

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

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

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

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

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

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

E. Funding Support

The MOLGEN grant is titled: MOLGEN: A Computer Science Application to Molecular Biology. It is NSF Grant MCS 78-02777. Current Principal Investigators are Edward A. Feigenbaum, Professor of Computer Science and Laurence H. Kedes, Investigator, Howard Hughes Medical Institute and Associate Professor of Medicine. The new grant (September 1980) will add Bruce G. Buchanan, Adjunct Professor of Computer Science, and Douglas Brutlag, Associate Professor Biochemistry as Co-PI's. MOLGEN is currently funded from 12/79-11/80 at \$153,959 including indirect costs and has had a total funding from 6/78-3/81 at \$294,476 including indirect costs.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

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

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

The combination of the excellent computing facilities and the instant communication with a large number of experts in this field has been a

determining factor in the success of the MOLGEN project. It has made possible the near instantaneous dissemination of MOLGEN systems to a host of experimental users in laboratories across the country. The wide-ranging input from these users has greatly improved the general utility of our project.

We find it very difficult to find fault with any aspect of the SUMEX resource management. It has made it easy for us to expand our user group, to give demonstrations (through the 20/20 adjunct system), and to disseminate software to non-SUMEX users overseas. We do find that we are running moderately close to machine capacity both in size and in speed since our user group has been rapidly expanding during the last year.

III. RESEARCH PLANS

A. Project goals and plans

We have proposed further MOLGEN research in several broad categories: representation, planning, knowledge base development, and immediate applications to molecular biology. As would be expected, there will be much interaction among those general areas.

Representation

As part of the MOLGEN effort, a new representation package, the Units System, has been developed and tested. Its basis was mainly theoretical; we now have the opportunity to improve it from the practical considerations of a large knowledge base containing many different types of information. We expect to learn which features are important and which are window-dressing. These findings will increase in importance as many other problem-solving systems using large domain-specific knowledge bases are developed.

The MOLGEN knowledge base will serve as a laboratory for this research. Among the issues we would like to explore are:

1. MOLGEN currently uses the hierarchy representation features of the Units System for both acquisition and design. Will this continue to be practical as the knowledge base grows, or will the two representation functions have to be divorced?
2. The Units System allows different types of knowledge, e.g. numbers and nucleic acid sequences, to be described and stored in different manners. How much diversity is useful, both from the viewpoint of the representation system and from the viewpoint of the user?
3. Will new features become necessary to make large knowledge bases "perusable" by the human expert describing his domain? Is there some point at which graphics are needed for the expert to have a good grasp of what the system already knows?

Planning

Both of the two problem solving methods developed in MOLGEN have shown promise. We plan to keep pushing their development until we know their respective limitations and until a practical laboratory tool results. As was previously mentioned, we will combine the two planning methods to produce a system which should produce substantially higher performance than either of its two components.

The current experiment design systems are not designed to take an already existing laboratory plan and determine if the plan will satisfy some stated goal. We have proposed using the knowledge base to simulate the result of applying each step of a plan in succession to see if the experiment goal really would be achieved. This sort of a plan verifier will serve to take scientist-designed plans and provide guidance on whether the plan will work before it is actually tried in the laboratory.

The plan verifying system will be extended to become first a plan optimizing system and then a plan debugging system. Plan optimization will involve both domain-specific heuristics about how particular steps interact and domain-free heuristics about what good experiment designs should look like. The plan optimizer will make minor changes and introduce subgoals in order to take an already working experiment design and make it more efficient, convenient, reliable, or inexpensive. The knowledge base already contains most of the raw information humans use to make optimization decisions. The research is in developing the proper methods to make automated use of this knowledge.

Plan debugging means taking a partially working experiment design and finding and fixing any errors in it. This involves aspects of both verification and optimization as well as new error-correction heuristics. According to Feitelson and Stefik, the serendipity of the experimental laboratory also contributes greatly to plan debugging. Extending the MOLGEN design systems to become execution monitoring systems that can note and take advantage of this serendipity will be a major research effort of about thesis level in magnitude.

Knowledge Base Acquisition and Development

The current MOLGEN knowledge base is the result of over a man-year of effort by Professors Douglas Brutlag and Laurence Kedes and Drs. Peter Friedland and John Sninsky. It will continue to grow and improve throughout the term of the new proposal with the full time work of Dr. Rene' Bach. By the end of the period covered in the proposal the knowledge base will be in itself a useful tool for teaching, information retrieval, and sequence analysis. It will be expert in some of the most important areas of molecular biology. It will be especially proficient in those judgmental heuristics that guide technique selection as an experiment is being designed.

A major new research goal is to provide a facility for self-improvement of the knowledge base. When the design system produces a plan that is especially efficient or innovative, it would be useful to

generalize and save that plan so that it can drive future problem-solving without having to be reinvented. The generalization and learning process has roots in the MACROPS work in STRIPS.

Having such a capability would mean that the experiment design system would be a learning system, able to continuously improve its knowledge base. There are two main research questions inherent in the problem: how to recognize when a plan is worth saving, and how to decide how general to make it while still retaining its utility.

There are several possible measures of plan "worthiness." One would be whether the plan performed dramatically better than previous plans (e.g. it may have decreased the time to perform an experiment by an order of magnitude). Another would be related to how difficult it was for the system to create the plan. In other words, the plan should be saved because it would take a long time to find it again. The question is an experimental one; the research will involve trying many heuristics and balancing the improvement in system planning performance against the growth of an unwieldy and overly constrained knowledge base.

The question of how general to make the plan and how to parameterize it should also be solved experimentally. There will be trade-offs between how frequently the plan is used and what percentage of the time it will lead to a useful instantiated experiment design.

Another research goal is to use the knowledge base and experiment design system as a testbed for an automated performance evaluation system. The goals of such a system are quite general: to determine exactly how well the system is making use of the knowledge base, and how suitable the knowledge base is for the task at hand.

Among the specific questions a performance evaluation system for MOLGEN might answer are:

1. Is the system overlooking skeletal plans that it should find?
2. Is it needlessly considering many poor alternative plans?
3. Is it poorly modelling the consequences of plan steps?
4. In what areas of the knowledge base are decision heuristics weak or missing?
5. What types of knowledge are hardly ever being used?

All of these questions should be generalizable to many other knowledge-based problem-solving systems. Since the construction of large, expert knowledge bases is such a difficult task, the feedback from the evaluation of the use of these knowledge bases will be invaluable to future system builders.

Applications to Molecular Biology

The direct applications of MOLGEN to the field of molecular biology fall into three categories: knowledge base development and experiment design, analysis of nucleic acid sequences, and miscellaneous tools.

Knowledge Base Development and Experiment Design

The original and principal goal of the MOLGEN project is to provide a sophisticated experiment planning program containing an extensive knowledge base in the domain of molecular biology. As described above, our progress towards this goal has succeeded in the development of an extensive outline of this broad domain with emphasis on the myriads of analytical laboratory techniques that exist in this field. Using this knowledge base, MOLGEN is now capable of designing a number of sophisticated analytical experimental procedures. The procedures designed by the system are those already utilized in the laboratory, indicating that the knowledge base contains the correct sorts of heuristics to produce at least competent experiment designs. The limited scope of the current knowledge base provides a constraint on the originality of plans that can be produced; the most novel plans designed by humans are those which draw from many different, perhaps unrelated, knowledge sources.

Another success of the knowledge base concerns the organization of the information about each experimental technique. Because of the great flexibility of the Unit System, it is easy for the domain experts to modify and expand the existing information about each entity. We are continuously fine tuning the type of information contained within the knowledge base, in both content and in organization, during the actual knowledge acquisition phase.

We now propose to attack problems in synthetic molecular biology. We feel that by focusing our efforts on this subject we can assure an extensive repertoire of knowledge for that particular type of problem. This will also allow the planning algorithms to develop more sophisticated plans in the particular area. We have chosen to develop a knowledge base dedicated to the problem of cloning specific genes by recombinant DNA techniques. We have chosen this problem for four reasons: it is one of the most widely used methods in molecular biology today; most of our existing knowledge base is relevant to this problem; both of our current planning algorithms have been successful on either this problem (Stefik's thesis) or closely related problems of analysis of recombinant DNAs (Friedland's thesis); and because the method can be readily divided into four limited subdomains. These include choice of vectors, method of linking foreign DNA to the vector, transformation of host cells with the recombinant DNAs, and selection of the recombinant DNA containing the gene of interest.

We will describe current methods for cloning genes in both eukaryotes and prokaryotes, using methods in which one can select either for the vector or the inserted gene, and we will describe all the known methods of selecting for genes including direct functional selection, hybridization methods and expression of specific gene products. In addition to specifying the starting population or DNA sample and the ultimate goal, we will allow the user to specify certain subgoals or substrategies.

Analysis of Nucleic Acid Sequences

Our goal is to provide powerful, but easily used programs for the problem of the recognition of biologically significant patterns within nucleotide sequences. To make a set of programs both powerful and easy for a novice to use they must be interactive, self-documenting, and have easy to understand output formats. It also helps tremendously if they are very rapid so that they may be utilized online with nearly instantaneous feedback concerning the progress of the comparison. For this reason we have chosen to utilize the search algorithm developed by Korn and Queen and to convert it to an interactive form. This program was originally designed to provide for speed of comparison of very long nucleotide sequences while still allowing a degree of sophistication within the matching procedure. The algorithm compares two sequences beginning at every position where they share at least a dinucleotide but only carries the comparison as far as certain criteria of matching are allowed. This method, while lacking the sophistication of algorithms that potentially simulate evolutionary steps in the divergence of two sequences or the energetics of the pairing of single-stranded regions of dyad symmetry, is capable of detecting all statistically significant homologies or dyad symmetries given any level of significance desired. Unfortunately it is not capable of comparing more than two sequences at a time nor giving a quantitative measure of the divergence or relatedness of those two sequences. It merely describes the probability of each homology in terms of that expected for a random sequence of a given length and base composition.

Our improvements to the program have included converting it into SAIL and making it interactive. Whenever a user is in doubt about the next step he merely enters a ? and his options at that point are explained. We have also considerably improved the statistical calculations so that the probabilities and expectation frequencies that are determined for a homologous region are based not only on the length of the sequences being compared, but also on the base composition and on the exact algorithm being used in the search itself. Finally we have markedly improved the output displays so that mismatches are indicated with stars and base pairs in dyad symmetries with bars. We have done all of this without any overhead in terms of execution time so that the program executes almost without delay in a time-sharing environment.

We propose to improve our current sequence analysis capabilities by implementing more sophisticated algorithms within the interactive framework. For instance the pattern recognition algorithm of Sellers is currently being implemented in C language at Rockefeller University by Dr. Bruce Erickson. We believe that this program would be a useful addition to our current armory in that it would allow us an accurate metric of relatedness of two sequences which is essential for building phylogenetic trees. This would be the first step towards the comparison of more than one sequence.

We would also like to develop methods for determining the secondary structure of single-stranded RNAs. The most commonly used methods are often limited to short nucleotide regions because of the complexity of the energy calculations for large numbers of comparisons. By first utilizing a

rapid method for finding homologous sequences or dyad symmetries, perhaps guided by statistical significance of very low stringency, one might be able to rapidly eliminate most of the fruitless comparisons. By then examining the resulting culled homologies by a set of heuristics concerning their additivity, extension, or exclusiveness, we could order them in terms of their biological significance. This would automate some of the tedious cutting and patching of homologies and dyad symmetries in which molecular biologists are now involved even after they have made comparisons with a computer. With respect to calculations of the thermal stability of symmetric regions it would reduce the total time of calculation by orders of magnitude. In other words, we would use a comparison algorithm based more on biological intuition than calculation in order to find the most profitable regions to apply the more quantitative methods of biophysics.

We would further hope to automate the development of phylogenetic trees utilizing these sequence comparison algorithms. Once quantitative measures of relatedness are obtained in all pairwise combinations, then the matrix methods for the generation of the trees and the lengths of the branches is rather straightforward. These calculations are not likely to need any intelligent heuristics for their determination since they are defined analytically and they are rapid compared to the calculations involved in determining the relatedness of the sequences in the first place.

Miscellaneous Tools

Restriction Digest Analysis

One of the best examples of the utility of the application of heuristics and production rules to problems of molecular biology is the GA1 program, developed in this project, for the analysis of restriction endonuclease digests. Determining restriction maps of even simple DNA structures from restriction enzyme digest data can require consideration of millions of possible structures. The application of heuristic methods simplifies the analysis by orders of magnitude allowing solutions to complex problems and even simplifying the amount of data that must be collected to ensure a unique solution. These methods have even resulted in the proposal of a new experimental method for the analysis of restriction data.

GA1 is a program which determines all possible organizations of restriction fragments based on restriction endonuclease digests with single, double, and triple combinations of enzymes. The program contains an intelligent hypothesis generator and a set of production rules which allow it to generate and evaluate hypothetical restriction maps which are consistent with all of the data. These rules dramatically reduce the total number of possible structural candidates that must be both generated or evaluated.

Modern laboratory methods for determining restriction maps include end labeling procedures and two dimensional cross hybridization procedures. In order to extend the program GA1 to cover this kind of data we propose to

be able to set up initial constraints on the locations of all restriction sites in certain local regions of the hypothetical restriction map. Such initial conditions (regional constraints) would be useful not only for entering data obtained from partial digestion of end labelled DNA segments, but would also be very useful if the complete nucleotide sequence were known for a particular region. Such conditions are often found in recombinant DNAs in which the nucleotide sequence of the vector is completely known.

Another improvement in GA1 which would both simplify and extend its use would be to allow the user to describe the complete restriction map determined previously for a limited number of restriction enzymes and then to enter digestion data for new enzymes, singly and in combination with the previously analyzed sites. These initial conditions would impose global constraints over the entire map. Global constraints will not be as readily implemented as the regional constraints described above.

If sufficient programming support is available we would also like to attempt to apply the hypothesis generating and production rule pruning approach to the analysis of two dimensional restriction data. In this method, radioactively labeled DNA segments generated from a DNA by a one restriction enzyme are hybridized to nonradioactive fragments generated by a second restriction enzyme thus indicating which pairs of fragments are homologous and hence overlapping. Currently the typical analysis is a data driven approach of finding a continuous path among all the overlapping DNA fragments cataloged by this experimental procedure. A model driven approach should extend this already powerful method. While the two dimensional cross-hybridization method only allows the generation of maps for two enzymes at a time, maps generated from all possible pairwise combinations of any set of enzymes are possible by analogy with the standard one dimensional method. Furthermore, by alternately labeling the fragments from either restriction enzyme and hybridizing those fragments to unlabeled fragments derived from the second enzyme in both directions, sufficient data should be obtained in order to overcome most mapping ambiguities which are usually the downfall of this method. Utilization of the model driven approach to the cross-hybridization procedure will also allow the generation of restriction maps of much longer DNAs than currently possible.

Synthesis of Specific Nucleic Acid Molecules

The MOLGEN knowledge base contains complete sequence information for all published and many unpublished nucleic acid molecules. It also knows about restriction endonucleases and their cutting sites and about ligation methods for rejoining nucleic acid fragments. We see potential use for this knowledge in designing synthetic pathways for the in vitro production of specific target molecules. This may actually be considered a part of the main experiment design effort, but the problem is important enough to make an independent specialized system desirable.

Currently, three major methods are used by molecular biologists to select specific sequences of interest from a recombinant DNA "library". The most widely used method uses isolated messenger RNA as radiolabeled

probe to detect complementary DNA sequences in the recombinant molecules. This requires prior isolation of the mRNA which, unfortunately, is not always easily obtained. Secondly, and perhaps having the most long-term potential, are methods to select by expression in the host cell of the sought for functions. Such an approach will necessarily be limited to genes that can be made to supplement or rescue host functions. The problems of expression of eukaryotic genes in prokaryotic hosts may never be soluble because of the gene-splicing dichotomy. The utility of eukaryotic host-vector systems is now established but selection will still depend on prior creation of host mutants or use of immunological colony (or plaque) screening techniques still to be developed.

A third approach has been to use relatively short chemically synthesized oligonucleotide segments that are complementary to the gene of interest. The probe is used to select genomic clones of recombinants containing specific protein coding sequences. In theory, if the amino acid sequence is known, appropriate probes can be constructed. The techniques for chemical oligonucleotide synthesis are difficult and laborious. We propose a different approach using the recombinatorics of the computer stored and generated nucleotide sequences of all known DNA molecules. If the amino acid sequence of the protein whose gene is desired is known, then a computer assisted search through those sequences will attempt to locate oligonucleotides that could code for a short segment of that protein. By taking advantage of third base degeneracy and knowledge of restriction endonuclease cutting and splicing, constructions of natural oligonucleotides will be suggested. An intelligent algorithm might locate more than just one or two short segments capable of forming molecular hybrids with the DNA sequences being sought and these might be linked in a spaced out manner to provide a more powerful probe.

B. Justification and requirements for continued SUMEX use.

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

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

scientific computing resources to a national audience of bio-medical research scientists.

9.1.5 MYCIN Project

MYCIN Project

Edward. H. Shortliffe, M.D., Ph.D.
Department of Medicine
Stanford University Medical School

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

I. Summary of Research Program

A. Project Rationale

The MYCIN Project is a set of subprojects, each devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The project retains the name of our first system, the MYCIN program, but has grown to involve five interrelated sub-projects (MYCIN, EMYCIN, CENTAUR, GUIDON, and ONCOCIN), each of which will be discussed in the sections that appear below.

Our first system, MYCIN, is an interactive consultation program which gives physicians antimicrobial therapy recommendations for patients with infectious diseases. The system must often decide whether and how to treat a patient before definitive laboratory results are available. It must recommend a therapeutic regimen which minimizes the risk of toxic side-effects while covering for all organisms which are likely to be causing the infection. The relevant knowledge is stored in production rules, and the system currently has rules for treating bacteremias (blood infections) and meningitis. There has already been early work on the codification of cystitis knowledge. The primary goal of the project has been to develop a program which can provide advice similar in quality to that given by a human infectious disease consultant. Formal evaluations of the program's recommendations for patients with bacteremia or meningitis have shown that this goal has been achieved. We have also sought to develop a system that is easy to use and acceptable to physicians. To accomplish this, numerous human engineering features have been incorporated into the consultation. There is also an extensive explanation facility which enables the system to explain its reasoning and to justify its recommendations.

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

- (1) to develop consultation systems for other domains (our generalized system-building tool is known as "Essential MYCIN", or EMYCIN, and has been applied in several new areas);
- (2) to explore other uses of the knowledge base (our tutoring system, GUIDON, uses the infectious disease knowledge in MYCIN

- to teach medical students about diagnosis and management of infections);
- (3) to continue to improve the interactive process, both for the developer of a knowledge-based system, and for the user of such a system (both EMYCIN and our newest system, ONCOCIN, have stressed simplified techniques for interacting with a knowledge base and entering data); and
 - (4) to experiment with using other knowledge representations in conjunction with the production rules used in MYCIN (our CENTAUR system is a modification to EMYCIN which uses prototypical descriptors of situations or disease states to guide and focus a consultative session).

B. Medical Relevance and Collaboration

The MYCIN program was designed to help alleviate the well-documented problem of antimicrobial misuse. We felt that MYCIN would be clinically useful when it was able to handle all major infections that are likely to be encountered in a hospital. Our success in developing a high performance program for meningitis and bacteremia has been documented in two articles by Dr. Yu listed in the publications section below. However, the system is not ready for clinical use because it does not have rules for the other areas of infectious disease. A very large investment in time and human resources is required to develop, test and formally evaluate a rule set for each major infection area.

By utilizing our EMYCIN system to collaborate on building the PUFF program, however, we learned that it is possible in a short period of time to develop a clinically useful consultation system using the domain-independent parts of MYCIN. EMYCIN has since been applied in a number of additional medical domains outlined below. Although EMYCIN was not used to build our new ONCOCIN program, the lessons learned in building prior production rule systems have allowed us to create a large oncology protocol management system in only eight months. Furthermore, we expect to have ONCOCIN used by Stanford oncologists before the end of 1980.

Finally, there is a growing realization that medical knowledge, originally codified for the purpose of computer-based consultations, may be utilized in additional ways that are medically relevant. Using the knowledge to teach medical students is perhaps foremost among these, and GUIDON continues to focus on methods for augmenting clinical knowledge in order to facilitate its use in a tutorial setting.

C. Highlights of Research Progress

MYCIN

Due to the departure of Dr. Victor Yu, the infectious disease expert who worked with us until recently, it has not been possible to expand the rule set into new areas of infectious disease. The 500 rules relating to

bacteremia and meningitis are sufficiently rich and complex, however, that they serve as a particularly challenging vehicle for testing the new computational methods we are developing. MYCIN is now totally implemented as an EMYCIN system. Hence, our active work on EMYCIN has been thoroughly tested using MYCIN and our extensive library of patient cases. Ongoing efforts to expand MYCIN or prepare it for clinical implementation, however, have been temporarily set aside to allow us to concentrate on the projects below.

EMYCIN

Much of the work in the past year has been devoted to improving EMYCIN's facilities for allowing a system builder to construct and debug a knowledge base for a consultation system. This has included extensive documentation of the concepts used in EMYCIN consultation systems, the support programs for developing the knowledge base, and features of a working consultation system.

A knowledge-base debugging package was developed to assist the system builder in the task of testing, refining, and validating the knowledge base. This package includes: 1) the EMYCIN explanation facility; 2) a program that automatically explains how the system arrived at the results of a consultation; 3) a program that reviews each result of a consultation, allowing the user to judge whether the result is correct, and assisting the user in refining the knowledge base in order to correct any errors noted in the result or in intermediate conclusions; and 4) a program that automatically compares the results of a consultation to stored "correct" results for the same case, and explains any errors in the conclusions.

An additional development in the last year is the EMYCIN "rule compiler." Once a consultation program is built, it becomes important that it perform efficiently. This is most noticeable in large programs such as MYCIN. Production rules, while convenient in their modularity, are not the best representation for speedy execution. We have thus developed a rule compiler as part of EMYCIN that transforms a program's production rules into a decision tree, eliminating the redundant computation inherent in a rule interpreter, and compiles the resulting tree into machine code. The program can thereby use an efficient deductive mechanism for running the actual consultation, while the flexible rule format remains available for acquisition, explanation, and debugging.

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

EMYCIN Applications

Several consultation systems have been written in EMYCIN. All but the most recent of these were developed in parallel with EMYCIN, and thus served to focus attention on certain features and shortcomings of the program to guide in its development. Their brief description here is intended to provide some indication of the range of potential applications of EMYCIN.

PUFF

The PUFF system performs interpretation of measurements from the pulmonary function laboratory. The project is a collaboration of a pulmonary physiologist, biomedical engineers, and Stanford computer scientists who had previous experience with the MYCIN program. The data from over 100 cases were used to create some 60 rules diagnosing the presence of pulmonary disease. These rules are used to create a complete report including the input measurements, other patient data, and the measurement interpretation. The system is a separate SUMEX project now, and is described in full elsewhere in this document.

HEADMED

The HEADMED program is an application of EMYCIN to clinical psychopharmacology. The system diagnoses a range of psychiatric disorders and can recommend drug treatment if indicated. Like PUFF, this project is a separate SUMEX project.

SACON

As a stronger test of domain independence, EMYCIN was applied to the completely non-medical domain of structural analysis. SACON (Structural Analysis CONsultation) provides advice to a structural engineer regarding the use of a large structural analysis program called Marc. The Marc program uses finite-element analysis techniques to simulate the mechanical behavior of objects. Engineers typically know what they want the Marc program to do, e.g., examine the behavior of a specific structure under expected loading conditions, but they do not know how the simulation program should be set up to do it. The goal of the SACON program is to recommend an analysis strategy; this advice can then be used to direct the Marc user in the choice of specific input data, numerical methods and material properties.

The performance of the SACON program matches that of a human consultant for the limited domain of structural analysis problems that was initially selected. To bring the SACON program to its present level of performance, about two man-months of the experts' time were required to analyze their task as consultants and formulate the knowledge base. About the same amount of time was required to implement and test the rules.

CLOT

A recent application of EMYCIN is CLOT, a system designed to diagnose disorders of the blood coagulation system of patients. It requests clinical evidence regarding an episode of bleeding, facts from the patient's general medical history, and the results of a battery of coagulation screening tests. From these data CLOT infers the presence and type of coagulation defect (if any) in the patient and then proceeds to make a refined diagnosis for any particular enzymatic deficiency or

platelet defect. These diagnoses can be used by a physician to estimate the severity and cause of a particular episode of bleeding, evaluate the effects of various anti-coagulation therapies on a patient, or estimate the pre-operative risk of a patient having serious bleeding problems during surgery.

CLOT was constructed by David Goldman, a medical student at the University of Missouri, with the help of James Bennett, a member of our Stanford group who is very familiar with EMYCIN. Following approximately 10 hours of discussion about the contents of the knowledge base, they entered and debugged in another 10 hours a preliminary knowledge base of some 60 rules. CLOT is now an ongoing project at the University of Missouri.

GUIDON

Bill Clancey's thesis (August '79) marked the completion of version one of the program. Key results include:

- (1) A language was developed for representing teaching expertise in the form of "Discourse Procedures"--sequences of rules that reflect dialogue patterns and are independent of the subject material to be taught. This representation was found to be suitable and convenient for incrementally developing a tutorial program.
- (2) Various teaching methods were demonstrated for carrying on a case method dialogue with a student who is solving a complex diagnostic problem. Meta-knowledge about the representation of the subject material made it possible to express these capabilities in a domain independent way.
- (3) The representation of subject material as modular production rules was studied and found wanting. Though rules conveniently separate relationships into readily accessible associations, an adequate knowledge base for teaching requires the addition of structural knowledge (clusters and patterns), support knowledge (underlying causal mechanisms), and strategical knowledge (managerial approaches).

Ongoing GUIDON research focuses on a number of issues:

The Student Model.

A revised student model has been designed to deal with the following questions:

- (1) Can the student USE the program? i.e., is he able to enter recognizable input?
- (2) Is the dialogue with the student COHERENT? i.e., are there recognizable patterns of student input and meaningful transitions between segments of behavior?

- (3) Is the student PASSIVE OR ACTIVE? i.e., does he use his own knowledge to solve the problem, or does he rely on the tutor's initiative and ability to provide help?
- (4) Does the student have a STRATEGY for solving the problem? i.e., is there some plan that organizes the student's data measurements and hypothesis selection?

Representation of Problem Solving Strategies.

One of the few formalized methods for teaching diagnostic strategies to medical students is a printed outline of data to collect. This outline is woefully inadequate as a teaching tool: it does not convey in itself the meaning or logic of the diagnostic process. Informal experiments with physicians have enabled us to formalize an ideal model of medical diagnostic strategy appropriate to our present domain of investigation (infectious meningitis). Work is underway to incorporate this model in MYCIN so that it "thinks like a clinician," and can thus be used to teach not only diagnostic rules, but human-usable methods for applying them.

Some surprising findings coming out of this investigation include the following:

- (1) Establishing the hypothesis space is accomplished by considering causal links that might be enabled in this patient (called "risk factors"). This can be considered to be a process of determining the topology of the problem--causal connections that may have a bearing on the disorder.
- (2) "Dropping back" is important to human problem solvers. In fact, hypothesis formation as we have observed it might be described as a process of maintaining a sense of the differential. Focusing and delving deeper is just a temporary phenomenon.

Acquisition of this strategical knowledge was greatly helped by analyzing protocols according to the structure/support/strategy framework we have established. This is one of the "knowledge engineering" results of our research.

CENTAUR

During the last year we have completed an implementation of PUFF using the augmented EMYCIN system known as CENTAUR. In this work, largely the effort of Jan Aikins, we have sought to strengthen the pure production rule representation of EMYCIN with additional focusing power provided by hypothesis "frames" or prototypes. CENTAUR now includes 24 prototypes and about 160 rules dealing with pulmonary disease. The system was tested on 100 cases from the files at Pacific Medical Center. CENTAUR agreed with two pulmonary physiologists 84 and 91 per cent of the time respectively on their diagnoses of pulmonary disease in the cases. (This was an improvement over PUFF, which had 74 and 85 per cent agreement with the two physiologists).

Basic AI research issues were also explored, such as the representation of control knowledge for computer consultations, and the explicit representation of the context in which knowledge is applied. Furthermore, the MYCIN explanation facility was expanded to include explanations of control processes, and to give explanations of the prototypes, as well as the rules.

Current CENTAUR research is concentrating on polishing and fine-tuning the PUFF implementation described above. Additional studies are contemplated to better define the precise reasons that CENTAUR has performed more accurately than PUFF on the 100 cases mentioned above. One expert collaborator, Dr. R. Fallat feels PUFF had performed less well because of the significant difficulties he has had in adding more rules and still keeping the knowledge base consistent. This was less difficult using the CENTAUR representation scheme.

Other research that will draw upon CENTAUR work includes the creation of additional applications systems using the CENTAUR prototype representation mechanism. One challenge will be to interface CENTAUR with the "context-tree" that is provided in EMYCIN, a problem that was not addressed in PUFF because it utilizes only a single context.

ONCOCIN

The oncology protocol management system, termed ONCOCIN after its domain of expertise and its historical debt to the MYCIN program, has achieved many of its early goals since work on the project began in July 1979. We are developing an interactive system to be used by oncology faculty and fellows in the Debbie Probst Oncology Day Care Center at Stanford University Medical Center. Our overall goals are:

- (1) to demonstrate that a rule-based consultation system with explanation capabilities can be usefully applied and gain acceptance in a busy clinical environment;
- (2) to improve the tools currently available, and to develop new tools, for building knowledge-based expert systems for medical consultation, and
- (3) to establish both an effective relationship with a specific group of physicians, and a scientific foundation, that will together facilitate future research and implementation of computer-based tools for clinical decision making.

The ONCOCIN research goals are directed both towards the basic science of artificial intelligence and towards the development of clinically useful oncology consultation tools. We have undertaken AI research with the following aims:

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

- (2) to extend the methods of rule-based consultation systems to interact with a large database of clinical information; and
- (3) to continue basic research into the following problem areas: mechanisms for handling time relationships, techniques for quantifying uncertainty and interfacing such measures with a production rule methodology, approaches to acquiring knowledge interactively from clinical experts, assessment of knowledge base completeness and consistency.

Our simultaneous clinical goal is to develop and implement a protocol management system, for use in the oncology day care center, with the following capabilities:

- (1) to assist with identification of current protocols that may apply to a given patient;
- (2) to assist with determining a patient's eligibility for a given protocol;
- (3) to provide detailed information on protocols in response to questions from clinic personnel;
- (4) to assist with chemotherapy dose selection and attenuation for a given patient;
- (5) to provide reminders, at appropriate intervals, of follow-up tests and films required by the protocol in which a given patient is enrolled;
- (6) to reason about managing current patients in light of stored data from previous visits of (a) the individual patients, or (b) the aggregate of all "similar" patients.

During the first year of our research, it has been our aim to develop a prototype of the ONCOCIN consultation system, drawing from the programs and capabilities of EMYCIN. We have also analyzed carefully the day-to-day activities of the Stanford oncology clinic in order to determine how to introduce ONCOCIN with minimal disruption of an operation which is already running smoothly. Finally, we have spent much of our time considering the most appropriate mode of interaction with physicians in order to optimize the chances for ONCOCIN to become a useful and accepted tool in this specialized clinical environment.

We chose the series of protocols for Hodgkin's and non-Hodgkin's lymphoma as the first detailed knowledge to be encoded in the ONCOCIN system. These were selected because they were developed at Stanford, because they are among our most commonly used protocols in light of our position as a major lymphoma treatment center, and because the protocols are complicated, with many subtle details depending upon the stage of disease, concomitant or preceding radiotherapy, and evidence for drug toxicity.

Although the program will eventually be used on a high-speed terminal with a specially designed interface (see below), we decided that the initial prototype should be a self-contained consultation system that would be modeled on the form of interaction used for EMYCIN consultation systems. We chose not to use EMYCIN itself to build the system, however, because we quickly encountered several special needs that were better handled using alternate representation and control schemes. Therefore, although there are portions of the EMYCIN code that we have been able to borrow, ONCOCIN is an entirely new program in which production rules are only one of several types of knowledge representation used.

Both our own experience, plus evidence in the medical computing literature, have suggested that physicians will be unlikely to use consultation systems if they fail to fit smoothly in the day's normal routine. With this in mind, we have carefully studied the current organization and flow of information within Stanford's oncology clinic. A detailed document has been prepared which describes the current clinic organization and the ways in which our system will interact with the current routine. Two principal concerns have been:

- (1) that ONCOCIN should initially have minimal impact on the current daily routine: record-keeping systems should not be altered, patient flow within the clinic should be unchanged, and the physicians working there should not be forced to depend on an operational computer system in order to get their work done;
- (2) that it should not take any EXTRA effort on the physicians' part for them to use the ONCOCIN system (other than the initial time required while they are trained how to use it); this implies that the use of ONCOCIN should replace some task that the physicians are currently doing.

Currently the clinic physicians are asked to fill out, by hand, the time-oriented flowsheets that are kept in the patient clinic records. These sheets are the basis for data analysis of all the clinical research that is based on chemotherapy protocols in the oncology clinic. All information needed by ONCOCIN is entered on this flowsheet. Thus we intend to capture the data needed for an ONCOCIN consultation by having the physician fill out the flowsheet at a computer terminal rather than by hand.

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

First we have decided to use high-speed CRT terminals (approximately 9600 baud) with auxiliary hard-copy devices. This will permit almost instantaneous screen filling and allow greater flexibility in the design of what is actually displayed. However, a program written in a powerful but slow language like INTERLISP is not able to service a high-speed terminal

adequately. For this reason, our interface program will be written in a faster compiled language (we are using PASCAL), and this program will need to communicate in turn with the INTERLISP reasoning program that comprises the rest of ONCOCIN. The design of this interprogram interaction is largely complete, but actual implementation of the ideas is just beginning.

Second, we want to minimize typing by the physician. EMYCIN systems have required a typewriter-compatible keyboard, but we do not feel this is reasonable if ONCOCIN is to be used on a daily basis by a large number of oncologists. Initially we examined light-pen and touch-screen technologies, but feel that these are either too expensive or too unreliable. Ultimately, working closely with experts in human factors, we developed a customized 21-character keypad which has been interfaced with a Datamedia terminal similar to those we have used for other development work. This keypad can be used by the physician to fill out the patient's flowsheet (which will be displayed on the screen at high speed), and there should be minimal if any need to use the terminal keyboard itself.

Finally, we want to maintain the explanation and justification capabilities which we have argued are crucial to the acceptance of clinical consultation systems. A specialized split-screen display has been designed which will enable the physician to enter patient data entries in one region while pertinent explanations are displayed in another.

D. Publications Since January 1979

- Kunz, J.C., Fallat, R.J., McClung, D.H., Votteri, B.A., Aikins, J.S., Nii, H.P., Fagan, L.M., Feigenbaum, E.A. Physiological rule-based system for interpreting pulmonary function test results. Memo HPP-78-154, Stanford Heuristic Programming Project, 1978. Also Proceedings of Computers in Critical Care and Pulmonary Medicine, IEEE Press, 1979.
- Yu, V.L., Buchanan, B.G., Shortliffe, E.H., Wraith, S.M., Davis, R., Scott, A.C., Cohen, S.M. Evaluating the performance of a computer-based consultant. Comput. Prog. Biomed. 9,95-102 (1979).
- Clancey, W.J. Tutoring rules for guiding a case method dialogue. Int. J. of Man-Machine Studies 11,25-49 (1979).
- Clancey, W.J. Dialogue management for rule-based tutorials. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, pp. 155-161, August 1979.
- Aikins, J.S. Prototypes and production rules: an approach to knowledge representation for hypothesis formation. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, Tokyo, Japan, August 1979.
- Fagan, L.M., Kunz, J.C., Feigenbaum, E.A., Osborn, J. J. Representation of dynamic clinical knowledge: measurement interpretation in the intensive care unit. Proceedings of the 6th Intl. Joint Conf. on Artificial Intelligence, Tokyo, Japan, August 1979.

van Melle, W. A domain-independent production-rule system for consultation programs. Proceedings of the 6th IJCAI, August 1979.

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

Yu, V.L., Fagan, L.M., Wraith, S.M., Clancey, W.J., Scott, A.C., Hannigan, J.F., Blum, R.L., Buchanan, B.G., Cohen, S.N. Antimicrobial selection by a computer -- a blinded evaluation by infectious disease experts. J. Amer. Med. Assoc. 242:1279-1282 (1979).

Shortliffe, E.H. Medical consultation systems: designing for doctors. To appear in Communication With Computers (M. Sime and M. Fitter, eds.), London: Academic Press, 1980.

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

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

Shortliffe, E.H. Clinical knowledge engineering: the MYCIN Project. Proceedings of the First Japanese Conference on Artificial Intelligence in Medicine, pp. 1-8, Tokyo, Japan, August 1979.

Clancey, W.J. Transfer of Rule-Based Expertise through a Tutorial Dialogue. Computer Science Doctoral Dissertation, Stanford University, August 1979.

Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. Knowledge engineering for infectious disease therapy selection. Proceedings of the Intl. Conf. on Cybernetics and Society, Denver, Colorado, October 1979.

Clancey, W.J., Shortliffe, E.H., and Buchanan, B.G. Intelligent computer-aided instruction for medical diagnosis. Proceedings of the Third Annual Symposium on Computer Applications in Medical Care, Silver Spring, Maryland, October 1979.

Fagan, L.M., Kunz, J.C., and Feigenbaum, E.A. Representation of dynamic clinical knowledge: measurement interpretation in the intensive care unit. Proceedings of the Third Annual Symposium on Computer Applications in Medical Care, Silver Spring, Maryland, October 1979.

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

Shortliffe, Edward H. Consultation systems for physicians: the role of artificial intelligence techniques (invited paper). Proceedings of the 3rd Annual Meeting of the Canadian Society for the Computer Simulation of Intelligence, Victoria, British Columbia, May 1980.

E. Funding Support

Grant Title: "Research Program: Biomedical Knowledge Representation"
Principal Investigator: Edward A. Feigenbaum
Co-Principal Investigator (ONCOCIN Project): Edward H. Shortliffe
Agency: National Library of Medicine
ID Number: 1 P01 LM 03395
Term: July 1979 to June 1984
Total award: \$497,420
Current award (1979-1980): \$99,484

Grant Title: "Knowledge-Based Consultation Systems"
Principal Investigator: Bruce G. Buchanan
Agency: National Science Foundation
ID Number: MCS-7903753
Term: July 1979 to June 1980 (plus 6 months)
Total award: \$146,152
Current award (1979-1980): \$73,659

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

Grant Title: "Symbolic Computation Methods For Clinical Reasoning" (RCDA)
Principal Investigator: Edward H. Shortliffe
Agency: National Library of Medicine
ID Number: NIH 1K04 LM00048
Term: July 1979 to June 1984
Total award: Dollar amount negotiated annually
Current award (1979-1980): \$39,285

Grant Title: "Explanatory Patterns In Clinical Medicine"
Principal Investigator: Edward H. Shortliffe
Agency: Kaiser Family Foundation
Term: July 1979 to December 1980
Total award: \$20,000

II. Interaction With the SUMEX-AIM Resource

A. Medical Collaborations and Program Dissemination Via SUMEX

A great deal of interest in both MYCIN and EMYCIN have been shown by the medical and academic communities. For two years in succession we have been invited by the American College of Physicians to demonstrate MYCIN at the organization's annual meeting (San Francisco, March 1979, and New Orleans, April 1980). The physicians have uniformly been enthusiastic

about the program's potential and what it reveals about one current approach to computer-based medical decision making. In both cases, the demonstrations were performed on-line using network access to the SUMEX computer. There has also been significant growing interest in medical AI and MYCIN from colleagues in Japan. We were asked to demonstrate MYCIN from Tokyo during the 6th International Joint Conference on Artificial Intelligence held in August 1979. Access to SUMEX via a trans-Pacific TYMNET link worked very well and permitted large numbers of Japanese and other conference attendees to observe MYCIN demonstrations and experiment with the program themselves. Then, for three weeks in November 1979, Dr. Shortliffe returned to Japan as a visitor at the Tokyo Metropolitan Institute of Medical Sciences. This visit permitted an intensive period of exchange regarding MYCIN, EMYCIN, and the related work being done by the Japanese.

Several teachers have also asked to use MYCIN in their computer science or medical computing courses. For example, Prof. Carl Page of Michigan State University, Dr. Peter Szolovits of MIT, and Dr. Steven Zucker of McGill University in Montreal have demonstrated the MYCIN program in their university classes. Dr. Harold Goldberger of MIT made extensive use of the MYCIN program in his study of medical AI programs. Dr. Ves Morinov of the Norwegian Computing Center has used the MYCIN program to demonstrate the benefits of using a rule-based representation for consultation systems. Dr. Martin Epstein used MYCIN as one of the representative systems he demonstrated to students who took the clinical elective on medical computing at the NIH during the summer of 1979.

GUEST users who have recently requested access to MYCIN have come from such diverse locations around the country as the Brain Research Institute (UCLA), University of Texas, Stevens Institute of Technology, University of New Mexico, Columbia University, Systems Science Institute (Louisville), Naval Postgraduate Institute (Monterey, Ca.), Texas Women's University, IBM Scientific Labs, and Alta Bates Hospital (Oakland, Ca.).

EMYCIN has also generated a great deal of interest in the academic and business communities. We have been in frequent contact with Bud Frawley and Philippe Lacour-Gayet of Schlumberger, Chuck Brodnax and Milt Waxman of the Hughes Aircraft Corporation, and Harry Reinstein from IBM Scientific Research Center. Two students at the Naval Postgraduate School in Monterey, working under the direction of Colonel Ronald J. Roland, have been developing an EMYCIN system in the domain of selecting decision aids for solving problems in business organizations. The CLOT system mentioned earlier was a joint effort involving members of our group but with the idea and domain expertise coming from members of Don Lindberg's group at the University of Missouri. At the University of Illinois, students working under Donald Michie and Alan Levy have used EMYCIN in two ways: one group developed a new EMYCIN application in tax advising, and the other developed a PASCAL implementation of the ideas used in EMYCIN. The latter program is now being used experimentally in an application involving emergency responses on off-shore drilling rigs. Finally, David Stodolsky at the Systems Science Institute at the University of Louisville has begun to experiment with EMYCIN in an application involving the psychology of interactions in large group conferencing.

B. Sharing and Interaction with Other SUMEX-AIM Projects

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

We have also interacted with members of the SECS project on SUMEX who have considered developing a question answering system for SECS similar to the one in MYCIN.

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

C. Critique of Resource Management

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

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

Special mention should be made of the remarkable role played by Tom Rindfleisch and his staff in helping to organize remote demonstrations of MYCIN and INTERNIST. In March 1979, when the American College of Physicians met in San Francisco, they rented a truck and drove to the City

with terminals and monitors. The installation they arranged worked well and provided a superb demonstration environment for the physicians who attended. In New Orleans in 1980, the greater distance prevented us from installing the equipment ourselves. SUMEX kindly offered to help orchestrate the New Orleans arrangements, though, and literally hours were spent locating terminals, arranging for telephone hookups, and finding the right kind of slave monitors. We salute SUMEX for their uncomplaining assistance in this regard, but also would like to note the need for a mechanism that is somewhat less ad hoc for facilitating the demonstration of SUMEX systems from remote locations.

Finally, we continue to feel the need for more computing power. Most of our research and development takes place in the hours from 7 p.m. to 10 a.m., but it is unreasonable to expect all our collaborators to adjust their own schedules around a computer. The existence of the 20/20 has been helpful in permitting demonstrations with good response time, and it will also allow us to introduce ONCOCIN in a real clinical environment within the next several months, but ongoing R&D on the main machine remains difficult much of the time. Even the evening hours are now seeing higher load averages than was once the case.

III. Research Plans (8/80-7/86)

A. Project Goals and Plans

EMYCIN

Our current plans call for four principal efforts related to EMYCIN. First, the knowledge acquisition component of the program, derived from the TEIRESIAS work of Davis, is being modified and expanded. Our concerns relate to both the inefficiencies and limited power of the current capabilities. The meetings during which the CLOT knowledge base was developed were recorded on tape and are forming the basis of an analysis of the knowledge acquisition process. Some early work implementing the ideas derived from those tapes is already under way.

We are also planning to prepare EMYCIN for "export" during the coming year. This will involve tightening up the code, maximizing efficiencies in space and time use, and improving the system's documentation. We do not intend to recode EMYCIN in a language other than INTERLISP, but do want to make it a stand-alone system that can be used for system building in a number of LISP environments. A key element of the documentation will be to better define those environments in which EMYCIN can be most effectively applied.

Now that the design and capabilities of EMYCIN are essentially fixed, we are also planning to develop a new application. Other EMYCIN systems have been developed in parallel with EMYCIN itself, and have therefore affected the program's design, but it is now appropriate to see how effectively a new system can be built within the current system