

II. Description of Program Activities

This section corresponds to the predefined forms required by the Division of Research Resources to provide information about our resource activities for their computerized retrieval system. These forms have been submitted separately and are not reproduced here to avoid redundancy with the more extensive narrative information about our resource and progress provided in this report.

II.A. Scientific Subprojects

Our core research and development activities are described starting in section III.A.2, our training activities are summarized in section III.A.2.7, and the progress of our collaborating projects is detailed starting in section IV.

II.B. Books, Papers, and Abstracts

The list of recent publications for our core research and development work is given in section III.A.2.5 and those for the collaborating projects are in the individual reports starting in section IV.

II.C. Resource Summary Table

The details of resource usage, including a breakdown by the various subprojects, is given in the tables starting in section III.A.2.8

III. Narrative Description

III.A. Summary of Research Progress

III.A.1 Resource Overview

This is an annual report for year 16 of the SUMEX-AIM resource (grant RR-00785), the third year of a 5-year renewal period to support further research on applications of artificial intelligence in biomedicine. For the technical and administrative reasons discussed in earlier reports, the SUMEX project now includes the continuation of work on the development and dissemination of medical consultation systems (ONCOCIN) that had been supported before 1986 as resource-related research under grant RR-01631. Progress on core ONCOCIN research is therefore now reported here as well.

The originally proposed research program (June 1985 renewal application) included an ambitious plan to:

- Continue our long-range core research efforts on knowledge-based systems, aimed at developing new concepts and methodologies needed for biomedical applications.
- Substantially extend ONCOCIN research on developing and disseminating clinical decision support systems.
- Develop the core systems technology to move the national SUMEX-AIM community from a dependence on the central SUMEX DEC 2060 to a fully distributed, workstation-based computing environment.
- Introduce these systems technologies into the SUMEX-AIM community with appropriate communications and managerial assistance to responsibly phase out the central resource and DEC 2060 mainframe in a manner that will support community efforts to become self-sustaining and to continue scientific interactions through fully distributed means.
- Maintain our aggressive efforts at training and dissemination to help exploit the research potential of this field.

III.A.1.1 SUMEX-AIM as a Resource

SUMEX and the AIM Community

Since the SUMEX-AIM resource was established in late 1973, computing technology and biomedical artificial intelligence research have undergone a remarkable evolution and SUMEX has both influenced and responded to these changing technologies. It is widely recognized that our resource has fostered highly influential work in biomedical AI — work from which much of the expert systems field emerged — and that it has simultaneously helped define the technological base of applied AI research.

The focus of the SUMEX-AIM resource continues to emphasize research on artificial intelligence techniques that guide the design of computer programs

that can help with the acquisition, representation, management, and utilization of the many forms of medical knowledge in diverse biomedical research and clinical care settings — ranging from biomolecular structure determination and analysis, to molecular biology, to clinical decision support, to medical education. Nevertheless, we have long recognized that the ultimate impact of this work in biomedicine will be realized through its synthesis with the full range of methodologies of medical informatics, such as data bases, biostatistics, human-computer interfaces, complex instrument control, and modeling. From the start, SUMEX-AIM work has been grounded in real-world applications, like systems for the interpretation of mass spectral information about biomolecular structures, chemical synthesis, interpretation of x-ray diffraction data on crystals, cognitive modeling, infectious disease diagnosis and therapy, DNA sequence analysis, experiment planning and interpretation in molecular biology, and medical instruction. Our current work extends this emphasis in application domains such as oncology protocol management, clinical decision support, protein structure analysis, and data base information retrieval and analysis. All of these research efforts have demanded close collaborations with diverse parts of the biomedical research community and the integration of many computational methods from those domains with knowledge-based approaches. Even though in the beginning the "AI-in-medicine" community was quite small, it is perforce no longer limited and easily-defined, but rather is spreading and is inextricably linked with the many biomedical applications communities we have collaborated with over the years. Driven both by the on-going diffusion of AI and by the development of personal computer workstations that signal the practical decentralization of computing resources, we must develop new resource communication and distributed computing technologies that will continue to facilitate wider intra- and inter-community communication, collaboration, and sharing of biomedical information.

The SUMEX Project has demonstrated that it is possible to operate a computing research resource with a national charter and that the services providable over networks were those that facilitate the growth of AI-in-Medicine. SUMEX now has a reputation as 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 and new resources, like the BIONET resource for molecular biologists, are closely patterned after the SUMEX example.

The projects SUMEX supports have generally required substantial computing resources with excellent interaction. Today, with the dramatic explosion of high-performance workstations that are more and more generally available, the need for a central source of raw computing cycles has significantly diminished. In place of being a distributor of CPU cycles, SUMEX has become a communications cross-roads and a source of AI and computer systems software and expertise.

SUMEX has demonstrated that a computer resource is a useful "linking mechanism" for bringing together electronically teams of experts from different disciplines who share a common problem focus. 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 and dissemination efforts of SUMEX have been able to bridge the gap in numerous cases. Over 40 biomedical AI application projects have developed in our national community and have been supported directly by SUMEX computing resources over the years — many more have benefitted indirectly through access to the software, information, and advice offered by the SUMEX resource.

The integration of AI ideas with other parts of medical informatics and their dissemination into biomedicine is happening largely because of the development in the 1970's and early 1980's of methods and tools for the application of AI concepts to difficult professional-level problem solving. Their impact was heightened because of 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 home of such well-known AI systems as DENDRAL (chemical structure elucidation), MYCIN (infectious disease diagnosis and therapy), INTERNIST (differential diagnosis), ACT (human memory organization), MOLGEN/BIONET (tools for DNA sequence analysis and molecular biology experiment planning), ONCOCIN (cancer chemotherapy protocol advice), SECS (chemical synthesis), EMYCIN (rule-based expert system tool), and AGE (blackboard-based expert system tool). Since 1980, our community has published a fifteen books that give a scholarly perspective on the scientific experiments we have been performing. These volumes, and other work done at SUMEX, have played a seminal role in structuring modern AI paradigms and methodology.

The Future of SUMEX-AIM

Given this background, what is the future need and course for SUMEX as a resource — especially in view of the on-going revolution in computer technology and costs and the emergence of powerful single-user workstations and local area networking? The answers remain clear.

Basic Research on AI in Biomedicine

At the deepest research level, despite our considerable success in working on medical and biological applications, the problems we can attack are still sharply limited. Our current ideas fall short in many ways against today's important health care and biomedical research problems brought on by the explosion in medical knowledge and for which AI should be of assistance. Just as the research work of the 70's and 80's in the SUMEX-AIM community

fuels the current practical and commercial applications, our work of the late 80's will be the basis for the next decade's systems.

The report of the panel on medical informatics¹, convened late in 1985 by the National Library of Medicine to review and recommend twenty-year goals for the NLM, listed among its highest priority recommendations the need to greatly expand and aggressively pursue an interdisciplinary research program to develop computational methods for acquiring, representing, managing, and using biomedical knowledge of all sorts for health care and biomedical research. Similar recommendations have been stated recently by the panel on Information Technology and the Conduct of Research of the National Academy of Science². These are precisely the problems which the SUMEX-AIM community has been working on so successfully and which will require work well beyond the five year funding period we have requested. It is essential that this line of research in the SUMEX-AIM community, represented by our core AI research, the ONCOCIN research, and our collaborative research groups, be continued.

The Changing Role of the Central Resource

At the resource level, there are changing, but still growing, needs for computing resources for the active AIM research community to continue its work over the next five years. The workstations to which we directed our attention in 1980 have now demonstrated their practicality as research tools and, increasingly, as mechanisms for disseminating AI systems as cost-effective decision aids in clinical settings such as private offices. The era of highly centralized general machines for AI research is nearly at an end and is being replaced by networks of distributed but heterogeneous single-user machines sharing common information resources and communication paths among members of the biomedical research community.

Most of our community groups have been able to take advantage of local computing facilities, with SUMEX-AIM providing a central cross-roads for communications and the sharing of programs and knowledge. In its core research and development role, SUMEX-AIM has its sights set on the hardware and software systems of the next decade. We expect major changes in the distributed computing environments that are just now emerging in order to make effective use of their power and to adapt them to the development and dissemination of biomedical AI systems for professional user communities. In its training role, SUMEX is a crucial resource for the education of badly needed new researchers and professionals to continue the development of the biomedical AI field. The "critical mass" of the existing

¹ *Long Range Plan*. Report of the Board of Regents, National Library of Medicine. National Institutes of Health. January 1987.

² *Information Technology and the Conduct of Research — The User's View*. Report of the Panel on Information Technology and the Conduct of Research, National Academy of Sciences. National Academy Press. 1989.

physical SUMEX resource, its development staff, and its intellectual ties with the Stanford Knowledge Systems Laboratory (KSL — see Appendix A for a summary of current KSL research activities), make this an ideal setting to integrate, experiment with, and export these methodologies for the rest of the AIM community.

We will continue our experimental approach to distributed systems, learning to build and exploit distributed networks of these machines and to build and manage graceful software for these systems. Since decentralization is central to our future, we must learn its technical characteristics.

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 and for the dissemination of AI technology. 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 researchers¹. Another of the key recommendations of the NLM medical informatics planning panel² was that high-speed network communication links be established throughout the biomedical research community so that knowledge and information can be shared across diverse research groups and that the required interdisciplinary collaborations can take place. Recent efforts to establish a national NSFNet³, largely to support the supercomputer projects funded by NSF but also to replace and upgrade part of the national research community linkage that the now aging ARPANET has supported, have made important progress. Still, these efforts do not encompass the broad range of biomedical research groups that need national network access and to date, the NIH has not played an aggressive role in the interagency Research Internet coordination efforts. We must work to build a stronger institutional support for a National Research Network.

SUMEX continues to be an important pathfinder to develop the technology and community interaction tools needed to expand community system and communication resources. Our community building effort is based upon the developing state of distributed computing and communications technology

¹ Lederberg, J. "Digital Communications and the Conduct of Science: the New Literacy." *Proc. IEEE*, 66(11):1314-1319, 1978.

Coulter, C. L. "Research Instrument Sharing." *Science*, 201(4354), 1978.

Newell, A., and Sproull, R. F. "Computer Networks - Prospects for Scientists." *Science*, 215(4534):843, 1982.

² *NLM Long Range Plan: Medical Informatics*. NLM Planning Panel 4. National Library of Medicine, National Institutes of Health. January 1987.

³ Marshal, E. "NSF Opens High-Speed Computer Network." *Science*. 243: 22-23, 1989.

and we have therefore turned our core systems research to actively supporting the development of distributed computing and communications resources to facilitate collaborative project research and continued inter-group communications.

Summary of Long-term Goals

- Maintain the synergistic relationship between SUMEX core system development, core AI research, our experimental efforts at disseminating clinical decision-making aids, and new applications efforts.
- Continue to serve the national AIM research community, less and less as a source of raw computing cycles and more and more as a transfer point for new technologies important for community research and communication. We will also continue our coordinating role within the community through electronic media and periodic AIM workshops
- Maintain our connections to national networks (e.g., NSFNet, ARPANET, and TELENET) and our local Ethernet and assist other community members to establish similar links by example, by integrating and providing enabling software, and by offering advice and support within our resources.
- Focus new computing resource developments on more effective exploitation of distributed workstations through better communication and cooperative computing tools, using transparent digital networking schemes.
- Enhance the computing environments of workstations so that only minimal dependency on central, general-purpose computing hosts remains and these mainframe time-sharing systems can be phased out eventually. Remaining central resources will include servers for communications, community information resources, and special computing architectures (e.g., shared- or distributed-memory symbolic multiprocessors) justified by cost-effectiveness and unique functionality.
- Incrementally phase in, disseminate, and evaluate those aspects of the local distributed computing resource that are necessary for continuing national AIM community support within this distributed paradigm. This will ultimately point the way towards the distributed computing resource model that we believe will interlink this community well into the next decade.
- Responsibly phase out the existing DEC 2060 machine as effective distributed computing alternatives become widely available. Because of severe budget pressures, the 2060 was taken out of routine service during this past year in a much more accelerated fashion than was planned or was comfortable for AIM users to acclimate to the new UNIX operating system environment. We are still finishing up a number of interim systems alternatives to discontinued 2060 services not available in standard UNIX environments.

- Continue the central staff and management structure, essentially unchanged in function during the five-year transition period, except for the merging of the core part of the ONCOCIN research with the SUMEX resource.

III.A.1.2 Significance and Impact in Biomedicine

Artificial intelligence is the computer science of representations of symbolic knowledge and its use in symbolic inference and problem-solving processes. Projects in the SUMEX-AIM community are concerned in some way with the application of AI to biomedical research and the resource has given strong impetus and support to knowledge-based system research in biomedicine. For computer applications in medicine and biology, this research path is crucial. Medicine and biology are not presently mathematically-based sciences; unlike physics and engineering, they are seldom 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.

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

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. AI methodology is beginning to fill that need. Heuristic knowledge can be elicited, studied, critiqued by peers, and taught to students.

The importance of AI research and its applications is increasing in general, without regard for the specific areas of biomedical interest. AI has been one of the principal fronts along which university computer science groups are expanding. The pressure from student career-line choices is great. Federal and industrial support for AI research and applications is vigorous, although support specifically for biomedical applications continues to be limited. All of the major computer manufacturers (e.g., IBM, DEC, TI, UNISYS, HP, Apple, and others) are using and marketing AI technology aggressively and many software companies are putting more and more products on the market. Many other parts of industry are also actively pursuing AI applications in

their own contexts, including defense and aerospace companies, manufacturing companies, financial companies, and others¹.

Despite the limited research funding available, there is also an explosion of interest in medical AI. The American Association for Artificial Intelligence (AAAI), the principal scientific membership organization for the AI field, has over 7000 members, several thousand of whom are members of the medical special interest group known as the AAAI-M. Speakers on medical AI are prominently featured at professional medical meetings, such as the American College of Pathology and American College of Physicians meetings; a decade ago, the words *artificial intelligence* were never heard at such conferences. And at medical computing meetings, such as the annual Symposium on Computer Applications in Medical Care (SCAMC) and the international MEDINFO conferences, the growing interest in AI and the rapid increase in papers on AI and expert systems are further testimony to the impact that the field is having.

AI is beginning to have a similar effect on medical education. Such diverse organizations as the National Library of Medicine, the American College of Physicians, the Association of American Medical Colleges, and the Medical Library Association have all called for sweeping changes in medical education, increased educational use of computing technology, enhanced research in medical computer science, and career development for people working at the interface between medicine and computing. They all cite evolving computing technology and (SUMEX-AIM) AI research as key motivators. At Stanford, we have a vigorous special graduate program in Medical Information Sciences for student training and research in AI. This program has many more applicants than available slots. Demand for these graduates, in both academic and industrial settings, is so high that students typically begin to receive solicitations one or two years before completing their degrees.

III.A.1.3 Summary of Current Resource Goals

The following outlines the specific objectives of the SUMEX-AIM resource during the current five-year award period begun in August 1986. It provides an overall research plan for the resource and the backdrop against which specific progress is reported. Note that these objectives cover only the resource nucleus; objectives for individual collaborating projects are discussed in their respective reports in Section IV. Specific aims are broken into five categories: 1) Technological Research and Development, 2) Collaborative

¹ Feigenbaum, E. A., McCorduck, P., and Nii, H. P. *The Rise of the Expert Company: How Visionary Companies are Using Artificial Intelligence to Achieve Higher Productivity and Profits*. Times Books, New York, NY, 1988.

Winston, P. H., and Prendergast, K. A. *The AI Business: Commercial Uses of Artificial Intelligence*. The MIT Press, Cambridge, MA, 1984.

Research, 3) Service and Resource Operations, 4) Training and Education, and 5) Dissemination.

Technological Research and Development

- NIH-based SUMEX funding and computational support for core research is complementary to similar funding from other agencies (including DARPA, NASA, NSF, NLM, private foundations, and industry) 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:
- Basic research on AI techniques applicable to biomedical problems. Over the next term we will emphasize work on very large multi-use knowledge bases, blackboard problem-solving frameworks and architectures, knowledge acquisition or learning, constraint satisfaction, and qualitative simulation.
- Investigate methodologies for disseminating application systems such as clinical decision-making advisors into user groups. This will include generalized systems for acquiring, representing and reasoning about complex treatment protocols such as are used in cancer chemotherapy and which might be used for clinical trials in other domains.
- Support community efforts to organize and generalize AI tools and architectures that have been developed in the context of individual application projects. This will include retrospective evaluations of systems like the AGE blackboard experiment and work on new systems such as BB1, CARE, EONCOCIN, EOPAL, Meta-ONYX, and architectures for concurrent symbolic computing. 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.
- Develop more effective workstation systems to serve as the basis for research, biomedical application development, and dissemination. We seek to coordinate basic research, application work, and system development so that the AI software we develop for the next 5-10 years will be appropriate to the hardware and system software environments we expect to be practical by then. Our purchases of new hardware will be limited to experimentation with state-of-the-art workstations as they become available for our system developments.

Collaborative Research

- Encourage the exploration of 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 consider exceptional applications justified by some other unique feature of SUMEX-AIM essential for important biomedical research. We

will continue to exploit community expertise and sharing in software development.

- Minimize administrative barriers to the community-oriented goals of SUMEX-AIM and direct our resources toward purely scientific goals. We will retain the current user funding arrangements for projects working on SUMEX facilities. User projects will fund their own manpower and local needs; actively contribute their special expertise to the SUMEX-AIM community; and receive an allocation of system resources under the control of the AIM management committees. We will progressively charge core SUMEX-AIM operations costs to Stanford users as DRR support for the central system (initially a DEC 2060) is phased out. Fees to national users will be delayed as long as financially possible.
- Provide effective and geographically accessible communication facilities to the SUMEX-AIM community for remote collaborations, communications among distributed computing nodes, and experimental testing of AI programs. We will retain the current ARPANET and TELENET connections for at least the initial term and will actively explore other advantageous connections to new communications networks and to dedicated links.

Service and Resource Operations

SUMEX-AIM does not have the computing or manpower capacity to provide routine service to the large community of mature projects that has developed over the years. Rather, their computing needs are better met by the appropriate development of their own computing resources when justified. Thus, SUMEX-AIM has the primary focus of assisting new start-up or pilot projects in biomedical AI applications in addition to its core research in the setting of a sizable number of collaborative projects. We do offer continuing support, when appropriate, for projects through the lengthy process of obtaining funding to establish their own computing base.

Training and Education

- Provide documentation and assistance to interface users to resource facilities and systems.
- Exploit particular areas of expertise within the community for assisting in the development of pilot efforts in new application areas.
- Accept visitors in Stanford research groups within limits of manpower, space, and computing resources.
- Support the Medical Information Science and other student programs at Stanford to increase the number of research personnel available to work on biomedical AI applications.
- Support workshop activities including collaboration with other community groups on the AIM community workshop and with individual projects for

more specialized workshops covering specific research, application, or system dissemination topics.

Dissemination

While collaborating projects are responsible for the development and dissemination of their own AI systems and results, the SUMEX resource will work to provide community-wide support for dissemination efforts in areas such as:

- Encourage, contribute to, and support the on-going export of software systems and tools within the AIM community and for commercial development.
- Assist in the production of video tapes and films depicting aspects of AIM community research.
- Promote the publication of books, review papers, and basic research articles on all aspects of SUMEX-AIM research.

III.A.2 Details of Technical Progress

This section gives an overview of progress for the nucleus of the SUMEX-AIM resource. A more detailed discussion of our progress in specific areas and related plans for further work are presented beginning in section III.A.2.2. Objectives and progress for individual collaborating projects are discussed in their respective reports in section IV. These collaborative 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 evolved in the past. Collaborating projects are autonomous in their management and provide their own manpower and expertise for the development and dissemination of their AI programs.

III.A.2.1 Key Areas of Progress

In this section we summarize highlights of SUMEX-AIM resource activities over the past year (May 1988 - April 1989), focusing on the resource nucleus. We have made continued significant progress in all of our areas of core research, including the ONCOCIN research on dissemination of clinical trial management tools, basic AI research, and distributed systems development.

Core ONCOCIN Research

- Our work has proceeded well along three main lines of research: 1) ONCOCIN, the therapy planning program and its graphical interface; 2) OPAL, a graphical knowledge entry system for ONCOCIN; and 3) ONYX, a strategic planning program designed to give advice in complex therapy situations. Each of these research components has in turn split into two parts: continued development of the cancer therapy versions of the system, and a generalization of each of the components for use in other areas of medicine (the prefix "E-" is added to the program names for the generalized versions). In addition, we have continued development of a generalized knowledge acquisition tool, named PROTÉGÉ, designed to encode descriptions of clinical trials. The system was the Ph.D. thesis work of Mark Musen, (who joined our faculty this year). The output of PROTÉGÉ is an OPAL-like input system designed for a target clinical area such as hypertension.
- Based on the success of our earlier ONCOCIN work, strong interest has developed, from such diverse quarters as the National Cancer Institute and the Stanford Hospital, for developing a fully operational version of ONCOCIN that can be broadly used in oncology clinics outside our research laboratory. This past year, the Stanford Hospital started a program to assist in the transfer of innovative medical technology out of the laboratory to patient care, committing approximately \$750K per year to seed this effort. ONCOCIN was selected as one of 10 projects to be funded from a large group of competing proposals. This presents a dilemma for the project that is still unresolved, namely, how to maintain a cohesiveness between ongoing research work to extend the various parts of ONCOCIN and generalize it for applications to other domains and at

the same time, meet the operational needs of a widely disseminated practical system. Much thought has gone into this problem this past year, including issues such as which of the modern workstation alternatives to select (Lisp machine, IBM PC, Apple Macintosh, SUN or NeXT UNIX workstation, ...), what language to pick (C, Lisp, ...), and can the research and operational systems really be consistent versions of a single system? In order to understand the scope and practical issues involved, we have begun an experiment to port ONCOCIN to a TI microExplorer running inside of a Mac II during the last six months. We have completed the translation of the Ozone object-oriented system, the temporal network and most of the reasoner. We will next approach the design of the user interface, which must be rewritten anew, since the current interface depends heavily on the graphical capabilities of the Xerox workstations. We are also starting a study of the overall design and specification of an "integrated" oncologist's workstation, under NCI sponsorship, that will lead to an attempt to coordinate federal, academic, and industrial efforts to implement such a system.

- Our E-ONCOCIN research has concentrated on understanding how protocols in medicine vary across subspecialties. We are examining several application areas: the intensive care unit, insulin treatment for diabetes, hypertension protocols, and both standard and complex cancer treatment problems. The diagnosis and therapy selection for patients in the intensive care unit is a natural application area because it is based on changing data and the need to determine the response to therapy interventions. We also felt that the area of insulin treatment for diabetes would be a good area to explore. Like cancer chemotherapy, the treatments for diabetes continues over a long period of time and has been the area of intensive protocol development. Unlike cancer chemotherapy, the treatment plan must handle multiple treatments in one day and deemphasises the use of multiple drugs (although there are a variety of types of insulin). Our initial experiments have shown that many of the elements of the ONCOCIN design are sufficiently general for other application areas, but that some specific elements (particularly the representation of temporal events) will have to be redesigned or extended. Another extension is to modify the framework so that it can work with established data base tools instead of the hand-tailored data base currently in use. In this work, we must be able to describe the changing clinical context and event intervals that show up in many diverse application areas. An example of a new area that we are exploring is the treatment of AIDS patients on clinical protocols. AIDS patients do not always follow the type of strict temporal schedules (e.g., regular visits to outpatient clinics) seen with oncology patients. They have a chronic disease with acute exacerbations of opportunistic infections. Furthermore, the medication schedule is interrupted by frequent hospitalizations and confounded by taking drugs not on the protocol.

Together, these factors will require a much more flexible model of the temporal dimension of treatment planning.

- We continued development of the OPAL system for graphical knowledge acquisition to facilitate protocol definition and knowledge base entry for the ONCOCIN oncology application area. A major accomplishment of this last year was to experimentally combine the OPAL and ONCOCIN programs into one working program, and to completely enter knowledge from OPAL using both the high level tools and lower level rule editors, but without needing to make changes at the ONCOCIN side of the system. Our experiments with OPAL, and our intention to generalize OPAL use outside of oncology protocols, suggests that we reorganize the OPAL program to use a relational data base to store its knowledge. We continue to explore the appropriate avenue for the connection of our knowledge acquisition systems to data bases, and have concentrated on the SQL query language to a relational data base using the client-server model (e.g., the physical data base may exist on a different machine than the knowledge acquisition tool — transmitting the query and the response over the network).
- With the current uncertainties in what workstation environment to use for future work, we began to explore alternative platforms for developing the interface for OPAL-like systems. We have begun experiments using HyperCard on the Mac II and Interface Builder on the NeXT machine. In order to build experience with the each of these possible platforms, we have re-implemented portions of OPAL system, and are analyzing the results. It is particularly hard to determine the best platform since the NeXT machine software is still in a rudimentary stage, and HyperCard on the Mac II has significant limitations including small "card size" and the inability to display multiple cards simultaneously.

We continue to work on the integration of speech-recognition technology into the interface to ONCOCIN. The project uses a commercially available continuous speech recognition product and a prototype ONCOCIN adaptation. The system uses the location of the cursor on the screen to provide a context for choosing candidate grammars with which to attempt recognition of a user's utterance. The system dynamically re-orders the list of candidate recognition grammars based on the dialog history. Albeit with limitations on the legal grammars, it is now possible to carry on most of the ONCOCIN data acquisition steps using speech alone or speech plus pointing with the mouse. We are also exploring a second medical record-keeping task — the creation of portions of a progress note that describes in textual form the changes in the patients status from week to week. We have also mounted the CMU SPHINX speech understanding system on our NeXT machines and are comparing its performance against the SSI hardware-based system we have been using.

Core AI Research

- In the last year, research has progressed on several fundamental issues of AI. As in the past, our research methodology is experimental, concentrating on building and analyzing actual systems. We have continued to explore the design and use of very large, multi-use knowledge bases with the hypothesis that both the problems of brittleness and over-specialization in current knowledge-based systems can be overcome. Some of the key directions for this work include knowledge representation, knowledge compilation, knowledge justification, model-based reasoning, and case-based reasoning. During the past year we have been exploring a variety of representations and the systems which employ them, including CYC from MCC, CLASS from Schlumberger, and QPE from Univ. of Illinois. In the study of knowledge compilation techniques, we note that effective problem solving is not typically carried out at the level of first principles, but rather at the level of more compact, efficient forms of knowledge, compiled from experience with specific tasks. We are developing an integrated scheme for using "first principles" knowledge of the physical world for simulation. Given a description of the structure of a device in terms of its constituent objects and their relations, the system identifies applicable physical laws, processes, types of matter, etc. and produces a set of equations to describe the behavior of the device. The equation model is then analyzed using the method of causal ordering to produce a model that reveals the dependency relations among the parameters of the model.
- Research has also progressed on our study of blackboard frameworks, especially as they relate to adaptive intelligent systems. Important questions for this work include: how can we design flexible control structures for powerful problem solving programs? How can we use these structures effectively in many problem domains? How can we represent processes and reason about their behavior, and perform intelligent actions under real-time requirements? This past year, we have begun or continued work on five domain-independent BB1 modules: the Focus module (provides a dynamic focus of attention); the ReAct module (provides time-sensitive problem detection and response capabilities); the ICE module (provides reasoning from first principles to handle complex or unfamiliar problems); the TPlan module (provides time-sensitive planning of coherent courses of action); and the TDB module (provides a temporally organized database of observed, expected, and intended models of external entities, and associated temporal reasoning functions).
- We have built upon earlier results in our parallel symbolic computing architectures project, including the SIMPLE CAD (Computer Aided Design) system for hierarchical, multiple level specification of computer architectures and the CARE parameterized, multiprocessor array emulator (specified in SIMPLE's specification languages and running on SIMPLE's simulator). These systems are in use by several research

groups at Stanford and have been ported to several external sites, including NASA Ames Research Center. A videotaped tutorial was held in June, 1988, attended by representatives from industry and government, which described the CARE/SIMPLE system, as well as the LAMINA programming interface. The attendees received instruction in use of the system for making measurements of the performance of various simulated multiprocessor applications. Due to rapidly growing interest in the SIMPLE/CARE system, a major effort is now underway to port it to wider class of hardware platforms. The system is currently being reimplemented in Common Lisp and the X window system, with Sun workstation as the initial target. During the past year, the research effort associated with SIMPLE/CARE has largely focussed on investigations of communication protocols and techniques for monitoring concurrent object-based applications. In other areas of our parallel architectures work, we have studied the measured speed-up of two different expert system applications, ELINT (a system for interpreting electronic intelligence signals) and AIRTRAC (a system for identifying and tracking aircraft based on diverse radar data). Our preliminary conclusions are that for relatively simple and well-structured applications such as ELINT, two (or possibly more) orders of magnitude speedup via parallel execution are possible. However, for complex and ill-structured applications such as AIRTRAC Path Association, speedup over a well-tuned serial program by using parallel execution is probably limited, at best, to an order of magnitude. Experiments are continuing to verify this preliminary conclusion.

The machine learning work has focussed this past year on explanation-based generalization and chunking work in the SOAR framework and inductive rule learning. This area of research is winding down due to the departures of Profs. Buchanan and Rosenbloom. During the past year finishing students extended the RL induction program to learn incrementally, that is from small sets of examples presented in sequence without benefit of looking at them all together. A front-end program was written to assist in the definition of RL's starting knowledge, the so-called "half-order theory". In our SOAR research, we completed a set of experiments evaluating a representational restriction on productions that guarantees an absence of expensive chunks, with encouraging results. We have applied our domain-independent abstraction mechanism to a set of problems in two domains (mobile robot and computer configuration), and evaluated its ability to reduce problem solving time, reduce learning time, and increase the generality of the rules learned. We have run a set of experiments which evaluate the ability of rules learned in medical diagnosis to transfer to related problems (done in a reduced-size version of NEOMYCIN-SOAR). In the area of theoretical developments and system building, we have extended our work on declarative learning to allow indexing off of arbitrary features, but in the process uncovered a new issue concerned with how to deal with multiple retrieval and discrimination.

Core System Development

- Because of budget cuts in our award, this has been a particularly busy and chaotic year in terms of changes to the orderly progression we had planned for the transition to a distributed environment. There were two immediate consequences of this cut: a) reducing our systems staff by two people and b) taking the DEC 2060 off of contract maintenance early in the grant year, thereby forcing us to close it down for routine use. These steps have had substantial impacts in forcing us to devote full energy to the 2060-to-SUN-4 transition mechanics approximately a year before we expected to be ready for it and in diverting staff from work on longer-term distributed computing problems. In spite of all this unplanned redirection of our energies, we have made substantial progress this past year as summarized below.
- Because of the necessary preoccupation of most of our staff with the premature 2060 transition this past year, we were not able to convene the visiting advisory group as was recommended by BRTP to help guide our long-term research efforts. As we finally close out the 2060 chapter this summer, we will plan to assemble such a group in the early fall (September or October) to reassess our plans for the coming two years.
- As detailed in our report last year, we have chosen Apple Macintosh II workstations as the general computing environment for researchers and staff, TI Explorer Lisp machines (including the microExplorer Macintosh coprocessor) as the near-term high-performance Lisp research environment, and a SUN-4 as the central network server replacement for the DEC 2060. We outlined there the many tasks facing us in making the transition from the central 2060 environment to the new distributed model, including selecting and integrating tools for text processing (editing, graphics, formatting, and bibliographic references), presentation graphics, printing, help facilities and distributed information access, interpersonal communication tools (EMail and BBoards), file management (storage, access, backup, and archiving), and system building tools (languages, development environments, and integration tools). Because of the high maintenance cost of the DEC 2060, we could not afford to continue its coverage in light of the large budget cut. Since this old-technology machine quickly becomes unreliable without regular maintenance, this forced us to transfer nearly all of our AIM community usage to the SUN-4/280 in October and November of 1988. The DECSystem-20 had been our major computer resource since February of 1983 and this machine, in turn, had replaced a KI-TENEX system in use for nine years earlier. Thus, our conversion to the UNIX based SUN-4/280 represented a major departure from a long-established approach to computing and for many, converting to the use of UNIX was a difficult transition. A significant and urgent effort went into developing a *UNIX Users Guide for TOPS-20 Users* which has provided substantial help in navigating through the most common of commands. In the process of

converting, we had to transfer about 400 user accounts. Most of the immediately-needed working files from the 2060 system were dumped to tape and loaded into the SUN-4. Most of this transfer was done during a four week period of intensive work. In addition, we had to orchestrate the transfer of "SUMEX-AIM" name from the 2060 to the SUN-4, and provide effective continuation of facilities such as EMail, BBoard, text processing, etc. Continuous and nearly compatible AIM community mail services were maintained through the transition by installing the Columbia University MM-C mail program on the SUN-4. This program closely duplicates the functions of the TOPS-20 COMAND JSYS under UNIX and presents the user with a mail reader/composer interface very similar to that of TOPS-20 MM. This system, coupled with a UNIX version of the EMACS text editor, called GNUEMACS, provided a relatively familiar setting for the most common computing functions used by AIM community members. In the succeeding months, we added bulletin board functionality to MM-C so that, from the user's perspective, mail access was nearly identical to the former system.

- Another major issue in the 2060-to-SUN-4 transition has been the need to provide our users with continued access to their large collection of archived files and to a set of permanent annual backup dumps (done January of each year) which have been collected and maintained since 1975. This has required very careful planning as the directory information for these tapes resides in Archive-Directory files (for TENEX) and the on-line File Descriptor Blocks (FDB's) of the TOPS-20 file system. These two sets of information must be converted to simple UNIX-compatible text files to provide users with continued facilities to review and access their collections of archived files. This work has been a major undertaking and is still in progress.
- In the move from TOPS-20 to UNIX, we have had to ensure continued access to "standard" services, such as file backup, archiving, a flexible and intuitive naming facility, and data interchange services (e.g., file transfer). UNIX has many of the needed facilities, e.g., backup, long names, hierarchical directory structure, some file property attributes, data conversion, and limited archival tools. We have worked on adapting a commercial system developed by UniTech to allow users to manage large file collections by moving files not needed on-line to and from off-line tape storage. This system also maintains a historical archive of files. The system is in beta test now and will be released to the entire community early this summer.
- Electronic mail continues as a primary means of communication for the widely spread SUMEX-AIM community. As reported last year, the move to workstations has forced a significant rethinking of the mechanisms employed to manage such mail in order to ensure reliable access, to make user addressing understandable and manageable, and to facilitate keeping the mail software distributed to workstations as simple, stable, and

maintainable as possible. We are following a strategy of having a shared mail server machine which handles mail transactions with mail clients running on individual user workstations. The mail server can be used from clients at arbitrary locations, allowing users to read mail across campus, town, or country. We have made significant progress this year in developing a Mac II version of the graphics-based MM-D/IMAP mail client reported on last year, including a complete rewrite of the InterLisp system into C and modifying the user reading and composing interface to be compatible with the Mac "look and feel". This system is nearly ready for alpha test starting early this summer.

One of the key issues in selecting the systems for our distributed computing environment was the performance of Common Lisp and to help make this evaluation, we have continued to expand an informal survey of the performance of two KSL AI software packages, SOAR and BB1, on a wide variety of machines. This study was completed this past year and a "final" report written (see Appendix B), recognizing that each month new workstations are announced that deserve additional evaluation.. Within a factor of two of the best performance, a considerable range of workstations based on stock microprocessor chips as well as specially microprogrammed Lisp chips have comparable performance. Even though performance gaps between microprogrammed Lisp systems and stock workstation implementations are narrowing, there still remains a significant difference in the quality of the development environments. We have attempted to distill the key features of the Lisp machine environments that would be needed in stock machine implementations in order to make them attractive in a development setting.

- This past year we acquired 2 NeXT workstations, primarily to understand and evaluate the power of the NeXT *Interface Builder* for AI software development. Integrating these prototype systems took a significant effort and they are now being used in several of our core research and applications projects. In addition, we have continued to support a limited number of other "standard" workstations for our work, including Mac II's, TI Explorers, and SUN's. We have continued to work toward a complete phase-out of our old Xerox and Symbolics Lisp machines.
- As reported last year, major changes in ARPANET service have been underway as ARPA has responded to its own budget pressures to reduce operating subsidy of the ARPANET. Starting late last spring, sections of the ARPANET serving university users were being shut down and replaced by connection to the NSFNET. Our own connection is just in the process of being decommissioned with our IMP scheduled to be removed sometime this summer. Our Internet access is now implemented through the Bay Area Regional Research Network (BARRNet) and the NSFNet. We continue to operate the Develcon gateway between our Ethernet environment and the TELENET network by which many AIM users gain access to SUMEX.

Other Resource Activities

- We have continued the dissemination of SUMEX-AIM technology through various media. The distribution system for our AI software tools (EMYCIN, AGE, and BB1) to academic, industrial, and federal research laboratories continues to work effectively. We have also continued to distribute the video tapes of some of our research projects including ONCOCIN, and an overview tape of Knowledge Systems Laboratory work to outside groups. Our group has continued to publish actively on the results of our research, including more than 45 research papers per year in the AI literature and a dozen books in the past 7 years on various aspects of SUMEX-AIM AI research. We assisted and participated actively in the AIM Workshop sponsored by AAAI and held at Stanford in 1988 and hosted a number of AIM community visitors at our Stanford research laboratory. Members of the Medical Computer Science group are participating in the early organization phases of another workshop during the spring of 1990.
- The Medical Information Sciences program, begun at Stanford in 1983 under Professor Shortliffe as Director, has continued its strong development over the past year. The specialized curriculum offered by the MIS program focuses on the development of a new generation of researchers able to support the development of improved computer-based solutions to biomedical needs. The feasibility of this program resulted in large part from the prior work and research computing environment provided by the SUMEX-AIM resource. As already reported, it has recently received enthusiastic endorsement from the Stanford Faculty Senate for an additional five years and has been awarded renewed post-doctoral training support from the National Library of Medicine with high praise for the training and contributions of the SUMEX-AIM environment from the reviewing study section. This past year, MIS students have published many papers, including several that have won conference awards.
- We have continued to recruit new user projects and collaborators to explore further biomedical areas for applying AI. A number of these 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. At the same time we have encouraged older mature projects to build their own computing environments thereby facilitating the transition to a distributed AIM community. A substantial number of projects have already moved to their own computing resources.
- SUMEX user projects have made good progress in developing and disseminating effective consultative computer programs for biomedical research. These systems provide expertise in areas like cancer chemotherapy protocol management, clinical diagnosis and decision-

making, and molecular biology. We have worked hard to meet their needs and are grateful for their expressed appreciation (see Section IV).