C. Critique of Resource Management

The SUMEX computer resource and Lisp workstations have been very important for the work to date, and the SUMEX staff has continued to be very cooperative with the REFEREE project.

III. RESEARCH PLANS

A. Goals & Plans

The overall objective of the REFEREE project is to use recent Artificial Intelligence techniques to build a system that helps the informed but statistically non-expert reader to evaluate critically the medical literature on randomized controlled trials (RCT's). This system will contain and be able to apply dynamically the detailed specialized knowledge of Dr. Byron W. Brown, a biostatistician expert in the design and evaluation of randomized controlled trials. We have divided our overall objective into two goals:

- Goal 1 is the construction of an expert system to help readers (e.g. medical students, medical researchers, clinicians, journal editors, or editorial assistants) assess the credibility of a *single* conclusion drawn from a *single* journal report of a randomized controlled trial. We have already made substantial progress toward this goal with the development of the prototype REFEREE system.

- Goal 2 is the expansion of REFEREE to an expert system that can be used by a similar range of readers to facilitate the evaluation of *multiple* reports based on randomized controlled trials. This expanded system, to be known as the REVIEWER, will thus perform meta-analysis.

The task of extending and refining the prototype REFEREE system in order to achieve these goals can be characterized in terms of three dimensions:

- Making the system more accessible to a variety of people by improving the user interface, validating the system's performance with different types of users, and providing an explanatory capability

- Expanding the knowledge base by continuing the knowledge acquisition process to cover additional types of RCT's

- Improving the inference engine to ensure consistency of the knowledge base and to focus the consultation process on questions relevant to the situation and the individual user.

The specific steps that are planned for the enhancement of the REFEREE system include the following:

- Critique individual clinical trials according to the methodological quality of the trial;

- Measure the efficacy of treatment as demonstrated in a randomized control trial;

- Compare and contrast the credibility and efficacy of treatment reported by multiple journal articles; and

- Combine the *qualitative* techniques of heuristic reasoning and the *quantitative* methods of statistical meta-analysis to extract a consensus opinion from multiple knowledge sources.

In addition, plans for Goal 2, the REVIEWER system to analyze multiple RCT's and form a consensus judgment, include:

- Complete a review of the available literature on meta-analysis and augment the REFEREE prototype to produce estimators for meta-analysis and incorporate expert knowledge on the appropriateness of these methods.

- Add explicit and heuristic knowledge needed for the calculation of robust, non-parametric estimators of effect size.

- Construct a prototype of a system that builds categorical models in the domain of meta-analysis, to perform autonomous investigations in the domain of statistical model-building. The REVIEWER will utilize expert knowledge in biostatistics to guide its search for meaningful models.

- Build a prototype of a system that can explore the domain of regression models for multiple RCT's that will use expert knowledge in its selection of predictor variables.

- Package the REVIEWER in a form suitable for use by physicians and their assistants.

- Verify the expertise of the REVIEWER system on a suite of papers drawn from clinical trials, similar to the validation of REFEREE above.

B. Justification for continued SUMEX use

The local area network maintained by the SUMEX staff is essential to the effective development and use of the REFEREE system on Lisp workstations. The availability of the Xerox workstations makes possible the evaluation of prototypes in that environment, and also facilitates the development of good user interfaces. The connections through the 2060 to local and national computer networks such as ARPANET are important for sharing ideas and results with other medical researchers.

C. Need for other computing resources

We are exploring the implementation of REFEREE on personal computers such as the Macintosh II and other high performance machines. We anticipate the need for at least two of these machines for transporting our system and developing new modes of interaction with both naive and experienced users.

IV.D. Pilot AIM Projects

Following is a description of the informal pilot projects currently using the AIM portion of the SUMEX-AIM resource, pending funding, full review, and authorization.

In addition to the progress report presented here, an abstract is submitted on a separate Scientific Subproject Form.

IV.D.1. Dynamic Systems Project

Decision Support for Time-Varying Clinical Problems

Lawrence E. Widman, M.D., Ph.D.
Division of Cardiology
University of Texas Health Science Center
San Antonio, Texas 78284

I. SUMMARY OF RESEARCH PROGRAM

A. PROJECT RATIONALE

Time-varying systems, which include many areas of medicine, science, economics, and business, can be described mathematically by differential equations. They are distinct from the pattern-matching and logic-based domains dealt with so successfully by existing expert system methods, because they can include feedback relationships. It is generally felt that they are best approached by enhancement of existing methods for deep model-based reasoning.

The goal of this project is to develop AI methods for capturing and using knowledge about time-varying systems. The strategy is to address general problems in model-based knowledge representation and reasoning. The intermediate objective is to develop methods which are powerful enough to work in selected realistic situations yet are general enough to be transportable to other, unrelated knowledge domains.

Our approach is to work on well-defined yet complex and interesting problems in the medical domain. We have, therefore, selected the human cardiovascular system as our prototype of a time-varying system, and are developing methods for representing and reasoning about its mechanical and electrical activities in the normal and diseased states.

A.1 Technical Goals

This project presently has two distinct tracks: hemodynamic modeling and cardiac arrhythmia interpretation.

1. Hemodynamic Modeling

The goals of this subproject are to develop:

(a) a knowledge-representation method using symbolic modeling which captures the qualitative and, when possible, the quantitative behavior of systems with feedback relationships. Preferably, the symbolic model should be translatable into the differential equations which describe the behavior of the system being modeled.

(b) a reasoning method based on the symbolic modeling tool created in subgoal (a) which permits the inference of differential diagnoses (a set of hypothesized diagnoses) from incomplete data.

(c) a reasoning method based on subgoals (a) and (b) which permits inference of the state of the model for each hypothesized diagnosis. This subgoal would be satisfied by an algorithm which specifies a self-consistent set of values for all variables in the model, for a given hypothesis based on a given set of data. Such sets of data would constitute initial conditions for differential equations derived from the model.

(d) a simulation method, based on the model and its equivalent differential equations together with the initial conditions derived from the differential diagnosis (steps a-c above), for predicting the expected time course of the system being modeled for each hypothesized diagnosis. This method could also be used to predict the effects of treatments being considered for recommendation by the program.

(e) a reasoning method, based on domain-independent properties of the model, for shrinking and/or expanding the model automatically to use a minimal model configuration to account for normal and abnormal data.

(f) an explanation facility for examining the model, the given data, the inferred hypothesized diagnoses, predicted behaviors, and modifications of the model, to answer user queries and to teach fundamental concepts.

2. Cardiac Arrhythmia Recognition

The goals of this subproject are to develop:

(a) a symbolic model of the electrical system of the human heart, including pertinent anatomic and electrophysiologic features of the normal and diseased heart. The electrophysiologic features would include deterministic characteristics (e.g., conduction velocities, refractory periods), stochastic features (e.g., behavior of automatic foci), and temporal interactions (e.g., competing pacemakers).

(b) a symbolic/numeric representation of the observable features of the electrical activity of the heart, both surface ECG and intracardiac recordings, including noise. This representation would be intended to allow a feature extraction module working on actual patient data to communicate with a symbolic reasoning module, and would be translatable directly into waveform display format.

(c) a reasoning method for extracting features (identifying waveforms) from raw, digitized signal data. This method would augment established signal processing techniques by using knowledge-based algorithms to improve detection of P and T-U waves and to improve rejection of noise. It should be noted that this is itself a major research undertaking in the signal processing domain.

(d) a reasoning method for inferring the cardiac rhythms consistent with a given disease state in the model, similar to the prediction of consequences of the hemodynamic model in the first subproject. The output of this method would be in the symbolic/numerical representation of subgoal (b).

(e) a reasoning method for inferring possible disease states in the model from a given feature-extracted recording of the electrical activity of the heart. This subgoal constitutes cardiac arrhythmia interpretation, and is itself a major research project.

(f) a categorization method for inferring hierarchies of diagnoses from elementary abnormalities. For example, "periods of atrial fibrillation up to 30 minutes at up to 150 beats/min, supraventricular tachycardia of up to 10 beats length at a rate of 130 beats/min, and sinus bradycardia with a minimum rate of 45, all consistent with the sick sinus ("tachy-brady") syndrome" and "two QRS morphologies are present: they are narrow at rates less than 120 and are wide at rates above 120, consistent with a rate-dependent bundle branch block".

(g) an explanation facility for examining the model, the input data, and the interpretations to answer user queries and to teach fundamental concepts.

B. MEDICAL RELEVANCE AND COLLABORATIONS

The two subprojects have related but separate medical goals:

1. Hemodynamic Modeling.

There are two subgoals in this subproject: model-based sensor integration and model-based care-giver assistance.

a. Model-based Sensor Integration

The long-range application of this subproject is the integration of patient-related data in the intensive care environment. Model- based real-time systems would allow the system to share a global understanding of the patient's condition with the human care-givers. Thus, it could interpret significant trends in key parameters and could draw attention to relationships which might otherwise escape attention in the constant flood of data common to these environments.

b. Model-based Care-giver Assistance

It could also serve as an assistant to the care-giver. In this mode, the human care-giver could evaluate the merits of proposed diagnostic and therapeutic measures in light of available data on the patient's condition.

Practical application of these concepts requires further development of the model and the reasoning algorithms, and extensive testing against real clinical scenarios. Refinement and quality control are presently the responsibility of the principal investigator, who is a board-certified cardiologist with subspecialty interest in cardiac electrophysiology.

Practical application also awaits general acceptance of standardized hospital data buses for automatic acquisition of important parameters now stored primarily on paper or on computers outside the intensive care setting, such as fluid inputs and outputs, medications, and results of invasive and non-invasive tests. Further, improved user interfaces will require better graphics and increased computer literacy on the part of care-givers.

2. Cardiac Arrhythmia Recognition.

The long-range application of this sub-project is in clinical devices such as intensive-care arrhythmia monitors, portable Holter monitors, and implantable cardioverter-defibrillators. There are two subgoals: recognition of surface electrocardiographic (ECG) recordings and recognition of intracardiac recordings.

a. Recognition of surface electrocardiographic recordings.

Substantial and well-recognized obstacles in signal processing will likely prevent non-AI algorithms from advancing beyond the current state of the art of interpretation of surface ECG recordings. These obstacles are primarily the problems of reliable detection of P and T-U waves, and rejection of noise.

We hope further that, by mimicking the behavior of expert human cardiologists, these obstacles can be bypassed if they cannot be overcome. We have enlisted as consultants Dr. William Long of M.I.T., who has long experience with modeling of the cardiovascular system and with symbolic approaches to ECG interpretation, and Dr. Charles Mead, a consultant formerly with the Division of Cardiology at Washington University - St. Louis who has 15 years experience with the development of algorithms for clinical arrhythmia recognition.

b. Recognition of intracardiac recordings.

Intracardiac recordings, which are taken from wires placed in the heart by percutaneous venous puncture or around the heart by surgery, are relatively free of P wave ambiguity and of noise. They are representative of the quality of signals available to implantable cardioverter-defibrillators.

Cardioverter-defibrillators are devices like pacemakers in that they monitor the heart rhythm in a patient to determine if an abnormality exists. They are capable of taking action (electrical countershock) if an appropriate abnormality is detected. Unlike ordinary pacemakers, these devices detect abnormalities characterized by rapid rates of heart activity, rather than excessively slow rates. They have been shown to reduce one-year mortality in high-risk patients from 30% to 2%, and they are expected to play an increasingly large role in treatment of such patients.

These relatively new devices currently use quite simple algorithms to detect abnormalities. The action they take consists of applying an electrical shock directly to the heart. This shock is frequently unpleasant to the patient. The problem is that the algorithms sometimes confuse innocent rapid heart rates, such as from exercise or atrial tachyarrhythmias, with lethal ventricular arrhythmias. This has proved troublesome enough to prompt repeated calls in the electrophysiology literature for improved algorithms for arrhythmia recognition in these devices.

The algorithms developed in this subproject would be suitable for this application when the computer power in the devices improves. Because these devices require powerful energy sources to perform repeated shocks over their lifetimes of 2-3 years, the power drain of more sophisticated computer chips is less important than it would be in ordinary pacemakers.

C. Highlights of Research Progress

1. Hemodynamic Modeling

Subgoals (a) through (d) have been accomplished in prototype form. The approach relies on a semi-quantitative representation (subgoal (a)) which assigns values by default if the user does not specify more detailed information. Constraint propagation using a dynamically generated semi-quantitative quantity space is performed by interpreting the model as a set of constraint equations. Domain- independent heuristics which recognize morphological features of the model are used to further constrain the propagation of constraints and to generate hypotheses when ambiguities arise. These heuristics generate a set of self-consistent hypotheses, each of which is a hypothesized diagnosis (subgoal b). Each hypothesized diagnosis is then refined by mathematical relaxation, in which the propagated values are treated as initial guesses, and the values are refined iteratively, again by interpreting the model as a set of constraint equations (subgoal c). In the several scenarios which have been examined, the value assignments achieved by hypothesis and iterative refinement have achieved correlation coefficients up to 0.90 with the values obtained by simulation of the same model. Phase (d) was accomplished by translating the model into a set of dynamical systems equations, which were then integrated in the standard manner.

We did not anticipate continuing this subproject after beginning the cardiac arrhythmia subproject. However, an opportunity has arisen to collaborate with Professor Benjamin Kuipers at the University of Texas at Austin on the development of algorithms which combine the advantages of classical qualitative simulation and of our semi- quantitative simulation approach. Professor Kuipers and a graduate student will spend the summer of 1988 in San Antonio for this purpose and to develop a case-based reasoning knowledge base of cardiovascular disorders.

2. Cardiac Arrhythmia Recognition

This subproject has been operational for almost one year. With support from the University of Texas Health Science Center at San Antonio and the American Heart Association, Texas Affiliate, we have acquired and integrated a Symbolics 3640 LISP workstation and an IBM- compatible microcomputer with attached optical scanner.

Progress has been made in two subprojects of this project:

a) Signal acquisition from paper records. Many standardized and interesting ECG recordings are available on paper but not in digitized form. To make these recordings machine-readable, we have developed software in TurboPascal (a popular microcomputer language) which digitizes the ECG signal on a paper ECG record by extracting the ECG line from the bit map of the image produced by the optical scanner. Algorithms have been developed with can extract the ECG line when grid lines are present and in the presence of varying intensity of the original image (as, for example, in xerox copies of paper records). The output of this software is a sequence of numbers representing the ECG voltage at sequential instants of time, equivalent to the output of an analog-to-digital converter. This set of numbers can then be used for development of ECG recognition software. In comparison to conventional analog-to-digital recordings, the optical scanner output has a root-mean-square error of 3.5-4.0% (Widman and Freeman), or about 1.5 times the thickness of a typical pen line on paper.

b) Identification of waveforms in raw signal data. Using digitized signals from electronic analog-to-digital conversion and the optical scanner conversion described above, we have implemented prototype software for waveform identification. It is organized by levels of abstraction. At the lowest level, conventional slope-analysis and filtering algorithms are used to identify the baseline and the QRS complexes. Knowledge-based slope-analysis algorithms then analyze the remaining signal to assign tentative labels of "P waves," "T waves," and "other" waves. Then, knowledge-based pattern recognition algorithms look for temporal and morphological relationships to confirm or correct the tentative labels. The prototype system is able to find P waves superimposed on T waves, a combination which cannot be detected by any commercial cardiac arrhythmia system.

D. List of Relevant Publications

1. Widman, L.E. Knowledge-Based Fault Identification and "What If" Simulation in Symbolic Dynamic Systems Models. Proceedings of the 1988 Society for Computer Simulation Multiconference on AI and Simulation, pp. 89-94, 1988.

2. Widman, L.E., Lee, Y.-B., and Y.-H. Pao. Diagnosis of Causal Medical Models by Semi-Quantitative Reasoning. In: Miller, P.L. (ed.). Topics in Medical Artificial Intelligence, Springer-Verlag (in press)

3. Widman, L.E. Semi-Quantitative "Close Enough" Systems Dynamics Models: An Alternative to Qualitative Simulation. In: Widman, L.E., Loparo, K.A., and N.R. Nielsen (eds.). Artificial Intelligence, Simulation, and Modeling. John Wiley & Sons, New York (in press).

4. Widman, L.E., Loparo, K.A., and N.R. Nielsen (eds.). Artificial Intelligence, Simulation, and Modeling. John Wiley & Sons, New York (in press).

5. Widman, L.E. and G.L. Freeman. A-to-D Conversion from Paper Records with a Desktop Scanner and a Microcomputer. (submitted to the American Journal of Physiology).

E. Funding Support

1. American Heart Association, Texas Affiliate
Grant-in-Aid Award.
Knowledge-Based Computer Algorithms for Arrhythmia Analysis.
Principal Investigator: Lawrence E. Widman.
Award period: July, 1987 - June, 1988.
Level: $24,850 direct costs.

2. Medical Research Service, Department of Medicine and Surgery, Veterans Administration.
Knowledge-Based Computer Algorithms for Cardiac Arrhythmia Interpretation.
Principal Investigator: Lawrence E. Widman.
Award period: July, 1988 - June, 1990.
Level: $33,097, year 1; $26,647, year 2. (direct costs)

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Sharing and interactions with other SUMEX-AIM projects

The major interaction with SUMEX-AIM this year has been communication with members of the AIM community, now that our work has been transferred to the Symbolics machine here.

SUMEX-AIM allows ready Email access to users of ARPANET, Bitnet and CSNet. This access has proved invaluable in communicating rapidly and effectively with co-workers at other institutions. The value of this access has been demonstrated particularly during the preparation of the book, Artificial Intelligence and Simulation, which will go into production in June, 1988. Email was used to communicate with contributors, to receive electronic copies of their manuscripts, and to issue unobtrusive, gentle but effective reminders when the manuscripts were late. Without Email access, the editorial process would have been much more difficult.

Review of the longer term history of this project shows that it would not exist had SUMEX-AIM not provided telecommunication support for the initial feasibility project in 1984-1985, which was carried out on the computers of the MIT Laboratory for Computer Science, Clinical Decision Making Group, and computational support during the development of the semi-quantitative approach in the years 1985-1987.

C. Critique of Resource Management

The service provided by SUMEX-AIM has been exemplary, largely because of prompt and effective response to difficulties as they arise. There has been a clear effort to assure that telecommunication access remained reliable during changes in commercial vendors, and the staff have responded to several technical questions promptly and accurately. Down-time has been minimal compared to that of other systems we have used, and is almost always scheduled several days in advance.

The reason we work within the AIM community is that it is the natural niche for our research interests. There is no short-term prospect that this project will reach commercial maturity or that it will lose sight of fundamental AI issues, and so we feel that it still belongs in the scientific AIM framework.

As noted in the previous section, the communication with other members of the AIM community has proved invaluable in the advancement of this project.

III. RESEARCH PLANS

A. Project Goals and Plans

The long range goals of this project are to develop intelligent comprehensive monitoring/alarm systems for intensive care unit settings; and intelligent arrhythmia recognition systems for monitors, Holter recorders, and implantable cardioverter/ defibrillators. The short term strategies for achieving these goals are discussed above.

The principle investigator is a staff cardiologist and Assistant Professor in the Division of Cardiology. He serves as an Attending Cardiologist during hemodynamic, angiographic and electrophysiological studies on selected patients. Substantial time is committed to research, and this project will constitute his major research emphasis.

B. Justification and requirements for continued SUMEX use

The justification for this project is its potential for advancing the state of the art of expert system technology in the area of temporal reasoning and deep causal modeling, and for demonstrating practical use of expert symbolic computing in potentially life-saving, knowledge-intensive environments.

The requirements for continued SUMEX-AIM use should be the same as currently: telecommunications support, ARPANET access, about 3 megabytes of disc space, and a small amount of CPU time.

C. Needs and plans for Other Computing Resources Beyond SUMEX-AIM

As predicted in last year's annual report, the main computational burden has been transferred to a local resource. SUMEX-AIM is needed primarily for communication and demonstration projects, as noted above.

D. Recommendations for Future Community and Resource Development

Our strong recommendation is that SUMEX-AIM be maintained as a national AIM resource for communication, development of software useful to the AIM community, and sharing of demonstration projects. SUMEX-AIM could also serve as a central source of advice for new workstation users who may be geographically isolated from experienced workstation users.

Additionally, we would strongly support retention of the current telecommunication support and enough computing power to support promising young investigators who would otherwise not have access to symbolic computing power.

IV.D.2. Pathfinder Project

Pathfinder Project

Bharat Nathwani, M.D.
Department of Pathology
University of Southern California

Lawrence M. Fagan, M.D., Ph.D.
Department of Medicine
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Our project addresses difficulties in the diagnosis of lymph node pathology. Five studies from cooperative oncology groups have documented that, while experts show agreement with one another, the diagnosis made by practicing pathologists may have to be changed by expert hematopathologists in as many as 50% of the cases. Precise diagnoses are crucial for the determination of optimal treatment. To make the knowledge and diagnostic reasoning capabilities of experts available to the practicing pathologist, we have developed a pilot computer-based diagnostic program called Pathfinder. The project is a collaborative effort of the University of Southern California and the Stanford University Medical Computer Science Group. A pilot version of the program provides diagnostic advice on 72 common benign and malignant diseases of the lymph node based on 110 histologic features. Our research plans are to develop a full-scale version of the computer program by substantially increasing the quantity and quality of knowledge and to develop techniques for knowledge representation and manipulation appropriate to this application area. The design of the program has been strongly influenced by the INTERNIST/CADUCEUS program developed on the SUMEX resource.

Pathfinder computer science research is focused on the exploration and extension of formal techniques for decision making under uncertainty. Research foci include (1) the assessment and representation of important probabilistic dependencies among morphologic features and diseases, (2) the representation of knowledge about the progression of disease over time, (3) the acquisition and use of independent expert knowledge bases, (4) the customization of the system's reasoning and explanation behaviors to reflect the expertise of the user, and, (5) the explanation of complex formal reasoning techniques.

Toward the pragmatic goal of constructing a useful pathology teaching and decision support system, Pathfinder investigators are attempting to use intelligent computation to substantially increase the quantity and quality of pathology knowledge available to pathologists. Important areas of this knowledge integration task involve ongoing research on the crisp definition important morphologic features and feature severities, the synthesis of information from multiple experts, and the translation among multiple pathology classification schemes.

A group of expert pathologists from several centers in the U.S. have showed interest in the program and helped to provide the structure of the knowledge base for the Pathfinder system.

B. Medical Relevance and Collaboration

One of the most difficult areas in surgical pathology is the microscopic interpretation of lymph node biopsies. Most pathologists have difficulty in accurately classifying lymphomas. Several cooperative oncology group studies have documented that while experts show agreement with one another, the diagnosis rendered by a "local" pathologist may have to be changed by expert lymph node pathologists (expert hematopathologists) in as many as 50% of the cases.

The National Cancer Institute recognized this problem in 1968 and created the Lymphoma Task Force which is now identified as the Repository Center and the Pathology Panel for Lymphoma Clinical Studies. The main function of this expert panel of pathologists is to confirm the diagnosis of the "local" pathologists and to ensure that the pathologic diagnosis is made uniform from one center to another so that the comparative results of clinical therapeutic trials on lymphoma patients are valid. An expert panel approach is only a partial answer to this problem. The panel is useful in only a small percentage (3%) of cases; the Pathology Panel annually reviews only 1,000 cases whereas more than 30,000 new cases of lymphomas are reported each year. A panel approach to diagnosis is not practical and lymph node pathology cannot be routinely practiced in this manner.

We believe that practicing pathologists do not see enough case material to maintain a high level of diagnostic accuracy. The disparity between the experience of expert hematopathology teams and those in community hospitals is striking. An experienced hematopathology team may review thousands of cases per year. In contrast, in a community hospital, an average of only ten new cases of malignant lymphomas are diagnosed each year. Even in a university hospital, only approximately 100 new patients are diagnosed every year.

Because of the limited numbers of cases seen, pathologists may not be conversant with the differential diagnoses consistent with each of the histologic features of the lymph node; they may lack familiarity with the complete spectrum of the histologic findings associated with a wide range of diseases. In addition, pathologists may be unable to fully comprehend the conflicting concepts and terminology of the different classifications of non-Hodgkin's lymphomas, and may not be cognizant of the significance of the immunologic, cell kinetic, cytogenetic, and immunogenetic data associated with each of the subtypes of the non-Hodgkin's lymphomas.

In order to promote the accuracy of the knowledge base development we will have participants for multiple institutions collaborating on the project. Dr. Nathwani will be joined by experts from Stanford (Dr. Dorfman), St. Jude's Children's Research Center -- Memphis (Dr. Berard) and City of Hope (Dr. Burke).

C. Highlights of Research Progress

C.1 Previous Accomplishments

Since the project's inception in September, 1983, we have constructed several versions of Pathfinder. The first several versions of the program were *rule-based* systems like MYCIN and ONCOCIN which were developed earlier by the Stanford group. We soon discovered, however, that the large number of overlapping features in diseases of the lymph node would make a rule-based system cumbersome to implement. We next considered the construction of a *hybrid system*, consisting of a rule-based algorithm that would pass control to an INTERNIST-like scoring algorithm if it could not confirm the existence of classical sets of features. We finally decided that a modified form of the INTERNIST program would be most appropriate. The original version of Pathfinder is written in the computer language MacLisp and runs on the SUMEX DEC-20. This was transferred to Portable Standard Lisp (PSL) on the DEC-20, and later transferred to PSL on the HP 9836 workstations. Two M.D./Ph.D (Stanford Medical Information Science Program) students, David Heckerman and Eric Horvitz, designed and implemented the program and are continuing to lead research on the project.

The prototype knowledge base was constructed by Dr. Nathwani. During the early part of 1984, we organized two meetings of the entire team, including the pathology experts, to define the selection of diseases to be included in the system, and the choice of features to be used in the scoring process.

During the last three years, we have focused on methodologies for more accurately representing expert beliefs. In particular, we have used *influence diagrams* to represent dependencies among features in the Pathfinder knowledge base. A great deal of effort has been devoted to assessing and representing the intricate relationships among features that exist in the domain. We believe that this process will help to overcome some of the limitations of medical diagnostic systems.

We have also focused on the problem of complex information-theoretic inference. The explanation of a systems diagnostic behavior has been found to be of extreme importance to physicians. Unfortunately, it is often difficult to explain reasoning based on optimal models of inference. We have worked on the use of a set of alternative abstraction hierarchies to control inference. Our current techniques enable us to trade off optimality for the transparency of reasoning. We are now studying the control of this tradeoff to optimize inference.

We are also currently developing specialized probabilistic inference techniques that take advantage of particular patterns of independence that occur among features in lymph node pathology. In addition, we are developing richer mechanisms for controlling uncertain reasoning within the Pathfinder expert system.

C.2 The Pathfinder knowledge base

The basic building block of the Pathfinder knowledge base is the disease profile or *frame.* Each disease frame consists of *features* useful for diagnosis of lymph node diseases. Currently these features include histopathologic findings seen in both low- and high-power magnifications. Each feature is associated with a list of exhaustive and mutually exclusive *values.* For example, the feature *pseudofollicularity* can take on any one of the values *absent*, *slight*, *moderate*, or *prominent*. These lists of values give the program access to *severity* information. In addition, these lists eliminate obvious interdependencies among the values for a given feature. For example, if pseudofollicularity is *moderate*, it cannot also be *absent.*

Qualitative dependencies among features for each disease are represented using the influence diagram methodology mentioned above. An influence diagram contains *nodes* and *arcs*. Nodes represent features and arcs represent dependencies among features. In particular, an arc is drawn from one feature to another when an expert believes that knowing one feature can change his beliefs that another feature will take on its possible values even when the diagnosis is known. Probabilities are used to quantitate the beliefs asserted by the expert.

C.3 Macintosh II Workstations

We are currently moving the Pathfinder system into the Macintosh II environment. Thus the SUMEX-AIM resource will only be used for file archival and retrieval and for communication between the Stanford and USC Pathfinder research teams.

D. Publications Since January 1984

1. Horvitz, E.J., Heckerman, D.E., Nathwani, B.N. and Fagan, L.M.: *Diagnostic Strategies in the Hypothesis-directed Pathfinder System, Node Pathology.* HPP Memo 84-13. Proceedings of the First Conference on Artificial Intelligence Applications, Denver, Colorado, Dec., 1984.

2. Heckerman, D. E., and Horvitz, E. J., "The Myth of Modularity in Rule-based Systems," in Uncertainty in Artificial Intelligence, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland, New York, 1987.

3. Horvitz, E.J., Heckerman, D.E., Nathwani, B.N. and Fagan, L.M.: *The Use of a Heuristic Problem-solving Hierarchy to Facilitate the Explanation of Hypothesis-directed Reasoning.* KSL Memo 86-2. Proceedings of MedInfo, Washington D.C., October, 1986.

4. Horvitz, E. J., "Toward a Science of Expert Systems," Invited Paper, Computer Science and Statistics: Proceedings of the 18th Symposium on the Interface, American Statistical Association, March, 1986, pgs. 45-52.

5. Heckerman, D.E., "An Axiomatic Framework for Belief Updates," in Uncertainty in Artificial Intelligence, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland, New York, 1987.

6. * Heckerman, D. E., and Horvitz, E. J., "The Myth of Modularity in Rule-based Systems," in Uncertainty in Artificial Intelligence, Vol. 2, J. Lemmer, L. Kanal, ed., North Holland, New York, 1987.

7. * Heckerman, D.E., and Horvitz, E.J., "On the expressiveness of rule-based systems for reasoning under uncertainty," Proceedings of the National Conference on Artificial Intelligence, Seattle, Washington, July, 1987.

8. * Horvitz, E. J., Heckerman, D. E., Langlotz, C. P., "A framework for comparing alternative formalisms for plausible reasoning," Proceedings of the AAAI," August, 1986, Morgan Kaufman, Los Altos, CA, 1986,

9. * Horvitz, E.J., Breese, J.S., Henrion, M., Decision Theory in Expert Systems and Artificial Intelligence, International Journal of Approximate Reasoning, Elsevier, N.Y. July, 1988.

10. * Heckerman, D.E., An Evaluation of Three Scoring Schemes, Proceedings of the 4th AAAI Workshop on Uncertainty in Artificial Intelligence, Minneapolis, MN., (to appear August 1988).

11. * Horvitz, E.J., A Multiattribute Utility Approach to Inference Understandability and Explanation, Tech. Report, KSL-28-87, Knowledge Systems Laboratory, Stanford, California, March, 1987.

E. Funding Support

Research Grant submitted to National Institutes of Health
Grant Title: "Computer-aided Diagnosis of Malignant Lymph Node Diseases"
Principal Investigator: Bharat Nathwani
Funding for three years from the National Library of Medicine
1 RO1 LM 04529

$766,053 (direct and indirect)

Professional Staff Association, Los Angeles County Hospital, $10,000.

University of Southern California, Comprehensive Cancer Center, $30,000.

Project Socrates, Univ. of Southern Calif., Gift from IBM of IBM PC/XT.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

Because our team of experts are in different parts of the country and the computer scientists are not located at the USC, we envision a tremendous use of SUMEX for communication, demonstration of programs, and remote modification of the knowledge base. The proposal mentioned above was developed using the communication facilities of SUMEX.

B. Sharing and Interaction with Other SUMEX-AIM Projects

Our project depends heavily on the techniques developed by the INTERNIST/CADUCEUS project. We have been in electronic contact and have met with members of the INTERNIST/CADUCEUS project, as well as been able to utilize information and experience with the INTERNIST program gathered over the years through the AIM conferences and on-line interaction. Our experience with the extensive development of the pathology knowledge base utilizing multiple experts should provide for intense and helpful discussions between our two projects.

The SUMEX pilot project, RXDX, designed to assist in the diagnosis of psychiatric disorders, is currently using a version of the Pathfinder program on the DEC-20 for the development of early prototypes of future systems.

C. Critique of Resource Management

The SUMEX resource has provided an excellent basis for the development of a pilot project. The availability of a pre-existing facility with appropriate computer languages, communication facilities (especially the TELENET network), and document preparation facilities allowed us to make good progress in a short period of time. The management has been very useful in assisting with our needs during the start of this project.

III. RESEARCH PLANS

A. Project Goals and Plans

Collection and refinement of knowledge about lymph node pathology

The knowledge base of the program is about to undergo revision by the experts, and then will be extensively tested. We will be extending the knowledge base to capture important probabilistic dependencies among features. We will also begin to evaluate the knowledge base and inference techniques.

Other possible extensions include: developing techniques for simplifying the acquisition and verification of knowledge from experts, and creating mapping schemes that will facilitate the understanding of the many classifications of non-Hodgkin's lymphomas. We will also attempt to represent knowledge about special diagnostic entities, such as multiple discordant histologies and atypical proliferations, which do not fit into the classification methods we have utilized.

Representation Research

We have been working to enhance the knowledge base by structuring features so that overlapping features are not incorrectly weighted in the decision making process, implementing new methods for scoring hypotheses, and creating appropriate explanation capabilities.

B. Requirements for Continued SUMEX Use

We are currently dependent on the SUMEX computer for file storage and archival, and for communication. We have transferred the program over to Portable Standard Lisp which is used by several users on the SUMEX system. While the switch to workstations has lessened our requirements for computer time for the development of the algorithms, we will continue to need the SUMEX facility for the interaction with each of the research locations specified in our NIH proposal. The SUMEX community version of the program is stored on the SUMEX mainframe for use by non-Stanford users.

C. Requirements for Additional Computing Resources

Most of our computing resources will be met by the 2060 plus the use of the Macintosh II workstations. We will need additional file space on the 2060 as we quadruple the size of our knowledge base through the construction of multiple knowledge bases. We will continue to require access to the 2060 for communication purposes, access to other programs, and for file storage and archiving.

D. Recommendations for Future Community and Resource Development

We encourage the continued exploration by SUMEX of the interconnection of workstations within the mainframe computer setting. We will need to be able to quickly move a program from workstation to workstation, or from workstation back and forth to the mainframe. Software tools that would help the transfer of programs from one type of workstation to another would also be quite useful. Until the type of workstations that we are using in this research becomes inexpensive, we will continue to need a machine like SUMEX to provide others with a chance to experiment with our software.

IV.D.3. Knowledge Engineering for Radiation Therapy

KNOWLEDGE ENGINEERING FOR RADIATION THERAPY

Ira J. Kalet, Ph.D.
Witold Paluszinski
University of Washington
Seattle, Washington

I. Summary of Research Program

A. Project Rationale

We are developing an expert system for planning of radiation therapy for head and neck cancers. The project will ultimately combine knowledge-based planning with numerical simulation of the radiation treatments. The numerical simulation is needed in order to determine if the proposed treatment will conform to the goals of the plan (required tumor dose, limiting dose to critical organs). The space of possible radiation treatments is numerically very large, making traditional search techniques impractical. Yet, with modern radiation therapy equipment, the design of treatment plans might be significantly aided by automatically generating plans that meet the treatment constraints. The project will result in systematization of knowledge about radiation treatment design, and will also provide an example of how to represent and solve design problems with a knowledge based system.

B. Medical Relevance and Collaborations

Radiation therapy has shown dramatic improvement in the cure rate for many tumor sites in the last two decades. Much of this can be attributed to the improved penetration capability of modern megavoltage X-ray machines. These high energy beams can deliver high tumor doses without overdosing surrounding tissue in many cases. However, they are typically used in very limited ways, because of the lack of suitable simulation systems to compute the dose distribution for any but a few narrow choices of treatment geometry. In the last few years these simulation systems have been extended to the full range of geometric treatment arrangement that any therapy machine is capable of. Thus it would be valuable to be able to generalize our knowledge of treatment technique by exploring these expanded possibilities. In addition, even treatments with standard geometries can be very complex, and it is tedious to explore all of them individually. A knowledge-based system can generate a few "best" plans which satisfy the constraints and allow more time for the physician to evaluate the options, or make minor adjustments for optimization.

Since cancer treatment is a multi-disciplinary approach involving surgery and chemotherapy as well as radiation, it is important to coordinate this work with knowledge-based program projects in those areas. Most significant is the ONCOCIN project, which addresses management of patients on chemotherapy protocols.

This project has some relevance to computer science as well, in that our approach, if successful, may contribute to a better understanding of design problem solving with knowledge-based systems.

C. Highlights of Research Progress

In the past year, we have made further additions to the rule database for details of head and neck cancer treatment. In addition we have examined the issues of control strategy associated with using prototypes in planning. The major effort has been to rewrite the entire program in COMMON LISP on a VAXstation II/GPX, and to use the Portable Common LOOPS (PCL) package, which appears to be on the way to becoming a standard (Common Lisp Object System - CLOS). Following the successful rewrite, we have added code to allow the lisp program to run the dose computation program as a subprocess, and exchange messages with that subprocess. The expert system can now create a plan, run the dose computation and receive results, analyzed into a form that will facilitate plan refinement.

Our expert system now has about two hundred rules, a two-level (agenda-based) control strategy, and about ten prototypes for plan construction. It is now implemented in COMMON LISP, on a VAX running the VMS operating system. This environment was chosen because it is also the environment used for a graphic simulation system that does radiation dose calculations for arbitrary treatment plans. The dose calculation is needed to determine whether a plan meets the treatment goals set by the system in its early phases of planning.

D. List of Relevant Publications

1. I. Kalet and W. Paluszynski: A Production Expert System for Radiation Therapy Planning. Proceedings of the AAMSI Congress 1985, May 20-22, 1985, San Francisco, California. Edited by Allan H. Levy and Ben T. Williams. American Association for Medical Systems and Informatics, Washington, D.C., 1985.

2. W. Paluszynski and I. Kalet: Radiation Therapy Planning: A Design Oriented Expert System. WESTEX-87 (Western Conference on Expert Systems), Anaheim, California, June 2-4, 1987.

3. I. Kalet and J. Jacky: Knowledge-based Computer Simulation for Radiation Therapy Planning. Proceedings of the Ninth International Conference on the use of Computers in Radiotherapy, Scheveningen, the Netherlands, June 1987. North Holland, 1987.

4. I. Kalet and W. Paluszynski: Knowledge-based Computer Systems for Radiotherapy Planning. American Journal of Clinical Oncology. In press.

II. Interactions with the SUMEX-AIM Resource

Our main use of the SUMEX-AIM resource has been as a means to be in contact with other researchers working on AIM projects. The existence of a mailbox at SUMEX-AIM has made it much easier for colleagues at other institutions to communicate with us, and has been valuable in assisting us with organizing the AIM Workshop for 1987.

We have had a great deal of contact with members of the ONCOCIN project and other groups. This has been valuable to us in stimulating creative approaches to our project.

III. Research Plan

A. Project Goals and Plans

We plan to continue to acquire rules and develop our current expert system. This includes solving problems of use of prototypes, satisfaction of constraints by some kind of backtracking search, and incorporating evaluation of plans by using the results of the dose computation. This last idea involves coupling the expert system with the dose computation system (written in PASCAL) in suitably efficient ways. We are implementing rules for plan improvement by modification of parameters, and by adding or deleting treatment plan components (treatment fields). Our long-term goal is to shape the user interface and improve the system performance to where it can provide assistance to clinicians in treatment design for patients in the normal course of treatment.

B. Justification and Requirements for Continued SUMEX use

We foresee continued need to be in touch with other members of the AIM community, particularly projects centered at SUMEX. While we do not expect to use the computing resources of SUMEX directly, some more extensive communication and involvement is likely to be useful.

C. Plans For Other Computing Resources

The main computing resources for our project will continue to be local. We continue to develop the expert system code in VAX Lisp, an implementation of Common Lisp on the DEC VAXstation. We are currently using a VAXstation II/GPX. This appears to be a good choice to satisfy our need for high performance graphic simulation and a reasonable Lisp system. However, the resources for the dose computation may not be adequate as we incorporate more sophisticated computation models. As this develops, we hope to experiment with distributed systems, in which the dose computation may run on a remote resource, which may or may not be at SUMEX.

D. Recommendations for Future Community and Resource Development

Two areas will be of increasing importance to us in the future: communication capabilities (electronic mail and file transfer) and centralized databases. By centralized databases, we refer to the need for better maintenance of mailing lists, information about projects, and possibly on-line reports. Dr. Kalet's experience in organizing the AIM Workshop for 1987 demonstrated that electronic communication is invaluable, even in its present state, but in order to create a list to send announcements to, we expended many hours of manually cutting and pasting messages containing past lists and searching for up-to-date electronic mail addresses.

If fees for use of SUMEX resources were imposed, the main impact on our project would be one of increased isolation, unless we could find grant support for the fees.

IV.E. AIM Communications Users

One of the key functions of the SUMEX-AIM resource has been inter-community communications. As the CPU-cycle-provider part of SUMEX's role phases out according to our plan for the distributed AIM community, this communications role will become increasingly important. In a recent action, the AIM Executive Committee approved broad guidelines for admitting communications users to SUMEX-AIM and we are implementing that policy.

Over the past year, several of the projects that had been using SUMEX-AIM as their primary source of computing resources have obtained local resources and now use SUMEX only for communications contact with others in the AIM community. Those projects include:

- *Hierarchical Models of Human Cognition (CLIPR)*, under Professors Walter Kintsch and Peter Polson at the University of Colorado.
- *Problem Solving Expertise (SOLVER)*, under Professors Paul Johnson and William Thompson at the University of Minnesota.
- *RXDX*, under Professor Robert Lindsay at the University of Michigan.
- *Computer-Based Exercises in Pathophysiologic Diagnosis*, under Professor J. Robert Beck at Dartmouth College.
- *SOAR*, under Dr. Paul Rosenbloom at USC/ISI in Los Angeles (formerly Stanford University).
- *Logic Group projects (DART, Intelligent Agents, and MRS)*, under Professor Michael Genesereth at Stanford University

Other projects will be undergoing similar transitions this coming year and we will report these in the next annual report.

The following is a list of users and groups from the National AIM community who use the SUMEX-AIM resource primarily for communication. All of the people listed have logged on to SUMEX over the past report year or have their mail forwarded to other sites. Many have been investigators on past projects that made use of SUMEX as a CPU provider and are now doing their research computing on machines at their home institutions. No detailed reports on their research are included.

<AAAI-OFFICE>	American Association for Artificial Intelligence (AAAI Hdqtrs.)
<ALTENBERG>	ALTENBERG, Lee, Ph.D. Department of Biochemistry Stanford University Medical Center
<AMAREL>	AMAREL, Saul, Ph.D. Department of Computer Science Rutgers University
<ANZALDI>	ANZALDI, Emmanuele Instituto di Informatica e Sistemistica Universita' di Pavia, Italy

<BAKER>	BAKER, William R., Jr., Ph.D. Knowledge Research Associates Bradenton, Fl. 34203-0035
<BECK>	BECK, J. Robert, M.D. Dartmouth College, School of Medicine
<BELLAIRS>	BELLAIRS, Keith University of Minnesota
<BRAND>	BRAND, Tony Chemistry Department Trenton State College, NJ
<CHANDRASEKARAN>	CHANDRASEKARAN, B., Ph.D. Department of Computer and Information Science Ohio State University
<DARDEN>	DARDEN, Lindley, Ph.D. Committee on the History and Philosophy of Science University of Maryland
<DAVIS>	DAVIS, Randall, Ph.D. Massachusetts Institute of Technology
<DAVISON>	DAVISON, Daniel, B., Ph.D. Department of Biochemical and Biophysical Sciences University of Houston
<DIETTERICH>	DIETTERICH, Tom, Ph.D. Department of Computer Science Oregon State University
<DROGERS>	ROGERS, David Research Institute for Advanced Computer Science NASA Ames Research Center
<DUMAS>	DUMAS, Jean-Pierre, Ph.D. France
<FELTOVICH>	FELTOVICH, Paul J. Southern Illinois University School Of Medicine
<FREKSA>	FREKSA, Christian Institut fuer Informatik Technische Universitaet Muenchen, West Germany
<FRISSE>	FRISSE, Mark, M.D. Department of Internal Medicine Washington University School of Medicine
<GOERZ>	GOERZ, Gunther Friedrich-Alexander-Universitat Erlangen, West Germany

<HAMM>	HAMM, Greg H. European Molecular Biology Laboratory Heidelberg, W. Germany
<HEDRICK>	HEDRICK, Charles L., Ph.D. Center for Computer and Information Services Rutgers University
<HENRION>	HENRION, Max Carnegie-Mellon University
<HORSWILL>	HORSWILL, Ian D. University of Minnesota
<HUMPHREY>	HUMPHREY, Susanne M. National Library of Medicine
<ICRF>	EXOMOLGEN Project Directory c/o Walter BODMER, Ph.D. Imperial Cancer Research Fund Laboratories London, England
<JARONSON>	ARONSON, Jules National Library of Medicine
<KAIHARA>	KAIHARA, Shigekoto, M.D. University of Tokyo Hospital Tokyo, JAPAN
<KDUNCAN>	DUNCAN, Karen Health Information Systems Los Altos, California
<KERN>	KERN, Kevin B. Lab for Computer Science Research Rutgers University
<KINGSLAND>	KINGSLAND, Lawrence C., III, Ph.D. National Library of Medicine
<KULIKOWSKI>	KULIKOWSKI, Casimir, Ph.D. Department of Computer Science Rutgers University
<LAGADIC>	LAGADIC Pascal Clinique Kerfriden Chateaulin, France
<LANGRIDGE>	LANGRIDGE, Robert, Ph.D. Computer Graphics Laboratory University of California at San Francisco
<LEDERBERG>	LEDERBERG, Joshua, Ph.D. The Rockefeller University

<LESGOLD>	LESGOLD, Alan M., Ph.D. Learning Research & Development Center University of Pittsburgh
<LINDBERG>	LINDBERG, Donald, M.D. National Library of Medicine
<MASINTER>	MASINTER, Larry Xerox Corporation, Palo Alto Research Center
<MAXAM>	Allan Maxam, Ph.D. Harvard Medical School
<MAZZETTI>	MAZZETTI, Claudia Executive Director, American Association for Artificial Intelligence
<MCDONALD>	MCDONALD, Clement J. University of Indiana Medical Center
<MKENT>	KENT, Marty Learning Research & Development Center University of Pittsburgh
<MOEN>	MOEN, James Center for Research in Human Learning University of Minnesota
<NOVAK>	NOVAK, Gordon, Ph.D. Computer Science Department University of Texas at Austin
<PEARSON>	PEARSON, William R., Ph.D. Department of Biochemistry University of Virginia
<PJOHNSON>	JOHNSON, Paul E., Ph.D. Center for Research in Human Learning University of Minnesota
<POLSON>	POLSON, Peter G., Ph.D. Department of Psychology University of Colorado
<RGREENES>	GREENES, Robert, M.D. Department of Radiology Harvard Brigham and Women's Hospital
<SIEBER>	SIEBER, Willi, Ph.D. Sandoz Ltd., Switzerland
<SLAGLE>	SLAGLE, James University of Minnesota

<SOTOS>	SOTOS, John Division of Cardiology Johns Hopkins Hospital
<STEFIK>	STEFIK, Mark, Ph.D. Xerox Palo Alto Research Center
<SUWA>	SUWA, Motoi Tsukuba, JAPAN
<SZOLOVITS>	SZOLOVITS, Peter Massachusetts Institute of Technology
<THOMPSON>	THOMPSON, William B., Ph.D. Computer Science Department University of Minnesota
<VANBEMMEL>	VAN BEMMEL, Jan Vrije Universiteit Amsterdam, Netherlands
<WEBBER>	WEBBER, Bonnie Dept. of Computer and Information Sciences University of Pennsylvania
<WEISS>	WEISS, Sholom, Ph.D. Department of Computer Science Rutgers University
<WIPKE>	WIPKE, W. Todd, Ph.D. Department of Chemistry University of California at Santa Cruz
<ZINK>	ZINK, Sandra, Ph.D. Lawrence Berkeley Labs
<ZUKER>	ZUKER, Michael, Ph.D. National Research Council Ottawa, Canada

Appendix A

Knowledge Systems Laboratory Brochure

ARTIFICIAL INTELLIGENCE RESEARCH IN THE KNOWLEDGE SYSTEMS LABORATORY

Stanford University
Department of Computer Science
Department of Medicine

January 1988