

WORKING PAPER 65

ARTIFICIAL INTELLIGENCE APPROACHES  
TO MEDICAL DIAGNOSIS

by

ANN D. RUBIN

Massachusetts Institute of Technology  
Artificial Intelligence Laboratory  
March, 1974

Abstract

The differential diagnosis of hematuria, blood in the urine, is studied from the point of view of identifying crucial structures and processes in medical diagnosis. The thesis attempts to fit the problem of medical diagnosis into the framework of other A.I. problems and paradigms and in particular explores the notions of pure search vs. heuristic methods, linearity and interaction, plausibility and the structure of hypotheses within the world of kidney disease.

Work reported herein was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contract Number N00014-70-A-0362-00

Working Papers are informal papers intended for internal use.

Writing a program, inventing a formalism or working out a method to solve a problem is an activity which can be viewed in two very different lights. Its most immediate goal is to produce a working program or simulation, which may be used in speech understanding, scene analysis, game-playing or medical diagnosis. This more immediate point of view is the one more often discussed in papers, which report on a finished or soon-to-be-finished product. More abstractly and importantly from an Artificial Intelligence point of view is regarding the problem-solving process as an exploration of alternative approaches to representation and control structure, as the instantiation or discovery of more general concepts and theories, whose details are of lesser importance. This point of view has been particularly emphasized in AI, a field whose goal it is to investigate general problem-solving strategies and wide-ranging insights into possible patterns of human thought.

This thesis studies the problem of medical diagnosis basically from the second point of view, although it recognizes the necessity of paying attention to some of the details in any complex problem domain. It attempts to fit the problem of medical diagnosis into the framework of other AI problems and paradigms and in particular explores the notions of pure search vs. heuristic methods, linearity and interaction, plausibility and the structure of hypotheses within the not-so-mini-world of kidney disease.

## 1.1 Why Medicine?

The practical importance of studying and developing computer aids for medical diagnosis is obvious. Doctors train for years to become expert diagnosticians; they carry heavy responsibility for the accuracy of their diagnoses and the effectiveness of their treatments. Yet with all their training, they often make mistakes because of the vast body of ever-increasing medical knowledge they must remember and access. In a computer, the problem of pure memory disappears, while effort focusses instead on methods of representation of knowledge, selection of relevant knowledge and proper use of the selected facts.

Several diagnosis programs have already been written for small areas of medicine such as bone tumors and acute renal failure <ref>; a group at Rutgers is currently analyzing the time course of glaucoma and using their model to place a patient at a point along the temporal progression of the disease and thus determine the prescribed treatment. <ref> Programs have been written as well to investigate treatment choices <ref - Gorry/Schwartz> and as clinical aids in adjusting therapies. Silverman is currently working on making a program to calculate digitalis doses more sensitive to the individual patient and capable of using his or her reaction to the initial dose to revise its suggestions. <ref>

More recent attempts at writing medical diagnosis programs have been more all-encompassing, attempting to incorporate capacities for dealing with wider varieties and larger numbers of diseases, offering

coherent explanation of diagnoses, and containing more general models of the time course of diseases. In addition, there has been growing interest in the psychological processes of hypothesis-generation and decision-making as they are reflected in medical diagnosis. Medical educators envision better instruction for students in diagnostic skills such as data organization and test selection as a possible result of such research.

Another group interested in the processes involved in medical diagnosis are the cognitive psychologists and AI researchers who are interested in the structure of medical knowledge and the processes by which it is manipulated as examples of general knowledge structures and problem-solving processes.

Medicine has many characteristics which make it well-suited for such theoretical exploration:

1. There is no question that the complexity of medical diagnosis is sufficient to make it a worthwhile topic. Certainly, the data itself is complicated (or at least massive) and even a cursory glance at the kind and amount of processing which must occur is enough to justify studying it further. That there is some kind of rich structure present at least in many doctors' minds, if not in the data itself, is evident if we assume that diagnostic and question-asking strategies proceed from the same data structure; no overly-simple structure will account for the complexities of that process.

2. The final goal of a medical diagnosis system is clear, at least on one level; we want a program which will produce the "correct"

diagnosis (i.e. the same one as an "expert" would arrive at) at the end of some reasonable amount of processing. This is in contrast to the problem of defining "understanding" in a (language) understanding system. Many attempts have been made to come up with a taxonomy of the indicators of understanding <ref - Newell> <ref - Card>, but the job is not a small one. On the other hand, we notice that debugging problems do have a more clearly-defined goal: the production of a program which performs according to some externally-stated standards. <ref - Goldstein> <ref - Sussman> Of course, in both medicine and debugging, it is the process of arriving at the solution in which we are ultimately interested and the standards for judging these processes are much less well-specified or understood (but see below, 3). Still, we have at least a first-order criterion by which to judge diagnostic programs.

3. As mentioned above, process is of primary interest in looking at problem-solving programs; one problem which many such theories have had is that there was a lack of natural data giving insight into that process. The "success" of a theory had to be judged by a comparison of its results with the "correct" results - and independently by some general criteria about plausible processes. In visual recognition or language understanding, for example, there are no intermediate points in the process about which people naturally verbalize or to which we have any other access. The medical diagnosis process, on the other hand, is one which occurs and is verbalized naturally; getting informal protocols requires only sitting in on clinical sessions or listening to discussions on rounds. More formal and complete protocols are also easily

obtainable, since public diagnostic sessions and CPC's (see section 1.2) are common occurrences in hospitals. In this respect, studying medical diagnosis contrasts with taking protocols of subjects solving cryptarithmic problems, which uses an artificial task in an artificial situation, as well as with language understanding or visual scene analysis, which are certainly natural tasks, but are decision processes to which we have no natural access.

4. Medicine contrasts with vision, although both have been treated as recognition problems (see section 1.2), in terms of the vocabulary available for each subject area. Much of the work which has gone into current vision systems has been devoted to coming up with a limited yet sufficient vocabulary to describe structures as simple as vertices and angles and as complex as textures, curves and complex shapes.<ref - Fahlman working paper on vocab.> Medicine, on the other hand, comes completely equipped with a large technical (and sometimes baroque) vocabulary, whose stated aim is, in fact, to allow exact and accurate communication among doctors. Thus, a lot of effort has already been devoted to making the necessary distinctions among symptoms and disease states. We have, unfortunately, found that medical vocabulary is sometimes more confused than one would hope - definitions may be unclear and diseases may overlap. The basic structure, however, has already been laid down.

5. Medical diagnosis is so large and varied a field that it allows the construction of many different mini-worlds, the exploration of each aiming toward the clarification of different issues. Thus a problem

we often face in AI, that of finding an area small enough to study completely, yet large enough to provide real challenge, seems to be well addressed by the choice of medical diagnosis. The subject matter in medicine can be cut along many different dimensions; most often it has been limited by the selection of a small class of diseases, tests and symptoms, as well as by focussing attention on the final diagnosis to the exclusion of process. In addition, complicating non-technical issues such as the representation of time were often excluded or dealt with using special ad hoc mechanisms. For example, the Rutgers group has limited their investigation to one disease - glaucoma - and is concentrating instead on determining the stage of the disease which a patient manifests; thus the time course of the disease is specifically and exclusively considered. <ref> Gorry, on the other hand, chose a larger class of possible diagnoses and handled the time of occurrence of symptoms as one example of a general concept of interaction between symptoms. <ref - Gorry thesis> This is not to suggest that the hard problem of modularization has been solved in the case of medical diagnosis - but merely to inject some hope; the sub-domains are there, if we can only find and isolate them.

## 1.2 Description of the Problem

The particular aspect of medicine with which this thesis will deal is the process of diagnosis within a limited set of diseases: those whose presenting symptom is hematuria, or blood in the urine. We can

conceptualize the problem as one of a class of recognition problems <ref - Fahiman thesis prop.> in which features of the situation (called the sample by Fahiman) act as clues to its complete description - to its recognition as an already-known entity. In particular, a medical system is presented with a group of symptoms, signs, facts, test results etc. and its job is to come up with a diagnosis, an identification of a disease or several diseases whose manifestations most closely match the condition of the patient. Choosing a treatment on the basis of the diagnosis will not be included in the analysis here.

Because of an interest, mentioned above, in process, the model of diagnosis which will be used here is one of the serial acquisition <ref - Gorry the.> of facts about the patient. Thus, we require a diagnosis system to have hypotheses at each point and expect that these hypotheses will change after the addition of each new piece of information. As a first approximation, a hypothesis can be thought of as a disease, but several examples later will make it clear that the structure of a hypothesis is more complicated, often including several related or independent diseases, some of which are connected by relationships like CAUSED-BY or COMPLICATED-BY.

A distinction is often made between two forms of data acquisition in diagnosis : active and passive. <ref?> An active approach includes a physician's asking a question in order to solicit each new piece of information from a patient; clearly his or her questions will rely heavily on the previous dialogue and the present hypothesis. A passive mode is one in which each new piece of information is offered to the



physician in a pre-determined order. In fact, such a technique is actually used by doctors, who call it a CPC (clinical pathological case); the facts of the case are pre-arranged (often in a misleading manner) and read to a doctor who, at each stage, offers his or her current hypotheses and the reasons behind them. CPC's, unfortunately, are artificial in that the data is organized in ways which are foreign to a real doctor-patient interaction and the ensuing process may be unrepresentative of a doctor's normal strategy in making diagnoses. Thus, I have chosen to use a variation of the active process in which all the data about the patient is immediately available if the physician asks for it. This avoids assigning risks and costs to various diagnostic procedures, hopefully simplifying the problem to some extent. In this thesis, I will concentrate on the hypothesis-generation and evaluation aspects of the diagnostic process. I will not consider the question-asking strategy in detail, except as it illuminates the more general topics of data organization and hypothesis generation. The protocol below (Chapter 2) was taken from a session in which the physician actively acquired data from the patient, although I have not included his questions in my analysis of the interaction.

### 1.3 The Basic Approach

Forgetting for a moment all the complexity in hypotheses hinted at above, we can regard the diagnosis problem as a conceptually simple one for which, in fact, we can come up with a complete solution. We have

a collection of symptoms and a collection of diseases; the problem in each case is to choose the disease which is most likely causing the particular symptoms observed. In more general terms, we have a collection of effects and a collection of causes; the task is to find the cause which most likely accounts for the effects present in each particular situation. Under certain assumptions (which I will discuss below), the solution is straightforward and represents an elementary example of the use of probabilities. With each (disease, symptom) pair is associated a number which represents the probability of a patient who has the disease exhibiting the symptom. For example, if 20% of all people suffering from the flu have aching muscles, then the number associated with (flu, aching muscles) would be .2. Obviously, the number implicitly associated with (flu, no aching muscles) would be .8. Then making a diagnosis necessitates only multiplying all the probabilities associated with present and absent symptoms for each disease - and comparing all the results. The disease with the highest associated product is the winner and claims the victim.

This method is obviously generalizable to any recognition problem for which the correlations are available - given a few conditions:

1. that the symptoms are independent, in the probabilistic sense and
2. that the diseases are mutually exclusive and exhaustive.

Obviously, neither of these is true in the medical diagnosis case; patients often have more than one disease and the presence of one symptom more often than not affects the probability of the occurrence of others. Both of these non-linearities can, theoretically, be handled in the

probabilistic framework by considering all possible combinations of diseases and symptoms in recording and combining probabilities. By now, the third and most important for rejecting the above-outlined complete theory should be obvious: the uncontrolled proliferation of hypotheses and associated probabilities and the explosion of computations necessary to choose the correct answer. Even if all the numbers necessary were available (which they're not), this situation could become computationally infeasible - and is certainly cognitively impossible. It doesn't take very subtle intuition to judge that doctors are not maintaining up-to-date "scores" on every possible diagnosis. In addition, when this approach is combined with similar methods for choosing tests, the amount of processing necessary quickly gets out of hand.

So the 'complete theory seems untenable; the next step is to search for ways to reduce the number of hypotheses actively entertained at any given time and to cut down the amount of computation necessary to keep the relative status of each hypothesis up-to-date. The emphasis of the coming chapters will be on two stages in the movement away from a complete but unrealistic theory toward a 'heuristic theory which seems to model more closely the processing which physicians probably use. A brief summary of those two notions follows.

### 1.3.1 Activation vs. Deactivation : the first cut-back

A first difference attacks mainly the route through which hypotheses are actively considered. The complete theory postulates all diseases as possibilities from the beginning, eliminating them as their associated probability products go to 0. An obvious way to have fewer active hypotheses is not to consider a disease until it is suggested by a relevant piece of data. This has the reassuring consequence that every current hypothesis has a 'reason for begin remembered' - instead of just lacking a reason for begin forgotten. The issues surrounding this switch in emphasis are closely related to the concepts of 'expectation and 'evidence, which are discussed in detail in Chapter 4.

### 1.3.2 Heuristics and Interaction : the second cut-back

Both the complete theory and the modification discussed above are 'uniform theories; that is, every disease and symptom is treated the same. Some of the most powerful methods for controlling the growth of the hypothesis space, however, are much more specialized and local. They reflect knowledge about the non-independence of symptoms and the amount of details pertaining to particular symptoms which is a prerequisite to using them as reasons for considering a hypothesis. Such local pieces of knowledge will be viewed as 'compiled information, as they are derivable by general principles from the primitive data base of diseaseSymptom probabilities, but are clearly more efficient and useful in their specialized form. Chapter 5 contains an inventory of such interactions between symptoms and the imperative information associated with them.

In order to keep the number of active hypotheses at a reasonable level, it is important in addition to stop considering those whose plausibility has reached a low level and to avoid adding new hypotheses on top of old ones which have not yet been discarded as useless. Such methods are clearly 'heuristic - that is, they don't always do "the right thing" - since any hypothesis we eliminate on heuristic grounds may eventually turn out to be the correct one after all. But it seems that physicians (and, most likely, all of us) must do everything they can to keep their minds uncluttered and their short-term memories from overflowing. 'Inertia and 'premature 'rejection are two such mechanisms which will be further explored and exemplified.

#### 1.4 Anticipations

Chapter 2 contains a protocol of a doctor-patient interaction which illustrates many of the processes described above. The doctor is an expert; thus, modeling his reasoning means modeling expertise and we can expect many examples of compiled heuristics and special techniques. Chapter 3 describes a representational structure which we have developed in looking at hematuria and the diseases in which it plays an important part; the explanation of this data structure more clearly identifies the objects and relationships in a basic medical data base. Chapter 4 discusses the issue of local evaluation of hypotheses, making a distinction between 'disease-centered 'information ('expectations) and 'symptom-centered 'information ('evidence) and speculating on the place

of each in a doctor's developing expertise. Chapter 5 catalogues some of the interactions between symptoms which contradict any strictly linear theory of evaluation - and which exemplify the compiled information mentioned above. Chapter 6 continues the movement from local toward global strategies by explicitly considering the structure of both simple and complex hypotheses and a theory of 'coherence designed to provide a way of comparing competing hypotheses and choosing the most promising ones. Chapter 7 summarizes the preceding view of medical diagnosis as a 'hypothesis 'generation and 'testing problem and includes some tentative thoughts on learning and further research. The Appendix contains the data on hematuria which was collected during this research and which forms the basis for the protocol and other examples quoted in the discussions.