context of MRS. We designed a method for representing temporal knowledge in ONCOCIN. Finally, Cooper's Ph.D. thesis on representing and using causal and probabilistic knowledge was published in this year.

[See KSL technical memos KSL-84-9, KSL-84-10, KSL-84-18, KSL 84-31, KSL-84-41, KSL-85-5.]

2. *Advanced Architectures and Control*: What kinds of software tools and system architectures can be constructed to make it easier to implement expert programs with increasing complexity and high performance? How can we design flexible control structures for powerful problem solving programs?

Much of our research in the past year has involved investigations with the Blackboard architecture begun in previous years. We have implemented our design in a working system called BB1.

[See KSL technical memos KSL-84-11, KSL-84-12, KSL-84-14, KSL 84-16, KSL 84-36.]

3. *Knowledge Acquisition*: How is knowledge acquired most efficiently -- from human experts, from observed data, from experience, and from discovery? How can a program discover inconsistencies and incompleteness in its knowledge base? How can the knowledge base be augmented without perturbing the established knowledge base?

Three Ph.D. theses (Fu, Greiner, and Dietterich) in the area of knowledge acquisition were completed in this year. Fu's work develops methods for learning by induction, where the target rules may have some associated degrees of uncertainty and may contain names of intermediate concepts. This work was demonstrated in the context of diagnosing causes of jaundice. Greiner's work examines learning by analogy. Dietterich's work elucidates methods needed in learning programs to deal with state variables and with problems of using a partially learned theory to interpret new data that will be used to learn new elements of the theory. In addition, we implemented the first parts of a program that can learn by watching an expert. And we implemented a prototype system that learns control heuristics from an expert using a problem solving program written in BB1.

[Preliminary results have been published in KSL-84-10, KSL-84-18, KSL-84-24, KSL-84-38, KSL-84-45, KSL 84-46, KSL-85-2, KSL-85-4.]

4. *Knowledge Utilization*: By what inference methods can many sources of knowledge of diverse types be made to contribute jointly and efficiently toward solutions? How can knowledge be used intelligently, especially in systems with large knowledge bases, so that it is applied in an appropriate manner at the appropriate time?

We completed the design of a system using Dempster's rule of propagating uncertainty, and we examined several other issues regarding the use of probabilistic information in expert systems. Dr. Jean Gordon, a mathematician and Stanford medical student, collaborated with Dr. Shortliffe on work that examines inexact inference using the Dempster-Shafer theory of evidence, demonstrating its relevance to a familiar expert system domain, namely the bacterial organism identification problem that lies at the heart of the MYCIN system, and presenting a new adaptation of the D-S approach with both computational efficiency and permitting the management of evidential reasoning within an abstraction hierarchy.

We examined the use of counter-factual conditionals in logic-based systems and completed an analysis of how procedural hints can be used by a problem solver.

[See KSL technical memos KSL-84-11, KSL-84-17, KSL-84-21, KSL-84-30, KSL-84-31, KSL-84-35, KSL 84-41, KSL-84-42, KSL-84-42, KSL-84-43.]

5. *Software Tools*: How can specific programs that solve specific problems be generalized to more widely useful tools to aid in the development of other programs of the same class?

We have continued the development of new software tools for expert system construction and the distribution of packages that are reliable enough and documented so that other laboratories can use them. These include the old rule-based EMYCIN system, MRS, and AGE. Progress has been made in making the BB1 instantiation of the blackboard architecture domain-independent. We have begun constructing and editing subsystems and have completed a first implementation of an explanation subsystem.

[See KSL technical memos KSL-84-16, KSL-84-39.]

6. *Explanation and Tutoring*: How can the knowledge base and the line of reasoning used in solving a particular problem be explained to users? What constitutes a sufficient or an acceptable explanation for different classes of users? How can knowledge in a system be transferred effectively to students and trainees?

A program for inferring a model of users was designed and implemented in the context of a tutoring system that aids in teaching algebra. A second user-modelling program was implemented in the context of NEOMYCIN to help understand how an expert solves problems. A survey of explanation capabilities in medical consultation programs was published.

A new project on knowledge-based explanations in a decision analysis environment is getting underway as the thesis research of Dr. Glenn Rennels. This work is actually a synthesis of artificial intelligence, decision analysis and statistics. The work concerns medical management, not diagnosis; diagnostic decisions identify underlying mechanisms of the illness, and group the patient's problems under a diagnostic label, whereas management decisions plan actions that will prevent undesirable outcomes and restore health. The intelligent behavior we want to emulate is (a) the identification of studies relevant to a given clinical case, and (b) interpretation of those studies for decision-making assistance.

[See KSL technical memos KSL-84-12, KSL 84-27, KSL-84-29.]

7. *Planning and Design*: What are reasonable and effective methods for planning and design? How can symbolic knowledge be coupled with numerical constraints? How are constraints propagated in design problems?

A major paper on skeletal planning was published in this year. And we published in the biochemistry literature some results of applying skeletal planning to experiment design in genetic engineering.

[See KSL technical memos KSL-84-33, KSL-85-6.]

8. *Diagnosis*: How can we build a diagnostic system that reflects any of several diagnostic strategies? How can we use knowledge at different levels of abstraction in the diagnostic process?

Research on using causal models in a medical decision support system (NESTOR) was published in this year. Using the domain of hypercalcemic disorders, NESTOR attempts to use knowledge-based methods within a formal probability theory framework. The system is able to score hypotheses with causal knowledge guiding the application of sparse probabilistic knowledge; search for the most likely hypothesis without

exploring the entire hypothesis space; and critique and compare hypotheses which are generated by the system, volunteered by the user, or both.

A second medical diagnosis program that uses causal models of renal physiology (AI/MM) was also published. In this system, analysis and explanation of physiological function is based on two kinds of causal relations: empirical "Type-1" relations based on definitions or on repeated observation and mathematical "Type-2" relations that have a basis in physical law. Inference rules are proposed for making valid qualitative causal arguments with both kinds of causal basis.

A working implementation of the PATHFINDER system was evaluated and its diagnostic strategies were analyzed. A taxonomy of diagnostic methods was completed and integrated into the NEOMYCIN system.

[See KSL technical reports: KSL-84-13, KSL-84-19, KSL-84-48, KSL-85-5.]

Relevant Core Research Publications

HPP 84-9   David H. Hickam, Edward H. Shortliffe, Miriam B. Bischoff, A. Carlisle Scott, and Charlotte D. Jacobs; *Evaluations of the ONCOCIN System: A Computer-Based Treatment Consultant for Clinical Oncology, (1) The Quality of Computer-Generated Advice and (2) Improvements in the Quality of Data Management*, May 1984.

HPP 84-10   Thomas G. Dietterich; *Learning About Systems That Contain State Variables*, June 1984. In *Proceedings of AAAI-84*, August 1984.

HPP 84-11   M. Genesereth, and D.E. Smith; *Procedural Hints in the Control of Reasoning*, May 1984.

HPP 84-12   Derek H. Sleeman; *UMFE: A User Modelling Front End Subsystem*, April 1984.

HPP 84-13   Eric J. Horvitz, David E. Heckerman, Bharat N. Nathwani, and Lawrence M. Fagan; *Diagnostic Strategies in the Hypothesis-Directed PATHFINDER System*, June 1984, submitted to the *First Conference on Artificial Intelligence Applications, Denver, CO.*, December 5-7, 1984.

HPP 84-14   Vineet Singh, and M. Genesereth; *A Variable Supply Model for Distributing Deductions*, May 1984.

HPP 84-15   Bruce G. Buchanan; *Expert Systems*, July 1984, *Journal of Automated Reasoning, Vol. 1, No. 1, Fall*, 1984.

HPP 84-16   *STAN-CS-84-1034* Barbara Hayes-Roth; *BB-1: An Architecture for Blackboard Systems That Control, Explain, and Learn About Their Own Behavior*, December 1984.

HPP 84-17   M.L. Ginsberg; *Analyzing Incomplete Information*, 1984.

HPP 84-18   William J. Clancey; *Knowledge Acquisition for Classification Expert Systems*, July 1984, *Proceedings of ACM-84*, 1984.

HPP 84-19   E.H. Shortliffe; *Coming to Terms With the Computer*, to appear in S.R. Reiser, and M. Anbar (eds.), *The Machine at the Bedside: Strategies for Using Technology in Patient Care*, Cambridge University Press, 1984.

HPP 84-20	E.H. Shortliffe; *Artificial Intelligence and the Future of Medical Computing*, in *Proceedings of a Symposium on Computers in Medicine, annual meeting of the California Medical Association, Anaheim, CA.*, February 1984.
HPP 84-21	E.H. Shortliffe; *Reasoning Methods in Medical Consultation Systems: Artificial Intelligence Approaches (Tutorial)*, in *Computer Programs in Biomedicine* January 1984.
HPP 84-22	*ONCOCIN Project: Studies to Evaluate the ONCOCIN System; 6 Abstracts*, February 1984.
HPP 84-23	Edward H. Shortliffe; *Feature Interview: On the MYCIN Expert System*, in *Computer Compacts, 1:283-289*, December 1983/January 1984.
HPP 84-24	B.G. Buchanan, and E.H. Shortliffe; *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, published with *Addison-Wesley, Reading, MA.*, 1984.
HPP 84-25	W.J. Clancey, and E.H. Shortliffe; *Readings in Medical Artificial Intelligence: The First Decade*, published with *Addison-Wesley, Reading, MA.*, 1984.
HPP 84-27	Edward H. Shortliffe; *Explanation Capabilities for Medical Consultation Systems (Tutorial)*, in D. Lindberg, and M. Collen (eds.), *Proceedings of AAMSI Congress 84, pp. 193-197, San Francisco*, May 21-23, 1984.
HPP 84-28	E.H. Shortliffe, and L.M. Fagan; *Artificial Intelligence: The Expert Systems Approach to Medical Consultation*, in *Proceedings of the 6th Annual International Symposium on Computers in Critical Care and Pulmonary Medicine, Heidelberg, Germany*, June 4-7, 1984.
HPP 84-29	David C. Wilkins, Bruce G. Buchanan, and William J. Clancey; *Inferring an Expert's Reasoning by Watching, Proceedings of the 1984 Conference on Intelligent Systems and Machines*, 1984.
HPP 84-30	M.L. Ginsberg; *Non-Monotonic Reasoning Using Dempster's Rule*, June 1984.
HPP 84-31	M.L. Ginsberg; *Implementing Probabilistic Reasoning*, June 1984.
HPP 84-32	Bruce G. Buchanan; *Artificial Intelligence: Toward Machines That Think*, July 1984, in *Yearbook of Science and the Future, pp. 96-112, Encyclopedia Britannica, Inc., Chicago*, 1985.
HPP 84-33	Rene Bach, Yumi Iwasaki, and Peter Friedland; *Intelligent Computational Assistance for Experiment Design*, in *Nuclear Acids Research*, January 1984.
MCS Thesis	Kunz, John C.; *Use of Artificial Intelligence and Simple Mathematics to Analyze a Physiological Model*, Doctoral dissertation, Medical Information Sciences, June 1984.
HPP 84-35	Jean Gordon, and Edward Shortliffe; *A Method for Managing Evidential Reasoning in a Hierarchical Hypothesis Space*, September 1984 and in *Artificial Intelligence, 26(3)*, July 1985.
HPP 84-36	Michael R. Genesereth, Matt Ginsberg, and Jeff S. Rosenschein; *Cooperation Without Communication*, September 1984.

HPP 84-38   Li-Min Fu, and Bruce G. Buchanan; *Enhancing Performance of Expert Systems by Automated Discovery of Meta-Rules*, September 6, 1984.

HPP 84-39   Paul S. Rosenbloom, John E. Laird, John McDermott, Allen Newell, and Edmund Orciuch; *R1-Soar: An Experiment in Knowledge-Intensive Programming in a Problem-Solving Architecture*, to appear in the *Proceedings of the IEEE Workshop on Principles of Knowledge-Based Systems*, October 1984.

HPP 84-41   *STAN-CS-84-1032* Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosenschein; *Solving the Prisoner's Dilemma*, November 1984.

HPP 84-42   Matthew L. Ginsberg; *Does Probability Have a Place in Non-Monotonic Reasoning?* submitted to the *IJCAI-85*, November 1984.

HPP 84-43   *STAN-CS-84-1029* Matthew L. Ginsberg; *Counterfactuals*, submitted to the *IJCAI-85*, December 1984.

HPP 84-45   Devika Subramanian, and Michael R. Genesereth; *Experiment Generation with Version Spaces*, December 1984.

HPP 84-46   Thomas G. Dietterich; *Constraint Propagation Techniques for Theory-Driven Data Interpretation*, PhD Thesis, to be published as a book by *Kluwer*, December 1984.

HPP 84-48   *STAN-CS-84-1031* Gregory F. Cooper; *NESTOR: A Computer-Based Medical Diagnostic Aid That Integrates Causal and Probabilistic Knowledge*, PhD Thesis, December 20, 1984.

KSL 85-2   *STAN-CS-85-1036* Barbara Hayes-Roth, and Michael Hewett; *Learning Control Heuristics in BB1*, submitted to the *IJCAI-85*, January 1985.

KSL 85-4   (Needs Authors Permission) Li-Min Fu, and Bruce G. Buchanan; *Learning Intermediate Knowledge in Constructing a Hierarchical Knowledge Base*, submitted to the *IJCAI Conference Proceedings for 1985*, January 1985.

KSL 85-5   (Needs Authors Permission) William J. Clancey; *Heuristic Classification*, March 1985.

KSL 85-6   Peter E. Friedland, and Yumi Iwasaki; *The Concept and Implementation of Skeletal Plans*, published in the *Journal of Automated Reasoning*, 1985.

KSL 85-7   Rene Bach, Yumi Iwasaki, and Peter Friedland; *Intelligent Computational Assistance for Experiment Design*, published in *Nucleic Acids Research*, 1985.

KSL 85-8   (Needs Authors Permission) M.G. Kahn, J. Ferguson, E.H. Shortliffe, and L. Fagan; *An Approach for Structuring Temporal Information in the ONCOCIN System*, March 1985.

Summary of Core Research Funding Support

We are pursuing a broad core research program on basic AI research issues with support from not only SUMEX but also DARPA, NASA, NSF, and ONR. SUMEX provides

some salary support for staff and students involved in core research and invaluable computing support for most of these efforts. Additional salary support comes from the **sources shown starting on page 36.**

Interactions with the SUMEX-AIM Resource

Our interactions with the SUMEX-AIM resource involve the facilities -- both hardware and software -- and the staff -- both technical and administrative. Taken together as a whole resource, they constitute an essential part of the research structure for the KSL. Many of the grants and contracts from other agencies have been awarded partly because of the cost-effectiveness of AI research in the KSL due to the fact that much of our computing needs could be more than adequately met by the SUMEX-AIM resource. In this way the complementary funding of this work by the NIH and other agencies provides a high leverage for incremental investment in AI research at the SUMEX-AIM resource.

We rely on the central SUMEX facility as a focal point for all the research within the KSL, not only for much of our computing, but for communications and links to our many collaborators as well. As a common communications medium alone, it has significantly enhanced the nature of our work and the reach of our collaborations. The existence of the central time-shared facility has allowed us to explore new ideas at very small incremental cost.

As SUMEX and the KSL acquire a diversity of hardware, including LISP workstations and smaller personal computers, we rely more and more heavily on the SUMEX staff for integration of these new resources into the local network system. The staff has been extremely helpful and effective in dealing with the myriad of complex technical issues and leading us competently into this world of decentralized, diversified computing. At the same time, the staff has provided a stable, efficient central time-shared machine running software that has been developed at many sites over many years. Without the dedication of the SUMEX staff, the KSL would not be at the forefront of AI research.

2.1.4.6. Dissemination Activities

Throughout the history of the SUMEX-AIM resource, we have made extensive efforts at disseminating the AI technology developed here. This has taken the form of many publications -- over 45 combined books and papers are published per year from the KSL; wide distribution of our software including systems software and AI application and tool software, both to other research laboratories and for commercial development; production of films and video tapes depicting aspects of our work; and significant project efforts at studying the dissemination of individual applications systems such as the GENET community (DNA sequence analysis software) and the ONCOCIN resource-related research project (see 209).

Books and Publications

A sampling of the recent research paper publications of the KSL was given in the previous section on core AI research progress. The following lists the major books published in the past 4 years from the KSL:

- *Heuristic Reasoning about Uncertainty: An AI Approach*, Cohen, Pitman, 1985.

- *Readings in Medical Artificial Intelligence: The First Decade*, Clancey and Shortliffe, Addison-Wesley, 1984.

- *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Buchanan and Shortliffe, Addison-Wesley, 1984.

- *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Feigenbaum and McCorduck, Addison-Wesley, 1983.

- *Building Expert Systems*, F. Hayes-Roth, Waterman, and Lenat, eds., Addison-Wesley, 1983.

- *System Aids in Constructing Consultation Programs: EMYCIN*, van Melle, UMI Research Press, 1982.

- *Knowledge-Based Systems in Artificial Intelligence: AM and TEIRESIAS*, Davis and Lenat, McGraw-Hill, 1982.

- *The Handbook of Artificial Intelligence*, Volume I, Barr and Feigenbaum, eds., 1981; Volume II, Barr and Feigenbaum, eds., 1982; Volume III, Cohen and Feigenbaum, eds., 1982; Kaufmann.

- *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project*, Lindsay, Buchanan, Feigenbaum, and Lederberg, McGraw-Hill, 1980.

Software Distribution

We have widely distributed both our system software and our AI tool software. We have no accurate records of the extent of distribution of the system codes because their distribution is not centralized and controlled. The recent programs such as the TOPS-20 file recognition enhancements, the Ethernet gateway and TIP programs, the SEAGATE AppleBus to Ethernet gateway, the PUP Leaf server, the SUMACC development system for Macintosh workstations, and our Lisp workstation programs are well-distributed throughout the ARPANET community and beyond.

We do have reasonably accurate records of the distribution of our AI tool software because the recipient community is more directly coupled to us and the distribution is centralized:

GENET
Prior to the establishment of the BIONET resource at IntelliCorp, we distributed 21 copies of the DNA sequence analysis programs and databases for both DEC-10 and DEC-20 systems.

EMYCIN
A total of 56 sites have received the EMYCIN [6, 68] package for backward-chained, rule-based AI systems.

AGE
The AGE [54] blackboard framework system has been sent out to 35 sites in versions for several machines.

MRS
The MRS [20] logic-based system for meta-level representation and reasoning has been provided to 76 sites.

Other Programs
Smaller numbers of copies of programs such as the SACON [2] knowledge base for EMYCIN, the GLISP [57] system (now distributed by Gordon Novak at the University of Texas), and the new BB1 [28, 27] system have been distributed.

A number of other software packages have been licensed or otherwise made available for commercial development including DENDRAL (Molecular Designs), MAINSAIL (Xidak), UNITS (IntelliCorp), and EMYCIN (Teknowledge and Texas Instruments).

Video Tapes and Films

The KSL and the ONCOCIN project have prepared several video tapes that provide an overview of the research and research methodologies underlying our work and that demonstrate the capabilities of particular systems. These tapes are available through our groups, the Fleischmann Learning Center at the Stanford Medical Center, and the Stanford Computer Forum and copies have been mailed to program offices of our various funding sponsors. The three tapes include:

- *Knowledge Engineering in the Heuristic Programming Project* -- This 20-minute film/tape illustrates key ideas in knowledge-based system design and implementation, using examples from ONCOCIN, PROTEAN, and knowledge-based VLSI design systems. It describes the research environment of the KSL and lays out the methodologies of our work and the long term research goals that guide it.

- *ONCOCIN Overview* -- This is a 30-minute tape providing an overview of the ONCOCIN project. It gives an historical context for the work, discusses the clinical problem and the setting in which the prototype system is being used, and outlines the plans for transferring the system to run on single-user workstations. Brief illustrations of the graphics capabilities of ONCOCIN on a Lisp workstation are also provided.

- *ONCOCIN Demonstration* -- This 1-hour tape provides detailed examples of the key components of the ONCOCIN system. It begins with a demonstration of the prototype system's performance on a time-shared mainframe computer and then shows each of the elements involved in transferring the system to Lisp workstations.

The GENET Dissemination Experiment

Beginning in early 1980, the MOLGEN project investigators at Stanford have made a new set of computing tools available to a national community of molecular biologists through a guest facility called GENET on the SUMEX-AIM resource. This experimental subcommunity was started to broaden MOLGEN's base of scientist collaborators at institutions other than Stanford and to explore the idea of a SUMEX-like resource to disseminate sophisticated software tools to a generally computer-naive community. The enthusiastic response to the very limited announcement of this facility eventually necessitated SUMEX placing severe restrictions on the scope of services provided to this community.

Three main programs were offered to assist molecular genetics users: SEQ, a DNA-RNA sequence analysis program; MAP, a program that assists in the construction of restriction maps from restriction enzyme digest data; and MAPPER, a simplified and somewhat more efficient version of the MOLGEN MAP program, written and maintained by William Pearson of Johns Hopkins University. Some of the other, more-sophisticated programs being developed through MOLGEN research efforts were not yet available for novice users. However, GENET users had access to the SUMEX-AIM programs for electronic messaging, text-editing, file-searching, etc.

The GENET experiment proved so successful that eventually that community was the single biggest consumer of processor cycles on SUMEX. This overload diverted our very limited computing resources away from our mainline goal of supporting projects developing new AI systems in the medical and biological sciences, including molecular biology. Efforts to secure funds to increase SUMEX capacity for the burgeoning GENET use failed. Thus, without any fair way to allocate a small resource to the growing GENET community and in order to restore the necessary emphasis on biomedical computer science research on SUMEX, it was necessary to phase out the GENET usage. We closed the GENET account at the end of 1982, with a mandate from an ad hoc GENET Executive Committee, and phased out all usage by spring of 1983. In the process, we developed procedures by which academic users could obtain their own copies of the GENET programs used at SUMEX and we provided a list of alternate sources for GENET-like computing services. As indicated above, SUMEX has supplied 21 systems to academic users with compatible machines.

Since the phase-out of GENET at SUMEX, IntelliCorp, a commercial AI company, submitted a proposal to the NIH Division of Research Resources for a BIONET resource and was successful in obtaining funding. The BIONET resource began operation in the summer of 1984.

2.1.4.7. Training Activities

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. The resource staff has spent significant effort in assisting users gain access to the system and use it effectively. We have also spent substantial effort to develop, maintain, and facilitate access to documentation and interactive help facilities. The HELP and Bulletin Board subsystems have been important in this effort to help users get familiar with the computing environment.

On another front, we have regularly accepted a number of scientific visitors for periods of several months to a year, to work with us to learn the techniques of expert system definition and building and to collaborate with us on specific projects. Our ability to accommodate such visitors is severely limited by space, computing, and manpower resources to support such visitors within the demands of our on-going research.

And finally, the training of graduate students is an essential part of the research and educational activities of the KSL. Currently 41 students are working with our projects centered in Computer Science and another 20 students are working with the Medical Computer Science program in Medicine. Of the 41 working in Computer Science, 25 are working toward Ph.D. degrees, and 16 are working toward M.S. degrees. A number of students are pursuing interdisciplinary programs and come from the Departments of Engineering, Mathematics, Education, and Medicine.

Based on the SUMEX-AIM community environment, we have initiated two unique and special academic degree programs at Stanford, the Medical Information Science program and the Masters of Science in AI, to increase the number of students we produce for research and industry, who are knowledgeable about knowledge-based system techniques.

The *Medical Information Sciences (MIS)* program is one of the most obvious signs of the local academic impact of the SUMEX-AIM resource. The MIS program received recent University approval (in October 1982) as an innovative training program that offers MS and PhD degrees to individuals with a career commitment to applying computers and decision sciences in the field of medicine. The MIS training program is based in School of Medicine, directed by Dr. Shortliffe, co-directed by Dr. Fagan, and overseen by a group of nine University faculty that includes several faculty from the Knowledge Systems Laboratory (Profs. Shortliffe, Feigenbaum, Buchanan, and Genesereth). It was Stanford's active ongoing research in medical computer science, plus a world-wide reputation for the excellence and rigor of those research efforts, that persuaded the University that the field warranted a new academic degree program in the area. A group of faculty from the medical school and the computer science department argued that research in medical computing has historically been constrained by a lack of talented individuals who have a solid footing in both the medical and computer science fields. The specialized curriculum offered by the new program is intended to overcome the limitations of previous training options. It focusses on the development of a new generation of researchers with a commitment to developing new knowledge about optimal methods for developing practical computer-based solutions to biomedical needs.

The program accepted its first class of four trainees in the summer of 1983 and a second class of five entered last summer. A third group of seven students has just been selected to begin during 1985. The proposed steady state size for the program (which should be reached in 1986) is 20-22 trainees. Applicants to the program in our first two years have come from a number of backgrounds (including seven MD's and five medical students). We do not wish to provide too narrow a definition of what kinds of

prior training are pertinent because of the interdisciplinary nature of the field. The program has accordingly encouraged applications from any of the following:

- medical students who wish to combine MD training with formal degree work and research experience in MIS;

- physicians who wish to obtain formal MIS training after their MD or their residency, perhaps in conjunction with a clinical fellowship at Stanford Medical Center;

- recent BA or BS graduates who have decided on a career applying computer science in the medical world;

- current Stanford undergraduates who wish to extend their Stanford training an extra year in order to obtain a "co-terminus" MS in the MIS program;

- recent PhD graduates who wish post-doctoral training, perhaps with the formal MS credential, to complement their primary field of training.

In addition, a special one-year MS program is available for established academic medical researchers who may wish to augment their computing and statistical skills during a sabbatical break.

With the exception of this latter group, all students spend a minimum of two years at Stanford (four years for PhD students) and are expected to undertake significant research projects for either degree. Research opportunities abound, however, and they of course include the several Stanford AIM projects as well as research in psychological and formal statistical approaches to medical decision making, applied instrumentation, large medical databases, and a variety of other applications projects at the medical center and on the main campus. Several students are already contributing in major ways to the AIM projects and core research described in this application.

Early evidence suggests that the program already has an excellent reputation due to:

- high quality students, many of whom are beginning to publish their work in conference proceedings and refereed journals;

- a rigorous curriculum that includes newly-developed course offerings that are available to the University's medical students, undergraduates, and computer science students as well as to the program's trainees;

- excellent computing facilities combined with ample and diverse opportunities for medical computer science and medical decision science research;

- the program's great potential for a beneficial impact upon health care delivery in the highly technologic but cost-sensitive era that lies ahead.

The program has been successful in raising financial and equipment support (almost $1M in hardware gifts from Hewlett Packard, Xerox, and Texas Instruments; over $200K in cash donations from corporations and foundations; and an NIH post-doctoral training grant from the National Library of Medicine).

The *Master of Science in Computer Science: Artificial Intelligence (MS:AI)* program is a terminal professional degree offered for students who wish to develop a competence in the design of substantial knowledge-based AI applications but who do not intend to obtain a Ph.D. degree. The MS:AI program is administered by the Committee for Applied Artificial Intelligence, composed of faculty and research staff of the Computer Science Department. Normally, students spend two years in the program with their

time divided equally between course work and research. In the first year, the emphasis is on acquiring fundamental concepts and tools through course work and and project involvement. During the second year, students implement and document a substantial AI application project.

2.1.4.8. Resource Community Management

Early in the design of the SUMEX-AIM resource, an effective management plan was worked out with the Biotechnology Resources Program (now Biomedical Research Technology 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 Section 2.3.3 on page 181 of the renewal plan. It has continued to function effectively as summarized below.

- The AIM Executive Committee meets regularly by teleconference to advise on new project 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 C on page 307 for a current listing of AIM committee membership).

- We have actively recruited new application projects and disseminated information about the resource. The number of formal projects in the SUMEX-AIM community still runs at the capacity of our computing resources. With the development of more decentralized computing resources within the AIM community outside of Stanford (see below), the center of mass of our community has naturally shifted toward the growing number of Stanford applications and core research projects. We still, however, actively support new applications in the national community where these are not able to gain access to suitable computing resources on their own.

- 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.

- 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.

- 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 fully authorized and pilot projects have been encouraged to develop their own computing resources separate from SUMEX or have been phased off of SUMEX as a result and more productive collaborative ties established for others.

- We have continued to provide active support for the AIM workshops. The last one was held at Ohio State University in the summer of 1984 and the next one will be in Washington, DC, hosted by the National Library of Medicine under Drs. Lindberg and Kingsland.

- 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.

• 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.

As indicated above, we have sought to retain SUMEX resources for new projects, those exploring new areas in biomedical AI applications and those in such an early state of feasibility that they are unable to afford their own computing resources. This policy has worked effectively as seen from the following lists of terminated projects and projects now using their own computing resources at other sites:

Projects Moved All or In Part to Other Machines:

Stanford Projects:

• GENET [Brutlag, Kedes, Friedland - IntelliCorp]

National Projects:

• Acquisition of Cognitive Procedures (ACT) [Anderson - CMU]

• Chemical Synthesis [Wipke - UC Santa Cruz]

• Simulation of Cognitive Processes [Lesgold - Pittsburgh]

• PUFF [Osborne, Feigenbaum, Fagan - Pacific Medical Center]

• CADUCEUS/INTERNIST [Pople, Myers - Pittsburgh]

• Rutgers [Amarel, Kulikowski, Weiss - Rutgers]

• MDX [Chandrasekaran - Ohio State]

• SOLVER [P. Johnson - University of Minnesota]

Completed Projects Summary

Stanford Projects:

• DENDRAL [Lederberg, Djerassi, Buchanan, Feigenbaum]

• MYCIN [Shortliffe, Buchanan]

• EMYCIN [Shortliffe, Buchanan]

• CRYSALIS [Feigenbaum, Engelmore]

• MOLGEN I [Feigenbaum, Brutlag, Kedes, Friedland]

• AI Handbook [Feigenbaum, Barr, Cohen]

- AGE Development [Feigenbaum, Nii]

National Projects:

- Ventilator Management [Osborne, Feigenbaum, Fagan – Pacific Medical Center]
- Higher Mental Functions [Colby – USC]

2.2. Planned Resource Activities

We have already summarized the overall aims of the SUMEX-AIM resource for the proposed 5-year renewal period on page 64. This section gives details of our research plans in pursuit of those aims for the major areas of our resource activities -- core research and development, collaborative research, service, training and education, and dissemination. To recap the overall scope and guiding goals of our new work:

- SUMEX-AIM is a national computing resource that develops and provides advanced computing facilities and expertise to support 1) a long-term program in basic research in artificial intelligence, 2) applying AI techniques to a broad range of biomedical problems by collaborative and user projects at Stanford and other universities around the country, 3) studying and developing methodologies for disseminating AI systems into the biomedical community, 4) experimenting with communication technologies to promote scientific interchange, and 5) developing better tools and facilities to carry on this research.

- Our applications, core research, and system development will be directed toward realizing and exploiting the computing environment that will be routinely available in the late 1980's and early 1990's, based on compact, decentralized, high-performance personal workstations that take advantage of the intelligent computing environments beginning to emerge from today's Lisp workstations. Consistent with these plans, we will immediately discontinue DRR subsidy for the DEC 2020 demonstration machine and for the shared VAX 11/780 time-sharing system. Also we will gradually and responsibly phase out DRR support for the DEC 2060 mainframe system that has been our chief shared resource and link to the past.

- There are consistent threads through our applications, system dissemination, core research, and computing environment development work. These threads are that our research work at all levels is driven by the real-world scientific applications that we undertake; that we choose applications that have a high impact on current medical and biological problems and that expose key underlying AI research issues; and that we seek to maximize the availability of the facilities for and results of this work in the biomedical community. This is seen, for example, in the coupling between our core research and development work and applications projects such as ONCOCIN and PROTEAN.

- We must continue to provide the computing resources for the growing Stanford biomedical AI research community and the national projects still dependent on us, to emphasize nurturing newly started AI applications, to serve as a communications cross-roads for the large and diverse AIM community, and to ensure broad dissemination of our research results and methods.

2.2.1. Core Research and Development

Reasoning in medicine and the biological sciences is knowledge-intensive. A recent article in *Science* [12], for example, discusses the role of information in the search for a cure for cancer. As the rate of explosion of knowledge continues to increase, clinicians and biomedical scientists must turn to computers for help in managing the information, and applying it to complex situations.

Artificial intelligence methods are particularly appropriate for aiding in the management and application of knowledge because they apply to information represented symbolically, as well as numerically, and to reasoning with judgmental rules as well as logical ones. They have been focused on medical and biological problems for over a decade with considerable success. This is because, of all the computing methods known, AI methods are the only ones that deal explicitly with symbolic information and problem solving and with knowledge that is heuristic (experiential) as well as factual.

Expert systems are one important class of applications of AI to complex problems -- in medicine, science, engineering, and elsewhere. Expert Systems draw on the current stock of ideas in AI, for example, about representing and using knowledge. They are adequate for capturing problem-solving expertise for many bounded problem areas. But the current ideas fall short in many ways, necessitating extensive further basic research efforts. Our core research goals are to analyze the limitations of current techniques, to investigate the nature of methods for overcoming them, and to develop tools to build and disseminate new and more effective biomedical expert systems.

Long-term success of computer-based aids in medicine and biology depend on improving the programming methods available for representing and using domain knowledge. That knowledge is inherently complex -- it contains mixtures of symbolic and numeric facts and relations, many of them uncertain; it contains knowledge at different levels of abstraction and in seemingly inconsistent frameworks; and it links examples and exception clauses with rules of thumb as well as with theoretical principles. Current techniques have been successful only insofar as they severely limit this complexity. As the applications become more far-reaching, computer programs will have to deal more effectively with richer expressions and much more voluminous amounts of knowledge.

2.2.1.1. ONCOCIN-Related Core Research

As mentioned earlier in this application, our research plan for the next five years includes merging the core research activities of the ONCOCIN project with other basic research activities coordinated by the SUMEX resource. The ONCOCIN project is now in its sixth year and has involved approximately 40 research staff and students, some of whom have worked full time on aspects of the program or its knowledge base. It is accordingly large and has elements that span a variety of basic and applied research issues. The project's elements have been summarized in some detail elsewhere in this application and in the SUMEX annual report.

Since 1983 the Biomedical Research Technology Program, through a resource-related grant (RR-01631), has supported the effort to convert ONCOCIN to run on professional workstations (the Xerox 1108 Lisp machine). When that grant terminates in 1986, ongoing research will include a mixture of applied activities (evaluation of the workstations in the Stanford clinic and experiments to implement ONCOCIN workstations in private oncology offices in Northern California) and more basic activities intended to generalize past ONCOCIN results for the AIM community. We propose to continue the basic aspects of this work as core research under the SUMEX grant, and use complementary support for the other aspects of the project from the National Library of Medicine and, if a pending application for a dissemination experiment is successful, jointly from the National Center for Health Services Research and the National Cancer Institute.

In this section we summarize the core research activities that we intend to pursue in the context of ONCOCIN. They fall into four principal categories: implementation of ONCOCIN workstations in the Stanford clinic, knowledge acquisition research (OPAL), research to generalize ONCOCIN for application in clinical trial domains other than medical oncology (E-ONCOCIN), and research on generalized approaches to strategic therapy planning (ONYX).

Background on The ONCOCIN Program

From the outset, the ONCOCIN research effort has been directed towards both basic research in artificial intelligence and the development of a clinically useful consultation tool. We initially sought to apply techniques developed during our earlier work on the MYCIN system and to extend those methods to interact with a large database of clinical information. More recently, however, the system has departed from the uniform production rule approach of MYCIN in several significant ways (e.g., introduction of heterogeneous knowledge structures and distributed control processes [50] in the workstation version of ONCOCIN). Our approach to these problems has been greatly influenced by the Lisp machine technology to which we were first exposed through the foresight of SUMEX when it acquired such experimental machines in the early 1980's.

The initial version of ONCOCIN, including its clinical implementation in our cancer clinic, runs on a time-shared DEC-20 computer and uses a customized video display terminal installed in our oncology clinic. Since May of 1981, the prototype has been used on a limited experimental basis by oncology faculty and fellows to obtain advice on the treatment of patients enrolled in protocols for the treatment of Hodgkin's disease and non-Hodgkin's lymphoma. In the past year, additional protocols for adjuvant chemotherapy of breast cancer were added to the system.

We are excited by the promise of this prototype version of ONCOCIN. Formal evaluation of the system has shown that ONCOCIN does very well in suggesting therapy, even in cases where complex attenuation or changes in drugs are required [33]. It has also had a significant effect on the completeness with which clinical trial data are captured and made available for analysis [35]. In addition, we are extremely

encouraged by the effectiveness of the interface program we have devised (the *Interviewer*) and the speed with which new users have been able to learn to use the system.

We believe that our current efforts to adapt the existing prototype for use on professional workstations will increase ONCOCIN's clinical acceptability. The use of a dedicated computer featuring high resolution graphics and mouse pointing devices to obviate typing should make the system even more attractive to busy physicians. As is described in the ONCOCIN progress report elsewhere in this proposal, we expect to have two Lisp machine (Xerox 1108) workstations in use in the Stanford oncology clinic by mid-1986. Thus, the continuation of ONCOCIN research in that clinic (knowledge base enhancement, software development in response to user feedback, and evaluations of the impact and acceptance of the workstation technology) will continue under the SUMEX umbrella after the merger of the SUMEX and ONCOCIN activities at the beginning of the next grant period. We should emphasize that, because of the moderate price of these computers, we look forward to transferring ONCOCIN for use in small clinics and physicians' offices. This will offer private physicians up-to-date decision support for the treatment of cancer patients (a recognized area of need) while allowing randomized clinical trials (RCTs) in oncology the benefit of greatly expanded access to appropriate patients. A four year experiment to install and test ONCOCIN in private offices has been proposed and is awaiting review and a site visit at this time.

Automated Knowledge Acquisition for RCTs

RCTs are based on rigidly structured therapy plans. Oncology protocols demonstrate this point nicely. RCT protocols are comprised of treatment *arms*, which in the case of oncology specify sequences of chemotherapy or radiotherapy. There is an explicit hierarchy of knowledge elements in these protocols which becomes important for knowledge acquisition. The hierarchy for a typical cancer chemotherapy protocol is shown in Fig. 6.
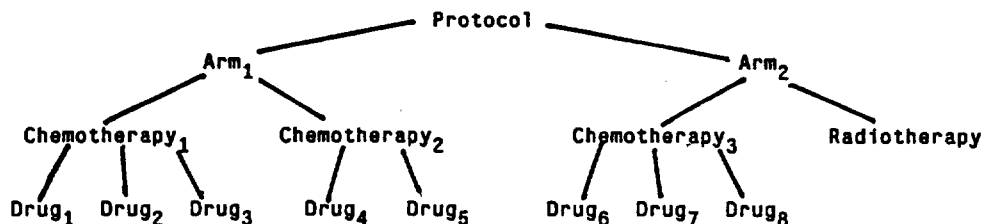

Figure 6: Sample Chemotherapy Protocol Hierarchy

ONCOCIN uses a variety of internal representations to store protocol knowledge. For example, in one arm of a protocol for small cell lung cancer, seven different drugs are used as part of two chemotherapies in a specific sequence over seven weeks. The sequence of chemotherapies is repeated five times, making the total duration of treatment 35 weeks. The names of the chemotherapies are POCC and VAM. Administering POCC requires that the patient make two separate clinic visits to receive medication during each treatment cycle. Hence, POCC is divided into two sub-cycles: POCC-A and POCC-B. After the second complete cycle of POCC, the patient is given cranial irradiation. The computer representation of this entire complex sequence is:

```
((POCC 1 A) (POCC 1 B) (VAM 1)
 (POCC 2 A) (POCC 2 B)
 (XRT CRANIAL)
 (VAM 2)
 (POCC 3 A) (POCC 3 B) (VAM 3)
 (POCC 4 A) (POCC 4 B) (VAM 4)
 (POCC 5 A) (POCC 5 B) (VAM 5))
```

This purely procedural knowledge can be extracted from protocol documents fairly easily; one need not understand oncology. However, much of the important knowledge in ONCOCIN is more judgmental and is represented in the form of production rules. ONCOCIN currently uses over 400 rules to determine:

- how to adjust specific drug dosages because of treatment-induced low blood counts or other adverse (toxic) reactions to therapy

- when to delay treatment or abort a therapy cycle

- how to modify therapy in light of a patient's changing clinical conditions or response to the protocol

- when to order certain laboratory tests and how to interpret their results.

Note that these issues are generic for all clinical trials, and similar rules could be written to assist with proper administration of treatment for RCTs in other medical domains.

An example of one such rule, drawn from the ONCOCIN system, is shown in Fig. 7. It was developed by examining a formal protocol and then further enhancing and validating the knowledge through discussions between an oncologist and a knowledge engineer.

```
To determine the current attenuated dose for patients with all lymphomas
    in CHOP chemotherapy for Cytoxan or Adriamycin:

If:  1. The blood counts warrant dose attenuation
     2. It patient did not receive chemotherapy
        before the last radiation therapy
     3. This is the first cycle after significant radiation
     4. This is not the first visit after an Abort cycle

Then:  Conclude that the current dose is 75% of the standard
       dose further attenuated by either the dose attenuation
       for low WBC or the dose attenuation for low platelets,
       whichever is less.
```

Figure 7:   Sample ONCOCIN Rule, Translated to English from Internal Format

The knowledge engineer then must convert this rule into a representation understandable by the computer. The rule format for computer use is generally unreadable to the clinician who helped to develop the rule in the first place. It is the translation shown in the figure that is created and reviewed by the clinician. The knowledge engineer's detailed understanding of the manner in which information is represented in the computer allows him or her to develop the corresponding machine-understandable format.

Because the knowledge engineering process is cumbersome and inefficient, we have recently embarked on work to develop a system, termed OPAL, that acquires new knowledge of oncology protocols directly from physicians while shielding them from technical details. As part of our SUMEX core research activities, we will seek to

generalize this approach for application in other medical domains in which RCTs are commonly used. The knowledge contained in protocols for oncology (and for other RCTs as well) has already been formalized in the protocol document. The most fundamental problems of conceptualizing and structuring the domain knowledge should therefore not be an issue in this work.

For example, detailed discussions with our oncology experts and review of dozens of protocol documents make it clear that the knowledge in protocols is both predictable and constrained by the very nature of oncologic clinical trials. For each concept that appears in oncology protocols, we can anticipate the general nature of most of its possible *values*. For example, we can assume that all drugs will have a *dose* that can be represented by an integer. All drugs will have a *route*--intravenous, intramuscular, or oral. Our knowledge of the field allows us to determine *a priori* what possible choices might be appropriate for most concepts. This has great implications for automated assistance in knowledge acquisition.

We have known for some time that it would be ideal to provide an environment so that the physicians can themselves enter and manipulate knowledge of a RCT protocol and related medical knowledge. However, since it is generally unrealistic to teach collaborators to become programmers or knowledge engineers, we are faced with the traditional problems of getting a computer to understand the *meaning* underlying unstructured phrases or sentences entered by a physician. TEIRESIAS had approached the problem by cleverly manipulating the context of an interaction with an expert, thereby simplifying the task of understanding entries [13]. However, problems in computer-based understanding of natural language (still a major research topic in artificial intelligence) prevented TEIRESIAS from becoming sufficiently robust for routine use. We have been unwilling to reopen the Pandora's box of natural language understanding for the ONCOCIN project, and therefore in the early years have had to resort to the LISP-based entry of knowledge.

Two factors have accounted for our decision to turn again to the problem of knowledge acquisition. The first has been a simple matter of *need*. As we have developed plans to adapt ONCOCIN for use on single-user machines in physicians' offices, and have contemplated the large numbers of protocols that must be available online for practical use of such a tool, we have been forced to acknowledge the necessity of an enhanced knowledge acquisition capability. Second, in transferring ONCOCIN to personal workstations and familiarizing ourselves with this new technology, we have become aware of the potential for using advanced graphics techniques to avoid problems of natural language understanding during entry of knowledge by a computer-naive user. To explore the possible use of the graphics capabilities of LISP machines to facilitate knowledge acquisition directly from experts, we have recently developed a prototype system for knowledge entry. OPAL was designed in close collaboration with oncologists who will be the eventual end users of such a system. To build the prototype version of OPAL we reviewed all of the concepts that had been required for each of the protocols that we entered by hand, and explored a large number of existing protocol documents that we hoped to enter into the completed system.

The OPAL prototype runs on the same professional workstation (the Xerox 1108 "Dandelion") on which the new version of ONCOCIN is being developed. Like the new ONCOCIN system, OPAL is designed to take advantage of the advanced graphics capabilities of the workstation and uses a mouse pointing device almost exclusively for input by the physician.

In developing OPAL, we attempted to organize the information to be entered by the physician in a manner similar to the structure of typical protocol documents. A constant consideration was to request knowledge from the physician in a manner consistent with the way oncologists tend to think about protocols. OPAL guides

protocol entry in a loose fashion; the expert is provided with an ability to change topics at his or her convenience. However, the program follows an orderly progression, first asking for general information about the scope of the protocol, principal investigators, and inclusion and exclusion criteria; next asking for the protocol "schema" -- a shorthand notation that describes the sequences of treatments; and finally requesting information on specific drugs, dose modifications, and diagnostic tests required by the protocol.

The questions for each of these categories are grouped into individual windows on the graphics display. These windows contain a number of "blanks" on the screen to be completed in order to provide pertinent protocol information. Most blanks can be filled in by selecting them with the mouse and then selecting an item from a menu that is displayed. Rarely the blanks are filled in by typing at the keyboard. The windows are not all displayed at once but rather are selected one at a time by the physician working his or her way through a protocol. Selecting a window brings it "into view". In the present OPAL prototype, most of the major windows are portrayed graphically as a stack of overlapping "file folders" on the screen. Using the "mouse" to select the "tab" of one of these folders brings the corresponding window into view. Special menu windows can be created for the entry of purely numerical data. For example, we have developed menus, called "registers", that appear either in the format of a 10-key calculator pad (for free-form digit selection) or else in a columnar format, akin to the front of an old-style cash register. In either case, the user indicates the appropriate digits sequentially using the mouse without needing to touch the keyboard. Several examples of the windows used for protocol entry are provided in the working paper by Differding included as an Appendix to this application.

The OPAL prototype presumes that the user will have no appreciation for how knowledge is stored in the computer for use by the reasoning elements in ONCOCIN; the user need only be able to understand oncology protocol documents. The system deals with chemotherapy knowledge at such a high level that the user is completely shielded from issues of knowledge base organization and format. The physician using OPAL needs to be concerned only with the actual knowledge in the protocol to be entered.

The preliminary version of OPAL consists of a series of windows that may be displayed on the screen of the 1108 workstation in any order. Each window represents a series of questions or blanks to be filled in for a specific portion of a protocol's knowledge. For example, one window asks questions about the names and standard dosages for the drugs to be used for a given chemotherapy; another asks what laboratory studies are required by the protocol; a third inquires what actions to take if certain toxicities develop.

For each possible "blank" in the window, information is entered automatically by the system if the corresponding data are already known because of previous responses (e.g., if a standard chemotherapy is chosen in one window, the individual drugs involved will then appear in all of the other windows that ask for drug information). Otherwise, selecting a blank with the mouse causes a menu with possible completions for that item to "pop up" on the screen. The mouse is then used to select the desired response from the menu.

The OPAL prototype has been tested by several physicians and all have found the system easy to use after a few minutes of training. Frequent feedback from our oncology collaborators has allowed us to make modifications, expanding the options in certain menus and improving the user interface. These modifications have been effected by reprogramming parts of the system. However, we plan to be able to make changes to OPAL eventually by editing *data structures*, rather than by having to update the actual computer *programs*.

When a protocol is entered using OPAL, the knowledge ultimately must be encoded in an internal form so that ONCOCIN can use it to give advice and manage the protocol data. We see this encoding occurring in a two stage process, with an intermediate data structure serving to insulate the interaction with OPAL from the detailed structure of the knowledge base. Thus OPAL will be used to enter protocol knowledge, it will be stored in an intermediate data structure (or IDS), and then further refined into a knowledge base for use by ONCOCIN. As is outlined in the next section, these ideas generalize to RCT advice systems in other clinical domains -- a generalized OPAL might be used to enter RCT guidelines, thereby creating a knowledge base for use by a generalized version of ONCOCIN.

Generalization of ONCOCIN: *E-ONCOCIN*

Most protocols in clinical medicine contain elements in common with oncology trials. We plan to build on our experience creating OPAL to apply the same methodology to develop expert systems for RCTs in other medical areas. This research to develop generalized knowledge acquisition programs like OPAL for other RCTs will be of great practical importance. However, we recognize that the work will address significant theoretical issues in the field of medical artificial intelligence. In fact, we expect that the Meta-OPAL work outlined below will constitute a Ph.D. dissertation for one of our Medical Information Sciences graduate students (Dr. Mark Musen).

What we propose is a high-level tool for use by knowledge engineers in conjunction with clinicians to define all the properties of a *knowledge acquisition system* (KAS) that may be used subsequently to enter the knowledge for a particular *class* of clinical trials. OPAL is an example of a KAS, one that is customized for the class of clinical trials relevant to clinical oncology. A KAS for another domain, such as hypertension or epilepsy management, might look very different. Certainly the display windows for protocol entry would bear little resemblance to those used in the current version of OPAL. This new high-level tool, *Meta*-OPAL, will take as its input the complete specifications for a KAS. It will produce as its output a data structure that will enable a second program, E-OPAL, to interact with a domain expert to capture and encode a whole class of new protocols. These encoded protocols can then be used for data management and consultation by a domain-independent version of ONCOCIN (the ONCOCIN inference engine, to be termed E-ONCOCIN)[1]. E-OPAL will be a version of OPAL stripped of all its built-in oncology knowledge. E-OPAL thus will rely on Meta-OPAL to provide all the information required to perform knowledge acquisition and management. The relationships of the various modules is diagramed in Figure 8.

The concept of a "knowledge acquisition system for knowledge acquisition systems" is attractive in many respects. First, many of the problems of a limited "world view" in a program such as OPAL will be readily overcome because all of the domain assumptions (e.g., beliefs about oncology, cancer protocols, or chemotherapy) will be explicitly declared at the Meta-OPAL level. For example, an implicit assumption built into the present OPAL prototype is that patients are treated with either chemotherapy or radiotherapy. The physician using OPAL is never asked to enter information regarding, say, *surgery* because knowledge about options for surgery is not currently within OPAL's "world view". Even by modifying OPAL to specify new parameters, no protocol that called for repeated surgical procedures could be satisfactorily encoded unless we had an ability to make even higher-level modifications to OPAL.

At present, we can make this sort of higher level modification to OPAL only by

---

[1] The names E-OPAL and E-ONCOCIN are inspired by the similar domain independent tool developed by our group in the 1970's. This program, EMYCIN or "Essential MYCIN", is the inference engine separated from the knowledge base of MYCIN