医療情報学 1－1
(1979. 8. 25)

CURRENT RESEARCH ON ARTIFICIAL INTELLIGENCE IN MEDICINE IN THE UNITED STATES

Casimir A. Kulikowski, Ph.D.

Department of Computer Science
Rutgers University
New Brunswick, New Jersey

Abstract

Recent years have seen the development of many new approaches to decision-making in medical consultation. These include both artificial intelligence (A.I.) methods and the introduction of structural constraints in the context of more classical probabilistic models.

The major problems that arise in designing a consultation program involve choices of knowledge representations, diagnostic interpretation strategies, and treatment planning strategies. The need to justify decisions and update the knowledge base in the light of new research findings places a premium on the modularity of a representation and the ease with which its reasoning procedures can be explained.

Most current A.I. consultation systems use either a semantic network (including special cases such as causal or taxonomic nets) or a frame schema to represent descriptive knowledge of disease processes and associated patient findings. The normative knowledge is usually expressed as a system of decision rules attached to the semantic net nodes or as logical constraint conditions attached to frames. Some major representational problems currently include: how to choose the level of abstraction at which hypotheses are to be made explicit (different levels are often needed for different diagnostic problems); how to maintain consistency between different schemas of disease taxonomy; how to integrate temporal information into the logical rules that describe diseases; and how to represent multiple competing points of view from different medical experts as well as "consensus knowledge."

In both diagnosis and treatment decisions, the relative advantages and disadvantages of different schemes for quantifying the uncertainty of inferences raises difficult issues of a formal logical nature, as well as many specific practical problems of system design. An important insight that has resulted from the design of several artificial intelligence systems is that robustness of performance in the presence of many uncertainty relationships can be achieved by eliciting from the expert a segmentation of knowledge that will also provide a rich network of deterministic relationships to interweave the space of hypotheses.

Artificial Intelligence Methods in Consultation

Several somewhat different A.I. methods have developed over the past several years. All of them can be contrasted to previous methods by the greater complexity of medical knowledge that they explicitly represent on the computer. They all have data structures that permit the expression of semantic relationships between the facts in their knowledge bases that go beyond the simple probabilistic or heuristic weights used as links by previous formal models. Yet they differ widely among themselves in their scopes and their representations of medical knowledge.

The sequence of evolution of the earliest, and by now most developed A.I. consultation systems is shown in Table 1, together with their institutional and project affiliations, and their major characteristics in terms of representation of knowledge, strategies of reasoning, and medical domain of application. Specific references to each of the systems are also listed. A comparison of the first three systems in terms of their reasoning methods can be found in (10) and (16). Since a detailed analysis and comparison is beyond the scope of the overview intended in the present paper, only a few of the more distinguishing aspects of the programs and their underlying methods are discussed. It is tural to seek connections between the stated goals and motivations of the system designers and the resulting choices of knowledge representation and reasoning strategies.

TABLE 1

CHARACTERIZATION OF "FIRST PHASE" A.I. CONSULTATION SYSTEMS

Year of First Prototype	Name of Medical System Domain of Application	Institution/Project	Representation of Knowledge Descriptive	Normative
1971	CASNET/Glaucoma	Rutgers Univ./ Rutgers Research Resource on Computers in Biomed.	Causal-Associational Network	Implicational links among findings, hypotheses, and treatments
1972	MYCIN/Infectious Diseases: (Bacteremias)	Stanford Univ./ SUMEX-AIM	Context tree and properties of data structures	Production rules
1973	DIALOG (renamed INTERNIST 1) /Internal Medicine	Pittsburgh Univ./ Clinical Decision Laboratory	Hierarchical (taxonomic) network with some causal links	Implicational links among findings and hypotheses
1974	PIP(Present Illness Program) /Internal Medicine- /(Renal Disease)	M.I.T.-Tufts	Frames for diseases [Long-term memory structure]	Logical Constraints within and between frames

System	Focusing	Reasoning Strategies Diagnostic	Prognostic	Treatment	Explanation	Ref.
CASNET	Global assessment over causal net	Elicitation of intermediate and high level hypotheses	Follows pathways in the Causal net	Elicitation of major treatment plans from diagnoses, modified through global preference score from the findings.	Listing of patient-specific causal pathways and contribution of evidence to each hypothesis.	[1] [2] [3] [4] [5]
MYCIN	Guided by context tree	Goal directed backward chaining thru production rules from high level to low level goals			Listing of patient specific production rules and their sequence of application	[6] [7] [8] [9]
INTERNIST	Partioning heuristic	Scoring heuristics to combine weights from evidence to hypotheses & vice-versa	---	---	---	[10] [11] [12]
PIP	Heuristics for activating hypotheses into short-term memory	Scoring heuristics combine weights from logical constraints	---	---	---	[13] [14] [15]

In order to build the knowledge base, or general model for a disease domain a program for knowledge acquisition can be used. It will guide the user in the construction of data structures that encode knowledge in forms acceptable to the other programs. An alternative is to specify a syntax for describing a consultative model, translate the expert's descriptions and rules into the syntax in a file, and have a parsing program interpret this 'model file'.

The task of producing a patient-specific interpretation is that of the consultation program proper. Its major components include: a data acquisition module or scheme for eliciting patient data from the program user; a focusing component which will direct the reasoning of the program at any time to concentrate on specific sub-areas of the knowledge base that are most relevant to the interpretation of the client's findings (or other sub-goals of the system); an inferencing component that applies strategies and tactics from the knowledge base in the appropriate combinations and sequences for interpretation; a planning component that does likewise for producing treatment recommendations; and an explanation component for providing explanations of the program's reasoning to the user.

A separate question answering program may also be present in some systems to permit the user to either query the knowledge base or obtain more detailed traces of the reasoning about a specific case.

In all the A.I. systems, the interpretation of a specific patient's data can be viewed as the generation of a patient-specific model, which is obtained by extracting from the knowledge base the subsets of general descriptive knowledge and the normative rules that are most applicable to the patient. This process of extraction is ultimately characterizable as some form of pattern matching between the incoming specific patient data and the possible patterns of data that may be expected under various hypotheses about the patient and the health care environment within which the data is produced. The generation of these hypotheses can be carried out in response to the formulation of a set of goals that the program may pursue in a top-down manner from higher level goals (such as recommending a satisfactory treatment for the patient) to the lower goals (such as finding a diagnosis, or elicting contraindications for the possible treatments)that enable the higher level ones to be finally satisfied. When the goals and their structure are explicitely described within the knowledge base it is possible to use simple strategies of inferencing and planning to produce patient-specific interpretations. This is the case with the MYCIN system, which concentrates its knowledge in the form of production rules that specify both general strategies and specific interpretation tactics. The inferencing scheme is then a simple backward cleaning mechanism that seeks to either prove (or extract directly from patient data) the preconditions (antecedents) that will satisfy the consequents of a rule that results in the completion of successively higher level goals.

There is relatively little flexibility in the focusing scheme of MYCIN: it is guided by a context tree that is pre-specified for the program. MYCIN does not have an explicit general descriptive model of disease mechanisms; though the level of abstraction at which hypotheses can be related is implicitly constrained by the choice of properties defined for each of the data structures in the knowledge base. One of the major characteristics of MYCIN is the emphasis placed on a separate question-answering program in addition to the explanation program that permits a careful tracing of the reasoning taking place in the consultation session. The simplicity of the backward chaining strategy is particularly valuable in facilitating such tracings.

In contrast, the CASNET, or causal associational network model, is deliberately designed to represent a generalized description of disease processes as they evolve over time. The causal relations express the mechanisms of a disease and their modifications under various regimens of treatment. Different patterns over the causal network are associated with the various elements in a classification scheme of diagnostic hypotheses, that can include degrees of severity and progression of a disease. Appropriate treatment plans can be associated with the diagnostic hypotheses, and specific treatments within the plans related to each other by constraints of how they cover for particular illnesses, how they may interact, etc. Normative knowledge is in the form of implicational rules that link intermediate hypotheses about pathophysiological states to patient findings,with similar rules for treatment planning.

The price paid for this greater explicitness of structural description is an increase in the complexity of specialized strategies in the inferencing and planning components of the consultation program. This is valuable when it helps to strengthen the chains of inference, which can happen in domains, such as glaucoma, where mechanisms of disease are reasonably well understood, and where their understanding is important in the planning of treatment. There are many situations in medicine where mechanisms are only poorly understood, if at all, and treatment is almost entirely empirical in nature. In such domains, causal relations will be sparse, and the major conceptual constructs remaining will be taxonomies of diseases based on anatomical locus (i.e. liver diseases), or a common systemic relationship (i.e. hematological diseases). Such relations can also be described in the CASNET formalism. From the point of view of strategies, there are distinct ways in which the explicit CASNET model can be used to guide focusing and perform inferences. Taking advantage of the partial ordering implicit in causal networks, the CASNET model can be precompiled, which permits rapid calculation of the propagation of confidence weights over the entire network (a global assessment) every time a new piece of evidence about a patient is given to the program. Unlike MYCIN, there is no explicit statement of strategies in the normative knowledge base - they are implemented in the consultation program itself which applies them selectively to the appropriate elements from the knowledge base. In general, inference strategies can be characterized as propagating confidences from findings to lower level hypotheses, which in turn propagate to higher level hypotheses. This is often referred to as bottom-up processing in artificial intelligence as opposed to the goal-oriented top-down reasoning of MYCIN. CASNET only places minor emphasis on focusing because its global assessment strategy obviates the need for it in restricted medical domains. Several elements of hypothesis-directed (top-down) reasoning enter into the selection of findings and the elicitation of treatment plans from the diagnoses.

A system that explicitly relies for its reasoning strategies on a hierarchial network of disease hypotheses is INTERNIST-I, which performs diagnostic consultation in a CPC mode over the field of internal medicine. INTERNIST-I is now reported to cover 80% of the diagnostic knowledge in internal medicine. It places strong emphasis on a focusing heuristic to narrow down the scope of hypotheses, that must be considered in the patient-specific-model from among the large number of possible diseases in internal medicine. This heuristic partitions the possible hypotheses into those that cover the subset of the patient's data that is being focused on. Various scoring methods are used to build up confidence in the different diagnostic hypotheses (using weights from both findings-to hypotheses and from hypotheses-to-findings), and once only a few are left in contention a discriminatory strategy is applied to arrive at a differential diagnosis.

The PIP system also has a strong descriptive component, characterizing diseases by the frame formalism of Minsky. In building a patient-specific-model, the consultation system has focusing heuristics that bring frames out of the knowledge base (also called long-term memory) and put into an "active" or "semi-active" status. These conditions enable various reasoning strategies to take advantage of the normative rules that are attached to the frames, and through a scoring algorithm, arrive at a ranking of diagnostic hypotheses. A major motivation of the PIP project was to develop a better understanding of clinical cognitive processes, and to this effect a protocol analysis was carried out. In the spectrum of complexity of structure, the PIP formalism is potentially the most complex. With it comes the attendant difficulties of specifying strategies that will take advantage of all the descriptive relations in the knowledge base. INTERNIST has as its goal the incorporation of the rules of interpretation and some descriptive knowledge structures of a single expert diagnostician. CASNET/Glaucoma provides a formalism where certain elements of the model (such as the causal network) may represent a consensus about disease mechanisms reached by a group of expert investigator clinicians. The rules of interpretation and treatment planning may, however, reflect several different viewpoints, which can be presented as alternatives. Both CASNET and INTERNIST have descriptive structures of an intermediate level of complexity with well-proven strategies that have allowed the programs to give expert-level advice in many complex cases. CASNET/Glaucoma has been demonstrated and tested against a panel of experts at the National Academy of Ophthalmology and Otolaryngology, with very satisfactory results (17) and INTERNIST-1 routinely handles the CPC cases from the New England Journal of Medicine.

MYCIN, which has the least descriptive structure, has also achieved performance that is comparable to that of a group of experts in infectious diseases (18). Thus, the first generation of A.I. consultation systems has shown the feasibility of obtaining both high performance and many elements of understanding (through the structured knowledge bases) in clinical consultation.

New A.I. Consultation Systems:

Recent developments have tended to generalize the methods and representations that were generated during the first phase of system investigation and development. The production rule formalisms of MYCIN have been studied (9) and generalized (19). The CASNET approach has also been generalized by replacing a strictly causal net by a partially ordered network description for diseases in the EXPERT system (20). There have been new strategies of focusing implemented in the INTERNIST-II program (12).

There has also been a movement to study knowledge acquistion and updating processes in greater detail and link consultation systems to data bases (21,22). Attempts at extracting general, simple, and streamlined modules from existing systems to be used as tools in structuring knowledge is the main theme of the emerging field known as knowledge engineering (23). The first development in these new directions are promising, since they show that many of the existing concepts and tools may indeed find application in a large variety of clinical problem-solving situations and medical domains.

References

1. Weiss, S., Kulikowski, C.A., and Safir, A., "Glaucoma Consultation by Computer," Computers in Biology and Medicine 8:25, 1978.

2. Kulikowski, C.A., and Weiss, S., Technical Report TR-3, Computers in Biomedicine, Rutgers University, 1971.

3. Weiss, S., "A System for Model Based Computer Aided Diagnosis and Therapy," Ph.D. Thesis in Computer Science, Rutgers University, 1974.

4. Weiss, S., et.al., "A Model-Based Consultation System for the Long-Term Management of Glaucoma," Proc. 5-IJCAI, Boston, August 1977, p. 826.

5. Weiss, S., et.al., "A Model-based Method for Computer-aided Medical Decision-Making," Journal of Artificial Intelligence, 1978.

6. Shortliffe, E.H., et.al., "An Artificial Intelligence Progam to Advise Physicians Regarding Antimicrobial Therapy," Computers in Biomedical Research 6:544, 1973.

7. Shortliffe, E.H., and Buchanan, B., "A Model of Inexact Reasoning in Medicine," Mathematical Biosciences 23:351, 1975.

8. Shortliffe, E.H., et.al., "Computer-based Consultations in Clinical Therapeutic Explanation and Rule Acquisition Capabilities of the MYCIN System," Computers in Biomedical Research 8:15, 1977.

9. Davis, R., Buchanan, B., and Shortliffe, E.H., "Production Rules as a Representation for a Knowledge-based Consultation Program," Artificial Intelligence 8:15, 1977.

10. Pople, H., "Artificial Intelligence Approaches to Computer-based Medical Consultation," Proceedings of the IEEE, Intercom 31/3, 1975.

11. Pople, H., Myers, J., Miller, R., "DIALOG: A Model of Diagnostic Logic for Internal Medicine," Proceedings of the 4th IJCAI, Tiblisi, 1975, p. 841.

12. Pople, H., "The Formation of Composite Hypotheses in Diagnostic Problem Solving," Proceedings of the 5th IJCAI, Boston, 1977, p. 1030.

13. Rubin, A., "Hypothesis Formation and Evaluation in Medical Diagnosis," AI-TR-316, MIT, Cambridge, 1975.

14. Rubin, A., "The Role of Hypotheses in Medical Diagnosis," 4th IJCAI, Tiblisi, 1975, p.856.

15. Pauker, S.G., et.al., "Towards the Simulation of Clinical Cognition," American Journal of Medicine 60:981, 1976.

16. Szolovits, P., and Pauker, S.G., "Categorical and Probablistic Reasoning in Medical Diagnosis," Artificial Intelligence (in press).

17. Lichter, P., and Anderson, D., Discussions on Glaucoma, Grune and Stratton, 1977.

18. Buchanan, B.G., Davis, R., Yu, V., and Cohen, S., "Rule-Based Medical Decision Making by Computer," Proc. MEDINFO '77, 1977.

19. Davis, R., "Interactive Transfer of Expertise: Acquisition of New Inference Rules," Proceedings of the 5th IJCAI, Boston, 1977, p. 321.

20. Kulikowski, C. and Weiss, S., Representation of Expert Knowledge for Consultation: The CASNET and EXPERT Projects, Proc. AAAS Annual Meeting, January 1979.

21. Mathew, R.J., Kulikowski, C.A., Kaplan, K.R., "A Multi-leveled Representation for Knowledge Acquisition in Medical Consultation Systems," Proceedings of MEDINFO '77, p. 151.

22. Weiss, S., et.al., "A System for the Interactive Analysis of a Time Sequenced Ophthalmological Data Base," Third Illinois Conference on Medical Information Systems, Chicago, 1976.

23. Feigenbaum, E.A., "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering," Proceedings of the National Computer Conference, 1978, p.227.

Acknowledgment

This work was supported in part by Grant RR643 of the NIH DRR-BRP.