I. THE PROGRAM PROPOSED

We begin this proposal with a description of the broad program contemplated, with rationale and justification of need, and a description of resources and facilities already available for the purpose.

Herein we propose a five-year program of research on knowledge representation, and the various problems associated with it in the design of knowledge-based computer programs. The Stanford University group will work collaboratively with a group from the University of Missouri's Health Care Technology Center, under the direction of Dr. Donald Lindberg. The program will be under the general direction of Professor Edward Feigenbaum of Stanford, who presently serves also as the Principal Investigator of SUMEX-AIM, the NIH-sponsored National Computer Resource for research on the application of Artificial Intelligence (AI) techniques to medicine and biology. This Resource will serve the computer needs of the proposed program.

The proposed program consists of four activities: three projects and a core research activity.

Projects One and Three address the problems of knowledge representation, acquisition, and utilization in specific medical/hospital settings.

In Project One, the clinical setting is the Oncology Day Care Clinic. The task that provides specificity and direction to the research is the construction of a consultation system regarding experimental protocols and selection of therapy for clinic outpatients. This project is led by Professor E.H. Shortliffe of the Stanford Medical School, the original developer of the MYCIN program for consultations regarding infectious disease diagnosis and therapy.

In Project Two, the transfer of such expertise to other places and to other medical applications can be viewed as the primary goal. One powerful way of cumulating the concepts and methods of an emerging branch of Computer Science is to cumulate them in working software packages that widely applicable and widely shared. This project aims at developing a number of such packages or "tools", constituting a computer-program "workbench" for further research on and application of knowledge-based systems. The packages emerge as generalizations of work done in the task-specific projects;constitute a very tangible type of result therefrom; and serve to amplify and accelerate future efforts. This project is under the direction of Professors Bruce Buchanan and Douglas Lenat of Stanford.

In Project Three, the setting is the Clinical Laboratory and the task is one of acquiring and representing the medical expertise that allows the laboratory expert (e.g. the Laboratory Director) to interpret test results and discuss these with the patient's clinical physician. This is the inter-university collaboration headed by Dr. Lindberg. An important subgoal of this project is the transfer of the Stanford expertise in knowledge based systems research to the Missouri Center.

The Core Research Activity will' investigate a variety of fundamental research questions whose answers will shape present and future developments in knowledge representation research. Such questions involve formalisms and data structures for representing various types of knowledge; various methods—some automatic, some interactive—for acquiring new knowledge in systems; new inferential methods for putting this knowledge to work; strategy-knowledge representations for reasoning about the domain specific knowledge; and so on. The Core Research Activity is under the direction of Professor Feigenbaum. Douglas Lenat of Stanford.

Lastly, it is an objective of the overall program to disseminate the findings of the research, and to provide training opportunities to others. This objective will be accomplished through publications, presentations of research results at scientific meetings, by making room in the operational sites and the core activity for visiting scientists and trainees, and by - participation in a special annual meeting. The meeting to discuss our research and similar projects in this field will either be a part of or be coordinated with the annual artificial intelligence in medicine meetings at Rutgers University. That is, in years when the Rutgers meeting agenda and housing facilities can accommodate this group and its audience, we will join with Rutgers. In years when this is not possible, we will sponsor a separate meeting addressed to the four principal objectives of this program.

The administrative arrangements for the Program will be these:

The Principal Investigators of the various program activities will collectively constitute an Executive Committee for the Program, under the chairmanship of the Program Director. The Executive Committee will meet routinely by telephone-conference and occasionally face-to-face.

An Advisory Group will be formed, consisting of colleagues at other institutions who share our motivations and scientific interests. This group will advise the Executive Committee on major decisions and will offer peer review as necessary. The kernel of the Advisory Committee will be drawn from the membership of the SUMEX-AIM Advisory Committee (for which Dr. Lindberg is currently chairman).

I.A. Rationale for the Program

I.A.1. What do we mean by knowledge?

Computer scientists have long recognized that a computer is a general symbol-manipulating device. Arithmetic constitutes a special case of this capability—the manipulation of those symbols that are numbers. In this proposal we will be discussing non-numeric symbol manipulation by computers. In thinking about non-numeric computation, it is useful to think about:

a. inference methods (as opposed to calculation and algorithms)

b. qualitative "lines of reasoning" (as opposed to quantitative formulations)

c. symbolic facts (not merely numeric parameters and formulas)

d. decision rules of expertise and judgment (as opposed to mathematical decision rules)

The use of the term "knowledge" in this proposal is intended to cover both (c) and (d) above. In common usage, the term "knowledge" does not usually include (d), because such judgmental and experiential knowledge is largely tacit knowledge and therefore not recognized (i.e. the knowledge is "private" and the expert is not aware of what he/she knows and is using in problem-solving). The knowledge is private not because the expert is unwilling to share it, but because he/she is unable to discover and verbalize it.

It is central to our view that such knowledge—the knowledge of "expertise"— is critical for competent practice in medicine and science, in fact constituting the bulk of the knowledge employed in such practice. We view as a matter of great importance that such knowledge be codified and given a concrete (and at least semi-formal) representation, so that it can be used, stored, transmitted to others, analyzed, discussed, and taught.

Every activity of this proposed program is aimed at developing the scientific concepts and methods by which this can be most expeditiously, carefully, and usefully done.

Symbolic computation, though general and powerful, has hardly begun to be exploited in real applications. The specialty within Computer Science that has studied complex methods of symbolic computation is "Artificial Intelligence Research."

I.A.2. Some Relevant Global and Local History

Early work in artificial intelligence aimed toward the creation of generalized problem solvers. Work on programs like GPS [by Newell and Simon] and theorem proving , for instance, was inspired by the apparent generality of human intelligence and motivated by the belief that it might prove possible to develop a single program applicable to all (or most) problems. While this early work demonstrated that there was a large body of useful general purpose techniques (such as problem decomposition into subgoals, and heuristic search in its many forms), these techniques did not by themselves offer sufficient power for expert levels of performance. Recent work has instead focused on the incorporation of large amounts of task specific knowledge is what have been called "knowledge-based" systems. Rather than non-specific problem solving power, knowledge based systems have emphasized high performance based on the accumulation of large amounts of knowledge about a single domain. A second successful focus in work on intelligent systems has been the emphasis on the utility of solving "real world" problems, rather than artificial problems fabricated in simplified domains. This is motivated by the belief that artificial problems may prove in the long run to be more a diversion than a foundation for further work, and by the belief that the field has developed sufficiently to provide techniques that can aid working scientists. While artificial problems may serve to isolate and illustrate selected aspects of a task, solutions developed for those selected aspects often do not generalize well to the complete problem.

There are numerous current examples of successful systems embodying both of these trends, systems which apply task-specific knowledge to real world problems.

The following are synopses of a variety of knowledge-based systems developed by the Stanford participants in this program over the past thirteen years:

DENDRAL: An intelligent assistant to an analytic and structural chemist. It infers the structures of complex organic molecules from structural constraints. These constraints are either supplied interactively by the user from his "private" knowledge and intuition, or are inferred automatically from instrument data, such as mass spectral data, nuclear magnetic resonance data, etc. For those families of molecules for which the knowledge base has been carefully elaborated, the DENDRAL program performs at levels equalling or exceeding the best human experts. The DENDRAL program now has a significant user community in university laboratories and in industry, and is being used to solve difficult real problems.

Meta-DENDRAL: This program is focused on the problem of elaborating DENDRAL's knowledge base for specific families of compounds. It infers an empirical theory (a body of fragmentation rules) of the mass spectrometry of specific families from recorded mass spectral data. It has not only "rediscovered" rules previously acquired from chemists, but has discovered novel rules for certain families—rules that have recently warranted publication in the chemical literature.

MYCIN: This program is an intelligent assistant to a physician diagnosing infectious diseases. In conjunction with its diagnoses, it recommends therapeutic action. It is capable of explaining its line-of-reasoning in any (and varying) level of detail to the user in English. It can accept new decision rules from the user in English. It keeps an updated model of its own knowledge base, which it uses to critique the introduction of new rules into the system. It is capable of acquiring and using measures of the uncertainty of the knowledge, and produces a "believability" index with each inference, i.e., it is capable of approximate implication. A version called EMYCIN, sans infectious disease knowledge, has been developed to extend the use of the system to other domains.

HASP: Project scientists working in a classified environment led the development of a signal-understanding program for continuous surveillance of certain objects of military interest. The program ran successfully in a number of highly varied test situations, and is being further developed in a currently-funded ARPA program. The program used a design for incremental hypothesis formation that was a modification of the HEARSAY design for the CMU speech-understanding system. Symbolic knowledge from a number of sources was used to aid the interpretations of the primary signal data. Time-dependent analysis was novel in this system and played an important role.

AM: This remarkable program conjectures "interesting" mathematical concepts. Its knowledge base encompasses the (usually private) knowledge of a mathematician as to what constitutes an "interesting" construct in mathematics. Starting with the simplest set-theory concepts, and hundreds of rules defining "interestingness" of mathematical concepts, it has conjectured such concepts as addition, multiplication, factorization, primes, unique factorization into primes (the fundamental theorem of arithmetic), and an almost unstudied concept in number theory called "maximally divisible numbers."

MOLGEN: (under development) This program is being designed to be an intelligent assistant to an experimental molecular geneticist in formulating plans for laboratory experiments involving the manipulation of short DNA strands with restriction enzymes. The program is concerned with representing knowledge about planning and with the automatic formulation of plans to the level of detail demanded by the user. The program's knowledge must be represented at various levels—biological, genetic, topological, and chemical—and these levels must be incorporated into the reasoning.

CRYSALIS: Crystallographic Image Interpretation: (under development) This program is being designed to interpret ambiguous, incomplete three-dimensional image data obtained in x-ray crystallography of protein structures. The image input data is the so-called electron density map and the answer desired is an approximately correct protein molecule (or portion thereof). As with HASP, many sources of symbolic data support the interpretation of the primary signal data. The HASP program organization has been imported as a test of its generality. The interpretation problem is difficult because the best wavelength available (x-rays) is too long to resolve atoms and interatomic separations; hence the need for additional sources of symbolic knowledge, e.g., the amino acid sequence of the protein.

PUFF: This program interprets data from the pulmonary function testing laboratory and provides for the Lab Director an interpretive summary of findings regarding airways obstruction, lung restriction, and the degree of severity; subtype, such as bronchitis; the corroborating evidence and its weight; treatment recommendations;etc. This knowledge-based system was built in collaboration with a pulmonary physiologist at Pacific Medical Center, and is in routine daily use.

VM: A program that offers the attending physician or nurse interpretations of streams of data monitored from a patient in Intensive Care; signals alarm conditions due to unexpected patient condition or possible instrument malfunction; and offers advice regarding the management of the patient's ventilator machine assistance. This is another collaboration with Pacific Medical Center.

SACON: A MYCIN-like consultation system that advises a structural engineer on the analysis plan necessary to compute the multitude of structural engineering design parameters needed for building a complex structure (such as an airplane wing or an off-shore oil drilling platform or a building). Interactively, in consultation, the user supplies the design specifications. The system was built in collaboration with structural engineers at the MARC Analysis Corporation. It was built rapidly using the EMYCIN package discussed later.

In short, as the capsule sketches above indicate, the main themes of our work involve: the acquisition and maintenance of knowledge bases; the utilization of this knowledge in a variety of ways for data interpretation, problem solving, and planning; and the representation of this knowledge for computer inference.

I.A.3. Knowledge Representation Issues and Designs--the MYCIN Experience

In lieu of further general discussion of knowledge representation, we have chosen to explicate in some depth our viewpoint and methodology by drawing upon the experience in design and development of just one of our programs, the well-known consultation system MYCIN. For us, this work has been seminal; hence the discussion of it that follows generalizes to most of the other Stanford-based efforts mentioned above.

I.A.3.a. Background

Several computer programs have been written that attempt to model a physician's decision making processes. Some of these have stressed the diagnostic process itself [27],[17]; others have been designed principally for use as educational tools [31],[36],[56]; while still others have emphasized the program's role in providing medical consultations [4],[29],[51],[57]. Actually, these applications are inherently interrelated since any program that is aimed at diagnosing disease has potential use for educating and counselling those who lack the expertise or statistical data that have been incorporated into the program. Consultation programs often include diagnosis as a major component, although their principal focus involves interactive use by the physician and/or the determination of appropriate advice regarding therapy selection.

In general, the educational programs designed for instruction of medical students and other professionals have met with more long-term success [60] than has been the case for the diagnostic and consultation programs. The relative success in implementing instructional programs may result because they deal only with hypothetical patients as part of an effort to teach diagnostic and therapeutic concepts, whereas the consultation programs attempt to assist the physician in the management of real patients in the clinical setting. A program making decisions that can directly affect patient well-being must fulfill certain responsibilities to the physician if he is to accept the computer and make use of its knowledge.

Physicians will, in general, reject a computer program designed for their use in decision making unless it is accessible, easy to use, forgiving of noncrucial errors from nonexpert typists, reliable, and fast enough to facilitate the physician's task without significantly prolonging the time required to accomplish it. They also require that the program function as a <u>tool</u> to the physician, not as an all-knowing machine that analyzes data and then states its inferences as dogma without justifying them.

Those who design computer programs to give advice to physicians must devise solutions to these requirements in an effort to combat the current lack of acceptance of computer-aided diagnosis by the medical profession [14],[24]. The physician is most apt to need advice from such a program when an unusual diagnostic or therapeutic problem has arisen. However, he may be unwilling to experiment with a program that does not meet the general requirements outlined above.

Considerations such as those mentioned here have in large part motivated the research of our group over the last half-decade. We felt it was important to devise a consultation program that was (1) useful, (2) educational when appropriate, (3) able to explain its advice, (4) able to understand and respond to simple questions stated in natural language, (5) able to acquire new knowledge interactively, and (6) able to be modified easily. Although we recognized that this list of design considerations was somewhat idealistic in light of the state of the art in computer science, we did feel that it provided a useful set of long-range goals. The program we developed, known as MYCIN, has had considerable success in achieving many of the goals stated.. The current research proposes to build on the MYCIN experience, both by expanding the basic computer science methodology to deal with recognized problems as yet unsolved, and by implementing a consultation system in a clinical setting where its usefulness and acceptability to physicians can be assessed.

I.A.3.b. The MYCIN Program

As medical knowledge has expanded in recent decades, it has become evident that the individual practitioner can no longer hope to acquire enough expertise to manage adequately the full range of clinical problems that will be encountered in his practice. Thus when a patient's problem clearly falls outside the area of the attending physician's expertise, consultations from experts in other subspecialties have become a well accepted part of medical practice. Such consultations are acceptable to doctors in part because they maintain the primary physician's role as ultimate decision maker. The consultation generally involves a dialog between the two physicians, with the expert explaining the basis for his advice and the nonexpert seeking justification of points he finds puzzling or questionable. A consultant who offered dogmatic advice he was unwilling to discuss or defend would find his opinions were seldom sought.

Fig. 1 shows a schematic view of the consultation process. Appendix A shows a detailed typescript of a sample consultation. The physician nonexpert gives information about his patient to the expert in response to questions and, in return, receives advice and explanations. Thus there are actually three kinds of information flow between the physician and his consultant. The MYCIN program models the consultative process by attending to all three kinds of information. It is our conviction that programs which ignore the explanation pathway will fail to be accepted by physicians because they will see in such systems too severe a departure from the human consultation process (in which the primary physician is provided with sufficient information to allow him to decide whether to follow the offered advice).

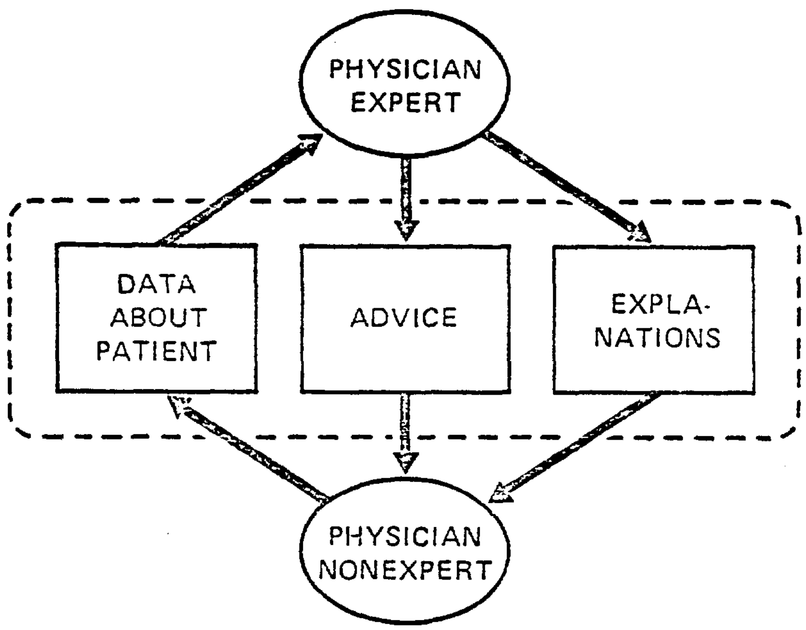

Figure 1 - Information Flow Between Physician And Consultant

MYCIN is a LISP program designed to serve as a clinical consultant on the subject of therapy selection for patients with serious infections. The program may be envisioned as interposed between the expert and nonexpert in much the way that the large box is positioned in Fig. 1. The difference is that the human expert can offer only general knowledge to the program, not patient-specific decisions. The program thus becomes the decision maker, using general medical knowledge from experts to assess a specific patient and to give advice plus explanations for its judgments.

Fig. 2 details the organization of MYCIN relative to the human consultation process depicted in Fig. 1. As before, the nonexpert offers data about his patient and in return receives both advice and, when desired, information via one of two internal explanation mechanisms (the general question-answerer or the reasoning-status checker). The basis for all decisions is domain-specific knowledge acquired from experts (static knowledge). A group of computer programs (the rule interpreter) uses this knowledge, and data about the specific patient, to generate conclusions and, in turn, therapeutic advice. It simultaneously keeps a record of what has happened, and this record is available to the explanation routines if the physician asks for justification or clarification of some conclusion that the program has reached. Although Fig. 2 is somewhat complicated, the following discussion should clarify the interrelationships among the various system components depicted in the diagram. Furthermore, Appendix A gives detailed examples of all the features described below.

Knowledge Representation

Static Knowledge

Static knowledge refers to all data that are constant in the program and unchanging from one consultation to the next.

Facts About The Domain. Much of the knowledge MYCIN requires is simple statements of fact about the domain. These can generally be represented as attribute-object-value triples

Overview Sec I.A.

uses this knowledge, and data about the specific patient, to generate conclusions and, in turn, therapeutic advice. It simultaneously keeps a record of what has happened, and this record is available to the explanation routines if the physician asks for justification or clarification of some conclusion that the program has reached. Although Fig. 2 is somewhat complicated, the following discussion should clarify the interrelationships among the various system components depicted in the diagram. Furthermore, Appendix A gives detailed examples of all the features described below.

Knowledge Representation

Static Knowledge

Static knowledge refers to all data that are constant in the program and unchanging from one consultation to the next.

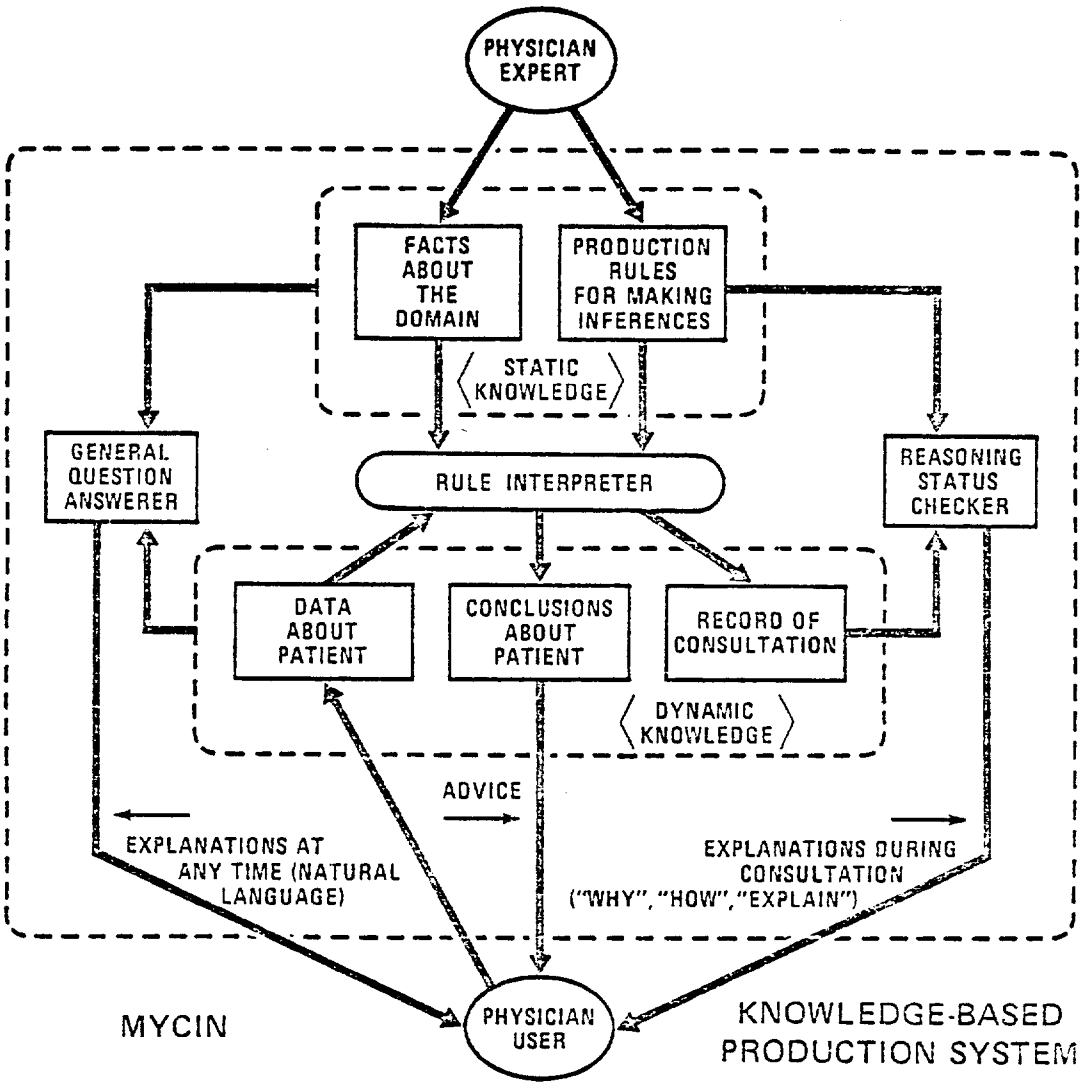

Figure 2 - Schematic Description Of MYCIN Related To Fig. 1

Production Rules. (Appendix A - Section I) In addition to simple facts, MYCIN requires judgmental knowledge acquired from experts and available for use in analyzing a new patient. Judgmental knowledge in MYCIN is expressed as production rules [16] which define certain preconditions (the PREMISE) that allow a conclusion to be reached (the ACTION) with a specified degree of confidence (the "certainty factor" [49]). Although such rules are stored as LISP list structures, a series of routines is available for translating them into English. For example:

```
PREMISE: If the stain of the organism is gramneg, and
         the morphology of the organism is rod, and
         the organism is anaerobic,
ACTION:  Then there is suggestive evidence (.7) that
         the identity of the organism is bacteroides.
```

Note that the _purpose_ of this rule is the determination of organism identity. Rules are classified and accessed in accordance with their purpose as described below.

Dynamic Knowledge

Dynamic knowledge refers to all data that are variable and change from one run of the program to the next.

Data About The Patient - Acquired From The User. MYCIN asks questions of the user, driven by a reasoning algorithm described below. These questions generally ask the user to fill in the "value" in an attribute-object-value triple (eg., "What is the patient's name?"), or to give the truth value of a predicate (eg., "Is the patient a compromised host?"). Thus these data may be represented, once acquired, in precisely the way that facts about the domain are represented in the static knowledge base (see above).

Data About The Patient - Generated By The Program. When the preconditions in the PREMISE of a rule are found to hold, MYCIN executes the ACTION portion of the rule and generates a new "fact" which can, once again, be represented as an attribute-object-value triple. As mentioned above, conclusions may also have a confidence value associated with them, thereby requiring that the triple be expanded to a quadruple:

```
the identity of ORGANISM-1 is
        bacteroides, with     -
        certainty factor of 0.7
                       (IDENTITY ORGANISM-1 BACTEROIDES .7)
```

Predicates may be similarly expanded. Furthermore, by generalizing this scheme to include representation of data acquired from the user, the physician may be asked to express his confidence in the answer he gives when MYCIN asks a question.

Maintenance Of A Record Of The Consultation. A history of the consultation is the third variety of dynamic knowledge. The details of representation need not be described here, but these data include records of which rules succeeded, which rules were tried but failed, how specific decisions were made, how information was used, and why questions were asked.

The Production System

The Rule Interpreter

This series of routines analyzes rules in the static knowledge base, determines whether they apply to the patient under consideration, and if so draws the conclusions delineated in the ACTION portions of the rules. This process would quickly become unmanageable as system knowledge grew if there were not a mechanism for selecting only the most relevant rules for a given patient. This is accomplished by a goal-oriented approach that we have described in detail [50],[51]. Briefly, as the rule interpreter examines the PREMISE of a rule, it notes whether the relevant data needed to determine the truth of each precondition are already known. If not, it digresses to examine those rules which make conclusions about the data that are needed by the first rule. The PREMISE conditions of those rules may, in turn, invoke additional rules, and in this way a reasoning network relevant to the first rule is formed. Since rules are classified according to their purpose, as previously described, it is easy to identify all rules that may aid in determining the truth of a specific precondition. The entire process is initiated by invoking a specific "Goal Rule" which defines MYCIN's task and is the only rule necessarily invoked for every consultation. When MYCIN can find no rules for determining the truth of a precondition, it asks the user for the relevant data. If the physician does not know the information either, the invoking rule is simply ignored.

Maintenance Of Initiative In The Hands Of The Physician

As was discussed above, a physician is not likely to accept a system such as MYCIN if the program simply asks a series of questions and then presents a piece of dogmatic advice as it terminates execution. The production system has therefore been provided with a series of "interrupts" that allow the physician to digress with questions of his own or to demand justification for the line of questioning on which MYCIN has embarked during the consultation. Whenever the program asks a question, the user can temporarily refuse to answer and instead call on the explanation capabilities described in the next section.

Explanations

The Reasoning-Status Checker (RSC) (Appendix A - Section IV)

This component of the explanation system deals with most questions that arise during the consultation session itself. Because the context of current reasoning about the patient is well-defined, the physician can be given a great deal of information on the basis of a few simple commands that do not require natural language processing. These commands are briefly described below; the details of their implementation have also been documented [48]. As shown in Fig. 2, the reasoning status checker (RSC) uses only the knowledge base of rules and the current record of the consultation; the general question-answerer (GQA) described below, on the other hand, has access to all static and dynamic knowledge.

The WHY Command. Whenever MYCIN asks a question, the physician may prefer not to answer initially and instead to inquire about the reasoning underlying the questioning. Thus he may simply respond with the command WHY (i.e., "Why do you think that the information you are requesting may be useful?"). Since all questions MYCIN asks are generated by rules, and since the rules are selected according to their purpose as previously mentioned, an English language translation of the rule under consideration generally serves as an adequate response to the WHY query. The RSC therefore responds by displaying the current rule. In addition, it places an identifying number before each of the preconditions in the PREMISE and indicates whether the condition is (a) already known to be true, or (b) still under investigation (note that one of the latter group of preconditions will have generated MYCIN's current question to the user). The physician can in turn inquire why the displayed rule was selected by asking WHY a second time, and the RSC will accordingly display the next rule in the reasoning network.

The HOW Command. As mentioned above, when MYCIN displays a rule in response to the WHY command, it labels each precondition in the PREMISE with a unique number. The physician may then respond to the displayed explanation by entering HOW followed by one of the identifying labels. If the reference condition is one that MYCIN has already concluded to be true, the RSC assumes that the physician is asking "HOW did you decide that the specified precondition is true?" and answers by citing the relevant rules used to make the decision. If, on the other hand, the cited condition has not yet been fully investigated, MYCIN assumes the physician is asking "HOW will you decide if the specified precondition is true?" and responds by citing the rules it intends to try, only some of which may actually succeed.

<u>The General Question-Answerer (GQA)</u> (Appendix A - Section V)

The general question-answerer (GQA) is a more comprehensive explanation system which, at any time during or after the consultation session, has full access to all static and dynamic knowledge in MYCIN (Fig. 2). Since it cannot make simple assumptions based on context, as the RSC can do, the GQA must accept and answer questions expressed in natural language. MYCIN's rule-based knowledge representation scheme, and some techniques borrowed from early work in computational linguistics [13],[30],[47], permit a straightforward but powerful approach to interpreting simple English questions without contending with several of the complex problems of natural language understanding. The details of this approach have been documented [76].

<u>Questions About Static Knowledge</u>. The ability to retrieve information from the static knowledge base gives the GQA a tutorial capability. Since the static knowledge is acquired from experts, the GQA can essentially act as an intermediary between an expert and a physician seeking general information about the infectious disease field. The user might ask simple questions of fact (eg., "Which culture sites are normally sterile?") or questions regarding judgments stored in rules. Questions of the second variety are termed "rule-retrieval" questions because they may be answered simply by identifying and displaying English versions of relevant rules from the knowledge base. Retrieval may be keyed to the rule PREMISE (eg., "How do you use the gram stain of an organism?"), the ACTION (eg., "When do you decide an organism might be a streptococcus?"), or to both the PREMISE and ACTION (eg., "Do you ever use the morphology of an organism to determine its identity?"). Furthermore, a question may deal with a specific rule (eg., "What is rule037?"). Note that none of these rules refers to a specific patient or consultation and thus requires no access to the dynamic knowledge base (Fig. 2).

Questions About Dynamic Knowledge. Although the RSC permits inquiries regarding the dynamic knowledge base, its scope is limited by the context of the current question being asked by MYCIN. If the physician wishes to ask more general questions regarding the status of MYCIN's reasoning, or if he wishes to review the program's decisions after the consultation is complete and MYCIN is no longer questioning him, the GQA gives him free access to all information about the specific consultation. Once again, the user might ask simple questions of fact (eg., "From what site was culture-2 obtained?") or questions regarding the basis for MYCIN's judgments. The second variety is again a rule-retrieval question, but is keyed to the consultation record in dynamic data rather than to the knowledge base of rules in static data (see Fig. 2). Thus questions may again reference the PREMISE (eg. "How did you use the gram stain of organism-1?"), the ACTION (eg., "What makes you think that organism-2 might be a streptococcus?"), or both (eg., "Did you use the morphology of organism-1 to determine its identity?"). Note that these questions parallel the examples given in the previous section but that they are consultation-specific and thus request the retrieval not of all relevant rules, but only those that were actually used successfully in the specified context. Finally, one may again wish to ask about a specific rule (eg., "Did you use rule037 when considering organism-1?").

Knowledge Acquisition

The only component of Fig. 2 not yet discussed is the crucial step of acquiring domain-specific knowledge from experts and coding it for storage in the static knowledge base. When MYCIN was first being developed, such knowledge was acquired by extensive meetings during which infectious disease experts and computer scientists discussed specific patients and attempted to analyze and extract the individual facts and rules that they were using. Recently extensive work has been devoted to the problem of automating the knowledge acquisition process in sessions involving clinical experts interacting with MYCIN directly (Appendix A - Section IX). This problem has been the subject of a doctoral dissertation by one member of our group [15].

Certainty Factors

Efforts to develop techniques for modeling clinical decision making have had a dual motivation. Their potential clinical significance has of course been apparent. The design of such programs also has required an analytical approach to medical reasoning that has in turn led to a distillation of decision criteria that in some cases had never been explicitly stated before. It is a fascinating and educational process for experts to reflect on the reasoning steps that they have always used when providing clinical consultations.

Several programs have successfully modeled the diagnostic process [27],[28],[55]. Each of these examples has relied upon statistical decision theory as reflected in the use of Bayes' Theorem for manipulation of conditional probabilities. Use of the theorem, however, requires either large amounts of valid background data or numerous approximations and assumptions. The successful performance of Gorry and Barnett's early program [27], for example, and a similar study by Warner using the same data [55], depended to a large extent upon the availability of good data regarding several individuals with congenital heart disease. Gorry [28] has had similar access to data relating the symptoms and signs of acute renal failure to the various potential etiologies.

Although conditional probability provides useful results in areas of medical decision making such as those mentioned, vast portions of medical experience suffer from so little data and so much imperfect knowledge that a rigorous probabilistic analysis, the ideal standard by which to judge the rationality of a physician's decisions, is not possible. It is nevertheless instructive to examine models for the less formal aspects of decision making. Physicians seem to use an ill-defined mechanism for reaching decisions despite a lack of formal knowledge regarding the interrelationships of all the variables that they are considering. This mechanism is often adequate, in well-trained or experienced individuals, to lead to sound conclusions on the basis of a limited set of observations.

We have examined the nature of such nonprobabilistic and unformalized reasoning processes, have considered its relationship to formal probability theory, and have proposed a model whereby the incomplete "artistic" side of the practice of medicine might be quantified. We have had to develop this model of inexact reasoning in response to MYCIN's needs; i.e., the goal has been to permit the opinion of experts to become more generally available to nonexperts. The model is, in effect, an approximation to conditional probability. Although conceived with MYCIN's problem area in mind, it is potentially applicable to any domain in which real world knowledge must be combined with expertise before an informed opinion can be generated. The model has been described in detail [75] and is based upon a scheme of weighted numbers we call "certainty factors". Although the model has been implemented in the MYCIN system, and in EMYCIN (see below), and although it has allowed the program to demonstrate impressive decision making performance, we still recognize many problems with the formalism. The model has generated considerable attention in the literature [1] and many important suggestions for further research have been forthcoming.

Evaluations Of MYCIN's Performance

Work on MYCIN to date has concentrated on the infectious disease subfields of bacteremia and meningitis. Formal evaluations have been undertaken which show that MYCIN compares favorably with infectious disease experts in selecting therapy for patients with bacteremia [62] or meningitis [63]. However, we have not undertaken a clinical implementation of MYCIN yet, and do not intend to do so in the near future. The reasons for this decision are important in that they explain part of the reason that we have turned from infectious diseases to oncology at this time.

First, we have felt it is crucial that MYCIN not be placed on the wards for clinical use if it does not already compare favorably with other forms of consultative advice available to primary care physicians. We have learned that this requires that MYCIN know about essentially <u>all</u> major infectious disease subfields since the various disease syndromes interrelate clinically in such important ways. In our evaluations of the program, it has tended to be in those cases in which a concomitant infection existed at some other site that MYCIN has failed to perform adequately. Yet the time required for us to develop the required knowledge bases for genitourinary infections, endocarditis, pneumonia, and pelvic infections would necessarily be at least as long as the period it has required to acquire and test the system's knowledge of bacteremia and meningitis. We therefore anticipate a considerable period of time before the program will be able to provide consistently reliable infectious disease consultations and hence be ready for ward implementation.

There are other problems as well that have been brought out by the complex decisions involved in infectious disease therapy selection. First, the truth model we have devised (see discussion of certainty factors above) has several recognized inadequacies that will require further research and testing. Secondly, no computer-based decision making program with which we are familiar has adequately managed time relationships amongst variables, and MYCIN is no exception. We see the need for continued research into the ways in which the production rule formalism can be suitably adapted to accommodate the need to represent time dependencies in clinical reasoning and to use such dependencies to make appropriate decisions. For example, trends in a fever or white count over time may be much more important in assessing an infected patient's illness than the actual values of these parameters at the precise time when the consultation is being requested.

Finally, in order to expand MYCIN's infectious disease knowledge into new problem areas, improved capabilities for knowledge acquisition would be extremely useful. Although we have made important initial steps in the development of this kind of complex capability [15], there is clearly much more to be done before an infectious disease expert who is a computer novice will be able to comfortably interact at a computer terminal in order to "teach" MYCIN the infectious disease judgmental knowledge that it needs to know.

I.B. Resources that exist to aid this project

The research work proposed herein will not stand alone or apart from other research already under way in the two sites. The personnel and facilities in place at the University of Missouri's Health Care Technology Center are described later in the appropriate Project section. At Stanford there is an interlocking set of existing grants and contracts supporting the work of a large group of scientists and students, the Heuristic Programming Project of the Stanford Computer Science Department. This group has, over the years, produced the various systems summarized earlier.

Historically the most significant sources of funding have been:

1. contracts from the Defense Advanced Research Projects Agency, the leading government agency for funding artificial intelligence research.

2. grants from the Biotechnology Resources Program of NIH for the SUMEX-AIM computer facility, without which it would have been very difficult to accomplish what was accomplished.

The other grants have had a short-term character. Some have been renewed, others not.

The proposed NLM grant is important to this complex of funding not only because it represents a significant amount of funding but most importantly because it represents stable funding over a five year period. It, therefore, like the ARPA funding, will constitute the stable base of support that will allow the work to advance steadily without personnel and funding fluctuations. The NLM-sponsored work will, in turn, benefit from the other supported work in the usual coordinated and synergistic way that significantly amplifies the effect of the NLM support.

The grant for the SUMEX-AIM computer resource ends in mid-1981. There is no reason now to believe that at renewal time the grant will face trouble. However such large facilities grants are always subject to a great deal of pressure, not always from peer-review. The need to service the research activities of an ongoing five year NLM research project will definitely add strength to the renewal application.

Finally, a resource of the greatest significance for the success of this work are the collaborative links that we have built over the years with medical scientists and clinicians at the Stanford Medical Center, the Pacific Medical Center, and the University of Missouri. It takes years to make such links work smoothly, but the resource is indispensable to a project on biomedical knowledge representation.

I.C. Significance

Collectively, we stand on the threshold of a new era in our understanding of the nature of medical and scientific knowledge, its distribution, and its effective use. Superficially, the cause of this has been the emergence of electronic symbol-processing and digital communication. More substantially, the reason for optimism is the emergence of knowledge-based computer systems research and application as a viable scientific and technical discipline.

We are now beginning to understand in a scientific and technical way what practitioners have always understood about their fields of learning and practice: that the bulk of the knowledge they employ is not the knowledge of textbooks and journals, but the informal and judgmental knowledge gained from long experience and practice. This knowledge is almost never codified, but is passed from mentor to apprentice by long periods of training and interaction, such as the internship, residency, and the Ph.D. graduate program

In the last decade there have been significant demonstrations that such heuristic knowledge can be explicated, represented, and put to use. Needed is an interdisciplinary team consisting of computer scientists, domain specialists, and various computer programs and computer-oriented methodology.

Once explicated, this knowledge can participate in the ordinary processes of cumulation of understanding in a field. For example, it can be subject to further analysis and be the basis for empirical studies and experimental investigation. It can be criticized by peer review. And it can be taught, or disseminated by library methods (electronic or otherwise).

In addition, the formal knowledge of a field can be coupled to the informal knowledge to produce computer programs that act as "intelligent agents" to assist practitioners in solving large numbers of routine problems, and even some of the more difficult problems, with which they are faced. Some methods of computer-based inference are available today to do this, and more are coming as research in this area matures. The concept is one of "active knowledge" available to work for users, in contrast to the passive knowledge of texts and articles (knowledge which is useless until "discovered" by the practitioner through library search and reading).

Such a prospect is not visionary. It demands our immediate attention. We have known for many decades that computers are general symbol processing devices, not merely calculators. We have known for two decades how to program them to infer lines-of-reasoning through complex problems of a symbolic nature. In the last decade we have learned how to make such reasoning powerful and useful—by supplying such programs with considerable bodies of knowledge about the problem domains. And we have had to learn how to represent the knowledge. Now microelectronics has brought the time of low-cost computing upon us. The electronic processing necessary to make the power of symbolic computing available to a wide community will be available. We should not allow ourselves to drop behind in the development of the concepts and methods necessary for the emergence of the applications.

There are also roles for knowledge-based symbolic computing that are visionary, but must be explored. The kind of "active" knowledge we have been discussing can be used to assist in the discovery of new knowledge. The very human process of discovery of new knowledge is a slow and halting process at best, done by very few and marked by very rare bursts of creative insight. It now seems possible (even plausible) that models of certain kinds of discovery can be formulated that will systematize for computer application the intertwined activities of inferential search and literature (i.e. knowledge) search. The Meta-DENDRAL program (that has formulated new rules of fragmentation in mass spectrometry) and the AM program (that conjectured some not-so-new objects and theorems in number theory) are demonstrable precursors of this type of knowledge-acquiring program.

We envision a National Library of Medicine that will be a living library of the knowledge of medicine and biology, not merely the repository of texts, journals, and articles and not merely the immense file of their electronic images available at terminals.

II. CODIFICATION AND USE OF MEDICAL KNOWLEDGE IN ONCOLOGY

II.A. Introduction

II.A.1. Objectives

The long term objective of our research effort is the development of tools for the representation and use of medical knowledge in computer-based clinical consultation systems. Such systems will provide useful assistance to primary care physicians while incorporating features that heighten the acceptability of the systems to their intended users. We also wish to increase our understanding of the logic of medical diagnosis and therapy planning through this work. To that end we propose a five year research effort with the following goals:

(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;

(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 basic research will build on our group's prior experience with a computer-based consultant, termed MYCIN, that uses production rule symbolic reasoning techniques to assist in therapy selection for patients with serious infections. The domain we have selected for the first clinical implementation of these techniques is the management of research therapy protocols for cancer outpatients at Stanford Medical Center's new oncology day-care center.

II.A.2. Background

This research builds on a long history of work on the MYCIN and EMYCIN projects directed principally by Shortliffe and Buchanan. Many of the persons developing those systems will be involved with the research proposed here. These two projects are described elsewhere and thus need not be described here as well.

II.A.2.a. Stanford Division Of Oncology

In the past decade chemotherapy has assumed a more important role in the treatment of patients with cancer. Some 2,000 patients are under the direct care of the five faculty physicians of Stanford's Division of Oncology in the Department of Medicine. Most patients are receiving care on an outpatient basis, either at the Debbie Probst Oncology Day Care Center in Stanford Hospital or at the Division's twice-weekly clinic at the Palo Alto Veterans Administration Hospital. Altogether, about 9,000 outpatient visits are made to the Division physicians each year.

Effective management of cancer often involves more than one therapeutic technique. Increasingly, the initial course of treatment utilizes a combined modality approach. Surgery and/or radiation may be followed by chemotherapy to control any remaining cancer. However, chemotherapy alone may be curative in some cases.

Refined programs (protocols) have been developed for the administration of radiation and chemotherapy for many forms of cancer. The Division has had particular success with those used against Hodgkin's disease (the sixth most common cancer) and other lymphomas. In designing and carrying out individual programs of treatment, the physicians of the Division of Oncology work closely with Stanford specialists in other areas, particularly radiotherapists, surgeons, pathologists, diagnostic radiologists, pharmacologists, and immunologists. Stanford's expertise in these many disciplines contributes to the high level of care received by patients in the Division of Oncology.

The Division is of course also involved in educating and training physicians on all levels, from medical students to practicing physicians. Among the trainees are nine clinical fellows in oncology who participate actively in both clinical research and patient care. Five physician specialists and private physicians are involved directly with patient care in the