PROBLEM SOLVING EXPERTISE

A Proposal Submitted to SUMEX-AIM by:

Paul E. Johnson, Ph.D.
Center for Research in Human Learning
University of Minnesota

William B. Thompson, Ph.D.
Department of Computer Science
University of Minnesota

January, 1981

A) MEDICAL AND COMPUTER SCIENCE GOALS

1) Describe the proposed research to be undertaken on the SUMEX-AIM resource.

Work done at Minnesota to date has focused upon the study of problem solving expertise in domains of science, law, and medicine. This work has included the development of a large LISP-based simulation program of expertise in the diagnosis of congenital heart disease (The program and related work were described at the Sixth Annual Work Shop on Artificial Intelligence in Medicine held at Stanford University this past summer). Currently, we envision using the SUMEX-AIM resource to create small-scale simulations of component reasoning processes in science and law, as well as medicine, and as a means of extending our methodology for the study of problem solving expertise in medicine into a new area that we call "knowledge capturing."

The work in science and law will be based upon research currently underway or planned and described in the attached proposal (The work in science was funded by NSF for a two-year period beginning in the Fall of 1979. The work in law, which involves collaboration with Professor Michael Johnson at the University of Tennessee, is currently under review by the NSF Law and Social Sciences program).

The proposed research in medicine is in two parts. One part, which has been funded for a number of years by the Minnesota Center for Research in Human Learning through grants from NICHD and NSF, extends our research on expertise in the diagnosis of congenital heart disease using knowledge engineering tools available through SUMEX-AIM. This work involves collaboration with Dr. James Moller, a pediatric cardiologist in the University of Minnesota Medical School and Dr. David Swanson at the American Board of Internal Medicine in Philadelphia. The second part of the work in medicine is new: a proposal is in preparation for submission to the National Library of Medicine (and possibly the

Sloan Foundation) in collaboration with Dr. James Jenkins, a psychologist at the University of Minnesota, and Dr. Donald Connelly, a pathologist, from the University of Minnesota Medical School, as well as colleagues at the New England Medical Center (Dr. Gerald Kassirer, a nephrologist, and Dr. Benjamin Kuipers, a computer scientist). The new work involves basic research on knowledge capturing - the development of a methodology for systematically exploring and obtaining the knowledge which experts in a domain use in solving problems. Both parts of the proposed research are described further below.

Expertise in Diagnosis of Congenital Heart Disease

Our prior research on expertise in diagnosis of congenital heart disease has resulted in a theory of diagnosis and an embodiment of that theory in the form of a computer simulation which diagnoses cases of congenital heart disease. At a macroscopic level, the theory identifies four categories of knowledge used by the expert physician (pediatric cardiologist) in diagnosis. First, the physician has clinical knowledge of disease. This knowledge is hierarchically structured, including categories of disease, specific diseases, and variants of the same disease that differ in presentation. For each element in the hierarchy, there is knowledge of the associated anatomy and physiology and the expected clinical manifestations. Second, the physician has deductive knowledge of disease - knowledge of principles of cardiovascular pathophysiology and the clinical manifestations useful in detecting underlying pathology. This causal knowledge deductively relates cardiovascular defects to hemodynamics and to expected patient data. This category of knowledge is not typically used in diagnosis, since the expert physician can simply recall expected clinical manifestations through clinical knowledge of disease, rather than deduce them. Third, the physician has heuristic knowledge of disease and of clinical findings useful in its diagnosis. One aspect of this knowledge provides indices from clinical manifestations to diseases and from disease to disease, useful in choosing diagnostic hypotheses to consider. Another aspect of heuristic knowledge is related to evaluation of diagnostic alternatives -- rules of thumb for ruling in and ruling out alternatives as correct or incorrect. Another form of heuristic knowledge screens abnormal from normal clinical findings and identifies subsets of findings which are likely to have the same underlying cause. All heuristics are aimed at basically the same end -- reducing the cognitive demands of diagnosis without loss of diagnostic accuracy. The fourth and last category of knowledge is <u>knowledge of data acquisition techniques</u>: interviewing methods, physical exam maneuvers, special procedures, and laboratory utilization.

Diagnosis is characterized as a heuristic search process, with four sources of knowledge involved. Clinical knowledge of disease is the hierarchical structure to be searched. Deductive knowledge of disease is useful in construction of missing pieces of that hierarchy, and in justifying the clinical information it contains. Heuristic knowledge aids in limiting the section of clinical knowledge to be searched and in providing simple evaluation functions for use in search. Data acquisition for knowledge is essential in obtaining patient information to be used in the search process.

The research for which we wish to use the SUMEX-AIM resource has two subcomponents, one directed at selective reimplementation of the diagnosis program and the other directed at extensions of the research more broadly. Reimplementation of the diagnosis program would be carried out first, probably using the AGE and UNITS knowledge engineering tools, and INTERLISP. The present program is a blend of a semantic network representation (for clinical knowledge of disease) in a production system/blackboard control structure (heuristic activation of portions of clinical knowledge). This structure is nicely congruent with the facilities provided through UNITS and AGE linked together. This relatively well specified reimplementation task will provide an appropriate environment for learning about SUMEX-AIM and the resources it provides. It will also provide an interesting comparison of performance of similar diagnostic programs implemented in different ways.

The second subcomponent of the proposed research will focus upon data collection and patient management skills of the expert physician. Simulation efforts thus far have focused solely on diagnosis. Such limitations were necessary initially because of our lack of familiarity with the medical domain and because of the paramount importance of methodological development, as discussed in the next section on knowledge capturing. In pediatric cardiology, we now have the substantive knowledge and methodological tools required to approach medical problem solving more broadly, including knowledge and procedures related to both collection of patient data for diagnostic purposes and patient management.

Knowledge Capturing

In the last fifteen years, great progress has been made in synthesizing the expertise required for solving extremely complex problems. Computer programs exist with competence comparable to human experts in diverse areas ranging from the analysis of mass spectrograph and nuclear magnetic resonance [DENDRAL] to the diagnosis of certain infectious diseases [MYCIN].

Design of an expert system for a particular task domain usually involves the interaction of two distinct groups of individuals, "knowledge engineers", who are primarily concerned with the specification and implementation of formal problem solving techniques, and "experts" (in the relevant problem area), who provide factual and heuristic information of use for the problem solving task under consideration. Typically, the knowledge engineer, after consulting with one or more experts, decides on a particular knowledge representational structure and inference strategy. Next, "units" of factual information are specified. That is, properties of the problem domain are decomposed into a set of manageable elements suitable for processing by the inference operations. Once this organization has been established, major efforts are required to refine representations and acquire factual knowledge organized in an appropriate form. Major research problems exist in developing more effective representations, improving the inference process, and in finding better means of acquiring information from either experts or the problem area itself.

We propose to study one aspect of these problems which we feel will have significant impact on the future development of expert systems. Specifically, we hope to investigate the "knowledge capturing" process that occurs in the early stages of the development of expert systems when problem decomposition and solution strategies are being specified. Several related questions will be addressed: What are the performance consequences of different organization approaches, how can these consequences be evaluated, and what tool can assist in making the best choice? How can organizations be determined which not only perform well, but are structured so as to facilitate knowledge acquisition from human experts?

Programs currently exist for empirical investigation of some of these questions for a particular problem domain [AGE, UNITS, RLL]. These tools allow the investigation of alternate organizations, inference strategies, and rule bases in an efficient manner. What is still lacking, however, is a theoretical framework capable of reducing dependence on the expert's intuition or on near exhaustive testing of possible organizations. Despite their successes, there seems to be a consensus that expert systems could be better than they are. Most expert systems embody only the limited amount of expertise that individuals are able to report in a particular, constrained language (e.g., production rules). If current systems are approximately as good as human experts, given that they represent only a portion of what individual experts know, then improvement in the "knowledge capturing" process should lead to systems with considerably better performance.

The usual means employed to acquire the knowledge necessary for building an expert system is to have an expert individual construct rules that represent the knowledge required to perform given tasks. Unfortunately not all of the knowledge of the human expert may be in a form that can be assessed in this fashion. Much of what individuals know, experts included, is not available to conscious awareness. Thus, in talking with experts, one often encounters the statement "But I don't know what knowledge I have that enables me to read the x-ray" or "I can't tell you how I perform clinical diagnosis because that is an art and not a science." What experts are telling us, we believe, is that more subtle methods must be employed if we are to discover the basis for their knowledge and skill.

The substantive area chosen for initial attack is medical problem solving. This area is chosen because of its high intrinsic interest, because our past work suggests that further efforts should be fruitful, and because it is accessible to perceptual, memorial, and problem solving studies, which can provide rich evidence concerning the adequacy of conceptualizations of knowledge.

The specific methodology of our research derives from the discipline of cognitive science, and from our study of expert problem solvers in pediatric cardiology. This methodology consists of: (1) extensive use of verbal thinking aloud protocols as well as other experimental data as a source of information from which to make inferences about underlying cognitive structures and processes; (2) development of computer models as a means of testing the adequacy of inferences derived from the protocol studies; (3) testing and refinement of the cognitive models based upon the study of human and model performance in experimental settings.

2) How is this research presently supported? Please identify application and award statements in which the contingency of SUMEX-AIM availability is indicated. What is the current status of any application for grant support of related research by any federal agency? Please note if you have received notification of any disapproval or approval, pending funding, withing the past three years. Budgetry information should be furnished where it concerns operating costs and personnel for computing support. Please furnish any contextual information concerning previous evaluation of your research plans by other scientific groups.

Work being done in scientific reasoning is sponsored under a current NSF (SE079-13036) grant to Paul Johnson (see attached proposal). The work in law has been supported by the Minnesota Center for Research in Human Learning and is described in a proposal which is currently under review by the NSF Law and Social Science Program. The work in medicine is supported by NICHD (T36-HD-07151 and HD-01136) and NSF (NSF/BNS-77-22075) grants to the Minnesota Center for Research in Human Learning of which Paul Johnson is principal investigator. Additional support is being sought from the National Library of Medicine by the principal investigator.

3) What is the relevance of your research to the AI approach of SUMEX-AIM as opposed to other computing alternatives?

Our work is complimentary to many of the current projects supported on SUMEX-AIM. We will be investigating certain aspects of expert problem solving in order to develop better organizational and knowledge acquisition strategies. Such work requires that we be able to build upon the extensive experience in knowledge engineering within the SUMEX-AIM communities. Specifically, we first need to investigate a number of existing programs in order to determine the degree to which they satisfy the design goals which we will be establishing. We then hope to use the program construction tools that are available in order to build prototype systems to illustrate our idea.

B) COLLABORATIVE COMMUNITY BUILDING

1) Will the programs designed in your research efforts have some possible general application to problems analogous to that research?

The concepts we develop will clearly be relevant to related problem areas since we are studying strategies for the development of problem solving systems. It is too early to say whether specific software that is developed will prove useful to other researchers.

2) What application programs already publicly available can you use in your research? Are these available on SUMEX-AIM or elsewhere?

EMYCIN and possibly AGE, UNITS, and RLL, if and when they are available.

3) What opportunities or difficulties do you anticipate with regard to making available your programs to other collaborators within a reasonable interval of publication of your work?

Most of our software, at least initially, will be prototype systems on which we can run experiments. Should these programs prove of interest to others we will make every effort to provide them in a distributable form.

4) Are you interested in discussing with the SUMEX staff possible ways in which other artificial-intelligence research capabilities might interrelate with your work?

Definitely. The major work done on expert systems has been carried out by the SUMEX-AIM group. The point of our research is to investigate strategies for discovering and specifying expert problem solving strategies, not to develop another expert system.

5) If approved as a user, would you advise us regarding collaborative opportunities similar to yours with other investigators in your field?

Of course!

C) HARDWARE AND SOFTWARE REQUIREMENTS

1) What computer facilities are you now using in connection with your research or do you have available at your institution? In what respect do these not meet your research requirements?

We are currently using a CDC Cyber 74 and a Cyber 172 for most of our work in problem solving. The Cyber computers are not well suited for interactive computing and have serious limitations with respect to address space and available support software. We expect delivery on a VAX-11/780 in June, 1981 at the University of Minnesota. Hopefully, much of our work will eventually be implementable on the VAX. However, this depends both on our ability to acquire additional peripheral hardware and on increased availability of AI software for the VAX. Until that time access to a TENEX facility is extremely desirable.

> 2) What languages do you either use or wish to use? Will your research require the addition of major system programs or languages to the system? Will you maintain them? If you are committed to systems not now maintained at SUMEX, what effort would be required for conversion to and maintenance on the PDP-10-TENEX system? What are the merits of the alternative plan of converting your application programs to one of the already available standards? Would the latter facilitate the objectives of Part B, Collaborative Community Building?

Initial resource requirements will be limited primarily to experimentation with existing programming systems using restricted knowledge bases. Thus, our secondary storage requirements should be modest, but we will require CPU resources sufficient to run reasonably large and complex programs. In the later stages of the project, our program development efforts will require additional computing resources. At that point, we will need to decide whether to continue the entire effort on the SUMEX-AIM system or to off-load at least part of the work onto the VAX (assuming software is available at that time). The utility of the SUMEX knowledge engineering tools in our work will probably play a major role in that decision.

We understand that the Stanford TENEX system is at or near saturation during prime time (9-5 California time). We would have no problem avoiding that time period given the two-hour time differential. Response time on a loaded TENEX system would push us toward early morning and/or night usage in any event.