Reports of panel

on

APPLICATIONS OF ARTIFICIAL INTELLIGENCE

Saul Amarel (Session Chairman), Rutgers University
John Seely Brown, BBN
Bruce Buchanan, Stanford University
Peter Hart, SRI
Casimir Kulikowski, Rutgers University
William Martin, MIT
Harry Pople, University of Pittsburgh

INTRODUCTORY REMARKS. SUMMARY OF ISSUES AND OPEN PROBLEMS IN AI APPLICATIONS (S. Amarel)

In recent years there has been a growing amount of applications-oriented AI work. The distinction between an applied AI project and a 'basic' project is mainly determined by the principal objectives of the effort. If the objective is to develop a high performance, expert, system for solving certain problems in a given domain, in a manner which is acceptable (and hopefully useful) to users in the domain, then the project is in the applied category. What characterizes, in addition, work in several recent applied AI projects is the focus on natural domains, where the tasks involve reasoning with empirical data in the context of bodies of empirical knowledge. Problems in medicine, biochemistry and signal processing are of this type. It is certainly possible to apply AI methods and tools to the development of expert systems in the context of formal task environments such as mathematics. Indeed, one of the most successful early applications of AI has been in this area - specifically, the MACSYMA project at MIT (Moses, 1971). It is a fact however that most of the efforts in AI applications over the last 5-10 years have concentrated on systems whose knowledge bases contain relatively large amounts of empirical information about pieces of the real world.

An important reason for this is the high potential usefulness (social value) of AI systems that could provide help to scientists and professionals in problems that involve the understanding and control of natural phenomena - especially in situations where available knowledge is complex, ill structured, and changing. Another essential reason is the expectation that work in real world problems will bring new challenges to AI and it will contribute significantly to the scientific development of the field.

At present, there are several applied AI projects in the country that have produced systems at an impressive level of expertise in limited real world task environments.

In analytical chemistry, the heuristic DENDRAL project at Stanford has been the pioneer effort in the area of computer-based expert assistants for problems in a scientific domain. The DENDRAL program can interpret mass spectra of organic molecules at a level of performance which equals, and in some cases surpasses, experts in the field (Buchanan et al, 1969; Smith and Carhart, 1976).

In medicine, the CASNET/GLAUCOMA system at Rutgers, which provides consultation in diagnosis and therapy of glaucomas (Kulikowski and Weiss, 1971; Weisss 1974), and the MYCIN system at Stanford which assists in the treatment of infectious diseases (Shortliffe et al, 1975) have now reached expert status. The INTERNIST system at the University of Pittsburgh, which provides clinical consultation in problems of internal medicine, (Pople et al, 1975; Pople 1975, 1977) is nearing expert status for diagnostic tasks in parts of its domain. Several other expert* systems are now being developed in medicine, biochemistry, genetics, psychology, instruction, mineral exploration, business management, language and speech processing, and computer programming.

Much experience has been accumulating in the course of developing these expert systems. Workers in the field found that many of the methods and techniques that grew from previous basic work in AI could be successfully adapted to the building of knowledge-based expert systems. Also, in their attempts to raise the levels of expertise and performance of systems in specific task environments, they identified several important AI problems on which more basic work is needed.

In general, we are still at the stage of learning how to build high-performance knowledge-based systems. With the possible exception of DENDRAL, the expert systems that have been built to date for specific natural domains have not had as yet a significant impact on the users in these domains. The main value of these developments has been to demonstrate that expert systems of certain types are feasible within the present state of AI knowledge. Also, the exploratory work on realistic expert systems has forced us to look more closely into many aspects of knowledge handling that are of fundamental importance within AI. Other basic topics whose study has been stimulated by these system development efforts include methods of plausible reasoning, strategies of planning under uncertainty and approaches to problems of interpretation and theory formation.

From the viewpoint of AI as a scientific discipline, work on AI applications is providing a valuable testing ground for existing AI schemes and tools, it is helping to identify their scope and limitations, and it is stimulating the development of new methodologies for the design and construction of AI systems. Furthermore, it is forcing us to direct more attention to scientific communications - both within the field and with experts in other fields. I expect that this will bring about a better understanding of the present state of knowledge in AI, and a better perspective of how this knowledge is related to other parts of computer science and to knowledge in other disciplines.

From the viewpoint of a discipline such as medicine in whose domain AI systems are being developed, the effort to structure knowledge in the domain - which is an essential prerequisite for achieving an expert system - may be of a much more fundamental significance than was initially realized. The clarification of basic concepts in the discipline, the identification of relationships between various pieces of knowledge, the assessment of the nature and status of underlying empirical knowledge, of theories and of beliefs, the recognition of gaps, and - most important - the providing of a convenient computer-based environment which stimulates and supports these activities, are of great scientific value. In the long run, this may have a much stronger practical impact on the discipline than the impact produced by the addition of a resource or a technological tool to assist users in a specific task area within the discipline.

Emphasis of panel. Relationship to AIM Workshop

The main emphasis of our panel will be on applications of AI to the development of knowledge-based systems in medicine. In addition, applications in other domains (mineral exploration, education) will be considered, and certain methodological/technical problems that are shared by all these projects will be examined.

In particular, the panel is intended to provide a forum for an assessment of work being done in the AIM (Artificial Intelligence in Medicine) community. The AIM project is a national resource sharing activity which is supported by the Biotechnology Resources Program (BRP) of DRR, NIH, and whose principal objective is to promote applications of AI to biological and medical problems. The main focus of AIM activity is at the Stanford SUMEX-AIM project, which provides computer shared resources to the AIM community via national networks. The Rutgers Research Resource on Computers in Biomedicine (Amarel, 1975) is one of the (BRP-supported) projects in the AIM community; included among its functions is the organization of annual AIM Workshops. The objective of the Workshops is to strengthen scientific interactions within the national AIM community, and to disseminate AI-based methodologies, tools, and specific systems that are relevant to AIM. The 3rd AIM Workshop will be held at Rutgers on July 6 to 8, 1977. All the panel participants (except for Peter Hart who will be represented by Dick Duda) are expected to attend the Rutgers Workshop. Points of view and conclusions developed at the AIM Workshop will be presented at the panel as part of the discussion of issues that are outlined in the present report.

The panel will provide an opportunity to summarize for a wide AI audience the highlights of certain key issues in the development of applied AI systems that will be discussed at the AIM Workshop.

Comments based on recent experience with AI applications

Experience with work in AI applications to date shows the following:

(a) The problem of acquiring specific knowledge in a domain, managing it in a AI system, modifying it, and using it appropriately is fundamental. The approach to most designs is incremental and responsive to the fact that the knowledge base in a domain is not stationary. Initially, a relatively low-performance AI system is created to provide the basis for subsequent stages of knowledge acquisition and improvement, which eventually leads to a high-performance, expert, system.

(b) Work on applications requires very close collaboration between AI experts and experts in the problem domain. Furthermore, special technical support facilities (e.g., computer networks) can play a significant role in establishing these collaborations and in sustaining their effectiveness. This is a key point which has important implications on organizational and shared resource aspects of applied AI projects.

(c) The development of an expert system within a reasonable time span requires more powerful technologies than those in use today - expecially when the knowledge bases will grow from the present 102-103 'facts' to more realistic situations with 104 - 105 'facts'. So far,system development times (from conception to expert level in a research environment) have been 4-8 years. To reduce this time span, or to keep it from growing too much as knowledge bases grow, we need more effective methods of knowledge acquisition and organization and more powerful program design environments. Related to this, we need better techniques for interfacing AI programs with experts and users. At a more basic level, we need better schemes for coordinating multiple knowledge bases and for handling information which is inconsistent and/or uncertain.

Main issues and open problems in AI applications

There has been considerable progress in the development of schemes for representing knowledge in AI systems. Production systems, semantic nets and frame systems are among the major schemes used in the projects discussed in this panel. Experience is accumulating steadily with types of representations that are appropriate in different situations. Much more work is needed, however, on how to represent knowledge of different types (form, completeness, validity) for various problem solving tasks, and how to shift representations in a manner that facilitates problem solving. At present the choice of representational framework for a task (the set of basic concepts and their relationships, the grain of knowledge and the form of knowledge to be associated with specific types of processes) is the key decision made by the builder of a AI system. The success of a system depends critically on this decision. This fundamental problem of problem representation (or of conceptualization, as Buchanan puts it in his remarks later on in this report) has been with us for almost a decade (Amarel, 1970); it is now gaining added significance in the context of knowledge-based system design. Briefly, in a dynamically evolving system where the initial design is based on a given choice of problem representation (conceptualization), it is possible that no increase in performance - via accretion of the knowledge base - can be achieved beyond a certain point unless a change in representation (a re-conceptualization) can take place. This has strong implications on methods of system design, and it may determine the limits of performance of specific applied AI system. Experience with this problem already exists in the MYCIN and CASNET projects - where knowledge-based consultation systems ot high levels ot expertise have been achieved in limited domains of medicine. Until more is known about how to choose appropriate representations and how to handle shifts among representations, it would be reasonable to start new AI application projects only in domains with fairly well established conceptual frameworks.

Knowledge bases are built through processes of knowledge acquisition and assimilation. There has been some progress to date in the development of system aids for knowledge acquisition from experts. In his remarks in this report, Pople raises important conceptual issues in this area. He points out that the process of acquisition of an expert's knowledge in a given domain is a theory formation process. This is an intellectually demanding activity which involves the formation of hypotheses about the expert's knowledge structure, and the representation of these hypotheses in a computer with the help of available AI methods and tools. A key role of the computer scientist who is working with the domain expert in the development of an AI system is to perform this empirical theory formation process. The question arises whether AI methods can be used to assist the computer scientist in this task. Some work in this area is underway at Rutgers; it involves the development of the AIMDS representational framework which is intended to facilitate the formulation and testing of a psychologist's theories in the domain-of belief systems (Schmidt et al, 1976). Much more work is needed in this area - both in problems of theory formation and in related issues of representation. Also, methods of automatic knowledge acquisition and assimilation - from a system's operating experience or from direct observation of empirical data - must await further progress in theory formation strategies. Thus, more work in formation problems is needed to strengthen the processes of incremental construction and improvement of complex knowledge-based systems.

Many applied AI projects face the problems of how to represent and use multiple knowledge sources. Each source may represent knowledge of different grain (level of resolution) or of different form. For example one body of knowledge in a medical consultation system may represent a qualitative model of pathophysiological relationships in a disease process, another may represent a detailed anatomical model underlying (parts of) the qualitative model of disease, and a third may express normative rules of action for treatment decisions. Approaches to the use of multiple knowledge sources are discussed in this report by Kulikowski and by Brown in the context of medical and instructional tasks respectively.

One of the key processes in medical reasoning and in scientific inquiry is the interpretation of empirical data in the light of a given body of theoretical knowledge. Much of the work discussed in this panel is concerned with these processes. Typically, a problem of interpretation involves reasoning about an individual case. Given evidence (data) about the case and a theory within which the evidence is to be understood/explained, find a hypothesis which explains the case on the basis of the theory. There has been considerable progress in the development of strategies for solving interpretation problems. However, there remain several open problems: under what conditions the interpretation process should be controlled by the specific 'low level' features of the case under consideration, or by possible 'high level' hypotheses and by expectations derived from these hypotheses; how to best represent and keep track of information about a special case, of general knowledge in terms of which the case is interpreted, and of alternative interpretations of the case as the process evolves with time. These issues are discussed further in this report by Buchanan, Martin and Hart.

Hart is also stressing the importance of methods of plausible reasoning in diagnostic problem solving. In addition to work reported by Hart in the context of a system for mineral exploration, considerable work on plausible reasoning was done in recent years in the context of the medical AI projects. Also, work by LeFaivre at Rutgers (Le Faivre, 1976) has been focusing on systems for representing, and experimenting with approximate reasoning methods. More controlled experimentation, as well as theoretical consolidation of existing methods, are needed in this area.

An important feature of current AI applications is the emphasis on explanation facilities. Explanation of the reasoning done by a knowledge-based system in special cases is needed both for system development/debugging purposes and for interaction between system and user. This point is stressed below by Buchanan, as well as by other panelists. The methodologies and techniques of explanation developed for applied AI systems are expected to have a broader impact on the problem of scientific communication in AI. We are still faced with a major problem of how to communicate effectively what is known in the AI field. This involves the difficult issue of communicating complex programs, principles underlying their designs, accounts of their behavior, and their properties. The problem is general for computer science, but it can be seen most forcefully in the context of AI work. There is reason to hope that current work on explanation facilities in applied AI system will result in the development of general AI methods for effective scientific communication of complex programs.

Experience shows that we need powerful design methods and system tools for building knowledge-based systems and for evolving them from the concept stage to a user's environment. This is bringing most AI applications-oriented efforts in closer contact with many other areas of computer science and technology. There is renewed concern with programming languages and other support software, machine architectures, networking, communications and interfaces. This situation will probably result in an enrichment of system building technologies - both within AI and outside it. More fundamentally, it will clarify further the essential concerns and scientific content of AI and its many relations with other parts of computer science.

Before concluding, let me stress that any complete assessmentof the state of AI applications must include two parts. The first should be concerned with the performance achievements and the levels of expertise attained by specific programs in various application areas. The second, must focus on design issues that go across specific applications, and on the identification of new AI problems and approaches. The thrust of this panel is on the second part, because this part is more intrinsically relevant to the status of main technical issues within AI - and thus it is more appropriate for an AI Conference. However, this should not obscure the fact that in 1977 we are at a point where substantial achievements have been made in AI applications, and the prospects for the future of high performance AI expert systems are bright.

Let me now summarize the areas of AI that I believe require more work in order to strengthen the basis for the design of useful knowledge-based systems: guidelines for representing and using knowledge of different types; schemes for shifting representations; methods for solving formation problems; flexible strategies of interpretation; methods for acquiring and managing large bodies of knowledge (HP facts)from experts and from empirical data; and programming techniques for AI system development and for system transfer to various types of user environments. In addition, in order to increase the impact of AI on applications, we must find effective ways of organizing what is known in the field and of communicating this knowledge to others.

In the following sections, members of the panel - who represent major AI applications efforts - will comment on several of the issues that I discussed above. Their comments are presented in the context of specific application projects with which they are associated. A list of references concludes this panel report.

REPRESENTATION AND USE OF EXPERT KNOWLEDGE IN MYCIN (B.Buchanan)

Successful applications of AI to Science and medicine require large amounts of specific knowledge. Yet this presents problems for the representation, acquisition and use of the knowledge by an AI program, as evidenced in the MYCIN program developed at Stanford (Shortliffe, 1974; Shortliffe et al, 1975; Davis et al., 1977). This tension is not wholly resolved in MYCIN but we believe that extensions to the methodology will alleviate it.

A major goal of the MYCIN system was to provide a computer-based therapeutic tool designed to be clinically useful. This requires development of a system that has a medically sound knowledge base, and that displays a high level of clinical competence in its field.

Since many clinicians are not likely to accept the advice provided by a computer-based system unless they can understand why the recommended therapy has been selected, the system has to do more than just give advice dogmatically. It should have the ability to explain the reasoning behind its decisions, and should be able to do so in terms that suggest to a physician that the program approaches the problem in much the same way that he does. This permits the user to validate the program's reasoning, and modify (or reject) the advice if he believes that some step in the decision process is not justified.

We have also attempted to provide the program with capabilities for experts in infectious disease therapy to augment or modify the knowledge base, in order to improve the validity of future consultations. The system therefore, includes a knowledge acquisition system which enables experts to update MYCIN'S knowledge base, without requiring that they know how to program a computer. A principal feature of MYCIN, central to these goals, is the format in which its knowledge is encoded. Knowledge used by MYCIN is contained in diagnostic and therapeutic decision rules formulated during extensive discussions of clinical case histories. The MYCIN knowledge base currently consists of approximately 400 such rules. Each rule consists of a set of preconditions (the 'premise') which, if true, justifies the conclusion made in the "action" part of the rule. For example: If 1) the gram stain of the organism is gram negative, and 2) the morphology of the organism is rod, and 3) the aerobicity of the organism is anaerobic, then there is suggestive evidence (.6) that the identity of the organism is bacteroides.

Such rules form modular "chunks" of knowledge about the domain, represented in a form that is comprehensible to a clinician.

The consultation system uses its collection of rules to make conclusions about the patient. If, for instance, it is attempting to determine the identity of an organism responsible for a particular infection, it retrieves the entire list of rules which, like the one above, conclude about identity. It then attempts to ascertain whether the conclusion of the first rule is valid, by evaluating in turn each of the clauses of the premise. Thus, for the rule above, the first thing to find out is its gram stain. If this information is already available in the data base, the program retrieves it. If not, determination of gram stain becomes a new subgoal and the program retrieves all rules which conclude about it, and tries to use each of them to obtain the value of gram stain. If, after trying all the relevant rules, the answer still has not been discovered, the program asks the user for the relevant clinical information which will permit it to establish the validity of the premise clause. Thus, the rules "unwind" to produce a succession of goals, and it is the attempt to achieve each goal that drives the consultation.

The use of a rule-based representation of knowledge makes it possible for the system to explain the basis for its clinical recommendations. For example, if the clinician asks "How did you determine the identity of the organism?" the program answers by displaying the rules which were actually used, and explaining, if requested, how each of the premises of the rules was established. This is something which the clinician can readily understand, and it provides a far more comprehensible and acceptable explanation than would be possible if the program were to use a simple statistical approach to diagnosis.

As work proceeds to expand the program's knowledge base, new "chunks" are added in much the same way that a clinician in training learns new pieces of knowledge about his field. This rule-based representation of knowledge means that the expert himself can offer new "chunks" of knowledge by expressing them in the same rule-based format. He can thus help make the program more competent, without having to know anything about computer programming. In addition, since the rules are independent of one another, and are used by the program as necessary in order to deal with the particular consultation underway, the addition of a new rule or modification of an existing rule does not require alteration of other items in the knowledge base, as is often necessary with systems using the decision-tree methodology.

In order to represent the medical knowledge of infectious disease diagnosis and therapy for MYCIN, the designer of the rule base (the "expert") must first decide on the set of concepts that the rules will relate. In domains that are not fully codified already, this is a significant problem because there are many ways to think about the world and many ways to describe the organization of the descriptive terms. The whole methodology depends on the expert's choice of an <u>adequate conceptualization.</u>

The "grain size" for conceptualizing the domain will change as the expert gathers more experience with the performance of the system. That is, new distinctions have to be made and, occasionally, separate concepts need to be merged into a single one. The representation of rules can support this now as long as the descriptive terms are organized hierarchically and the finer grain terms are added to the lower ends of the hierachy. Reorganizing the whole tree is more difficult than I would like.

We have found also that experts give the program rules that are nice general principles, in a desire to codify the domain, but that these general rules have so many exceptions that high performance is impossible to achieve until the exceptions are also codified. This means that the expert has to be reminded to put in all the exceptions to the general principles, and has to spell out the exceptions in tedious detail. Moreover, it means that the representation for the rules mut tolerate many special case rules and exceptions and still be efficiently interpreted.

As the amount of detail required for good performance grows, the amount of time the system needs for obtaining the information will also grow. This may show up either as increased I/O time, to ask the user for the details required for making new inferences, or as increased computation time for reasoning through longer inference chains. We find that we are caught between the necessity of asking users for more and more data and the necessity of keeping the whole consultation short enought for users with little time and patience.

The formalism used by MYCIN can now include rules that are triggered by events (in particular by new information), but it is not yet possible to set up expectations of future events that will confirm or disconfirm a current hypothesis. This is symptomatic of the difficulty we have in representing time dependent relations in static rules. For example, much of a physician's reasoning is of the form "If you see X today then look for Y tomorrow, and take appropriate action if you fail to observe Y (or if you do)."

These problems, among others, are topics of current research. Some extensions to the methodology will be necessary, but we still believe in the basic idea of goal-directed reasoning from an expert's set of rules.

<u>ON THE KNOWLEDGE ACQUISITION PROCESS IN APPLIED A.I. SYSTEMS (H. Pople)</u>

The process of knowledge acquisition, as this concept applies to the development of knowledge-based AI programs, can be considered from several different perspectives. There is, for example, the -knowledge acquisition process employed by the computer scientist who is called on, typically to develop a model of the expertise required to function productively in some problem domain. If we take such a model as given, there is then the process of knowledge-base accretion, which expands the knowledge base in accordance with the fixed precepts of information structure and process underlying the given model. I would like to concentrate my remarks on the former of these two modes of knowledge acquisition, then discuss briefly what implications this carries with respect to the general concept of computer-based "knowledge engineering".

What I consider to be the primary mode of knowledge acquisition for knowledge-based systems was brought into focus recently during an exposition of INTERNIST. This is the system we have developed at the University of Pittsburgh to provide cognitive support in the formation and solution of difficult clinical problems in internal medicine (Pople, et al. 1975; Pople 1975, 1977). Dr. Jack Myers revealed that when he and I embarked on the task of building a computer-based expert consultant for internal medicine some six or seven years ago, he knew nothing of computers and I knew even less of medicine. This prompted a noted philosopher of science in the audience to wonder aloud how we ever managed to get started. How were we able to ask useful questions of one another and make use of the answers provided?

On the basis of my experience in this mode of knowledge acquisition, I am inclined to characterize the modeling of expertise by a naive inquirer interacting directly with an expert as essentially an empirical process. One begins by acquiring a small number of facts, on the basis of which some hints may emerge as to the expert's structuring of knowledge. These structures in turn may provide cues as to the process by which this knowledge is accessed and used in the course of reasoning and problem solving. Further process cues may be obtained from direct testimony of the expert, from a study of protocols and the expert reasoning aloud, and from the inquirer's own introspections concerning the structure and process of knowledge in other domains.

The next step is to fashion from these empirical findings a testable hypothesis; and here, for the computer scientist, there is a considerable armamentarium including but not limited to the models and methods of AI that can be brought into play in the construction of a working model. It should perhaps be emphasized that what results from this stage is at best a model of the inquirer's concept, which may or may not bear much resemblance to that of the expert. Nonetheless, if reasonably faithful to the inquirer's emerging concept of the domain, the model can serve to guide and sharpen the further search for more subtle aspects of the expert's information structure and process.

There are a number of implications that flow from this characterization of the knowledge modeling process that I would like to put forth as topics for consideration by this panel.

First, the term "applications" often used to characterize this sort of investigation is something of a misnomer. It suggests the existence of a general theory which is being instantiated by a "knowledge engineer", whereas experience suggests that the modeling of expertise is primarily an empirical theory-formation activity.

A corollary of this proposition is that the models and methods of AI, which often prove useful at various stages of the modeling process, should be regarded merely as tools of the investigator - not theories. One should not set about to fit a model of expertise to the models of AI. If would be better to devise new methods and techniques even at the risk of being called "ad hoc" if this is necessary in order to deal with the essential nature of the problem domain being investigated.

It also follows that one must be prepared to throw over any or all of a given model when further investigation reveals subtleties of expertise that cannot be represented within that framework. One must take whatever comfort there is in knowing that the model has served it's purpose if it has sharpened the investigator's awareness of the structures and processes that underlie expertise and then move on to the more sophisticated models and experiments enabled by this new level of understanding.

Tt is clear that there are certainaspects of knowledge acquisition currently amenable to computer-based support. For example, at any stage of model development, there is typically the need to have encoded a sizable corpus of domain specific information so that realistic experiments may be carried out. Many techniques can be employed to facilitate this knowledge-base accretion process through direct interaction with experts in the domains being modeled, through induction on cases, or simply through batch entry of prestructured data. While these methods of knowlege accretion are important, they constitute only a small part of the knowledge modeling process that we have here characterized as "empirical theory formation". Until we are able to model the role of naive inquirer and incorporate this critical aspect of the knowledge acquisition process in a computer-based system,it seems clear that the most significant mode of knowledge exchange in the development of knowledge based systems will continue to be person to person.

REPRESENTATION OF KNOWLEDGE FOR REASONING IN MEDICAL CONSULTATION SYSTEMS (C. Kulikowski)

Building a flexible and sophisticated knowledge based expert consultation system in medicine is a formidable task because of the complexity and heterogeneity of medical knowledge and our very limited understanding of clinical reasoning processes. During the past five years several artifical intelligence approaches have been taken in building descriptions of patients and diseases that combine knowledge from a variety of sources with a diversity of structural representations, and experiments have been carried out with an equally varied array of inferential problem solving strategies. (Kulikowski and Weiss, 1971; Amarel $ Kulikowski, 1972; Shortliffe 1974, 1975; Pople et al, 1975; Rubin, 1975; Silverman, 1974; Pauker, et al. 1976).

In the CASNET program which we developed at Rutgers (Kulikowski § Weiss, 1971; Weiss, 1974), the knowledge of a group of clinical researchers in the domain of a disease is structured in the form of a causal-network representation which describes the mechanisms and evaolution of disease processes. The variability of individaul observations is accounted for by postulating a related associational structure of observations. The causal network then stands as its underlying conceptual model.

The CASNET/GLAUCOMA system is an application of CASNET in the domain of the glaucomas. The system can be utilized in a variety of reasoning modes to provide diagnostic, prognostic, and therapeutic recommendations, together with explanations and references to diverse, alternative expert opinions.

A novel characteristic of the CASNET/GLAUCOMA system is that it can, for a particular case, simultaneously present alternative opinions and reasoning derived from different consutlants. To provide the system with a variety of opinions we have established a computer-based network of collaborating glaucoma expert consultants who share in the development and testing of the programs. This ophthalmological network (ONET) includes glaucoma researchers at the Mt. Sinai School of Medicine, Johns Hopkins University, Washington University, the University of Illinois at Chicago and the University of Miami.

By representing in the computer detailed patterns of disease evolving with the passage of time, we are able to deal with multiple follow-up visits. Sequences of suggested therapies for the various types and stages of progression of glaucoma have also been incorporated in the system. Specific knowledge of disease is continually added by the ONET members. This is done by presenting the computer program with a variety of difficult clinical cases and weighing its performance against the judgment of the panel of consultants. Their suggestions are used to refine the diagnostic and therapeutic recommendations, improve the systems' assessment of signs and symptoms, and to perfect specific techniques of acquiring and displaying the clinical data. Currently, the program runs and cases are stored on the SUMEX-AIM computer at Stanford, which is accessed via TYMNET from Rutgers and from the ONET sites. The stored cases form a data base which serves as a source for clinical studies on prognostic indicators and treatment evaluation. Selected results from such studies can then be used to improve the model of disease. The development of the glaucoma system with the help of ONET has demonstrated the feasibility of collaboration between geographically remote medical investigators working on a common computer-based research project.

The CASNET/GLAUCOMA system was demonstrated and tested at the 1976 Meeting of the Anerican Academy of Ophthalmology and Otolaryngology. In addition, the system was part of the symposium on glaucoma which was held at this meeting, where its recommendations were compared to those of a panel of experts. The results of this demonstration and evaluation were most successful. Most of the ophthalmologists who tested the glaucoma consultation system judged it to have reached a very competent-to-expert status(95%). From the comments received at the meeting, it is clear that the scientific and health care significance of the system is widely recognized by now.

More recently, we have explored several different representations for medical knowledge that extend the range of concepts and relationships beyond those defined in the CASNET formalism. Our motivation has been dual: a) to provide more modular descriptive capabilities so that explanations of reasoning can be more explicit than in CASNET; b) to broaden the scope of reasoning strategies as they apply to a greater variety of conceptual elements. The first alternative representation that we used was the <u>semantic network</u>, which permitted us to define relationships other than causality at the process level (preconditions, complications, triggering effects, etc.). The strategies for decision-making were formulated as <u>normative rules</u> for propagating information and confidence judgments between nodes in the semantic net in response to patient data. This provides a very general and flexible means of tracing the flow of reasoning for a particular case. A system called IRIS, embodying these ideas is now at an advanced state of implementation (Trigoboff ft Kulikowski, 1977).

A more recent development involves the formulation of a template-based description, which is well suited for building of the anatomical models that are associated with the existing process models. We have most recently begun to develop facilities for aiding in knowledge acquisition from the expert using the semantic net and template approaches. A facility for building a model in neuro-ophthalmology (visual pathways, and related anatomical structures together with the physiology of neural transmission) is being developed (Mathew, Kulikowski, and Kaplan, 1977).

In conclusion, we believe that the introduction of knowledge-based systems in medicine may help to <u>reduce the gap between science and practice</u> in this area. Medicine is primarily concerned with diagnostic decisions and with the design of courses of action for the treatment of patients. Knowledge of different types, some of which comes from research in the life sciences, enters into the reasoning processes of the medical decision maker.By studying medical problem solving in a computer, in the context of models that embody physiological/pathological knowledge of body processes, we expect to obtain a better understanding of how this knowledge can be used in the reasoning processes that lead to medical decisions. If we establish systematic ways of using such models in medical problem solving, then new knowledge about a body process, a basic mechanism of a disease, etc.,can be related to clinical observations and actions and thus it can be used and tested in medical practice. The contribution of the computer system in these situations is in handling the complexity of the relationship between a large and changing body of relevant general knowledge and decisions in special cases.

The work on applications of AI to the development of medical consultation systems is part of our research in the Rutgers Research Resource on Computers in Biomedicine which is supported by the Biotechnology Resources Program of DRR, NIH.

REMARKS ON KNOWLEDGE-BASED PROGRAMS (W. Martin)

In the past few years we have worked on a number of knowledge-based programs at the M.I.T. Laboratory for Computer Science. These include:

1) A present illness program focussing on kidney disease. (Pauker et al, 1976; Szolovitz and Pauker, 1976)
2) A digitalis therapy advisor (Silverman, 1975; Swartout, 1977).
3) An acid/base electrolyte disequilibrium advisor.
4) A program for the design of procurement systems. (Bosyj, 1976).

5) A program to assist in simple financial projections.
6) A program to write simple programs.

The work related to medicine has been funded by the Bureau of Health Manoower and involved Professors G.A.Gorry (Now at Baylor Med.Center), and P. Szolovitz and Drs. W. Schwartz, S. Pauker and J. Kassirer at Tufts New England Medical Center The work in management and programming technology has been funded by ARPA and involved Prof. A. Hax and me.

One can divide such interactive systems, or functions of an eclectic system into three types:

A) Operator based -- The system defines data objects and operators which may be applied to them. If is up to the user to determine a usefule sequence of operators. A high level language I take to be a weak form of operator based system.
B) Model based -- The system provides one or more models which can be parameterized and run against the user data.
C) Knowledge based -- The system is operator and/ or model based, but it also 1) can exhibit self knowledge in the form of explantation, 2) is able to help the user formulate his problem because it has knowledge of situations which often can be cast in a form to which its operators or models would apply.

There are several unresolved alternatives concerning the best architecture *for* knowledge based systems. Consider first the problem of explanation. In the digitalis therapy advisor we have employed what might be called "semantic programming". All functions and variables have been called by the names they would have in English and the programs have been structured, with details supressed into subroutine calls. It is then possible to write a simple routine which translates a subroutine into English - high level subroutines give an overview of the processing and lower level ones give more details. By describing a routine or its execution in English the system can answer "How would you" or "How did you" questions. One advantage of this approach is that we are guaranteed that the explanation is not out of date with the code. There are two disadvantages. First, it is necessary that the user understand the terms and the model that is actually used in the calculations. Since there is common terminology among users of digitalis this is not a problem with that program, but in experimenting with managers querying a data base we found a wide range of terminology for the same model concept. (Malhotra, 1975). A group at IBM Yorktown have opted to build a separate knowledge base for interpreting the models to the users. This is obviously a difficult strategy to implement. Even the strategy employed in the digitalis advisor forces a careful structuring of the program for clarity of explanation. Those espousing production rules sometimes claim that even this can be avoided. The work of abstraction and organization turned over to the system. The best approach is not clear.

On a different but related dimension we have observed that our present illness program exhibits bad medical style. It often asks apparently unneeded questions; it sometimes runs on and on pursuing ever less plausible dead ends, and when it works correctly, it occasionally does things in a nonstandard order which is confusing and annoying.The standard style of operation of some of today's most sophisticated medical diagnostic programs(Pauker et al.,1976;Shortliffe,1974; Pople et al.,1975) involves the application of a single complex cyclical processing algorithm to a uniformly structured database of medical knowledge. The processing algorithm repeatedly performs the following steps: 1) select a disease hypothesis, network node, or heuristic rule to invesitgate; 2) get some relevant fact from the database or program's user; 3) update the state of the database according to the interpretation of the new fact; 4) re-iterate this process

"In its responsiveness to new information this scheme is *very* attractive. However, it is very difficult to make a program structured in this way exhibit a natural clinical style. Because the database contains information only about diseases and their symptoms, and not about the state of the consulting program, the responsibility for both correct medical conclusions and appropriate medical style rest on the purely medical knowledge and its use. The expert physician seems to have a great deal of specific experience which suggests to him definitive ways of accomplishing a diagnostic goal. What we need is a way to add such strategies to our program. One approach is to view these specific expert strategies as replacements for parts of the general diagnostic strategy, which replacements apply only in specific circumstances. Thus, before running a step of the general diagnostic engine, the system would first look to see if there were a more specific strategy pertinent. The recognition of when a strategy is applicable is a classic problem of AI. Here it will probably require that the data be characterized at different levels of abstraction. Currently, data are abstracted only to a small number of classes - disease, symptom, etc. - used by the single diagnostic cycle.

A further, nagging thought is the apparent orthogonality of the abstractions suggested by Schank to those required to cast all data into a particular scientific model or user point of view. While Schank seeks a uniform way to represent all knowledge which is problem independent, all useful systems arc problem dependent - their special model of the problem is at the heart of what they have to offer. To date, Schank's abstractions have had litle, if any, impact *on* knowledge based system - yet they do seem to capture important ideas about the representation of knowledge. What is going on here? Is it that Schank's techniques provide a level of generality to which we have not yet aspired, should not aspire, or will soon need to aspire?

OBSERVATIONS ON THE DEVELOPMENT OF EXPERT KNOWLEDGE-BASED SYSTEMS (P. E. Hart)

The last five years have seen an accelerated interest in applications of artificial intelligence with what to my mind are beneficial effects both scientifically and politically. The most immediately practical of these applications use AI techniques that are well-understood. For example, the commercial picture processing and word recognition systems available today are based directly on techniques developed in the late sixties and early seventies. Other applications, which may have equally practical goals, require more profound use of the methods of artificial intelligence. These applications are more interesting from a scientific point ofview, and I shall examine a few issues that they raise in the following.

Among the most prominent of these application themes is the development of "expert knowledge-based systems." These systems — DENDRAL (Buchanan, et aL, 1969), MYCIN (Shortliffe, 1974), INTERNIST (Pople, et al., 1975) and others -- have forced AI technology ahead because of the need to represent and deploy substantial amounts of real-world, or domain-specific, knowledge. Further, because the computer scientists involved in the design are usually inexperienced in the subject domain, there is a heightened need to develop methodologies for extracting relevant knowledge from those who are experienced. This is in sharp contrast to an earlier period of AI research, when the designer, programmer, and domain expert could Be one and the same person.

An expert system under development at SRI furnishes a number of instructive examples of how current applications efforts are advancing AI technology. We first briefly describe the systme, and then focus on some specific examples. PROSPECTOR (Duda et al., 1976; Duda et al, 1977) is a developmental system aimed at aiding geologists in their search for mineral resources. In some respects, the mineral exploration process is similar to the process of diagnosing diseases. A body of observations, often uncertain in nature, must be interpreted with the aid of a knowledge base that typically supports plausible reasoning but not strict logical inference. It is not surprising, therefore, that PROSPECTOR bears some resemblance to existing medical diagnosis systems,particularly to MYCIN.

Domain specific knowledge is given to PROSPECTOR as a set of inference rules of the form A IMPLIES B. Two numbers associated with the rule measure the degrees to which A is necessary and sufficient for B. A prior probability is also associated with the conclusion B.

The inference rules are represented internally as a partitioned semantic network along the lines suggested by Hendrix. Such networks when used in their general form have all of the expressive power of predicate calculus, and additionally encourage the exploitation of subset-element relations, inverted indexing on terms, and a number of other such features. We are currently using a pecialized form of semantic network with which we can easily represent both an important subset of predicate calculus formulas, and the inference rules themselves.

The top-level nodes of the network correspond to top-level hypotheses about the presence of the various types of ore deposits. Lower-level nodes may correspond to directly observable geologic data, or to intermediate concepts that cannot be directly observed but that can be inferred from observables. A principal task of the system is to infer probabilities for the top-level hypotheses on the basis of available observations.

The special characteristics of the mineral exploration process force the system designer to consider a number of issues that have not, to my knowledge, received much attention to date, but that could well have relevance to existing and future systems.

Questions of Existence

The hallmark of exploration is the search for the existence of objects with specified properties. Thus, the quintessential question for the user (i. e., the geologist) is "Can you find an object X satisfying P (X)?" The issue raised here by PROSPECTOR is in interesting contrast to medical diagnosis systems, for which the existence of a culture, a test, or indeed a patient is eldom in question. On philosophical grounds, it is always hard to rule out the possibility of finding something yet undiscovered; on practical grounds, a committment to the absence of X sometimes needs to be made, leading to a need to handle the statement "There does not exist an X satisfying P(X)". PROSPECTOR currently handles this by denying the existence of any "real-world binding" for a formal object satisfying P. However, although this statement is logically equivalent to the statement "For all X not P(X)" PROSPECTOR cannot handle the very similar statement "For all X Q(X)" except by setting Q equal to not P. This stems from a representation that highlights the possible existence of objects, rather than one that stresses logical completeness.

Problems of the existence of objects also lead to interesting choices about how to ask questions of the user. For example, a user may doubt (but not rule out) the existence of an X satisfying P(X). If we want to ask him about other properties of X, we may find ourselves in the position of asking "Does X,which probably doesn't exist (but may), also satisfy Q(X)?" PROSPECTOR handles this issue by simply resorting to a threshold probability for X below which X is deemed to be non-existent, but this expedient has little theoretical justification.

Antecedent reasoning

A typical consultation with a field geologist would begin by having the geologist tell the system about the significant geological features of the case, or "prospect" of interest. Ideally, the system should be able to use this volunteered information to help it focus on hypothesesof interest. It is therefore, important to be able to reason in the forward direction from volunteered features to the top-level hypothese that they suggest (even weakly). The structure of PROSPECTOR supports antecedent reasoning through the use of an explicit network representation of inference rules, and through the use of a taxonomy discussed below. It is interesting to note that Pople's INTERNIST system accomplishes the same aim in a very different system framework.

Taxonomies

A surprising amount of geologic reasoning is done through taxonomies of objects. For example, chalcopyrite is a copper-iron sulfide, which is a copper sulfide, which may be significant evidence for a certain type of copper deposit. PROSPECTOR easily handles such taxonomies using the semantic net mechanism. However, the current state of geology does not support the use of taxonomies indiscriminately since, for example, many rocks can be class-