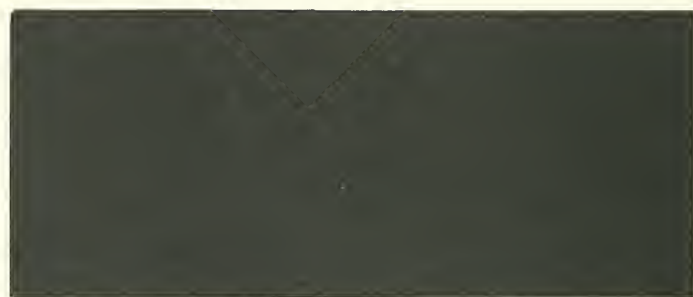WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

COMPUTER-ASSISTED CLINICAL DECISION-MAKING PROJECT

George

by G. Anthony Gorry

June 8, 1972

WP605-72

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

COMPUTER-ASSISTED CLINICAL DECISION-MAKING PROJECT

George

by G. Anthony Gorry

June 8, 1972

WP605-72

Computer-Assisted Clinical Decision-Making

During the past few years, I have been conducting some preliminary research into the use of computers to augment the decision making abilities of physicians. A few months ago, I decided to increase my efforts in this area. In what follows, I will outline briefly the motivation for this work, the results obtained to date, and the plan for pursuing this research in the future. The plan for future research is not completely clear, and it undoubtedly can be improved through the active interest and criticism of both computer people and physicians.

I. Motivation for the Research[1]

In the past few years, there have appeared in the literature many discussions of the use of computers in the health care system, and the way in which they might improve the efficiency of that system. Such improvements are seen as arising from a wide variety of computer-based activities such as scheduling of hospital admissions, control of laboratories, and the maintenance of medical records. Although these activities (and others as well) can undoubtedly benefit from the introduction of well-designed computer systems, more fundamental problems remain. There is an increasing shortage of physician manpower and a geographical maldistribution because new doctors are reluctant to practice in rural or depressed urban communities. Also these discussions fail to indicate how a high level of

---

[1]The bulk of this section is drawn from an article by Dr. William B. Schwartz "Medicine and the Computer: The Promise and Problems of Change." Dr. Schwartz is a collaborator in the current research.

of physician competence can be maintained in the face of a continued expansion of medical knowledge. The gap between what a doctor should know and what he can retain and utilize is continually widening.

As Schwartz has noted: "The computer thus remains (in the light of conventional projections) as an adjunct to the present [health care] system, serving a palliative function, but not really solving the major problems of that system."

There is, in fact, little reason to believe that any of the current proposals for solving these problems, technologic or other, will do more than mitigate their severity. Despite plans to reorganize patterns of medical care and efforts to enlarge medical school capacity and create new classes of "doctor's assistants," the physician shortage promises to be with us for decades and to pose a serious obstacle to health planning. The problem of maintaining and improving quality appears equally knotty since there is little indication that current programs in postgraduate education will be adequate to the challenge.

If conventional remedies will not meet the demands imposed by society's broad commitment to extensions of health care, it is clear that new, even heretical strategies must be devised. One intriguing possibility is to use the computer as an "intellectual" or "deductive" instrument—a consultant that is built into the very structure of the health care system and augments or replaces many of the traditional activities of the physician. Once can envision an ongoing dialogue between the physician and the computer with the latter continuously taking note of history, physicial findings,

laboratory data, and the like alerting the physician to probable diagnoses and suggesting possible courses of action. One may hope that the computer, well equipped to store a large volume of information and ingeniously programmed to assist in decision-making, will help free the physician to concentrate on the application of bedside skills, the management of the emotional aspects of disease, and the exercise of good judgment in the non-quantifiable aspects of clinical care.

The computer, used in this manner, might also open the way to quite different means of employing nonphysician manpower. Use of the computer as an intellectual resource in diagnosis and treatment could well be coupled to the development of new types of highly specialized allied health personnel who could perform functions of a scope well beyond those currently considered feasible for doctor's assistants. Computer-supported "health-care specialists," aided by a variety of automated devices for history taking, blood analysis and other procedures, and trained to perform a careful physicial examination, might take over a large segment of the responsibility for the delivery of primary medical care. Guided by the computer, constrained from exceeding his capacities by instructions built into the computer programs, and linked to regional consulting centers by appropriate display devices, the new breed of "health-care specialist" could make a major contribution to the resolution of the seemingly insoluble problem of maldistribution and shortage of physician manpower.

While such visions of the future are heady stuff, a serious consideration of the problems to be solved is immediately sobering. Clearly considerable

intellectual and technological resources must be marshalled and a long term research commitment must be made if such a scenario is to become a reality.

The work discussed in the next section constitutes a very modest investigation of one aspect of this problem. The focus of this work in on the decision making aspects of clinical medicine. The original hope was to embody in a computer program a normative procedure for diagnostic and therapeutic decision making that could be applied to a variety of clinical problems [2]. Although this work was only a partial success, it proved a very valuable exercise from which a number of new ideas have been gained. A discussion of these ideas will be postponed until the discussion of the new research plan. The discussion in the next section has not been "edited" to reflect the new (and hopefully better) view of the problem.

II. Review of Past Research

Introduction

The purpose of this section is to review my research on the use of a computer to solve diagnostic and treatment problems in medicine. A major result of this research has been the development of a computer program which is intended to serve as a consultant in a number of medical problem areas. Here the considerations which underlie the program are discussed. The basic functions of the program are outlined in a non-technical way, and an example of the use of the program is given. Then the results of the use of the program for several different medical problems are reviewed. Finally, an attempt is made to ascertain the potential of programs such as this in the delivery of appropriate medical care. Detailed reports on various aspects of this research are available in the literature ([1], [2], [3], [4]), and so the emphasis here will be on providing a general overview of the work and results obtained to date.

Modelling the Diagnostic and Treatment Problem

The use of digital computers in the selection of good diagnostic and treatment strategies has received increased attention in recent years. One reason for this interest is the general desire to improve the ability of the clinician to deal with the difficult problems which can arise in the management of a patient. A significant portion of the difficulty stems from the fact that the physician must sort out numerous possibilities and develop hypotheses about the state of health of the patient. The ability

of the computer to store extremely large amounts of data, to enumerate many possibilities, and to perform complex logical operations suggests its potential value in this problem solving process. Before a computer can be used to significant advantage in analyzing diagnostic and treatment strategies, however, precise procedures must be formulated for the means of inference required to deduce the clinical state of the patient from observed signs and symptoms, and a formalized capability must be developed for the prediction and assessment of possible therapeutic measures. In other words, the problem of performing diagnostic inference and weighting therapeutic strategies must be reduced to a problem of computation.

In order to better understand the requirements, a model of the diagnostic-treatment problem was formulated. The model is a mathematical one, but its principal characteristics can be discussed in terms of the way a physician deals with this problem. Although it should be noted that the model was not developed as a description of the way in which physicians operate. The purpose of the model is to permit the exploitation of the particular capabilities of a computer. Hence, in the next several paragraphs, when I am discussing the way in which a physician deals with the problem, I am using "physician" instead of "model" for convenience, and are not presenting a theory of human problem solving in the medical area. (The relationship of the model to the actual problem solving behavior of physicians in discussed in [6].)

In general, a doctor confronted with a potentially ill patient initially does not have sufficient information about the patient to decide on a diagnosis or on a therapeutic policy. The information he does have, however, in addition to his general medical knowledge and experience enables him to formulate some tentative hypotheses about the state of health of the patient. This opinion will exert a considerable effect on the strategy which the doctor will employ in dealing with the patient. For convenience, let us say that the options available to the physician are tests and treatments. By test we mean any means for obtaining additional information about the patient ranging from simple questions to laboratory procedures to certain surgical procedures. He employs those tests which he expects to provide results of significant value in improving his current view of the patient's problem. The term treatment will be used to refer to any means at the doctor's disposal to correct the health state of the patient. Treatments range from drugs to a variety of surgical procedures. The selection of an appropriate treatment for a given patient is strongly dependent on the correctness of the doctor's opinion about the patient's problem. The selection of the wrong treatment, for whatever reason, can have very serious consequences for the patient.

The value of the information obtained from a test is determined by the contribution which this information makes to improving the doctor's current view of the patient's problem and hence to reducing the risk of misdiagnosis with its associated cost. Hence the doctor is inclined to perform many tests. On the other hand, the tests available to him generally

are not without some cost in terms of patient discomfort, time of skilled persons, money, etc. Thus there is a conflicting tendency to hold the number of diagnostic tests to a minimum.

As is discussed in [3], the doctor resolves these conflicting tendencies by performing sequential diagnosis. At a particular point in time, given his current view of the patient's problem, he can evaluate the choices available to him. The basic choice is to employ a test to obtain more information, or to select a treatment in the hopes of curing the patient.

If he elects to cease testing and to make a diagnosis, the choice of a treatment implies a certain risk of mistreatment through a misdiagnosis. On the other hand, he can perform some test in the hopes of gaining additional information upon which to base his diagnosis and the resulting choice of treatment. In this case, he incurs the cost (in some terms) of the test selected. When the results of the test are known, and when they have been incorporated into his current view of the problem, he is faced with a decision problem of exactly the same form as the one which he has just solved. Thus a doctor can be thought of as solving a sequence of similar decision problems. At each stage of the process, he balances the cost of further testing against the expected reduction in the cost of treatment which the test results will permit. When, in the opinion of the physician, no tests possesses the property that is expected to reduce the risk of treatment by an amount which exceeds its cost, he will cease testing, make a diagnosis, and treat the patient. If the physician repeatedly updates his current view of the problem in keeping with the latest

information available to him, and if he has sufficient knowledge, he is able to develop effective diagnostic and therapeutic strategies.

Although this description of the manner in which a physician deals with diagnosis-treatment problems is simplified and somewhat artificial, it does emphasize the fundamental role that sequential decision making plays in the process. It seemed clear that it was necessary for a computer program to exploit an analogous capability (framed in terms suitable for a machine) in solving more general problems of this type.

The Development of the Computer Program

In this section, the basic components of a computer program to assess diagnostic and therapeutic strategies are discussed. These components directly reflect the view of the required problem solving process outlined in the preceding section. The discussion of the program in non-technical. Readers interested in the technical details are referred to [1] and [2].

The program has three basic components. The first is called the information structure, and it constitutes the medical experience of the program. By changing the information structure, one can convert the program for use in a new problem area. This is the only part of the program which changes from one application to the next.

In addition to the diseases, signs, symptoms, tests and treatments, the information structure contains two types of information: probabilities and utilities. The probabilities relate signs and symptoms to diseases. For example, one probability might be the conditional probability of red blood cell casts in the urine given that the patient has acute tubular

necrosis. The program's understanding of various diseases is entirely in terms of the conditional probabilities which relate the variety of signs and symptoms and treatment consequences to those diseases.

The utilities of the tests, treatments, and treatment consequences are thought of as the subjective preferences of an expert. The utility of a test reflects the pain associated with the test, the cost of the test, the time of a skilled person required for the test, the risk of the test to the patient, etc. Similar factors are reflected in the utilities of the treatments and the treatment consequences. Utility can be thought of as the common denominator in terms of which all these diverse factors are measured. Utility assessment will be considered in more detail later. Here we only note that if the program is to make comparisons of factors such as risk and cost, a common scale must be established for seemingly diverse outcomes.

The second major segment of the program is called the inference function. Basically the task of the inference function is to establish the diagnostic significance of a particular test result. In a typical situation, a doctor confronted with a particular diagnostic problem must interpret the available evidence (observed signs and symptoms, etc.) in terms of his general medical experience. In other words, he employs a method of deduction which can accommodate both his general understanding of diseases and the individual instance represented by the patient before him. The inference function of the program is the analogue of this capability in the physician. It uses probabilistic inference based on

Bayes rule [1],[4] to obtain a probability distribution for the likelihood of each disease given the evidence to date and general medical experience. The latter is incorporated in the information structure of the program. It is this probability distribution, then, which constitutes the current view taken by the program of the given problem. This view is updated whenever any new evidence is made available to the program. The updated probability distribution is one of the major factors which influence the strategy chosen by the program for dealing with a given patient.

The third component of the program is called the test/treatment selection function. Its purpose is to select at each stage in the problem solving process an appropriate test or treatment for use on the patient. By considering the probability distribution associated with the current view of the problem and the utilities of the various treatment consequences, this function can determine the best treatment to perform assuming that no further tests are to be used. The treatment chosen is the one which minimizes the expected risk, and it provides the standard used in evaluating the potential value of further testing.

In evaluating the potential usefulness of a particular test, the program considers the current view, the utilities of the various tests, and the likelihood of the possible test results. For each possible result of a test, the program can simulate the change in the current distribution which would occur if this result were obtained. The expected risk of treatment can be estimated for this new distribution. For each result of a test, the expected risk of treatment given the result is weighted by the

likelihood of obtaining that result, and the sum of these products is added to the utility of the test to obtain the overall measure. A schematic representation of the factors considered in evaluating a test is presented in Figure 1. By analyzing decision trees such as the one shown, the program attempts to select the best test or treatment at each stage of the analysis.

In Figure 2, an actual dialogue between a user and the program is presented.[1] The problem being considered is the diagnosis of a case of congenital heart disease. At the outset of the discussion, the program is essentially passive, simply accumulating whatever evidence the user offers and using the inference function to update its current view of the problem. When the user has completed his initial description of the patient, the test/treatment function is invoked to determine the best diagnosis-treatment policy. In this case, no treatments were considered, and the problem was only one of diagnosis. The example, however, does give a basic impression of the use of the program.

Experience with the Program

The program has performed extremely well in the medical problems to which it has been applied. It was used as a diagnostic program (no treatments were considered) in two areas: primary bone tumors and congenital heart disease. The bone tumor problem involved the classification of patients into one of nine disease groups based on radiographic evidence and other physical findings. Studies made with the program indicated that it performed

---

[1]The ELIZA program, developed by Professor Joseph Weizenbaum of MIT, was used to facilitate user input.
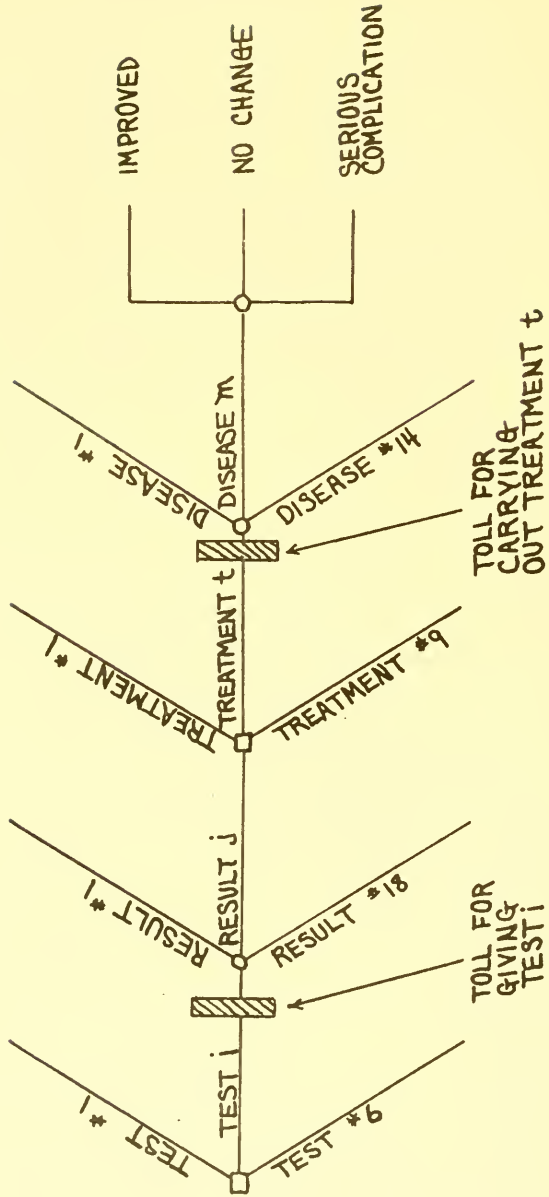
SAMPLE DECISION TREE

FIGURE 1

Figure 2

Example of a Session with the Program*

```
PLEASE TELL ME ABOUT THE PATIENT.
Well he is about 65 years old.
PLEASE CONTINUE.
No cyanosis.
O.K.  PLEASE GO ON.
What is your opinion of the 3 most probable diseases now.
     2 WITH PROBABILITY .68
    18 WITH PROBABILITY .10
    24 WITH PROBABILITY .04
ALL RIGHT.  GO ON PLEASE.
Name disease 4 for me.
ATRIAL SEPTAL DEFECT WITH PULMONARY HYPERTENSION.
PLEASE CONTINUE.
I meant the name of disease 24.
VENTRICULAR SEPTAL DEFECT WITH PULMONARY FLOW LE. 1.4 SYSTEMIC FLOW
O.K.  PLEASE GO ON.
What more do you need to know.
CHECK THE HILAR VESSELS.  ARE THEY INCREASED OR DECREASED.
That information is not available.
LOOK FOR A LEFT SECOND SYSTOLIC MURMUR.
What is your current view.
MY CURRENT VIEW IS
     DISEASE      PROB
         2         .68
        18         .10
        24         .04
        25         .04
        34         .03
         5         .03
        12         .02
NOW LET'S GO ON.
LOOK FOR A LEFT SECOND SYSTOLIC MURMUR.
Say there is a left second systolic murmur.
I SEE.
IS THE AXIS IN THE ECG RIGHT, LEFT, OR NORMAL.
What is the probability of atrial septal defect now.
IT IS 85.
ALL RIGHT.  GO ON PLEASE.
IS THE AXIS IN THE ECG RIGHT, LEFT, OR NORMAL.
Left.
ALL RIGHT.
CHECK FOR A LEFT FOURTH SYSTOLIC MURMUR.
What are the 3 most probable diseases now.
     5 WITH PROBABILITY .43
     2 WITH PROBABILITY .39
    18 WITH PROBABILITY .05
PLEASE CONTINUE.
CHECK FOR A LEFT FOURTH SYSTOLIC MURMUR.
What is the medical name of disease 5.
ATRIO-VENTRICULAR COMMUNIS.
O.K. PLEASE GO ON.
CHECK FOR A LEFT FOURTH SYSTOLIC MURMUR.
Say there is one.
O.K.
IS THERE AN ACCENTUATED OR DIMINISHED P2.
What is the probability of disease 5 now.
IT IS 71.
ALL RIGHT. GO ON PLEASE.
IS THERE AN ACCENTUATED OR DIMINISHED P2.
Not available.
I RECOMMEND A DIAGNOSIS OF DISEASE 5.
What are the 3 most probable diseases.
     5 WITH PROBABILITY .71
     2 WITH PROBABILITY .23
    24 WITH PROBABILITY .01
I'm finished.  Thanks for your help.
I HOPE THAT I HAVE BEEN OF SERVICE.  GOODBYE.
```

*User inputs in lower case; program responds in capitals.

at the level of an experienced radiologist given the same evidence [1].
The second medical problem considered was the diagnosis of congential heart
disease. Here there are 35 diseases in question. Again the results showed
that the program performed at the level of an expert [1], [2]. Both the
bone tumor problem and the congenital heart disease problem had certain
characteristics which made them unsatisfactory tests of the program's
capability, however. First, in neither area were the costs of the tests
sufficiently high to make the sequential aspect of the program particularly
important. Because tests were so cheap, it mattered little how many were
performed or in what order. Second, the treatment problem was not
considered, because doctors to help define these treatments and their
consequences were not available.

Because of these difficulties, a third medical area was considered, the
diagnosis and treatment of acute renal failure. The management of the acute
renal failure syndrome is an important medical problem. Although the
incidence of the problem is relatively small in the context of all disease
treatment problems, the potential risk to the patient's life is sufficiently
great to give the problem special significance. In dealing with this
problem, the physician must account for this risk factor as well as the
possibility that the tests which he could employ to gain further information
about the patient can contribute to the medical problem if improperly used.

The definition of the acute renal failure problem used in this study
included fifteen diseases. As in the previous problems of bone tumors and
congenital heart diesase, the information structure for the program included
the relevant probabilities. In this case, however, no attempt was made to

obtain these probabilities from an analysis of historical data. Rather, the opinion of an expert was used in establishing each probability. In addition, special attention was paid to the assessment of the required utilities. Again the opinion of an expert renal specialist served as the basis for these numbers.

The precise manner in which these judgments were obtained from the expert and the way in which they were converted to utilities is discussed in [5]. Here we want to briefly outline the procedure. The renal expert was given a series of hypothetical decision problems. Each problem required him to make a choice between a particular event for certain (such as curing the patient by performing a certain operation) and accepting a chance in a lottery. If he chose the lottery, a given event would be chosen for him with probability "p," and some other event would be chosen with probability "1-p." Before making his choice, the expert is told exactly what the two events in the lottery are and what the value of "p" is. With the theory discussed in [5], a series of these decision problems can be used to establish the utilities of tests, treatments, and consequences required by the program.

With the information structure for the renal failure problem developed in this way, the program duplicated the diagnostic-treatment decisions of expert renal specialists in over 90 percent of the cases tested. Furthermore, when the information structures from two experts were used, the program agreed more closely with the expert whose judgments it was using than did the other expert.

III. <u>Plan for Further Research</u>

To provide a context for a discussion of my plan for further research in this area, I want to offer a criticism of the work to date. Without going into detail, let me say that the evaluations of the program were strongly biased in favor of the program. The number of diseases, their rigid definitions, and the types of tests and treatments used all combined to make simple search an effective strategy. Thus the program did quite well compared to the experts, but the method it employed differed from the ones they used. Although I cannot characterize precisely the methods used by the experts, it is clear that these methods can accomodate the greater complexity of real clinical situations. The potential usefulness of search as the primary decision procedure for the program, however, is open to question. In this regard, it is instructive to consider some of the failures of the program in the experiments described above.

One such case was a patient with acute glomerulonephritis (AGN), a common cause of acute renal failure. Patients with AGN seldom have severe hypertension, but the patient presented to the physicians and the program did. The program obtained the correct diagnosis, but the treatment it recommended differed from that proposed by the doctors. Although both the physicians and the program chose the same treatment for AGN, the physicians recognized the need to deal with the patient's hypertension and hence recommended a second treatment as well.

Clearly, the program could be modified to check for this problem and to make the appropriate decisions. The same could be done for several

other problems of this type which were identified. Similar modifications would be required to obtain the appropriate interpretation of certain signs and symptoms. For example, hematuria (red blood cells in the urine) is an important diagnostic finding in acute renal failure. On the other hand, a patient with an indwelling catheter will generally have hematuria regardless of his intrinsic disease. Hence the interpretation of this finding should reflect this fact. Again either the program or the data it uses must be changed.

Although these particular problems could easily be solved within the context of the existing problem, they raise an important question. How many such "minor modifications" will be required for the program to have practical use in the clinical management of acute renal failure?

For a period of several months, I have investigated the amount and type of knowledge possessed by two acknowledged renal experts. Although much more work needs to be done, I can offer certain tentative conclusions. These conclusions provide motivation for a change of direction in this research.

1) Although detailed knowledge of physiology and pathophysiology is sometimes useful in clinical decision making, gross knowledge of this kind coupled with a large number of experiential facts and mini-decision procedures forms the primary basis of clinical judgment in renal disease.

2) The knowledge used by the experts is both factual and procedural. Their experience has provided them with a rich repetoire of ideas of

the form "if x is present and y is absent, then a good trial hypothesis is D." Such rules allow them to focus their attention on relatively few diagnoses or treatments. Of course these rules are heuristics, but many of them are of considerable value in dealing with their decision-making problems. By remembering large numbers of such patterns or rules, they avoid search to a large extent.

3) This experiential knowledge is not framed in deterministic terms, but is associated with various degrees of certainty.

4) The renal experts can specify only part of this knowledge a priori. A large part of this knowledge can be elicited only in response to apparent misconceptions on my part (or as embodied in the program).

5) Although there are very many "pieces" of knowledge involved, these experts seem able to state them clearly when the occasion arises.

The physicians I have been working with are acknowledged experts in renal disease, and their performance in this field far surpasses that of a very large fraction of the doctors who treat patients with this problem.[1] It is important,then, to get as much of their knowledge as possible in distributable form (i.e. a program).

The original program was based on a particular normative view of clinical decision making. The judgments of experts could be added only to the extent that these judgments could be expressed as simple probabilistic relationships or as utilities. Procedural knowledge was added through

---

[1] This is not a condemnation of the latter group. It is simply a reflection of the fact that most people with kidney disease do not have access to the experts and resources of a major teaching hospital.

reprogramming. Thus the addition of knowledge was either implicit (setting probabilities or utilities to cause the program to arrive at a conclusion that a physician could obtain more directly) or laborious (reprogramming). Unfortunately, I am convinced that for the foreseeable future, the desire to add knowledge will be great, and an attempt to maintain the program (perhaps for its simple, aesthetic appeal) will prove frustrating at best.

Although this discussion has been brief, it indicates the general tenor of the problems I foresee with the approach I had been using. Decision analysis is a useful tool when the problem has been reduced to a small, well-defined action selection one. It cannot be the sole basis of a program to generally assist clinicians in an area such as renal disease.

A New Program for Renal Disease

Several months ago, I began the development of a prototype program for use in the problem of acute renal disease. This program is currently in a most rudimentary form. Therefore I will be discussing here, not an existing program as much as some goals toward which I am working. My short term goal is to produce a version of this prototype which can be used by renal specialists in an informal way as a means to assess the potential of the ideas on which it is based.

Recent developments by people in the Artificial Intelligence laboratory at M.I.T. have opened the way for the exploration of new approaches to computer assimilation of knowledge. The developments comprise both a way of looking at the problem of machine knowledge and some very high level programming systems [7, 8]. The prototype system incorporates

some of these new ideas, and as a result is better able to accept experiential knowledge directly from the user. The details of the new program are beyond the scope of this paper (and many change significantly over time). Here, I will restrict myself to the conceptual framework within which this program is being built.

A simple language has been implemented to permit renal experts to give advice to the program regarding facts or ways to proceed in a particular circumstance. Examples of such statements are the following:

1) In acute glomerulouephritis, if hematuria is gross then red blood cell casts are very likely.

and

2) If protehuria is heavy and hematuria is gross and red blood cell casts are present and diagnosis is acute renal failure then diagnosis of glomerulitis is very likely.

The basic functions of the program and: 1) to accept such statements; 2) to note appropriate associations among various statements; and 3) to use to statements deductively when appropriate to draw conclusions about diagnosis or management.

It must be emphasized that the new program is very primitive as yet. The new technology mentioned above has greatly facilitated its development, however, and it seems likely that a much improved program can be implemented. The real question is whether sufficient improvement can be realized to make the program useful. At present, I cannot answer this question, but I can indicate the chief problems areas to be explored.

Problems for Investigation

1. Concept Identification

I intend to continue to try to identify the important concepts in renal disease. By this, I mean the identification of the central, problem-specific ideas in terms of which the experts organize their knowledge. One example is the concept of renal function. There are several approaches to inferring renal function and assessing whether it is stable or changing. This determination is very important in diagnosis and in choosing management strategies. From the experts, it is possible to obtain the procedure by which they infer a value for renal function. Further many statements about the interpretation of changes in renal function can be made. To capture the knowledge embodied in these statements, some computer realization of the concept of renal function must be developed.

Already it is clear that there are many such concepts. I will be trying to identify the most important ones and to develop reasonable ways to represent them in the program. Needless to say, a major question will be how many such concepts are required in the program, and the complexity of their realization. One possibility is that the number is so large as to be impossible to deal with at present. Another is that the individual concepts are based on an implicit assumption of enormous knowledge about the world.

I believe that the number of important concepts indeed is large, but not beyond our capabilities. For example, a very large portion of the basic knowledge about kidney disease is contained in one book (admittedly a large one). Further the expert clinicians believe that big chunks of that book are unnecessary for the support of clinical activities.

The issue of how much common sense is assumed in these concepts is also important. On the one hand, it could be argued that to understand these concepts, a program must understand a tremendous amount about the world. On the other hand, the relatively precise language of medicine may be the key here. The program may know many facts about streptococcal infection and its role in acute renal failure without understanding the concept of germs, etc. The physician using the program may have little need to ask the program for the latter. More generally, he will have considerable knowledge organized in terms of fairly well-defined words and phrases. The knowledge of the program can be expressed in these terms to assist him. More detailed knowledge on the part of the program may be unnecessary.

Already it is clear that there are many such concepts, but not all are of great importance. I will be trying to identify the most important ones and to develop reasonable ways to represent them in a program.

2. Language Development

Because I believe that the continual addition of knowledge is critical, I will be working on the development of a language within which experts can express this knowledge to the program. An understanding of the important concepts in renal disease, of course, is a prerequisite for the design of such a language. In general terms, what I am seeking is an automatic programming capability so experts can "program" the machine directly. At present, I can envision three languages involved in this process.

First, at the lowest level there will be the computer language in which the concepts are realized. At a higher level will be a language in which statements concerning these concepts are made without explicit recognition

of the details of the lower level realization. Such a language may well be an extension of the simple "IF-THEN" type language already implemented. By maintaining this separation, the problems arising from changes in the particular realization of the concepts in the machine may be lessened.

The third level language will be English. I am hoping to use Winograd's program [8] to translate statements made by the experts (in a subset of English) into the intermediate language mentioned. The second level language can be viewed as a cannonical representation of the subset of English which can be accepted. Such a translation will require an interaction with both the lower level languages, but I can say little in detail about this process. I do believe, however, that whatever the realization, language will be critical if the knowledge of experts is to be captured. Also I believe that they must be given some form of English for input and inquiry. Hence the tasks of concept identification and language development will have highest priority.

One question is worth raising here, although at present I do not know the answer. This question concerns the necessity for English. With experts dedicated to the project as the sole source of knowledge input, there might be little need for English; they could be taught to use the second level language. On the other hand, if interaction with other clinicians proves to be important (and I believe it will) then English may be very important. The question of how much is to be gained from English is one that will be considered carefully.

3. Explanation

The other side of the coin is explanation. If experts are to use and improve the program directly, then it must be able to explain the reasons for its actions. Furthermore, this explanation must be in terms the physicians can understand. The steps in a deduction and the facts employed must be identified for the expert so that he can correct one or more of them if necessary. As a corollary, the user must be able to easily find out what the program knows about a particular subject.

A Comment on Goals

The original aim of this research was to produce a decision-making program. Although this is still the long term goal, I believe the time to achieve this goal is sufficiently long to require the establishment of some short term goals. Presently I consider a reasonable (but somewhat vague) goal to be the construction of a program which can accept knowledge and answer simple requests for parts of that knowledge. Because there will be many cases where the program will lack knowledge relevant to a particular clinical situation, it should not make pronouncements but rather suggestions of things to consider and the assumptions on which its suggestions are based.

REFERENCES

[1]   G. Anthony Gorry, "A System for Computer-Aided Diagnosis,"
       Project MAC Technical Report, MAC-44, MIT, 1967.

[2]   _____, "Strategies for Computer-Aided Diagnosis,"
       Mathematical Biosciences, 2 (1968), 293-318.

[3]   G. Anthony Gorry and G. Octo Barnett, "Sequential Diagnosis by
       Computer," The Journal of the American Medical Association,
       205 (1968), 849-854.

[4]   _____, "Experience with a Model of Sequential Diagnosis,"
       Computers and Biomedical Research, 1 (1968), 490-507.

[5]   Norman E. Betaque and G. Anthony Gorry, "Automating Judgmental
       Decision Making for a Serious Medical Problem," Management
       Science, 17 (1971), B421-B434.

[6]   G. Anthony Gorry, "Modelling the Diagnostic Process," Journal
       of Medical Education, 45 (1970), 293-302.

[7]   G. J. Sussman, T. Winograd, and E. Charmiak, "MICRO-PLANNER
       Reference Manual," Memo No. 203A, The Artificial Intelligence
       Laboratory, M.I.T., Cambridge, Mass.

[8]   T. Winograd, "Computer Program for Understanding Natural Language...,"
       Project MAC Technical Report, TR-84, M.I.T., 1971.