file restoration automation.

5) File system verification - improvements to the CHECKDSK program for detecting file system integrity problems after a crash to allow better notification to users of the names of files that might have been lost or damaged.

6) System and crash analysis - improvements to the program developed to assist in sorting through the complex interlinked monitor tables when unraveling a core dump to analyze the cause of a crash. Also develop several display programs to observe the dynamic operation of individual job structures or network connections.

7) Ethernet/PUP service - import and adapt to the SUMEX system the Xerox user-level service programs for file transfer, terminal connections, mail forwarding, gateway routing, etc.

8) 2020 conversions - on-going conversion of useful KI-10 programs to run in the TOPS-20 environment.

9) TENEX/TOPS-20 compatibility package - we have made substantial extensions to a compatibility package, PA-2040, that was originally written at USC-ISI. This package now emulates many of the TOPS-20 unique JSYS's. We have added the monitor mode instruction emulation software written initially for the SUMEX GTJFN development so that unique TOPS-20 monitor JSYS code can be run directly from user space. This allows JSYS's without TENEX equivalents to be emulated directly. There are still TOPS-20 JSYS definition changes that cannot be handled by means of a compatibility package.

User Subsystems

We have continued to assemble (develop where necessary) and maintain a broad range of user support software. These include such tools as language systems, 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 user information/help assistance programs.

1) new installations or versions of subsystems essential to users have been brought up with varying requirements for local adaptation to run on the SUMEX KI-10's. New or updated subsystems include MLAB

and OMNIGRAPH from NIH; FORTRAN, CCL, COBOL, BACKUP, MACRO, LINK10, GLOB, and a new set of utility routines used by many of the DEC CUSP's from DEC; INTERLISP from Xerox PARC; ESSEX-BCPL from the University of Essex in England; PASCAL and SAIL from Rutgers University (C. Hedrick); PUB (a text formatting program) from IMSSS (M. Hinckley) and SUMEX; MSG (a mail reading program) from BBN (J. Vittal); and TEX (a text publication system) from Stanford (D. Knuth).

2) upgrade the crt display package in the TV text editor to support many additional terminals. TV now handles Teleray-1061, Heath H-19, and a locally modified version of the Hazeltine 1500. Support will soon be available for the NIH Delta Data 5200, Infoton 400, and Visual 200. We are also incorporating enhancements made recently by C. Hedrick at Rutgers to allow improved search and text relocation facilities.

3) import and support the EMACS text editing system from MIT. Substantial effort has gone into developing macro packages that improve the human engineering features of EMACS and providing introductory documentation for new users. This has been closely coordinated with similar efforts at SRI and MIT. A community of EMACS users is now developing at SUMEX.

4) add features to allow attaching batch jobs that have an initial interactive phase that has to be run from a user terminal but which can then be turned over to batch operation for background or deferred running. Also improve batch efficiency and help facilities.

5) add facilities to the spelling corrector to replace misspelled words with phrases, remember the names of subdictionaries loaded, and override misspellings to do simple translations.

Communications Subsystems

Of key importance for our community effort is a set of tools for inter-user communications. We have built up a group of programs to facilitate many aspects of communications including interpersonal electronic mail, a "bulletin board" system for various special interest groups to bridge the gap between private mail and formal system documents, and tools for terminal connections and file transfers between SUMEX and various external hosts. Recent developments include:

1) TTYFTP - A system for file transfers usable over any circuit that appears as a terminal line to the operating system (hardline, dial-up, TYMNET, etc.) and incorporating appropriate control protocols and error checking. The design is derived from the DIALNET protocols developed at the Stanford AI Laboratory with extensions to allow both user and server modules to run as user processes without operating system changes. TTYFTP is written in MAINSAIL and is implemented for TENEX, TOPS-20, RT-11, and RSX-11M.

2) Bulletin Board - BBD has been extended to allow remote posting of bulletins via communication network and has improved efficiency.

3) VTTY - we have combined outbound (TELNET) terminal access protocols for TYMNET, SCIT (Stanford IBM facility), SUMEX 2020, and pseudoteletypes in a single virtual terminal program. VTTY provides typescript services to record sessions.

4) Electronic mail - improve the mail facilities for guests and allow reediting of all message fields (i.e., addressees, subject, and body) in SNDMSG. Also import the more efficient protocols for network mail developed by K. Harrenstien at MIT.

Software Sharing

At SUMEX-AIM we are committed to 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 availability 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.

1) We continue to import 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.

2) We have 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 9.1.3 on page 149). We have provided him with the non-licensed programs requested. We are working on a similar arrangement for a group interested in the MOLGEN program.

User Assistance and Documentation

The SUMEX resource exists to facilitate biomedical artificial intelligence applications from program development through testing in the target research communities. This user orientation on the part of the facility and staff has been a unique feature of our resource and is responsible in large part for our success in community building.

1) We have tailored resource policies to aid users whenever possible within our research mandate and available facilities. Our approach to system scheduling, overload control, file space management, etc. all attempt to give users the greatest latitude possible to pursue their research goals consistent with fairly meeting our responsibilities in administering SUMEX as a national resource.

2) The resource staff has spent significant effort in assisting users gain access to the system and use it effectively. We respond promptly to questions by telephone, terminal link, or electronic mail. We also exercise great care in managing system file integrity and assisting users in recovering files lost through user error or system malfunction.

3) We have worked hard to assist projects achieve their goals in setting up an appropriate computing environment on the system including directory groups, collaborator and guest facilities, file space allocations, and special software subsystems.

4) We have solicited and acted upon user recommendations for system development goals. A "gripe" system is available to users for general comments as well as electronic mail to individual staff members responsible for particular aspects of the system.

5) We have spent substantial effort to develop, maintain, and facilitate access to 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. We have worked to review the existing documentation system, reorganize it for easier access and maintenance, create command and documentation summaries where appropriate for new users, and update on-line and hardcopy documents for compatibility with the programs now running. We have collected useful comparisons and difference summaries between the KI-TENEX and 2020 systems to assist users in moving easily between them. Maintenance of accurate and useful documentation is a continuing task.

5.3.1.3 Network Communication Facilities

A highly important aspect of the SUMEX system is effective communication with remote users. 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. These issues become even more important with the emerging computing technology that will make increasing decentralization possible. Networks will be crucial for maintaining the collaborative scientific and software contacts built up. A detailed description of our network connections can be found in Appendix D on page 376. Recent milestones include:

1) We continue our connection to TYMNET as the primary means for access to SUMEX-AIM from research groups around the country and abroad. There has been no significant change in user service or network performance. Very limited facilities for file transfer exist and no improvements appear to be forthcoming soon. Services continue to be purchased through the NLM contract and we have elected "dedicated port" pricing as the most cost effective. We continue to have serious difficulties getting needed service from TYMNET for debugging network problems. See Figure 18 on page 379 for a recent list of TYMNET access nodes.

2) We continue our advantageous connection to the Department of Defense's ARPANET, now managed by the Defense Communications Agency (DCA). Terminal access restrictions are in force so that only users affiliated with DoD-supported contractors may use TELNET facilities. ARPANET is the primary link between SUMEX and other machine resource such as Rutgers-AIM. Current ARPANET geographical and logical maps are shown in Figure 19 and Figure 20 on page 380.

3) We implemented an experimental connection to TELENET via a TP-2200 interface with 12 asynchronous lines to SUMEX and one 4800 baud line connecting to the network backbone. In spite of potential economic advantages, this experiment was unsuccessful. Users complained of poor node reliability, intolerable delays in response, uneven flow of terminal output, and poor operational management of the network. Similar problems existed from the system standpoint. Other half-duplex users (e.g., the NLM MEDLINE system) have reported more successful connections. Because of funding limitations, we had to abandon our TELENET link for the time being. See Figure 21 on page 382 for a recent list of TELENET access nodes.

5.3.1.4 Resource Management

Early in the design of the SUMEX-AIM resource, a rather elaborate management plan was worked out with the Biotechnology Resources Program at NIH to assure fair administration of the resource for both Stanford and national users and to provide a framework for recruitment and development of a scientifically meritorious community of application projects. This structure is described in some detail in Appendix E on page 383. It has continued to function effectively as summarized below.

1) The AIM Executive Committee meets regularly by teleconference to advise on new project access applications, discuss resource management policies, plan workshop activities, and conduct other community business. The Advisory Group meets together at the annual AIM workshop to discuss general resource business and individual members are contacted much more frequently to review project applications. (See Appendix I on page 399 for a current listing of AIM committee membership).

2) Effective July 1, 1978, Prof. Edward Feigenbaum, Chairman of the Stanford Department of Computer Science became SUMEX principal investigator after Prof. Joshua Lederberg assumed the presidency of The Rockefeller University. This transition took place smoothly because of Prof. Feigenbaum's role as co-Principal Investigator of SUMEX from its start and his long standing collaboration with Prof. Lederberg. Close scientific and administrative ties are maintained with the Stanford medical community through Prof. Edward H. Shortliffe, who is one of the key designers of MYCIN and co-Principal Investigator of SUMEX. The project will become administratively part of the Stanford Department of Medicine, effective August 1980. As part of the largest clinical medicine department at Stanford, SUMEX will have increased visibility and opportunity to broaden its local scientific collaborations.

3) We have actively recruited new application projects and disseminated information about the resource. The number of formal projects in the SUMEX-AIM community has nearly quadrupled since the start of the project (see Figure 6 on page 331). Here, for example, are just some recent efforts to broaden outside awareness of work in the AIM community and to encourage new projects: the CONGEN workshop at Stanford (1978); the AGE workshop at Stanford (1980); an AI session at the Fourth Illinois Conference on Medical Information Systems (1979); INTERNIST and MYCIN participation in a course on AI computing at NIH (1979); an AI session at the Association for Information Science meeting (1979); an AI session at the Sixth International Joint Conference on AI (1979); an extensive lecturing tour among Japanese university, government, and industrial research groups; and MYCIN and INTERNIST program demonstrations at the American College of Physicians meetings (1979 and 1980).

4) With the advice of the Executive Committee, we have awarded pilot project status to promising new application projects and investigators and where appropriate, offered guidance for the more

effective formulation of research plans and for the establishment of research collaborations between biomedical and computer science investigators.

5) We have welcomed a number of visiting investigators at Stanford who were able to pay their own expenses, so they could see first hand how AI applications programs are formulated and get acquainted with the computing tools available. Funds for such visiting scientists were deleted from our previous grant award.

6) We have allocated limited "collaborative linkage" funds as an aid to new projects or collaborators with existing projects to support terminals, communications costs, and other justified expenses to establish effective links to the SUMEX-AIM resource. Executive Committee advice is used to guide allocation of these funds.

7) We have carefully reviewed on-going projects with our management committees to maintain a high scientific quality and relevance to our biomedical AI goals and to maximize the resources available for newly developing applications projects. Several pilot projects have been terminated as a result and more productive collaborative ties established for others.

8) We have continued to provide active support for the AIM workshops. The last one was held in May 1979. It was organized by MIT-Tufts and Rutgers and was devoted to clinical diagnosis programs. We also have supported individual project workshops such as those held for CONGEN and AGE. The next AIM workshop will be held at Stanford in August 1980 together with several tutorial sessions on AI for physicians. Prof. Shortliffe is the program chairman for this workshop.

9) We have continued our policy of no fee-for-service for projects using the SUMEX resource. This policy has effectively eliminated the serious administrative barriers that would have blocked our research goals of broader scientific collaborations and interchange on a national scale within the selected AIM community. In turn we have responded to the correspondingly greater responsibilities for careful selection of community projects of the highest scientific merit.

5.3.2 Core Research

Since the last report 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. Specific areas of current effort include:

1) The AI Handbook, under Prof. Feigenbaum and Mr. Avron Barr: a compendium of knowledge about the field of artificial intelligence being compiled by students and investigators at several research facilities across the nation. 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 first two volumes covering heuristic search, knowledge representation, natural language and speech understanding, AI languages, various applications domains, and automatic programming are complete. All completed sections are published as Stanford Computer Science Department technical reports. Work on a third volume is progressing well. [see Section 9.1.2 on page 145 for a more detailed report and Appendix G on page 392 for an outline of the handbook contents)

2) The AGE project: 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 first version of the AGE system has been completed. It uses the "blackboard model" control structure for coordinating multiple expert sources of knowledge for the solution of problems. The UNITS package [9] for a "frame-oriented" representation of knowledge is now being incorporated. AGE provides a general structure and an interactive facility for implementing knowledge-based systems. A workshop to introduce AGE to the AIM community was held at Stanford in February 1980. [see Section 9.1.1 on page 137 for a more detailed report].

3) The MAINSAIL project: an effort to design and demonstrate a machine-independent, ALGOL-like language system to facilitate software transportability between different machine/operating system environments. We successfully completed the design and a demonstration of the MAINSAIL language system as a tool for software portability [14, 15]. A common compiler, code generators, and runtime support for TENEX, TOPS-10, TOPS-20, RT-11, and RSX-11 have been developed as part of this demonstration system and numerous applications programs written by collaborating research groups. Further work past this demonstration phase will be done independently of SUMEX through a private company, XIDAK, formed to continue the development, dissemination, and maintenance of MAINSAIL. Work is under way to develop MAINSAIL for the VAX and a number of other target machines. [See Appendix H on page 398 for a more detailed summary of the final phases of this project].

It should be noted that SUMEX provides only partial support for the AI Handbook and the AGE projects with complementary support coming from an ARPA contract to the Heuristic Programming Project. Other portions of our original proposal for core research in knowledge acquisition, planning, and generalized explanation systems have not been supported for lack of resources following council reduction of this section of our budget.

5.3.3 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 9 on page 135). 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.

6 Methods of Procedure

This section details our approach to achieve the goals summarized in Section 3.3 on page 18 during the next five year period. As indicated earlier, objectives and plans for individual collaborating projects are discussed in Section 9 beginning on page 135.

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.

6.1 Resource Operations Plans

6.1.1 Resource Hardware

6.1.1.1 Rationale for Future Plans

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 in Appendix B on page 355 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 Appendix F on page 390. 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.

An "obvious" solution to our capacity needs (but not the address space limitations) is to buy additional large machine resources that are software compatible with the existing community KI-10 and PDP-20 systems. By placing these nodes at user sites, an improvement in communication bandwidth would be possible to enhance the human interactive support. The addition of more DEC 2060 or larger machines to the SUMEX community is not cost-effective, however.

An alternative and more feasible approach to meet community needs is to explore the use of smaller, less expensive machines as satellites (some remote) to the main resource. A variety of technologies are now becoming available as machines that we can buy and use. These could have a number of advantages:

1) A relatively small investment in capital equipment is required for each incremental capacity augmentation.

2) New architectures directly support larger program address spaces.

3) Possible location close to individual research groups allows better human engineering of user interfaces by using higher speed

communication, improved display technology, and other modalities for human interaction such as speech and touch.

4) System capacity can be allocated more flexibly and efficiently by having to satisfy fewer simultaneous scheduling constraints and by being more easily dedicatable to operational demonstrations.

This approach poses a number of possible disadvantages stemming primarily from the distributed nature of the computing resources:

1) Each such machine would have a relatively small capacity. These may be sufficient for many computing tasks of a local user group. It would be difficult to aggregate such dispersed capacity, however, when needed for a single computing-intensive task except through multiprocessing. This would be made difficult by geographic remoteness. Such intensive computing needs will likely still be best handled by shared specialized central resources.

2) Decentralizing the computing resources places an increased centrifugal force on community interactions. Effective network communications must be maintained to allow continued collaborative interactions, software sharing, access to common knowledge and data bases, message exchange, etc.

3) Geographically distributed computing tends to encourage costly duplication of similar operations and maintenance functions for system hardware and software support. These added costs are lessened when distributed over clusters of systems near SUMEX-AIM community nodes.

These trade-offs, coupled with the developing new computer technology, suggest a continuing need for a spectrum of resource configurations 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.

6.1.1.2 Summary of Proposed Hardware Acquisitions

As discussed in more detail in later sections, we plan to acquire the following additional hardware

yr 1 - Add 256K words of core to the existing KI-10 AMPEX memory to reduce page swapping overhead.

- Buy a VAX 11/780 with 2M bytes of memory and minimal disk and tape peripherals to provide large address space INTERLISP facilities, to experiment with AI program export, to support development of VAX system software for the community, and to alleviate congestion in the Stanford 40% of the SUMEX resource.

- 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. The server will be based on a PDP-11 with 630M bytes of disk storage initially and tape facilities for backup and archives.

yr 2 - Add 2M bytes of memory to the VAX purchased in year 1.

- Add 630M bytes to the file server purchased in year 1.

- Buy 5 single-user "professional workstations" (PWS) based on the Zenith-MIT NU system (or equivalent) to develop and experiment with this means for AI program development, export, and human interface enhancements. These machines will be distributed within the Stanford community initially to facilitate development and will be coupled by Ethernet with the main resource.

yr 3 - Add a second VAX 11/780 for general community support with large address space INTERLISP. This machine will be managed for program testing in a way similar to the existing 2020.

- Add 2 PWS systems to be distributed within the AIM community under Executive Committee control.

yr 4 - Add 3 PWS systems to be distributed within the AIM community under Executive Committee control.

- Add 630M bytes to the central file server to meet expected growth in community file storage needs.

yr 5 - Add 3 PWS systems to be distributed within the AIM community under Executive Committee control.

6.1.1.3 Existing Hardware Operation

The current SUMEX-AIM facilities represent a large existing investment. The KI-10 facility has operated at capacity for more than three years, even with periodic augmentation. Significant augmentation to any of the present hardware configuration cannot be done without major upgrades to the mainframe and memory components. A factor of 5-10 increase in throughput could be achieved by replacing the KI-10's with a DEC 2060 or the projected new 2080 processor. This would maintain software compatibility in the same sense as the 2020 (TENEX vs TOPS-20) but would cost $500 - 1000K. We do not believe the funding for such an upgrade would be forthcoming. It also would not attack the INTERLISP addressing limitations or the needs for higher performance interactive support. Whereas this magnitude of capacity augmentation within the AIM community would indeed be welcome, we feel that SUMEX as a research resource should invest its efforts in exploring newer technologies that offer solutions to current needs with broader long range impact.

For these reasons, 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. We do propose to augment the KI-10 AMPEX memory box purchased in 1977 in order to reduce page swapping overhead. During peak loads, an average of 15-20% of system capacity is lost to pager traps and a substantial additional load comes from drum service interrupt handling. The AMPEX will physically hold another 256K words or 512 pages of memory. Since our current configuration has a net of 852 pages available to users, this increment would provide 60% more physical user space at a cost of only $65,000. We feel this will measurably improve efficiency and smooth out interactive response at high loads.

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.

6.1.1.4 Large Address Space Machines

As indicated in Appendix F on page 390, the user address space limitations imposed by the architecture of the PDP-10/20 systems have been increasingly felt in building large knowledge-based systems for biomedicine. After considerable study, the ARPANET INTERLISP community has started active projects to convert INTERLISP to run on the DEC VAX and to extend the UNIX operating system for VAX to support demand paging and to take advantage of the 31-bit address space. VAX was also the preferred choice as an export machine for the DENDRAL project to support the biomolecular characterization community. Their choice of VAX was made to provide the best match with machines increasingly available in the biochemistry laboratory environment and able to run the programs being developed by DENDRAL (including CONGEN recently converted from INTERLISP to BCPL). Whereas other machines (e.g., PRIME) offer a comparable address space capability and are cost competitive, a comparable software community does not exist on which to base not only AI program development but also the extensive utility software packages for interactive user support necessary to the AIM community.

For these reasons we feel VAX is an ideal candidate for augmenting the SUMEX resource to experiment with large address space LISP systems, to provide added capacity to support software export efforts like DENDRAL, and to alleviate the congestion of the Stanford aliquot of the current system. We propose a modest configuration initially to support developmental efforts to integrate the VAX into the SUMEX resource during the first year of the continuation grant (see Figure 4 for a configuration diagram). This machine can be expected to support 8-10 users initially. In year 2 we plan to increase the memory size by 2 Mbytes to allow more efficient use of the VAX capacity, increasing the users supported to 15-20. In year 3, we plan to add a second VAX to make large address space LISP available more broadly in the community to support future program testing akin to the purpose of the 2020 system. We tentatively plan for another 11/780 system although by then newer models may be available.

6.1.1.5 Single-User Professional Workstations

Motivated by the development of AI programs that are truly useful to their target communities, another major thrust of our research plans for the coming term is the investigation of single-user "professional workstations" (PSW) as a vehicle for exporting AI programs and providing computing power local to the user so that high bandwidth human interactions can be supported (e.g., bit mapped displays for high quality video and graphics, touch, and speech). Emerging VLSI technology promises increasingly capable and cost-effective computing tools through denser packing of microelectronic circuits and reduced development costs to produce relatively specialized systems. Packing density increases by four orders of magnitude may be expected over the next five to ten years [16]. Such hardware advances make the cost-effective marketing of complex AI systems a coming reality.

Prototype single-user professional workstation systems based on current technology such as the Motorola MC-68000 or other special microprocessors are being developed and will begin to be delivered within the year. We must begin now to develop our software systems to take advantage of the improved computing environments these provide for biomedical AI programs. We propose an active role in integrating these systems into the SUMEX-AIM community so that user projects can exploit them for developing, testing, and disseminating their programs.

Current candidates as experimental single-user PWS's include the "PERQ" by Three Rivers Computer Corporation [17], the "D-Class" machines by Xerox Corporation [18, 19], the MIT-developed "CADR" LISP machine by Symbolics, Inc. [20], the MIT-developed "NU" system by Zenith [21], and the "Jericho" system by BBN. Details of the design of most of these systems are still proprietary but deliveries of PERQ, CADR, and NU are expected within a year with continued active development based on user community needs. Characteristically, these systems are intended to be high performance, single user computers with local disk storage, bit-mapped display, and connection to a contention network such as Ethernet or MIT's Chaosnet.

Considerable hardware and system software development work remains on these machines, but by year 2 (1982), we expect them to be relatively well established and we plan to purchase 5 for integration into the Stanford community. We budget $30,000 per machine based on projected pricing of the NU system. The NU will be produced by Zenith from a design by S. Ward at MIT around the MC-68000 microprocessor. This machine supports 23-bit addressing, 32-bit internal data and address registers, 16-bit asynchronous bus, and will soon have facilities for virtual memory management. These will be allocated with 2 machines for Heuristic Programming Project development work, 1 for the experimental ONCOCIN system, 1 for Prof. Shortliffe's research work in MYCIN, and 1 for development work within the SUMEX staff. Our efforts will be to tailor AI performance programs to these systems to provide improved and cost effective expert assistance to biomedical professionals. This first batch of machines will be limited to the Stanford community to allow close access for developing software and tailoring network connection facilities as well as easy maintenance. We will work during that year to tune the software systems on these machines for AIM community use.

In years 3-5, we play to acquire an additional 2-3 machines per year to be allocated among the user community based on Executive Committee advice. We will establish necessary communication links to couple these machines to other AIM resources using leased telephone lines, dial-up services, or commercial network links as appropriate.

6.1.1.6 File Server

An equally important resource to SUMEX-AIM community development is file storage. We have reported frequently in the past on the effects of file storage limitations for our existing resource. As AI programs develop larger knowledge and data bases, as the community of application projects grows, and as more and more external users gain access to test working programs, significantly increased file storage capacity will be needed to support interactive work. It makes little sense to duplicate expensive file storage facilities for each of the machines contemplated in the SUMEX-AIM resource and community.

We expect users to work between several machines in the course of their research and many of the files will be common. Similarly there are many system and documentation files common between the KI-10 and 2020 systems as will be the case between other clusters of similar machines (VAX's and professional workstations).

Thus, a more efficient approach is to implement for each machine only the amount of storage needed to support the currently active users together with a community file service coupled to each machine through a high speed local network (Ethernet). Such a "file server" has worked effectively in the Xerox Alto/Ethernet environment and is a natural approach for the evolving SUMEX-AIM environment. By centralizing file storage, we can minimize equipment costs and file backup, archiving, and operations costs. Such a system even makes selective redundancy for reliability possible and thereby makes users more immune to failures in individual machines.

We plan to implement a basic file server for the SUMEX-AIM community in the first year. It will be based initially on a PDP-11/34 computer with two 317M byte disk drives and two tape drives for backup. The choice of the PDP-11 is based on the ready availability of disk/tape systems for these machines. In years 2 and 4, we plan to add an additional 2 drives each year to bring the total capacity to 2000M bytes.

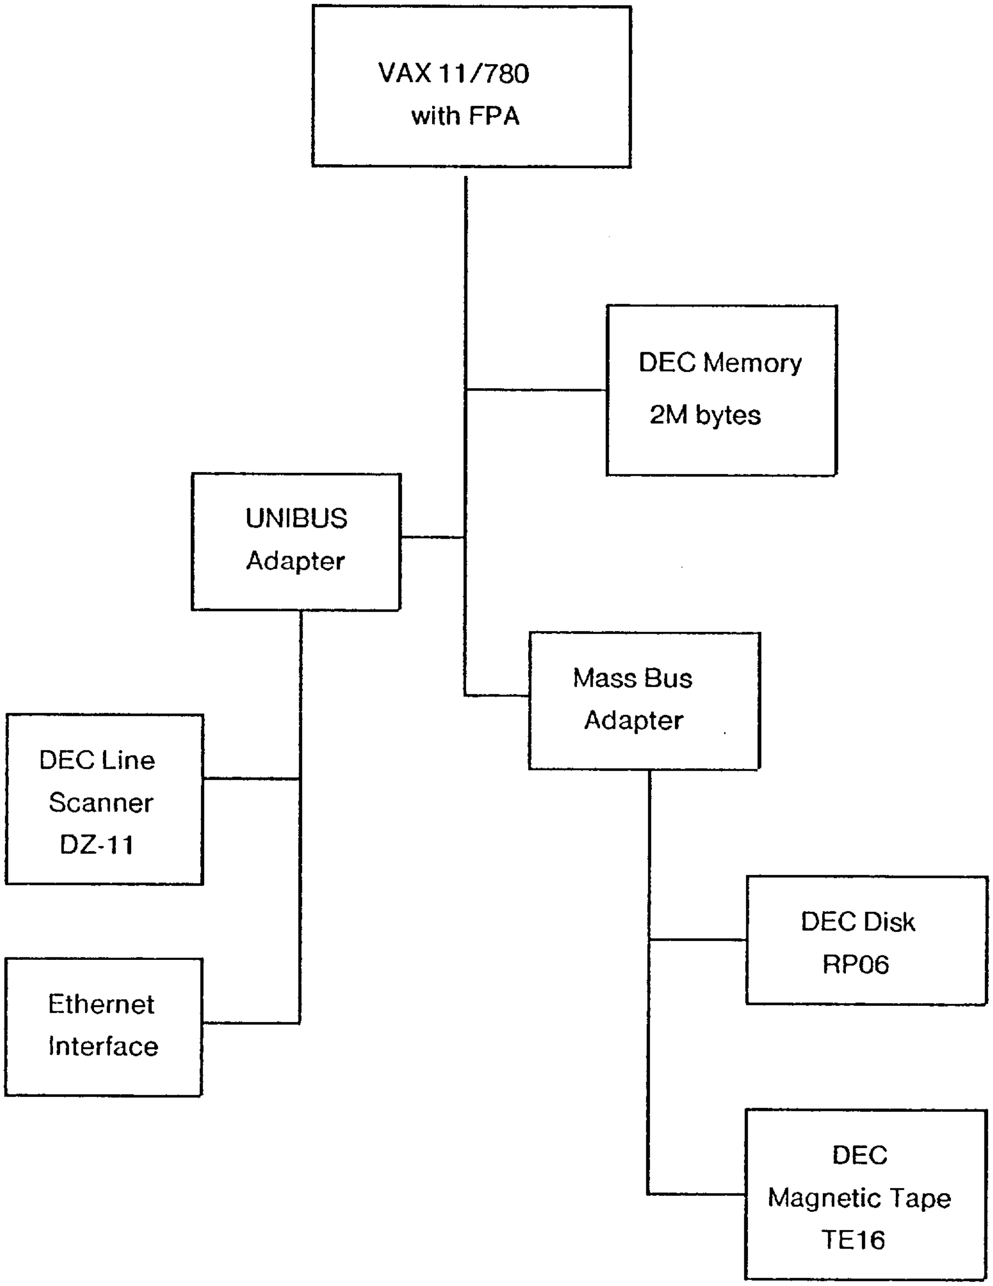
VAX 11/780
with FPA
DEC Memory
2M bytes
UNIBUS
Adapter
Mass Bus
Adapter
DEC Line
Scanner
DZ-11
Ethernet
Interface
DEC Disk
RP06
DEC
Magnetic Tape
TE16

Figure 4. Proposed VAX configuration

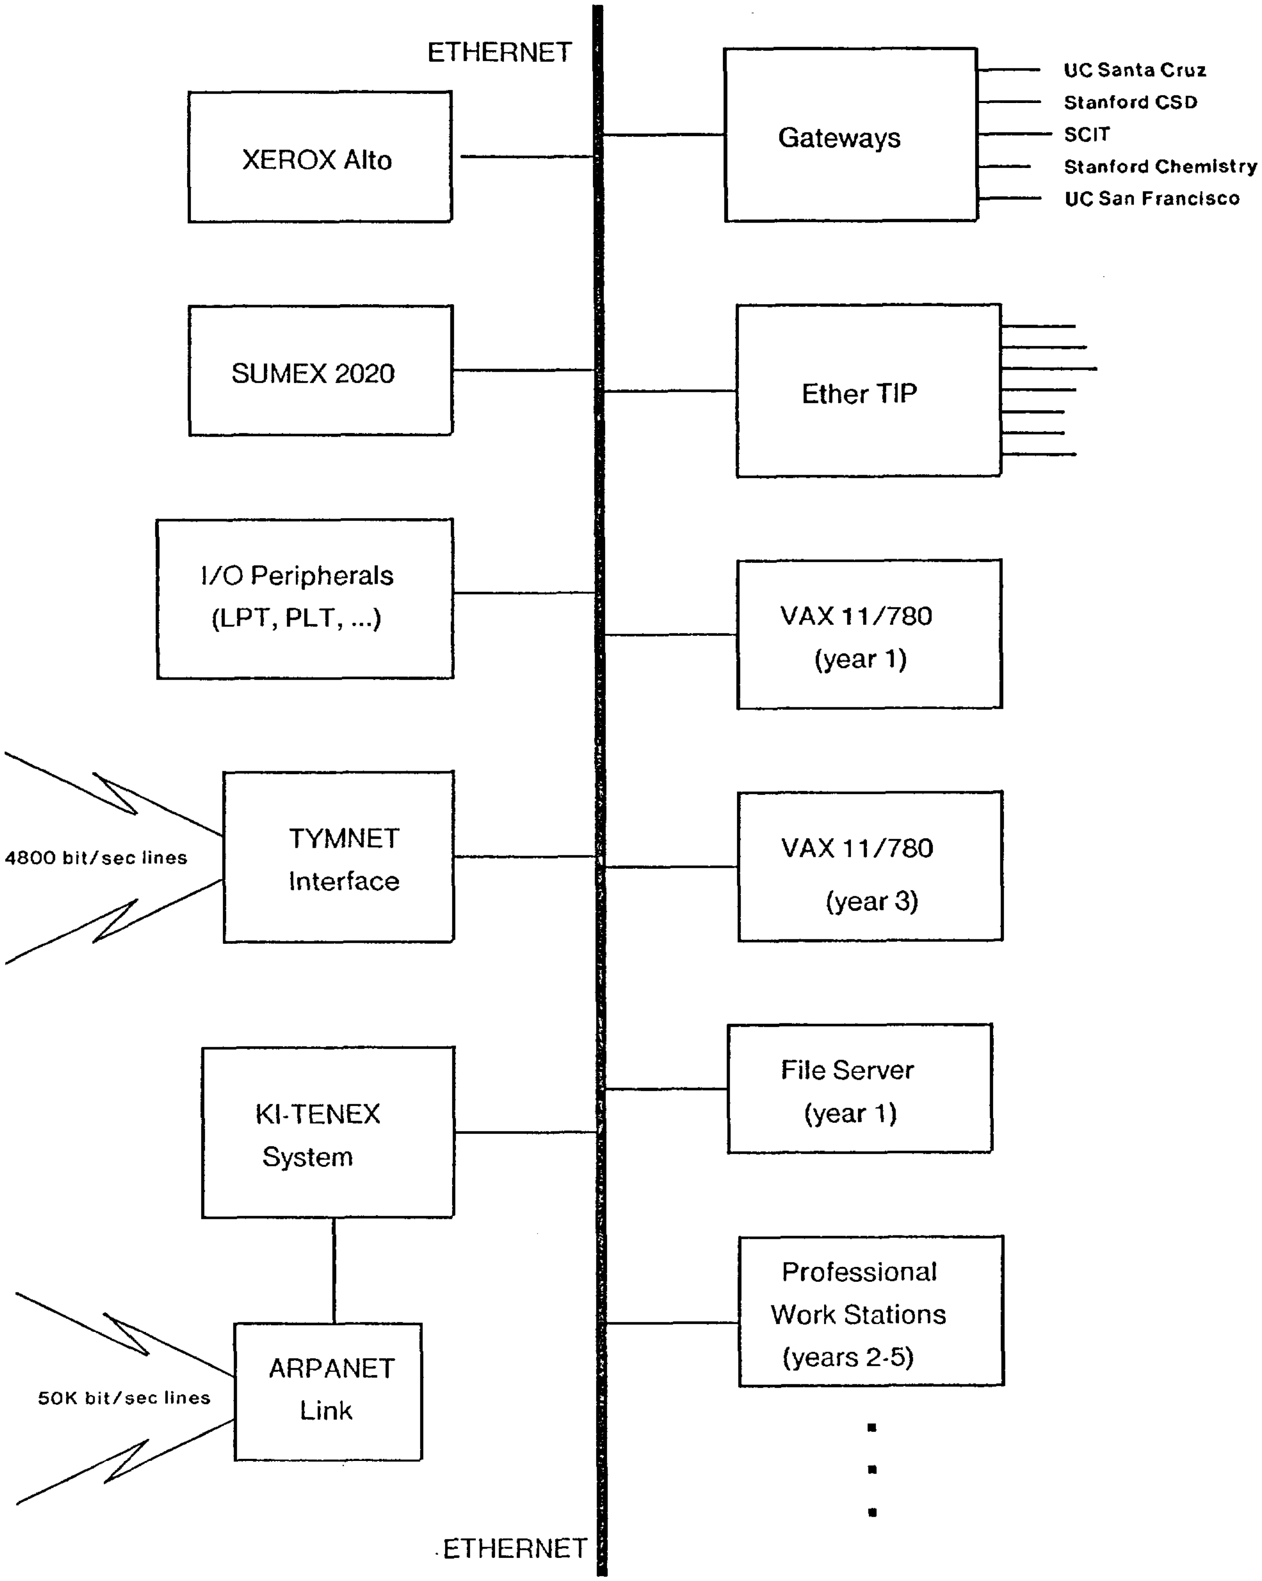
ETHERNET
XEROX Alto
SUMEX 2020
I/O Peripherals
(LPT, PLT, ...)
4800 bit/sec lines
TYMNET
Interface
KI-TENEX
System
50K bit/sec lines
ARPANET
Link
Gateways
UC Santa Cruz
Stanford CSD
SCIT
Stanford Chemistry
UC San Francisco
Ether TIP
VAX 11/780
(year 1)
VAX 11/780
(year 3)
File Server
(year 1)
Professional
Work Stations
(years 2-5)
ETHERNET

Figure 5. Planned Ethernet System to Integrate System Hardware

6.1.2 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.

6.1.2.1 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.

6.1.2.2 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 [10, 13] 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 2020 system with the KI-10 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 5 on page 61 and includes the following major elements:

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

2) 2020 interface. Complete the hardware and software connection of the 2020 using the UNIBUS adapter.

3) 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.

4) 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.).

5) 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.

6) 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.)

6.1.3 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 have 6 bugs partially characterized 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.

6.1.4 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 at Stanford and 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.

Fee-for-Service?

We have pondered the possibilities of a fee-for-service approach for allocation of the resource in the coming period. We believe that this would be inappropriate for an experimental research resource of national scope like SUMEX for several reasons:

1) We have based the development of the national SUMEX-AIM resource entirely on experimentation with tools for new AI research and inter-community scientific collaborations. If obliged to recover some portion of the overall facility cost, these goals may become diluted with administrative and financial impediments, and commitments to paying users, that are tangential to our main research efforts. There is little doubt that a facility of the quality of SUMEX could be tailored to attract paying users (we have turned down numerous such potential users already because they were not aligned with our AI research goals). However, there is little point in demonstrating once again that a computing resource can pay for itself. Rather we should judiciously allocate the available resources to encouraging new medical AI research efforts and stimulating scientific collaborations that cannot always be financially justified at these early stages.

2) A key element in our management plan for SUMEX is to encourage mature projects to acquire computing resources of their own, as soon as justified, and to couple them through communications tethers to SUMEX. This preserves the limited capacity of the central resource for new research efforts and applications. Maturing projects (those able to pay a fee) have every incentive to obtain separate facilities since they cannot obtain sufficient resources from the heavily loaded central resource. In this way such projects effectively pay a "fee" in securing their own facilities and freeing up part of the central facility.

3) A fee structure would impose substantial additional administrative overhead on the project, compounded by its national character. We would face problems of accountability for the transfer of funds from one institution to another. Also SUMEX is a evolving research resource based on changing experimental facilities. Any fee schedule would need to change frequently to fairly respond to developments in the system. Put simply, it would be an administrative nightmare.

For these reasons, 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.

6.2 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.

6.3 Core Research Plans

Motivation

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 each of the projects described in the proposal 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 problems of the area, 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.

Research Plan

In the following descriptions, planned core research efforts are grouped under the five major headings listed above, although it should be clear that the boundaries are frequently crossed. Knowledge utilization and tool construction are grouped together.

6.3.1 Knowledge Representation

6.3.1.1 RLL -- The Representation Language Language

A framework for constructing new representation languages, called RLL (for "Representation Language Language"), is under development within the HPP. RLL explicitly represents (i.e., contains schemas for) the components of representation languages, including itself. The primitive building blocks of representation languages are larger and more abstract than the primitives of general programming languages in order to make them easier and more natural to use. Building blocks of a representation language include such things as control regimes (agendas, backward chaining, etc.), methods of associating procedures with relevant knowledge (footnotes, demons, etc.), fundamental access functions (put/get, assert/match, etc.), automatic inference mechanisms (inheritance of various kinds), and specifications of intended semantics of the components (consistency constraints, etc.).

RLL is designed to help manage these complexities by providing (1) an organized library of such representation language components and (2) tools for manipulating, modifying, and combining them. Rather than produce a new representation language as the "output" of a session with RLL, it is rather the RLL language itself, the environment the user sees, which changes gradually in accord with his commands.

The RLL system needs to be developed into a usable package, and experimented with. Only through multiple usages will directions for future research be revealed. Several systems are already planned for (some layer of) RLL, including: a new version of the system for diagnosis of pulmonary function disorders; a program for guiding a physician in constructing a new expert system automatically; and a few non-medical applications.

Already, we have isolated several core research issues, which will govern the direction of our research during the next five years. This agenda of issues includes:

(1) Incorporating the representational schemes of other researchers into RLL. For instance, the user should be able to specify that he or she wants a KRL-like environment, or a MYCIN-like environment, and the bundle of "organ-stops" which must be adjusted should change immediately.

(2) Codifying knowledge about representation. This includes refining our taxonomies of inheritance modes, control structures, etc.

(3) Building up our stock of ideas about fundamental representation issues: dealing with nested quantification, mass nouns, time, intensional objects, counterfactual conditionals, etc.

(4) Easier knowledge acquisition. One approach to this is to improve the interface to an expert user, who must transfer his knowledge into a program. For example, the knowledge acquisition program mentioned

above can direct its knowledge acquisition process because it possess a detailed model of what comprises such a session. A second, and currently underexplored, approach is to have the program automatically discover the knowledge for itself. This may appear much more costly, but recall that "expert knowledge" breaks down into facts and heuristics. The latter are almost never articulated by experts; it is easier to induce them from examples. This leads us to study:

(5) Automatically discovering new domain-dependent heuristics. This was the critical lack in an earlier discovery system, AM [22], which had some success in automatically discovering new (albeit elementary) concepts, by combining old ones. Our work in the past two years has indicated that powerful heuristics can be found as simple patterns in the values of slots, provided the system has very useful domain-specific slots. Thus this is pointing us to the problem which follows:

(6) Automatically discovering new domain-dependent slots which prove useful. Our approach, as usual, is to explicate and codify. We are building a taxonomy of slots; i.e., of useful relations between concepts (units). Already the number of slots is in the hundreds, and over the next five years we expect this number of different kinds of slots worth distinguishing to increase by an order of magnitude. This in itself will raise several new issues.

(7) Ultimately, tackling the problem of automatically discovering new representations of knowledge. Currently, our only plan to attack this problem is to represent each type of representation (e.g., graphical, schematized, linguistic,...) as a unit, organize these into a hierarchy, and see if the domain-independent heuristics are adequate to guide the search for new and better representation schemes.

6.3.1.2 Research on Planning

In many situations, solving problems by trial and error can be prohibitively expensive or impossible. One of the characteristics of an intelligent problem solver is the ability to formulate a "plan" before acting. Consider, for example, a physician ordering costly or risky tests or a chemist designing an experiment or a businessman trying to get to the airport. Much of the research in Artificial Intelligence has been concerned with formalizing this idea of planning in the form of intelligent computer programs. One approach has been to concentrate on techniques applicable in all task domains. Another approach has emphasized the importance of techniques specific to particular task domains. The experience in building high performance programs like DENDRAL, MYCIN, and MACSYMA has shown us the value of the latter, "knowledge-based" approach to system construction. However, even within this approach there are many domain independent questions yet to be answered. Consequently, we propose to explore some fundamental issues of planning that promise to increase performance and facilitate the construction of expert systems.

Research using SUMEX has demonstrated the value of the knowledge-based approach to planning program construction in the MOLGEN program. By extending the technology, we expect to enable similar success in other types of experimental design. The basic planning research we are proposing here will mesh nicely with a collateral effort to create an "intelligent agent" that has mastered the facts and lore of using complex computer systems and can use this knowledge to facilitate a user's interactions with the system (ARPA-funded research).

In our research on planning, we view problem solving as a three step process. Given a goal to satisfy, the problem solver uses information about the actions it can perform in synthesizing a plan to solve the problem. Then it executes the plan, possibly monitoring its performance to confirm success or detect failures. In the event of a failure (perhaps due to unforeseen circumstances), the problem solver must rectify any undesirable consequences and create a new plan.

A. Plan Formulation

The outstanding problems in plan formulation involve the use of strategical (or meta-level) control of planning, the development of a good representation for plans and planning methods, and the encoding of powerful planning techniques.

A.1 Meta-Planning

The operation of many planning programs can be described in terms of the "queue and process" model. The program maintains a data structure representing a partially designed plan and a queue of operations to perform on this data structure. An interpreter selects an operation from the queue and executes it, thereby expanding or refining the plan and possibly adding new operations to the queue. The problem of deciding the order in which to select members from the queue is a strategical one, and a variety of techniques have been proposed to solve it. One approach is to annotate each operation with a number reflecting its cost and probability of success. A more powerful approach is to view the selection task as a problem in its own right, on which the full power of the problem solver can be brought to bear. This latter approach is usually termed "meta-planning".

Stefik [23] has recently shown the power of this approach in the design of laboratory experiments. The MOLGEN program uses a level of strategical planning to direct the operation of the basic plan generator in designing gene-cloning experiments. His meta-planning techniques allow the program to choose between constraint propagation or a guess and backup approach.

We propose to consolidate Stefik's success by extracting the domain independent skeleton of MOLGEN and by developing additional techniques. We are interested particularly in the following questions.

(1) What is the appropriate structure for a system with meta-planning capabilities? Also, how many meta-levels are ideal?

(2) What techniques are there in addition to those used by Stefik?

A.2 Representations for Plans and Planning Methods

Sacerdoti's work on the NOAH system [25] has pointed out the importance of a flexible representation for plans that does not force one to make premature decisions about the order of actions or the identity of essentially arbitrary objects. Nevertheless, his "procedural net" formalism has several limitations, e.g. there is no way of representing conditionals, loops, or actions with parameters. We propose to develop an extension of his formalism that remedies many of these deficiencies.

Sacerdoti describes a procedural net as "a network of actions at varying levels of detail structured into a hierarchy of partially ordered time sequences. An action at a particular level of detail is represented by a single node in the network." [25] Each node may represent a "primitive action" or may point to a subnet of more detailed subactions. When executed in proper order, the subactions achieve the effects of their parent.

A "parameterized procedural net", or PPN, is a procedural net in which each node has associated with sets of input and output objects and sets of "prerequisites" and "postrequisites" (conditions that must be true for the action to succeed and those that become true after its execution). In Sacerdoti's formulation, each action node is described in terms of specific objects in the task domain. In a PPN the dataflow into and out of an action node is described in terms of other actions without naming any specific domain objects. Thus, a PPN is like a partial program. We believe the PPN formalism is an adequate representation for plans that allows complete flexibility in the specification of action type, control flow, dataflow, etc. However, we would like to study its formal properties and expressiveness, and we need to develop an appropriate interpreter to execute plans in this representation.

The extreme flexibility of the PPN notation makes it an excellent choice for encoding planning techniques as well as the plans they produce. The only difficulty is that its two-dimensional character makes it more difficult to use than the strings of characters used in conventional programming languages. Consequently, we propose to develop a one-dimensional version, similar to LISP except with multiple return values, nondeterministic function calls, and explicit representation of prerequisites and postrequisites.

A.3 Basic Planning Methods

The AI literature describes many domain independent planning techniques. Consider, for example, Newell and Simon's means-end analysis, prerequisite achievement, and Sussman's and Stefik's constraint propagation techniques. We would like to assemble a library of such planning techniques all encoded within our planning formalism, to be put at the disposal of a system builder. The library should also include domain specific techniques. Friedland [24] has already made some progress in this direction in the domain of molecular genetics. We would like to continue this work and strive to provide convenient methods for users to develop and edit these libraries.

B. Execution Monitoring

Once a plan is formulated, it can be executed. In many domains there is the possibility of failure due to unforeseen circumstances, e.g. a chemical synthesis has an unacceptably low yield or a computer system runs out of disk space. In other cases, a failure may occur due to partial or incorrect knowledge about the domain (as described in the last subsection). To trap such problems, the problem solver must monitor the execution of the plan. Execution monitoring is an important problem that has not yet received adequate attention. The key questions that must be answered include the following.

(1) How does one monitor the execution? In many cases, devising the test is as difficult a problem as solving the original problem. What special techniques are necessary in this regard?

(2) What aspects should be monitored? Which have the greatest diagnostic value? Which aspects are most likely to be violated? Given that testing is not free, how does one decide when to monitor?

The reason for monitoring each step of a plan rather than just the final outcome is that failures can propagate and that one can miss the opportunity to take corrective action.

C. Recovery from Execution Failure

Once a problem solver observes a failure, there are a variety of actions it can take. If adequate monitoring has been performed, the source of the failure should be immediately evident. If not, the program must localize the problem in order to correct it.

Once the failure is diagnosed, several responses are possible. The problem solver may scrap its efforts and start from scratch; or it may pick up from the failed state and produce another plan. In some cases, the original plan may be used but only after the effects of the failed attempt are undone. We would like to explore the techniques for dealing with failure in a variety of settings from the laboratory to the computer system, and we would like to study the trade-offs involved where several approaches are possible.

6.3.1.3 Causal Models

Medical and scientific reasoning depend on exploiting causal relationships. We have encoded causal knowledge in production rules and other representations without separating it from empirical associations. This was a successful pragmatic approach. However, we recognize the importance of representing and manipulating causal models as a separate kind of knowledge in our reasoning programs, knowledge acquisition systems and tutoring programs. We propose using the VM program as a test-bed and point of takeoff for this research.

The VM program provides real-time interpretation of the clinical significance of measured data in the ICU. The project has considered the relation between three related sets of abstract clinical data: physiological information, measurements provided by a monitoring system, and diagnostic parameters used in patient care. For example, VM uses a patient's respiratory rate, a measured parameter, in interpreting his effort of breathing, a diagnostic parameter. VM now includes a limited set of physiological parameters, those directly related to measured data. Associations between measured, diagnostic parameters and physiology are now represented in VM when the association between these parameters is close and apparent.

```
                 diagnostic
                 parameters
                 /        \
                /          \
     measurements    ----    physiological
    from monitoring           information
         system
```

The proposed research will increase the number of physiological parameters used in the system and increase the number of interactions among the three kinds of parameters which are represented in the system. Specifically, this research will attempt to represent causal relations among physiological parameters (based on a physiological model) and among measurement parameters (based on a model of instrument function).

6.3.2 Knowledge Utilization and Tools for Building Expert Systems

6.3.2.1 Attempt to Generalize (AGE)

The AGE system is currently supported under SUMEX core research and its progress and future plans are described in Section 9.1.1 on page 137.

6.3.2.2 AI Handbook

The AI Handbook is also supported currently under SUMEX core research and its progress and future plans are described in Section 9.1.2 on page 145.

6.3.2.3 Research in Automated Consultation about Expert Systems

One of the drawbacks of knowledge based systems is that they are often difficult to use. Consider, for example, a scientist trying to solve a problem with a computer system he does not fully understand, and assume that he has encountered a problem due to his lack of knowledge of that system (say MACSYMA or MOLGEN). For example, he may be unaware of the capabilities available, not know the system's vocabulary, or he might get an answer he didn't expect. The simplest way for him to acquire just the information he needs is to ask a consultant for help. Consultation is a method widely used in computer centers; and, as complex programs become more pervasive and more complex, the need for consultants will grow. Unfortunately, consultants are scarce, expensive, and often unavailable when needed.

One partial solution to this dilemma is in the form of automated consultation about the use of complex programs. Recent work by Genesereth in building an automated consultant for MACSYMA demonstrates the feasibility of the approach, but there are many problems to be solved before such consultants are put into general use.

Genesereth's program deals primarily with a user's violated expectations about a system and tries to uncover and correct the misconception responsible for those expectations. It assumes that the user's actions are rational, i.e. that he has some plan for achieving his goal. This plan explains why he chose the operations he did in terms of his beliefs about those operations. The key to the identification of the user's misconception is the recognition and debugging of this plan. In the years ahead, we propose to extend and develop this "plan recognition" approach and apply it to some of the computer capabilities available in the SUMEX resource.

The importance of automated consultation should not be overlooked. In general, a consultant is necessary whenever one is faced with a problem solving situation in a domain one does not fully understand. The lack of knowledge may be incidental, as it is when the domain is fairly simple but time constraints make it impossible for the individual to learn all that is necessary (e.g. with a simple text editor). Or, it may be essential, as when the domain is very complex and the user can't possibly learn everything (e.g. chemistry or MACSYMA).

6.3.2.4 EMYCIN

EMYCIN is a tool for building consultation systems within a backward-chaining framework. For small domains in which experts' judgmental knowledge is expressible in conditional rules, EMYCIN can provide rapid feedback on the adequacy of a rule set for providing reliable consultations. The INTERLISP version of EMYCIN is ready for use by others now. Future work includes the following:

(1) Translation of EMYCIN for broader export

(2) Incorporation of strategic and structural information to integrate the information needed for tutoring and for acquiring new knowledge

The development of GUIDON, a case method tutor for EMYCIN knowledge bases, has given us a new perspective on the nature of the expertise that we have captured in our programs, and suggests guidelines for both representation and acquisition of knowledge. In particular, we have found that rules conveniently separate relationships into readily accessible associations, but an adequate knowledge base for teaching and acquisition requires the addition of structural knowledge (clusters and patterns), support knowledge (underlying causal mechanisms), and strategical knowledge (managerial approaches).

The strategical model is expressed in terms of rules in which the goal or action part is a task to carry out and the premise part consists of steps for achieving the goal. The strategical model will provide the foundation for a new version of EMYCIN which will encourage EMYCIN clients to incorporate in their specialized consultation systems the knowledge we have found to be useful for teaching.

We believe the strategical rules will be useful for controlling inferences as well as for teaching. Implementation of this idea requires that MYCIN's rule interpreter be modified slightly to recognize that some rules describe "tasks" that may be done repetitively, unlike "inference rules" which it only considers once for any case. With the addition of structural knowledge described below, MYCIN's backward-chaining interpreter can then be used to do hypothesis formation with focusing and non-exhaustive search.

Structural knowledge consists of clusters and patterns of rules and parameters--distinctions made by the strategical rules in their control of diagnostic reasoning. The central form of structural knowledge will be a taxonomic classification of the problem space. In MYCIN this will take the form of parameters that are hierarchically related to one another and share properties. One portion of the classification is shown below.

```
                  INFECTIOUS-DISEASE
                     /          \
                    /            \
                   /              \
            MENINGITIS         other infections...
             /        \
            /          \
           /            \
          /              \
    TYPE.ACUTE          TYPE.CHRONIC
     /    \               /  |   \
    /      \             /   |    \
BACTERIAL  VIRAL      FUNGAL TB  PARTIALLY-TREATED-BACTERIAL
   |                   /    \
   |                  /      \
<node for each    CRYPT.   COCCI.
  bacterium>
```

Forming this classification involves regrouping the existing rule set, creating a new parameter for each node in the hierarchy. This design of taxonomic organization and inheritance of properties will make MYCIN's representation more "frame-like," while preserving the use of rules to make judgmental associations among the parameters.

Because the strategical rules embody a weak model of diagnostic behavior, we believe that they constitute a backbone that will be useful for multiple problem areas. In particular, the strategical backbone could be used to structure a knowledge acquisition dialogue. In addition to encouraging a taxonomic classification of parameters, the strategical rules indicate what other kinds of knowledge the expert building an EMYCIN system will have to specify. For example, it is important to detail the knowledge that suggests a broad category of problems that merit attention in a particular case ("triggering associations") and knowledge to adequately discriminate a case on the basis of the taxonomic distinction.

Representing the diagnostic and strategical knowledge in a uniform formalism of rules and parameters, and using an accepted backbone of strategical knowledge, will enable us to use GUIDON for teaching from any new EMYCIN-based program without needing to reorganize the consultation knowledge base. The teaching program will be able to teach a student how to approach cases, while the consultation program will direct its problem-solving according the same approach, one that might be more acceptable to physicians because it is patterned after their methods for solving problems.

6.3.3 Knowledge Acquisition

Our research on knowledge acquisition to date has largely focused on adding new knowledge to an existing knowledge base. A long-term effort is proposed in which we focus on acquiring the structure and contents of a whole knowledge base.

The keystone of our approach to knowledge acquisition is the belief that there is a substantial overlap in the knowledge of many different task domains. We are not referring here to superficial facts and rules (say of the physical world) but rather to the abstract structure implicit in even quite disparate domains. For example, the notion of a hierarchy is found in biological taxonomy, the classification of geologic time, and business organization charts. The advantage of recognizing such abstract structures is that they often possess efficient representations and efficient algorithms for reasoning about them.

In the past this commonality has not been exploited. One reason is the difficulty of representing these abstract structures in a form directly useable in different domains. Another problem is the difficulty of finding and piecing together the structures appropriate to a novel domain. We believe there is an elegant solution for these problems via the notions of abstraction and simulation structure described below, and we propose to develop a library of useful abstractions together with their specialized representations and algorithms from which a knowledge engineer can pick and choose in assembling expert programs.

More specifically, we propose to expend our effort in four major directions: (1) encoding useful abstractions and simulation structures, (2) exploring the use of abstractions in checking the consistency and completeness of knowledge bases, (3) automated selection of simulation structures, (4) the use of abstractions in understanding analogy and the use of analogies in identifying abstractions.

(1) A Library of Abstractions and Simulation Structures

There are an infinite number of possible abstractions. What motivates us to talk of a finite library is the fact that certain abstractions have data representations or algorithms that are particularly efficient or powerful. Some examples are trees, partial orders, rings, groups, and monoids. We propose to differentiate simulation structures on the basis of their representational economy and deductive power. For some structures, this economy and power outweighs the uniformity of semantic networks and frames. We intend to include only those abstractions for which this is the case.

A certain amount of theoretical work must precede the construction of this library. We must devise an adequate language for describing simulation structures and develop data and algorithm representations that facilitates their interface and direct application in new domains. The recent work on abstract operations by Barton, Genesereth, Moses, and Zippel should help in this effort:

(2) The Use of Abstractions in Checking Consistency and Completeness

An abstraction prescribes a set of axioms that must be satisfied by all its models. These axioms can be used to check the consistency and completeness of the assertions a knowledge engineer makes in describing his task domain. For example, if a knowledge representation system suspected that a group of assertions was intended to describe a hierarchy, it could detect inconsistent data, such as cycles or multiple parents, and incomplete data, such as nodes without parents.

The abstractions appropriate to the task domain are determinable from a number of sources. The user may directly name the abstraction or describe it with an analogy; or the system may be able to infer it from partial information.

(3) Modeling

The use of models is a time-renowned problem solving technique. For example, architects and ship builders use models to get answers that would be too difficult to obtain using purely formal methods. We would like to draw an analogy between the architect's use of a physical model and the expert system's use of a simulation structure. In both cases the advantages to be gained are power and efficiency in reasoning about their domains.

Most knowledge representation systems store assertions in a uniform, domain-independent formalism like predicate calculus or semantic networks or frames. While there are advantages to uniformity and domain independence, these representations are in many cases considerably less efficient than specialized data structures, and the associated algorithms are often less efficient and less powerful. We are proposing to develop a systematic way of describing when well-known data representations and well-known algorithms are applicable and to devise a program able to employ simulation structures automatically in representing knowledge, given the abstractions it satisfies.

(4) Analogies

Many analogies are best understood as statements that the situations being compared share a common abstraction. For example, when one asserts that the organization chart of a corporation is like Linnaean taxonomy, what he is saying is that they are both hierarchies.

This view of analogy can be turned around and used to help novice users of our abstraction library in finding appropriate entries. Imagine an engineer describing the classification of time in geology (epochs, eras, periods, etc.) who can tell the system that his knowledge base is like that of biological taxonomy and have it infer and use the hierarchy abstraction.

In order to realize this goal, a number of problems must first be solved. The fundamental problem is completing a partial interpretation of an abstraction. Once we have a method for completing interpretations, analogy understanding (or at least the bit of it we are considering) becomes easy. The system merely checks each of the abstractions of the comparison domain, testing to see whether it is applicable.

Sometimes the system may not have a suitable prestored abstraction, and this process will fail. Understanding an analogy in this situation requires the invention of a new abstraction. We are interested in applying and extending the concept formation techniques of Hayes-Roth, Mitchell, and Dieterrich and Michalski in building a program to formulate new abstractions automatically. Of course, a new abstraction will not initially have any specialized data structures or algorithms, but it can provide the next system builder with the techniques developed by the originator.

6.3.4 Explanation

Our motivation for making explanation a primary focus of our research is a belief that expert systems will not be accepted by physicians or scientists unless the systems are able to justify the decisions they make. When important real world domains are involved, human decision makers are loathe to consult machines unless they understand and agree with the basis for the advice. This constraint not only forces us to consider mechanisms for generation of explanations, but it also impacts on the design of the underlying reasoning and representation techniques used by the rest of the consultation system.

In the case of MYCIN and its descendents, we have been able to generate intelligible explanations by taking advantage of our rule-based representation. 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. MYCIN's explanation capabilities have been thoroughly discussed elsewhere [26].

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 totally 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.

During the second half on 1979 we formed a weekly seminar group to analyze the characteristics of good explanations. We generally tried to keep 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 proposed work, then, is to build an explanation system which explicitly addresses all four of these topics. We shall briefly discuss each point:

(1) 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 domain of pharyngitis (sore throats). 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 will provide the basis for further discussions in Year 1 of the proposed work. 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 thyroid disease, or of the pituitary adrenal axis. Both these domains are 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. During Year 02 we will encode this knowledge base in detail and begin experiments on the generation of explanations using the kinds of techniques outlined above.

(2) 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."

(3) 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.

(4) 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.

7 Available Facilities

The existing SUMEX-AIM computer and communications facilities have been described in earlier sections. The number of personnel to support this follow-on work will remain at approximately the same level as before so no additional office space will be required. The additional equipment (VAX's, file server, and PWS's) will be accommodated in the existing SUMEX machine room, a portion of the Pine Hall machine room allocated to Prof. Feigenbaum, and in existing individual office areas. Technician support and hardware development for this equipment will be housed in the existing SUMEX electronics laboratory.

8 Literature Cited

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, Stanford University Department of Computer Science, forthcoming.

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. Stefik, M., An Examination of a Frame-Structured Representation System, Proceeding of the Sixth International Joint Conference on Artificial Intelligence, Vol. 2, 845, August 1979.

10. Metcalfe, R.M. and Boggs, D.R., Ethernet: Distributed Packet Switching for Local Computer Networks, Comm. ACM, Vol. 19, No. 7 (July 1976).

11. 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.

12. Taft, E.A., Implementation of PUP in TENEX, Internal XEROX PARC memorandum, June 1978.

13. 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.

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

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

16. Mead, C. and Conway, L., Introduction to VLSI Systems, Addison-Wesley Publishing Co. (1980).

17. Rosen, B., PERQ: A Commercially Available Personal Scientific Computer, COMPCON 1980.

18. Thacker, C. P., McCreight, E. M., Lampson, B. W., Sproull, R. F., and Boggs, D. R., ALTO: A Personal Computer, Computer Structures: Readings and Examples (Siewiorek, Bell, and Newell, eds.), 1979.

19. McDaniel, The Dorado: A Compact High-Performance Personal Computer for Computer Scientists, COMPCON 1980.

20. Greenblatt, R., MIT's LISP Machine, COMPCON 1980.

21. Ward, S. and Terman, C., An Approach to Personal Computing, COMPCON 1980.

22. Lenat, D., "AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search, Ph.D. Dissertation, Stanford University, July 1976.

23. Stefik, M.J., Planning With Constraints, Ph.D. dissertation, Stanford University, January 1980.

24. Friedland, P.E., Knowledge-Based Experiment Design In Molecular Genetics, Ph.D. dissertation, Stanford University, October 1979.

Literature Cited

25. Sacerdoti, E.D., Problem Solving Tactics, Invited Lecture, Proceedings of the Sixth International Joint Conference on Artificial Intelligence, IJCAI-79. Available from Computer Science Dept., Stanford University, August, 1979.

26. Scott, A.C., Clancey, W., Davis, R., and Shortliffe, E.H., Explanation Capabilities of Knowledge-Based Production Systems, American Journal of Computational Linguistics, Microfiche 62, Knowledge-Based Consultation Systems, 1977.

Biographical Sketches

The following are biographical sketches for all professional personnel contributing to the SUMEX-AIM resource project. These do not include sketches for any of the individual collaborating project investigators.

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME ACHENBACH, Michael W.	TITLE System Programmer	BIRTHDATE *(Mo., Day, Yr.)* August 2, 1952
PLACE OF BIRTH *(City, State, Country)* Los Angeles, California, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Stanford University	B.S.	1974	Physics
Stanford University	M.A.	1975	Education

HONORS

MAJOR RESEARCH INTEREST Network communications, Small machines	ROLE IN PROPOSED PROJECT System Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1978 - present	System Programmer, SUMEX Computer Project, Department of Genetics, Stanford University School of Medicine
1975 - 1978	Scientific Programmer, Instrumentation Research Laboratories, Department of Genetics, Stanford University School of Medicine
1975	Scientific Programmer, Institute for Mathematical Studies in the Social Sciences, Stanford University

PUBLICATIONS

Smith, D.H., Achenbach, M., Yeager, W.J., Anderson, P.J., Fitch, W.L., and Rindfleisch, T.C.: Quantitative Comparison Gas Chromatographic/ Mass Spectrometric Profiles of Complex Mixtures, Anal. Chem., 49, 1623, 1977.

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161:3024

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME AIELLO, Nelleke T.G.K.	TITLE Scientific Programmer	BIRTHDATE *(Mo., Day, Yr.)* March 21, 1949
PLACE OF BIRTH *(City, State, Country)* Amsterdam, The Netherlands	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☐ Male ☒ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
University of California, Santa Cruz	B.A.	1971	Mathematics
University of California, Santa Cruz	B.A.	1971	Information and Computer Science
University of Utah, Salt Lake City	M.S.	1972	Computer Science

HONORS

Departmental Honors, Information and Computer Science, University of California

Crown College Honors, University of California

MAJOR RESEARCH INTEREST Building intelligent systems Knowledge engineering	ROLE IN PROPOSED PROJECT Scientific Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1977 - present	Scientific Programmer, Heuristic Programming Project, Computer Science Department, Stanford University
1972 - 1977	Programmer, Bolt, Beranek and Newman, Inc.
1973 - 1975 Summers	Teaching Assistant, Structured Programming, University of California Extension
Summer 1972	Teaching Assistant, Compiler Writing, University of California Extension
1971	Programmer, Shell Benelux Centre, De Hage, The Netherlands

PUBLICATIONS (See continuation page)

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161:3024

BIOGRAPHICAL SKETCH - AIELLO, Nelleke T.G.K.
PUBLICATIONS

1. Aiello, N: An Analysis of Notations for Music Applicable to the Digital Control of Electronic Musical Instruments. Masters Thesis, University of Utah, 1972.
2. Collins, A.M., Warnock, E.L., Aiello, N., and Miller, M.L.: Reasoning from Incomplete Knowledge, In D. Bobrow and A.M. Collins (Eds.) REPRESENTATION AND UNDERSTANDING STUDIES IN COGNITIVE SCIENCE. New York, Academic Press, Inc., 1975.
3. Nii, H.P. and Aiello, N.: AGE (Attempt to Generalize): Profile of the AGE-0 System. Stanford Heuristic Programming Project Memo HPP-78-5 (Working Paper), June 1978.
4. Nii, H.P. and Aiello, N.: AGE: A knowledge-based program for building knowledge-based programs. Proc. of IJCAI-6, pp 645-655, 1979.
5. Aiello, N., Nii, H.P. and White, W.C.: The Joy of AGE-ing: An Introduction to AGE-1. Stanford Heuristic Programming Project Memo (work in progress), May 1980.

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME BUCHANAN, Bruce G.	TITLE Adjunct Professor Computer Science	BIRTHDATE *(Mo., Day, Yr.)* July 7, 1940
PLACE OF BIRTH *(City, State, Country)* St. Louis, Missouri, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Ohio Wesleyan University	B.A.	1961	Mathematics
Michigan State University	M.A.	1966	Philosophy
Michigan State University	Ph.D.	1966	Philosophy

HONORS

(see continuation page)

MAJOR RESEARCH INTEREST Artificial Intelligence	ROLE IN PROPOSED PROJECT Technical Director of Core Research

RESEARCH SUPPORT *(See instructions)*

(see continuation page)

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1976 - present	Adjunct Professor, Computer Science Department, Stanford University
1972 - 1976	Research Computer Scientist, Computer Science Department, Stanford University
1966 - 1971	Research Associate, Artificial Intelligence Project, Stanford University

PUBLICATIONS (see continuation page)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161:3024

BIOGRAPHICAL SKETCH - BUCHANAN, Bruce G.

RECENT HONORS

- Editorial Board, Artificial Intelligence: An International Journal
- American Association for Artificial Intelligence - Organizing Committee, Program Committee and Membership Chairman
- Chairman of Program Committee, IJCAI-79 (International Joint Conference on Artificial Intelligence, Tokyo, 1979)
- Invited Colloquium Speaker:
  - University of Maryland
  - Carnegie-Mellon University
  - Rutgers University
  - University of California at Berkeley
  - Michigan State University
- Invited Speaker:
  - AISB Annual Conference (Amsterdam, July 1980)
  - Workshop on the Logic of Discovery and Diagnostics in Medicine (Pittsburgh, October 1978)
  - Douglass College Seminars for Faculty (Rutgers University, 1978)
  - Workshop on Pattern Directed Inference Systems (Honolulu, 1977)
- Recipient, National Institutes of Health Career Development Award (1971-1976)

MEMBERSHIPS

- American Association for Artificial Intelligence (AAAI)
- Cognitive Science Society
- Association for Computing Machinery (ACM), SIGART
- Philosophy of Science Association

RESEARCH SUPPORT

		Funding			
Grant No.	Title of Project	Current Year	Project Period	% of Effort	Grant Agency
1P01 LM 03395-01	Research Program : Biomedical Knowledge Representation	$ 99,484 (7/79-6/80)	$497,420 (7/79-6/84)	10	NLM
MCS-7903753	Knowledge-Based Consultation Systems	$ 73,659 (7/79-6/80)	$ 73,659 (7/79-6/80 + 6 months)	10	NSF
N00014-79-C-0302	Exploration of Tutoring and Prob. Solv. Strategies in Intelligent Computer-Aided Instruction	$396,325 (3/79-3/82)	$396,325 (3/79-3/82)	10	ONR
MDA 903-80-C-0107	Heuristic Programming Project	$496,256 (10/79-9/80)	$1,613,588 (10/79-9/82)	40	ARPA
5R24 RR00612-10	Resource-Related Research - Computers and Chemistry	$221,255 (5/80-4/81)	$641,419 (5/80-4/83)	5	NIH

Selected Publications

Edward H. Shortliffe, Bruce G. Buchanan, and Edward A. Feigenbaum, "Knowledge Engineering for Medical Decision Making: A Review of Computer-Based Clinical Decision Aids," Proceedings of the IEEE, September, 1979.

Bruce G. Buchanan, "Issues of Representation in Conveying the Scope and Limitations of Intelligent Assistant Programs." In J.E. Hayes, D. Michie, and L.I. Mikulich (eds.), Machine Intelligence 9: Machine expertise and the human interface. New York: John Wiley, 1979.

Bruce G. Buchanan and Edward A. Feigenbaum, "DENDRAL and Meta-DENDRAL: Their Applications Dimension," Artificial Intelligence 11, 5, 1978.

Bruce G. Buchanan, Tom M. Mitchell, Reid G. Smith and C. Richard Johnson, Dr., "Models of Learning Systems," in J. Belzer (ed.), Encyclopedia of Computer Sciences and Technology, New York: Marcel Dekker, Inc., 1978, Vol 11.

Randall Davis and Bruce G. Buchanan, "Meta-Level Knowledge: Overview and Applications," Proceedings of the Fifth IJCAI,1,920, August 1977.

Bruce G. Buchanan and Tom Mitchell. "Model-Directed Learning of Production Rules," in D.A. Waterman and F. Hayes-Roth (eds.), Pattern Directed Inference Systems, New York: Academic Press, 1978.

Bruce G. Buchanan and Dennis Smith, "Computer Assisted Chemical Reasoning," in E.V. Ludena, N.H. Sabelli and A.C. Wahl (eds.), Computers in Chemical Education and Research, New York: Plenum Press, 1977, p. 401.

Randall Davis, Bruce Buchanan, Edward Shortliffe, "Production Rules as a Representation of a Knowledge-Based Consultation Program," in Artificial Intelligence, 8, 1, February 1977.

Bruce G. Buchanan, D.H. Smith, W.C. White, R.J. Gritter, E. Feigenbaum, J. Lederberg, and C. Djerassi, "Application of Artificial Intelligence for Chemical Inference XXII. Automatic Rule Formation in Mass Spectronomy by Means of the Meta-DENDRAL Program," Journal of the American Chemical Society, 98, 6168, 1976.

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME FEIGENBAUM, Edward A.	TITLE Professor and Chairman Computer Science Department	BIRTHDATE *(Mo., Day, Yr.)* January 20, 1936
PLACE OF BIRTH *(City, State, Country)* Weehäwken, New Jersey, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Carnegie Institute of Technology, Pittsburgh, Pennsylvania	B.S.	1956	Electrical Engineering
Carnegie Institute of Technology, Pittsburgh, Pennsylvania	Ph.D.	1959	Industrial Administration

HONORS

MAJOR RESEARCH INTEREST Artificial Intelligence	ROLE IN PROPOSED PROJECT Principal Investigator

RESEARCH SUPPORT *(See instructions)*

(See continuation page)

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1976 - present	Professor (by Courtesy) Department of Psychology, Stanford University
1976 - present	Chairman, Department of Computer Science, Stanford University
1969 - present	Professor of Computer Science, Stanford University
1965 - present	Principal Investigator, Heuristic Programming Project, Stanford Universtiy
1965 - 1968	Associate Professor of Computer Science, Stanford University
1965 - 1968	Director, Stanford Computation Center, Stanford University
1964 - 1965	Associate Professor, School of Business Administration, University of California, Berkeley
1960 - 1963	Assistant Professor, School of Business Administration, University of California, Berkeley
1961 - 1964	Research Appointment, Center for Human Learning, University of California, Berkeley
1960 - 1964	Research Appointment, Center for Research in Management Science, University of California, Berkeley
1968 - 1972	Member, Computer and Biomathematical Science Study Section, National Institutes of Health, Bethesda, Maryland
1977 - 1978	Member , Committee on Mathematics in the Social Sciences, Social Science Research Council, New York, New York
1977 - present	Member, Computer Science Advisory Committee, National Science Foundation
1979 - present	Member, Advisory Committee on Mathematics in Naval Research, NRC/ONR

Professional Societies, Consultantships, Publications (see continuation pages.)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

BIOGRAPHICAL SKETCH - FEIGENBAUM, Edward A.

RESEARCH SUPPORT

Grant No.	Title of Project	Funding		% of Effort	Grant Agency
		Current Year	Project Period		
MCS78-02777	MOLGEN: A Computer Science Application to Molecular Genetics	$153,959 (12/79-11/80)	$294.476 (6/78-3/81)	5	NSF
1 P01 LM 03395-01	Research Program; Biomedical Knowledge Representation	$ 99.484 (7/79-6/80)	$497,420 (7/79-6/84)	10	NLM
MDA 903-80-C-0107	Heuristic Programming Project	$496.256 (10/79-9/80)	$1,613,588 (10/79-9/82)	25	ARPA
MCS 7923666	The Automation of Scientific Inference: Heuristic Computing Applied to Protein Crystallography	$54.469 (12/79-11/81)	$54.469 (12/79-11/81)	0	NSF

BIOGRAPHICAL SKETCH - FEIGENBAUM, Edward A.

PROFESSIONAL SOCIETIES

American Association for Artificial Intelligence (President-Elect, 1979-80)
Cognitive Science Society (member, Governing Board, 1979-)
American Psychological Association
American Association for the Advancement of Science
Association for Computing Machinery (member of National Council of ACM, 1966-68)

CONSULTANTSHIPS

Information Sciences Intstitute of University of Southern California
The RAND Corporation
Schlumberger, Inc.
Jaycor, Inc.

BOOKS AND MONOGRAPHS

Handbook of Artificial Intelligence, co-editor with A. Barr, (in final preparation).

Computers and Thought, co-editor with Julian Felman, McGraw-Hill, 1963.

Information Processing Language V Manual, Englewood Cliffs, N.J., Prentice-Hall, 1961 (with A. Newall, F. Tonge, G. Mealy et al).

An Information Processing Theory of Verbal Learning, Santa Monica, The RAND Corporation Paper P-1817, October 1959 (Monograph)

SOME RECENT AND SELECTED PAPERS:

Edward H. Shortliffe, Bruce G. Buchanan, Edward A. Feigenbaum, "Knowledge Engineering For Infectious Disease Therapy Selection" in Proceedings of the IEEE, Vol. 67, No. 9, September 1979.

L. Fagan, J. Kunz, E. Feigenbaum, CSD Stanford University
J.J. Osborn from PMC, San Francisco
"Knowledge Engineering for Dynamic Clinical Settings: Giving Advice in the Intensive Care Unit," submitted to Sixth International Conference on Artificial Intelligence, 1979, February 1979.

E. H. Shortliffe, B.G. Buchanan, E. A. Feigenbaum, "Knowledge Engineering for Medical Decision Making: A Review of Computer-Based Clinical Decision Aids," appeared in the Proceedings of the IEEE, September 1979.

J.C. Kunz, R.J. Fallat, D.H. McClung, J.J. Osborn, B.A. Votteri, H.P. Nii, J.S. Aikins, L.M. Fagan, E.A. Feigenbaum, "A Physiological Rule Based System for Interpreting Pulmonary Function Test Results," Stanford Heuristic Programming Project Memo (144) HPP-78-19.

B.G. Buchanan and E.A. Feigenbaum, "DENDRAL and Meta-DENDRAL: Their Applications Dimension," Artificial Intelligence, 11(1,2)5(1979). (Also Stanford Heuristic Programming Project Memo (126) HPP-78-1).

Feigenbaum, E.A.: The Art of Artificial Intelligence: I. Themes and Case Studies of Knowledge Engineering. Proceedings of the IJCAI, 1977.

Feigenbaum, E.A., Engelmore, R.S. and Johnson, C.K.: A Correlation between Crystallographic Computing and Artificial Intelligence Research. Acta Cryst., A33 (Jan 1): 13-18, 1977. (Also Stanford Heuristic Programming Project Memo (102) HPP-77-15.)

Nii, H.P. and Feigenbaum E.A.: Rule-based Understanding of Signals. Proceedings of the Conference on Pattern-directed Inference Systems, 1977. (Also Stanford Heuristic Programming Project Memo (94) HPP-77-7 and Computer Science Department Memo STAN-CS-77-612.)

Feigenbaum, E.A.: Computer Applications: Introductory Remarks. IN Proceedings of Federation of American Societies for Experimental Biology 33, 2331 (1974) also IN W. Siler and D.A.B. Lindberg (Eds.) Computers in Life Science Research, Plenum Press, 49-51 (1975). (Also Stanford Heuristic Programming Project Memo (57) HPP-74-4.)

Buchanan, B.G., Feigenbaum E.A. and Sridharan, N.S.: Heuristic Theory Formation; Data Interpretation and Rule Formation. IN Machine Intelligence 7, Edinburgh University Press (1972). (Also Stanford Heuristic Prgramming Project Memo (38) HPP-72-2.)

Buchanan, B.G., Feigenbaum, E.A. and Lederberg, J.: A Heuristic Programming Study of Theory Formation in Science. IN Proceedings of the Second International Joint Conference on Artificial Intelligence , Imperial College, London (September, 1971). (Also Stanford Artificial Intelligence Project Memo No. 145, and Heuristic Programming Project Memo (35) HPP-71-4.)

Feigenbaum, E.A., Buchanan, B.G. and Lederberg, J.: On Generality and Problem Solving: A Case Study Using the DENDRAL Program. IN B. Meltzer and D. Michie (Eds.) Machine Intelligence 6, Edinburgh University Press (1971). (Also Stanford Artificial Intelligence Memo No. 131, Heuristic Programming Project Memo (30) HPP-70-5, and Computer Science Memo STAN-CS-176.)

Feigenbaum, E.A.: Artificial Intelligence: Themes in the Second Decade. IN Final Supplement to Proceedings of the IFIP 68 International Congress, Edinburgh, August 1968. (Also Stanford Artificial Intelligence Project Memo No. 67, August 1968, and Heuristic Programming Project Memo (11) HPP-67-3.)

Lederberg, J. and Feigenbaum, E.A.: Mechanization of Inductive Inference in Organic Chemistry. IN B. Kleinmuntz (Ed.), Formal Representtions for Human Judgment (Wiley, 1968). (Also Stanford Artificial Intelligence Project Memo No. 54, August 1967, and Heuristic Programming Project Memo (11) HPP-67-3.)

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME GENESERETH, Michael R.	TITLE Acting Assistant Professor Computer Science	BIRTHDATE *(Mo., Day, Yr.)* October 15, 1948
PLACE OF BIRTH *(City, State, Country)* Philadelphia, Pennsylvania, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Massachusetts Institute of Technology	B.S.	1972	Physics
Harvard University	M.S.	1974	Computer Science
Harvard University	Ph.D.	1978	Applied Mathematics

HONORS

MAJOR RESEARCH INTEREST Computer Science/Artificial Intelligence	ROLE IN PROPOSED PROJECT Core research

RESEARCH SUPPORT *(See instructions)*

(see continuation page)

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1979 - present	Acting Assistant Professor, Department of Computer Science, Stanford University
1978 - 1979	Research Associate and co-Group Leader, Department of Electrical Engineering and Computer Science, M.I.T.
1973 - 1978	Research Assistant, Department of Electrical Engineering and Computer Science, M.I.T.
1971 - 1973	Programmer, Mathlab Group, M.I.T.

PUBLICATIONS (see continuation page)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

BIOGRAPHICAL SKETCH - GENESERETH, Michael R.

RESEARCH SUPPORT

Grant No.	Title of Project	Funding		% of Effort	Grant Agency
		Current Year	Project Period		
MDA 903-80-C-0107	Heuristic Programming Project	$496,256 (10/79-9/80)	$1,613,588 (10/79-9/82)	10	ARPA
MCS-7903753	Knowledge-Based Consultation Systems	$ 73,659 (7/79-6/80 + 6 months)	$ 73,659 (7/79-6/80 + 6 months)	33	NSF
1P01 LM 03395-01	Biomedical Knowledge Representation	$ 99,484 (7/79-6/80)	$497,420 (7/79-6/84)	32	NLM

Selected Papers:

"The Role of Plans in Intelligent Teaching Systems"
- in *Intelligent Teaching Systems*, edited by Derek Sleeman, Academic Press, 1980.
- STAN-CS-784, Stanford Computer Science Dept, Mar. 1980.

"The Use of Semantics in a Tablet-Based Program for Selecting Parts of Mathematical Expressions"
- in Proc. of the Second MACSYMA Users' Conference, M.I.T., June 1979.

"The Canonicality of Rule Systems"
- in Proc. of the European Symposium on Symbolic and Algebraic Manipulation, Springer-Verlag, June 1979.

"Artificial Intelligence Techniques in MACSYMA"
- in *AI Handbook*, edited by Feigenbaum and Barr.

"Automated Consultation for Complex Computer Systems"
- doctoral dissertation, Harvard University, Nov. 1978.

"The Difficulties of Using MACSYMA and the Functions of User Aids"
- Proc. of the First MACSYMA Users' Conference, June 1977.

"A Fast Inference Algorithm for Semantic Networks"
- Memo No. 4, M.I.T. Mathlab Group, 1977.

Invited Talks:

"An Automated Consultant for MACSYMA"
- Stanford Research Institute, April 1979.
- University of Maryland, April 1979.
- Worcester Polytechnic Institute, Jan. 1979.
- M.I.T., Apr. 1978.

"The Role of Plans in Automated Tutors and Consultants"
- Harvard University, Nov. 1978.

"Algebraic Simplification Using MACSYMA"
- White Sands Missile Range, July 1978.
- Sigma Xi Lecture, David W. Taylor Naval Ship R&D Center, Feb. 1978.

"The Simplification of Mathematical Expressions"
- Los Alamos Scientific Laboratory, July 1978.

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME GILMURRAY, Frank S.	TITLE System Programmer	BIRTHDATE *(Mo., Day, Yr.)* July 20, 1948
PLACE OF BIRTH *(City, State, Country)* Brooklyn, New York, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Polytechnic Institute of Brooklyn, New York	B.S.	1970	Electrical Engineering
University of Pittsburgh, Pennsylvania Graduate School (1970-71)	None	--	Computer Science

HONORS

MAJOR RESEARCH INTEREST Operating Systems	ROLE IN PROPOSED PROJECT System Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1977 - present	System Programmer, SUMEX Computer Project, Department of Genetics, Stanford University School of Medicine
1976 - 1977	System Programmer, On-Line Systems, Inc., Pittsburgh, Pennsylvania
1971 - 1976	System Programmer, Computer Center, University of Pittsburgh

PUBLICATIONS (none)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME LENAT, Douglas B.	TITLE Assistant Professor Computer Science	BIRTHDATE *(Mo., Day, Yr.)* September 13, 1950
PLACE OF BIRTH *(City, State, Country)* Philadelphia, Pennsylvania, U.S.A	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
University of Pennsylvania	B.A.	1972	Mathematics
University of Pennsylvania	B.A.	1972	Physics
University of Pennsylvania	M.S.	1972	Applied Mathematics
Stanford University	Ph.D.	1976	Computer Science

HONORS

MAJOR RESEARCH INTEREST Computer Science/Artificial Intelligence	ROLE IN PROPOSED PROJECT Core Research

RESEARCH SUPPORT *(See instructions)*

(see continuation page)

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1979 - present	Consultant to IBM Yorktown, on Maurice Karnaugh's Automatic Programming Effort
1978 - present	Assistant Professor, Computer Science Department, Stanford University
1978	Instructor at General Electric's Program for Modern Managers, Saratoga Springs, N.Y.
1978 - present	Consultant to Schlumberger Oil Co., Ridgefield, Conn.
1978 - present	Consultant to Xerox-PARC's Systems Science Laboratory, Palo Alto, Calif.
1977 - present	Consultant to NIH, as member of their Special Study Section on Biotechnology Resources
1977	Consultant to BBN, Boston, on John Seely Brown's CAI project
1976 - 1978	Assistant Professor, Computer Science Department Carnegie-Mellon University
1976 - present	Consultant to RAND Corp., Santa Monica, Ca., on their "Intelligent Terminal" project

PUBLICATIONS (see continuation page)

BIOGRAPHICAL SKETCH - LENAT, Douglas B.

RESEARCH SUPPORT

Grant No.	Title of Project	Funding Current Year	 Project Period	% of Effort	Grant Agency
1P01 LM 03395-01	Research Program: Biomedical Knowledge Representation	$ 99,484 (7/79-6/80)	$497,420 (7/79-6/84)	10	NLM
MCS78-02777	MOLGEN: A Computer Science Application to Molecular Genetics	$153,959 (12/79-11/80)	$294,476 (6/78-3/81)	20	NSF
MDA 903-80-C-0107	Heuristic Programming Project	$496,256 (10/79-9/80)	$1,613,588 (10/79-9/82)	20	ARPA

BIOGRAPHICAL SKETCH - LENAT, Douglas B.

Published Papers

[1] *Progress Report on Program-Understanding Systems*, Memo AIM-240, CS Report STAN-CS-74-444, Artificial Intelligence Laboratory, Stanford University, August, 1974. Co-authored with Green, Waldinger, Barstow, Elshlager, McCune, Shaw, and Steinberg.

[2] *Synthesis of Large Programs from Specific Dialogues*, Proceedings of the International Symposium on Proving and Improving Programs, IRIA, Le Chesnay, France, July, 1975.

[3] *Duplication of Human Actions by an Interacting Community of Knowledge Modules*, Proceedings of the Third International Congress of Cybernetics and Systems, Bucharest, Romania, August, 1975.

[4] *BEINGS: Knowledge as Interacting Experts*, Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, September, 1975.

[5] *AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search*, Ph.D. Thesis, Stanford A. I. Lab Memo Memo AIM-286, CS Report No. STAN-CS-76-570, and Heuristic Programming Project Report HPP-76-8, Stanford University, July, 1976.

[6] *Designing a Rule System That Searches for Scientific Discoveries*, (Lenat and Harris), invited paper for the conference in Honolulu, May, 1977; published in (Hayes-Roth and Waterman, eds.) Proceedings of the Conference on Pattern-Directed Inference, Academic Press, 1977. Also issued as a CMU technical report, April, 1977.

[7] *Automated Theory Formation in Mathematics*, Fifth IJCAI, Cambridge, Mass., August, 1977.

[8] *Less Than General Production System Architectures*, (Lenat and J. McDermott,) Fifth IJCAI, Cambridge, Mass., August, 1977.

[9] *The Ubiquity of Discovery*, the 1977 Computers and Thought Lecture (invited talk at the Fifth IJCAI). Preliminary version published in the proceedings of that conference; final version printed in the Journal of A.I. Repeated as an invited talk at NCC (Anaheim, June, 1978).

[10] *On Automated Scientific Theory Formation: A Case Study Using the AM Program*, invited paper presented at the Ninth Machine Intelligence workshop in Leningrad, USSR, April, 1977. Forthcoming publication in (Michie, ed.) Machine Intelligence 9, 1978.

[11] Programs that Acquire Expert Knowledge: Two AI Approaches, (Davis & Lenat), McGraw Hill, 1978.

[12] *Pattern Directed Inference Rules the Waves*, Journal of the AISB (Artificial Intelligence Society of Britain), October, 1977, 8-12. Reprinted in SIGART, 1978.

[13] *Rule Based Computation: Some Syntheses*, (Hayes-Roth, Waterman, and Lenat), concluding chapter for (Hayes-Roth and Waterman, eds.) Proceedings of the Conference on Pattern-Directed Inference, Academic Press, 1977.

[14] *Artificial Intelligence and Natural Statistics*, invited paper at "Computer Science and Statistics: Eleventh Annual Symposium on the Interface", University of North Carolina at Raleigh, March 6, 1978.

[15] Unscripted interview on AI & Problem Solving, broadcast over the BBC, as part of the Open University's 32 week course on Cognitive Psychology. Taped at CMU on Feb. 22, 1978, by Clive Holloway, Open University, Milton Keynes, England.

[16] *On Astrophysics and Superhuman Performance* (an invited commentary), Journal of the Behavioral and Brain Sciences, Vol. 1, No. 1, 1978..ss(Societies/committees/awards)

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME LEVINTHAL, Elliott C.	TITLE Adjunct Professor of Genetics Dir., Instrumentation Res. Lab.	BIRTHDATE *(Mo., Day, Yr.)* April 13, 1922
PLACE OF BIRTH *(City, State, Country)* Brooklyn, New York, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Columbia College, New York	B.A.	1942	Physics
Massachusetts Institute of Technology	M.S.	1943	Physics and Math
Stanford University	Ph.D.	1949	Physics and Math

HONORS

Public Service Medal, awarded by NASA, April, 1977, for exceptional contributions to the success of the Viking project

MAJOR RESEARCH INTEREST	ROLE IN PROPOSED PROJECT
Medical instrumentation research	AIM Liaison

RESEARCH SUPPORT *(See instructions)*

Grant No.	Title of Project	Funding Current Year	Funding Project Period	% of Effort	Grant Agency
NSG 7538	Mars Data Analysis	$102,689 (10/79-9/80)	$144,781 (4/79-9/80)	50%	NASA

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1974 - present	Adjunct Professor, Department of Genetics, Stanford University, Director, Instrumentation Research Laboratory, Department of Genetics, Stanford University
1970 - 1973	Associate Dean for Research Affairs, Stanford University School of Medicine
1961 - 1974	Senior Scientist/Director, Instrumentation Research Laboratories, Department of Genetics, Stanford University
1953 - 1961	President, Levinthal Electronic Products
1952 - 1953	Chief Engineer, Century Electronics
1950 - 1952	Research Director/Member of Board of Directors, Varian Associates
1949 - 1950	Research Physicist, Varian Associates
1946 - 1948	Research Associate, Nuclear Physics, Stanford University
1943 - 1946	Project Engineer, Sperry Gyroscope Company, New York
1943	Teaching Fellow in Physics, Massachusetts Institute of Technology

PUBLICATIONS (See continuation page)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

BIOGRAPHICAL SKETCH - LEVINTHAL, Elliott C.
PUBLICATIONS (Selected)

1. Levinthal, E.C., Lederberg, J. and Hundley, L.: Multivator - A Biochemical Laboratory for Martian Experiments. Life Sciences and Space Research II, COSPAR (Committee on Space Research), 1964.

2. Halpern, B., Westley, J.W., Levinthal, E.C. and Lederberg, J.: The Pasteur Probe: An Assay for Molecular Asymmetry. Life Sciences and Space Research, COSPAR (Committee on Space Research), 1966.

3. Levinthal, E.C.: Space Vehicles for Planetary Missions. In Biology and the Exploration of Mars, Nat. Acad. Sci., National Research Council.

4. Levinthal, E.C.: Prospects for Manned Mars Missions. In Biology and the Exploration of Mars, Nat. Acad. Sci., National Research Council.

5. Levinthal, E.C., Lederberg, J. and Sagan, C.: Relationship of Planetary Quarantine to Biological Search Strategy. Presented at COSPAR Meeting (Committee on Space Research), London, 1967.

6. Sagan, C., Levinthal, E.C. and Lederberg, J.: Contamination of Mars. Science 159:1191-1196, 1968.

7. Levinthal, E.C.: The Role of Molecular Asymmetry in Planetary Biological Exploration. Presented at Gordon Research Conferences, Nuclear Chemistry Section, 1968.

8. Mutch, T.A., Binder, A.B., Huck, F.O., Levinthal, E.C., Morris, E.C., Sagan, C. and Young, A.T.: Imaging Experiment. Icarus 16:92, 1972.

9. Levinthal, E.C., Green, W.B., Cuts, J.A., Jahelka, E.D., Johnsen, R.A., Sander, M.J. Seidman, J.B., Young, A.T. and Soderblom, L.A.: Mariner 9 - Image Processing and Products. Icarus 18:1088, 1973.

10. Sagan, C., Veverka, J., Fox, P., Dubisch, R., French, R., Gierasch, P., Quam, L., Lederberg, J., Levinthal, E., Tucker, R., Eross, B. and Pollack, J.B.: Variable Features on Mars, 2, Mariner 9 Global Results. J. Geophysical Research 78, No. 20, p. 4163-4196, 1973.

11. Lederberg, J., Feigenbaum, E., Levinthal, E. and Rindfleisch, T.: SUMEX - A Resource for Application of Artificial Intelligence in Medicine. Proc. Ann. Conference, Association for Computing Machinery, November, 1974.

12. Levinthal, E.C., Carhart, R.E., Johnson, S.M. and Lederberg, J.: When Computers Talk to Each Other. Industrial Research 17(12):35-42, 1975.

13. Mutch, T.A., Binder, A.B., Huck, F.O., Levinthal, E.C., Liebes, S. Morris, E.C., Patterson, W.R., Pollack, J.B., Sagan, C. and Taylor, G.R.: The Surface of Mars: The View from the Viking I Lander. Science 193(4255):791-801, 1976.

14. Mutch, T.A., Arvidson, R.E., Binder, A.B., Huck, F.O., Levinthal, E.C., Liebes, S., Morris, E.C., Nummedal, D., Pollack, J.B. and Sagan, C.: Fine Particles on Mars: Observations with the Viking I Lander Cameras. Science 194(4260): 87-91, 1976.

15. Mutch, T.A., Arvidson, R.E., Aurin, P., Binder, A.B., Huck, F.O., Levinthal, E.C., Liebes, S., Morris, E.C., Pollack, J.B., Sagan, C. and Saunders, R.: The Surface of Mars: The View from Lander 2. Science 194(4271):1277-1283, 1976.

16. Levinthal, E.C., Green, W., Jones, K.L. and Tucker, R.: Processing the Viking Lander Camera Data. Jour. Geophys. Res., No. 28, 30 Sept. 1977.

17. Levinthal, E.C., Jones, K.L., Fox, P. and Sagan, C.: Lander Imaging as a Detector of Life on Mars. Jour. Geophys. Res. 82, No. 28, 30 Sept. 1977.

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME NII, H. Penny	TITLE Research Associate Computer Science	BIRTHDATE *(Mo., Day, Yr.)* October 6, 1939
PLACE OF BIRTH *(City, State, Country)* Tokyo, Japan	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☐ Male ☒ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Tufts University, Jackson College Medford, Massachusetts	B.S.	1962	Mathematics
Stanford University	M.A.	1973	Computer Science

HONORS

MAJOR RESEARCH INTEREST Knowledge-based computer systems design	ROLE IN PROPOSED PROJECT Core Research

RESEARCH SUPPORT *(See instructions)*

		Funding			
Grant No.	Title of Project	Current Year	Project Period	% of Effort	Grant Agency
MDA 903-80-C-0107	Heuristic Programming Project	$496,296 (10/79-9/80)	$1,613,588 (10/79-9/82)	20	ARPA

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1977 - present	Research Associate, Heuristic Programming Project, Department of Computer Science, Stanford University
1976 - 1977	Scientific Programmer,Heuristic Programming Project, Department of Computer Science, Stanford University
1973 - 1975	Associate Investigator for Computer Science, HASP Project, Systems Control, Inc., Palo Alto, California
1967 - 1968	Systems Engineering Advisor, International Business Machines Corporation, Tokyo, Japan
1962 - 1967	Research Staff Programmer. International Business Machines Corporation, Thomas J. Watson Research Center. 1965-67 Project Leader, Electronic Coding Pad (ECP) System 1965-66 Assistant Manager, Man-Computer Interaction Group 1963-64 Programmer, World's Fair Lexical Processing System 1962-63 Programmer, applications ranging from text processing to linear programming problems

RECENT PUBLICATIONS (See continuation page)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161:3024

RECENT PUBLICATIONS

Nii, H. P. and Aiello, N., "AGE: A Knowledge-based Program for Building Knowledge-based Programs," Proc. of IJCAI-6, 1979, pp.645-655.

Kunz, J.C., Fagan, L.M., Fallat, R.J., McClung, D.H., Aikins, J.S., Nii, H.P., Feigenbaum, E.A., Osborn, J.J., "Use of Artificial Intelligence for Interpretation of Physiological Measurments: Pulmonary Function Diagnosis and I.C.U. Ventilator Management," (to be published); abstract in Proc. of NCC, 1978, pp. 260-261.

Nii, H.P. and Feigenbaum E.A., "Knowledge-based Understanding of Signals", in Pattern-Directed Inference Systems, D.A. Waterman and F. Hayes-Roth (eds.), NY: Academic Press, 1978.

Engelmore, R.A. and Nii, H.P., "A Knowledge-based System for the Interpretation of Protein X-ray Crystallographic Data", Heuristic Programming Project Memo; HPP-77-2, (also STAN-CS-77-589), January 1977.

Feigenbaum, E.A., Nii, H.P., et al., "HASP (Heuristic Adaptive Surveilance Program) Final Report, Vols. I-IV, Technical Report under ARPA Contract M66314-74-C-1235, Systems Control, Inc., Palo Alto, CA., 1975 (Classified document).

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME RINDFLEISCH, Thomas C.	TITLE Senior Research Associate	BIRTHDATE *(Mo., Day, Yr.)* December 10, 1941
PLACE OF BIRTH *(City, State, Country)* Oshkosh, Wisconsin, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Purdue University, Lafayette, Indiana	B.S.	1962	Physics
California Institute of Technology, Pasadena	M.S. Ph.D.	1965 Thesis to be completed; all course work and examinations completed.	Physics

HONORS

Graduated with Highest Honors, Purdue University
NSF Fellowship, Caltech
Sigma Xi

MAJOR RESEARCH INTEREST Computer science applications in medical research; image processing and artificial intelligence	ROLE IN PROPOSED PROJECT Facility Manager

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

Stanford University:

1978 - present	Senior Research Associate, Computer Science Department
1976 - present	Senior Research Associate, Genetics Department, School of Medicine
1974 - present	Director, SUMEX Computer Project, Genetics Department
1971 - 1976	Research Associate, Genetics Department:
	1974 - 1976 SUMEX Computer Project
	1971 - 1976 Mass Spectrometry, Instrumentation Research Labs.

Jet Propulsion Laboratory, California Institute of Technology, Pasadena:

1969 - 1971	Supervisor, Image Processing Development and Applications Group
1968 - 1969	Mariner Mars 1969 Cognizant Engineer for Image Processing
1962 - 1968	Engineer, design and implement image processing computer software

PUBLICATIONS (see continuation page)

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161.3024

BIOGRAPHICAL SKETCH - RINDFLEISCH, Thomas C.

PUBLICATIONS

1. Rindfleisch, T. and Willingham, D.: A Figure of Merit Measuring Picture Resolution. JPL Technical report 32-666, September, 1965.
2. Rindfleisch, T.: A Photometric Method for Deriving Lunar Topographic Information. JPL Technical Report 32-786, September, 1965.
3. Rindfleisch, T. and Willingham, D.: A Figure of Merit Measuring Picture Resolution. Advances in Electronics and Electron Physics, Vol. 22A, Photo-Electronic Image Devices, Academic Press, 1966.
4. Rindfleisch, T.: Photometric Method for Lunar Topography. Photogrammetric Engineering, March, 1966.
5. Rindfleisch, T.: Generalizations and Limitations of Photoclinometry. JPL Space Science Summary, Vol. III, 1967.
6. Rindfleisch, T.: The Digital Removal of Noise from Imagery. JPL Space Science Summary 37-62, Vol. III, 1970.
7. Rindfleisch, T.: Digital Image Processing for the Rectification of Television Camera Distortions. Astronomical Use of Television-Type Image Sensors. NASA Special Publication SP-256, 1971.
8. Rindfleisch, T., Dunne, J., Frieden, H., Stromberg, W. and Ruiz, R.: Digital Processing of the Mariner 6 and 7 Pictures. J. Geophysical Research, Vol. 76, No. 2, January, 1971.
9. Pereira, W.E., Summons, R.E., Reynolds, W.E., Rindfleisch, T.C. and Duffield, A.M.: The Quantitation of Beta-Aminoisobutyric Acid in Urine by Mass Fragmentography. Clinica Chimica Acta, 49, 1973.
10. Summons, R.E., Pereira, W.E., Reynolds, W.E., Rindfleisch, T.C. and Duffield, A.M.: Analysis of Twelve Amino Acids in Biological Fluids by Mass Fragmentography. Analytical Chemistry, Vol. 46, No. 4, April, 1974.
11. Pereira, W.E., Summons, R.E., Rindfleisch, T.C. and Duffield, A.M.: The Determination of Ethanol in Blood and Urine by Mass Fragmentography. Clin. Chim. Acta, 51, 1974.
12. Pereira, W.E., Summons, R.E., Rindfleisch, T.C., Duffield, A.M., Zeitman, B, and Lawless, J.G.: Stable Isotope Mass Fragmentography: Quantitation and Hydrogen-Deuterium Exchange Studies of Eight Murchison Meteorite Amino Acids. Geochem. et Cosmochim. Acta, 39, 163, 1975.

PUBLICATIONS (continued)

13. Dromey, R.G., Stefik, M.J., Rindfleisch, T.C. and Duffield, A.M.: Extraction of Mass Spectra Free of Background and Neighboring Component Contributions from Gas Chromatography/Mass Spectrometry Data. Analytical Chemistry, 48, 1368, 1976.

14. Smith, D.H., Achenbach, M., Yeager, W.J., Anderson, P.J., Fitch, W.L. and Rindfleisch, T.C.: Quantitative Comparison of Combined Gas Chromatographic/Mass Spectrometric Profiles of Complex Mixtures. Anal. Chem., 49, 1623, 1977.

15. Smith, D.H., Rindfleisch, T.C. and Yeager, W.J.: Exchange of Comments: Analysis of Complex Volatile Mixtures by a Combined Gas Chromatography-Mass Spectrometry System. Anal. Chem., 50, 1585, 1978.

16. Rindfleisch, T.C. and Smith, D.H.: Chapter 3. In G.R. Waller (Ed.) Biomedical Applications of Mass Spectrometry. (in press)

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME SHORTLIFFE, Edward H.	TITLE Assistant Professor Medicine Computer Science (by courtesy)	BIRTHDATE *(Mo., Day, Yr.)* August 28, 1947
PLACE OF BIRTH *(City, State, Country)* Edmonton, Alberta, Canada	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Harvard College, Cambridge, Massachusetts	B.A.	1970	Applied Math and Computer Science
Stanford University School of Medicine	Ph.D.	1975	Med. Info. Sciences
Stanford University School of Medicine	M.D.	1976	

HONORS

(see continuation page)

MAJOR RESEARCH INTEREST Computer-based Medical Consultation Systems	ROLE IN PROPOSED PROJECT Co-Principal Investigator

RESEARCH SUPPORT *(See instructions)*

(see continuation page)

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1979 - present	Assistant Professor (by courtesy), Department of Computer Science, Stanford University, Stanford, California
1979 - present	Assistant Professor of Medicine (General Internal Medicine) Stanford University School of Medicine, Stanford, California
1977 - 1979	Resident in Medicine, Stanford University School of Medicine
1976 - 1977	Intern in Medicine, Massachusetts General Hospital, Boston, Mass.
1971 - 1975	Doctoral Research, Medical Scientist Training Program, Stanford University School of Medicine, Stanford, California
1970 - 1971	Research assistant, Drug Interaction (MEDIPHOR) Project, Stanford University School Of Medicine, Stanford, California

PUBLICATIONS (see continuation page)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE 1977–241-161:3024

BIOGRAPHICAL SKETCH - SHORTLIFFE, Edward H.

HONORS

Graduation Magna Cum Laude, Harvard College, June, 1970.

Medical Scientist Training Program, Traineeship, September 1971 - June 1976.

Grace Murray Hopper Award (Distinguished computer scientist under age 30), Association for Computing Machinery, October 1976.

Recipient of Research Career Development Award, National Library of Medicine, July 1979 - present.

RESEARCH SUPPORT

Grant No.	Title of Project	Funding Current Year	 Project Period	% of Effort	Grant Agency
NLM LM03395	Research Program: Biomedical Knowledge Representation	$ 99,484 (7/79-6/80)	$497,420 (7/79-6/84)	50	NLM
----	Explanatory Patterns In Clinical Medicine	$ 20,000 (7/79-12/80)	$ 20,000 (7/79-12/80)	25	KAISER

To support the 75% research time above:

Grant No.	Title of Project	Current Year	Project Period	% of Effort	Grant Agency
NLM LM00048	Symbolic Computation Methods for Clinical Reasoning (RCDA)	$ 39,285 (7/79-6/80)	$196,425 (7/79-6/84)	—	NLM

BIOGRAPHICAL SKETCH - SHORTLIFFE, Edward H.
PUBLICATIONS (Selected)

BOOK

Shortliffe, E.H. Computer-Based Medical Consultations: MYCIN , Elsevier/ North Holland, New York, 1976.

JOURNAL ARTICLES

Shortliffe, E.H., Axline, S.G., Buchanan, B.G., Merigan, T.C., and Cohen, S.N. "An artificial intelligence program to advise physicians regarding antimicrobial therapy". Comput. Biomed. Res. 6:544-560 (1973).

Shortliffe, E.H. and Buchanan, B.G. "A model of inexact reasoning in medicine." Math. Biosci. 23:351-379 (1975).

Shortliffe, E.H., Davis, R., Axline, S.G., Buchanan, B.G., Green, C.C., and Cohen, S.N. "Computer-based consultations in clinical therapeutics: explanation and rule-acquisition capabilities of the MYCIN system." Comput. Biomed. Res. 8:303-320 (1975).

Davis, R., Buchanan, B.G., and Shortliffe, E.H. "Production rules as an approach to knowledge-based consultation systems." Artificial Intelligence 8:15-45 (1977).

Scott, A.C., Clancey, W., Davis, R., and Shortliffe, E.H. "Explanation capabilities of knowledge-based production systems." Amer. J. Computational Linguistics, Microfiche 62, 1977. Also available as TR HPP-77-1, Heuristic Programming Project, Stanford University, March 1977.

Wraith, S.M., Aikins, J.S., Buchanan, B.G., Clancey, W.J., Davis, R., Fagan, L.M., Hannigan, J.F., Scott, A.C., Shortliffe, E.H., vanMelle, W.J., Yu, V.L., Axline, S.G., and Cohen, S.N. "Computerized consultation system for selection of antimicrobial therapy." Amer. J. Hosp. Pharm. 33:1304-1308 (1976).

Yu, V.L., Buchanan, B.G., Shortliffe, E.H., Wraith, S.M., Davis, R., Scott, A.C., Axline, S.G., and Cohen, S.N. "Evaluating the performance of a computer-based consultant." Comput. Prog. Biomed. 9:95-102 (1979).

Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. "Knowledge engineering for medical decision making: A review of computer-based clinical decision aids." Proceedings of the IEEE, 67:1207-1224 (1979).

Shortliffe, E.H. "The computer as clinical consultant" (editorial). Arch. Int. Med. 140:313-314 (1980).

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

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME SWEER, Andrew J.	TITLE System Programmer	BIRTHDATE *(Mo., Day, Yr.)* March 12, 1945
PLACE OF BIRTH *(City, State, Country)* Washington, D.C., U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
University of Pittsburgh, Pennsylvania	B.S.	1965	Mathematics
University of Pittsburgh, graduate school (1965-66)	None	--	Mathematics, Computer Science

HONORS

MAJOR RESEARCH INTEREST Operating systems	ROLE IN PROPOSED PROJECT System Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1976 - present	Head System Programmer, SUMEX Computer Project, Department of Genetics, Stanford University
1974 - 1975	Senior Systems Designer, ILLIAC IV Project, Evans and Sutherland
1970 - 1974	Systems Analyst Supervisor, Computer Center, University of Pittsburgh
1968 - 1969	Computer Specialist, Office of Personnel Operations, Department of the Army, Headquarters the Pentagon
1966 - 1968	Systems Programmer/Analyst, Computer Center, University of Pittsburgh

PUBLICATIONS (none)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME TUCKER, Robert B.	TITLE System Programmer	BIRTHDATE *(Mo., Day, Yr.)* June 12, 1940
PLACE OF BIRTH *(City, State, Country)* Seattle, Washington, U.S.A.	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
Stanford University	B.S.	1962	Mathematics

HONORS

MAJOR RESEARCH INTEREST Network Communications Digital Image Processing	ROLE IN PROPOSED PROJECT System Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

Department of Genetics, Stanford University School of Medicine:

1977 - present	System Programmer, SUMEX Computer Project
1965 - 1977	Scientific Programmer, Instrumentation Research Laboratories

PUBLICATIONS (see continuation pages)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977—241-161:3024

PUBLICATIONS

Tucker, Robert B. "A Mass Spectrometer Data Acquisition and Analysis System." Stanford Inst. Res. Lab. Tech. Report IRL-1063, NASA CR-94919, CFSTI Accession N-68-25743, 1968.

Reynolds, W., Bridges, J., Tucker, R. and Coburn, T. "Computer Control of Mass Analyzers." 16th Annual Conference on Mass Spectrometry and Allied Topics, ASTM Committee E-14, NASA CR-96821, 1968.

Reynolds, W., Bacon, V., Bridges, J., Coburn, T., Halpren, B., Lederberg, J., Levinthal, E., Steed, E., and Tucker, R. "A Computer Operated Mass Spectrometer System." Analytical Chemistry, vol 42, pp 1122-1129, Sept. 1970.

Quam, L., Liebes, S., Tucker, R., Hannah, M., and Eross, B., "Computer Interactive Picture Processing." Stanford Artificial Intelligence Project Memo. AIM-166." 1972.

Sagan, C., Veverka, J., Fox, P., Dubisch, R., Lederberg, J., Levinthal, E., Quam, L., Tucker, R., Pollack, J. and Smith, B. "Variable Features on Mars: Preliminary Mariner 9 Television Results." Icarus, vol 17, pp 346-372, 1972.

Quam, L., Tucker, R., Eross, B., Veverka J. and Sagan, C. "Mariner 9 Picture Differencing at Stanford." Sky and Telescope, vol 46 no. 2, August 1973.

Sagan, C., Veverka, J., Fox, P., Dubisch, R., French, R., Gierasch, P., Quam, L., Lederberg, J., Levinthal, E., Tucker, R., Eross, B. and Pollack, J. "Variable Features on Mars, 2, Mariner 9 Global Results." Journal of Geophysical Research, vol 70, no. 20, pp 4163-4196, 1973.

Veverka, J., Sagan, C., Quam, L., Tucker, R. and Eross, B. "Variable Features on Mars III: Comparison of Mariner 1969 and Mariner 1971 Photography." Icarus, vol 21, pp 317-368, 1974.

Sagan, C., Veverka, J., Steinbacher, R., Quam, L., Tucker, R. and Eross, B. "Variable Features on Mars IV. Pavonis Mons." Icarus, vol 22, pp 24-47, 1974.

Veverka, J., Noland, M., Sagan, C., Pollack, J., Quam, L., Tucker, R., Eross, B., Duxbury, T. and Green, W. "A Mariner 9 Atlas of the Moons of Mars." Icarus, vol 23, no. 2, pp 206-289, 1974.

Veverka, J., Sagan, C., Quam, L., Tucker, R. and Eross, B. "The Changing Surface of Mars." Astronomy, vol 3, no. 6, June 1975.

Mutch, T. A., et al. "The Surface of Mars: The View from the Viking 2 Lander." Science, vol 194, pp 1277-1283, 17 Dec. 1976.

PUBLICATIONS (continued)

Levinthal, E., Green, W., Jones, K., Tucker, R. "Processing the Viking Lander Camera Data." Journal of Geophysical Research, vol 82, no. 28, Sept. 1977.

Tucker, Robert B. "More on the Viking Mission." Keyboard, 1977/2 pp 1-4, 1977 (Hewlett-Packard).

Tucker, Robert B. "Viking Lander Imaging Investigation Picture Catalog of Primary Mission Experiment Data Record." NASA Reference Publication 1007, 568 pp, 1978.

SECTION II – PRIVILEGED COMMUNICATION

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME VEIZADES, Nicholas	TITLE R&D Engineer Instrumentation Research Labs.	BIRTHDATE *(Mo., Day, Yr.)* August 25, 1932
PLACE OF BIRTH *(City, State, Country)* Larissa, Greece	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)* U.S. Citizen	SEX ☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
City College of San Francisco, California (1954-55)			
University of California, Berkeley	B.S.	1958	Electrical Engineering
Stanford University	M.S.	1961	Engineering Science

HONORS

MAJOR RESEARCH INTEREST Electronic circuit design	ROLE IN PROPOSED PROJECT Electronics Engineer

RESEARCH SUPPORT *(See instructions)*

		Funding			
Grant No.	Title of Project	Current Year	Project Period	% of Effort	Grant Agency
RR-00612	Resource Related Research - Computers and Chemistry (DENDRAL)	$221,255 (5/80-4/81)	$641,419 (5/80-4/83)	5	NIH

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1962 - present	Electronics Engineer, Department of Genetics, Stanford University School of Medicine: 1978 - present SUMEX Computer Project 1962 - 1978 Instrumentation Research Laboratories
1961 - 1962	Project Engineer, Fairchild Semiconductor (Instrumentation), Division of Fairchild Instrument and Camera Company, Palo Alto, California
1958 - 1961	Senior Engineer, Link Division, General Precision, Inc., Palo Alto, California

PUBLICATIONS (none)

NIH 398 (FORMERLY PHS 398)
Rev. 1/73

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161:3024

BIOGRAPHICAL SKETCH

(Give the following information for all professional personnel listed on page 3, beginning with the Principal Investigator. Use continuation pages and follow the same general format for each person.)

NAME	TITLE	BIRTHDATE *(Mo., Day, Yr.)*
YEAGER, William J.	System Programmer	June 16, 1940

PLACE OF BIRTH *(City, State, Country)*	PRESENT NATIONALITY *(If non-U.S. citizen, indicate kind of visa and expiration date)*	SEX
San Francisco, California, U.S.A.	U.S. Citizen	☒ Male ☐ Female

EDUCATION *(Begin with baccalaureate training and include postdoctoral)*

INSTITUTION AND LOCATION	DEGREE	YEAR CONFERRED	SCIENTIFIC FIELD
University of California, Berkeley	B.A.	1964	Mathematics
California State University, San Jose	M.A.	1967	Mathematics
University of Washington, Seattle Doctoral studies (1969-70)	None	--	Mathematics

HONORS

MAJOR RESEARCH INTEREST	ROLE IN PROPOSED PROJECT
Network communications	System Programmer

RESEARCH SUPPORT *(See instructions)*

RESEARCH AND/OR PROFESSIONAL EXPERIENCE *(Starting with present position, list training and experience relevant to area of project. List all or most representative publications. Do not exceed 3 pages for each individual.)*

1978 - present	System Programmer, SUMEX Computer Project, Department of Genetics, Stanford University School of Medicine
1975 - 1978	Scientific Programmer, Instrumentation Research Laboratories, Department of Genetics, Stanford University School of Medicine
1971 - 1975	Programmer, Bendix Field Engineering, Moffett Field, California
1970 - 1971	Programmer, WELLSCO Data Corp., San Francisco, California
1968 - 1969	Mathematics Instructor, Gavilan Jr. College, Gilroy, California
1967 - 1968	Mathematics Instructor, California Western Univ., San Diego
1966 - 1967	Mathematician/Programmer, Applied Physics Laboratory, Seattle, Washington
1966	Systems Representative, Burroughs Corp., San Jose, California

PUBLICATIONS

Smith, D.H.,Achenbach, M., Yeager, W.J., Anderson, P.J., Fitch, W.L., Rindfleisch, T.: Quantitative Comparison of Combined Gas Chromatographic/Mass Spectrometic Profiles of Complex Mixtures. Anal. Chem,, 49, 1623, 1977.

Smith, D.H., Rindfleisch, T.C. and Yeager, W.J.: Exchange of Comments: Analysis of Complex Volatile Mixtures by a Combined Gas Chromatography-Mass Spectrometry System. Anal. Chem., 50, 1585,1978.

☆U.S. GOVERNMENT PRINTING OFFICE: 1977–241-161.3024

9 Collaborative Project Reports

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 and those using the Rutgers-AIM facility (these 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 331.

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 (8/81 forward)
  B. Justification and requirements for continued SUMEX use
    (This section will be of special importance to the study section and council review of the SUMEX-AIM renewal application)
  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.

9.1 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.

9.1.1 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

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 used in these programs. However, the long-range aspiration is to integrate methods and techniques 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.

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.

Highlights of Research Summary

In addition to the framework for building programs based on the Blackboard model that was available last year, we have added the following additional tools:

1. Framework for building programs that use backward-chained production rules: Backward chaining of production rules is an inference generating mechanism that is used in the MYCIN program (and its offshoots). A simple framework has been implemented in AGE that can be used by itself (i.e. to write MYCIN-like programs) or as a part of a Blackboard based program.

2. Interface to the Units Package: There are kinds of knowledge for which the production rule representation is not suitable. We have augmented the rule-based representation in AGE with frame-like representation, as implemented in the Units package.

The Units data base can be used from the left-hand-sides of rules or can be modified by the right-hand-sides of rules. This combination, in addition to providing another representational form for the frameworks in AGE, provides inference mechanism for Units in the form of rules and other control mechanisms available in AGE.

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.

In addition, to acquaint a variety of users in the use of AGE, three documents are being prepared. They will be available July 1, 1980.

1. "Introduction of Knowledge Engineering, Blackboard Model, and AGE." A high level introduction to knowledge engineering and to the formulation of problems using the Blackboard model.
2. "The Joy of AGE-ing: A User's Guide to AGE-1." An introduction to the use of AGE-1 system.
3. "AGE Reference Manual." A detailed documentation.

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

A three-day workshop was conducted on the week of March 4, 1980 for a limited number of people who had requested access to AGE. Without exception, the attendees represented organizations that wish to build knowledge-based programs, but could not do so because of lack of qualified staff. The aim of the workshop was to familiarize the user with AGE, and for each participant to implement a running program (even if a simple one) related to his own problem. The names of the organizations represented and brief descriptions of the problems for possible implementation on AGE are listed below:

Information Science Group, University of Missouri-Columbia

Interpretation of test results for determining the cause of blood coagulation problems in patient with excessive bleeding. If the interpretation problem can be successfully implemented, they will go on to implement a program that recommend anti-coagulation therapy.

Institute of Medical Electronics, University of Tokyo

Diagnosis of cardiovascular diseases using diverse data and knowledge, and therapy recommendation with re-evaluation diagnosis. In general, this group is interested in building programs that serve as research tools rather than as applied clinical tools.

Department of Psychology, University of Colorado

This groups is using the Blackboard framework in AGE to build a psychological model of prose comprehension. They have been using AGE for about one year.

Oak Ridge National Laboratory

Interpretation of physical signals--non-medical application.

Schlumberger-Doll Research Center

Interpretation of physical signals--non-medical application.

In the process of building AGE, we have used it to write some programs to serve as test programs. Three different versions of PUFF [Feigenbaum 1977; Kunz 1978]--one using the Event-driven control macro, one using the Expectation-driven control macro [Nii 1978], and another using backward-chained productions rules [Shortliffe 1977] were implemented. Since the domain-specific knowledge for PUFF already existed and was implemented in EMYCIN, each AGE version took about a week to bring up--time needed to reorganize the rules into KSs and to rewrite them in the AGE rule syntax. We have also tested a variety of small programs, including programs for cryptogram analysis, determining a bidding strategy for the game of hearts, and a graph traversal problem.

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 user 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 module and editor that asks 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. However, 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 which can be executed with various tracing modes. AGE currently provides no special debugging tools beyond what is available in Interlisp.

EXPLANATION: AGE has enough information to replay its execution steps, and it has reasonable justifications for the actions within the various framework. However, AGE is totally ignorant of the user's task domain and has no means of conducting a dialogue about the task domain. A detailed history of the execution steps is available to the user. The HISTORYLIST can be used in a variety of ways, including the construction of explanations.

```
     SYSTEM KNOWLEDGE      SUBSYSTEM           RESULT
                              |
      +-----------+     +-----V------+
      |  MANUAL   |....>|  BROWSE    |
      +-----------+...  +-----+------+
                     .        |
      +-----------+  .  +-----V------+     +------------+
      |  DESIGN   |....>|  DESIGN    |....>|USER SYSTEM |
      |  SCHEMA   |...  |            |     |   DESIGN   |
      +-----------+  .  +-----+------+     +------+-----+
                     .        |<.................|
      +-----------+  .  +-----V------+     +------------+
      |COMPONENTS |....>| ACQUISITION|....>|    USER    |
      |           |     |            |     |   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

AGE-1 system is now complete, and will be released for general use on July 1. The research and development plan for AGE-2 include the following:

1. Improving the Front-end

Although the current Design subsystem provides specification functions that allow the user to interactively specify the knowledge of the domain and control structure, it does not (aside from simple advice) provide the user any help in the designing process. For example, AGE should be able to provide some heuristics on what kind of inference mechanisms and representation are appropriate for different kinds of problems. We have begun collecting knowledge-engineering heuristics, but much more work is needed in building a design aid that will be useful.

2. Adding More Tools

Our concept of a software laboratory is a facility by which the users are provided with a variety of preprogrammed components that can be combined into problem-solving frameworks--similar in spirit to designs of prefabricated houses. The user can augment and modify a framework to develop his own programs. We currently provide tools for developing programs that use the Blackboard framework and framework for backward-chained inference rules. We have also integrated the Units Package (described elsewhere) to be used within the Blackboard framework. Given the current set of components, other frameworks can, and need to be defined; i.e. other combinations of components that would be useful in solving a wide range of problems. Another inference mechanism, the heuristic search paradigm also need to be added.

3. Performance Test

Although various users have attempted to use the AGE system, it has not been tested for its power and flexibility. For the next three to five years, we will add to our task the development of an application problem complex enough to exercise the variety of components available in the current system.

Computing Resources and Management

I believe the computing and communication resources provide by the SUMEX Facility is 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.

9.1.2 AI Handbook Project

Handbook of Artificial Intelligence

E. A. Feigenbaum and A. Barr
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 Appendix G on page 392).

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 Volume I of the Handbook, which will be going through the final stages of manuscript preparation in the Spring and Summer of 1980. The material in Volume I 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

Most of the chapters in Volume I of the AI Handbook have already appeared in preliminary form as Stanford Computer Science Technical Reports, authored by the respective chapter-editors:

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, Principle 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 our tentative schedule for completion and publication of the AI Handbook:

Spring and Summer, 1980 - Volume I will go through final editing, computer typesetting, and printing.

Fall, 1980 through Spring, 1983 - Volume I will be published. Research for Volume II will be started and draft material will go through the external review process.

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. Your needs and plans for other computational resources

We will use document preparation programs at SUMEX and a xerographic output device at the Stanford Computer Science Department to produce the final copy of the AI Handbook.

D. Recommendations for future community and resource development

None.

9.1.3 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.

I.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.

Our research grant has recently been renewed for a three-year period beginning May 1, 1980. This renewal has come at a particularly opportune time in the development of computer aids to structure elucidation. We are beginning to push our techniques for spectral interpretation, structure generation (e.g., CONGEN) and spectral prediction to their limits within the confines of topological representations of molecular structure. Even so, 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). In order to proceed further in providing to the community 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 will emphasize in the new grant period development of stereochemical extensions to CONGEN, existing related programs and the proposed new programs GENOA and SASES, including machine representations and manipulations of configuration and conformation and constrained generators for both aspects of stereochemistry.

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).

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. The capability for representation of stereochemistry is required for any comprehensive treatment of: 1) interpretation of spectroscopic data; 2) prediction of spectroscopic data; 3) induction of rules relating known molecular structures to observed chemical or biological properties. 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.

I.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.

I.C. Highlights of Research Progress

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

I.C.1. Past Year

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 previous grant year we have completed an exportable version of the CONGEN program and have begun to export 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. Even though we have been exporting versions of CONGEN for only six months, already the program has been used for new structures and recent results have formed the basis for at least four formal lectures by users of CONGEN at remote sites.

2) Version I of the GENOA program for structure generation with overlapping atoms. GENOA 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, because the necessity for manual intervention in the assignment of substructures is no longer required as it was for CONGEN.

3) Exhaustive and irredundant generation of stereoisomers. During the current grant period we have 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. 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, this new development is crucial to use of computational techniques in structural studies. It is interesting to note that this particular problem remained unsolved, until the present work, since it was originally proposed by Van't Hoff more than 100 years ago.

I.C.2. Research in Progress

1) 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.

The stereochemical substructure descriptors are absolutely essential, especially in magnetic resonance data, for either interpretation or prediction. 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.

The data bases can be used to interpret spectral data to obtain substructures to be used in CONGEN and GENOA, the structure generating programs. Automation of this aspect of structure elucidation could 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.

This effort marks the beginnings of the SASES system, a general, automated system for computational assistance in several phases of structure elucidation.

2) Constrained Generation of Configurational Stereoisomers. We have just completed an experimental version of a program, designed to be used with the structure generation programs CONGEN and GENOA, capable of constrained generation of stereoisomers. This means that, 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 (see Highlight 1) 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.

I.D. List of Recent Publications

(1) D.H. Smith and R.E. Carhart, "Structure Elucidation Based on Computer Analysis of High and Low Resolution Mass Spectral Data," in "High Performance Mass Spectrometry: Chemical Applications," M.L. Gross, Ed., American Chemical Society, 1978, p. 325.

(2) T.H. Varkony, D.H. Smith, and C. Djerassi, "Computer-Assisted Structure Manipulation: Studies in the Biosynthesis of Natural Products," Tetrahedron, 34, 841 (1978).

(3) D.H. Smith and P.C. Jurs, "Prediction of 13C NMR Chemical Shifts," J. Am. Chem. Soc., 100, 3316 (1978).

(4) T.H. Varkony, R.E. Carhart, D.H. Smith, and C. Djerassi, "Computer-Assisted Simulation of Chemical Reaction Sequences. Applications to Problems of Structure Elucidation," J. Chem. Inf. Comp. Sci., 18, 168 (1978).

(5) D.H. Smith, T.C. Rindfleisch, and W.J. Yeager, "Exchange of Comments: Analysis of Complex Volatile Mixtures by a Combined Gas Chromatography-Mass Spectrometry System," Anal. Chem., 50, 1585 (1978).

(6) 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).

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

(8) 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," G.R. Waller, Ed., in press.

(9) T.C. Rindfleisch and D.H. Smith, in Chapter 3 of "Biomedical Applications of Mass Spectrometry," G.R. Waller, Ed., in press.

(10) 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).

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

(12) 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.*, in press (1980).

(13) 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.*, in press (1980).

(14) J.G. Nourse, D.H. Smith, and C. Djerassi, "Computer-Assisted Elucidation of Molecular Structure with Stereochemistry," *J. Am. Chem. Soc.*, submitted for publication.

(15) 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).

(16) 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.

(17) 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.

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

(19) N. A. B. Gray, A. Buchs, D. H. Smith, and C. Djerassi, "Computer-Assisted Structural Interpretation of Mass Spectral Data," *Helv. Chim. Acta*, submitted for publication.

(20) 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.*, in preparation.

(21) 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.*, in preparation.

I.E. Funding Support

I.E.1. Title

RESOURCE RELATED RESEARCH: COMPUTERS IN CHEMISTRY (grant)

I.E.2. 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

I.E.3. Funding Agency

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

I.E.4. Grant Identification Number

RR-00612-11

I.E.5. Total Award and Period

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

I.E.6. Current Award and Period

Current - 5/1/80 - 4/30/81 -------- $221,255

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 program. 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 new grant period we have identified the SUMEX-AIM resource 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. In a sense, then, relating our new research to SUMEX formalizes a relationship which already exists. However, such a formalization seems much more relevant now than in the past because of our broader emphasis on software tools and new capabilities for sharing the results of our research. The relationship 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 program and future developments in the DENDRAL Project to yield more generally useful programs will simply magnify the problem.

We will, therefore, under the new relationship between SUMEX-AIM and our project, 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.

II.A. Scientific Collaboration and Program Dissemination

II.A.1. Scientific Collaborations

Several of our research goals involve problems in structural analysis whose solution is of interest to other research groups with specific, health-related problems in structural biochemistry. 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. Dr. David Cowburn, The Rockefeller University. A very likely application for CONGEN enhanced with a conformation generator would be to the field of conformational analysis. This is the problem of determining the conformation of a structure with known constitution and configuration and is a general problem in describing the structures of molecules. The description of the conformation(s) of molecules of biological origin or of those possessing biological activity is of considerable importance in establishing more clearly the relationship of structure to function in the actions of drugs,hormones, and neurotransmitters on their natural receptors, the mechanism of enzyme action, and the rational design of new drugs. We will develop this application in collaboration with Professor David Cowburn and his coworkers at the Rockefeller University in New York. Professor Cowburn is actively engaged in determining peptide conformations using principally nuclear magnetic resonance studies of specifically designed and synthesized isotopic isomers of peptide hormones. These studies use the stable isotopes - deuterium, carbon-13, and nitrogen-15 [91]. Dr. Cowburn now has an account at SUMEX and would use the program remotely, at least at first. It is hoped that an effective collaboration can be developed in which Dr. Cowburn will investigate techniques for effectively rejecting chemically unreasonable conformations as they are generated. Those strategies that may be generally useful will then be adapted for CONGEN and incorporated. These techniques will be related either to general considerations(e.g. insufficient degrees of freedom for cyclization of a particular ring system, from a partially generated conformational state) or to the specific molecules being examined (e.g. restrictions stemming from experimental data such as nmr vicinal coupling constants). Some research using small programs outside CONGEN would be expected to be useful in investigating this area. CONGEN equipped with a conformation generator, would likely be useful to Prof. Cowburn's research in at least three ways:

a) The program would be able to generate all the possible conformations for a given problem with input constraints based on NMR couplings. Such a generation is a difficult task for,e.g, compounds containing large rings. The value of CONGEN would be to provide assurance of exhaustion and to explicitly construct all the possibilities.

b) The program would be able to generate all possible isotopic isomers for a given constitution and configuration. if a pruning technique was available, then the generated list would be extremely useful to Dr. Cowburn in considering the strategies of synthesis and nmr experimentation. The avoidance of particularly costly or time consuming steps is of considerable importance in that experimental work.

c) In conjunction with the spectral interpretation and planning modules proposed, CONGEN may be able to generate strategies for patterns of enrichment or for nmr experiments which are optimum for conformational determination. Some additional programming would probably be necessary to accomplish this.

2. Dr. Gilda Loew, Stanford Research Institute and The Rockefeller University. Since our conformation generator will output structures with internal (torsional angle) coordinates, it is possible to obtain further information about these structures by doing quantum mechanical energy calculations. By developing a link to these methods, the usefulness of CONGEN should be considerably increased. Since a great deal of work has been done by others on such methods it is not necessary for our group to develop programs of this kind. Instead we will develop this link by collaborating with Prof. Gilda Loew and her group. Professor Loew's work has involved the use of semi-empirical quantum mechanical energy calculations to derive structure-activity for a variety of drug types. The first step in such a collaboration would be to construct the interface necessary to link the CONGEN output structures with the input for the PCILO (Perturbation Configuration Interaction using Localized Orbitals) program. This program requires as input, structures with internal coordinates. This will be the form of the output from the proposed conformation generator with an assumption of bond lengths and angles.

Once this link has been made then we see at least two areas where CONGEN might be helpful to Professor Loew's ongoing research.

a) It will be possible to generate systematically variants of a structure with respect to its constitution, configuration, and conformation. Each such structure would then be given to PCILO for an energy calculation, the results of which are used to help explain potency variations [92]. The advantage of using CONGEN in this way is that an exhaustive generation can be guaranteed which assures no possibilities are overlooked.

b) Professor Loew has been considering the conformational variations caused by the intercalation of ethidium into nucleic acids. The observed stability of such intercalated structures has been related to conformational changes in parts of the DNA structure, in particular, the sugar moieties. The application of CONGEN to such a study would again be a systematic variation of possibilities with particular emphasis on the more difficult cyclic structures.

3. 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.

4. 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.

5. 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. Recent use of our programs has aided her in correcting erroneous assignments of 13C resonance shifts to known structure and aided in the solution of the structures of new diterpenoid alkaloids.

6. 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.

7. 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.

II.A.2. Program Dissemination

Because one of our goals is dissemination of our programs to a wide community of collaborators, we have made use of several of the mechanisms provided by SUMEX-AIM to introduce new investigators to our work and to encourage close collaboration in the study of important structural problems. Generally speaking, introduction of new persons and the development of collaborative projects has followed the course outlined below:

1) GUEST Access. The GUEST account mechanism of SUMEX-AIM is normally used when persons from the outside community contact us to learn more about our programs. We provide to them a special packet of information on network access and connection to the GUEST account, together with documentation of specific programs in which they are interested. This is a simple way of performing a "try it and see" experiment to determine the utility of the programs to the individual investigator. The following persons have used this method of access the past year:

Dr. Robert Adamski - Alcon Labs

Dr. A. Bothner-by - Carnegie Mellon University

Dr. Reimar Bruening - Institut fur Pharmazeutische Arzneimittellehre der Universitaet, West Germany

Dr. William Brugger - International Flavors and Fragrances

Dr. Raymond Carhart - Lederle Laboratories

Dr. Robert Carter - University of Lund, Sweden

Dr. Francois Choplin - Institut Le Bel, France

Dr. Jon Clardy - Cornell University

Dr. Mike Crocco - American Hoechst Corp.

Dr. V. Delaroff - Roussel UCLAF, France

Dr. Dan Dolata - University of California at Santa Cruz

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

Dr. Y. Gopichand - University of Oklahoma

Ms. Wendy Harrison - University of Hawaii at Manoa

Dr. Richard Hogue - University of California at Santa Cruz

Dr. David Lynn - Columbia University

Dr. In Ki Mun - Cornell University

Dr. Koji Nakanishi - Columbia University

Dr. Suba Neir - Washington University, St. Louis

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

Dr. Joseph SanFilippo - Rutgers University

Dr. Babu Venkataraghavan - Lederle Laboratories

Dr. W.T. Wipke - University of California at Santa Cruz

Dr. Michael Zippel - Institut fur Biochemie Zentrale
Arbeitsgruppe Spectroskopie, Germany

2) EXODENDRAL Accounts. SUMEX-AIM has set aside a special account group called EXODENDRAL designed to give each collaborator, whose initial GUEST experience has proven fruitful, an account of his or her own. These accounts facilitate both access to a variety of our experimental programs (not generally available through GUEST) and communication using the various message and bulletin board programs. For persons who use exportable versions of our programs on their own computer facilities, EXODENDRAL accounts are used primarily for rapid contact and exchange of messages.

Dr. Jean-Claude Braekman - Universite Libre de Bruxelles,
Belgium

Dr. Hartmut Braun - Organische-Chemisches Institut der
Universitaet Zurich, Switzerland

Dr. Roy Carrington - Shell Biosciences Laboratory, England

Dr. David Cowburn - The Rockefeller University

Dr. Douglas Dorman - Lilly Research Laboratories

Dr. Andre Dreiding - Organische-Chemisches Institut der
Universitaet Zurich, Switzerland

Dr. Janet Finer-Moore - University of Georgia

Dr. Kenneth Gash - California State College at Dominguez Hills

Dr. Steven Heller - Environmental Protection Agency

Dr. Martin Huber - Ciba-Geigy, Switzerland

Dr. Peter W. Milne - CSIRO Division of Computing Research,
Australia

Dr. James Shoolery - Varian Associates

Dr. William Sieber - Sandoz Ltd., Switzerland

Dr. Mark Wood - Rutgers University

3) Program Export. SUMEX-AIM is also the facility which is used to develop and perform experiments with exportable versions of our programs. Wherever possible we encourage collaborators to run our programs on their own computers to decrease the computational burden on SUMEX-AIM as much as

possible. This year 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

Dr. Hartmut Braun - Organische-Chemisches Institut der Universitaet Zurich, Switzerland

Dr. Raymond Carhart - Lederle Laboratories

Dr. Roy Carrington - Shell Biosciences Laboratory, England

Dr. Robert Carter - University of Lund, Sweden

Dr. Daniel Chodosh - Smith, Kline & French Laboratories

Dr. Douglas Dorman - Lilly Research Labs

Dr. Martin Huber - Ciba-Geigy, Switzerland

Dr. Carroll Johnson - Oak Ridge National Laboratory

Dr. G. Jones - ICI Pharmaceuticals, England

Dr. Peter W. Milne - CSIRO Division of Computing Research, Australia

Dr. James Morrison - Latrobe University, Australia

Dr. Fred W. McLafferty - Cornell University

Dr. David Pensak - E.I. duPont de Nemours and Company

Dr. Gretchen Schwenzer - Monsanto Agricultural Products Co.

Dr. William Sieber - Sandoz, Ltd., Switzerland

Dr. M.D. Sutherland - University of Queensland, Australia

Dr. R.O. Watts - Australian National University

4) 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 have begun solicitation of interested industrial organizations to participate in a DENDRAL Project Industrial Affiliates Program. 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 may be 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. As of the writing of this portion of the SUMEX-AIM renewal proposal we have not had any organizations formally take up membership.

II.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.

II.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 proposed 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

III.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).

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;

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;

C) Capabilities for prediction of spectral (and biological) properties to rank-order candidates on the basis of agreement between predicted and observed properties.

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 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);

B) proton magnetic resonance (1HMR);

C) infrared spectroscopy (IR);

D) mass spectrometry (MS)

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. 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.

4) Develop a constrained generator of stereoisomers, 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 and via the National Resource for Computation in Chemistry and the NIH/EPA Chemical Information System;

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.

III.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. 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 SUMEX is our sole source of computational facilities, 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 preliminary 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 this reason, we strongly support the effort of SUMEX to acquire a VAX and other small machines in future years, for all the reasons mentioned above. 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.

III.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 all 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). Until recently, the computational facilities at the NRCC have not been suitable for running interactive programs. Recently, however, the NRCC has obtained a VAX system and we will investigate whether or not the community as a whole will have access to that system. The NRCC is currently under review for continued funding. Obviously that review will have to be favorable for the NRCC to represent an alternative for access to our programs.

III.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.

9.1.4 MOLGEN Project

MOLGEN - A Computer Science Application to Molecular Biology

Profs. E. Feigenbaum, L. Kedes, and D. Brutlag, Dr. P. Friedland
Department of Computer Science
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. Several companies (Genentech, Cetus, Biogen) have already formed to exploit the commercial potential of the burgeoning technology.

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 many others throughout the country (University of Utah, Syracuse University, NIH, Johns Hopkins, Yale, Rockefeller University, and others) 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, for example).

C. Highlights of Research Progress

Accomplishments

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.

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.

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.

Friedland's system is based on the observation that human scientists rarely plan experiments from scratch. They start with an abstracted or "skeletal" plan which contains the entire design in outline form. The major design task is in instantiating or detailing the steps by finding tools that will work best in the given problem environment. This system has roots in classic problem-solving work dating back to Polya, and also in the Scripts language understanding work of Schank and Abelson. It is heavily dependent upon large amounts of domain specific knowledge, especially upon good heuristics for choosing among alternatives for plan-step instantiation.

Stefik's system emphasizes the role that interactions between steps in a plan should have when the plan is being designed. It uses an approach called "constraint posting" to make the interactions between subproblems explicit. Constraints are dynamically formulated and propagated during hierarchical planning and are used to coordinate the solution of nearly independent subproblems. The system also formalizes the problem of control during planning (what to do next) within a structure called "meta-planning". See Appendix B for an annotated example of the system at work.

Knowledge Base Construction

With the experiment design research as an impetus and the Unit System as a tool, a large knowledge base has been constructed by several Stanford molecular biologists--Prof. Douglas Brutlag, Prof. Laurence Kedes, Dr. John Sninsky, and Rosalind Grymes. This knowledge base is near-expert in several areas (enzymatic methods, nucleic acid structures, detection methods) and contains pointers and references to almost all areas of modern molecular biology. Its design and construction will soon be taken over by a full-time molecular biologist.

Besides its use as a fundamental part of an experiment design system, the knowledge base is proving useful for applications in teaching, in automated nucleic acid sequence analysis (see below), and as an intelligent "encyclopedia" for providing information about technique selection in the laboratory.

Other Applications of Symbolic Computation to Molecular Biology

Along with the central research in representation and planning, considerable work has been devoted to the construction of tools that are immediately useful to molecular biologists. Most of these tools were developed at the request of the various domain scientists working on the MOLGEN project and are being used by several dozen scientists both at Stanford and elsewhere through the facilities of the SUMEX computer system.

Interactive tools for nucleic acid sequence analysis--a multi-purpose program for analysis of primary sequence data has been made interactive with full help facilities. The program has also been improved to correctly calculate the expected probability of symmetries and homologies, and to properly allow for GU and GT bonding. A series of smaller programs for similar tasks has also been made interactive on the SUMEX system.

Sequence analysis through the knowledge base--some of the representational tools developed during the process of knowledge base construction (see above), have proven useful for computer-assisted sequence analysis. Facilities are available for building and displaying restriction maps and region information, and for writing rules which cause this information to be automatically updated as new enzymes or structures are added to the knowledge base.

A program for restriction mapping,the GA1 program constructs restriction maps using data from total and partial restriction enzyme digests.

A program was written which aids in enzyme selection for gene excision. The SAFE program takes amino acid sequence data and predicts those restriction enzymes which are guaranteed not to cut within the gene.

A ligase simulation program was written. It is based on a kinetic theory of ligation which helps scientists select time of reaction and concentrations of reaction components to produce single inserts into vectors.

Research in Progress

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.

Experiment Design

Both Friedland's and Stefik's experiment design system have already achieved modest success in producing reasonable plans for a variety of synthetic and analytic problems in molecular biology. Friedland's system can provide technically competent designs for about twenty different types of analytical problems. Stefik's system provides more innovative planning for a single type of synthetic experiment.

We intend to begin to integrate the two systems; Stefik's system will serve as a "front-end" that supplies the skeletal plans that drive Friedland's system. The combination of the two methods should provide a synergistic effect that facilitates both efficiency and innovation.

A second area of improvement in experiment design lies in providing the design systems with a deeper "theory of the domain." We would like design decisions to be made on the basis of mechanism whenever possible; e.g. to denature a molecule pick the best hydrogen bond-breaker, rather than the best pre-stored denaturation method. The current first step in making this improvement is in giving the representation formalism the power to work with sequence and topology of molecules, as described below.

An added benefit of the work on sequence and topology is in giving the planning system the ability to carry out certain steps of experiment designs. Many problems involve one or more steps that can be solved by use of the sequence analysis tools described in the previous section. The design system can make use of these tools directly and sometimes find faster and better solutions than can be achieved in the laboratory.

For example, the sub-problem of finding the right restriction enzymes to excise a gene for cloning can be solved by laborious experimental effort or by a few seconds of automated comparison of the gene with the cutting sites of all of the available restriction enzymes.

Knowledge Base Construction

The current knowledge base contains information about some three hundred laboratory methods and thirty strategies (skeletal plans) for using those methods. It also contains the best currently available data on about forty common phages, plasmids, genes, and other known nucleic acid structures.

We have recently concentrated on providing rules that allow the knowledge base to be automatically updated as new techniques or structures are added (for example, automatically revising restriction maps when a new restriction endonuclease is described). We are also working on mechanisms for facilitating the description of restriction sites and functional regions within molecules. After we are satisfied that our representation method is adequate, rules that model the changing structure of nucleic acid structures during the course of an experiment will be added to the knowledge base.

The knowledge base work to date has all been accomplished with the limited time of several expert molecular biologists, particularly Professors Douglas Brutlag and Laurence Kedes. We have just completed a search for an expert to carry on the knowledge base improvement full time and have hired Dr. Rene' Bach for this role. He will begin work on the MOLGEN project sometime early this summer.

Sequence Analysis

The sequence analysis methods described in the previous section have proven useful to a varied group of users throughout the country over the SUMEX-AIM facility. We will continue to improve these powerful tools and plan to make them available to the scientific community at large on the SUMEX-AIM national resource. If this test is successful, it will demonstrate the need for a full-scale national facility for sequence storage and analysis, and also the ability of MOLGEN to fill that need.

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 MCS 78-02777. Current Principal Investigators are Edward A. Feigenbaum, Professor of Computer Science and Laurence H. Kedes, Investigator, Howard Hughes Medical Institute and Associate Professor of Medicine. The new grant (September 1980) will add Bruce G. Buchanan, Adjunct Professor of Computer Science, and Douglas Brutlag, Associate PRofessor Biochemistry as Co-PI's. MOLGEN is currently funded from 12/79-11/80 at $153,959 including indirect costs and has had a total funding from 6/78-3/81 at $294,476 including indirect costs.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

All system development has 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.

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 combination of the excellent computing facilities and the 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. We do find that we are running moderately close to machine capacity both in size and in speed since our user group has been rapidly expanding during the last year.

III. RESEARCH PLANS

A. Project goals and plans

We have proposed further MOLGEN research in several broad categories: representation, planning, knowledge base development, and immediate applications to molecular biology. As would be expected, there will be much interaction among those general areas.

Representation

As part of the MOLGEN effort, a new representation package, the Units System, has been developed and tested. Its basis was mainly theoretical; we now have the opportunity to improve it from the practical considerations of a large knowledge base containing many different types of information. We expect to learn which features are important and which are window-dressing. These findings will increase in importance as many other problem-solving systems using large domain-specific knowledge bases are developed.

The MOLGEN knowledge base will serve as a laboratory for this research. Among the issues we would like to explore are:

1. MOLGEN currently uses the hierarchy representation features of the Units System for both acquisition and design. Will this continue to be practical as the knowledge base grows, or will the two representation functions have to be divorced?

2. The Units System allows different types of knowledge, e.g. numbers and nucleic acid sequences, to be described and stored in different manners. How much diversity is useful, both from the viewpoint of the representation system and from the viewpoint of the user?

3. Will new features become necessary to make large knowledge bases "perusable" by the human expert describing his domain? Is there some point at which graphics are needed for the expert to have a good grasp of what the system already knows?

Planning

Both of the two problem solving methods developed in MOLGEN have shown promise. We plan to keep pushing their development until we know their respective limitations and until a practical laboratory tool results. As was previously mentioned, we will combine the two planning methods to produce a system which should produce substantially higher performance than either of its two components.

The current experiment design systems are not designed to take an already existing laboratory plan and determine if the plan will satisfy some stated goal. We have proposed using the knowledge base to simulate the result of applying each step of a plan in succession to see if the experiment goal really would be achieved. This sort of a plan verifier will serve to take scientist-designed plans and provide guidance on whether the plan will work before it is actually tried in the laboratory.

The plan verifying system will be extended to become first a plan optimizing system and then a plan debugging system. Plan optimization will involve both domain-specific heuristics about how particular steps interact and domain-free heuristics about what good experiment designs should look like. The plan optimizer will make minor changes and introduce subgoals in order to take an already working experiment design and make it more efficient, convenient, reliable, or inexpensive. The knowledge base already contains most of the raw information humans use to make optimization decisions. The research is in developing the proper methods to make automated use of this knowledge.

Plan debugging means taking a partially working experiment design and finding and fixing any errors in it. This involves aspects of both verification and optimization as well as new error-correction heuristics. According to Feitelson and Stefik, the serendipity of the experimental laboratory also contributes greatly to plan debugging. Extending the MOLGEN design systems to become execution monitoring systems that can note and take advantage of this serendipity will be a major research effort of about thesis level in magnitude.

Knowledge Base Acquisition and Development

The current MOLGEN knowledge base is the result of over a man-year of effort by Professors Douglas Brutlag and Laurence Kedes and Drs. Peter Friedland and John Sninsky. It will continue to grow and improve throughout the term of the new proposal with the full time work of Dr. Rene' Bach. By the end of the period covered in the proposal the knowledge base will be in itself a useful tool for teaching, information retrieval, and sequence analysis. It will be expert in some of the most important areas of molecular biology. It will be especially proficient in those judgmental heuristics that guide technique selection as an experiment is being designed.

A major new research goal is to provide a facility for self-improvement of the knowledge base. When the design system produces a plan that is especially efficient or innovative, it would be useful to generalize and save that plan so that it can drive future problem-solving without having to be reinvented. The generalization and learning process has roots in the MACROPS work in STRIPS.

Having such a capability would mean that the experiment design system would be a learning system, able to continuously improve it knowledge base. There are two main research questions inherent in the problem: how to recognize when a plan is worth saving, and how to decide how general to make it while still retaining its utility.

There are several possible measures of plan "worthiness." One would be whether the plan performed dramatically better than previous plans (e.g. it may have decreased the time to perform an experiment by an order of magnitude). Another would be related to how difficult it was for the system to create the plan. In other words, the plan should be saved because it would take a long time to find it again. The question is an experimental one; the research will involve trying many heuristics and balancing the improvement in system planning performance against the growth of an unwieldy and overly constrained knowledge base.

The question of how general to make the plan and how to parameterize it should also be solved experimentally. There will be trade-offs between how frequently the plan is used and what percentage of the time it will lead to a useful instantiated experiment design.

Another research goal is to use the knowledge base and experiment design system as a testbed for an automated performance evaluation system. The goals of such a system are quite general: to determine exactly how well the system is making use of the knowledge base, and how suitable the knowledge base is for the task at hand.

Among the specific questions a performance evaluation system for MOLGEN might answer are:

1. Is the system overlooking skeletal plans that it should find?

2. Is it needlessly considering many poor alternative plans?

3. Is it poorly modelling the consequences of plan steps?

4. In what areas of the knowledge base are decision heuristics weak or missing?

5. What types of knowledge are hardly ever being used?

All of these questions should be generalizable to many other knowledge-based problem-solving systems. Since the construction of large, expert knowledge bases is such a difficult task, the feedback from the evaluation of the use of these knowledge bases will be invaluable to future system builders.

Applications to Molecular Biology

The direct applications of MOLGEN to the field of molecular biology fall into three categories: knowledge base development and experiment design, analysis of nucleic acid sequences, and miscellaneous tools.

Knowledge Base Development and Experiment Design

The original and principal goal of the MOLGEN project is to provide a sophisticated experiment planning program containing an extensive knowledge base in the domain of molecular biology. As described above, our progress towards this goal has succeeded in the development of an extensive outline of this broad domain with emphasis on the myriads of analytical laboratory techniques that exist in this field. Using this knowledge base, MOLGEN is now capable of designing a number of sophisticated analytical experimental procedures. The procedures designed by the system are those already utilized in the laboratory, indicating that the knowledge base contains the correct sorts of heuristics to produce at least competent experiment designs. The limited scope of the current knowledge base provides a constraint on the originality of plans that can be produced; the most novel plans designed by humans are those which draw from many different, perhaps unrelated, knowledge sources.

Another success of the knowledge base concerns the organization of the information about each experimental technique. Because of the great flexibility of the Unit System, it is easy for the domain experts to modify and expand the existing information about each entity. We are continuously fine tuning the type of information contained within the knowledge base, in both content and in organization, during the actual knowledge acquisition phase.

We now propose to attack problems in synthetic molecular biology. We feel that by focusing our efforts on this subject we can assure an extensive repertoire of knowledge for that particular type of problem. This will also allow the planning algorithms to develop more sophisticated plans in the particular area. We have chosen to develop a knowledge base dedicated to the problem of cloning specific genes by recombinant DNA techniques. We have chosen this problem for four reasons: it is one of the most widely used methods in molecular biology today; most of our existing knowledge base is relevant to this problem; both of our current planning algorithms have been successful on either this problem (Stefik's thesis) or closely related problems of analysis of recombinant DNAs (Friedland's thesis); and because the method can be readily divided into four limited subdomains. These include choice of vectors, method of linking foreign DNA to the vector, transformation of host cells with the recombinant DNAs, and selection of the recombinant DNA containing the gene of interest.

We will describe current methods for cloning genes in both eukaryotes and prokaryotes, using methods in which one can select either for the vector or the inserted gene, and we will describe all the known methods of selecting for genes including direct functional selection, hybridization methods and expression of specific gene products. In addition to specifying the starting population or DNA sample and the ultimate goal, we will allow the user to specify certain subgoals or substrategies.

Analysis of Nucleic Acid Sequences

Our goal is to provide powerful, but easily used programs for the problem of the recognition of biologically significant patterns within nucleotide sequences. To make a set of programs both powerful and easy for a novice to use they must be interactive, self-documenting, and have easy to understand output formats. It also helps tremendously if they are very rapid so that they may be utilized online with nearly instantaneous feedback concerning the progress of the comparison. For this reason we have chosen to utilize the search algorithm developed by Korn and Queen and to convert it to an interactive form. This program was originally designed to provide for speed of comparison of very long nucleotide sequences while still allowing a degree of sophistication within the matching procedure. The algorithm compares two sequences beginning at every position where they share at least a dinucleotide but only carries the comparison as far as certain criteria of matching are allowed. This method, while lacking the sophistication of algorithms that potentially simulate evolutionary steps in the divergence of two sequences or the energetics of the pairing of single-stranded regions of dyad symmetry, is capable of detecting all statistically significant homologies or dyad symmetries given any level of significance desired. Unfortunately it is not capable of comparing more than two sequences at a time nor giving a quantitative measure of the divergence or relatedness of those two sequences. It merely describes the probability of each homology in terms of that expected for a random sequence of a given length and base composition.

Our improvements to the program have included converting it into SAIL and making it interactive. Whenever a user is in doubt about the next step he merely enters a ? and his options at that point are explained. We have also considerably improved the statistical calculations so that the probabilities and expectation frequencies that are determined for a homologous region are based not only on the length of the sequences being compared, but also on the base composition and on the exact algorithm being used in the search itself. Finally we have markedly improved the output displays so that that mismatches are indicated with stars and base pairs in dyad symmetries with bars. We have done all of this without any overhead in terms of execution time so that the program executes almost without delay in a time-sharing environment.

We propose to improve our current sequence analysis capabilities by implementing more sophisticated algorithms within the interactive framework. For instance the pattern recognition algorithm of Sellers is currently being implemented in C language at Rockefeller University by Dr. Bruce Erickson. We believe that this program would be a useful addition to our current armory in that it would allow us an accurate metric of relatedness of two sequences which is essential for building phylogenetic trees. This would be the first step towards the comparison of more than one sequence.

We would also like to develop methods for determining the secondary structure of single-stranded RNAs. The most commonly used methods are often limited to short nucleotide regions because of the complexity of the energy calculations for large numbers of comparisons. By first utilizing a rapid method for finding homologous sequences or dyad symmetries, perhaps guided by statistical significance of very low stringency, one might be able to rapidly eliminate most of the fruitless comparisons. By then examining the resulting culled homologies by a set of heuristics concerning their additivity, extension, or exclusiveness, we could order them in terms of their biological significance. This would automate some of the tedious cutting and patching of homologies and dyad symmetries in which molecular biologists are now involved even after they have made comparisons with a computer. With respect to calculations of the thermal stability of symmetric regions it would reduce the total time of calculation by orders of magnitude. In other words, we would use a comparison algorithm based more on biological intuition than calculation in order to find the most profitable regions to apply the more quantitative methods of biophysics.

We would further hope to automate the development of phylogenetic trees utilizing these sequence comparison algorithms. Once quantitative measures of relatedness are obtained in all pairwise combinations, then the matrix methods for the generation of the trees and the lengths of the branches is rather straightforward. These calculations are not likely to need any intelligent heuristics for their determination since they are defined analytically and they are rapid compared to the calculations involved in determining the relatedness of the sequences in the first place.

Miscellaneous Tools

Restriction Digest Analysis

One of the best examples of the utility of the application of heuristics and production rules to problems of molecular biology is the GA1 program, developed in this project, for the analysis of restriction endonuclease digests. Determining restriction maps of even simple DNA structures from restriction enzyme digest data can require consideration of millions of possible structures. The application of heuristic methods simplifies the analysis by orders of magnitude allowing solutions to complex problems and even simplifying the amount of data that must be collected to ensure a unique solution. These methods have even resulted in the proposal of a new experimental method for the analysis of restriction data.

GA1 is a program which determines all possible organizations of restriction fragments based on restriction endonuclease digests with single, double, and triple combinations of enzymes. The program contains an intelligent hypothesis generator and a set of production rules which allow it to generate and evaluate hypothetical restriction maps which are consistent with all of the data. These rules dramatically reduce the total number of possible structural candidates that must be both generated or evaluated.

Modern laboratory methods for determining restriction maps include end labeling procedures and two dimensional cross hybridization procedures. In order to extend the program GA1 to cover this kind of data we propose to be able to set up initial constraints on the locations of all restriction sites in certain local regions of the hypothetical restriction map. Such initial conditions (regional constraints) would be useful not only for entering data obtained from partial digestion of end labelled DNA segments, but would also be very useful if the complete nucleotide sequence were known for a particular region. Such conditions are often found in recombinant DNAs in which the nucleotide sequence of the vector is completely known.

Another improvement in GA1 which would both simplify and extend its use would be to allow the user to describe the complete restriction map determined previously for a limited number of restriction enzymes and then to enter digestion data for new enzymes, singly and in combination with the previously analyzed sites. These initial conditions would impose global constraints over the entire map. Global constraints will not be as readily implemented as the regional constraints described above.

If sufficient programming support is available we would also like to attempt to apply the hypothesis generating and production rule pruning approach to the analysis of two dimensional restriction data. In this method, radioactively labeled DNA segments generated from a DNA by a one restriction enzyme are hybridized to nonradioactive fragments generated by a second restriction enzyme thus indicating which pairs of fragments are homologous and hence overlapping. Currently the typical analysis is a data driven approach of finding a continuous path among all the overlapping DNA fragments cataloged by this experimental procedure. A model driven approach should extend this already powerful method. While the two dimensional cross-hybridization method only allows the generation of maps for two enzymes at a time, maps generated from all possible pairwise combinations of any set of enzymes are possible by analogy with the standard one dimensional method. Furthermore, by alternately labeling the fragments from either restriction enzyme and hybridizing those fragments to unlabeled fragments derived from the second enzyme in both directions, sufficient data should be obtained in order to overcome most mapping ambiguities which are usually the downfall of this method. Utilization of the model driven approach to the cross-hybridization procedure will also allow the generation of restriction maps of much longer DNAs than currently possible.

Synthesis of Specific Nucleic Acid Molecules

The MOLGEN knowledge base contains complete sequence information for all published and many unpublished nucleic acid molecules. It also knows about restriction endonucleases and their cutting sites and about ligation methods for rejoining nucleic acid fragments. We see potential use for this knowledge in designing synthetic pathways for the in vitro production of specific target molecules. This may actually be considered a part of the main experiment design effort, but the problem is important enough to make an independent specialized system desirable.

Currently, three major methods are used by molecular biologists to select specific sequences of interest from a recombinant DNA "library". The most widely used method uses isolated messenger RNA as radiolabeled probe to detect complementary DNA sequences in the recombinant molecules. This requires prior isolation of the mRNA which, unfortunately, is not always easily obtained. Secondly, and perhaps having the most long-term potential, are methods to select by expression in the host cell of the sought for functions. Such an approach will necessarily be limited to genes that can be made to supplement or rescue host functions. The problems of expression of eukaryotic genes in prokaryotic hosts may never be soluble because of the gene-splicing dichotomy. The utility of eukaryotic host-vector systems is now established but selection will still depend on prior creation of host mutants or use of immunological colony (or plaque) screening techniques still to be developed.

A third approach has been to use relatively short chemically synthesized oligonucleotide segments that are complementary to the gene of interest. The probe is used to select genomic clones of recombinants containing specific protein coding sequences. In theory, if the amino acid sequence is known, appropriate probes can be constructed. The techniques for chemical oligonucleotide synthesis are difficult and laborious. We propose a different approach using the recombinatorics of the computer stored and generated nucleotide sequences of all known DNA molecules. If the amino acid sequence of the protein whose gene is desired is known, then a computer assisted search through those sequences will attempt to locate oligonucleotides that could code for a short segment of that protein. By taking advantage of third base degeneracy and knowledge of restriction endonuclease cutting and splicing, constructions of natural oligonucleotides will be suggested. An intelligent algorithm might locate more than just one or two short segments capable of forming molecular hybrids with the DNA sequences being sought and these might be linked in a spaced out manner to provide a more powerful probe.

B. Justification and requirements for continued SUMEX use.

The MOLGEN project is dependent 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 are already large list of pilot exo-MOLGEN users on SUMEX.

SUMEX is currently 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.

9.1.5 MYCIN Project

MYCIN Project

Edward. H. Shortliffe, M.D., Ph.D.
Department of Medicine
Stanford University Medical School

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

I. Summary of Research Program

A. Project Rationale

The MYCIN Project is a set of subprojects, each devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The project retains the name of our first system, the MYCIN program, but has grown to involve five interrelated sub-projects (MYCIN, EMYCIN, CENTAUR, GUIDON, and ONCOCIN), each of which will be discussed in the sections that appear below.

Our first system, MYCIN, is an interactive consultation program which gives physicians antimicrobial therapy recommendations for patients with infectious diseases. The system must often decide whether and how to treat a patient before definitive laboratory results are available. It must recommend a therapeutic regimen which minimizes the risk of toxic side-effects while covering for all organisms which are likely to be causing the infection. The relevant knowledge is stored in production rules, and the system currently has rules for treating bacteremias (blood infections) and meningitis. There has already been early work on the codification of cystitis knowledge. The primary goal of the project has been to develop a program which can provide advice similar in quality to that given by a human infectious disease consultant. Formal evaluations of the program's recommendations for patients with bacteremia or meningitis have shown that this goal has been achieved. We have also sought to develop a system that is easy to use and acceptable to physicians. To accomplish this, numerous human engineering features have been incorporated into the consultation. There is also an extensive explanation facility which enables the system to explain its reasoning and to justify its recommendations.

The success of the MYCIN 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);

(2) to explore other uses of the 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 our newest system, ONCOCIN, have stressed simplified techniques for interacting with a knowledge base and entering data); and

(4) to experiment with using other knowledge representations in conjunction with the production rules used in MYCIN (our CENTAUR system is a modification to EMYCIN which uses prototypical descriptors of situations or disease states to guide and focus a consultative session).

B. Medical Relevance and Collaboration

The MYCIN program was designed to help alleviate the well-documented problem of antimicrobial misuse. We felt that MYCIN would be clinically useful when it was able to handle all major infections that are likely to be encountered in a hospital. Our success in developing a high performance program for meningitis and bacteremia has been documented in two articles by Dr. Yu listed in the publications section below. However, the system is not ready for clinical use because it does not have rules for the other areas of infectious disease. A very large investment in time and human resources is required to develop, test and formally evaluate a rule set for each major infection area.

By utilizing our EMYCIN system to collaborate on building the PUFF program, however, 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 outlined 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 in only eight months. Furthermore, we expect to have ONCOCIN used by Stanford oncologists before the end of 1980.

Finally, 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.

C. Highlights of Research Progress

MYCIN

Due to the departure of Dr. Victor Yu, the infectious disease expert who worked with us until recently, it has not been possible to expand the rule set into new areas of infectious disease. The 500 rules relating to bacteremia and meningitis are sufficiently rich and complex, however, that they serve as a particularly challenging vehicle for testing the new computational methods we are developing. MYCIN is now totally implemented as an EMYCIN system. Hence, our active work on EMYCIN has been thoroughly tested using MYCIN and our extensive library of patient cases. Ongoing efforts to expand MYCIN or prepare it for clinical implementation, however, have been temporarily set aside to allow us to concentrate on the projects below.

EMYCIN

Much of the work in the past year has been devoted to improving EMYCIN's facilities for allowing a system builder to construct and debug a knowledge base for a consultation system. This has included extensive documentation of the concepts used in EMYCIN consultation systems, the support programs for developing the knowledge base, and features of a working consultation system.

A knowledge-base debugging package was developed to assist the system builder in the task of testing, refining, and validating the knowledge base. This package includes: 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.

An additional development in the last year is the EMYCIN "rule compiler." Once a consultation program is built, it becomes important that it perform efficiently. This is most noticeable in large programs such as MYCIN. Production rules, while convenient in their modularity, are not the best representation for speedy execution. We have thus developed a rule compiler as part of EMYCIN that 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, an extensive EMYCIN user's document has been drafted. 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.

EMYCIN Applications

Several consultation systems have been written in EMYCIN. All but the most recent of these were developed in parallel with EMYCIN, and thus served to focus attention on certain features and shortcomings of the program to guide in its development. Their brief description here is intended to provide some indication of the range of potential applications of EMYCIN.

PUFF

The PUFF system performs interpretation of measurements from the pulmonary function laboratory. The project is a collaboration of a pulmonary physiologist, biomedical engineers, and Stanford computer scientists who had previous experience with the MYCIN program. The data from over 100 cases were used to create some 60 rules diagnosing the presence of pulmonary disease. These rules are used to create a complete report including the input measurements, other patient data, and the measurement interpretation. The system is a separate SUMEX project now, and is described in full elsewhere in this document.

HEADMED

The HEADMED program is an application of EMYCIN to clinical psychopharmacology. The system diagnoses a range of psychiatric disorders and can recommend drug treatment if indicated. Like PUFF, this project is a separate SUMEX project.

SACON

As a stronger test of domain independence, EMYCIN was applied to the completely non-medical domain of structural analysis. SACON (Structural Analysis CONsultation) provides advice to a structural engineer regarding the use of a large structural analysis program called Marc. The Marc program uses finite-element analysis techniques to simulate the mechanical behavior of objects. Engineers typically know what they want the Marc program to do, e.g., examine the behavior of a specific structure under expected loading conditions, but they do not know how the simulation program should be set up to do it. The goal of the SACON program is to recommend an analysis strategy; this advice can then be used to direct the Marc user in the choice of specific input data, numerical methods and material properties.

The performance of the SACON program matches that of a human consultant for the limited domain of structural analysis problems that was initially selected. To bring the SACON program to its present level of performance, about two man-months of the experts' time were required to analyze their task as consultants and formulate the knowledge base. About the same amount of time was required to implement and test the rules.

CLOT

A recent application of EMYCIN is CLOT, a system designed to diagnose disorders of the blood coagulation system of patients. It requests clinical evidence regarding an episode of bleeding, facts from the patient's general medical history, and the results of a battery of coagulation screening tests. From these data CLOT infers the presence and type of coagulation defect (if any) in the patient and then proceeds to make a refined diagnosis for any particular enzymatic deficiency or platelet defect. These diagnoses can be used by a physician to estimate the severity and cause of a particular episode of bleeding, evaluate the effects of various anti-coagulation therapies on a patient, or estimate the pre-operative risk of a patient having serious bleeding problems during surgery.

CLOT was constructed by David Goldman, a medical student at the University of Missouri, with the help of James Bennett, a member of our Stanford group who is very familiar with EMYCIN. Following approximately 10 hours of discussion about the contents of the knowledge base, they entered and debugged in another 10 hours a preliminary knowledge base of some 60 rules. CLOT is now an ongoing project at the University of Missouri.

GUIDON

Bill Clancey's thesis (August '79) marked the completion of version one of the program. Key results include:

(1) A language was developed for representing teaching expertise in the form of "Discourse Procedures"--sequences of rules that reflect dialogue patterns and are independent of the subject material to be taught. This representation was found to be suitable and convenient for incrementally developing a tutorial program.

(2) Various teaching methods were demonstrated for carrying on a case method dialogue with a student who is solving a complex diagnostic problem. Meta-knowledge about the representation of the subject material made it possible to express these capabilities in a domain independent way.

(3) The representation of subject material as modular production rules was studied and found wanting. Though rules conveniently separate relationships into readily accessible associations, an adequate knowledge base for teaching requires the addition of structural knowledge (clusters and patterns), support knowledge (underlying causal mechanisms), and strategical knowledge (managerial approaches).

Ongoing GUIDON research focuses on a number of issues:

The Student Model.

A revised student model has been designed to deal with the following questions:

(1) Can the student USE the program? i.e., is he able to enter recognizable input?

(2) Is the dialogue with the student COHERENT? i.e., are there recognizable patterns of student input and meaningful transitions between segments of behavior?

(3) Is the student PASSIVE OR ACTIVE? i.e., does he use his own knowledge to solve the problem, or does he rely on the tutor's initiative and ability to provide help?

(4) Does the student have a STRATEGY for solving the problem? i.e., is there some plan that organizes the student's data measurements and hypothesis selection?

Representation of Problem Solving Strategies.

One of the few formalized methods for teaching diagnostic strategies to medical students is a printed outline of data to collect. This outline is woefully inadequate as a teaching tool: it does not convey in itself the meaning or logic of the diagnostic process. Informal experiments with physicians have enabled us to formalize an ideal model of medical diagnostic strategy appropriate to our present domain of investigation (infectious meningitis). Work is underway to incorporate this model in MYCIN so that it "thinks like a clinician," and can thus be used to teach not only diagnostic rules, but human-usable methods for applying them.

Some surprising findings coming out of this investigation include the following:

(1) Establishing the hypothesis space is accomplished by considering causal links that might be enabled in this patient (called "risk factors"). This can be considered to be a process of determining the topology of the problem--causal connections that may have a bearing on the disorder.

(2) "Dropping back" is important to human problem solvers. In fact, hypothesis formation as we have observed it might be described as a process of maintaining a sense of the differential. Focusing and delving deeper is just a temporary phenomenon.

Acquisition of this strategical knowledge was greatly helped by analyzing protocols according to the structure/support/strategy framework we have established. This is one of the "knowledge engineering" results of our research.

CENTAUR

During the last year we have completed an implementation of PUFF using the augmented EMYCIN system known as CENTAUR. In this work, largely the effort of Jan Aikins, we have sought to strengthen the pure production rule representation of EMYCIN with additional focusing power provided by hypothesis "frames" or prototypes. CENTAUR now includes 24 prototypes and about 160 rules dealing with pulmonary disease. The system was tested on 100 cases from the files at Pacific Medical Center. CENTAUR agreed with two pulmonary physiologists 84 and 91 per cent of the time respectively on their diagnoses of pulmonary disease in the cases. (This was an improvement over PUFF, which had 74 and 85 per cent agreement with the two physiologists).

Basic AI research issues were also explored, such as the representation of control knowledge for computer consultations, and the explicit representation of the context in which knowledge is applied. Furthermore, the MYCIN explanation facility was expanded to include explanations of control processes, and to give explanations of the prototypes, as well as the rules.

Current CENTAUR research is concentrating on polishing and fine-tuning the PUFF implementation described above. Additional studies are contemplated to better define the precise reasons that CENTAUR has performed more accurately than PUFF on the 100 cases mentioned above. One expert collaborator, Dr. R. Fallat feels PUFF had performed less well because of the significant difficulties he has had in adding more rules and still keeping the knowledge base consistent. This was less difficult using the CENTAUR representation scheme.

Other research that will draw upon CENTAUR work includes the creation of additional applications systems using the CENTAUR prototype representation mechanism. One challenge will be to interface CENTAUR with the "context-tree" that is provided in EMYCIN, a problem that was not addressed in PUFF because it utilizes only a single context.

ONCOCIN

The oncology protocol management system, termed ONCOCIN after its domain of expertise and its historical debt to the MYCIN program, has achieved many of its early 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.

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

Our simultaneous clinical goal is 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.

During the first year of our research, it has been our aim to develop a prototype of the ONCOCIN consultation system, drawing from the programs and capabilities of EMYCIN. We have also analyzed carefully 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. Finally, we have spent much of our time considering 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.

We chose the series of protocols for Hodgkin's and non-Hodgkin's lymphoma as the first detailed knowledge to be encoded in the ONCOCIN system. These were selected because they were developed at Stanford, because they are among our most commonly used protocols in light of our position as a major lymphoma treatment center, and because the protocols are complicated, with many subtle details depending upon the stage of disease, concomitant or preceding radiotherapy, and evidence for drug toxicity.

Although the program will eventually be used on a high-speed terminal with a specially designed interface (see below), we decided that the initial prototype should be a self-contained consultation system that would be modeled on the form of interaction used for EMYCIN consultation systems. We chose not to use EMYCIN itself to build the system, however, because we quickly encountered several special needs that were better handled using alternate representation and control schemes. Therefore, although there are portions of the EMYCIN code that we have been able to borrow, ONCOCIN is an entirely new program in which production rules are only one of several types of knowledge representation used.

Both our own experience, plus evidence in the medical computing literature, have suggested that physicians will be unlikely to use consultation systems if they fail to fit smoothly in the day's normal routine. With this in mind, we have carefully studied the current organization and flow of information within Stanford's oncology clinic. A detailed document has been prepared which describes the current clinic organization and the ways in which our system will interact with the current routine. Two principal concerns have been:

(1) that ONCOCIN should initially have minimal impact on the current daily routine: record-keeping systems should not be altered, patient flow within the clinic should be unchanged, and the physicians working there should not be forced to depend on an operational computer system in order to get their work done;

(2) that it should not take any EXTRA effort on the physicians' part for them to use the ONCOCIN system (other than the initial time required while they are trained how to use it); this implies that the use of ONCOCIN should replace some task that the physicians are currently doing.

Currently the clinic physicians are asked to fill out, by hand, the time-oriented flowsheets that are kept in the patient clinic records. These sheets are the basis for data analysis of all the clinical research that is based on chemotherapy protocols in the oncology clinic. All information needed by ONCOCIN is entered on this flowsheet. Thus we intend to capture the data needed for an ONCOCIN consultation by having the physician fill out the flowsheet at a computer terminal rather than by hand.

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.

First we have decided to use high-speed CRT terminals (approximately 9600 baud) with auxiliary hard-copy devices. This will permit almost instantaneous screen filling and allow greater flexibility in the design of what is actually displayed. However, a program written in a powerful but slow language like INTERLISP is not able to service a high-speed terminal adequately. For this reason, our interface program will be written in a faster compiled language (we are using PASCAL), and this program will need to communicate in turn with the INTERLISP reasoning program that comprises the rest of ONCOCIN. The design of this interprogram interaction is largely complete, but actual implementation of the ideas is just beginning.

Second, we want to minimize typing by the physician. EMYCIN systems have required a typewriter-compatible keyboard, but we do not feel this is reasonable if ONCOCIN is to be used on a daily basis by a large number of oncologists. Initially we examined light-pen and touch-screen technologies, but feel that these are either too expensive or too unreliable. Ultimately, working closely with experts in human factors, we developed a customized 21-character keypad which has been interfaced with a Datamedia terminal similar to those we have used for other development work. This keypad can be used by the physician to fill out the patient's flowsheet (which will be displayed on the screen at high speed), and there should be minimal if any need to use the terminal keyboard itself.

Finally, we want to maintain the explanation and justification capabilities which we have argued are crucial to the acceptance of clinical consultation systems. A specialized split-screen display has been designed which will enable the physician to enter patient data entries in one region while pertinent explanations are displayed in another.

D. Publications Since January 1979

Kunz, J.C., Fallat, R.J., Mcclung, D.H., Votteri, B.A., Aikins, J.S., Nii, H.P., Fagan, L.M, Feigenbaum, E.A. Physiological rule-based system for interpreting pulmonary function test results. Memo HPP-78-154, Stanford Heuristic Programming Project, 1978. Also Proceedings of Computers in Critical Care and Pulmonary Medicine, IEEE Press, 1979.

Yu, V.L., Buchanan, B.G., Shortliffe, E.H., Wraith, S.M., Davis, R., Scott, A.C., Cohen, S.N. Evaluating the performance of a computer-based consultant. Comput. Prog. Biomed. 9,95-102 (1979).

Clancey, W.J. Tutoring rules for guiding a case method dialogue. Int. J. of Man-Machine Studies 11,25-49 (1979).

Clancey, W.J. Dialogue management for rule-based tutorials. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, pp. 155-161, August 1979.

Aikins, J.S. Prototypes and production rules: an approach to knowledge representation for hypothesis formation. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, Tokyo, Japan, August 1979.

Fagan, L.M., Kunz, J.C., Feigenbaum, E.A., Osborn, J. J. Representation of dynamic clinical knowledge: measurement interpretation in the intensive care unit. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, Tokyo, Japan, August 1979.

van Melle, W. A domain-independent production-rule system for consultation programs. *Proceedings of the 6th IJCAI*, August 1979.

Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. Knowledge engineering for medical decision making: a review of computer-based clinical decision aids. *Proceedings of the IEEE*, 67:1207-1224 (1979).

Yu, V.L., Fagan, L.M., Wraith, S.M., Clancey, W.J., Scott, A.C., Hannigan, J.F., Blum, R.L., Buchanan, B.G., Cohen, S.N. Antimicrobial selection by a computer -- a blinded evaluation by infectious disease experts. *J. Amer. Med. Assoc.* 242:1279-1282 (1979).

Shortliffe, E.H. Medical consultation systems: designing for doctors. To appear in *Communication With Computers* (M. Sime and M. Fitter, eds.), London: Academic Press, 1980.

Shortliffe, E.H. The computer as clinical consultant (editorial). *Arch. Int. Med.* 140:313-314 (1980).

Fagan, L.M., Shortliffe, E.H., and Buchanan, B.G. Computer-based medical decision making: from MYCIN to VM, *Automedica*, March 1980 (in press).

Shortliffe, E.H. Clinical knowledge engineering: the MYCIN Project. *Proceedings of the First Japanese Conference on Artificial Intelligence in Medicine*, pp. 1-8, Tokyo, Japan, August 1979.

Clancey, W.J. *Transfer of Rule-Based Expertise through a Tutorial Dialogue*. Computer Science Doctoral Dissertation, Stanford University, August 1979.

Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. Knowledge engineering for infectious disease therapy selection. *Proceedings of the Intl. Conf. on Cybernetics and Society*, Denver, Colorado, October 1979.

Clancey, W.J., Shortliffe, E.H., and Buchanan, B.G. Intelligent computer-aided instruction for medical diagnosis. *Proceedings of the Third Annual Symposium on Computer Applications in Medical Care*, Silver Spring, Maryland, October 1979.

Fagan, L.M., Kunz, J.C., and Feigenbaum, E.A. Representation of dynamic clinical knowledge: measurement interpretation in the intensive care unit. *Proceedings of the Third Annual Symposium on Computer Applications in Medical Care*, Silver Spring, Maryland, October 1979.

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

Shortliffe, Edward H. Consultation systems for physicians: the role of artificial intelligence techniques (invited paper). *Proceedings of the 3rd Annual Meeting of the Canadian Society for the Computer Simulation of Intelligence*, Victoria, British Columbia, May 1980.

E. Funding Support

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: 1 P01 LM 03395
Term: July 1979 to June 1984
Total award: $497,420
Current award (1979-1980): $99,484

Grant Title: "Knowledge-Based Consultation Systems"
Principal Investigator: Bruce G. Buchanan
Agency: National Science Foundation
ID Number: MCS-7903753
Term: July 1979 to June 1980 (plus 6 months)
Total award: $146,152
Current award (1979-1980): $73,659

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: "Symbolic Computation Methods For Clinical Reasoning" (RCDA)
Principal Investigator: Edward H. Shortliffe
Agency: National Library of Medicine
ID Number: NIH 1K04 LM00048
Term: July 1979 to June 1984
Total award: Dollar amount negotiated annually
Current award (1979-1980): $39,285

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

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 have been shown by the medical and academic communities. For two years in succession we have been 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. There has also been significant growing interest in medical AI and MYCIN from colleagues in Japan. We were asked to demonstrate MYCIN from Tokyo during the 6th International Joint Conference on Artificial Intelligence held in August 1979. Access to SUMEX via a trans-Pacific TYMNET link worked very well and permitted large numbers of Japanese and other conference attendees to observe MYCIN demonstrations and experiment with the program themselves. Then, for three weeks in November 1979, Dr. Shortliffe returned to Japan as a visitor at the Tokyo Metropolitan Institute of Medical Sciences. This visit permitted an intensive period of exchange regarding MYCIN, EMYCIN, and the related work being done by the Japanese.

Several teachers have also asked to use MYCIN in their computer science or medical computing courses. For example, Prof. Carl Page of Michigan State University, Dr. Peter Szolovits of MIT, and Dr. Steven Zucker of McGill University in Montreal have demonstrated the MYCIN program in their university classes. Dr. Harold Goldberger of MIT made extensive use of the MYCIN program in his study of medical AI programs. Dr. Ves Morinov of the Norwegian Computing Center has used the MYCIN program to demonstrate the benefits of using a rule-based representation for consultation systems. Dr. Martin Epstein used MYCIN as one of the representative systems he demonstrated to students who took the clinical elective on medical computing at the NIH during the summer of 1979.

GUEST users who have recently requested access to MYCIN have come from such diverse locations around the country as the Brain Research Institute (UCLA), University of Texas, Stevens Institute of Technology, University of New Mexico, Columbia University, Systems Science Institute (Louisville), Naval Postgraduate Institute (Monterey, Ca.), Texas Women's University, IBM Scientific Labs, and Alta Bates Hospital (Oakland, Ca.).

EMYCIN has also generated a great deal of interest in the academic and business communities. We have been in frequent contact with Bud Frawley and Philippe Lacour-Gayet of Schlumberger, Chuck Brodnax and Milt Waxman of the Hughes Aircraft Corporation, and Harry Reinstein from IBM Scientific Research Center. Two students at the Naval Postgraduate School in Monterey, working under the direction of Colonel Ronald J. Roland, have been developing an EMYCIN system in the domain of selecting decision aids for solving problems in business organizations. The CLOT system mentioned earlier was a joint effort involving members of our group but with the idea and domain expertise coming from members of Don Lindberg's group at the University of Missouri. At the University of Illinois, students working under Donald Michie and Alan Levy have used EMYCIN in two ways: one group developed a new EMYCIN application in tax advising, and the other developed a PASCAL implementation of the ideas used in EMYCIN. The latter program is now being used experimentally in an application involving emergency responses on off-shore drilling rigs. Finally, David Stodolsky at the Systems Science Institute at the University of Louisville has begun to experiment with EMYCIN in an application involving the psychology of interactions in large group conferencing.

B. Sharing and Interaction with Other SUMEX-AIM Projects

We have continued collaboration with the EMYCIN-based projects RX, HEADMED and PUFF. 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.

We have also interacted with members of the SECS project on SUMEX who have considered developing a question answering system for SECS similar to the one in MYCIN.

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, to be held at Stanford in August 1980, have been 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 are aware that we have taxed 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 if 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 MYCIN and INTERNIST. In March 1979, when the American College of Physicians met in San Francisco, they rented a truck and drove to the City with terminals and monitors. The installation they arranged worked well and provided a superb demonstration environment for the physicians who attended. In New Orleans in 1980, the greater distance prevented us from installing the equipment ourselves. SUMEX kindly offered to help orchestrate the New Orleans arrangements, though, and literally hours were spent locating terminals, arranging for telephone hookups, and finding the right kind of slave monitors. We salute SUMEX for their uncomplaining assistance in this regard, but also would like to note the need for a mechanism that is somewhat less ad hoc for facilitating the demonstration of SUMEX systems from remote locations.

Finally, we continue to feel the need for more computing power. Most of our research and development takes place in the hours from 7 p.m. to 10 a.m., but it is unreasonable to expect all our collaborators 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 will also allow us to introduce ONCOCIN in a real clinical environment within the next several months, 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.

III. Research Plans (8/80-7/86)

A. Project Goals and Plans

EMYCIN

Our current plans call for four principal efforts related to EMYCIN. First, the knowledge acquisition component of the program, derived from the TEIRESIAS work of Davis, is being modified and expanded. Our concerns relate to both the inefficiencies and limited power of the current capabilities. The meetings during which the CLOT knowledge base was developed were recorded on tape and are forming the basis of an analysis of the knowledge acquisition process. Some early work implementing the ideas derived from those tapes is already under way.

We are also planning to prepare EMYCIN for "export" during the coming year. This will involve tightening up the code, maximizing efficiencies in space and time use, and improving the system's documentation. We do not intend to recode EMYCIN in a language other than INTERLISP, but do want to make it a stand-alone system that can be used for system building in a number of LISP environments. A key element of the documentation will be to better define those environments in which EMYCIN can be most effectively applied.

Now that the design and capabilities of EMYCIN are essentially fixed, we are also planning to develop a new application. Other EMYCIN systems have been developed in parallel with EMYCIN itself, and have therefore affected the program's design, but it is now appropriate to see how effectively a new system can be built within the current system constraints. We are just beginning work, in conjunction with IBM Scientific Labs, to develop an EMYCIN consultation package for electronic fault diagnosis.

GUIDON

A plan for further development of GUIDON is described in terms of a partial ordering of research problems. Improving the student model will receive priority.

```
                interruption/assistance/evaluation
                     teaching strategies
                     /                  \
                    /                    \
                   /                      \
        dialogue planning                  \
                 |                          \
                 |                           \
                 |                            \
        case selection                         \
                 |    \                         \
                 |     \---------------- student model
                 |
         case differences/
       genetic epistemology
```

Implementation of the strategical methods is now proceeding. There are several tasks (corresponding to the managerial and operational considerations) organized hierarchically. These tasks will be expressed in rule form (if <proc> then <task>).

Structural knowledge will serve to hook these domain independent strategical rules into a particular rule set like MYCIN's. This will involve adding a taxonomic problem classification to the knowledge base and regrouping rules and parameters according to this classification.

Besides using the strategical model for guiding a dialogue with a student, we are investigating the possibility of reconfiguring MYCIN's rule set so that the strategy rules direct a consultation. The result will be a knowledge base of rules and parameters, just like MYCIN's, that does hypothesis formation with focusing by the same backward chaining interpreter we have always used. Even without this step, by formalizing (on paper) a strategical model in terms of production rules, we are led to conclude that it is the exhaustive, depth-first character of MYCIN's search that is different from hypothesis formation, not backward chaining. The strategical rules are meta-rules that modify MYCIN's search. Subgoaling by backward chaining of rules is compatible with both depth-first search and hypothesis formation.

Missing knowledge aside, we find that many of MYCIN's rules are too detailed to be learned by people. We find that people just don't think about the fine-line, statistically-based distinctions that MYCIN rules record. We have developed a way to encode what an expert actually knows by overlaying qualifications on top of MYCIN's rules. This takes the form of a functional statement (e.g., csf-protein is proportional to intensity and duration of illness) and ranges of discrimination (<100 means viral; >250 means chronic or bacterial; otherwise "it could be anything"). These summary statements capture what the student should learn; they will be used in quizzes based on the rules, as well as for selecting cases.

In a related development, we are trying to record aphorisms and mnemonics that experts use for remembering strategical and mechanistic principles, e.g., "when you hear hoof beats think of horses, not zebras" and "csf glucose is low for bacterial meningitis because bacteria eat the glucose for food" (this is wrong, but physicians remember it and generally don't realize or care that it is wrong!). We find that causal knowledge in our domain serves as a cue for remembering associations; actual diagnosis generally occurs at a level higher than causal mechanism.

ONCOCIN

In the three months remaining in the current year, we expect to have completed the PASCAL interface program that will respond to the special keypad on the Datamedia terminal. We also intend to codify the rules for one more chemotherapy protocol (probably oat cell carcinoma of the lung) in order to verify the generality and flexibility of the representation scheme we have devised. In the coming year, our plans include the following:

(1) To develop the software protocols for achieving communication between the PASCAL interface program and the INTERLISP reasoning program.

(2) To coordinate the printing routines needed to produce hardcopy flowsheets, patient summaries, and encounter sheets.

(3) To install the new terminal and hard copy device in the Oncology Day Care Center for final testing and debugging.

(4) To begin offering the ONCOCIN system for use by oncology faculty and fellows in the chemotherapy clinics (three mornings per week) in which most of the lymphoma patients receive their treatment.

(5) To codify and implement additional protocols contingent upon adequate progress with the steps outline above.

Throughout this work 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 other 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 20/20 which will enable us to use a 9600 baud interface. Most of us use SUMEX as the only computer on which we work.

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, Great Britain).

C. Requirements for Additional Computing Resources

The recent acquisition of the 20/20 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), additional small computers would 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 would 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 during the daytime hours, and is beginning to be limiting on occasion in the evenings as well. Any acquisitions that would provide additional cycles or permit off-loading of some users from the PDP-10 would significantly benefit the SUMEX research community.

The continued growth of our research project, with MYCIN space still required, GUIDON growing, and ONCOCIN now a new and large system, has resulted in some moderate problems with disk allocation as well. We have managed to shuffle allocations reasonably effectively until now, but there is no longer much flexibility and an additional allocation of approximately 2500 pages would greatly relieve the pressure.

D. Recommendations for Future Community and Resource Development

We have two principal recommendations for new SUMEX developments. First, the acquisition of several small machines, linked to the main processor through the ethernet, and each able to run INTERLISP, would allow important experiments in bringing the more mature AIM systems closer to being exportable for use outside of strict research environments.

Second, we propose the formal establishment of a mechanism for providing hardware and communications equipment for SUMEX demonstrations at a distance. There are beginning to be enough invitations for the older AIM systems to be shown at meetings and to funding agencies, that a dedicated system of demonstration equipment and personnel seems appropriate at this time.

9.1.6 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.

C. Progress summary

We have completed a major cycle of design review and program reorganization, resulting in the system described in publication number three below. The system now has a completely 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 completed the Initialization task, implemented a tracing task, and implemented a task to split group toeholds. Further details of these tasks and their content can be found in publication number three.

We have also continued our efforts to improve the power of our data representations. Towards this end we have implemented a new preprocessor to assign functional labels to segments. This program consists of heuristics that attempt to capture the knowledge a human uses when he visually examines a skeletonized EDM. We find the use of labeled segments greatly aids the main CRYSALIS program by allowing rules to be written in terms much closer to those which humans use rather than the language in which the EDM skeleton is defined.

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. Whereas the previously-implemented tracing task attempts to grow an "island of certainty" in the hypothesis in a non-directed manner, we are now working on a task that specifically tries to link two such islands. In addition to this new task, we are augmenting the system's tracing knowledge to deal with small sidechains that seldom appear in the data. The final addition to the knowledge base is an effort to incorporate 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) 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)

2) 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)

3) Robert Engelmore and Allan Terry, "Structure and Function of the CRYSALIS System", Proc. 6IJCAI, 1979. pp250-256 (Alternative identification: HPP-79-16)

4) 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

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 and UCSF). 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. Use of SUMEX during the remaining grant period (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 about a dozen 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.

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.

9.1.7 RX Project

The RX Project: Deriving Medical Knowledge from Time-Oriented Clinical Databases

Robert L. Blum, M.D.
Division of Clinical Pharmacology
Department of Internal Medicine
Stanford School of Medicine

Gio C. M. Wiederhold, Ph.D.
Departments of Computer Science and Electrical Engineering
Stanford University

I. Summary of Research Program

I.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 1967. 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 10,000 patients with a variety of rheumatologic diagnoses. Over 30,000 patient visits have been recorded, accounting for 20,000 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 discover 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.

I.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 a year, and, hence the project is very much in a developmental stage. The primary issues being addressed at this stage are those concerned with the specifics of knowledge representation and flow of control, rather than with the testing of specific hypotheses in chronic disease management.

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.

Highlights of Research Progress

1 July 1979 to 1 April 1980

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 250 INTERLISP functions accounting for 75 disk pages of code. The KB is about 30 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. In order to do this we use the SPSS package (Statistical Package for the Social Sciences) available on SUMEX. A collection of RX programs create SPSS "source decks" containing card images of the appropriate commands along with the extracted data. RX then calls the operating system and runs SPSS on the source file. The human-readable listing is then searched for important results which are automatically extracted and interpreted.

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.

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 and the development of expanded capabilities for interfacing RX to on-line statistical packages. 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 large matrices of simple and partial correlation coefficients.

Publications

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

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

Wiederhold, Gio: Databases in Healthcare. To be published in a 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

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

II.A. Collaborations

Since our project is 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

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.

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.

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.

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.

SUMEX as a Resource

To discuss SUMEX as a resource for program development, one need only compare it to the environment provided by our other resource, the IBM 370/168 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.

Programs:

RX is written in INTER-LISP. To increase our useable address space, we actually use a stripped-down version prepared by William VanMelle of the MYCIN Project. To run statistical data RX calls SPSS in an inferior fork. The text editor, TVEDIT, is also called from an inferior exec fork.

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 and to transfer statistical source matrices back to SCIP to be run on that machine.

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.

9.2 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.

9.2.1 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:

This last year has seen developments in two main directions. We are completing developing and documenting a system (ACTF) that is capable of a relatively rich variety of cognitive learning and we are completing an application to the modelling of the acquisition of proof skills in high school students.

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.

A manual has been under construction describing these changes. We plan to have a final version of the ACTF system by the end of May and the manual should be finished by the end of the summer.

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; 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, an 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. To appear 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. Becoming expert at a cognitive skill. To appear 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. Learning to Plan in Geometry. To appear in J.R. Anderson (Ed.), Cognitive Skills and their Acquisition. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1981.

E. Funding Support:

A Model for Procedural Learning,
John R. Anderson, Principal Investigator,
Office of Naval Research (N00014-77-C-0242)
$175,000 September 1, 1978 - September 30, 1980

II. Interaction With the SUMEX-AIM Resource

A. & B. 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.

C. 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.

Our more immediate goals (for the next year or two) involve application of the ACTF system, whose development we have finished, to three domains. First, we hope to complete the development of a simulation of geometry learning in the system. Second, we are starting to embark on an effort to model the acquisition of programming skills in LISP. This will serve as another test of the ideas that we have developed in geometry about learning and planning. The third application will be the modelling first language acquisition. This is a more radical departure from our work in problem-solving and so will provide a rather different test of the learning theory.

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 increased computational power, to increase the address space of INTERLISP (e.g. VAXes), or to create significantly more space efficient versions of INTERLISP.

9.2.2 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, CA. 95064

Coworkers:

D. Dolata	(Grad student)
R. Lasater	(Grad Student)
D. Rogers	(Grad Student)
J. Chou	(Postdoctoral)
P. Condran	(Postdoctoral)
T. Moock	(Postdoctoral)
T. Blume	(Programmer)

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 focused on basic research into representation of strategies, facilities for user-defined transforms, revision of our ALCHEM language for better debugging of transforms and extension of capabilities for representing complex reactions. In addition we hoped to improve capabilities for remote teletype usage of SECS and to initiate the formation of a world-wide SECS Users Group for sharing chemical transforms.

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 6) 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. 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 in the process of moving to a newly renovated room with raised floor in the same building and same floor as the synthetic organic laboratories at Santa Cruz so the environment is excellent.

I. C. Highlights of Research Progress

1. 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 CONGEN layout program of Carhart has been modified and incorporated in SECS for clean teletype output and simplified teletype input for users without graphics terminals.

The synthesis tree plotting program for hard copy has been rewritten to give more compact trees which are faster to plot on the plotter. This generates better plots in less time and can also be used with XENO.

The ALCHEM language which we developed for representing chemical reactions has undergone extensive revision to make it easier to represent absolute stereochemistry and some of the complex reactions in heterocyclic chemistry. Part of this revision now enables SECS to explain to the chemist which ALCHEM statements are being used and the results of their interpretation via a new decompiler for ALCHEM. A complete manual on ALCHEM and a manuscript on the revisions has been written.

A User Defined Transform (UDT) module has been added to bridge the gap between program knowledge and user knowledge. This allows the chemist, during a synthetic analysis, to graphically specify a reaction which SECS doesn't know, and continue without interrupting the analysis. The SECS database is also still expanding as a result of contributions from our group and from the SECS Users Group.

A META-SECS top-level plan generator has been outlined 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 an inference processor is currently in design stages. The explicit representation of synthetic strategies will be an interesting exploration which we feel other synthetic chemists will benefit from, even through manual use of these strategies. Hand simulation of this program is in progress.

2. 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 program is operational, but both the program and the data base need improvement for field use.

The teletype input and output has been improved by incorporating a modified version of Carhart's teletype plot module from CONGEN so the program can be accessed remotely via teletype or graphics terminal.

The second phase of XENO which evaluates potential biological activity is currently being developed. Currently XENO can check each metabolite generated by exact match against a library of compounds and thus if a match is found, pull out the biological activities. Our plans however are to allow extrapolations beyond known compounds and for that we are pursuing several approaches using chemical pattern recognition and chemical similarity.

Collaborations with experimental metabolism researchers have begun in order that XENO can make predictions for compounds actively being studied in the laboratory. We hope to get feedback regarding the usefulness of this methodology and to accumulate a list of verified predictions for publication. These collaborators include scientists from NIH, FDA, EPA, ICI Pharmaceutical, Upjohn Co., and UCSF Medical School. 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.

J.D. Andose, E.J.J. Grabowski, P. Gund, J.B. Rhodes, G.M. Smith, and W.T. Wipke, "Computer-Assisted Synthetic Analysis: The Merck Experience," in Computer-Assisted Drug Design, ed Olson and Christoffersen, ACS Symposium Series 112, pp 527-552, 1979.

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, in press.

Manuscripts describing our work on symmetry, similarity, and ALCHEM are currently in the review process.

E. Funding Status

1. Resource-Related Research: Biomolecular Synthesis
   PI: W. Todd Wipke, Associate Professor, UCSC
   Agency: NIH, Research Resources
   No: RR01059-03S1
   7/1/80-2/28/81 $ 36,949 TDC

2. Computer-Aided Prediction of Metabolites for Carcinogenicity Studies
   PI: W. Todd Wipke
   Agency: NIH, National Cancer Institute
   No: N01-CP-75816
   1/1/80-12/31/80 $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. A complete synthesis tree was generated for Prof. William Dauben, UC Berkeley of isocomene which was analyzed in detail by his students. They were impressed by the magnitude of the number of synthetic approaches and that all known syntheses were found by the computer. Similarly an analysis of several insect pheremones was done and sent to Prof. A.C. Oehlschlager, Dept of Chemistry, Simon Fraser University, British Columbia, Canada. Other visitors for whom we have done analyses include Dr. M. Onozuka, A. Tomonaga and H. Itoh, Kureha Chemical Co, Tokyo Japan, Dr. Rhyner, Director of research, Ciba-Geigy, Basel. A synthesis of vellerolactone, a substance found to be toxic and teratogenic was generated for Prof. R.E. Carter, Univ. Lund Sweden. A conformational study of substituted hydroazulenes was performed for Clayton Heathcock, Berkeley (Synthesis of Isoprenoid Antitumor Lactones, NIH CA 12617). The XENO project is working on metabolism of diallylmelamine N-oxide, a hypotensive compound in collaboration with Dr. John M. McCall of Cardiovascular Diseases Research, The Upjohn Co.

Dr. Wipke has also used several SUMEX programs such as CONGEN in his course on Computers and Information Processing in Chemistry. 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.

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 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. Lastly our disk space (8000 pages) is too small for our current research projects and staff.

D. Collaborations and Medical Use of Programs via Computers other than SUMEX. Arrangements are currently being made to place SECS 2.7 on several computer networks so anyone can access it without having to convert code for their machine. This has proved very useful in the past as a method of getting people to try this new technology. SECS 2.0 has resided on the First Data network since 1974 and has been used extensively in the US and abroad.

III. Research Plans (8/80-7/86)

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 adding a 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 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.

Currently the similarity module requires a special version of SECS. We plan in the next year to incorporate this module into the standard version of SECS so that the bonds that if broken could lead to identical or similar fragments can be used to create a goal to guide SECS toward such efficient syntheses, even though there may not be a reaction capable of doing that rejoining step.

We will incorporate the Aldrich catalog of available chemicals, both to recognize when a precursor is available and to explore strategies based on available starting materials. The process must be efficient for the library contains 20,000 compounds.

We have now a PDP-10, a Univac, and an IBM version of SECS. We hope to compare these and create one version which will run on these and other machines to facilitate sharing of new modules among collaborators.

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 an expansion of our disk allocation from 8000 pages to 10000 pages for the growth of our programs, databases, and personnel. We are currently tightly constrained spacewise and are hampered in research because of inability to keep needed files. We also would like to have the overlay loader fixed so that an overlaid program can retain its symbol 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 heavily utilized, 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. We further feel that it would be worthwhile if discussions regarding the future expansion of SUMEX and the community could include the remote users who depend on SUMEX. SUMEX can not currently handle additional people from 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.

A trivial suggestion but also important is that TV-EDIT be improved to not leave null characters in files which cause problems with compilers both at SUMEX and at other sites when the files are sent to another machine. This suggestion has been made many times by many people but the situation still exists.

9.2.3 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

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).

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; these people include Alan Lesgold of SUMEX's SCP project. 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.

Progress Summary

The prose comprehension project has completed an early version of a comprehension model that has now been used by several different researchers (Miller & Kintsch, 1980a). This model has been applied to twenty different texts, and has yielded quite reasonable predictions of recall and readability. We are currently expanding on the premises of this model toward a system 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. We are about to start development of a program to partially automate this modelling process.

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.

Polson, P. G., Jeffries, R., Turner, A., & Atwood, M. E. The process of designing software. To appear in J. R. Anderson (Ed.), Learning and Cognition. 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. Paper presented at the American Educational Research Association meetings, April, 1980.

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
    9/1/79 - 8/31/80: $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
    6/1/79 - 5/31/80: $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 - 12/31/80: $85,000

II. Interactions with the SUMEX-AIM Resource

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. It is likely that the initial comprehension model (Miller & Kintsch, 1980a) will be used by Dr. Lesgold and other researchers at the University of Pittsburgh, as well as researchers at Carnegie-Mellon University and the University of Manitoba.

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 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)

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.

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.

Needs and Plans for Other Computational Resources

We currently use three other computing systems, two of which are local to the University of Colorado. One is the Department of Psychology's CLIPR system, which is a Xerox Sigma 3 used primarily for the real-time running of 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. Thirdly, the planning group has been using a PRIME computer located at Science Applications, Incorporated for the storage and analysis of protocols.

CLIPR is about to replace the Sigma 3 with a VAX 11/780. 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.

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.

Given our imminent acquisition of a VAX, we would 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.

9.2.4 HMF - Higher Mental Functions

Higher Mental Functions Project

Kenneth Mark Colby, M.D.
Professor of Psychiatry and Computer Science
Neuropsychiatric Institute
University of California at Los Angeles

I. Summary of Research Program

A. Project rationale

The rationale of this project is to contribute new knowledge and instruments to the fields of psychiatry, neurology, and communication disorders using the concepts and methods of artificial intelligence. The project is involved in studies of paranoid conditions, psychiatric taxonomy, intelligent speech prostheses, ideographics for language generation, and computer enhancement of patient outcomes in large mental hospitals.

B. Medical relevance and collaboration.

As can be seen from the above description, the project has clear medical relevance. The project collaborates with psychiatrists, neurologists, speech pathologists and biomedical engineers. Besides working at the UCLA Neuropsychiatric Institute, the project collaborates with the Northridge Hospital Foundation, Northridge, California.

C. Highlights of research progress.

In collaboration with three psychiatrists and four psychologists we are working out a new taxonomy for the "neuroses", a category which is notoriously unreliable in the psychiatric classification scheme. In this pilot study we are collecting data on 50 patients and 70 controls. One segment of data is provided by the subjects' self-accounts which are analyzed by a large program run on the SUMEX facility. This program finds the key ideas in the subject's account and assigns him a profile. The profiles will be clustered into groups and the groups compared to those formed on the basis of the other data-collections in the study. During the past year, the project has developed intelligent speech prostheses (ISPs) which (a) utilize a lexical-semantic word-finding algorithm for anomic aphasias and (b) utilize ocular control for the generation of synthesized speech. These devices serve as aids to nonvocal patients handicapped by strokes, tumors, cerebral palsy, and tracheostomies.

The word-finding algorithm is dynamically re-organized by the user's selection of words. It is currently being tested on a 54-year-old man with an almost complete anomia due to a stroke in the left hemisphere. The algorithm needs a larger memory to accommodate at least 5,000 English words. The large dictionary on the SUMEX facility is of great help in constructing the lexical-semantic memory.

We have just begun to test the use of ocular control of an ISP. The patient wears specially designed spectacles which can detect where the eye is directed on a small TV screen. Thus the patient spells out words by looking at letters on the screen. Signals from the spectacles are sent to the ISP which generates the utterance of the words thus spelled.

Although we have ceased to work on the paranoid PARRY program, due to lack of funding, it is available for demonstration and study by those interested in modelling psychiatric syndromes.

We are in the planning stages of developing a computer ideographic writing system for language generation by nonspeaking patients who cannot spell. If they can learn ideographic symbols which stand for certain concepts and construct the symbols on a graphics terminal by pressing keys, a translating program will convert the symbols into English words which in turn will be spoken by an ISP. We are also beginning to design a type of computerized "recreational-educative" therapy for patients in large mental hospitals with such a shortage of professional manpower that the patients' treatment is limited mainly to custodial care.

D. List of Relevant Publications.

Colby, K. M., Christinaz, D., Graham, S. 1978. A computer-driven personal, portable, and intelligent speech prosthesis. Computers and Biomedical Research, 11: 337-343.

Colby, K. M. 1979. Computer simulation and artificial intelligence in psychiatry. In Methods of Biobehavioral Research E. A. Serafetinides, (ed.), New York: Grune and Stratton.

Colby, K. M. 1980. Computer psychotherapists. In Technology in Mental Health Care Delivery Systems, J. B. Sidowski, J. H. Johnson, T. A. Williams (Eds.). Norwood, New Jersey: Ablex Publishing Corporation.

Heiser, J. F., Colby, K. M., Faught, W. S., Parkison, R. C. 1980. Can psychiatrists distinguish a computer simulation of paranoia from the real thing? The limitations of Turing-like tests as measures of the adequacy of simulations. Journal of Psychiatric Research, Vol. 15, No. 3.

Parkison, R. C. 1980. An effective computational approach to the comprehension of purposeful English dialogue. Stanford University, Ph.D. dissertation, (forthcoming).

Colby, K. M., Christinaz, D., Graham, S., Parkison, R. C. A word-finding algorithm using a dynamic lexical-semantic memory for patients with anomia. (In press)

E. Funding Support.

1. Titles of grants
   a) Intelligent Speech Prosthesis
   b) Ocular control of Intelligent Speech Prosthesis.

2. Principal Investigator

   Kenneth Mark Colby, M.D.
   Professor of Psychiatry and Computer Science
   Neuropsychiatric Institute
   University of California at Los Angeles

3. Funding agencies

   a) Intelligent Systems Program, Division of Mathematics and Computer Science, National Science Foundation.

   b) Science and Technology to Aid the Handicapped Program, National Science Foundation.

4. Grant numbers

   a) NSF-MCS 78-09900

   b) NSF PFR - 17358

5. Total award period

a)	6/1/78 - 11/30/80	$135,260.
b)	10/1/79 - 3/31/81	$318,368.

6. Current period

   (see 5. above)

II. Interactions with the SUMEX-AIM Resource

A. The project communicates and collaborates with the Communication Enhancement Project at Michigan State University, John Eulenberg, Principal Investigator.

B. The project communicates with the SUMEX project at the University of Texas at Galveston, John F. Heiser, M.D., Principal Investigator, who experiments with and demonstrates the PARRY program.

C. Critique of resource management. The SUMEX staff is still excellent and responsive to our needs. Our only problems are with the telephone company portion of our communications link with SUMEX.

III. Research Plans (8/80 - 7/86)

A. Project goals and plans

1. Near-term

We plan to continue to work on the problems described above. Further clinical experience is necessary in testing and developing the word-finding algorithm and the ocularly-controlled ISP. These efforts should be completed in about two years.

2. Long-range

It will take years to solve the problems of psychiatric taxonomy, computer ideographic writing systems, and computer enhancement of hospitalized patient outcome. Our work in these areas will depend upon obtaining the requisite funding.

B. Justification for continued SUMEX use.

All the problems we work on involve natural language in some form or other. We analyze natural language input and generate natural language output. These efforts require large dictionaries and large LISP programs which run at SUMEX. No comparable facilities are available at UCLA. Hence we are heavily dependent upon SUMEX for the continuation of this research.

C. Needs and plans for other computer resources.

An ISP consists of a microprocessor interfaced with a speech synthesizre. We have constructed 3 ISPs, building two of the microprocessors ourselves. We expect to purchase another microprocessor and a graphics terminal.

D. Recommendations for future development.

The SUMEX system is often heavily loaded during daytime hours. The batch facility permits us to run some large production jobs overnight unattended, but the daytime loading is often so great that it discourages even small interactive jobs, such as text editing. It would be very helpful to have more computing power during the daytime, if funding is available.

9.2.5 INTERNIST Project

INTERNIST 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, along with an extensive medical data base now encompassing almost 500 diseases and more than 3,000 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 system which we now call INTERNIST-I was begun about eight 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, 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 INTERNIST 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 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 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 INTERNIST-II, 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 INTERNIST-II.

C. Highlights of research progress

Accomplishments this past year

During the past year, the R & D activities of the INTERNIST project have concentrated on three major problem areas associated with the original implementation of INTERNIST. These areas are:

a) restructuring of the underlying diagnostic logic of INTERNIST to conform more closely to the expectations of clinician users of the system. The primary goal in developing a new model of diagnostic reasoning is to achieve a concurrent problem formation capability in order that improved scoring methods and attention to the principle of parsimony might be exploited in focusing the attention of INTERNIST on regions of the problem space having the greatest potential for yielding a solution. Moreover, the new approach has the potential for improved modes of interaction with the user, as it can reveal at any point in its analysis the multiple partial characterizations that have been postulated, and expose the space of alternative complex descriptions that can be generated by combining these partial characterizations. The potential for providing justification and explanation of the system's behavior is thereby greatly enhanced.

b) development of a friendlier user interface, enabling use of the system by clinicians unfamiliar with the specifics of the INTERNIST vocabulary. One of the barriers to successful implementation of the original INTERNIST system in a ward setting is the language of discourse used in that system for specifying the positive and negative findings in a clinical case. The number of possible findings that might be entered now numbers more than three thousand; thus some means for convenient browsing among these possible entries, and some convenient means for communicating the selected items to INTERNIST had to be found. We have developed for this purpose a menu-selection front end system, that comprises a network of approximately 1000 frames designed to permit selection of pertinent facts that might be revealed by any of a host of information acquisition procedures. Convenient escape mechanisms have been provided to permit the user to alternate between the interactive data entry and analytical components of the system.

c) incorporation of additional disease profiles and related medical information in the INTERNIST knowledge base, to approach the critical mass required for effective field tests of the system.

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.

2) The completion and implementation of the improved diagnostic consulting program, which has been designed to overcome certain performance problems identified during the past four years' experience with the original INTERNIST program.

3) Institution of field trials of INTERNIST 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 INTERNIST to subserve educational purposes and the evaluation of clinical performance and competence.

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, Proceeding IEEE Intercon, New York, 1975.

E. Funding support

1. Title of grant.
   Clinical Decision Systems Research Resource.
2. Harry E. Pople, Jr., Ph.D.
   Associate Professor of Business

   Jack D. Myers, M.D.
   University Professor (Medicine)
   University of Pittsburgh
3. Division of Research Resources
   National Institutes of Health
4. 5 R24 RR01101-03

5. 07/01/77-06/30/78
   $160,414

   07/01/78-06/30/79
   $178,414

6. 07/01/79-06/30/80
   $200,414

II. Interactions with the SUMEX-AIM Resource

A, B. Collaborations and Medical Use of Program Via SUMEX

INTERNIST 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 INTERNIST. 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 (8/80-7/86)

A. Project Goals and Plans

We expect that the conversion of INTERNIST knowledge structures to the form required by INTERNIST-II will be reasonably complete by the next fiscal year (June 30, 1981). Shortly thereafter, provided adequate hardware resources are available, we intend to commence formal field trials of INTERNIST at the Presbyterian-University Hospital of Pittsburgh. This local phase of the clinical evaluation will continue for approximately one year.

Beginning in July 1982, we intend to extend the clinical trials to collaborating institutions, with the addition of one additional user group approximately every six months through June 1984.

B. Justification and Requirements for Continued SUMEX Use

In order to provide the level of computer services required by the expanded level of R & D activity in the near term, and to support the schedule of field trial studies envisioned during the current five year planning horizon, we have requested NIH support for a dedicated INTERNIST machine to be acquired during the next fiscal year.

If this hardware support becomes available, we would not expect to make additional demands on SUMEX-AIM for computing services. However, we would continue to look to SUMEX for software support and for the communications network that so effectively bridges the far-flung AIM community.

Until such dedicated resources are in place, we would expect to make use of the SUMEX-AIM facilities at a moderately increased level of utilization.

9.2.6 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 A. Feigenbaum, Ph.D.
Computer Science Department
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.

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 proposal for supplemental funding, submitted 1 February 1979, is pending.

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

Knowledge base:

PUFF is implemented on the PDP-10 in a version of the MYCIN 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).

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.

The PMC and SUMEX computers will be linked by telephone. The physiological measurements are generated every 2-10 minutes by the PMC computer system. It will be provided to VM in real time using the phone link. Information, suggestions to the clinicians, and/or requests for additional information will be sent back to the ICU for action.

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 has been demonstrated using actual patient data recorded on magnetic tape. The input to VM is the values of 30 physiological measurements provided on a 2- or 10-minute bases by a automatic monitoring system. The output is in the form of suggestions to clinicians and periodic summaries (see example case below).

Example Case

The following case demonstrates the current state of development 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.

An extended description of the VM program can be found in a Ph.D. thesis to be available shortly as a Stanford technical memo.

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 (forthcoming).

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.)

II. Research Plans

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).

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. This link is required for real-time testing of VM. SUMEX staff had the necessary resources to design and implement this vital link mechanism. The link is now undergoing final testing, and it will dramatically contribute to the effectiveness of the research environment for VM.

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. All 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 process of validation then will require running VM in real time so that PUFF/VM researchers can compare system interpretations of patient state with the actual state as determined by careful concurrent clinical evaluation. We believe that we can effectively use 3-4 hours per day of running VM in a real time test mode during the initial validation period. As the system operation becomes more predictable in 1981, longer running times will be required to identify system problems, and we predict the need to run the system for a full eight hour shift each day on an intermittent basis.

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 compliment the technical service facilities currently provided by the SUMEX staff.

9.2.7 Simulation of Cognitive Processes

Simulation of Cognitive Processes

James G. Greeno
Alan M. Lesgold
Learning Research and Development Center
University of Pittsburgh

SUMMARY OF RESEARCH PROGRAM

Project Rationale

Our goal continues to be contribution to increased theoretical understanding of basic cognitive processes involved in reading, problem solving, and other tasks requiring cognitive skills. The form in which we theorize is computer simulation of human performance. Models of cognitive processes stand as hypotheses about the components of human information processing and the ways in which they interact in significant cognitive tasks.

Medical Relevance

Increased understanding of basic cognitive processes is relevant to medical needs in two ways. One form of relevance involves performance of tasks in the practice of medicine. One of us (Lesgold) collaborates in research on cognitive processes in radiology. Understanding of the nature and organization of these processes, and those in other domains of medical practice, should provide principles useful in the design of medical training and the arrangement of conditions for more efficient delivery of medical services. 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 sever disablement. Improved understanding of these basic skills should provide principles useful in improvement of diagnosis and therapy for learning disabilities.

Highlights of Research Progress

--Accomplishments this past year

Progress was made in the study of basic processes in reading skill, where preliminary findings suggest that children whose speed of vocalizing words develops slowly are destined to be slow in acquiring poor reading skill (Lesgold, 1979). Progress in Anderson's ACT system and in our own empirical work will enable more computer simulation work on reading in the coming year. Comprehension of quantitative concepts was studied, with development of a hypothesis relating the outcome of problem understanding and choice of an arithmetic operation through an abstract representation of a quantitative action (Heller, 1980). Developmental changes in quantitative understanding were identified in a study of children from 5 to 8 years of age (Riley & Robinson, 1980). Children's understanding of computational procedures was studied, and instruction based on procedural analogies was found to be helpful in remedying systematic procedural flaws in children's performance (Resnick, 1979). A simulation model and a formal analysis of preschool children's counting skills were developed, providing some progress on the question of what constitutes understanding of a general principle relevant to a cognitive procedure (Greeno, Gelman & Riley, 1978). A theory of problem-solving set and constructions was developed in the domain of high school geometry proof problems, based on an idea of schematic knowledge (Greeno, Magone & Chaiklin, 1979).

--Research in progress

Research is continuing on all these topics.

We will briefly describe four research projects that depend on SUMEX and are most directly relevant to AIM goals.

(1) Work has begun on a study of the acquisition of radiological skills. The general strategy is to start with empirical data and proceed to simulations of novice (first-year residents) and expert cognitive processing during film reading. The final stage will be the development of learning mechanisms that transform novice models into expert models. We use protocols of beginning residents, fourth-year residents, and senior radiologists gathered in relatively naturalistic film-reading situations along with eye movement data and studies of what subjects see in the first seconds of examining a film. Current work on the computer simulation part of the project is directed at development of an anatomy database to underpin representations of films and to provide a language for describing feature analyzers and higher level knowledge structures and their outputs. In general, it is expected that the novice model will be similar in form to that of HEARSAY-II, while the expert model will have a somewhat more "compiled" form, perhaps looking more like some of the diagnostic programs on SUMEX. At this point, only the novice model has been considered in any detail.

(2) Another study of spatial information processing is focused on alternative cognitive representations of information in diagrams. Venn diagrams are presented along with verbal keys that indicate probabilities of events. In solving simple computational problems, subjects identify figures in the diagrams that they use in organizing the numerical information needed for calculations. Subjects differ in the level of complexity of forms that they identify, indicating that individual differences in spatial information processing affect performance in this task in a fundamental way. Simulation models are being constructed using Anderson's ACT program in the SUMEX system. These models represent alternative forms of spatial information available to a problem solver, and permit investigation of the consequences of alternative forms of spatial information for the inferential processes required of a problem-solving system.

(3) A collaborative project with John Anderson is focussed on learning of problem-solving skills. Anderson and Greeno are developing simulation models, using ACT, representing different stages in the

acquisition of cognitive procedures for solving geometry proof problems. Greeno's contribution to the project involves simulation of learning new procedural skills that make use of previously known schemata that are used in representing problems. Problems being addressed include (a) acquisition of productions for instantiating a schema in a new context, thus making available the problem-solving procedures previously learned in different contexts; (b) acquisition of new procedural attachments required for solving new kinds of problems in a familiar domain; and (c) acquisition of complex schemata formed by combining components of simpler schemata that were known previously.

(4) We are conducting a formal analysis of acquisition of the syntax of simple arithmetic sentences. The problem is a form of the language acquisition problem, and we are developing a system patterned after Anderson's (1976) Language Acquisition System (LAS), which depends on semantic representations of the referents of sentences in acquiring syntactic parsing rules. Our project involves an extension of this idea, since the referents of arithmetic sentences are sequences of actions, rather than spatial arrangements of objects as Anderson used. The programming in this project is done in SAIL through the SUMEX system.

(5) Work continues on the longitudinal study of children's development of reading skill. This work is expected to facilitate modelling of different forms of reading acquisition problems by providing examples of different children's progress in acquiring various components of effective word recognition.

List of Relevant Publications

Greeno, J. G., Gelman, R., & Riley, M. S. Young children's counting and understanding of principles. Paper presented at meetings of the Psychonomic Society, San Antonio, November, 1978.

Greeno, J. G., Magone, M. E., & Chaiklin, S. Theory of constructions and set in problem solving. *Memory and Cognition*, 1979, *7*, 445-461.

Heller, J. I. The role of "focus" in children's understanding of arithmetic word problems. Paper presented at meetings of the American Educational Research Association, Boston, April, 1980.

Lesgold, A., & Curtis, M.E. Learning to read words efficiently. In A.M. Lesgold & C.A. Perfetti (Eds.), *Interactive processes in reading*, Hillsdale, NJ: Erlbaum, forthcoming.

Lesgold, A.M., Curtis, M.E., Roth, S.F., Resnick, L.B., & Beck, I.L. A longitudinal study of reading. Paper presented at the Annual Meeting of The American Educational Research Association, Boston, April, 1980.

Riley, M. S. & Robinson, M. A theoretical framework for word problem research in arithmetic. Paper presented at meetings of the American Educational Research Association, Boston, 1980.

Funding Support

National Institute of Education

1. Title: Research on Learning and Schooling
2. Principal Investigators: Robert Glaser, University Professor and Co-Director of Learning Research and Development Center, and Lauren B. Resnick, Professor of Psychology and Co-Director of Learning Research and Development Center, University of Pittsburgh
3. Funding Agency: National Institute of Education
4. Grant Number: NIE-G-80-0114
5. Total Award: 1 Dec 1979 to 30 November 1982, $7,879,729.
6. Current Period: 1 Dec 1979 to 30 Nov 1980, $2,625,520 (During the current period, $150,000 of the above has been allocated for Greeno's Research and $67,000 for Lesgold's).

Office of Naval Research and Advanced Research Projects Agency

1. Title: Cognitive and Instructional Factors in the Acquisition and Maintenance of Skill
2. Principal Investigators: Robert Glaser, University Professor and Co-Director of Learning Research and Development Center, and Alan M. Lesgold, Research Assistant Professor of Psychology, University of Pittsburgh
3. Funding Agency: Office of Naval Research (through funds currently provided by the Advanced Research Projects Agency)
4. Contract Number: N00014-79-C-0215
5. Total Award: 1 Jan 1979 to 30 Sep 1981, $1,265,272.
6. Current Period: 1 Oct 1979 to 30 Sep 1980, $420,000.

National Science Foundation and National Institute of Education

1. Title: Invention and Understanding in the Acquisition of Computation
2. Principal Investigator: Lauren B. Resnick, Professor of Psychology and Co-Director of Learning Research and Development Center, University of Pittsburgh.
3. Funding Agencies: National Science Foundation and National Institute of Education
4. Total and current Funding: 1 Dec 1978 to 31 May 1981, $161,238.

Office of Naval Research

1. Title: Analysis of Formal and Informal Reasoning in Problem Solving
2. Principal Investigator: James G. Greeno, University Professor, University of Pittsburgh
3. Funding Agency: Office of Naval Research
4. Contract Number: N00014-78-C-0022
5. Total Award: 1 Oct 1977 to 30 sep 1980, $274,419.
6. Current Period: 1 Oct 1979 to 30 Sep 1980, $92,293.

INTERACTIONS WITH THE SUMEX-AIM RESOURCE

Medical Collaborations and Program Dissemination 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.

Sharing and Interactions with Other SUMEX-AIM Projects

Two of the five projects described in Section 1.3 involve use of Anderson's ACT system in SUMEX. The skill acquisition project involves direct collaboration and programming using the ACT system in Anderson's directory. The project on spatial information processing with diagrams is also programmed in ACT. Access to Anderson's programs through SUMEX has allowed us to avoid costly duplication of his system, which would require translation from INTERLISP into another dialect as well as unnecessary duplication of disk files on another system. The reading work also involves access to ACT, currently for development work and later for actual building of models of reading.

Critique of Resource Management

None

RESEARCH PLANS

Project Goals and Plans

In the near term, we will complete our analysis of learning arithmetic syntax, the analysis of acquiring geometry problem-solving skill, and the analysis of spatial information processing with diagrams. Work on reading will continue for several years. We expect to use new versions of ACT that permit partial matching of production conditions to simulate one or more different types of low-reading-achievement children.

The radiology diagnosis modelling work is expected to continue through an initial phase of novice modelling in interaction with empirical work on chest film reading, after which we will proceed to the expert model and to specification of learning mechanisms. Those mechanisms are expected to include some of the mechanisms proposed by Anderson in his current work as well as mechanisms that take particular account of the need to not have all of the film viewing process excessively concentrated on the highest-probability hypotheses. That is, we see a need to understand how good radiologists come to be able to check films for unexpected pathology (such as tumors) even when seeing evidence for entirely different disorders. We hope that this work will lead to a better sense of how to teach radiologists to exercise this additional care.

Another long-term project is development of a theory of learning elementary arithmetic. Arithmetic is a relatively well structured domain. We now have a considerable body of empirical and theoretical knowledge about the cognitive structures and processes that constitute knowledge of elementary arithmetic. The development of a system that can acquire this knowledge appears to be a feasible goal. At the same time, the task of building a system that acquires both procedural skill and conceptual understanding and integrates these aspects of knowledge raises theoretical questions that seem nontrivial. Therefore, development of a learning system for elementary arithmetic appears to be a productive project for our research program during the next few years. As noted above, we will also use ACT for the reading modeling work.

Justification and Requirements for Continued SUMEX Use

We anticipate continued use of SUMEX in development of simulation programs, particularly in shared use of Anderson's ACT system. Anderson is presently developing the learning capabilities of ACT in a systematic way, and this is very likely to be an important resource for our long-range project on the learning of arithmetic.

Needs and Plans for Other Computing Resources

We depend on SUMEX for a relatively modest, albeit significant, share of our computing needs. We have installed a VAX-11/780 at LRDC and hope to benefit from SUMEX-AIM's experience to continue to improve the cognitive science resources we have locally. We are also exploring possibilities for involvement in any cognitive science network that may develop. In any event, though, having direct access to resources such as ACT _as they are developing_ plays a major role in allowing our work to proceed at the current pace. Complete detachment from ACT would produce a major setback and would waste a lot of staff time in re-inventing the work others already have in place.

9.2.8 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.

The Resource community includes 85 researchers and professionals - 37 members, 11 associates, 28 collaborators and 9 users. Members are mainly located at Rutgers. Collaborators are located in several distant sites and they interact, via the SUMEX-AIM and RUTGERS/LCSR facilities, with Resource members on a variety of projects, ranging from system design/improvement to clinical data gathering and testing of expert systems. Our collaborations are described further in section B below. Resource users are located at Harvard University, John Hopkins University, Ohio State University, University of Pennsylvania, University of Pittsburgh, Stanford University and the NIH campus.

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)

(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 fifth AIM Workshop, organized by the MIT - Tufts Clinical Cognition Project was held in Vermont in May, 1979. The Sixth Workshop is being organized by the SUMEX-AIM Resource and is to be held at Stanford University in August, 1980.

B) Medical Relevance; Collaborations:

During 1979-80 we continued the development of a versatile system for building consultation programs, called EXPERT. This system is being used extensively in the development and study of several medical consultation models - in collaboration with clinical investigators from several specialties.

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 a consultation models should follow 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.

In a single year we progressed from an initial model that represents a framework of major findings and diagnostic categories for diffuse connective tissue diseases to a refine model with a broad spectrum of well-defined observational and decision criteria. It is now being validated and further developed through a national network of rheumatology specialists organized by our medical collaborators at the University of Missouri. This work is directly contributing to the organization of clinical knowledge in rheumatology. It has been a notable achievement to have been able to reach a performance of over 90 percent of correct diagnoses on difficult cases of disease at each step of model design.

In ophthalmology, the CASNET/Glaucoma knowledge base was translated into the new EXPERT formalism. The development of the glaucoma knowledge base built in conjunction with the investigators of ONET (ophthalmological network) was supplemented by knowledge of Japanese variants of the disease and the decision rules embodying the clinical judgment of Japanese glaucoma experts. A model for neuro-ophthalmological consultation is being built in collaboration with Dr. William Hart of the Washington University School of Medicine, which is related to the automated interpretation of visual field measurement.

Another collaboration has been in the area of endocrinology, where a thyroid consultation knowledge base was developed in conjunction with Dr. R. A. Nordyke of the Pacific Health Research Institute.

All the above applications have shown the versatility of the basic EXPERT representation scheme for rapidly developing medical knowledge bases. By continued testing and development of various domain models, the current boundaries of applicability of the EXPERT formalism are being explored, and new facilities added as required to improve the consultative performance of the programs developed.

In addition to the direct medical collaboration, we have continued investigating problems of modeling in enzyme kinetics with Dr. David Garfinkel of the University of Pennsylvania.

C) Highlights of Research Progress

1) Medical Modeling and Decision-Making

Research activities during the past year have concentrated on the development and testing of the generalized consultative system scheme (knowledge representation and associated strategies of inference), called EXPERT, and its application to a number of different medical domains.

The structure of knowledge in EXPERT involves two data types: findings and hypotheses. The hypotheses (diagnostic, prognostic and treatment selection) are organized as a partially ordered network (PON) using hierarchical and causal relationships. The findings are organized according to observational constraints. Production rules are used to encode inferences among findings, between findings and hypotheses, and among hypotheses. Because of the PON organization of hypotheses, the knowledge base can be pre-compiled with attendant space and time efficiencies in the performance of the consultation programs that call on the knowledge base for decision-making advice.

Knowledge bases in ophthalmology, rheumatology, and endocrinology have served to test the versatility of the EXPERT formalism. (see I.B above)

There have been a number of significant generalizations of the EXPERT scheme during the current period, which fall in the following categories:

1.1) Representations:

a) The context for Hypotheses-to-Hypothesis rules can be defined (anchored) to include matching against a pattern of other hypotheses as well as findings. This permits both very data-specific contexts, as well as global contexts of disease domains and consultation environments, which are used as triggers for the sets of applicable production rules.

b) Multiple visits of a patient can now be handled by the same scheme. A time representation is being developed currently.

c) Internal functional and logical variables can now be defined for use by the reasoning schemes, hence allowing specification of clinical indices, discriminant functions, transformations among variables, etc.

d) Extension of the logical selector operation in the syntax to apply to hypotheses as well as findings.

e) The current version of EXPERT has been expanded to handle large amounts of knowledge, up to approximately 600 hypotheses, 3,000 findings and 20,000 rules, while retaining its processing efficiency.

1.2) Strategies of Reasoning:

a) A focusing capability, which permits the system to concentrate on only a preselected set of conclusions at a given time. Repeated application of the focusing command gives the user direct control over the "shifting of attention" in the reasoning sequence of the system, which may be an attractive alternative to the program's usual control strategies. It is also a powerful tool to test the effects of hypothesis-induced partitions in the sets of production rules.

b) Strategy selection capabilities have been added, which permit the user to pre-specify the type of scoring strategies used by the system in assessing the effects of propagated uncertainties throughout the space of hypotheses, while interpreting of a given patient's consultation results. This has proven to be a useful tool in adjusting the scoring method to match the degree of structural specification in a model. In particular, for applications where few interdependencies among findings and hypotheses are known and included in a model, it is desirable to strengthen the cumulative components of a scoring function.

1.3) Facilities for Model Updating and Explanations:

a) Compiler of benchmark case changes: The compiling program, XP, has been designed so that on each revision of the model for which it is invoked, it will provide a summary of all significant changes in the conclusions of a set of stored benchmark cases. These can then be retrieved and analyzed for unexpected effects of changes in the rules or the descriptive knowledge structure.

b) The explanation facilities have been generalized, so that supportive evidence or chains of reasoning leading up to a particular hypothesis or conclusion can be traced and assessed. The user can specify the range of uncertainty weights for which he wishes to obtain explanations of the conclusions (i.e., only those hypotheses that are strongly confirmed, only those strongly denied, or any other alternative). Rules can also be ordered according to weight criteria.

1.4) Knowledge-base Transfer Experiments with EXPERT

Experiments were carried out to test the facility with which knowledge bases constructed for other representations can be transferred into the EXPERT formalism. They primarily involved the CASNET model of glaucoma, and demonstrated that the causal structure was, as designed, representable in the EXPERT scheme. It allowed for the explicit specification of hierarchical relationships and rules for the inference of intermediate hypotheses from evidence and final conclusions from patterns of intermediate hypotheses. There were some specialized features of CASNET/Glaucoma that have yet to be added to the new system (use of symmetry relations for binocular findings and hypotheses and the visit-to-visit logic), that are currently under design.

Another experiment, carried out by Dr. Kitazawa in Japan, took part of the CASNET knowledge base and transferred it into the EXPERT representation, and then added new elements (findings, hypotheses, and rules) to adapt the model to his clinical environment. Examination of the INTERNIST-I representation showed the compatibilities between some of its components and those expressible in the EXPERT scheme.

Problems of updating a knowledge base and learning decision rules from a data base of case records are two other areas of investigation. A program for rule learning by five different fuzzy-logic heuristic methods was developed and tested using allergy case study data. Problems of the transferability of large-scale consultation programs to a minicomputer environment have also been investigated.

Clinical investigations in thyroid disease and hypertension (by investigators at Pacific Health Research Institute and the Johns Hopkins School of Medicine, respectively) have been aided by Resource support and development of the BRIGHT system.

2) Modeling of Belief Systems and Commonsense Reasoning

The central role of commonsense reasoning in human thinking makes it a particularly important form of reasoning to study and describe. However, the theoretical frameworks, research methodologies, and analytical tools that have been developed within psychology are not adequate for this task. Consequently, over the past 10 to 15 years, psychologists interested in investigating human reasoning in such "knowledge-rich" domains, have increasingly looked to the research in artificial intelligence in the hope that the tools and research strategies of this discipline can be borrowed, customized, or extended to aid the psychologist in the investigation of human reasoning.

At the broadest level, our research in commonsense reasoning represents one of a handful of research projects that are exploring this intersection of AI and cognitive psychology. We have borrowed, customized and extended the conceptual tools of AI as a result of trying to state, justify and test a knowledge-based theory of one aspect of commonsense reasoning, namely, the problem of how persons recognize the plans and intentions that guide the actions of another person.

Progress toward the achievement of the general goal of applying and/or developing AI approaches for use in developing a cognitive science in psychology can be attained only by providing, at some level of approximation, solutions to four problems. First, a general system framework must be developed within which a knowledge-based psychological theory of some aspect of human cognition can be expressed. The AIMDS system has been the continually evolving framework which we have developed and used to express our theories about commonsense reasoning.

The second and most visible problem is that of representing the knowledge and processes that constitute the psychological theory of plan recognition, person perception and belief attribution. We invented the term BELIEVER in order to distinguish, at least in our mind, those aspects of the implemented code and architecture that represent information processing structures and mechanism that constitute a psychological theory of aspects of human cognitive performance. There have, of course, been many versions of BELIEVER and there is a sense in which the code that constitutes BELIEVER represents several theories. The process of plan recognition requires that mechanisms of retrieval, matching, hypothesis revision, plan generation, inference, categorization, concept specialization or customization and question answering be specified. In non-AI based research in psychology, these various aspects are typically treated as independent areas in which theories are formed and tested. In an AI-based theory such as BELIEVER, there is a "unified" theory of these phenomena in the sense that the way in which these processes interface and communicate with each other has been worked out to yield a functioning system.

The third problem that must be faced is that of developing experimental paradigms which can yield a set of observations that are rich enough to constrain and test an information processing theory of this type. Data from typical psychological experiments yield only a few observations on a single subject and rarely attempt to speak to the way in which various processing mechanisms interface with each other to produce the observed behavior.

Finally, not only must promising experimental paradigms be identified and response protocols collected, but procedures must be devised for representing such data and evaluating it against theoretically interesting hypotheses. This is a very difficult task since the assumptions that underlie the standard statistical tools used in psychology for this purpose are usually at variance with the underlying assumptions of information processing theories.

Over the lifetime of this project we have made forays into each of these problem areas and have learned a good deal about the terrain of each of these problems areas--that of devising the system framework, AIMDS, and that of theory construction, BELIEVER. In the remaining two problem areas our outposts are still at the fringes of the terrain to be searched although we do feel that we at least now have an understanding of the topographies of these areas. This emphasis reflects both our assumptions about the precedence ordering that naturally falls over these problem areas as well as the nature of the collaborative effort that exists between AI and cognitive psychology within the Resource.

3) Artificial Intelligence; Representations and Systems Development

A major part of our effort in this core area continued to be directed to collaborations with investigators in the other applications - oriented projects of the Resource. These collaborations are having an impact on the application areas of the Resource, and they are stimulating work on basic AI issues that are related to designs of knowledge-based systems.

The following problems are providing foci of collaboration with investigators in the Medical Systems area: (i) Develop a natural language interface between a computer consulting system and a medical user; (ii) Find methods for representing and effectively using several related bodies of medical knowledge at various levels of resolution (anatomical, physiological, causal-associational) for decision making in diagnosis and therapy; (iii) Develop computer tools and design frameworks for facilitating the construction and improvement of expert systems.

The joint intensive work between investigators in this core area and researchers in the Belief Systems area is continuing. Developments in the AIMDS system and in the BELIEVER theory are proceeding in parallel and they are continuing to influence each other.

During this period we put new emphasis on basic problems of expertise acquisition and on related problems of theory formation; we increased our effort in problems of representation, interpretation and model-guided control of natural processes; we continued basic work on problems of knowledge acquisition in the context of a language learning task; and we continued a modest level of effort in programming language development - to provide a supportive programming environment for our research.

Our research on natural language processing has continued with the objective to develop methods that facilitate communication between people (domain experts, users, designers) and computers. We have taken a fresh look at the problem of developing a convenient man machine interface for a glaucoma consultation system. Building on our previous work in this area, we have added several novel features to the design of our interface processor.

We are continuing to study problems of language acquisition/learning to gain insight into the general problem of knowledge acquisition in expert AI systems. However, we have shifted emphasis this year to an approach which assumes a more active teacher-learner dialogue in the language acquisition process. This led to the identification of rules that govern such a dialogue, and to the design of acquisition processes that embody these rules.

Our commitment to a strong AI programming environment resulted in improvements of the Rutgers/UCI LISP system, as well as in other systems programming developments. These efforts are strengthening the tools for design and experimentation that are available to Resource investigators on the RUTGERS/LCSR computer facility.

D) Up-to-Date List of Publications

The following is a list of books, papers and abstracts published in 1978 and 1979 by the Rutgers Resource:

Amarel, S., (1978) "Introduction and Overview for AI in Science and Medicine," in 'Session on AI in Science and Medicine, National Computer Conference 1978', AFIPS Conference Proceedings, Vol. 47, 1978, AFIPS Press.

Amarel. S. (1979) Invited participation in panel on "History of AI: 1956-1961," with E. Feigenbaum (Chm.), J. McCarthy, H.A. Simon, IJCAI-79, Tokyo, Japan, August 1979; also, S. Amarel chaired two sessions on Representations at the IJCAI meeting

Biesel, P., C. Kulikowski, S. Weiss, and Z. Aviur, (1979) "Computer-Aided Acquisition and 3-Dimensional Display of Visual Field Data, Computers in Ophthalmology Conf. St. Louis, April (1979).

Ciesielski, V., (1979) "Natural Language Input to a Glaucoma Consultation System," Presented at the Annual Meeting of the Association for Computational Linguistics, San Diego, August 11-12, 1979.

Ciesielski, V. [ed.] (1979) "Proceedings of the Fourth Annual AIM Workshop, "Computer Science Report CBM-TR-104, Rutgers University, August, 1979.

Ciesielski, V., D. Smith, and P. Biesel, (1979) "Artificial Intelligence in Medicine: Contributions of the Rutgers Resource to the A.I. Handbook," Computer Science Report,

Goldberg, R. (1979) "BRIGHT User's Guide - Version 3.04", Report CBM-TR-95, Department of Computer Science, Rutgers University, January 1979.

Hall, J.S. and N.S. Sridharan, (1978) "Modeling Actions and Processes in AIMDS: An Example", Report CBM-TM-81, Department of Computer Science, and Medicine, Vol. 11, (1 and 2), August 1978.

Kulikowski, C., (1979) "Expert Consultation Systems: Designs for Generality", Hawaii Int. Conf. on Systems Science,

Kulikowski, C. and S. Weiss, (1979) "Representation of Expert Knowledge for Consultation: The CASNET and EXPERT Projects", Proc. AAAS Annual Meeting, Houston, January 1979, (in press) [also Rutgers Computer Science Report CBM-TR-98].

Kulikowski, C., (1979) "Representation of Expert Knowledge for Consultation; The CASNET and EXPERT Projects," Report CBM-TR-96, Department of Computer Science, Rutgers University, January 1979

Kulikowski, C., (1979) Current Research in Artificial Intelligence in Medicine in the United States", Proc. Japan-AIM Workshop, August 1979.

Kulikowski, C., (1979) "Computer-based Medical Decision-Making", Computers in Ophthalmology Conf., St. Louis, April, 1979.

Kulikowski,C., (1978) "Artificial Intelligence Approaches to Medical Consultation: A Tutorial Review, American Society of Clinical Pathology, October, 1978.

Kulikowski, C. and S. Weiss, (1978) "Laboratory Computers in the Year 2000", Medical Laboratory Observer, July, 1978, pp. 150-160.

Kulikowski, C.A. and Weiss, S.M., (1978) "The Evaluation of Performance in Empirical and Theoretical Models of Medical Decision-Making", Proc. IEEE Computer Society Workshop in Pattern Recognition and Artificial Intelligence, Princeton, April 1978.

Kulikowski, C.A. (1978) "Artificial Intelligence Approaches to Medical Consultation", Proc. fourth Illinois Conference on Medical Information Systems, May, 1978.

Kulikowski, C.A. (1978) "Strategies of Glaucoma Treatment Planning", National Computer Conference-78, AFIPS Conference Proceedings, Vol. 47, 1978, AFIPS Press, Anaheim.

Mitchell, T.M., (1979) "An Analysis of Generalization as a Search Problem", IJCAI, Tokyo, Japan, August 1979, [also Rutgers Computer Science Report DCS-TR-78]

Mitchell, T.M., (1978) "Version Spaces: An Approach to Concept Learning" PhD Dissertation, Stanford University, December 1978, (Also, Stanford Computer Computer Science Report STAN-CS-78-711, HPP-79-2).

Mizoguchi, M.K., Maruyama, T. Yamada, Y. Kitazawa, M. Saito, and C. Kulikowski, (1979), "A Case Study of EXPERT Formalism", Proc. IJCAI, Tokyo, August 1979, pp. 589-588

Morgenstern, M. (1978) "Transferring Technology from Research Systems to Users", Proc. The Third Jerusalem Conference on Information Technology, Jerusalem, August, 1978.

Nagel, D., (1979) "An Experimentation in Extracting Some Properties of Binary Relations", Department of Computer Science, Rutgers University, CBM-TM (forthcoming)

Nagel, D. and N.S. Sridharan, (1978) "A Test Procedure for Checking Combinations of Relational Flags in AIMDS", Report CBM-79, Department of Computer Science, Rutgers University, September 1978.

Politakis, P., (1979) "Evolution of an EXPERT Consultation System in Rheumatology", CBM-TR-99, September, 1979.

Politakis, P., S. Weiss, and C. Kulikowski, (1979) "Designing Consistent Knowledge Bases for Expert Consultation Systems", CBM-TR-100, September 1979.

Sangster, B.C. (1978) "Natural Language Dialogue with Data Base Systems: Designing for the Medical Environment", Proceedings of the Third Jerusalem Conference on Information Technology, August 1975

Schmidt, C.F., and N.S. Sridharan, (1978) Abstract, BELIEVER Project, report on the Workshop on "Social Plans and Language", P. Cohen and B. Bruce (eds.), SIGART Newsletter, No. 6, p. 5, August 1978.

Smith, D. and R. Smith, (1979) "Rules of Student/Teacher Interaction in Language Acquisition: A Computational Model", Report CBM-TR-93, DCS, Rutgers University, August 1979.

Sridharan, N.S. and Smith, D., (1978) "Design for a Plan Hypothesizer", Proceedings of the AISB/GI Conference, Hamburg, July 1978; (also, CBM-TR-85, DCS, Rutgers University).

Sridharan, N.S. and Schmidt, C.F., (1978) "Knowledge-Directed Inference in BELIEVER", In D. Waterman & F. Hayes-Roth (eds.), Pattern Directed Inference Systems New York, Academic Press, 1978, pp. 360-379.

Sridharan, N.S. (1978) "Guest Editorial", Special Issue on Applications of AI to the Sciences and Medicine, Artificial Intelligence Vol 11, No. 1, 2, August 1978.

Sridharan, N.S. and K.N. Venkataraman, "A Planning System Used for Plan Recognition in a Common Sense Domain", Report CBM-TM-83, Department of Computer Science, Rutgers University.

Sridharan, N.S., (1978) "Some Relationships between BELIEVER and TAXMAN", Report LRP-TR-1, Department of Computer Science, Rutgers University, December 1978.

Trigoboff, M.L. (1978) "IRIS: A Framework for the Construction of Clinical Consultation Systems", PhD Dissertation Rutgers University, May 1978; (also, CBM-TR-90, DCS, Rutgers University)

Van der Mude, A. and A. Walker (1978) "On the Inference of Stochastic Regular Grammars" Information and Control, (to appear).

Venkataramen, K.N. and N.S. Sridharan, (1978) "An AIMDS Implementation of a Plan Hypothesizer", Report CBM-TM-80, Department of Computer Science, Rutgers University, December 1978.

Weiss, S.M., Kern, K.B., Kulikowski, C.A. and Pincus, W., (1978) "An Interactive System for the Design of Classifiers in Diagnostic Applications", Proc. International Conference on Cybernetics and Society, Tokyo, November 1978.

Weiss, S.M., Kulikowski, C.A., and Safir, A., (1978) "Glaucoma Consultation by Computer", Computers in Biology and Medicine, 8:25-40, January 1978.

Weiss, S.M., Kulikowski, C.A., Amarel, S., and Safir, A. (1978) "A Model-Based Method for Computer-Aided Medical Decision Making", Artificial Intelligence Special Issue on Applications to the Sciences and Medicine (11, 1978); also, CBM-TR-63, DCS Rutgers University April 1978).

Weiss, S.M., Kulikowski, C.A., Mizoguchi, F. and Kitazawa, V. (1978), "A Computer-Based Comparison of Japanese and American Decision-Making" Proc. International Conference on Cybernetics and Society, Tokyo, November 1978.

Weiss, S., (1979) "The EXPERT and CASNET Consultation Systems", Proc. Japan-AIM Workshop, Tokyo, August 1979.

Weiss, S. and C. Kulikowski, (1979) "EXPERT: A System for Developing Consultation Models", Proc. IJCAI, Tokyo August 1979, pp. 942-947, [also Rutgers Computer Science Report CBM-TR-97].

Weiss, S., C. Kulikowski, and B. Nudel (1979) "Learning Production Rules for Consultation Systems, Proc. IJCAI, Tokyo, August 1979, pp 948-950.

Weiss, S., K. Kern, and C. Kulikowski, (1978) "A Guide to the Use of the EXPERT Consultation System", CBM-TR-94, November 1978.

E) Funding Support

The Rutgers Resource is funded through an NIH grant: Research Resource on Computers in Biomedicine. The NIH grant number is P 41 RR643. The Director and Principal Investigator is Dr. Saul Amarel of Rutgers--The State University of New Jersey.

This grant is in its third renewal extending for three years from December 1977 through November 1980. The total amount of the award for this period is $1,426,598 in direct costs. In the current year, December 1979 through November 1980, the funding level of direct costs is $451,383.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Dissemination; Interactions

The SUMEX-AIM facility provides one of the nodes where a good part of our collaborative program development and testing takes place (the other facility is RUTGERS/LCSR). These medical collaborations are described in I.B. above.

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.

As part of our collaborative activities with SUMEX-AIM in this area, we have continued our contribution to the preparation of the AI Handbook.

In order to increase the dissemination of AIM work within specialty fields of medicine, we have also presented tutorial papers at relevant conferences.

1) Fifth AIM Workshop (1979)

The fifth Annual AIM Workshop differed from the ones preceding it. It was a mini-workshop devoted to a single sub-area of AIM research: medical consultation problems and systems. This year, responsibility for the organization of the miniworkshop rotated to Dr. Peter Szolovits of MIT, with the intention that a system of rotation for the hosting will now develop (next year's will be held at Stanford). The Rutgers Resource retained a coordinating and funding role: Dr. Kulikowski worked with Dr. Szolovits in the organization of the miniworkshop, and the Resource covered the travel expenses of a number of AIM groups and individual participants.

The miniworkshop took place at the Talbot House, S. Pomfret, Vermont on May 7-20, 1979, and attendance was limited by a number of people that could be accommodated. A deliberate attempt was made to include only those investigators from medicine and computer science involved in the day-to-day activities of AIM research. Participation by graduate students was encouraged. there were 32 attendees, including 11 computer scientists, 14 graduate students and 7 physicians. In contrast to previous workshops, the concentration on a single research area allowed greater depth of discussion, and encouraged more informal interchange of ideas and opinions. It also served an important training function for graduate students and junior AIM researchers, who were not able to attend IJCAI because of its distant location (Tokyo) this year.

The general structure of the miniworkshop was as follows:

a) Five technical sessions on computer science issues of the representation of knowledge, explanation and justification, knowledge base acquisition and maintenance, and the computational methods of INTERNIST-11.

b) Three general AIM sessions on the assessment of the current state of progress of AIM methods and programs; the sources of difficulties in bringing these programs to clinical application; and an experimental discussion of knowledge-acquisition problems centered around a simulated CPC-style protocol analysis.

The miniworkshop format was successful in encouraging a more intimate exchange of ideas among workers in this AIM subfield. It would be desirable to hold other miniworkshops in the future in application areas such as psychology and biochemistry as the interest of the AIM researchers dictates. The format of a general workshop for all groups should probably be reserved for less frequent meetings.

2) AIM-Japan Workshop:

A one-day workshop was organized by Drs. Kaminuma, Kurashina, Kaihara and Mizoguchi to follow the IJCAI meeting in Tokyo. It was held on August 25, 1979, and was attended by over 100 biomedical researchers and clinical investigators. The workshop consisted of presentations by U.S. AIM researchers and Japanese investigators - two of them collaborators of the Resource.

Prior to the presentations, Drs. Kulikowski and Weiss gave demonstrations of the CASNET/Glaucoma and EXPERT/Rheumatology programs and Dr. Shortliffe demonstrated the MYCIN system. These were available for testing and study by the participants. In addition, Dr. Mizoguchi demonstrated the EXPERT/Glaucoma program that incorporates in its knowledge base a core of CASNET/Glaucoma and the new structures introduced by Dr. Kitazawa. The AIM-Japan Workshop had the effect of disseminating knowledge of AIM research among a large group of Japanese clinical investigators, and the attendance of a number of other IJCAI attendees lent an international character to the meeting.

3) IJCAI - International Joint Conference on Artificial Intelligence (Resource Participation):

Included among the other AIM activities that took place at the conference, Drs. Amarel and Mitchell chaired sessions, Drs. Kulikowski and Weiss presented two joint papers, and a third one was presented by our ophthalmology collaborators in Japan, Drs Kitazawa and Mizoguchi, who presented results of their experiments with the EXPERT knowledge-based methods. Drs. Kulikowski and Weiss demonstrated the CASNET/Glaucoma, EXPERT/Rheumatology and EXPERT/Thyroid consultation programs during the special demonstration periods at the conference. We connected by TYMNET via satellite to the Rutgers/LCSR DEC-20 and had excellent response time. This demonstrated the feasibility of remote collaborations and sharing of resources for the future. Dr. Mizoguchi demonstrated the EXPERT/Glaucoma program using Dr. Kitazawa's knowledge base, on the FUJIMIC DEC-20 in Tokyo, illustrating the practicality of knowledge base transfer methods.

4) National AIM Projects at Rutgers

The national AIM projects approved by the SUMEX-AIM executive committee were increased during the 1979 period of th Resource. A project using the BRIGHT system developed within the Resource and the NIH was continued in its application to various problems of clinical research by the group headed by Dr. W. Gordon Walker at the Johns Hopkins University and Hospital. We have projects that have their primary locus of activities on the SUMEX system and also use the Rutgers Resource for development, testing, and back-up functions. These include the MAINSAIL and CONGEN projects from Stanford University, and the INTERNIST project at the University of Pittsburgh. A project on medical knowledge representations at the Ohio State University was initiated on the Rutgers Resource Computer, as was a project in Artificial Intelligence models of clinical reasoning developed by Dr. R. Greenes at Harvard University. Dr. David Garfinkel at the University of Pennsylvania developed programs for his project on metabolic pathway modeling on the Resource Computer.

B) Critique of SUMEX-AIM Resource Management

We have now reached a steady state level of SUMEX-AIM usage--at least for the foreseeable future; we estimate it to remain at about 750 connect hours per year with an average compute to connect ratio of 1:25.

Since December 1979, the RUTGERS/LCSR computer is connected to the ARPANET again. Access to the ARPANET is facilitating close interactions between the Rutgers and Stanford AIM facilities, and in particular between the system staffs.

We continue to find the people support at SUMEX-AIM first rate and extremely helpful. On the technical side, we find communications via TYMNET of questionable reliability; and the SUMEX-AIM computer too heavily loaded.

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 and ophthalmology. Work on belief systems and commonsense reasoning will continue with emphasis on the psychology of plan recognition and handling of stereotypes. 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.

In October 1979 we submitted to NIH a renewal proposal for the Rutgers Resource. Our proposal for a five-year continuation (December 1, 1980 to November 30, 1985) was reviewed by a special study section earlier this April. A decision by NIH is expected in late May.

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 computing facilities (e.g. PERQ, DORADO, LISP machines, etc.) 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

Future development for hardware should be in the direction of smaller machines which could ultimately be acquired at or transferred to user's sites (e.g. VAXes or the larger personal computers). Special efforts in networking small machines and in developing methods of using small computers would be desirable. In particular, methods and technology for system transfer from large machine environments to small machines would be increasingly useful to the AIM community.

We continue to consider community developments as one of the significant goals of the national AIM project. The program of AIM Workshops should continue and new arrangements involving a program of lectures/seminars and working visits by AIM scientists should be encouraged.

9.2.9 Decision Models in Clinical Diagnosis [Rutgers-AIM]

A Goal-Oriented Model of Clinical Decision-Making Incorporating Decision Thresholds

Robert A. Greenes, M.D., Ph.D.
Harvard Medical School
Department of Radiology
Peter Bent Brigham Hospital
Boston, Massachusetts

I. Summary of Research Program

A. Project Rationale

The major objective of this project is to increase understanding of the way in which probability of diagnosis, and costs and benefits of contemplated actions, interact in the selection of the most appropriate actions. Actions include therapeutic as well as diagnostic procedures. The initial problem area being considered is the management of a patient with upper abdominal pain. The decision problem is modeled as a goal-oriented search process, where actions available for selection represent the system's goals. Costs and benefits of the actions are incorporated in the heuristic device of a probability threshold, which must be exceeded for the action to be taken. Production rules, modified to incorporate evaluation of probabilities in relation to thresholds, are used to embody the decision conditions.

Top-level patient management goals usually require for their adoption that diagnostic thresholds be exceeded, which in turn require specific groups of diagnostic tests to be obtained. Selection of these tests may recursively require still other diagnostic thresholds to be exceeded, involving other preliminary diagnostic tests. A forward-chaining search strategy is used, which restricts consideration to diagnoses exceeding minimal "rule-out" thresholds. When a group of diagnostic tests is eligible to be done, based on the above search, its members are placed in an eligibility list. Concurrent search may find several eligible test groups. Selection among eligible tests is based on heuristics such as: greatest likelihood of enabling a higher level goal to be reached, highest number of goals to which the test relates, and least overall cost or risk.

At the initiation of a session, information already known about a patient is entered. This, together with the system's estimates of prevalence of the various diseases considered, is used to "prime" the differential diagnosis probability distribution by means of Bayes theorem. The system identifies a next test to perform. As results on tests become available and are entered into the system, Bayes theorem is again used to update the distribution.

B. Medical Relevance and Collaboration

This model is somewhat unique in its attempt to incorporate both the decision analytic view of costs and benefits, and heuristics which make the problem more tractable. A probability threshold may be viewed as the indifference point for a decision maker, where the relationship between cost and benefit is equal for any of the available options. It can thus be used as a device for collapsing and summarizing an entire distal decision tree.

By maintaining the differential diagnosis in probabilistic terms, rather than using various surrogates for probability, the model is able to relate diagnostic probability to action, through the translation provided by the threshold concept. This permits behavior of the model to be analyzed and tuned either in terms of the accuracy of its probability estimates or the suitability of thresholds. Further, we believe that the use of a "rule-out" threshold, to permit the model to focus on limited subsets of the possible diagnoses, bears a resemblance to the focusing process which medical problem solving actually exhibits.

This project involves internal collaboration between the Departments of Radiology and Medicine at Peter Bent Brigham Hospital, with consultation by the Department of Biostatistics at Harvard School of Public Health, Bolt Beranek and Newman, Inc., and the Office of Medical Education, Michigan State University.

C. Highlights of Research Progress

In this first year of effort, major tasks have involved: (1) development of prototype programs for acquisition of decision rules, construction of probability tables, output of rules and probability data, and execution of the decision-making system; and (2) elaboration of specific decision rules and probability estimates for the management of patients with upper abdominal pain. An operational prototype of each of the programs described above now exists. In the application to abdominal pain, we consider approximately 90 diagnostic entities, 40 management goals, and 60 individual tests. Thus far, we have concentrated on the subset of patients suspected of having gastrointestinal obstruction, involving 6 diagnostic entities. The rules and probabilities are derived subjectively by periodic sessions with a gastroenterologist, T. E. Bynum, M.D.

D. Relevant Publications

[1] Greenes RA: A goal-directed model for investigation of thresholds for medical action. Proc Sympos on Computer Applications in Medical Care, Washington, DC, October, 1979, IEEE, pp 47-51.

[2] Greenes RA, McNeil BJ: The use of statistical measures as an aid to selection of appropriate diagnostic procedures. Proc 65th Scientific Assembly, RSNA, Atlanta, GA, November, 1979, p 257 (abstract).

[3] Greenes RA: The diagnostic test order decision problem. Proc of the 6th Illinois Conf on Medical Information Systems, Champaign-Urbana, Il, April, 1980 (in press).

[4] Greenes RA: Medical decision-making research: the role of academic radiology. Third Int Sympos on the Planning of Radiological Departments. Amsterdam, Holland, June, 1980.

E. Funding Support

This project is part of a Program Project, "Investigations in Clinical Decision-Making", supported by the National Library of Medicine, grant NLM 1 P01 LM03401, Robert A. Greenes, M.D., principal investigator. Total award (7/1/79-6/30/84), $1,177,582. First year (7/1/79-6/30/80), $235,582.

II. Interactions with the SUMEX-AIM Resource

A. Medical Collaboration and Program Dissemination Via SUMEX

With this project only in its early stages, no significant medical collaboration or dissemination has yet occurred. Because of the specific features of our model, it was not considered to be readily implemented within the framework of other extant decision-making systems.

B. Sharing and Interactions with Other SUMEX-AIM Projects

The project utilizes the PDP-20 AIM resource at Rutgers University. In making the decision to utilize the Rutgers rather than the SUMEX facility, much assistance and documentation was provided by the technical directors of both facilities. Our choice of Rutgers was based primarily on the expectation that the response time and communication support through TYMNET for an east-coast user would be likely to be better.

We have participated in the site visit to the AIM resource at Rutgers in April, 1980, and will also be participating in the AIM Workshop, and Artificial Intelligence in Medicine Continuing Education Tutorial, at Stanford, in August, 1980.

C. Critique of Resource Management

We are most pleased by the personal interest shown, and assistance provided, by the technical directors of both AIM resources. We have had no serious problems with the use of the Rutgers facility.

III. Research Plans (8/80-7/86)

A. Project Goals

Near term goals involve (a) expansion of the decision rules and probability matrix to include the entire range of diagnostic entities considered in our model of upper abdominal pain, (b) incorporation of the time duration involved in diagnostic tests into the selection process, and

(c) evaluation of model performance. Initially, our criteria for evaluation will be agreement with an expert regarding the management goals selected, and the tests utilized.

Long-range goals include (a) refinement of the human interaction with the system, (b) the incorporation of capabilities for explaining its decisions, (c) evaluation of sensitivity of its conclusions to estimates of the various probabilities and thresholds, and (d) incorporation of empirical probability data into the model when available. Ultimately, we would like to evaluate its suitability as a consultant.

B. Justification and Requirements for Continued SUMEX Use

We expect the complexity of our programs to grow considerably during the next 2-3 years. The knowledge and data bases will grow also, and we anticipate moderately large storage requirements. Continued availability of the SUMEX-AIM resource is thus highly desirable in terms of our need for LISP programming capabilities.

In addition, we would hope that closer interaction with other AIM users in the evaluation of our model and other approaches will become possible, as we both familiarize ourselves with the characteristics of other systems, and further develop the capabilities of our system.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM.

This project has no present need for other computing resources, although some of the probability data incorporated into the model are derived from studies carried out on other computer systems.

D. Recommendations for Future Community and Resources Development

As microcomputers become increasingly powerful, inexpensive, and capable of supporting at least single-user LISP programs, we expect that a natural evolution toward such systems will occur. Efforts to ensure compatibility, portability, and the ability to interface such systems to the AIM network will, thus, be highly desirable.

9.2.10 Heuristic Decisions in Metabolic Modeling [Rutgers-AIM]

Heuristic Decisions in Metabolic Modeling

David Garfinkel, Ph.D.
Moore School of Electrical Engineering
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, Pennsylvania

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

This research is concerned with developing methods of constructing computer models of complex metabolic systems, and applying thereto artificial intelligence and other relevant computer science techniques.

B. Medical Relevance

Most of our work is concerned with modeling cardiac metabolism and the effects of ischemic heart disease. There has been some collaboration with appropriate cardiac physiologists and cardiologists. We are trying to extend our research to diabetes and hematology, in collaboration with experts in those particular fields.

C. Research Progress

We constructed cardiac metabolism models, emphasizing the effects of acidosis, fatty acid metabolism, and the glycogenolysis cascade; this work is still in progress. We have developed methods for sensitivity analysis and a new program for building these models. Current efforts emphasize completing the glycogenolysis model, correcting some of the existing models, designing relevant experiments, and writing up a mass of completed models. We are extending our model-building software, and making it "friendly" enough for unsophisticated users.

D. List of Relevant Publications

M.J. Achs and D. Garfinkel, Metabolism of the Acutely Ischemic Dog Heart. I. Construction of a Computer Model. Am. J. Physiol. 236, R21-R30 (1979).

D. Garfinkel and M.J. Achs, Metabolism of the Acutely Ischemic Dog Heart. II. Interpretation of a Model. Am. J. Physiol. 236, R31-R39 (1979).

D. Garfinkel, M.J. Achs, M.C. Kohn, and L.S. Menten, Modeling of Complex Metabolic Systems as a Composite of Standard Computer Science Techniques, Proceedings of Summer Computer Simulation conference, 1979.

L. Garfinkel, M.C. Kohn, and D. Garfinkel, Computer Simulation of the Fructose Bisphosphatase-Phosphofructokinase Couple in Rat Liver, Eur. J. Biochem. 96, 183-192 (1979).

M.C. Kohn, L.E. Menten, ad D. Garfinkel, A Convenient Computer Program for Fitting Enzymatic Rate Laws to Steady-State Data, Comput. Biomed. Res. 12, 461-469 (1979).

M.C. Kohn, M.J. Achs and D. Garfinkel, Computer Simulation of Metabolism in the Pyruvate-Perfused Rat Heart. I. Model Construction. Am. J. Physiol. 237, R153-158 (1979).

L.E. Menten, M.C. Kohn and D. Garfinkel, A Convenient Computer Program for Estimation of Enzyme and Metabolite Concentrations in Multienzyme Systems (in progress).

E. Funding Support

1.) "Computer Simulation in Cardiology" HL 15622, 3 years, Dec. 1, 1977-Nov. 30, 1980; current year $111,051. Competing renewal now pending, $162,744; $181.060; $199,548; $220,647 (for four years).

2.) "Computer Modeling Methods for Metabolic Systems," (GM 16501-11A1), $60,598; $63,860 (two years) April 1, 1980 - March 31, 1982.

3.) "Computer-Aided Study of Glycolysis in Pancreatic Islets," submitted to NIH as a part of the Diabetes Center renewal proposal. Direct costs $38,853; $41,652; $44,322 for 3 years.

4.) "Metabolism of Malarial, Aged, and Normal Erythrocytes" (including experimental subcontract) submitted to NIH, $144,283; $158,267; $169,609 (3 years).

II. Interactions with the SUMEX-AIM Resource

There has been no interaction with SUMEX directly, and none is anticipated. There has been considerable interaction with the Rutgers resource, especially in the form of a personal collaboration with Prof. Kulikowski, and moderate usage of their computer. The program we have developed and are writing up was developed there. I have also been a regular attendant at AIM workshops over the last few years, and have made valuable contacts and acquired an understanding of the field that would not have been possible otherwise. Overall, interactions through the Rutgers resource have been of considerable importance to our development.

III. Research Plans

We are now at the stage where our technique research is concerned with knowledge representation and acquisition, with intelligent (and fuzzy) data bases having some of the characteristics of a knowledge base, with representation of knowledge in our subject area of interest, and with overcoming incompleteness. We cannot realistically expect much help in these matters from the subject-matter experts, and need the help of the AI community. We expect to continue using the Rutgers machine. Within the next year and a half we will try to refine our existing program and other model-building methods, to make them more manageable, modularized, efficient, and faster, and also friendly enough to be operated by biological experts directly without too much help from programmers or professional simulationists. We will also devise methods (using symbolic manipulation) to break large complex models down into workably small pieces. In the next few months we hope to start attacking the data-base (and later knowledge-base) aspects of this work.

Long-range research goals will be critically dependent on other funding, so we cannot give details now. We hope to be able to build a good model of cardiac ischemia which can be used to make predictions, to design experiments for areas in metabolism (especially multi-enzyme systems) which are now inefficient. This modeling process must be _fast enough to be of use_ and _rigorous enough to be reliable_. This implies development of the techniques mentioned above to the point where they can do most of the necessary work.

We do not expect to use the SUMEX computer, but do expect to make considerable use of the Rutgers computer and the expertise of the department, since such expertise is not otherwise available to us. It is conceivable that several years from now we may want to link in a personal computer under the collaborative linkage described in the RENEWAL RATIONALE, but this is sufficiently far in the future that a justification for one cannot be given at the present time.

9.3 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.

9.3.1 Ultrasonic Imaging Project

Ultrasonic Imaging Project

James F. Brinkley, M.D.
W.D. McCallum, M.D.
Depts. 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 plan to use is the reconstruction of an organ from a series of ultrasonic cross-sections taken in an arbitrary fashion. A real-time ultrasonic scanner will be 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 will be used to record a series of scans over the organ being modelled. The scans will be recorded on a video disk which is controllable by the microcomputer. 3D position information will be stored on a floppy disk file. 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 then 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 plan to develop this system in phases, starting with an earlier version developed at the University of Washington. During the first phase the previous system will be adapted and extended to run in the SUMEX environment. A clinical study will then be carried out to determine its effectiveness in predicting fetal weight. At the same time computer vision techniques will be used 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 its use in predicting fetal weight, this system could be used to determine other organ volumes such as that of the left ventricle. 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 as yet 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 essentially completed the first phase of this project which was to implement and adapt the previous system to the SUMEX environment. The accomplishments related to that goal are:

1. Completion of a microprocessor based data acquisition system

The following hardware has been obtained and integrated into the system--

a) A Toshiba real-time ultrasonic phased array scanner, in routine clinical use at the Dept. of Obstetrics.

b) A Sony video tape recorder and Hitachi monitor, for use in recording the scans prior to their being outlined with the light pen.

c) A custom built acoustic position locating system for determining the position of the scan plane in space, supplied to us by W.E. Moritz at the University of Washington.

d) A Datamedia computer terminal for communicating with SUMEX and controlling the procedure.

e) A microprocessor-based video graphics system supplied to us by Varian Corporation. This system includes a light pen, dual floppy disks and video display memory.

A large amount of software for the data acquisition system has been written and is now working. This software consists of routines to direct the patient exam, during which time scans are recorded on video tape and position information is stored on floppy disk. Additional software directs that the scans be outlined with the light pen (not digitized in this first phase) and stored with the position information. Finally, a program has been written which converts the microprocessor into a video graphics terminal. Characters are passed back on forth as with an ordinary terminal but special command sequences cause graphics to be displayed and a file transfer to take place. The file transfer, called File Transfer to Micro, is packet oriented and should prove useful to anyone wanting to do file transfers between SUMEX and a microcomputer. It is also flexible enough so that it can form the basis of a system for sending commands from SUMEX to do local image processing functions.

2. Completion of SUMEX high level routines

The SUMEX software for this first phase includes procedures to transfer the data from floppy disk to SUMEX (via the file transfer protocol), to build a 3D reconstruction using simple interpolation, and to display the images on a graphics terminal.

3. Initial tests and first patients

The completed first phase system has been used to build 3D displays of simple models such as cylinders in a water tank. We have also tried it on 2 patients and have obtained 3D plots of the fetal head and trunk and of the placenta inside the uterus.

The research currently in progress relates to testing the system:

1. An engineering study is being carried out on cylinders, balloons and point targets to determine the bench accuracy of the system.

2. A clinical protocol is being established on several obstetrics patients. Once we have gained enough experience we will begin our clinical study to determine the ability of the method to predict fetal volume and weight.

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.

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

"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

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 discussing the the applicability of the current system to the heart 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 completed we expect to interact much 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 plan to implement 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 and test its ability to predict fetal weight.

This system has been developed and is now undergoing testing. We plan to carry out engineering and clinical studies in order to test the ability of the current system to predict fetal and cardiac volume. If successful the system may have clinical impact as it stands. 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. For this aspect of our research the SUMEX-AIM community should be especially useful.

3. Develop program, as outlined in the introduction, with several limitations--

Only a simple organ will be modelled at first, i.e. 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

The only additional requirements we have at present are for an additional file directory and for a little more time in the afternoon. At present we only have one directory which must be shared by the system developer and an additional person conducting the engineering and clinical studies. An additional directory could be designated for users of the current implementation of the system while the present directory could be used for new developments.

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.

3. 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.

9.4 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.

9.4.1 Coagulation Expert Project

Coagulation Expert Project

Donald Lindberg, M.D.
University of Missouri
Columbia, Missouri

I. SUMMARY OF RESEARCH PROGRAM

A. Project rationale

Preliminary experiment in attempting to form a clinical consultant program based on a formal representation of medical knowledge of the blood coagulation (or clotting) 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 the University of Missouri-Columbia consists of the following individuals:

Lamont Gaston, M.D.
David Goldman
Lawrence C. Kingsland III
Donald A. B. Lindberg, M.D.
Haruki Ueno, Ph.D.
Anthony Vanker. Ph.D.

Dr. Gaston is a consulting hematologist, director of a coagulation laboratory, and co-director of a blood-banking service.

Expertise in the field as well as clinical laboratory and patient records are being provided by UMC to build and test the consultant. In the future we plan to incorporate the views of external experts as well.

A formal research proposal to NIH is planned for fall, 1980, based on the studies performed on SUMEX.

C. Highlights

-Accomplishments

Use of UNITS/AGE: an initial model has been created on SUMEX.

Experimental use of EMYCIN: a feasibility test with a text book 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 full coagulation screen (of 6 tests), then the partial coagulation screen (of 3 tests), and finally the individual determinations. In all these cases, laboratory and clinical features will be taken into account.

-Research in progress

Currently we are testing the initial models against actual clinical records for 270 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. In any event, the data gathering is almost complete and testing is about to begin.

D. List of relevant publications

None

E. Funding support

This preliminary research phase is being supported from two sources:

1. USPHS Grant No. T15 LM 07006, "Training Program in Medical Information Science". Full funding is $162,410/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 is being devoted to this project.

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, N.L.M. Training Programs, in May at Columbus, Ohio. He also plans to demonstrate our AGE-implemented model during the meeting. 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 (AI) systems in development at SUMEX. In fact, with the help of members of the SUMEX-AIM community, he was 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 in the study group are planning to attend the AIM-tutorial at Stanford in August, 1980.

Since the AI systems in which we are interested are in some stage of development on the SUMEX computer, and since partial documentation does exist, 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 received prompt and helpful information from various members of the SUMEX-AIM community via the network electronic message system.

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 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 have been 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 thus far.

III. RESEARCH PLANS

A. Plans for Summer, 1980

1. Continue assembling the knowledge data base, with emphasis on documentation of the primary literature sources for the knowledge sources (KS).

2. Continue learning the various aspects of UNITS/AGE and EMYCIN.

3. Continue comparing the two potentially complex models with the inherently simpler microprocessor version.

4. Appoint a clinical test panel for consultation on development of the next features.

5. Prepare the application to NIH.

Our long range plans are to develop consultation systems in anticoagulation therapy and in the interpretation of other hematological laboratory results.

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 designed to build other knowledge-based programs. To be deprived of the ability to study how these systems handle the knowledge and the actual consultation problems would be a serious impediment to our long range plans.

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 better understand 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.

The question of our continued use of both EMYCIN and AGE raises the serious problem of our exceeding our storage allocation. We are prepared either to settle on one, or to propose a more formal comparison of the two systems. The choice and the mechanism for the comparison will be made in concert with SUMEX management. We feel the comparison would help us to gain insight into these systems but would make these requests only if it were clear that such a study would be of interest to others on the SUMEX resource.

9.4.2 Communication Enhancement Project

Communication Enhancement Project

John B. Eulenberg and Carl V. Page
Michigan State University

I) Summary of Research Program.

A) Technical goals.

The major goal of this research is the design of intelligent speech prostheses for persons who experience severe communication handicaps. Essential subgoals are:

(1) Design of input devices which can be used by persons whose movement is greatly restricted.

(2) Development of software for text-to-speech production.

(3) Research in knowledge representations for syntax and semantics of spoken English in restricted real world domains.

(4) Development of micro-computer based portable speech prostheses.

B) Medical Relevance and Collaboration.

Members of our group are in touch with Dr. Kenneth Colby and his group at UCLA who have been working on similar problems for people who have aphasia.

The need for such technology in the medical area is very great. Millions of people around the world lead isolated existences unable to communicate because of stroke, traumatic brain injury, cerebral palsy or other causes. The availability of inexpensive micro-processors and voice synthesizers allows development of complex experimental systems to study human communication. The knowledge gained from these experimental systems should lead in a few years to prototypes of very low cost which will permit many people to engage in the vital acts of communication required for a "normal" life in human society.

Despite the importance of the problems in this area, it has been difficult to coordinate the many professions which are involved. We believe that both research and the support of research in this area suffers from the lack of an identifiable community of workers. To alleviate this problem, we have joined with the Trace Center of the U. of Wisconsin to publish the first newsletter for dissemination in this area Called "Communication Outlook", the first issue was published in April, 1978. There are now over 1100 paid subscribers. Subscribers and contributors to the Newsletter come from a vide variety of disciplines and from many countries. John B. Eulenberg helped to organize the first Federal workshop for governmental agencies who have some interest in funding work in these areas. Represented were the Bureau of Education for the Handicapped, The Veterans Administration, The Civil Service Commission, NIH, NSF, and others. We have also been in touch with United Cerebral Palsy associations at the state and national levels. Much of our effort has been in educating those medical, educational, and governmental communities with an interest in this area on the available technology since most of them are not accustomed to funding the development of high-technology systems.

C) Progress summary.
Although some facets of the research have been underway at MSU for several years, we have been using SUMEX-AIM for three years, having received our password in March, 1977.

During the last past three years, we have:

1) Organized a research team of 4 students possessing background in artificial intelligence lead by Dr. Carl V. Page to start a semantics-speech generator. This group had a very primitive prototype (written in Sail) running in June, 1977. The system uses statistical, grammatical and semantic information to generate sentences by anticipation. A similar group was organized in 1978 but it produced well documented but not fully debugged programs.

2) Converted a large program (Orthophone) for English text to speech synthesizer codes to SAIL from Algol.

3) Obtained local support for terminals and space to use the SUMEX-AIM facility. At present, the lack of a dedicated tie-line from East Lansing to Tymshare in Ann Arbor or Detroit is a problem for us during 0600 to 0900 PST.

4) In 1978, Dr. Reid of our project designed and built a wheel-chair portable personal communication system for a 10 year old boy who has cerebral palsy. It is micro-computer based and can accept inputs via an adaptive switch from a series of menus displayed on a TV screen, via Morse code, or by a keyboard. Its outputs can be TV display, hard copy, spoken English, Morse code, or musical sounds. As the memory available for small systems will soon be substantial, we will need to specify the content and connection of the choice menus using the knowledge gained in our SUMEX-AIM project. Although our prototype for semantic generation has not run satisfactorily, it has influenced the design of the next system, the "SAL" board for wheel-chairs described below.

5) During 1979, a communication aid using knowledge sources has been built into a lap board. Called the "SAL" prosthesis (Semantically accessible language), it uses a magnetic input to translate Bliss symbols into spoken language. Some ideas from the grammatical portions of our SUMEX-AIM project have found their way into the SAL system. The SAL system consists of an aluminum encased lap tray with an array of 252 reed switches arranged in a 12 row by 21 column matrix. Spacing between switches is one inch. They are activated by a small magnet held by the user on a mitt or a finger splint. The keyboard is interfaced to a Southwest Technical

Products 6800 computer possessing 8K of EPROM and 8K of RAM. Voice output is from a Votrax VS-6 sound synthesizer while visual output is provided by a LED array. The current system allows 512 lexical items. Frame cells provide a choice of syntactic frame, which the user may specify at the inception of the formation of a sentence to supply structural information. Each syntactic frame is a skeletal syntactic phrase marker representing a class of sentence structures. After choosing a given syntactic frame, the user goes on to choose the lexical items. The generation of appropriate pronouns depends on their role in the sentence. Thus the Bliss symbol for the speaker will come out "I" or "me" depending on the role. The system uses syntactic, phonetic, and orthographic features of previous inputs in order to generate its outputs. We expect to gain experience from our SUMEX-AIM prototype to guide the choice of semantics for the successors of this system. Ten SAL boards are being built for students now.

6) Dr. John Eulenberg began his Sabbatical leave in Palo Alto beginning in September, 1979. He has been associated with the Children's Hospital at Stanford and Telesensory Systems Inc. We have found in the past that SUMEX-AIM has provided us with a means to communicate with other members of our project when they were California. It is very important for the many ongoing projects which we have to be able keep Dr. Eulenberg in close communication with the rest of our project during his leave.

7) We have built and tested a myoelectric interface and used it (together with a miniature FM transmitter) for input of changing muscle potentials into a computer. There is reason to believe that this means of input may provide a higher bit rate than other known means for those people who possess severe cerebral palsy.

8) We continue to develop basic educational software for severely impaired persons. For example we have developed a "talking" system for drilling students in Bliss symbolics. Another system we have developed teaches spelling using a voice synthesizer and TV screen. A classroom in a Northville, Michigan public school now contains a Nova 2/10 for the evaluation of our systems.

D) Up-to-date list of publications. (1976 to date)
By John B. Eulenberg

"Technical Systems Development, Headin", Interim Report, April, 1976, Experimental Applications of Two-Way Cable Delivery, NSF Grant No. APR 75-14286.

"Interactive New Hired Information Access System with Both Voice and Hard Copy Output: User's Guide to NHQUERRY", April 11, 1976 (With Steven Kludt and Jerome Jackson (Artificial Language Laboratory Report AEB 041176))

"Language Individualization in a Computer-Based Speech Prosthesis System", National Computer Conference, New York, June 9, 1976.

"Individualization in a Speech Prosthesis System", Proceedings of 1976 Conference on Systems and Devices for the Disabled, June 10, 1976.

"The LEAF Language", Interim Report, September, 1976, NSF Grant No. APR 75-14286.

"Microprocessor-Based Artificial Language for Communication Prostheses", with M. R. Rahimi, Proc. of the National Electronics Conference, Vol. XXXI, October, 1977.

"A programmable Multi-Channel Modem Output Switch", September 22, 1976, with Joseph C. Gehman and Juha Koljonen (Artificial Language Laboratory Report AEB 092276)

"SMPTE Time Code Interface and Computer-Controlled Video Switcher", with Michael Gorbutt and Dennis Phillips, Interim Report, March, 1977 NSF Grant APR 75-14286.

"Representation of Language Space in Speech Prostheses", with R. Reid and M. Rahimi, Proc. of Fourth Annual Conference on Systems and Devices for the Disabled, June, 1977.

"Administration and Management of a Computer-Based Communication Enhancement Program", with M. R. Rahimi and L. Neiswander, Proc. of Amer. Acad. for Cerebral Palsy and Developmental Medicine, October, 1977. "When [-VOICE] becomes [+VOICE]- The Phonological Competence of People Who Cannot Speak", with Carol Myers Scotton, Proceedings of the Annual Confer. of the Linguistic Soc. of America, December, 1977.

"Toward a Semantically Accessible Communication Aid", (With M. A Rahimi) Proceedings of the National Electronics Conference, Vol. xxxii, Chicago, Illinois, October, 1978.

By Carl V. Page:

"Heuristics for Signature Table Analysis as a Pattern Recognition Technique", IEEE Transactions on Systems, Man and Cybernetics,Vol. SMC-7, No. 2, February 1977.

"Discriminant Grammars, an Alternative to Parsing". with Alan Filipski, Proceedings of the IEEE Workshop on Picture Processing, Computer Graphics, and Pattern Recognition, April 22, 1977.

"Pattern Recognition and Data structures". Chapter in "Data Structures in Computer Graphics and Pattern Recognition" Edited by Allen Klinger, Academic Press, 1977.

"A Survey of Artificial Intelligence in Computer-Aided Instruction", with Alice Gable (To appear in the International Journal of Man-Machine Systems, 1979)

"Economic Consequences of Robots Possessing Computer Vision Systems", Proceedings of the Upper-Midwest Small College Computer Conference, St. Cloud State University, St. Cloud, Minn. April, 1979.

E) Funding Status.

1) Current funding.

SOURCE	AMOUNT	PERIOD
Division of Engineering Research, Michigan State University	$45,000	Sept. 1, 79- Aug. 31, 80
Northville Public Schools	$90,000	Sept. 1, 79- Aug. 31, 80
Jackson County Schools	$10,000	Sept. 1, 79- Aug. 31, 80
Jackson Foundation	$10,000	Sept. 1, 79- Aug. 31, 80
United Cerebral Palsy Foundation	$42,000	Sept. 1, 79- Aug. 31, 80
National Science Foundation (Eng-7907753)	$98,000	Sept. 1, 79- Aug. 31, 80
State of Michigan Vocational Rehabilitation	$15,000	Sept. 1, 79- Aug. 31, 80

Commitments in the grants have prevented us from using very much of these funds to support long range goals such as those communicated to SUMEX-AIM. However, the special communication devices, student and other research facilities provide the critical mass which will allow us to do the work that we have proposed. The main value of SUMEX-AIM to us is to allow experimentation with AI technology in order to develop the experience to design intelligent speech prostheses.

2) Pending applications and renewals.

Oakland County Intermediate School District - $100,000.
Genessee County Intermediate School District - $100,000.

Tuscola County Intermediate School District - $20,000.
Livingston County Intermediate School District- $50,000.

II. INTERACTION WITH SUMEX-AIM RESOURCE

A. Collaborations and medical use of programs via SUMEX.

We have shown Mycin and Puff to physicians and clinical staff and discussions continue with them concerning possible research. During a visit to our campus in October 1978, Dr. Bruce Buchanan lectured on Mycin and stimulated some of our Medical School faculty to explore research opportunities with us. As a consequence, Dr. Carl V. Page has participated in a proposal to NSF with Dr. Su-Wah Chan (principal investigator) titled "A Structural Analysis of Problem Complexity in Information Processing Behaviors as Related to Human Problem Solving". We hope that some other research possibilities derive from this effort.

B. Sharing and interactions with other SUMEX-AIM projects.
During the past year we have had personal contact with the SUMEX-AIM staff. Dr. Eulenberg attended the 1978 Workshop in the Summer. Dr. Page used the facility while working in California as a means keep in touch with the project in East Lansing. The communication aspect of the project has been useful for us in the past and will continue to be so in the future inasmuch as Dr. Eulenberg is spending his Sabbatical in Palo Alto.

C. Critique of resource management.

We have found the staff to be professional and helpful. We have not used the system enough to comment on the management of the facility except to say that we have become somewhat disillusioned with the SAIL compiler.

III. RESEARCH PLANS (8/79-7/81)

A. Long Range project goals and plans.

We will continue to explore the interactions of different knowledge sources in the problem of generation of language. Such information as is learned will be scaled down so that it can be used in the design of portable, intelligent, speech prostheses.

B. Justification and requirements for continued SUMEX use.

We do not require any more resources than we have had in the past. Unfortunately our SUMEX research has not had the priority with us that it deserves. In one sense, our SUMEX research represents to us the future of work in this area, but we are involved with commitments for systems for communication enhancement that must be delivered soon. We expect to change the pattern of our funding to emphasize the kinds of problems we have addressed to SUMEX, beginning the process next year. Our prototype system on SUMEX has been built by volunteer student effort rather than our financial support. We hope to change this policy when pressing needs are satisfied. Our prototype has already has had some influence on the design of a wheel-chair portable system, the SAL prosthesis mentioned above. We have planned to incorporate at least one Ph. D. thesis into this research area. One of our former employees, Mr. Douglas Appelt has been doing his thesis in this area at Stanford and we believe that it is a good area. However, before we can advise a a student to start a thesis dependent on the system, we need assurance that we will have access to SUMEX for at least two years at some reasonable level comparable to what we have now.

C. Other Computational needs.

We use minicomputers and the central computers at MSU in addition to SUMEX. We have no plans to secure any additional equipment.

D. Recommendations for future community and resource development.

1. We have not heard much lately about the KRL language. If it is available or can be made available, we would be interested in considering it for our project.
2. We would be interested programs to help scale down a system developed on SUMEX-AIM to smaller machines.
3. We are interested in programs to facilitate the hardware design process for microcomputer based systems.

9.4.3 A Computerized Psychopharmacology Advisor

A Computerized Psychopharmacology Advisor

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

I. Summary Research Program

A. Technical Goals.

We are developing a computer-based automated system for education and consultation in clinical psychopharmacology. Our technical goals are envisioned in three phases:

- To develop a theory of expert teaching, consulting and decision-making in clinical psychopharmacology.
- To model this theory on a computer system which responds in real time and communicates in natural language.
- To evaluate this theory and model as a representation of psychiatric knowledge by analyzing both the performance of the system and the effort required for the system's development.

B. Medical Relevance and Collaboration.

1. Medical Relevance.

For many years, it has been recognized that potent psychopharmacological agents are frequently used in an unsystematic manner. There are at least 50 discrete syndromes currently identified in clinical psychiatry which have unique hierarchies of plausible pharmacological treatments. Each therapeutic regimen in each hierarchy may involve several classes of drugs which can often be preferentially ranked. A particular member of a class of drugs may be recommended on the basis of a patient's medical history, family history, response to previous treatments, current physical status, or current mental status. In addition, each treatment program has its own set of potential side effects, adverse reactions and drug-drug, drug-host, drug-age, drug-gender, drug-state of health, and drug-other treatment interactions.

Conventional sources of information for education or verification (books, journals, lectures, and seminars) are seldom quickly accessible or specifically pertinent. A traditional alternative is to consult a specialist. In addition to availability, reliability and validity, a good consultant has the ability to understand questions in their proper context and sequence, to give advice which can be explained or documented as needed, and to provide follow-up consultations which incorporate new information from clinical developments or additional expertise.

Our research on the Clinical Psychopharmacology Advisor is directed towards implementing all of the characteristics of a good consultant, which have only been outlined above, in a functional computer program. To our knowledge, no other computer program currently available, or under development, is pursuing all of these goals in clinical psychopharmacology.

2. Collaboration.

2.1 Principal Investigator: Jon F. Heiser, M.D., Associate Professor, Department of Psychiatry and Behavioral Sciences

2.2 Co-principal Investigator: Ruven E. Brooks, Ph.D., Assistant Professor, Department of Psychiatry and Behavioral Sciences

2.3 Pharmacist, University of Texas Medical Branch:
Carla Maria Brandt, B.S. (January 1979-present)

2.5 National Advisory Panel:

John M. Davis, M.D.
Illinois State Psychiatric Institute
1601 West Taylor Street
Chicago, Illinois 60612

Max Fink, M.D.
Department of Psychiatry
State University of New York at Stony Brook
Stony Brook, Long Island, New York 11794

Neal R. Cutler, M.D.
National Institute of Mental Health
9000 Rockville Pike, Building 10, Room 3S205
Bethesda, Maryland 20014

John H. Greist, M.D.
Department of Psychiatry
University of Wisconsin
Madison, Wisconsin 53706

Leo E. Hollister, M.D.
Departments of Medicine and Psychiatry, Stanford University
Veterans Administration Hospital
Palo Alto, California 94302

James W. Jefferson, M.D.
Department of Psychiatry
University of Wisconsin
Madison, Wisconsin 53706

Donald F. Klein, M.D.

New York Psychiatric Institute
722 West 168th Street
New York, New York 10032

George M. Simpson, M.D.
Psychopharmacology Unit, University of Southern California
Metropolitan State Hospital
Norwalk, California 90650

Robert L. Spitzer, M.D.
New York Psychiatric Institute
722 West 168th Street
New York, New York 10032

Zebulon C. Taintor, M.D.
Rockland Research Institute
Rockland State Hospital
Orangeburg, New York 10962

C. Progress Summary.

Our initial goal has been to develop a small, but fully functioning, Clinical Psychopharmacology Advisor. Approximately 250 rules, utilizing about 120 clinical parameters, were developed and used to diagnose and recommend therapy. The system, affectionately called HEADMED, had sound knowledge about the differential diagnosis of the major affective disorders and schizophrenia. The Psychopharmacology Advisor had perfunctory information concerning paranoid disorders and personality disorders. HEADMED also had skeletal knowledge about neuroses, behavior disorders, substance abuse, organic brain disorders, including both the type of brain disorder (e.g. delirium or dementia), and the cause of brain disorders (e.g. intoxication or trauma). The program has never known anything about child psychiatry, sexual disorders and other psychiatric conditions.

The HEADMED software had the capability of recommending a drug treatment, if indicated, and of cautioning about potentially harmful interactions with a compromised host and with other chemical substances. The system also could print out advice concerning dose and duration of therapy, pharmacokinetics, warnings about common side effects and possible adverse reactions.

Having been satisfied with the feasibility of using EMYCIN as a language for performing consultations in clinical psychopharmacology, our interest shifted to critically evaluating this application of EMYCIN and to modifying data structures and control mechanisms so that a a consultation process which is more natural, complete, and accurate occurs. We had planned concentrating our attention on psychiatric disorders whose management might include prescription of a tricyclic antidepressant, one of the major classes of psychotherapeutic medications, and on outputting individual case-oriented advice and precautions concerning management and monitoring of a patient receiving a tricyclic antidepressant medication (see reference I.D.5 below). Thus we have begun to develop knowledge structures which can utilize this information to compute diagnostic formulations and therapeutic plans which are highly specific to the unique properties and circumstances of a particular patient.

We have discovered what we believe is an essential design problem for medical expert systems, that of controlling the amount and the type of information which the system requests from the user. This problem is inherent in medical expert systems because of the nature of the distribution of clinical states, and the nature of the training and the background of physicians. The problem also exists for human consultants, and a complete and general solution for computer systems is probably not achievable. However, several techniques show promise for reducing the magnitude of the problem in various clinical domains. These include system use of dynamic and static domain models, user control over sophistication level, and user access to the rationales behind information requests.

A prolonged illness of the Principal Investigator prevented significant progress during the past year.

D. List of Relevant Publications.

1. Heiser, J.F., Brooks, R.E., and Ballard, J.P. Artificial Intelligence in Psychopharmacology. Abstracts - VI World Congress of Psychiatry, Honolulu, Hawaii, 28 August - 03 September 1977, page 135.

2. Heiser, J.F., Brooks, R.E., and Ballard, J.P. A Computerized Psychopharmacology Advisor. Continuing Medical Education Syllabus and Scientific Proceedings in Summary Form, The 131st Annual Meeting of the American Psychiatric Association, Atlanta, Georgia, 8-12 May 1978, American Psychiatric Association, Washington D.C., 1978, page 216.

3. Heiser, J.F., Brooks, R.E., and Ballard, J.P. A Computerized Psychopharmacology Advisor. The 11th C.I.N.P. Congress, Collegium Internationale Neuro-Psychopharmacologicum, 9-14 July 1978, Vienna, Austria, Book of Abstracts, page 233, c/o INTERCONVENTION, Kinderspitalgasse 5, A-1095, Vienna, Austria.

4. Heiser, J.F. and Brooks, R.E. Design Considerations for a Clinical Psychopharmacology Advisor. In Orthner, F.H. (ed.), Proceedings, The Second Annual Symposium on Computer Applications in Medical Care, 5-9 November 1978, Washington, D.C., Institute of Electrical and Electronics Engineers, Inc., New York, 1978, pages 278-286.

5. Cutler, N.R. and Heiser, J.F. The Tricyclic Antidepressants. Journal of the American Medical Association, vol. 240, pages 2264-2266, 1978. (Editorial comment on page 2287. Error retraction: "Corrections", vol.241, page 566.

6. Taintor, Z.C., Laska, E.M., Siegel, C., Hedlund, J., Williams, T. and Heiser, J.F. Panel: Automated Data Systems in Psychiatry. Continuing Medical Education Syllabus and Scientific Proceedings in Summary Form: The 131st Annual Meeting of the American Psychiatric Association, Washington, D.C., page 293, 1978.

7. Heiser, J.F., Brooks, R.E. A Computerized Psychopharmacology Advisor. Proceedings of the Fourth Annual AIM Workshop, Washington, D.C., page 68, 1978.

8. Heiser, J.F. and Cutler, N.R. Tricyclic Antidepressants. Letters to the Editor, Journal of the American Medical Association, vol. 242, pages 511-513,1979.

9. Heiser, J.F., Colby, K.M., Faught, W.S., and Parkison, R.C. Can Psychiatrists Distinguish a Computer Simulation of Paranoia from the Real Thing? Journal of Psychiatric Research, volume 15, pages 149-163, 1979.

10. Brooks, R.E. and Heiser, J.F. Decision-making Approaches in Medicine. Letter to the Editor, American Journal of Psychiatry vol. 136, pages 857-858, 1979, regarding the article Roberts, B.: A Look at Psychiatric Decision Making, American Journal of Psychiatry, volume 135, pages 1384-1387, 1978.

11. Mittel, N.S., Cole, J., Heiser, J.F., LaBrie, R., Taintor, Z. Panel: Computerized Psychotropic Drug Monitoring. Continuing Medical Education Syllabus and Scientific Proceedings in Summary Form, The 132nd Annual Meeting of the American Psychiatric Association, Washington, D.C., pages 317-318, 1979.

12. Brooks, R.E. and Heiser, J.F. Controlling Question Asking in a Medical Expert System. Proceedings of of the 6th International Joint Conference on Artificial Intelligence, Tokyo, Japan 1: pages 102-104, 1979.

13. Brooks, R.E. and Heiser, J.F. Some Experiences with Transferring the MYCIN System to a New Domain. IEEE Trans on Pattern analysis, Machine Intel (Special Issue Bio. Pattern Analysis), in press, 1980.

14. Brooks,R.E. and Heiser, J.F. Transferability of a Rule-based Control Structure to a New Knowledge Domain. In Dunn, R.A. Proceedings: The Third Annual Symposium on Computer Applications in Medical Care, IEEE Computer Society, pages 56-63, 1979.

E. Funding Support Status.

1. The Principal Investigator, Co-Principal Investigator, and Pharmacist are full-time employees of the University of Texas Medical Branch at Galveston, and have participated in this research as part of their assigned duties or in their spare time. Specifically, the Principal Investigator and Co-principal Investigator are assigned to work half time on this project.

2. Additional support in the form of Office and Laboratory Space, Clerical Assistance, Peripheral Data Processing Equipment, Supplies and Expenses for Traveling to Professional Meetings has also been provided by the University of Texas Medical Branch.

3. From October 1977 through June 1978, Mr. Holthus was employed half-time by the University of California, Irvine as a Research Technician in the Department of Psychiatry and Human Behavior. Mr. Holthus was assigned to work on this project and was paid $3.67 per hour.

4. A modest amount of additional support was obtained from: Title: A Computerized Psychopharmacology Advisor Principal Investigator: Jon F. Heiser, M.D. Funding Agency: Anne R. Issler Endowment Fund Department of Psychiatry and Human Behavior University of California, Irvine Total Award: $552.50 Date: January-June 1978

5. A grant application submitted to the National Institute of Mental Health in November 1977 was rejected. An application for a Career Development Award for the Principal Investigator submitted to the Veterans Administration in January 1978, was funded but rejected by the Principal Investigator in favor of accepting his current position with the University of Texas Medical Branch. Copies of these grant proposals are available upon request.

6. Two grants were submitted to the National Institute of Mental Health in November 1978. The titles are "A Computerized Psychopharmacology Advisor" and "Rule-based Tricyclic Antidepressant Knowledge System". A grant entitled "Transferability of a Rule-based Control Structure to a New Knowledge Domain" was submitted to the National Science Foundation in May 1979. All three of these grants were rejected. Copies are available upon request.

7. The Director of Professional Services, E.R. Squibb and Sons Pharmaceutical Company, has offered to support Professional Collaboration through Squibb's panel of distinguished consultants.

II. Interactions with the SUMEX-AIM Resource.

A. Collaborations and Medical Use of Programs via SUMEX.

1. The MYCIN group has collaborated with our group since work on the Psychopharmacology Advisor began. The MYCIN group supplies invaluable software support to the EMYCIN program. Our group has participated in writing documentation of the EMYCIN software which presumably is useful to all EMYCIN users.

B. Sharing and Interactions with Other SUMEX-AIM Projects.

1. Collaboration with Kenneth Mark Colby, M.D. and members of the Higher Mental Functions Project, begun two years ago, has continued in the form of performing and experiment and publishing a paper reporting a "Turing Test" which was performed on-line on SUMEX, with the psychiatrist-judges located at the University of California, Irvine, the patient-person at the University of California, Los Angeles (UCLA) and PARRY at SUMEX. Copies of this paper (see I.D.6. above: Heiser et al. Can Psychiatrists Distinguish a Computer Simulation of Paranoia from the Real Thing?) are available upon request. In addition, demonstrations of the PARRY and

DOCTOR programs have been given on-line, using SUMEX, to various groups of mental health professionals, computer scientists and other qualified and interested individuals.

C. Critique of Resource Management.

We continue to find the SUMEX resource a hospitable environment. We feel that the choice of operating system and associated utilities was an unusually good one, and it has become a standard against which we judge other systems.

III. Research Plans

A. Long Range Project Goals and Plans.

1. Evaluation of the Psychopharmacology Advisor.

When the performance of the Psychopharmacology Advisor approaches an optimal level in the judgment of the Principal Investigators and the Advisory Panels, a formal evaluation will be performed. Elaborate plans have been made for three types of evaluation: as a simulation of the Principal Investigator; as a national expert; and as an actual psychopharmacology advisor. In each evaluation the system will be tested on two sets of cases: one which represents the population of patients likely to be encountered in practice, thereby measuring whether HEADMED can do well what it must do most often; and one which represents unusual or exceedingly complicated cases, thereby measuring whether the program can do well in situations where usual practices may not suffice. Details of the evaluation plans are available upon request.

In order to evaluate the EMYCIN formalism regarding both its inherent properties as a consulting algorithm and its appropriateness for the domain of clinical psychopharmacology, we are seeking the answers to five questions:

1) Is it beneficial to capture knowledge and control structure in the same formalism?

2) Are certainty factors a useful way in which to encode uncertain information?

3) Can the needed input be captured through the parameter/value system?

4) Are the rules really modular?

5) Is the backward chaining rule structure appropriate?

B. Justification and Requirements for Continued SUMEX Use.

As mentioned in the preceding section, we consider the use of the EMYCIN software as integral to our project, at least for the next two years, or until we have learned enough about the domain of clinical psychopharmacology to know how to supersede the EMYCIN formalism.

C. Our Needs and Plans for Other Computational Resources, beyond SUMEX/AIM.

Our only immediate need for other computational resources beyond SUMEX/AIM continues to be for local, high-speed printing, preferably combined with local file storage. Our current slow-speed printing is unsuitable for listings of large rule sets or of system code. The planned acquisition of a 1200 baud printing terminal may substantially reduce the problem.

Our future plans will depend greatly on the outcome of our current effort. If the EMYCIN formalism proves suitable for our domain, we may find the conversion effort sufficiently worthwhile to transport EMYCIN to our local environment. If we discover that a major redesign is needed, we will make our future computing plans in light of that design.

D. Recommendations for Future Community and Resource Development.

We support the request for additional computing power for the SUMEX resource. We also welcome the proposal to experiment with sophisticated single-user machines.

9.4.4 Computer-Aided Refinement of Medical Knowledge

Computer-aided Refinement of Medical Knowledge

Allan H. Levy, M.D.
School of Basic Medical Sciences
College of Medicine
University of Illinois
Urbana, Illinois

I. Summary of Research Program

A. Project rationale

The structure and organization of any representation for medical knowledge determines the effectiveness of that representation for human or automated problem solvers. The construction of a consulting system requires the description of the relevant medical expertise at a high level of specificity and completeness. Thus, the act of constructing an automated problem solver can actually result in a "refinement" of medical knowledge as inconsistencies are weeded out and gaps are filled in. At present, such refinement is directed toward improving the performance of the automated consulting program. Although such improvement will continue to be a major goal for refinement, this project will emphasize processes which result in a presentation of the refined medical knowledge in forms suitable for assimilation by human problem solvers.

B. Medical Relevance and Collaboration

The management of cancer patients by the administration of chemical agents is a complex and increasingly successful medical procedure. The extensive use of cancer chemotherapy protocols provides both a source of medical knowledge for inclusion in an automated problem solver and a vehicle for introducing refined knowledge back into medical practice. Our initial thrust will deal the problem of protocol construction (either for a research trial or for an individual patient). Initially, our efforts will help to produce individual protocols, but our long term refinement goals are to provide additional insights into the process of protocol construction and to patient selection. These areas of expertise will become increasingly important as the routine management of chemotherapy ceases to be the exclusive domain of the clinical oncologist and moves into the realm of internal machine.

We are working closely with two practicing oncologists:

John Schmale, Hematologist-Oncologist, Christie Clinic, Champaign, Illinois; Clinical Assistant Professor, School of Clinical Medicine, University of Illinois.

Allen Hatfield, Head, Department of Oncology, Carle Clinic Association, Urbana, Illinois; Assistant Professor, School of Clinical Medicine, University of Illinois.

C. Research in Progress

We are just beginning this research effort and are still assembling staff and support. We are making some preliminary "pencil and paper" knowledge bases in order to assess the applicability of specific refinement techniques. As our efforts increase in scope, we will call upon consultants from outside the University of Illinois. We will work with Emil Freireich, David Wirthschafter, and John Laszlo both to acquire knowledge for refinement and to distribute the refined knowledge back into the medical community.

D. Relevant Publications

Although this specific research project has not yet produced any publications, the three articles listed below underlie much of the proposed research:

Baskin, A. B., "Logic Nets -- Semantic Networks and Variable-valued Logic," to be published in Int. J. of Man-machine Studies, 12, 1980.

Michalski, R. S., "Knowledge acquisition by encoding expert rules versus computer induction from examples: a case study involving soybean pathology," to be published in Int. J. of Man-machine Studies, 12, 1980.

Michie, Donald, "A theory of advice," in Machine Intelligence 8, (Elcock and Michie, eds.) pp. 131-168, Ellis Horwood, New York: John Wiley, 1977.

E. Funding support

We are in the process of attempting to secure both internal and external support for this research. A research proposal to the National Library of Medicine is pending a final review by the Board of Regents (as of 5/6/80). Internal University of Illinois support and the proposals listed below will provide major support for this research project:

> Computer-aided Refinement of Medical Knowledge -- a grant application under consideration by the National Library of Medicine, Allan H. Levy, principal investigator, University of Illinois, proposed budget: 8/1/80 - 7/31/85, $1,232,053 in direct costs.

We expect the participation of several of our physician in computer science trainees in this research project. These are physicians who are enrolled in the Master of Science in Computer Science program at the University of Illinois. Physician trainee support is provided by:

> Physician Computer Science Training -- (1 T15 LM 07011) National Library of Medicine, National Institutes of Health, 7/1/76-6/30/81, $804,262 in direct costs, Allan H. Levy, Principal Investigator.

II. Interactions with SUMEX-AIM Resource

Because we are still assembling staff for this research project, our interactions with the SUMEX facility and its staff have been minimal. We have begun to use the system on a pilot basis and we expect to be using the system extensively by the end of August, 1980. We anticipate extensive consultation with the staff at SUMEX as we begin to use the Resource. We have begun to establish collaboration with Dr. Shortliffe. Our refinement techniques are intended to complement the problem solving techniques being developed by the ONCOCIN project. We view the possibility of collaboration with members of the SUMEX-AIM community as an important part of the Resource and a principal reason for our joining it.

We expressly intend to develop tools for knowledge refinement and to make them available to the SUMEX-AIM community. The existence of a number of knowledge-based systems on SUMEX reduces the logistical problems which would otherwise result from an attempt to share the tools we develop.

III. Research Plans

A. Project goals and plans

The application of the computer as an aid in medical problem solving is fundamentally dependent on the formalization of the knowledge of medical experts. Currently, this formalization is a tedious process requiring a collaboration between a medical expert and a knowledge engineer with little help from the computer. This research project addresses the problem of using computers to improve and simplify the formulation and organization of medical knowledge. Specifically, the goals of this research are:

1) to investigate the theoretical problems of knowledge refinement and to develop an experimental system of programs for computer-aided knowledge refinement. In particular, the system will provide tools for improving knowledge representations, for testing correctness, for removal of errors and omissions, for detection of inconsistencies, and for simplification and generalization of medical decision rules. Such refined knowledge will be presented in both an automated and a printed form.

2) to integrate the resulting tools into a package available to the SUMEX-AIM community.

3) to establish a knowledge base for use in the management of patients receiving cancer chemotherapy. Such a knowledge base will be valuable both as a test bed for the refinement system and for its own merit. Knowledge concerning cancer chemotherapy is rapidly increasing in complexity as new drugs and treatment schedules are evaluated.

4) to analyze the strengths and weaknesses of different knowledge refinement techniques and to test the adequacy of various knowledge representation schemes including the integration of rule-based systems and semantic network representations of structured knowledge.

The definition of precise near-term goals will have to await the final determination of levels of available internal and external support. We intend to begin the migration of knowledge representation and refinement tools which exist at the University of Illinois onto the SUMEX-AIM system during the next year. These programs will serve as a basis from which new and extended tools for knowledge refinement will evolve.

B. Access to SUMEX-AIM

The resource management and communication facilities supported by SUMEX-AIM are essential to the research in this project. Without the support for knowledge engineering already provided by SUMEX-AIM, extensive redundant local development and support efforts would be necessary. Although local computing facilities can easily handle the computational load of the proposed research, the existing tools already developed on SUMEX-AIM will allow the research to progress more rapidly.

C. Additional computing resources

As outlined in our pending research grant application, we view SUMEX-AIM as a development medium and not as a final service delivery vehicle. For this reason, we anticipate the need for a locally dedicated machine on which we can use our knowledge refinement tools in a routine production mode. We are separately seeking funds to support this local machine. We recognize that all of the techniques we propose to use will not scale down to a small computer (e.g. LSI 11/23), but many will. In addition, the effective dissemination of refined knowledge may require the use of small "personal information managers." We have begun research in this area and will continue it with our own funds. Our local effort with "personal" machines will complement that proposed as a mainstream SUMEX-AIM activity.

D. Future requirements

If our knowledge base construction and refinement efforts are successful, our requirements for bulk storage attached to SUMEX will increase dramatically over the next five years. Accordingly, although our own needs alone would not justify the proposed extensions to bulk storage, we strongly endorse the hardware bulk storage purchases outlined for the Resource. The availability of bulk stores with the performance proposed will significantly facilitate our research project in the years 1983-1986.

Programs and data bases already developed at the University of Illinois could be most effectively integrated into our project on SUMEX-AIM if a high data rate connection to a commercial carrier existed. In addition, the deployment of subsets of our knowledge base on cancer chemotherapy will require an ability to transfer large files from SUMEX-AIM. Although these operations can be performed with magnetic tapes, a high speed data link is much more desirable. We will coordinate our efforts in this area with the efforts of the SUMEX-AIM staff so that we are able to benefit from high data rate access to SUMEX when it can be provided.

9.4.5 Interactive Statistical Package Advisor

An Interactive Statistical Package Advisor

Ruven Brooks
Department of Psychiatry and Behavioral Sciences

Harvey Bunce III
Division of Biometry
Department of Preventive Medicine and Community Health

The University of Texas Medical Branch
Galveston, Texas 77550

I. Abstract

The availability of multivariate statistical analysis packages such as BMDP, SPSS, and SAS has made the use of multivariate analyses an essential part of research in many areas. For researchers in these areas, the problems these packages pose are, first, to select from the analyses available those that are appropriate to the particular experimental questions being investigated and, second, to set up the control language necessary to run the analyses. The work proposed here is an aid to accomplishing these tasks in the form of an interactive statistical package advisor computer program. The user will describe to the program the data that has been collected and the experimental questions that are to be answered, and the advisor program will respond with suggestions as to the analysis and, if desired, will create the necessary control language to run the analysis.

In the initial work on this project, it will be assumed that the user has already consulted a statistician on design of the study so that the collected data is appropriate to the questions.

The proposed architecture of the system consists of a semantic parser for interpreting user input and a rule-based reasoning system for deducing the kind of analysis and the necessary control language. The rule-based reasoning system would be the EMYCIN programming system (VanMelle. 1979). The semantic parser would be similar to one already constructed for parsing natural language case summaries into parameter values in the MYCIN system (Bonnet, 1979).

This proposal is currently under consideration by both the National Science Foundation and the National Library of Medicine.

II. On Doing A.I.M. Research Outside of a Major A.I. Center

A major requirement of Artificial Intelligence in Medicine is to have sufficient access to both medical and artificial intelligence expertise and facilities so that the product of such research is relevant to both fields.

Rarely will one site combine access to both kinds of knowledge. My own situation is an illustration of this problem. I am located in the nation's fifth largest medical center, within a five minute walk of over 650 medical science specialists, so that I have excellent access to medical expertise. On the other hand, the nearest computing facility which makes some attempt to support A.I. systems and languages and which sells time to outsiders is located in Austin, Texas, a mere (Texas-style) 300 miles away. (Because of anomalies in intrastate versus interstate phone rates, it is actually cheaper in the evenings to call San Francisco than Austin.) While there are machines, such as the Cyber 173 in Houston, which could potentially support A.I. research, the responsibility for installing and maintaining A.I. languages would be largely mine. Similarly, the costs of programming support make purchase of my own hardware prohibitively expensive. If A.I.M. research is to be conducted in medical centers other than those fortuitously located near a major computer science department, then continued support of national centers such as SUMEX will be needed.

9.4.6 Conceptual Structures for Medical Diagnosis [Rutgers-AIM]

Conceptual Structures for Medical Knowledge Representation

Principal Investigator: Prof. B. Chandrasekaran, Ph.D.
Department of Computer and Information Science
The Ohio State University
2036 Neil Avenue
Columbus, Ohio 43210

I. SUMMARY OF THE RESEARCH PROGRAM

I.A. Project Rationale

ABSTRACT

This project investigates a new approach to medical knowledge representation in the computer so that efficient and effective diagnosis, consultation, and information storage and retrieval systems can be designed. The approach is based on the notion that knowledge representation should be organized around the deep conceptual structure of a field of medicine. We study the principles governing the organization of the conceptual structure, which can be viewed as a way of organizing the conceptual specialists working together as a community of experts. We are building new diagnosis, consultation, patient data assistant and Radiology Specialist systems based on this approach. We also plan (at a future date) to examine the feasibility of using diagnosis systems as an experimental tool to evaluate the value of different symptoms, test data and theories in concluding a successful diagnosis. This is deemed to be a novel and potentially far-reaching application of biomedical computer systems.

RATIONALE DESCRIPTION

The objective of this research is to develop principles of knowledge representation in the computer to enable efficient and natural interaction with the system to store and retrieve medical knowledge. There should be no conceptual barrier between the human and the machine. Our main thesis is that this knowledge organization must have a deep structure that corresponds to the conceptual structure of the particular field of knowledge, and that there are definite principles that determine the form and content of these conceptual structures. We believe that whatever the purpose of knowledge representation--diagnosis, data base organization, information storage and retrieval, reading and storing of imaging data, etc.--the conceptual structure organizes the knowledge in such a way that effective, efficient access to knowledge fragments can be achieved, The efficiency of this structure is a major reason for the power of problem-solving that experts demonstrate.

The heart of the project is the development of a medical diagnosis system for Cholestasis. Since we feel that a good test of any knowledge representation scheme is in diagnosis and consultation, the near term goal is to demonstrate the power of the ideas by designing diagnosis and consultation systems for nontrivial subdomains of medicine. As part of this research we shall also develop principles of patient data organization, and a structure to represent imagery data at various levels of descriptive abstractions. The long term goal of the program is to develop and apply the conceptual structures for a variety of important information processing tasks in medicine: fact-finding, diagnosis, consultation, data base organization, interpretive reporting, etc.

There is another interesting, and potentially far reaching, application that should be mentioned here. In diagnostic medicine, there is often a question about the value of certain symptoms of lab data in concluding a successful diagnosis. One way to test this would be to assign the same case with and without the particular set of data to the same (or similarly trained) clinician and examine the result of the diagnosis. This should be repeated with a large number of different cases for statistically reliable conclusions. Obviously this is hard to do in view of the fact that physician time is a very valuable commodity, and further it is hard to eliminate effects of memory of the same case with a different set of data. We believe that computer-based diagnosis systems are promising technological tools in this problem area. The effects of different data can be ascertained by subjecting the computer-based diagnosis system to various inputs. This notion can be carried even further. Suppose there are contending theories or conceptualizations of a field of medicine (as in cholestasis). The problem-solving efficiency of the contending theories can be evaluated by embodying each of them in the conceptual structure representation, designing the diagnosis system around it and testing each with data in real cases.

If properly developed this could be an important application area for biomedical information systems. Knowledge representation research clearly plays a crucial role in the development of such systems.

I.B. Medical Relevance and Collaboration

The medical relevance is indicated by the following aspects of the project:

(1) The development of MDX, a medical diagnosis system in the domain of cholestasis.

(2) The development of PATREC, a patient data base system.

(3) The development of RADEX, a computer-based radiology consultant.

(4) Plans for the development of a system to evaluate the effectiveness of various types of information in arriving at a diagnosis.

Our medical collaboration takes place through the following channels.

(1) Jack Smith, M.D., is a resident in pathology at Ohio State University Hospitals. He is also working on his Ph.D. in Artificial Intelligence in Medicine, and the research on this project will constitute his dissertation. He is our most active medical collaborator.

(2) Douglas Levin, M.D., a prominent Columbus hepatologist, and a clinical faculty member at Ohio State University College of Medicine, is our source of medical expertise in cholestasis.

(3) Carl Speicher, M.D., Director of Clinical Laboratories at OSU Hospitals, is our collaborator in the Patient Database and the associated interpretive reporting work.

(4) Joseph Schultz, M.D., a radiologist at Riverside Hospital has been our collaborator in the imaging interpretation and representation aspects of our project.

I.C. Highlights of research progress

ACCOMPLISHMENTS

The design and implementation of a working medical diagnosis program, called MDX, which was started in the last quarter of 1978, was completed in this past year. MDX, which is still limited to a small domain of medicine, has been successfully tested on a set of cases from medical journals, local university hospital and private practice.

The MDX system contains two major sub-systems, in addition to the diagnosis component, which are important in their own right. A patient data management system, based on a conceptual representation of medical data entities, has been implemented as part of the MDX system. It has two purposes. First, it manages all patient data for the diagnostic system in MDX and answers questions about the patient data. Second, this patient data management system is being used as a vehicle for research into design of AI-based data models and development of flexible and easy-to-use information management systems.

The second major sub-system in MDX is a radiology consultant, called RADEX. It provides consultation on the radiological data obtained from different xrays and cholangiograms. It also maintains an anatomical and physiological model of the patient.

RESEARCH IN PROGRESS

(i) Extending the domain of MDX to include all intra-hepatic diseases (including drug-related disorders). The goal for the next two years is to build conceptual experts for all liver diseases.

(ii) The data base model is being expanded to represent temporal information. This would allow temporal questions to be answered and enable causal inferences to be made.

(iii) A separate sub-system for preparing interpretive reports of clinical lab data is under development. The interpretive reports would present a highlighted summary of major clinical data and provide diagnostic suggestions.

I.D. List of Relevant Publications

(1) B. Chandrasekaran, F. Gomez, S. Mittal and J. Smith, "An approach to medical diagnosis based on conceptual structures," Proc. International Joint Conference on AI, Tokyo, Japan, Aug. 1979.

(2) S. Mittal, B. Chandrasekaran and J. Smith, "Overview of MDX - a medical diagnosis system," Proc. III Annual Symposium on Computer Applications in Medical Care, Washington, D.C., October 1979.

(3) B. Chandrasekaran, S. Mittal and J. Smith, "RADEX - towards a computer-based radiology consultant," invited paper to appear in Pattern Recognition in Practice, Gelsema and Kanal, eds, North Holland, 1980 (exp).

(4) F. Gomez and B Chandrasekaran, "Knowledge organization and distribution for medical diagnosis," to appear in IEEE Trans SM&C.

(5) S. Mittal and B. Chandrasekaran, "Conceptual representation of patient databases," to appear in J. of Medical Systems.

(6) S. Mittal and B. Chandrasekaran, "Temporal Organization of events in a medical data base," to appear in Proc. Am. Conf. Cybernetics & Society, Boston, 1980.

I.E. Funding Support

1. Graduate Student Support: Jack Smith, M.D., is supported by NIH/NLM Biomedical Computer and Information Science Training Grant to the Ohio State University, Grant no. LM07023-02, total direct costs for the entire year of 79-80 is 87,959 (direct), but only a portion goes to support AIM Training for graduate students. The renewal proposal for this Training Grant has been approved by NLM, and we are awaiting funding.

2. NLM Research Support. A research proposal on "Conceptual Structures for Medical Knowledge Representation," B. Chandrasekaran, P.I., Application no. 1 R01 LM03500-01, submitted to the NLM Computers in Medicine Program has been approved by the Review Committee in its March, 80 meeting. We are awaiting funding. The approved level of funding is 80/81: 71,370, 81/82: 75,711 & 82/83: 80,111.

II. INTERACTIONS WITH THE SUMEX AIM RESOURCE

II.A. Medical Collaboration and program dissemination via SUMEX.

Most of our actual research is being conducted on the RUTGERS resource. We are aware that some researchers at Rutgers and elsewhere have studied the program in detail and exercised it. The ideas behind the program have thus been comprehended more concretely. In particular, several researchers at RUTGERS and our group have sat down at extended sessions in the TALK mode and have run our program, analyzed and critiqued it. Such an effort would be impossible without these resources.

II.B. Sharing and Interaction with other SUMEX-AIM Projects.

(1) Prof. Chandrasekaran participated in the Vermont NIH-AIM Workshop last year. This participation has been a major intellectual boost to our efforts.

(2) Prof. Chandrasekaran, Dr. Jack Smith and Sanjay Mittal will be attending the forthcoming NIH-AIM Workshop at Stanford in August 1980. They will be presenting a demonstration of the MDX system at the workshop.

(3) An important benefit of our use of the Rutgers LCSR system has been the availability of software like the Rutgers Lisp system, screen text editors, documentation preparation programs etc. Some of the programming language features we need for research into building experts-based systems will be provided in the new Lisp system under development at Rutgers. Our limited resources would have made it difficult to make such extensions to lisp.

(4) The mail facilities of SUMEX and RUTGERS resources are invaluable to us in exchanging ideas, and keeping track of AIM activities. There is simply no cost-effective substitute for the benefits of such communication for our research.

(5) There is a sense of cohesiveness and common purpose fostered by this resource. We have been able to make a number of professional contacts through the resource and engage in intellectual dialogs about problems of common interest in medical knowledge representation.

II.C. Critique of Resource Management.

The Rutgers facility which is our major computer is excellently managed and provides us adequate service. Our impression of the Stanford machine is similar, even though it is not based on as extensive a use.

The major weakness and source of anxiety for us is the unreliability and slowness of the TYMNET nodes. The unreliability of TYMNET has often created problems in continuous use of the SUMEX facilities. The very slow speed (currently only 300 baud) has been a major impediment in our research effort from the point of view of acceptable turnaround time. Anyone who has tried to use run editors like EMACS; run programs with lots of output; or print out large files, at 300 baud, would agree that it is a maddeningly slow process.

III. RESEARCH PLANS (8/80 - 7/86)

III.A. Project goals and plans.

A.i. Near-term (8/80 - 7/81)

(1) Expansion of the MDX system to larger domains of medicine, first to the whole of cholestasis, then to the domain of all liver diseases.

(2) Creation of facilities to understand and organize the temporal aspects of patient data and case histories.

(3) Implementation of MDX-II, embodying a more advanced problem-solving strategy, which will result in a more coherent unified diagnosis. This will be done for the domain of liver diseases.

A.ii. Long-term (8/81-7/86)

(1) Increasing the capabilities of MDX for consultation, including a) explanation of diagnosis, and b) ordering tests.

(2) Evaluation of MDX as a tool in a clinical environment.

(3) Extension of the RADEX system to store and retrieve imaging information over a much larger domain.

(4) Extension of conceptual patient data bases, and interpretive reporting facilities.

(5) "Learning" of new diagnostic information by the system. We would like to investigate how MDX can acquire new knowledge and skills, either from episodic information, or productive problem solving using underlying medical and commonsense knowledge structures.

(6) Use of MDX to evaluate the usefulness of particular symptoms, manifestations and data in the diagnostic process. This can currently be done only by a large expenditure of expert clinician time.

III.B. Justification and Requirements for Continued SUMEX Use.

I will talk more broadly of SUMEX/RUTGERS use.

a. Computing Facility: Our research will simply be crippled without it. AIM work of the kind that we do requires very good LISP support and a large machine. Our Department will shortly have a DEC20/20 which is altogether too small for the kind of work that we need to get done. Rutgers DEC20/50 has sufficient power for our immediate needs, even though there are problems about disc space even there. Of course both SUMEX and RUTGERS resources will probably need considerable enlargement if the number of users and the sizes of the projects grows, as they almost surely will.

b. Training Facility: We not only need access to these resources for research, but we need it for purposes of training graduate students in AIM issues. For instance most of the graduate students in our AIM activity as well as those associated with the NLM Biomedical Computing Training Grant can now get hands on experience with other AIM programs through our access to these resources. Thus as an educational tool in AIM, this resource is an essential one for our Training activity also. It should be emphasized that there are no viable alternatives to networking in this regard. AIM programs are large, experimental and need substantial supporting resources.

c. Community Building: We need to continue our interaction with the centers of AIM activity, as well as individual researchers such as Ruven Brooks at the University of Texas Medical Center (SUMEX-AIM user). We need to feel part of this common AIM purpose, otherwise research groups like ours which are very active but small would be rather isolated.

III.C. Needs and plans for other computing resources, beyond SUMEX-AIM

In the next five years, we anticipate a fairly substantial demand for computing resources for our research projects. Our anticipated requirements would be as follows:

1. Memory size - Within the next 2-3 years we would need a computer system which can provide program memory size of the order of 2M bytes or more. Unless DEC makes extensions to the Dec-10/Dec-20 architecture, we would have to explore alternate systems, such as the VAX or IBM 370.

2. Computer usage - Our usage of the computer system at the Rutgers LCSR has been steadily growing since we started using their system eight months back. However, it is not clear how much more we would be able to load their system in the next few years, without degrading the system performance substantially. Our own department at Ohio State is currently installing a new Dec 2020 system, which would help in providing additional computer resource for our research, though only in a limited way. We would be interested in exploring ways to upgrade this Dec2020 system with SUMEX support.

3. Disk storage - Our experience of last year has shown that development of knowledge-based systems in medicine require large amounts of secondary disk storage, for keeping different versions of source codes, compiled code, documentation etc. Currently we are using about 1500 pages (1 page= 2.5K bytes) on the Rutgers system. We expect this to increase by at least 100 percent in the next two years.

III.D. Recommendations for future community and resource development

Our experience in the past eight months as a member of the SUMEX community has been quite rewarding, both in terms of availability of computer and software resources, as well as contact with other researchers.

Our biggest source of frustration has been the slow communication access to the SUMEX computing systems. The Tymnet node at Columbus provides only a 300 baud line, which has proved to be a serious hindrance in our effective use of the Rutgers system. We would be very interested in exploring ways to upgrade this to at least 1200 baud, and if possible higher. Any help from SUMEX, both technical and financial would be very useful. We recommend that SUMEX explore ways to setup a nation wide network access to SUMEX computing systems, which provide the following:

a. High speed access to the SUMEX computing nodes.

b. File transfer capability between the SUMEX facilities and the local facilities of other members of the community.

The new technology in computer communications may enable the public carriers to be used for fashioning such a network.

We feel that the dissemination of ideas and software among different members of the AIM community may be enhanced by the introduction of a quarterly SUMEX newsletter. Such a newsletter could be used to announce new programs and utilities and provide a forum for discussion of work-in-progress and other ideas not yet ready for formal publication.

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 6 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.

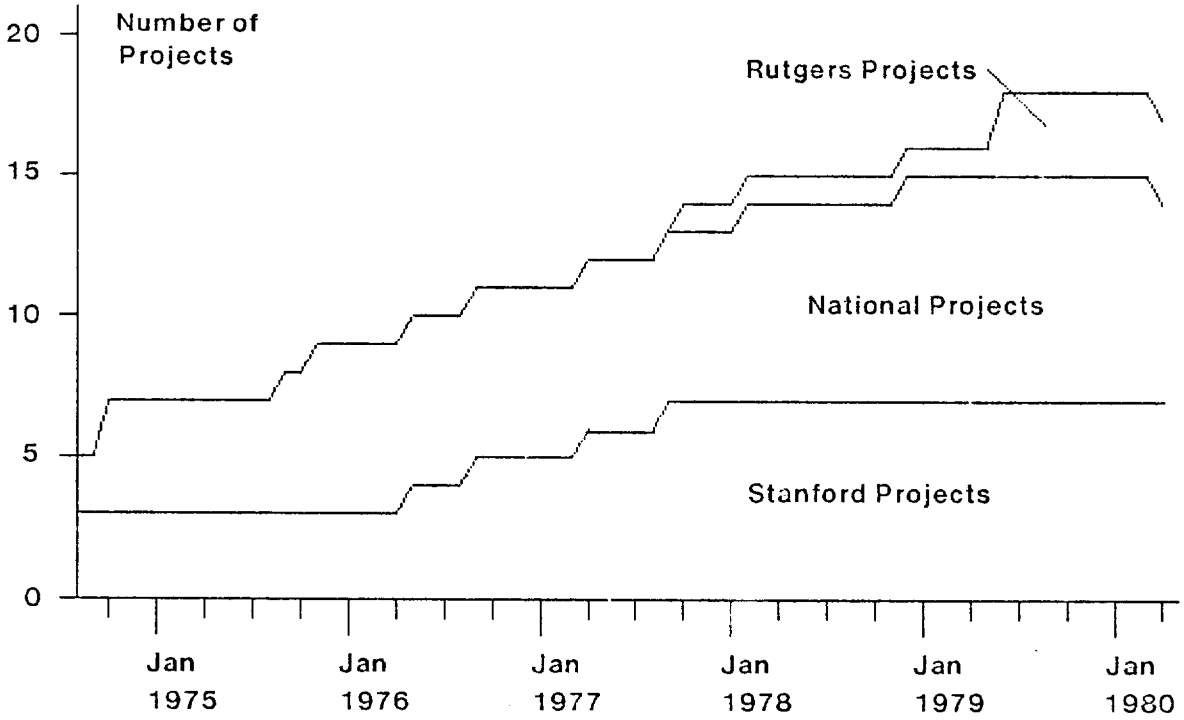

Figure 6. SUMEX-AIM Growth by Community

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:	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.

Rutgers AIM Project: A CLINICAL DECISION-MAKING MODEL INCORPORATING GOAL-SEEKING AND FOCUSING STRATEGIES

Principal Investigator: Robert A. Greenes, M.D.
Department of Radiology,
and Information Science
Peter Bent Brigham Hospital
Harvard Medical School
Boston, Massachusetts 02115
(617) 732-6281 (GREENES@RUTGERS)

Clinical decision-making in the model is viewed as a process which involves: 1) formulation of a set of diagnostic hypotheses and estimation of their likelihoods; 2) posing of patient management goals appropriate to the diagnostic hypotheses; and 3) selection of tests to perform, based on the relationship between the likelihoods of diagnosis and the certainty levels, or threshold probabilities, required for adoption of corresponding management goals. Focusing on particular diagnoses is accomplished by requiring minimum certainty levels for consideration of the diagnostic hypotheses. Probabilities are revised by Bayesian methods. The choice among suitable tests involves a heuristic scoring process. The system is being applied initially to the evaluation of upper abdominal pain.

SOFTWARE AVAILABLE ON SUMEX

Programs are in a developmental stage on the RUTGERS-AIM system and not yet available for use.

REFERENCES

Greenes, R.A.: Investigations in clinical decision-making. NLM program project grant application 1 P01 LM03401-01, 1979, pp. B27-31.

Greenes, R.A.: A goal-directed method for investigation of thresholds for medical action. Proc. Third Annual Symposium Computer Applications in Medical Care, Washington, D.C., 1979, pp. 47-51.

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 estimates of the text's recall and readability.

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.

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: HIGHER MENTAL FUNCTIONS (HMF)

Principal Investigator: Kenneth M. Colby, M.D.
Departments of Psychiatry
and Computer Science
Neuropsychiatric Institute
University of California at Los Angeles (UCLA)
Los Angeles, California 90024
(213) 825-4626 (COLBY@SUMEX-AIM)

Contact: William FAUGHT@SUMEX-AIM
Roger PARKISON@SUMEX-AIM

The HMF Project contributes new knowledge and instruments to the fields of psychiatry and neurology using concepts and techniques of artificial intelligence. The research includes a model of paranoid behavior, a cognitive psychiatric taxonomy, and the development of intelligent speech prostheses for nonspeaking patients.

SOFTWARE AVAILABLE ON SUMEX

PARRY--An interactive program which can be interviewed in unrestricted natural language and responds linguistically the way paranoid patients respond in an initial psychiatric interview.

REFERENCES

Colby, K.M.: Mind models: An overview of current work. Mathematical Biosciences 39:159-185, 1978.

Colby, K.M., Christinaz, D. and Graham, S.: A computer-driven personal, portable and intelligent speech prosthesis. Computers and Biomedical Research 11:337-343, 1978.

Colby, K.M.: Computer simulation and artificial intelligence in psychiatry. IN E.A. Serafetinedes (ed.), Methods of Biobehavioral Research. Grune and Stratton, New York, 1979.

National AIM Project: 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 INTERNIST 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 INTERNIST 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 an INTERNIST 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: 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.

EXPERT--System for designing and applying consultation models using a relatively simple language to describe the models.

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)

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). 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.

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, multi-disciplinary 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 10 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: 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:	MYCIN--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 (by courtesy) Stanford University Stanford, California 94305 (415) 497-5821 (SHORTLIFFE@SUMEX-AIM)
Subproject Directors:	MYCIN: Dr. Shortliffe, and A. Carlisle Scott EMYCIN: Dr. Buchanan, and William van Melle GUIDON: Drs. Buchanan, and William J. Clancey ONCOCIN: Dr. Shortliffe, and A. Carlisle Scott

The MYCIN Project is 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 physicians to use in the near future. There is also a heightened interest in gearing representation, knowledge acquisition, and explanation more to the way that an expert actually thinks.

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. The program continues to provide a powerful research environment for developing new approaches to the basic questions involved in knowledge engineering.

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 assist oncologists with the management of cancer chemotherapy protocols. Because the knowledge in this domain is already well-specified, the research emphasis is on human engineering and achieving 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 Conference Computer Applications in Medical Care, Silver Spring, Maryland, October, 1979, pp. 175-183.

Shortliffe, E.H., Buchanan, B.G. and Feigenbaum, E.A.: Knowledge engineering for medical decision making: A review of computer-based clinical decision aids. Proc. IEEE 67:1207-1224, 1979.

van Melle, W.: A domain-independent production rule system for consultation programs. Proc. Sixth IJCAI, Tokyo, August, 1979, pp. 923-925.

Yu, V.L., Fagan, L.M., Wraith, S.M., et al: Antimicrobial selection by a computer - A blinded evaluation by infectious disease experts. JAMA 242:1279-1282, 1979.

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

Resource Operations and Usage 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) Diurnal loading data
5) Network usage data
6) 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.

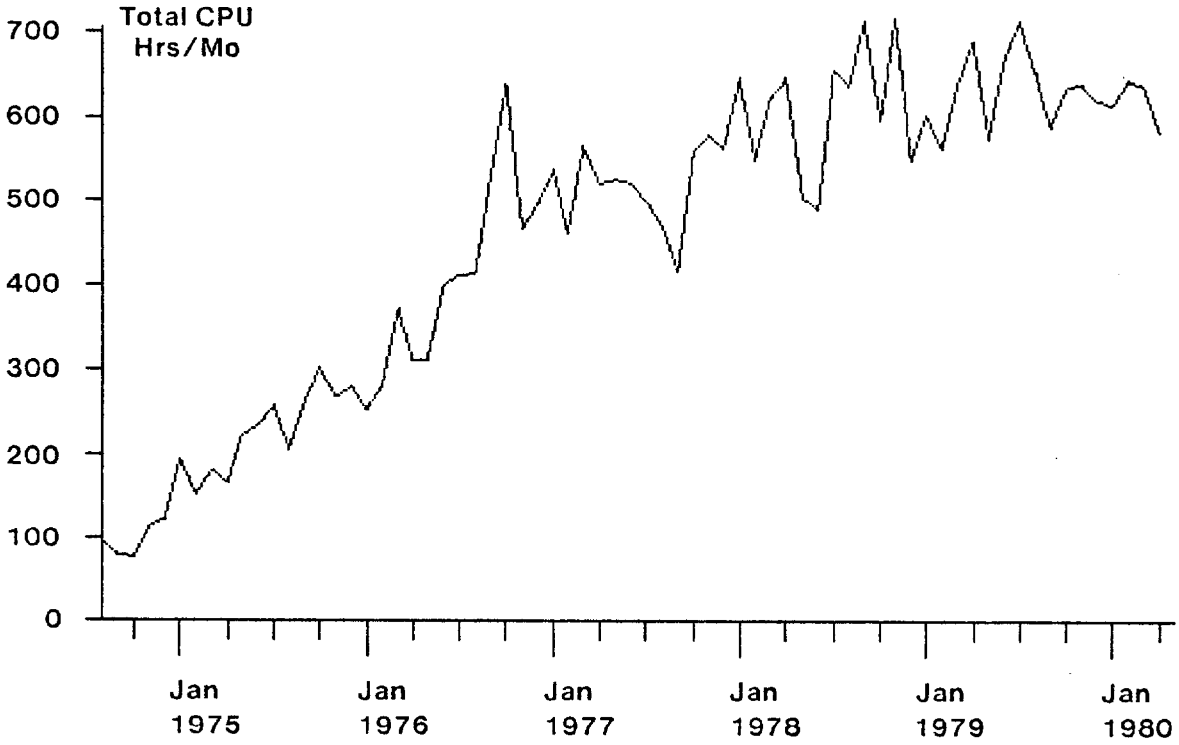

Figure 7. Total CPU Time Consumed by Month

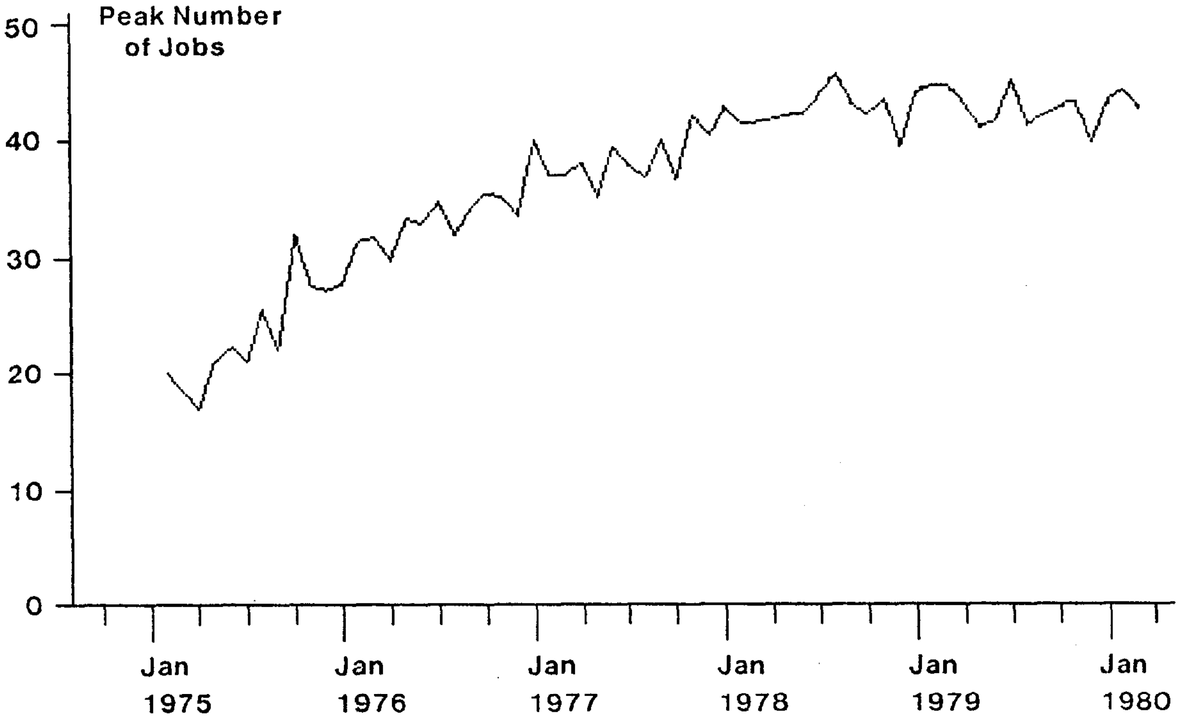

Figure 8. Peak Number of Jobs by Month

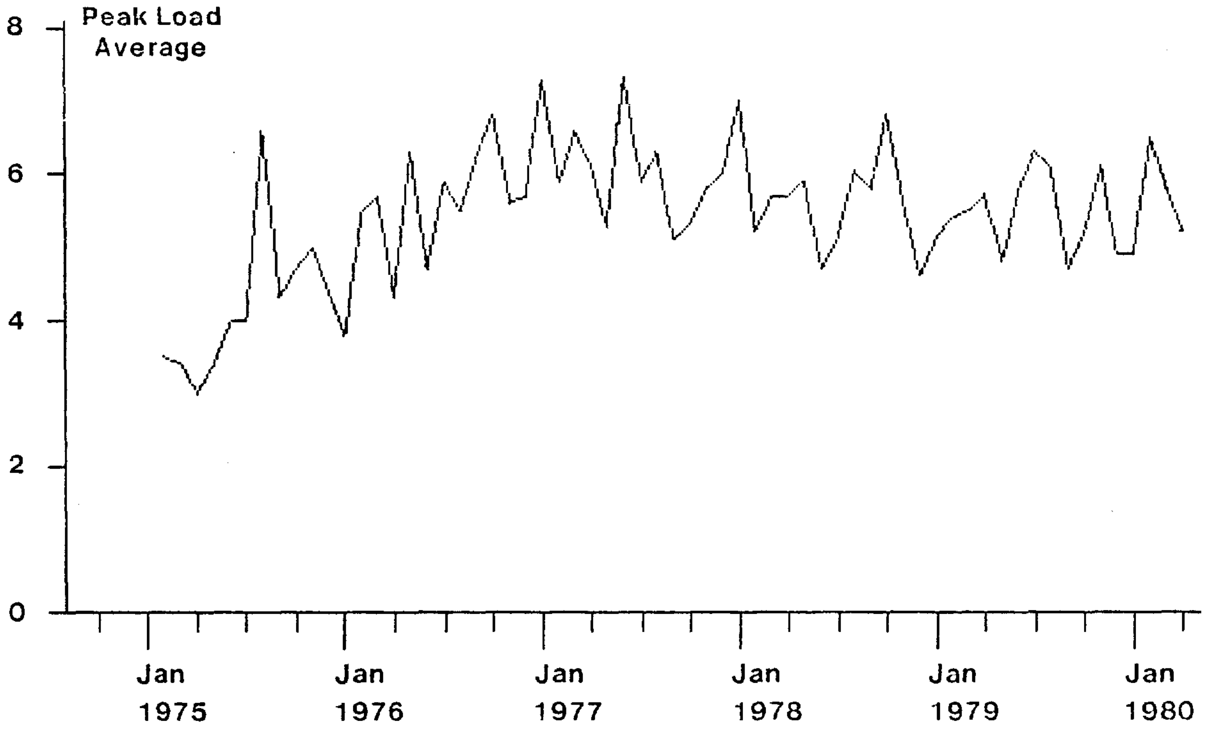

Figure 9. Peak Load Average by Month

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 10 and Figure 11. Terminal connect time is shown in Figure 12. 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.

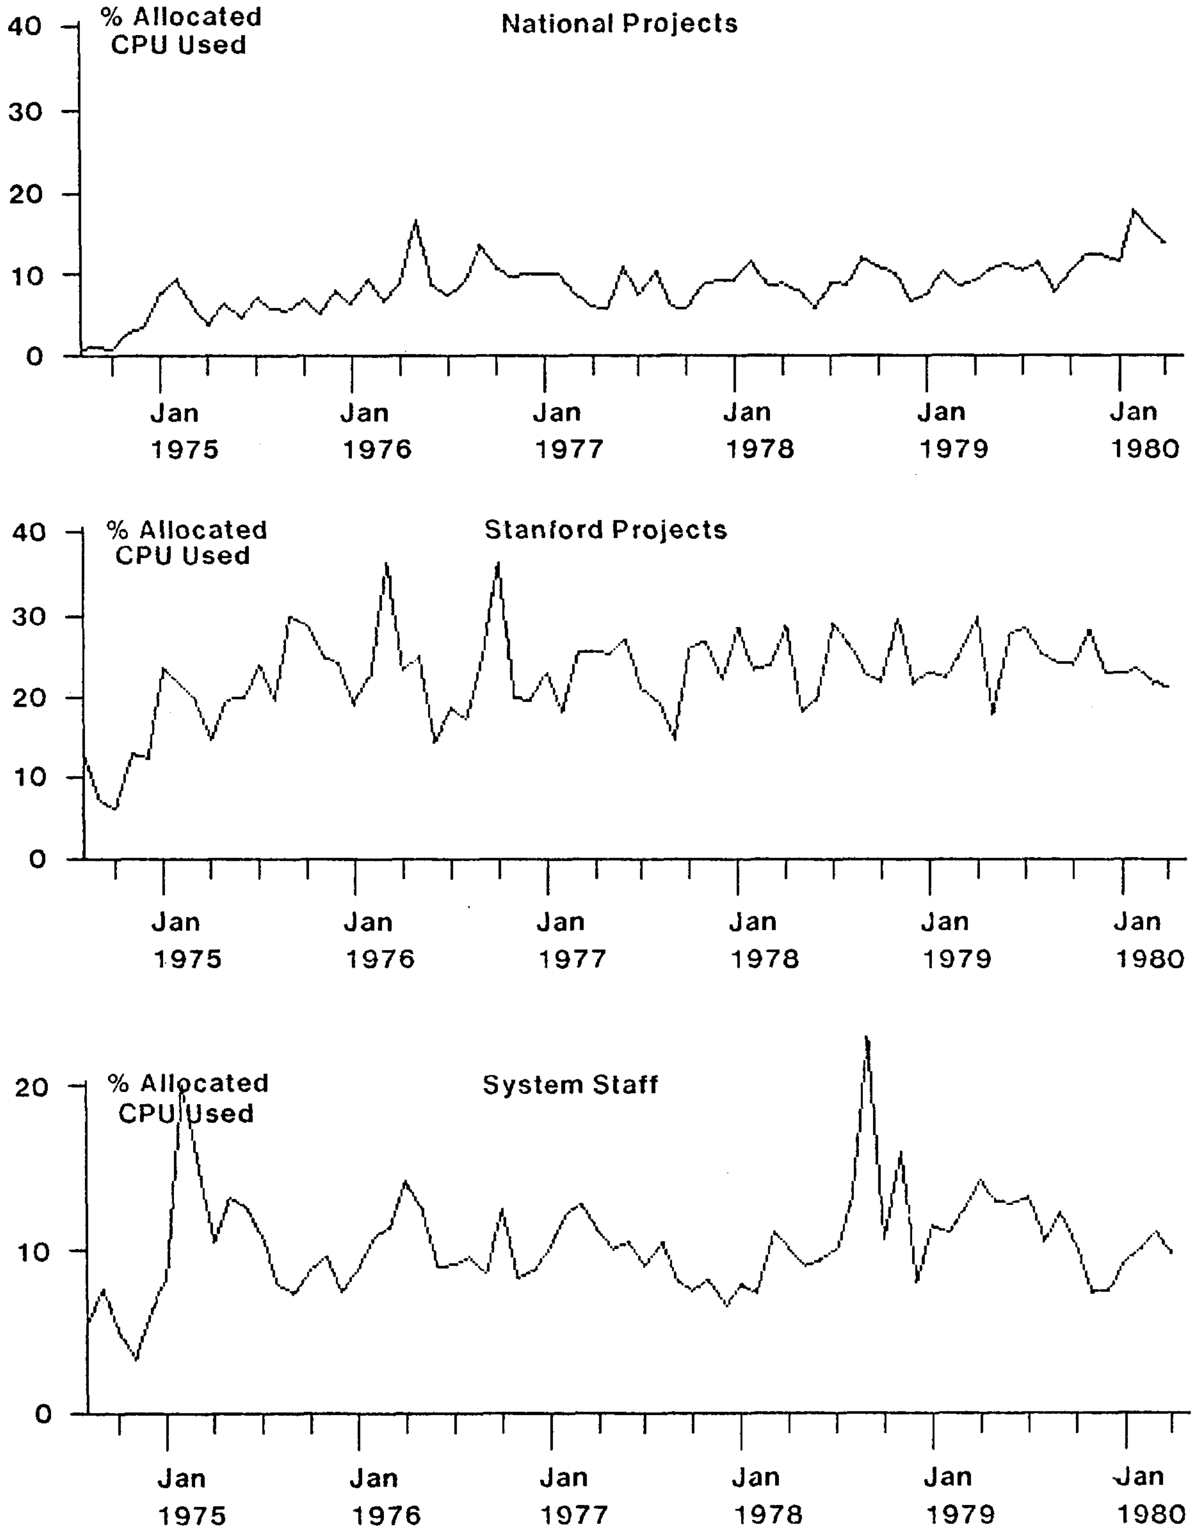

Figure 10. Monthly CPU Usage by Community

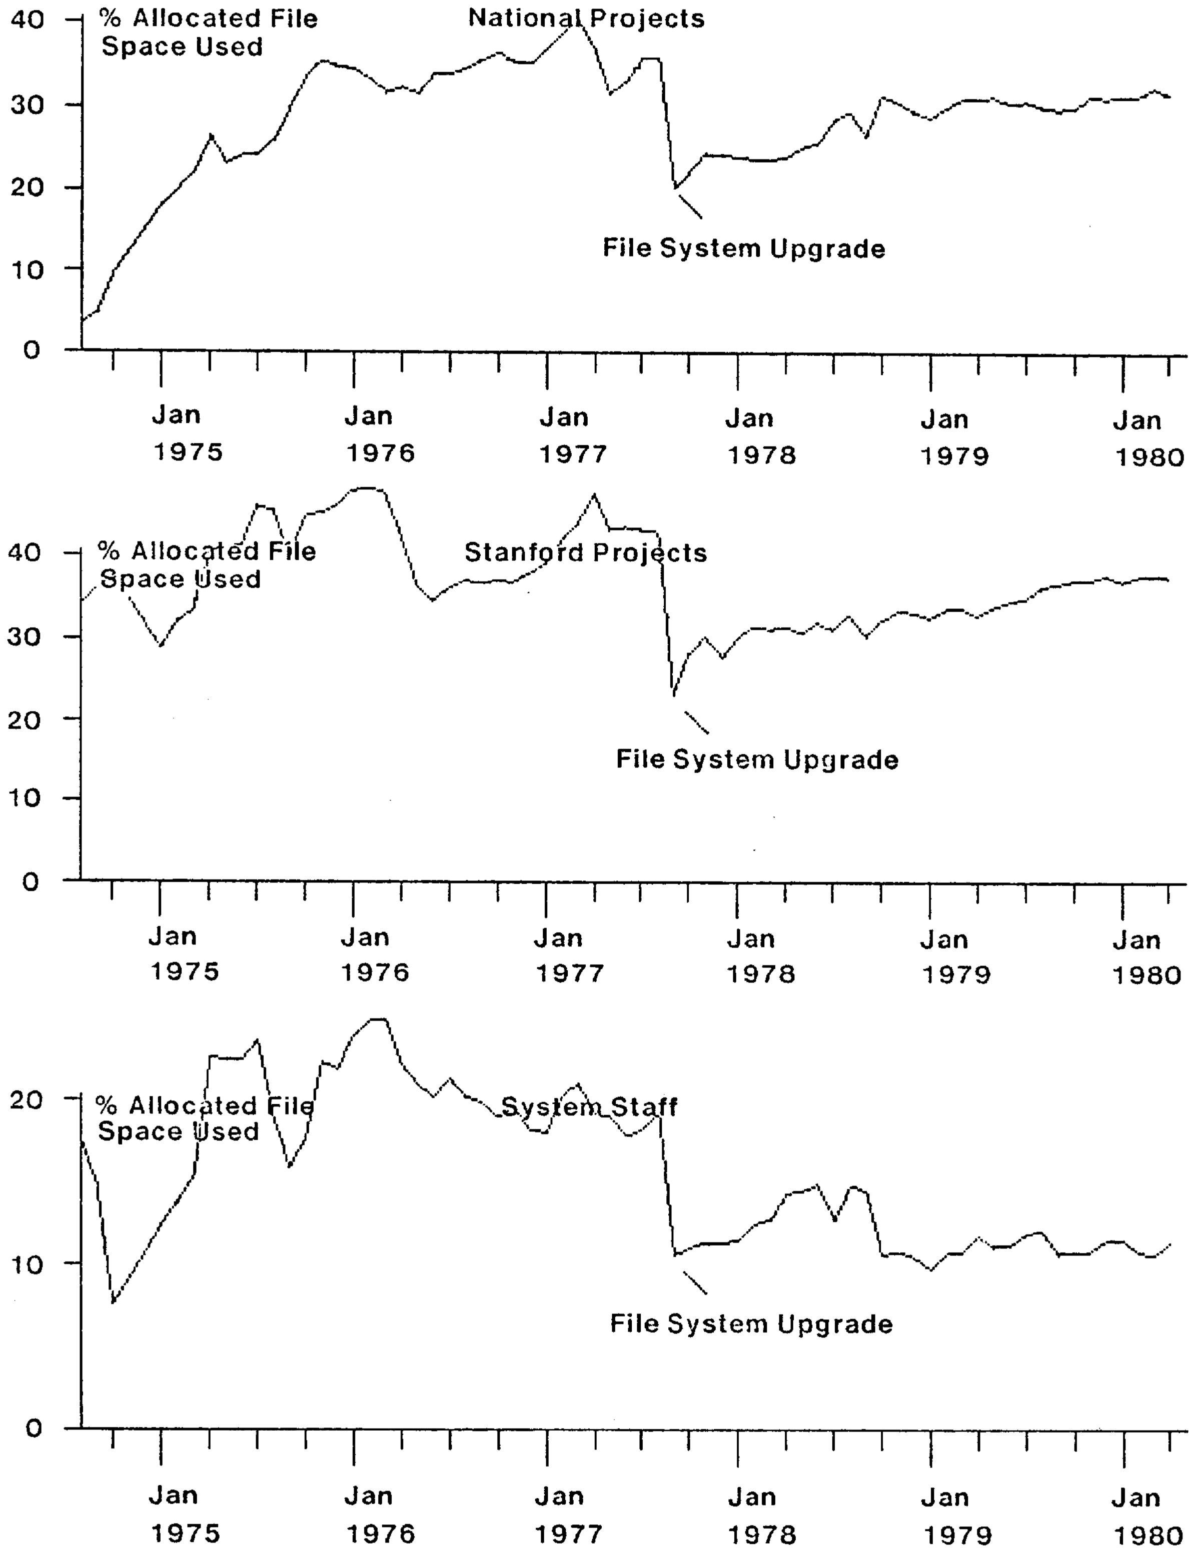

Figure 11. Monthly File Space Usage by Community

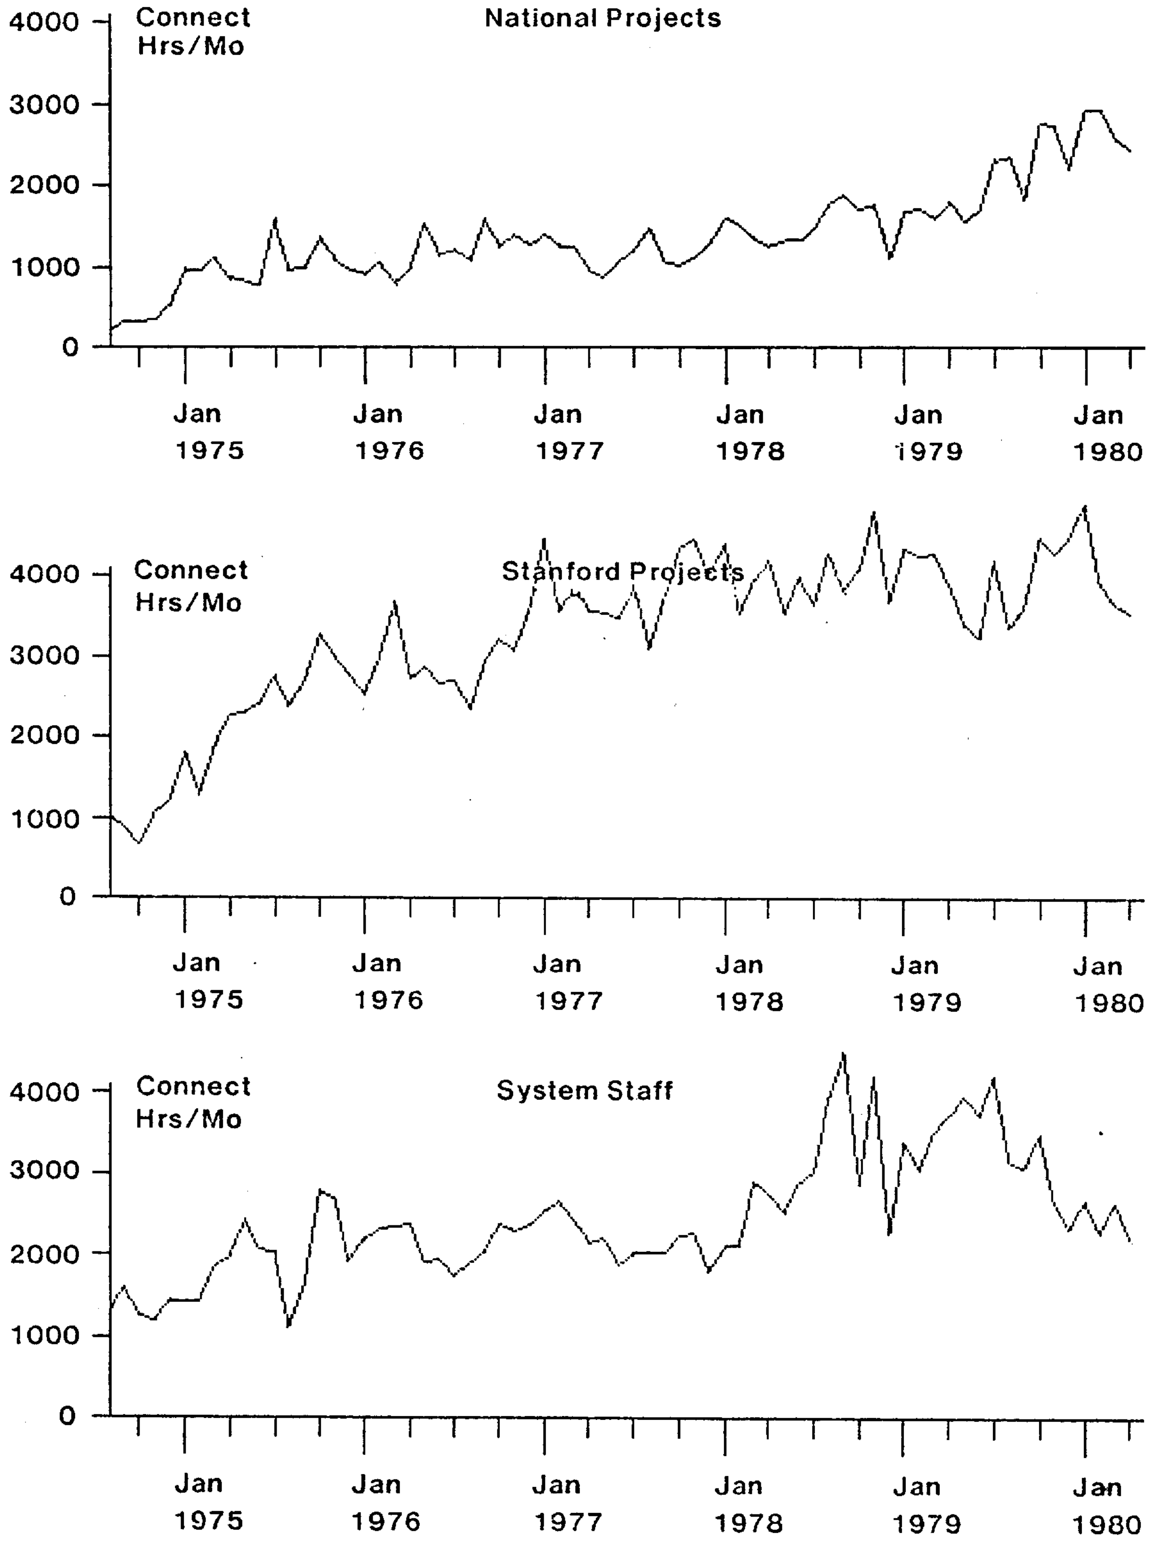

Figure 12. Monthly Terminal Connect Time by Community

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 1979 and April 1980.

Several of the projects newly admitted to the National AIM community use the Rutgers-AIM resource as their home base. These projects are listed in the tables to fully document the scope of the AIM community and are noted with the flag "[Rutgers-AIM]".

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 is charged to the national community in the "AIM USERS" category, however.

Resource Use by Individual Project - 5/79 through 4/80

National AIM Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) ACT Project "Acquisition of Cognitive Procedures" John Anderson, Ph.D. Carnegie-Mellon Univ. ONR N00014-77-C-0242 9/78-9/80 $175,000	106.50	1197.90	2634
2) SECS Project "Simulation & Evaluation of Chemical Synthesis" W. Todd Wipke, Ph.D. U. California, Santa Cruz NIH RR-01059-03S1 (3.7 yrs. 7/77-2/81) 7/80-2/81 $36,949 NIH/NCI NO1-CP-75816 (2 yrs. 1/79-12/80) 1/80-12/80 $74,394	538.31	9943.77	8389
3) Mod Human Cogn Project "Hierarchical Models of Human Cognition" Peter Polson, Ph.D. Walter Kintsch, Ph.D. University of Colorado NIE-G-78-0172 (3 yrs. 9/78-8/81) 9/79-8/80 $46,537 NIMH MH-15872-9-13 (5 yrs. 6/76-5/81) 6/79-5/80 $32,880 ONR N00014-78-C-0433 6/78-5/80 $68,315 6/80-5/81 $60,000 ONR N00014-78-C-0165 (2 yrs. 1/78-12/80) 1/80-12/80 $85,000	119.61	2696.51	712
4) Higher Mental Functions "Intelligent Speech Prosthesis" Kenneth Colby, M.D. UCLA NSF MCS-78-09900 6/78-11/80 $135,260 NSF PFR-17358 10/79-3/81 $318,368	20.65	637.47	2810

5)	INTERNIST Project "DIALOG: Computer Model of Diagnostic Logic" Jack Myers, M.D. Harry Pople, Ph.D. University of Pittsburgh NIH RR-01101-03 (3 yrs. 7/77-6/80) 7/79-6/80 $200,414	215.76	3756.99	7755
6)	PUFF/VM Project "Biomedical Knowledge Engineering in Clinical Medicine" John Osborn, M.D. Inst. Medical Sciences, San Francisco Edward Feigenbaum, Ph.D. Stanford University NIH GM-24669 9/78-8/81 $164,000 (*) Supplement pending	125.39	4669.66	3196
7)	SCP Project "Simulation of Cognitive Processes" James Greeno, Ph.D. Alan Lesgold, Ph.D. University of Pittsburgh NIE-G-80-0114 (3 yrs. 12/79-11/82) 12/79-11/80 $217,000 ONR/ARPA N00014-79-C-0215 (1.8 yrs. 1/79-9/81) 10/79-9/80 $420,000 NSF/NIE 12/78-5/81 $161,238 ONR N00014-78-C-0022 (3 yrs. 10/77-9/80) 10/79-9/80 $92,293	21.07	648.56	764
8)	*** [Rutgers-AIM] *** Rutgers Project "Computers in Biomedicine" Saul Amarel, D.Sc. NIH RR-00643 (3 yrs. 12/77-11/80) 12/79-11/80 $451,383	23.56	513.01	9204

9)	*** [Rutgers-AIM] *** Decision Models in Clinical Diagnosis Robert Greenes, M.D. Harvard University NLM LM-03401 (5 yrs. 7/79-6/84) 7/79-6/80 $235,582	.00	.00	0
10)	*** [Rutgers-AIM] *** Heuristic Decisions in Metabolic Modeling David Garfinkel, Ph.D. Univ. Pennsylvania HL-15622 (3 yrs. 12/77-11/80) 12/79-11/80 $111,051 GM-16501-11A1 (2 yrs. 4/80-3/82) 4/80-3/81 $60,598 Proposals pending	.00	.00	0
11)	AIM Pilot Projects			
	Coagulation Expert	9.67	207.67	480
	Commun. Enhancement	1.99	85.45	361
	KRL Demonstrations	.36	11.53	523
	MISL Project	2.40	115.23	1132
	Psychopharm. Advisor & Statistical Advisor	6.11	183.49	818
	Refinement of Med. Know.	.00	.00	0
	Struct. for Med. Diag. [Rutgers-AIM]	.45	10.85	0
		-----	-------	----
	AIM Pilot Totals	20.98	614.22	3320
12)	AIM Administration	16.67	661.64	4668
13)	AIM Users on Stanford Projects			
	AGE	3.26	57.94	35
	DENDRAL	140.61	1930.07	1592
	MOLGEN	20.53	293.11	106
	MYCIN	11.75	391.32	167
	Guest (all projects)	34.28	362.49	209
	Other	1.78	27.31	230
		------	-------	----
	AIM User Totals	212.21	3062.24	2342
		------	--------	-----
	Community Totals	1420.71	28401.97	45794

Stanford Community	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) AGE Project (Core) "Generalization of AI Tools" Edward Feigenbaum, Ph.D. ARPA MDA-903-80-C-0107 (**) (partial support)	341.46	3103.55	3277
2) AI Handbook Project (Core) Edward Feigenbaum, Ph.D. ARPA MDA-903-80-C-0107 (**) (partial support)	69.34	2149.53	2611
3) DENDRAL Project "Resource Related Research Computers and Chemistry" Carl Djerassi, Ph.D. NIH RR-00612-11 (3 yrs. 5/80-4/83) 5/80-4/81 $221,255	957.35	10625.34	15112
4) MOLGEN Project "Experiment Planning System for Molecular Genetics" Edward Feigenbaum, Ph.D. Laurence Kedes, M.D. NSF MCS-78-02777 12/79-11/80 $153,959 (*)	409.20	9229.85	7242
5) MYCIN Project "Computer-based Consult. in Clin. Therapeutics" Bruce Buchanan, Ph.D. Edward Shortliffe, M.D., Ph.D. NLM LM-03395 (5 yrs. 7/79-6/84) 7/79-6/80 $99,484 NSF MCS-79-03753 7/79-12/80 $146,152 ONR/ARPA N00014-79-C-0302 3/79-3/82 $396,325 (*) NLM LM-00048 (5 yrs. 7/79-6/84) 7/79-6/80 $39,285 Kaiser Fdn. 7/79-12/80 $20,000 (*)	632.87	11594.76	12809

6)	Protein Struct Modeling "Heuristic Comp. Applied to Prot. Crystallog." Edward Feigenbaum, Ph.D. NSF MCS-79-23666 12/79-11/81 $35,318	85.84	1631.07	4443
7)	RX Project Robert Blum, M.D. Gio Wiederhold, Ph.D. Pharm. Mfr. Assn. Fdn. 7/78-6/80 $32,500 NLM New Invest. 7/79-6/82 $90,000 NCHSR 4/79-3/81 $35,000 Proposal pending	21.03	817.30	1256
8)	Stanford Pilot Projects			
	Genetics Applic.	55.98	938.73	377
	Hydroid	30.34	1202.01	1037
	Ultrasonic Imaging	18.67	331.44	309
	Miscellaneous	.00	.00	8
	Stanford Pilot Totals	104.99	2472.18	1732
9)	Stanford and HPP Assoc.	211.83	6710.84	6827
	Community Totals	2833.91	48334.42	55309

	SUMEX Staff	CPU (Hours)	Connect (Hours)	File Space (Pages)
1)	Staff	900.39	27809.06	9731
2)	MAINSAIL Development	272.64	5526.13	3493
3)	Staff Associates, misc.	45.42	1850.25	3249
	Community Totals	1218.45	35185.44	16473

System Operations	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	2088.14	84421.51	75174
	=======	=========	======
Resource Totals	7561.21	196343.34	192750

* Award includes indirect costs. All other awards are reported as total direct costs only.

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

	Current Year (10/79-9/80)	Total Award (10/79-9/82)
Heuristic Programming Project	$ 496,256	$1,613,588
VLSI/CAD Network	248,918	685,374
	-----------	----------
Total award	$ 745,174	$2,298,962

4. System Diurnal Loading Variations

The following figures give a picture of the recent variations in diurnal SUMEX system load, taken during April 1980. The plots include:

Figure 13	- Total number of jobs logged in to the system
Figure 14	- System load average (average number of simultaneously runnable jobs)
Figure 15	- Percent of total CPU time used by logged in jobs (maximum is 200% for dual processor capacity)

The abscissa for these plots is broken into 20 minute intervals throughout the day. The ordinate for each interval is the average of all the daily measurements for that interval over the weekdays during April 1980. A daily measurement for a given 20 minute interval is in turn an average of the appropriate statistic sampled every 10 seconds. Since these plots display overall average data, they give representative illustration of the general characteristics of diurnal loading. There are, of course, substantial fluctuations in the quantities measured from day to day as well and for some, also on time scales shorter than the intervals displayed in the figures. For example in Figure 14, the number of runnable jobs shows a fairly smooth curve peaking at 5.2 jobs. On both a scale of minutes and from day to day, however, the number of runnable jobs will vary from only a few to 12 or more. These fluctuations are not shown in these average plots but also play an important role in the responsiveness of the system.

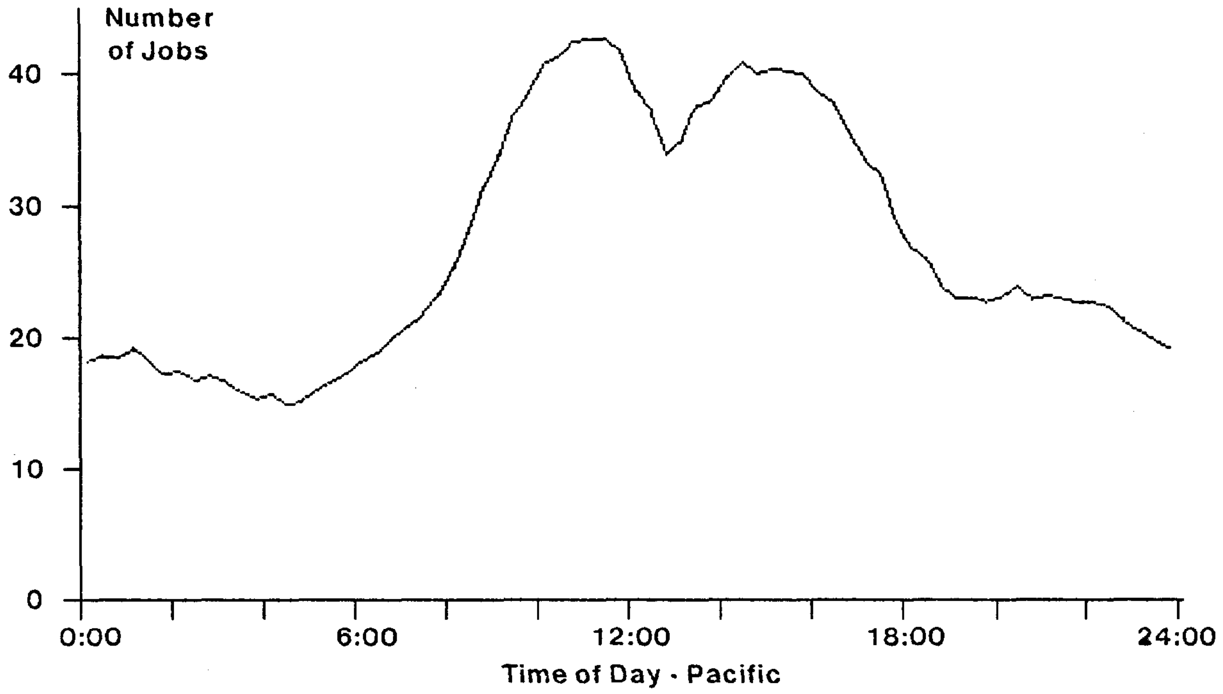

Figure 13. Average Diurnal Loading (4/80): Number of Jobs

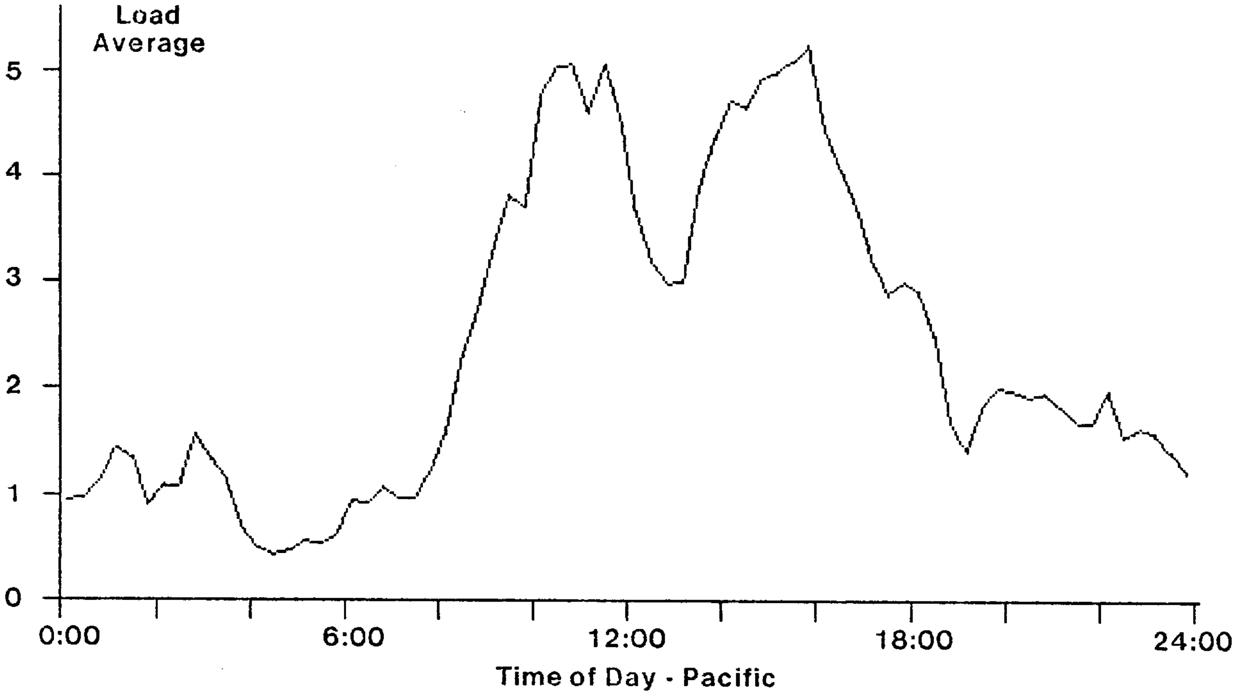

Figure 14. Average Diurnal Loading (4/80): Load Average

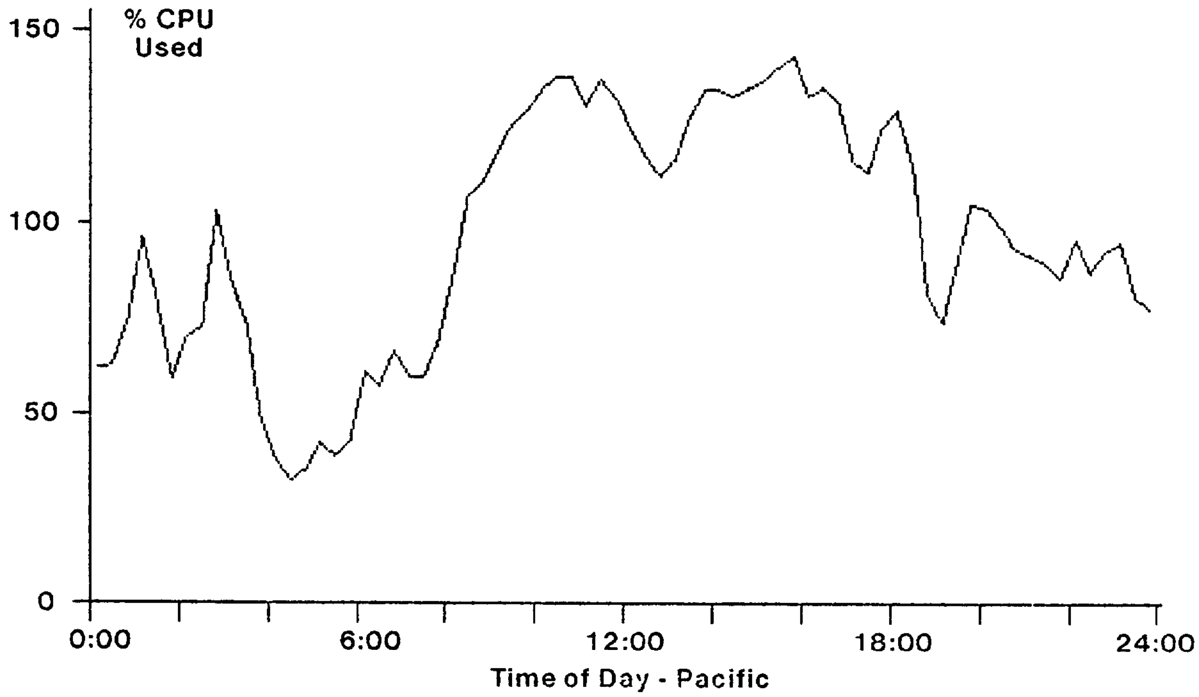

Figure 15. Average Diurnal Loading (4/80): Percent Time Used

5. Network Usage Statistics

The plots in Figure 16 and Figure 17 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.

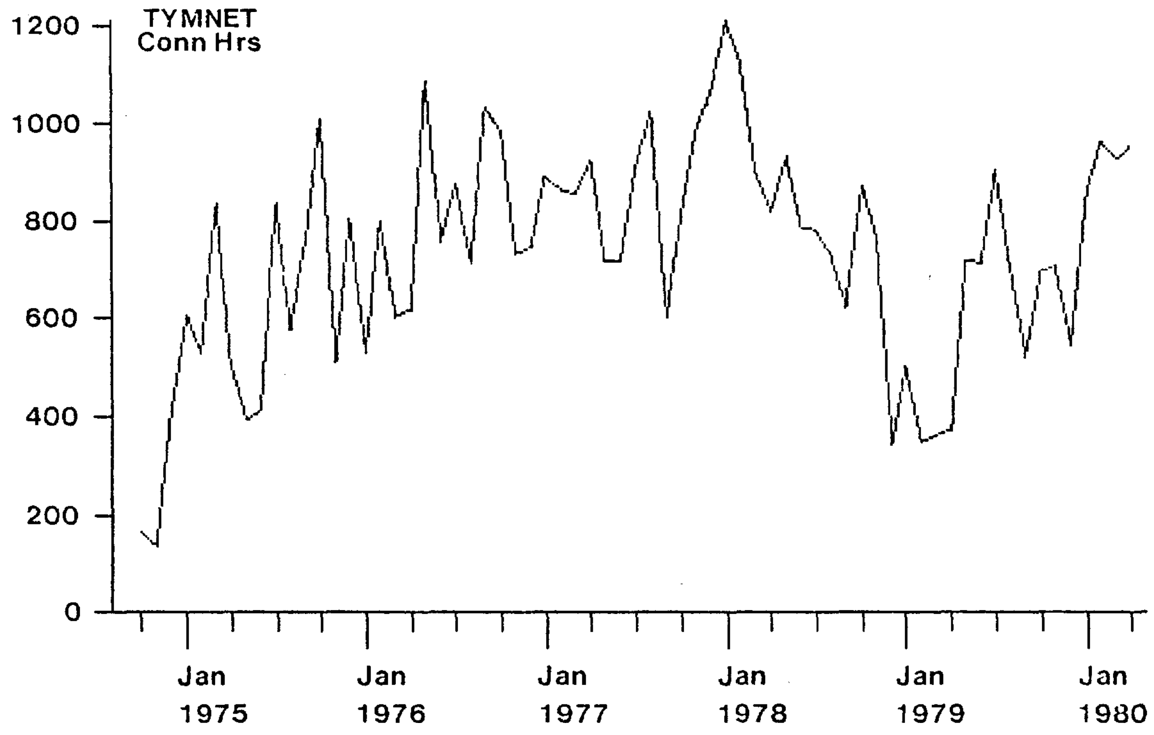

Figure 16. TYMNET Terminal Connect Time

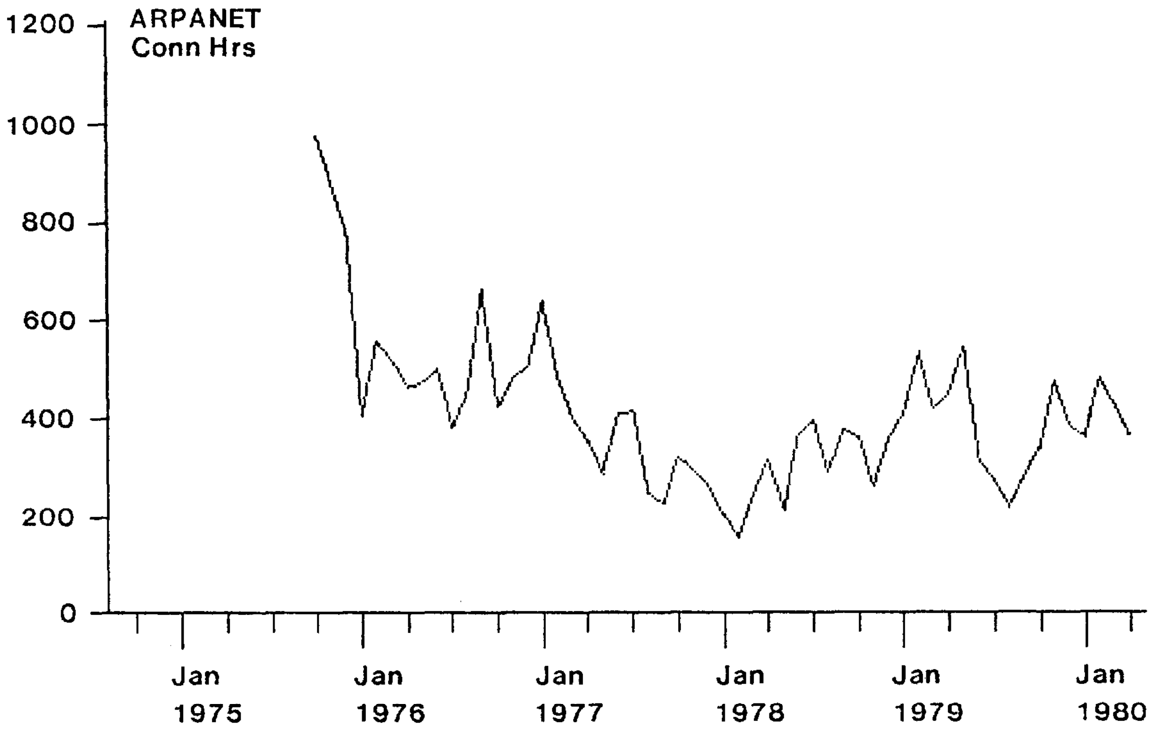

Figure 17. ARPANET Terminal Connect Time

6. 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.

				1979						1980		
	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	FEB	MAR	APR
RELOADS												
Hardware	12	4	1	15	1	0	13	6	4	2	9	8
Software	0	3	5	1	7	3	2	1	4	0	6	4
Environmental	1	1	0	0	0	1	0	1	0	1	0	2
Unknown Cause	0	2	3	1	3	0	0	2	0	0	0	1
	--	--	--	--	--	--	--	--	--	--	--	--
Totals	13	10	9	17	11	4	15	10	8	3	15	15
DOWNTIME (Hrs)												
Unscheduled	38	18	15	12	18	4	33	28	8	4	14	38
Scheduled	19	28	20	35	38	29	19	15	28	27	41	23
	--	--	--	--	--	--	--	--	--	--	--	--
Totals (Hrs)	57	46	35	47	56	33	52	43	36	31	55	61

TABLE 1. System Reliability by Month

Appendix C

Local Network Integration

The introduction of satellite machines into the SUMEX facility raises important issues about how best to integrate such systems with the existing machines. We seek to minimize disruptions to the operational resource with the addition of new machines, the duplication of peripheral equipment, and the interdependence among machines that would increase failure modes. We also require high-speed intermachine file transfer capabilities and terminal access arrangements allowing a user to connect flexibly to any machine of choice in the resource.

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 buses. With the addition of 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 interfaces, it would be inconvenient for a user to reconnect his terminal from one host to another. He would have to break his existing connection and go through another connect/login process to get to another machine. Since we want to facilitate user movement between various machines in the SUMEX resource, this process needs to be as simple as possible - in fact a user may have jobs running simultaneously on more than one machine at a time.

Similarly, we need to be able to quickly transfer files between any two machines in the resource, connect common peripheral devices (e.g. printer or plotter) to any machine desiring to use them, and allow any host to access other remote resources such as Stanford campus printers or terminal clusters. If we were to establish direct connections pairwise between machines and devices, the number of such connections would go up quadratically with the number of devices.

A more effective solution lies in the implementation of a local network in which all devices (host CPU's, peripheral devices, network gateways, etc.) are tied to a shared communications medium and can thereby establish logical connections as needed between any pair of nodes. Such network systems have been under development for a number of years, taking on various topological configurations and control structures depending on bandwidth requirements and interdevice distances. A very attractive design for a highly localized system configuration from the viewpoint of simplicity, reliability, and bandwidth is the Ethernet which has been under development for several years at Xerox Palo Alto Research Center [10]. The

Ethernet utilizes a fully distributed control structure in that each device connected to the net can independently decide to send a message to any other device on the net depending on the functions it is actively performing. Of course, decisions about which devices need to communicate with each other at a given time and what the precise message content is are determined by higher level system activities and requests, for example to implement a file transfer, mail forwarding, teletype connection, printer output, etc. Current Ethernets operate at 3 Mbits/sec and realize over 90% effective capacity utilization under heavy load [11]. Protocols exist is handle "collisions" between two devices trying to gain control of the network at the same time and to interconnect Ethernet with other networks.

The Stanford Computer Science Department is one of three recipients of grants from Xerox that includes Ethernet connection, terminal, and graphics printer equipment. Since the Computer Science Department systems are integrally connected with one of the major user groups on SUMEX (the Heuristic Programming Project) and since the Ethernet design is ideal for the the integration of new satellite machines with the existing SUMEX facility, we have chosen it as the model for our planned facility changes.

A diagram of the on-going Ethernet implementation for SUMEX is shown in Figure 3. Plans include developing interfaces for each host machine, the TYMNET, the local teletype scanner, other peripheral devices, and a gateway to other local networks (e.g., the Computer Science Department machine and planned terminal clusters). We already have the KI-10's connected through an I/O bus interface and are almost ready to debug the 2020 interface. These both use the Xerox interface board designed for PDP-11's. We are also working on a more efficient connection for the KI-10's through a direct memory access device and on connections for the other resources.

Appendix D

Remote Network Communication Facilities

Limitations for Interactive Work

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 fully 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 telephone lines to procure widespread geographical coverage at substantially reduced cost. 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 special problem for network users and is one of the reasons that coming more local computing systems will be especially important to improve the human interfaces to our AI programs. The key technological components to improved human engineering (high-speed bit-mapped displays, touch, and speech) all involve requirements for high bandwidth communications that can only be effectively implemented locally.

This does not diminish the importance of networks in our community, but rather enhances their role for facilitating remote scientific contacts, allowing remote access to regionalized resources, and sharing programs and knowledge bases. These are tasks for which national networks are ideally suited.

TYMNET

TYMNET provides broad geographic coverage for terminal access to SUMEX, spanning the country and also increasingly accessible from foreign countries (see Figure 18 on page 379). 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 installed and 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. A new tariff provision, based on "dedicated port" pricing, is advantageous to us though. This allows purchase of a number of logical network ports at the host for a fixed cost per month, independent of connect time or number of characters transmitted. We have implemented that option with BRP and save approximately $1,000 per month in service charges. We will continue to work closely with NIH-BRP and NLM to achieve the most cost-effective purchase of these services. The total use of TYMNET dropped during the TELENET experimental connection described below (see Figure 16) but has increased again since the TELENET service was dropped.

Technical aspects of our connection to TYMNET have remained unchanged since the last report and have continued to operate reasonably reliably. We have fixed several bugs in the TYMNET service related to handling editing terminals. Also we have had problems with incomplete closure of connections that can accumulate and leave us with all ports effectively blocked after long periods of uptime. The evidence points to a bug in TYMNET's interface code and we have had serious problems getting adequate support from them to fix the problem.

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 19 and Figure 20 on page 380. Consistent with 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 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.

The reconnection of the Rutgers resource to ARPANET has reopened our valuable scientific contacts with that subcommunity. In fact their efforts to justify reconnection may provide a basis for broader NIH use of the ARPANET and hence better network support for our collaborators.

TELENET

Initially SUMEX based its remote communication services on two networks - TYMNET and ARPANET. These were the only networks existing at the start of the project which allowed foreign host access. A third commercial network system, TELENET, is now competitively operational and offers a growing selection of services. Since our last review and with the advice and approval of the AIM Executive Committee and NIH-BRP, we established an experimental connection to TELENET to evaluate its technical and economic advantages relative of our existing connections. This initial experiment was unsuccessful but since then TELENET has been acquired by General Telephone and Electronics to provide a larger capital base. They have an aggressive program for augmenting network services and a reconnection may be of advantage sometime in the next grant term. A current TELENET network map is shown in Figure 21 on page 382.

Our experimental connection was via a TP-2200 interface with 12 asynchronous lines to the SUMEX host and one 4800 baud line connecting to the network proper. TELENET has many attractive features in terms of a symmetry analogous to that of the ARPANET for terminal traffic and file transfers and being a commercial network, it does not have the access restrictions of the ARPANET. Its tariff schedule also affords lower costs than TYMNET for comparable service volume.

However, despite system changes we made to optimize TELENET performance (Xon/Xoff facilities to improve traffic flow), users felt a substantial degradation in service when using TELENET as opposed to TYMNET. We insisted that users use TELENET whenever possible between November 1978 and May 1979 to maximize user accommodation so that problems arising from differences in access conventions would not cloud judgements of services. Complaints included poor node reliability, intolerable delays in response, uneven flow of terminal output, and poor operational management of the network in keeping users informed of network and host status. From the system viewpoint at SUMEX, we detected similar problems. We received ineffective system engineering support in trying to tune network parameters to optimize performance for our user community and poor or erroneous feedback about network failures and problem resolution. In practice, TELENET offered no service advantages over TYMNET, since no file transfer connections above 1200 baud were allowed, no facilities to control local versus remote echoing existed, and no electronic mail system existed to facilitate communication between network operations staff and host nodes. Also company financial problems portended substantial delays in remedying these problems.

Because of grant budget limitations, we were forced to decide between the TYMNET and TELENET connections. Based on the distinct user preference expressed for TYMNET, we decided to terminate the TELENET connection as of May 1, 1979. We will continue to monitor TELENET developments (and those of other potential national network servers, e.g., AT&T, IBM, and Xerox) and may recommend a reevaluation of an alternative source for network services in the future.

Figure 18a. TYMNET Network Access

TYMNET® | Domestic Access Locations

State	High density locations	Low density locations	Foreign exchange locations
Alabama		Birmingham	
Arizona		Phoenix* Tucson*	
Arkansas			Little Rock
California	El Segundo* Los Angeles* Mountain View Newport Beach* Oakland* Palo Alto Riverside/Colton San Francisco* San Jose/Cupertino* Ventura/Oxnard	Hayward Sacramento San Diego* Santa Rosa	Alhambra Burlingame Fresno Marina del Rey Norwalk San Clemente San Pedro Santa Barbara Van Nuys
Colorado	Denver*		Colorado Springs
Connecticut		Darien* Hartford*	Bridgeport Danbury New Haven Waterbury
Delaware			Wilmington
District of Columbia	Washington*		
Florida		Jacksonville Miami* Orlando* St. Petersburg Tampa	Ft. Lauderdale Pensacola West Palm Beach
Georgia		Atlanta*	Savannah
Idaho		Boise	
Illinois	Chicago*	Freeport Rockford Springfield*	Peoria
Indiana		Indianapolis* South Bend	Evansville Ft. Wayne Marion

*1200-baud access

September 1979

Figure 18a. TYMNET Network Access
(continued)

State	High density locations	Low density locations	Foreign exchange locations
Iowa		Des Moines Iowa City	Cedar Rapids
Kansas		Shawnee Mission* Wichita*	Topeka
Kentucky		Lexington Louisville*	
Louisiana		Baton Rouge* New Orleans*	Lafayette Shreveport*
Maryland	Baltimore*		
Massachusetts	Boston/Cambridge*	Springfield	
Michigan	Ann Arbor Detroit* Plymouth Southfield	Jackson Kalamazoo	Grand Rapids St. Joseph
Minnesota		Minneapolis*	
Mississippi			Jackson
Missouri		Kansas City* St. Louis*	
Nebraska		Omaha	
Nevada		Reno/Carson City*	Las Vegas
New Hampshire			Manchester Nashua
New Jersey	Englewood Cliffs Lyndhurst* Newark/Union* Piscataway Wayne	Princeton/ South Brunswick	Moorestown
New Mexico			Albuquerque*
New York	New York City*	Buffalo* Corning Rochester* Syracuse White Plains*	Albany Hempstead L.I. Huntington L.I. Niagara Falls

Figure 18a. TYMNET Network Access
(continued)

State	High density locations	Low density locations	Foreign exchange locations
North Carolina		Raleigh/Durham Winston-Salem*	Charlotte Greensboro
Ohio		Akron Cincinnati Cleveland* Columbus Dayton	Toledo
Oklahoma		Oklahoma City* Tulsa*	
Oregon		Portland*	
Pennsylvania	Philadelphia*	Erie Pittsburgh* Valley Forge	Allentown Harrisburg York
Rhode Island			Providence
South Carolina			Columbia Greenville
Tennessee		Chattanooga Memphis* Nashville	Knoxville
Texas	Houston*	Austin* Dallas* El Paso* Midland San Antonio	Baytown Beaumont Corpus Christi Ft. Worth Longview Lubbock Odessa
Utah		Salt Lake City*	
Vermont			Burlington
Virginia	Arlington*		Norfolk Richmond
Washington		Richland* Seattle*	Spokane*
West Virginia			Charleston
Wisconsin		Madison Milwaukee*	Oshkosh

TYMNET® International Access Locations

- **Argentina**
  - Buenos Aires
- **Australia**[*]
  - Sydney
- **Austria**
  - Vienna
- **Bahrain**
- **Belgium**
  - Brussels
- **Bermuda**[†]
- **Brazil**
  - Rio de Janeiro
- **Canada**
  - All Datapac cities
- **Denmark**
  - Copenhagen
- **Finland**
  - Helsinki
- **France**[*]
  - Paris
- **Germany**[*]
  - Frankfurt
- **Hong Kong**
- **Israel**
  - Tel Aviv
- **Italy**
  - Milan
  - Rome
- **Japan**[†]
  - Tokyo
- **Mexico**
  - Mexico City
- **Netherlands**
  - Amsterdam
- **New Zealand**[†]
  - Wellington
- **Norway**
  - Oslo
- **Philippines**
  - Manila
- **Portugal**
  - Lisbon
- **Puerto Rico**
  - San Juan
- **Singapore**
  - Singapore
- **Spain**[*]
  - Madrid
- **Sweden**
  - Farsta
- **Switzerland**[*]
  - Berne
- **United Kingdom**
  - London
- **United States**[‡]
  - Anchorage
  - Honolulu
  - Juneau

[*] Access can be made throughout the country with a local call.

[†] Projected for 1980.

[‡] Noncontinental.

October 1979

Figure 18b. TYMNET Network Access

Figure 19. **ARPANET GEOGRAPHIC MAP, APRIL 1980**

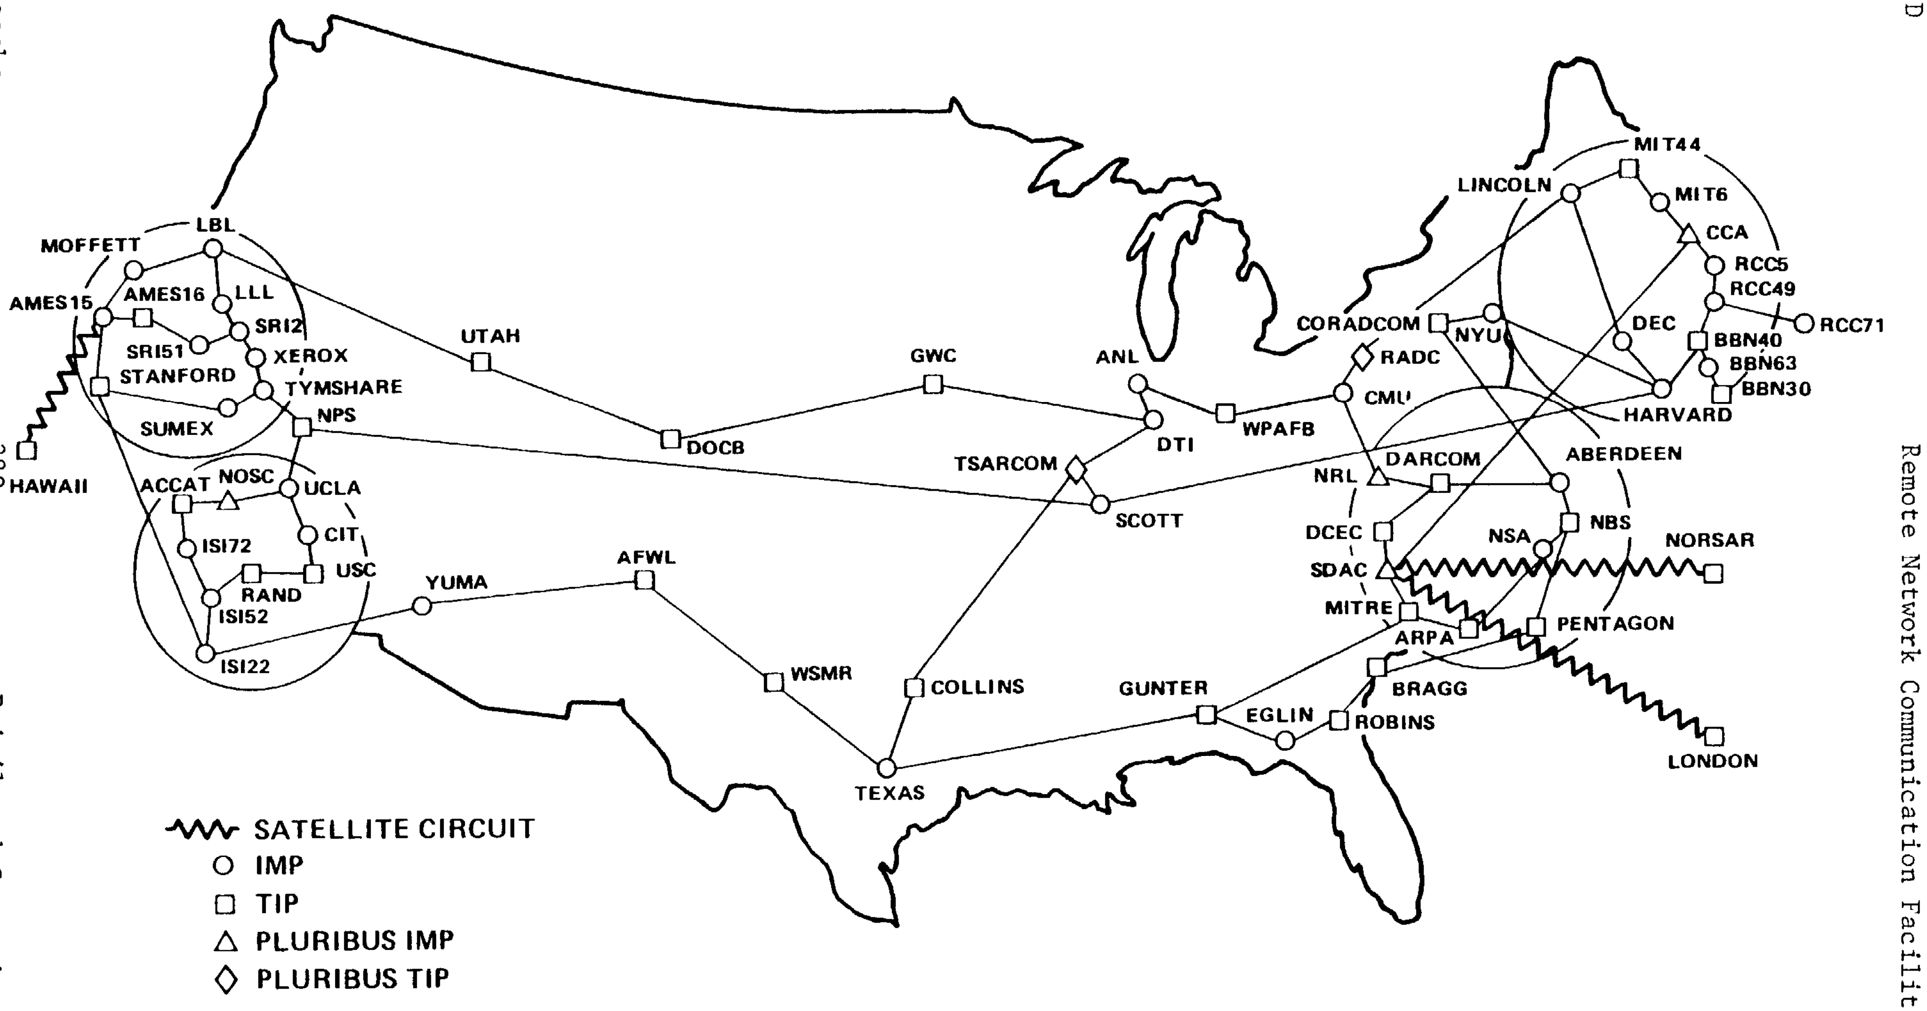
MIT44
LINCOLN
MIT6
CCA
RCC5
RCC49
RCC71
DEC
BBN40
BBN63
BBN30
HARVARD
CORADCOM
NYU
RADC
CMU
ABERDEEN
NRL
DARCOM
DCEC
NBS
NSA
NORSAR
SDAC
MITRE
ARPA
PENTAGON
BRAGG
ROBINS
LONDON
GUNTER
EGLIN
TEXAS
COLLINS
TSARCOM
SCOTT
ANL
DTI
WPAFB
GWC
DOCB
UTAH
WSMR
AFWL
YUMA
LBL
MOFFETT
AMES15
AMES16
LLL
SRI2
SRI51
XEROX
STANFORD
TYMSHARE
SUMEX
NPS
HAWAII
ACCAT
NOSC
UCLA
ISI72
CIT
RAND
USC
ISI52
ISI22
SATELLITE CIRCUIT
IMP
TIP
PLURIBUS IMP
PLURIBUS TIP

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

ARPANET LOGICAL MAP, MARCH 1980

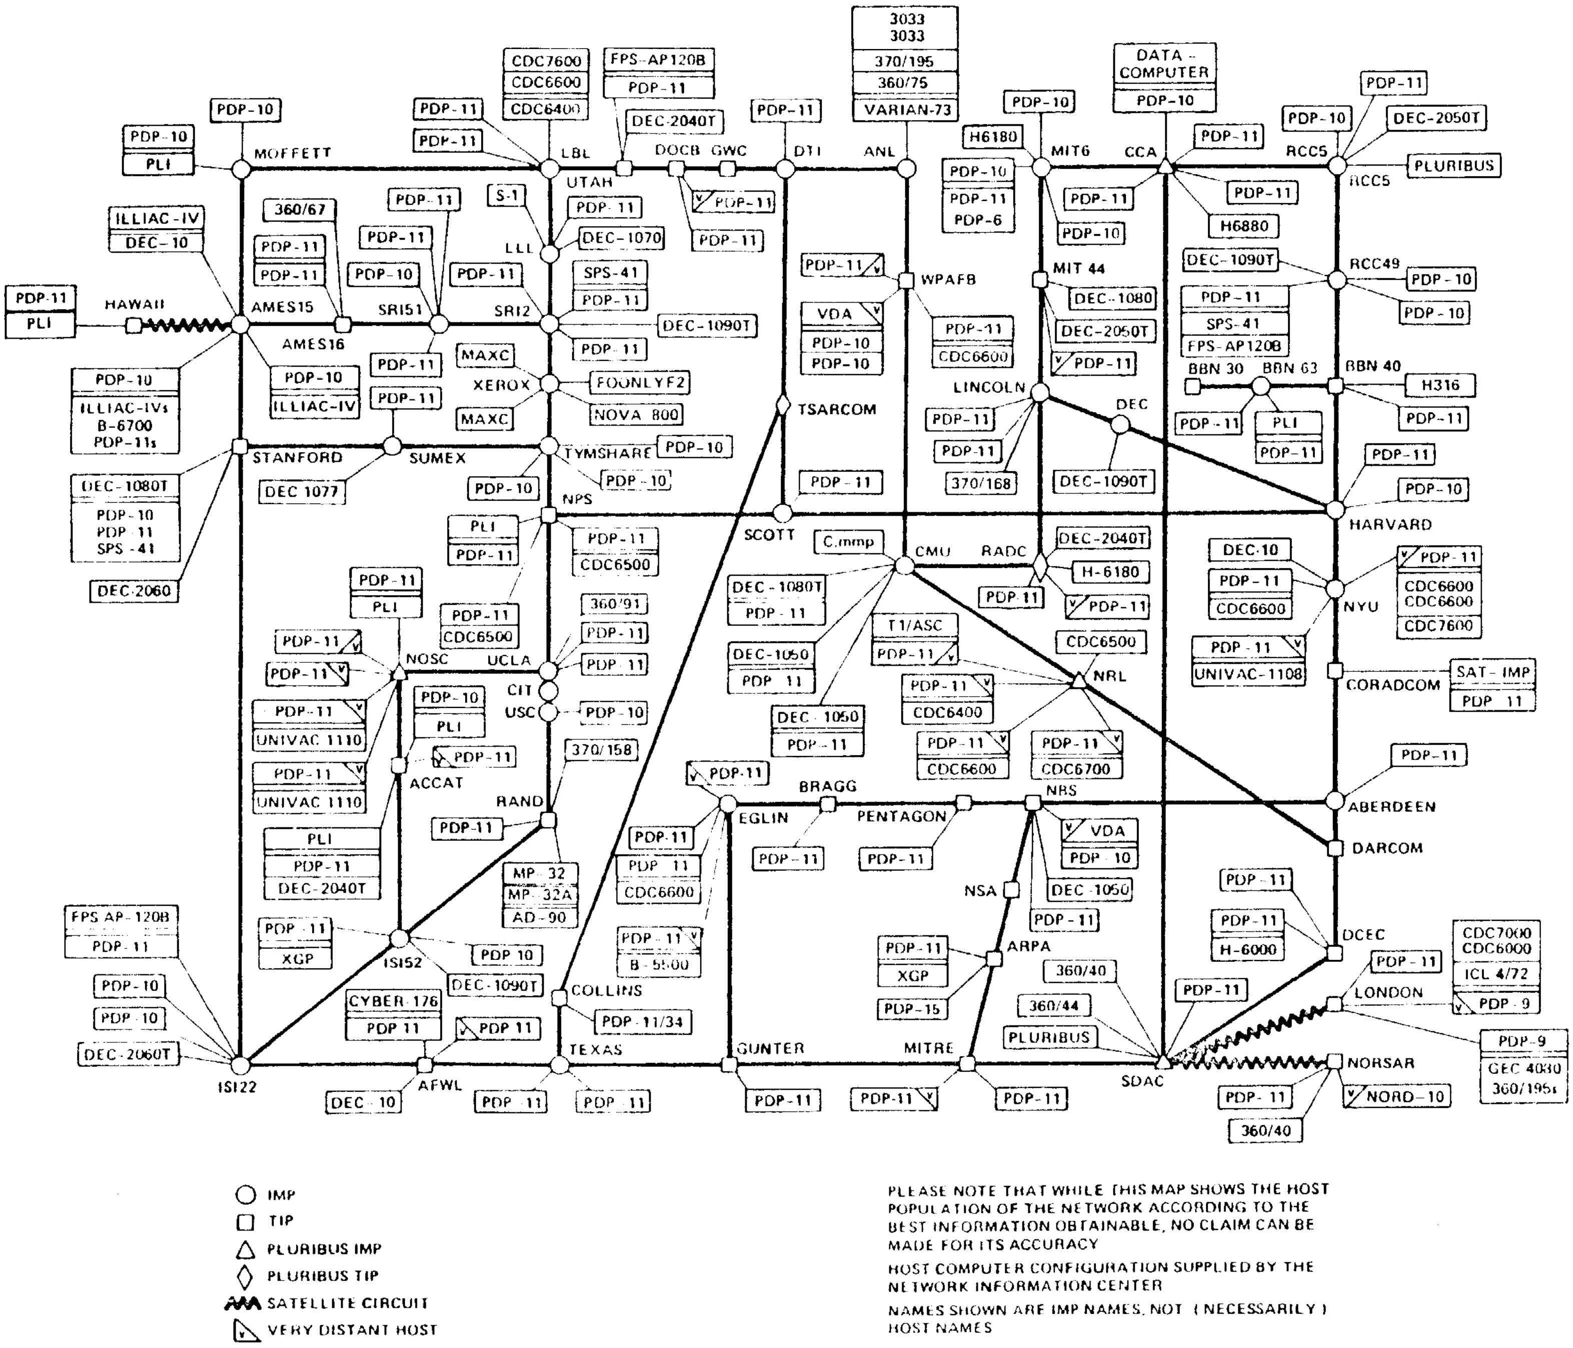
MOFFETT
PDP-10
PDP-10
PLI
LBL
UTAH
CDC7600
CDC6600
CDC6400
PDP-11
PDP-11
FPS-AP120B
PDP-11
DEC-2040T
DOCB
GWC
PDP-11
PDP-11
DTI
PDP-11
ANL
3033
3033
370/195
360/75
VARIAN-73
MIT6
PDP-10
H6180
PDP-10
PDP-11
PDP-6
PDP-10
CCA
DATA - COMPUTER
PDP-10
PDP-11
PDP-11
PDP-11
H6880
RCC5
PDP-10
PDP-11
DEC-2050T
PLURIBUS
ILLIAC-IV
DEC-10
360/67
PDP-11
PDP-11
PDP-11
PDP-11
PDP-10
PDP-11
S-1
LLL
PDP-11
DEC-1070
PDP-11
PLI
HAWAII
AMES15
AMES16
SRI51
SRI2
SPS-41
PDP-11
DEC-1090T
PDP-11
WPAFB
PDP-11
VDA
PDP-10
PDP-10
PDP-11
CDC6600
MIT 44
DEC-1080
DEC-2050T
PDP-11
RCC49
DEC-1090T
PDP-10
PDP-10
PDP-11
SPS-41
FPS-AP120B
BBN 30
BBN 63
BBN 40
H316
PDP-11
PDP-11
PLI
PDP-11
PDP-10
ILLIAC-IVs
B-6700
PDP-11s
PDP-10
ILLIAC-IV
PDP-11
PDP-11
MAXC
XEROX
MAXC
FOONLYF2
NOVA 800
TSARCOM
LINCOLN
PDP-11
PDP-11
370/168
DEC
DEC-1090T
STANFORD
SUMEX
TYMSHARE
PDP-10
PDP-10
DEC 1077
DEC-1080T
PDP-10
PDP-11
SPS-41
PDP-10
NPS
PDP-11
SCOTT
PDP-11
HARVARD
PDP-11
PDP-10
PLI
PDP-11
PDP-11
CDC6500
DEC-2060
C.mmp
CMU
RADC
DEC-2040T
H-6180
PDP-11
PDP-11
DEC-1080T
PDP-11
DEC-10
PDP-11
CDC6600
NYU
PDP-11
CDC6600
CDC6600
CDC7600
PDP-11
PLI
360/91
PDP-11
PDP-11
CDC6500
PDP-11
PDP-11
NOSC
UCLA
PDP-11
CIT
USC
PDP-10
PDP-10
PLI
T1/ASC
PDP-11
DEC-1050
PDP-11
CDC6500
NRL
PDP-11
UNIVAC-1108
CORADCOM
SAT - IMP
PDP-11
PDP-11
UNIVAC 1110
PDP-11
UNIVAC 1110
PDP-11
ACCAT
370/158
PDP-11
CDC6400
DEC 1050
PDP-11
PDP-11
CDC6600
PDP-11
CDC6700
PDP-11
ABERDEEN
PLI
PDP-11
DEC-2040T
RAND
PDP-11
PDP-11
EGLIN
BRAGG
PENTAGON
NBS
VDA
PDP-10
DARCOM
MP-32
MP-32A
AD-90
PDP-11
PDP-11
CDC6600
PDP-11
PDP-11
DEC 1050
PDP-11
NSA
PDP-11
PDP-11
H-6000
DCEC
FPS AP-120B
PDP-11
PDP-11
XGP
ISI52
PDP 10
DEC-1090T
PDP-11
B-5500
PDP-11
XGP
PDP-15
ARPA
360/40
360/44
PLURIBUS
PDP-11
PDP-11
LONDON
CDC7000
CDC6000
ICL 4/72
PDP-9
PDP-10
PDP-10
DEC-2060T
CYBER 176
PDP 11
PDP 11
COLLINS
PDP-11/34
ISI22
AFWL
DEC-10
TEXAS
PDP 11
PDP 11
GUNTER
PDP-11
MITRE
PDP-11
PDP-11
SDAC
NORSAR
PDP-11
360/40
NORD-10
PDP-9
GEC 4080
360/195

○ IMP
☐ TIP
△ PLURIBUS IMP
◇ PLURIBUS TIP
SATELLITE CIRCUIT
VERY DISTANT HOST

PLEASE NOTE THAT WHILE THIS MAP SHOWS THE HOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY

HOST COMPUTER CONFIGURATION SUPPLIED BY THE NETWORK INFORMATION CENTER

NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES

Figure 20. ARPANET Logical Network Map

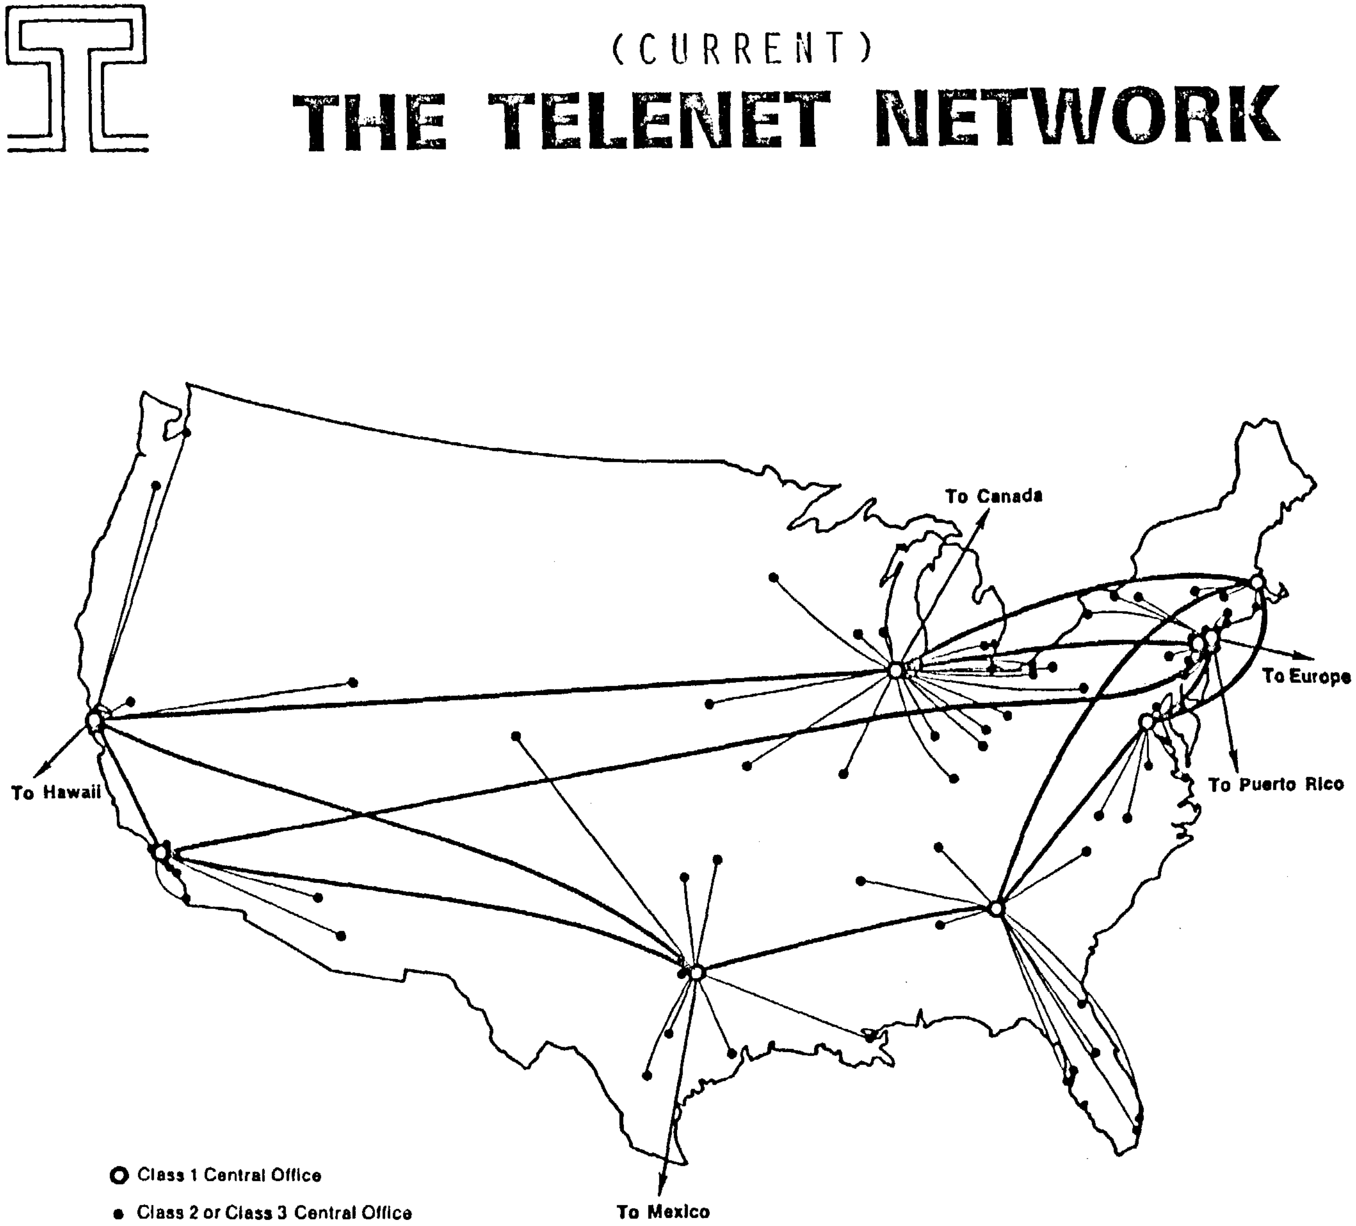
(CURRENT)
THE TELENET NETWORK
To Canada
To Europe
To Puerto Rico
To Hawaii
To Mexico
Class 1 Central Office
Class 2 or Class 3 Central Office

Figure 21a. TELENET Geographical Network Map

Figure 21b. (MID 1980)

THE TELENET NETWORK

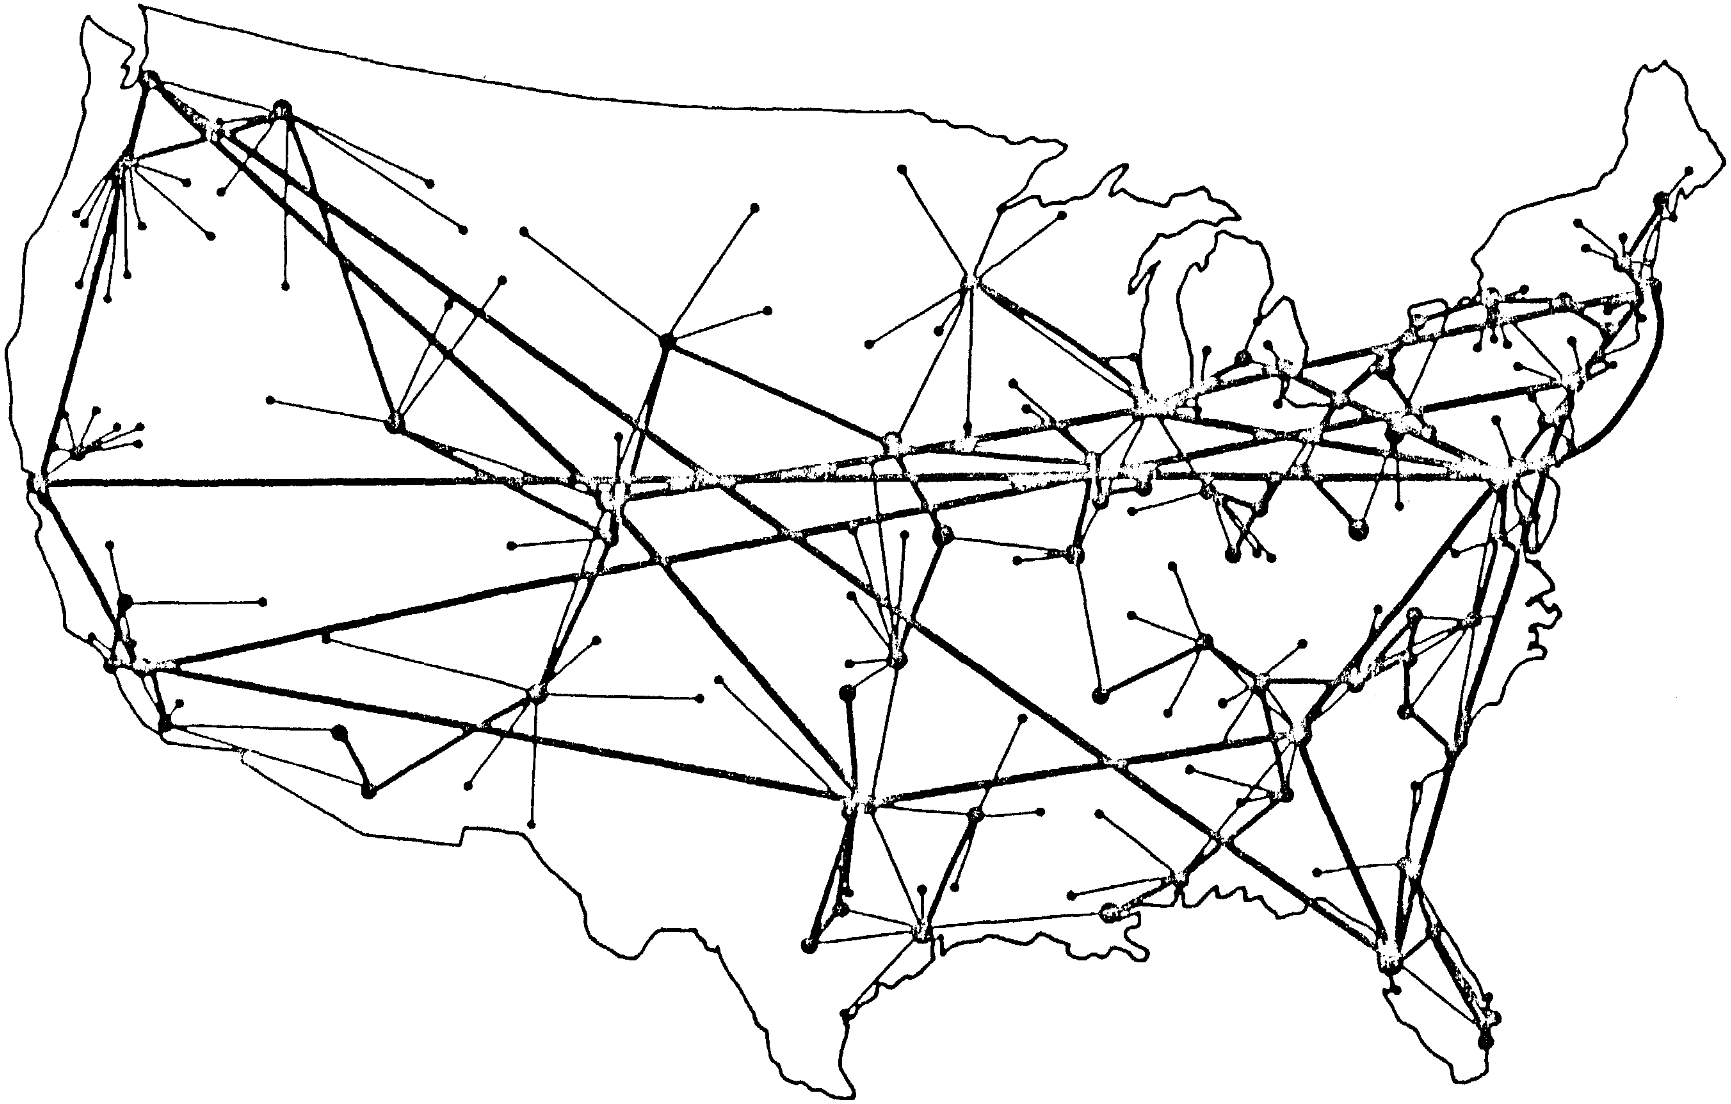

Appendix E

Resource Management Structure

Philosophy of Management

One way to administer a national resource is by subcontract to a fee-compensated, neutral agent under a governing body that could speak to the technical and quality-control interests of the served constituency. Appropriate in some circumstances, this model would separate the administration of the resource from active participation in the on-going research and development. An approach expected to foster greater creativity is to couple the resource closely with an active user-center. This of course can lead to manifest conflicts of interest that must be addressed and avoided if the resource is to be available fairly on a regional or national basis.

SUMEX-AIM has been based on the latter approach with a charter that spells out the underlying objectives and responsibilities of the program, and which establishes incentives, resources, and obligations for proper performance. Our resource design, incorporating all of these ingredients, has made the development of the procedural framework a matter of simple common-sense logic. It will be plain that the convergence of local self-interest with peer and contractual responsibility offers the best assurance that the programmatic goals will be respected and simplifies the tasks of surveillance and accountability.

The self-interest part of this equation stems from our original motivation in requesting the resource: the need for specialized computing facilities to support intense, interdisciplinary studies in applications of AI at Stanford University Medical School. Comprising several departments (Chemistry, Medicine, Genetics, and Computer Science), interwoven projects (DENDRAL, MYCIN, MOLGEN, Heuristic Programming), and principal faculty (Professors Feigenbaum, Lederberg, Djerassi, Shortliffe, and Buchanan), a substantial body of research has progressed and evolved over many years. Successful, stable collaborations of this scope are not readily found. This history both depends upon and contributes to the doctrine of resource-sharing that underlies the SUMEX-AIM effort.

One premise of the management plan is therefore the charter allocation of half the user-available capacity of the SUMEX facility to the Stanford complex of projects, subject to a local committee chaired by Professor Feigenbaum. This principle clearly defines the local benefit of the resource, minimizes anxiety and conflict-of-interest, and enables the local group to respond quite objectively to the allocations that are made by an Executive Committee for the "national" or non-Stanford aliquot (see the section on "Management Committees" below). Another important contribution to the success of the plan is the welcome participation of an NIH-BRP representative on the Executive Committee. What would be inappropriate meddling in the conduct of a narrower research project funded by NIH, is a communication channel and source of detached judgment that has been invaluable in expediting the innumerable decisions about which NIH must and should be consulted in the week-to-week business of the resource. The efficacy of this principle, as is appropriate to acknowledge here, has been validated and enhanced by the style and energy that Dr. William Baker has brought to this task.

Further consequences of the charter principles are the conscientious cultivation of the "national" community for the most efficacious use of its aliquot, and the further growth of distributed facilities in due course. In summer of 1977, a computing facility at Rutgers University was established, coupled to SUMEX-AIM via the ARPANET and with 15% of the user-available capacity allocated for AIM use with the advice of the AIM Executive Committee. An increasing number of projects are using that resource as reported in Section 9.

Finally, the recognition in the charter that SUMEX-AIM is not merely a retail-store for computer cycles, but the means of building a community, is a necessary basis for the morale of the whole operation and the rationale for no fee-for-service.

The remainder of this section will summarize the way in which these responsibilities are handled bureaucratically.

Organization and Procedures

The SUMEX-AIM resource is administered between the Departments of Medicine and Computer Science of Stanford University. Its mission, 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, psychological and affective behavior modeling, instrument data interpretation, and tool building to facilitate the development of new AI applications.

In July 1978, Professor Lederberg, the original SUMEX Principal Investigator, became president of The Rockefeller University. Professor Feigenbaum, chairman of the Stanford Department of Computer Science, took over as Principal Investigator of the SUMEX project. Because of Prof. Feigenbaum's role as co-Principal Investigator of SUMEX from its start and his long standing collaboration with Prof. Lederberg, the management transition took place very smoothly. The SUMEX-AIM community continues to function with the same high level of vitality as before and has continued to grow. Professor Lederberg retains an active role in the SUMEX-AIM community as chairman of the AIM Executive Committee and on a more frequent basis through the system message facilities.

Close scientific and administrative ties are retained with the Stanford medical community. Immediately following Prof. Lederberg's departure, Professor Stanley Cohen, new chairman of the Department of Genetics, provided this liaison. In recognition of the growing scope and significance of the clinical applications being pursued at SUMEX, we have recently significantly strengthened our contacts within the Stanford community in that area. Professor Edward H. Shortliffe, one of the key designers of MYCIN, has assumed the role of co-Principal Investigator of SUMEX and the project will become administratively part of the Stanford Department of Medicine, effective August 1980. As part of the largest clinical medicine department at Stanford, SUMEX will have increased visibility and opportunity to broaden its local scientific collaborations.

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 has 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 I.

For the national community, two committees serve complementary functions. An Executive Committee oversees the operations of the AIM resources (SUMEX and the AIM portion of the Rutgers facility) 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 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 I.

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 I.

These committees have functioned actively 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.

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 nearly quadrupled since the start of the project; others are working tentatively as pilot projects or are under review. Reports for the various projects can be found in Section 9 and a graphical summary of community growth in Appendix B.

In the recent past we have made numerous efforts to broaden outside awareness of work in the AIM community and to encourage new research projects including:

1) CONGEN workshop at Stanford, December 1978.

2) AGE workshop at Stanford, February 1980.

3) AI session in the Fourth Illinois Conference on Medical Information Systems, 1979.

4) INTERNIST participation in a course on AI computing at NIH, 1979.

5) AI session in the Association for Information Science meeting, 1979.

6) AI session at Sixth International Joint Conference on AI, August 1979 and extensive lecture tour among Japanese university and industrial research projects.

7) MYCIN and INTERNIST program demonstrations at the American College of Physicians meetings in 1979 and 1980.

We have prepared a variety of materials for prospective new users ranging from general information 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. Dr. E. Levinthal has prepared a questionnaire 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 support already formulated programs and have provided guidance for new investigators and projects to formulate new biomedical AI applications and establish appropriate collaborations between medical and AI scientists. The AIM Executive and Advisory Committees have also played important roles in suggesting to pilot efforts ways in which their research programs could be strengthened through better collaborative ties.

We have welcomed a number of visiting investigators at Stanford who were able to pay their own expenses, so they could see first hand how AI applications programs are formulated and get acquainted with the computing tools available. 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.

Stanford Community Building

The Stanford community has undertaken several internal efforts to encourage interactions and sharing between the projects centered here. Numerous classes and seminars have been held over the years including ones to introduce chemistry students to the DENDRAL programs and to develop the early versions of the AI Handbook 5 articles. We also hold weekly informal lunch meetings (SIGLunch) between community members to discuss general AI topics, concerns and progress of individual projects, or system problems as appropriate as well as having frequent outside invited speakers.

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 biomedical AI goals and to maximize the resources available for newly developing applications projects. At the last full AIM workshop, meetings of the AIM Advisory Group and Executive Committee were held to review the national AIM projects. These groups recommended continued access for all formal projects then on the system. They also recommended phasing out the Organ Culture pilot project.

In the fall of 1978, meetings of the Stanford Advisory Group were held to review projects supported out of the Stanford aliquot. The recommendation of this group was to phase out support for the Hydroid Project, pending work more directly applicable to SUMEX-AIM goals. The group also recommended phasing out the Quantum Chemistry and Genetics Applications pilot projects unless stronger AI relevance were established immediately. The Quantum Chemistry project has since developed close collaboration with the DENDRAL stereochemistry effort. The Genetics Applications project has transferred their work to other systems to continue their calculations on genetic demographic data and has stopped using SUMEX.

AIM Workshop Support

The Rutgers Computers in Biomedicine resource (under Dr. Saul Amarel) has organized a series of workshops devoted to a range of topics related to artificial intelligence research, medical needs, and resource sharing policies within NIH. Until recently, meetings have been held regularly at Rutgers.

In May 1979, a mini-AIM workshop devoted to clinical diagnosis programs was organized by MIT-Tufts and Rutgers and held in Vermont. This meeting was small (about 25 attendees) and emphasized detailed technical discussions about system designs and the strengths and weaknesses of various approaches. Many of the attendees were graduate students in order to maximize the benefit of personal contacts and discussions for on-going research projects. Topics covered in the discussions included state-of-the-art in explanation, causality in reasoning, strategies of focusing and dealing with multiple diagnostic problems, issues of representation and grain of description, creating and updating a knowledge base, planning strategies, issues of time representation, and inexact reasoning.

In August 1980, the AIM workshop will be held at Stanford as part of an extensive series of meetings. The workshop will be followed by a two-day series of tutorials for medical scientists to introduce them to AI computing goals and capabilities. This in turn will be followed by the first annual conference of the American Association for Artificial Intelligence devoted to a broad range of scientific issues in AI research.

The SUMEX facility has served as a communications base for workshop planning and provided support for workshop demonstrations when requested. We expect to continue this support for future workshops. The AIM workshops provide much useful information about the strengths and weaknesses of the performance programs both in terms of criticisms from other AI projects and in terms of the needs of practicing medical people. We plan to continue to use this experience to guide the community building aspects of SUMEX-AIM.

Resource Capacity Planning and Allocation Policies

As the SUMEX-AIM community has grown, the facility has become increasingly loaded and 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 implemented user-oriented policies in trying to give users the greatest latitude possible to pursue their research consistent with fairly meeting our responsibilities in managing SUMEX as a national resource.

We have described the details of our allocation procedures in earlier reports. These have been implemented to attempt to maintain the 40:40:20 balance in system use between Stanford, National, and staff communities. The initial complement of user projects justifying the SUMEX resource was centered to a large extent at Stanford. As the number of national has grown, so has the Stanford group of projects matured and in practice the 40:40 split between Stanford and non-Stanford projects is not ideally realized (see Appendix B). 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 new overload controls, the Stanford community still experiences mutual contentions and delays while the AIM group has relatively open access to the system. 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 actively encouraged mature projects to apply for their own machine resources in order to preserve the SUMEX-AIM resource for new AI applications. In the recent past, several projects have submitted proposals for such facilities including DENDRAL (see Section 9.1.3 on page 149). In spite of favorable reviews of the research project itself (resulting in a 3-year renewal), the study section did not want to see the DENDRAL project divert its energies to run a separate machine resource. Rather they felt such an augmentation should be coordinated and implemented by the SUMEX resource in conjunction with the DENDRAL group. Such a relationship is feasible in the case of the local DENDRAL project and we feel can serve as a model for further distribution of resources to advanced projects. We cannot effectively operate such resources for all the projects in our community but through experimentation with new machines, we can lay the groundwork for packaged systems that other groups may be able to acquire and easily operate. This mandate through the DENDRAL review is one of the bases for our long term plans for the coming renewal period.

Appendix F

LISP Address Space Limitations

In recent years, the program address space limitations imposed by the architecture of the PDP-10/20 systems have been increasingly felt in building large knowledge-based systems for biomedicine and in other application areas. Each user has access to a 256K 36-bit virtual address space (slightly more than 1M byte). For many conventional programs, this is adequate but the large language and program structures required for expert systems easily consume this space.

Current systems have used many approaches to compress their address space requirements including compiling established static code so it can be swapped between the main LISP space and an inferior fork and reorganizing dynamic code and data structures so they can be swapped between memory and hash-coded files. For example, space is now a critical problem for GUIDON because it is itself a large system built on top of another large system, MYCIN. In MYCIN, the dictionary, tables of facts (drugs/organism relations), and static properties that consume string space have already been moved off to disk in the form of hash files. In GUIDON, even this is not enough; MYCIN's rules must be hashed as well. For the short term, it appears that more of GUIDON's code will have to be non-resident ("recognized files"), thus trading time for space. Since response time is crucial for consultative programs, this trade-off is not acceptable.

Early in the development of Internist-I it became obvious that the 18 bit address space of INTERLISP imposed a severe limitation on the size of the knowledge base. The limit was on both atom and list space. To make matters worse there was no room left for the dynamic data structures (mostly lists) that are established by the diagnostic program. To get around this problem the INTERNIST group invested approximately 2 man years to develop a disk-oriented knowledge base that fetched and overlayed knowledge structures on demand. As a result all but the most trivial changes on knowledge structures are prohibitive, the system is not portable, and they still see an occasional case for which there is insufficient list space to be used by the diagnostic program.

Similar problems are anticipated in the development of Internist-II. The plan, at present, is to employ LISP hash files for the larger and/or infrequently accessed structures.

In both AGE and Meta-DENDRAL, it is not possible to load all the information on the system files into a single save file. This is handled by having different specialized environments that contain different system information, e.g., system execution and system development. In Meta-DENDRAL, all of the executing code will not fit in a single address space, so a system of selective loading is used based on dynamic demand. This reduces memory requirements for code but increases system overhead. In addition, DENDRAL has used a greatly stripped down version of LISP (also used by INTERNIST) in order to have sufficient data space to handle meaningful problems. They are still are constrained in problem complexity by the limited space to store data structures.

Similarly in MOLGEN, the address space in INTERLISP was sufficiently tight that the knowledge base would not fit in core, even at a very early stage in the project. To remedy this, they added a "virtual memory" system to the Units representation system which paged from a disk file on a demand basis. This patch basically made the PDP10 usable at a cost in execution time.

While the 18-bit address limit has not stopped research, it has stifled it by increasing overhead and causing users to scale down the scope of their research efforts. In order to minimize the cost of knowledge-base and program overlays, each project has had to tune their approach to the particular program structure. Even fairly modest ambitions push tolerance and system capacity to the limits. Much effort has gone into solving this problem in the ARPANET INTERLISP community. Address extensions for the PDP-10/20 class machines (including Foonly, Inc. machines) based on memory segmentation schemes do not lend themselves to a LISP environment since there is no intrinsic difference between program and data and the added overhead of keeping track of the extended address constructs with software becomes prohibitive. Thus, the solutions under active consideration include moving either to general purpose machines with larger logical address spaces (e.g., Prime or DEC VAX) or to special purpose LISP machines.

One of our objectives for the renewal period is to add facilities to the SUMEX-AIM resource that will provide a uniform and effective solution to these problems.

Appendix G

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 eight Chapters are expected to appear in Volume I. A tentative list of the all of articles in each Chapter follows.

I. Introduction
II. Search
III. Representation of Knowledge
IV. Natural Language Understanding
V. Speech Understanding
VI. AI Programming Languages
VII. Applications-oriented AI Research: Science
VIII. Applications-oriented AI Research: Medicine
IX. Applications-oriented AI Research: Education
X. Automatic Programming
XI. Information Processing Psychology
XII. Theorem Proving
XIII. Vision
XIV. Robotics
XV. Learning and Inductive Inference
XVI. Planning, Reasoning, and Problem Solving

I. INTRODUCTION

- A. The AI Handbook (intent, audience, style, use, outline)
- B. Overview of AI
- C. History of AI
- D. An Introduction to 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
    - b. Alpha-beta pruning
    - c. Heuristics in game tree search
- D. Example search programs
  - 1. Logic Theorist
  - 2. GPS
  - 3. Gelernter's geometry theorem-proving machine
  - 4. Symbolic integration programs
  - 5. STRIPS
  - 6. ABSTRIPS

III. Representation of Knowledge

- A. Issues and problems in representation theory
- 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. Natural Language Understanding

- A. Overview - History and issues
- B. Early attempts at mechanical translation
- C. Grammars
    - 1. Review of formal grammars
    - 2. Transformational grammars
    - 3. Systemic grammars
    - 4. Case grammars
- D. Parsing
    - 1. Overview of parsing techniques
    - 2. Augmented transition nets, Woods
    - 3. CHARTS - The GSP system
- E. Text generating systems
- F. Natural language processing systems
    - 1. Early NL systems
    - 2. Wilks' machine translation work
    - 3. MARGIE
    - 4. LUNAR
    - 5. SHRDLU
    - 6. SAM and PAM
    - 7. LIFER

V. Speech Understanding Systems

- A. Overview
- B. Some early ARPA speech systems
    - 1. DRAGON
    - 2. HEARSAY I
    - 3. SPEECHLIS
- C. Recent Speech Systems
    - 1. HARPY
    - 2. HEARSAY II
    - 3. HWIM
    - 4. SRI-SDC System

VI. AI Programming Languages

- A. Historical overview
- B. AI programming language features
    - 1. Overview and comparison
    - 2. Data structures
    - 3. Control structures
    - 4. Pattern matching
    - 5. Programming environment
- C. Major AI programming languages
    - 1. LISP
    - 2. PLANNER and CONNIVER
    - 3. QLISP
    - 4. SAIL
    - 5. POP-2

VII. Applications-oriented AI Research: Science and Mathematics
    A. Overview
    B. TEIRESIAS - Issues in expert systems design
    C. Applications in chemistry
        1. Applications in chemical analysis
        2. The DENDRAL Programs
            a. DENDRAL
            b. CONGEN and its extensions
            c. Meta-DENDRAL
        4. CRYSALIS
        5. Applications in organic synthesis
    D. Applications in mathematics
        1. MACSYMA
        2. AM
    F. Miscellaneous science applications research
        1. The SRI Computer-Based Consultant
        2. PROSPECTOR

VIII. Applications-oriented AI Research: Medicine
    A. Overview
    B. Medical systems
        1. MYCIN
        2. CASNET
        3. INTERNIST
        4. Present Illness Program
        5. Digitalis Advisor
        6. IRIS

IX. Applications-oriented AI Research: Education
    A. Historical overview
    B. Issues in ICAI systems design
    C. ICAI Systems
        1. SCHOLAR
        2. WHY
        3. SOPHIE
        4. WEST
        5. WUMPUS
        6. BUGGY
        7. EXCHECK

X. Automatic Programming
    A. Overview - Methods of program specification
    B. Basic approaches
    C. AP Systems
        1. PSI
        2. SAFE
        3. Programmer's Apprentice
        4. PECOS
        5. DAEDALUS
        6. PROTOSYSTEM-1
        7. NLPQ
        8. LIBRA - Program Optimization

XI. Information Processing Psychology
- A. Overview
- B. GPS
- C. Cognitive development
- D. EPAM
- E. Semantic network models
  - a. Quillian's network
  - b. LNR's MEMOD
  - c. HAM
  - d. ACT
- F. Belief systems

XII. THEOREM PROVING
- A. Overview
- B. Logic
- C. Resolution theorem proving
  - 1. Basic resolution method
  - 2. Syntactic ordering strategies
  - 3. Semantic and syntactic refinement
- D. Non-resolution theorem proving
  - 1. Overview
  - 2. Natural deduction
  - 3. Boyer-Moore
  - 4. LCF
- E. Applications of theorem proving
  - 1. Use in question answering
  - 2. Use in problem solving
  - 3. Theorem proving programming languages
  - 4. Man-machine theorem proving
  - 5. Use in automatic programming
- F. Proof checkers

XIII. VISION
- A. Overview
- B. Image-level processing
  - 1. Overview
  - 2. Edge detection
  - 3. Texture
  - 4. Region growing
  - 5. Overview of pattern recognition
- C. Spatial-level processing
  - 1. Overview
  - 2. Stereo information
  - 3. Shading
  - 4. Motion
- D. Object-level processing
  - 1. Overview
  - 2. Generalized cones and cylinders
- E. Scene level processing

- F. Vision systems
    - 1. Polyhedral or Blocks World vision
        - a. Overview
        - b. COPYDEMO
        - b. Guzman
        - c. Falk
        - d. Waltz
        - e. Navatya
    - 2. Robot vision systems
    - 3. Perceptrons

XIV. Robotics

- A. Overview
- B. Robot planning and problem solving
- C. Arms
- D. Present-day industrial robots
- E. Robotics programming languages

XIII. Learning and Inductive Inference

- A. Overview
- B. Simple inductive tasks
    - 1. Sequence extrapolation
    - 2. Grammatical inference
- C. Pattern recognition
    - 1. Character recognition
    - 2. Other recognition tasks
- D. Learning rules and strategies of games
    - 1. Formal analysis
    - 2. Examples of game-learning programs
- E. Single concept formation
- F. Multiple concept formation: Structuring a domain (AM, Meta-DENDRAL)
- G. Interactive cumulation of knowledge (TEIRESIAS)

XIV. Problem Solving, Planning & Reasoning by Analogy

- A. Overview of problem solving
- B. Planning
    - 1. Overview
    - 2. STRIPS (see IID5)
    - 3. ABSTRIPS (see IID6)
    - 4. NOAH
    - 5. HACKER
    - 6. INTERPLAN
    - 7. Rieger's causal reasoning system
    - 8. Rutgers work
    - 7. QA3 (see IXE1)
- C. Reasoning by analogy
    - 1. Overview
    - 2. Evans's ANALOGY program
    - 3. ZORBA
    - 4. Winston's learning system
- D. Contraint relaxation
    - 1. Waltz
    - 2. REF-ARF
- E. Game playing

Appendix H

MAINSAIL System Demonstration

As of July 30, 1979, the MAINSAIL project has successfully designed, demonstrated, and documented an ALGOL-like language system for machine-independent software design. This system includes the compiler, code generators, and run-time support for a range of target machine environments including TENEX, TOPS-20, TOPS-10, RT-11, and RSX-11. The designs for other environments have been studied but resources have not allowed more extensive implementations. Within Council-approved funding and manpower limits and the AI charter of the SUMEX resource, we do not have access to the more extensive resources that would be required to continue effective development and export of this system beyond this initial research and demonstration phase. The principal individuals involved (Messrs. Wilcox and Jirak and Ms. Dageforde) have formed a small private company, XIDAK, to support and continue development of MAINSAIL under license from Stanford University. XIDAK has almost completed a VAX implementation of MAINSAIL and is pursuing interests from a growing group of potential users, including a microprogrammed implementation for the PERQ computer. The following is a brief summary of recent work in this final demonstration phase of the MAINSAIL effort. Detailed reports on the language manual and design description can be found in references 14 and 15.

1) The compiler has undergone major reexamination and improvement with a substantial reduction in the size of data structures. As a result, it is now able to run on 16-bit machines with small address spaces (e.g., 32K words).

2) The runtime systems were thoroughly reexamined for optimizing execution efficiency and memory utilization. The garbage collection facility, used in the dynamic storage allocation system, was also substantially improved.

3) A new approach to code generation was introduced utilizing tree structures for the intermediate representation, rather than the more primitive triples or quadruples.

4) Facilities for managing "module libraries" of executable MAINSAIL modules were implemented.

5) At the conclusion of the demonstration phase, there were three sites using the TENEX version, six using the TOPS-10 version, and five using the TOPS-20 version.

6) A research project based on MAINSAIL is underway, aimed at an efficient program execution and development environment implemented on a microcoded "MAINSAIL machine" which directly executes a tailor-made MAINSAIL instruction set. This is the basis of Wilcox's Ph.D. thesis.

Appendix I

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

LEVINTHAL, Elliott C., Ph.D.
Department of Genetics, S047
Stanford University Medical Center
Stanford, California 94305
(415) 497-5813