Report 83-23
Stanford -- KSL



Computer-Assisted Decision Making in Medicine.

John C. Kunz, Edward H. Shortliffe, Bruce G. Buchanan, Edward A. Feigenbaum,

Apr 1983

Heuristic Programming Project
Report No. HPP-83-23

April 1983

# Computer-Assisted Decision Making in Medicine

John C. Kunz
Edward H. Shortliffe
Bruce G. Buchanan
Edward A. Feigenbaum

Heuristic Programming Project
Department of Computer Science & Medicine
Stanford University
Stanford, CA 94305

Reprinted with permission from the Journal of Philosophy and Medicine, April 1983, K. Schaffner, Ed.

Memo HPP-83-23

# COMPUTER-ASSISTED DECISION MAKING IN MEDICINE

John C. Kunz*

Edward H. Shortliffe**

Bruce G. Buchanan

Edward A. Feigenbaum

Heuristic Programming Project
Departments of Computer Science and Medicine
Stanford University
Stanford, Ca. 94305

Submitted for publication in the Journal of Philosophy and Medicine,
K. Schaffner, Editor

April 1983

*Current address: IntelliGenetics
124 University Avenue
Palo Alto, Ca. 94301

**Dr. Shortliffe is a Henry J. Kaiser Family Foundation Faculty Scholar in General Internal Medicine and recipient of research career development award LM00048 from the National Library of Medicine.

## Abstract

This article reviews the strengths and limitations of five major paradigms of medical computer-assisted decision making (CADM): 1) clinical algorithms, 2) statistical analysis of collections of patient data, 3) mathematical models of physical processes, 4) decision analysis, and 5) symbolic reasoning or artificial intelligence (AI). No one technique is best for all applications, and there is recent promising work which combines two or more established techniques. We emphasize both the inherent power of symbolic reasoning and the promise of artificial intelligence and the other techniques to complement each other.

# 1 INTRODUCTION

During the past decade, there has been a dramatic increase in both the capabilities of computer systems and in physicians' interest in their potential medical uses. Attitudes about their routine clinical use has evolved from early skepticism to curiosity and, increasingly, to to the presumption that all physicians will in time incorporate computer-based tools into their daily practices. A statement in a major medical journal provides evidence for this trend: "... computers may never replace physicians, but they can and will assist their decisions" [59].

A variety of techniques have been applied to computer-assisted decision making (CADM), accounting for at least 1000 references in the clinical and computing literature [51]. While progress has been made toward providing effective computer support for CADM in medicine, major research challenges remain. This paper updates an earlier review of the relevant methodologies [43], cites recent progress, and identifies some of the continuing research issues. The earlier article discussed important issues that account for both the multiplicity of approaches to the problem and the limited clinical success of most of the systems developed to date. It also contained an extensive bibliography which is revised in this article.

This article focuses on the representation and utilization of knowledge in computer programs, as developed within the field of applied AI. AI research is concerned with symbolic reasoning. Applied to practical problems, these AI techniques have been used to develop programs, termed "expert systems", which attempt to solve problems which are normally thought to require human specialists for their solution. Expert systems can make decisions about problems such as diagnosis and therapy, and they are designed to serve as consultants for human decision-makers. A recent Science article discusses research in expert systems [11].

We discuss some inadequacies of data-intensive techniques which have led to recent interest in symbolic reasoning approaches. In addition, we review some research projects which have combined two or more CADM techniques, discuss the way in which different techniques complement each other, and suggest that AI offers an accommodating paradigm under which other appropriate techniques can be merged.

## 1.1 Perspectives for comparing CADM systems

The models on which computer systems base their advice can be compared along three dimensions:

(1) inputs (i.e., the patient-specific data they use): Complexity of CADM system input can range from minimal data from a single patient to large amounts of data from many patients.

(2) processing (i.e., their use of medical and information processing knowledge): Processing knowledge can be very domain-specific, such as knowledge about managing a particular kind of patient in a local setting, or it can be domain-independent, such as knowledge about how to use statistical methods to assess the likelihood of a patient having a particular disease, given statistics on disease incidence in a relevant patient population. Some CADM systems use small amounts of domain-independent knowledge, while others use large amounts of domain-specific knowledge.

(3) outputs (i.e. both the form and content of their specific advice on a case): CADM system output may range from the quantitative value of a single parameter to presentation of a ranked list of multiple possible diagnoses or therapy plans, with explanations of conclusions.

When based on relatively limited amounts of domain-independent knowledge, decision-making systems typically use rich bases of domain-specific data. For example, statistical methods are often used to analyze the contents of large data banks in order to describe a class of patients or to make a prognosis about an individual patient. The output of such systems is typically a score or a quantitative value for a statistical probability. In contrast, some knowledge is relatively specific and applicable to a limited domain, such as the heuristic knowledge of an expert diagnostician. When based on large amounts of domain-specific knowledge, powerful decision-making systems typically use very small amounts of data. AI techniques have been used to make inferences based on large amounts of domain-specific knowledge; these systems can sometimes provide diagnoses and explanations of their decisions.

If there is a general chronology to the field over the last 30 years, it is that there has been progressive trend from techniques based on general knowledge, such as statistical methods, toward emphasis on techniques based on domain-specific symbolic knowledge, such as diagnostic inference rules. This trend has resulted in an emphasis on systems which interpret and explain the clinical significance of their findings, rather than simply produce another number for the user to interpret.

We include with domain knowledge a category of "judgmental knowledge" which reflects experience and opinions of experts regarding issues for which formal data may be fragmentary or nonexistent. Since many decisions made in clinical medicine depend upon this kind of judgmental expertise, it is not surprising that investigators should begin to look for ways to capture and utilize the knowledge of experts in decision making programs. Symbolic processing allows explicit representation and manipulation of heuristic judgmental knowledge. Thus, symbolic processing approaches allow exploiting judgmental knowledge acquired from experts. These systems can describe their reasoning process. By explaining its line of reasoning, a system provides credibility for decisions which are correct, and it helps the user to identify and to modify decisions which are incorrect.

Knowledge-based symbolic reasoning systems derive their power from making inferences which are consistent with, but not explicitly described in the knowledge base. Inference is the process of concluding a fact which was not explicitly described. Making inferences is a key issue in knowledge-based programming. Inferences can be made given appropriate knowledge, such as a diagnostic rule. Simple data, however, such as simple records of observations, provide no inherent basis for making inferences. Knowledge and data are closely related since inferences are made using both general knowledge of how to interpret data and specific data.

Approaches to symbolic reasoning characterize the distinctions between AI programs and programs using conventional calculations. For example, symbolic processing programs solve problems by pursuing a line of reasoning: the individual inference steps and the whole chain of reasoning may also form the basis for explanations of decisions. A major concern in expert systems is clear separation of the medical knowledge in a program from the inference mechanism that applies that knowledge to individual cases. One goal of this paper is to identify, in the strengths and weaknesses of earlier work, those issues which have motivated several current research groups to investigate the knowledge engineering approach to the development of CADM systems.

## 2 Clinical Algorithms and Automation

### 2.1 Overview

Clinical algorithms, or protocols, are structured decision making flowcharts to which a diagnostician or therapist can refer when deciding how to manage a patient with a specific clinical problem [41]. Like recipes, algorithms are simple, short, and relatively easy to use with limited prior training. Clinical algorithms take small amounts of input: a few selected observations of an individual patient. Their processing is simple flow-chart logic, based on highly distilled representation of the clinical knowledge of a domain. Their output is a recommended action: an additional test, a diagnosis or a suggested therapy. Algorithms provide no explanations per se, but the sequence of decisions used in arriving at a conclusion allows a user to understand the overall logic, if not tha basis for the individual steps.

Typically, clinical algorithms have been designed by expert physicians. Goldman recently reported on a large study which used statistical analysis of the contents of a data bank to generate an algorithm for managing patients who complain of chest pain [18]. Wirtschafter reports improved adherence to a protocol and fewer complications when physicians used an algorithm for chemotherapy treatment of Hodgkin's Disease [57]. Reporting that physician assistant users are more likely to follow the suggestions of a computer-based algorithm than those of a paper algorithm, Cannon suggests that use of paper algorithms may alienate users, while the use of computer-based algorithms computer elicits better compliance [7]. McDonald [29] hypothesizes that physicians have an acute problem with information overload, and that using algorithms can help to reduce the effects of this problem. He reports significant improvement in specific indices of the quality of care when physicians used algorithms. Further, he reports that the improvements tend to disappear quickly when physicians stop using the algorithms, indicating that the improvements are due to use of the algorithm and not due to educational effect.

### 2.2 Discussion of the Methodology

Although clinical algorithms are among the most widespread and accepted of the decision aids described in this article, the simplicity of their logic precludes effective use of the technique in most medical domains. Decision points in the algorithms are generally binary (i.e., a given sign or symptom is present or absent), and there tend to be many circumstances that can arise for which the user is advised to consult the supervising physician (or specialist). Thus the complex decision tasks are left to experts, and there is generally no formal algorithm for managing the case from that point on. It is precisely the simplicity of the algorithmic logic, and the supervising expert "escape valve", which has permitted many algorithms to be represented on one or two sheets of paper and has obviated the need for direct computer use in many applications. However, the methodology provides a limited basis for extension to the complex decision tasks to be discussed in the following sections.

## 3 Statistical Analysis of Collections of Patient Data

### 3.1 Overview

Making decisions based on analysis of patient records has been a major research concern since the earliest days of medical computing. The input to such statistically-based CADM systems includes either selective or exhaustive data for an individual patient. Input may also include data from other patients in a population to which the individual belongs. Population data may be used either in raw or summarized forms. In comparison with the other techniques discussed in this paper, statistical CADM systems typically use relatively large amounts of input data. These systems use different statistical analysis methods, including discriminant analysis, multi-attribute statistical pattern recognition [34], and Bayesian analysis [10]. These processing methods are uniformly domain-independent, and their output is usually numeric: a probability or a statistical index, possibly including a statistical measure of confidence. Because these systems contain no explicit knowledge about medicine, hey cannot explain the basis for their conclusions, either in terms of clinical knowledge or in terms of pathophysiology. The investigator typically selects a group of patients who belong to a common population. The statistical properties of the population are then characterized. In the individual case, the inference is made that if a patient is a member of a population whose properties are known, then the patient has the properties of the population, with some statistical measure of confidence.

Bayesian analysis is used to aid diagnosis by calculating the probability of a disease using a formula which explicitly considers both the a priori probability of the disease and the conditional probabilities relating the observations to the diseases in which they may occur. The appeal of Bayes' Theorem is clear: it offers an exact method for computing the probability of a disease based on observations and on data regarding the frequency with which these observations are known to occur for a specified set of diseases. In several domains the technique has been shown to be exceedingly accurate. For example, deDombal reports that a computer system, using Bayesian statistics, outperformed physicians in diagnosing causes of abdominal pain (e.g., appendicitis, diverticulitis, perforated ulcer, etc.) [10]. The proportion of correct diagnoses in their study was 91% for the computer, and it varied from 42% to 81% for different physician groups. In addition, the fraction of surgeries for non-diseased appendices dropped from 25% to 7%.

Among the most commonly recognized problems with the Bayesian approach is the large amount of data required to determine all the conditional probabilities needed in the rigorous application of the formula. A variety of additional assumptions must be made. For example: (i) the diseases under consideration are assumed to be mutually exclusive and exhaustive (i.e., the patient is assumed to have exactly one of the diseases considered), (2) in order to make the number of required conditional probabilities manageably small, the assumption is usually made that conditional probabilities are independent [47], and (3) the incidence of the symptoms of a disease is assumed to be static (i.e., the model generally does not allow for changes in disease patterns over time).

### 3.2 Discussion of the Methodologies

Statistical analysis systems offer powerful capabilities to individual clinical decision makers. Furthermore, medical computing researchers recognize the potential value of large databanks in supporting many of the other decision making approaches we discuss in subsequent sections. There are important additional issues regarding analysis of statistical data:

(1) A continuing issue is recognizing the context in which analysis of data is appropriate. Statistical analyses -- whether based on analysis of confidence limits, Bayesian analysis, multivariate analysis or some other technique -- depend ultimately on classifying an individual patient record as an instance of data drawn from a well characterized population. Diagnosis or prognosis is based on the inference that the properties of the patient are those of the population, with some statistical measure of confidence. However, knowledge of statistical techniques provides no basis for assessing the clinical importance of the differences between a patient record and the summary characteristics of a population. For example, continuing problems with use of pure statistical methods include deciding whether the presence of an arbitrary second disease or the recent use of a drug should change the interpretation that a patient's data belong to some population. A recent editorial discusses an extreme but important case in which the issue of analysis of context is a crucial issue in interpreting the results of a statistical analysis [33]. The issue is complicated by the need to change context, for example in making a local decision that some area of an X-ray is inside or outside the heart, and then making a global decision about whether heart volume is abnormal for the particular patient.

(2) Data acquisition remains a major problem. Many systems have avoided direct physician-computer interaction but have then been faced with the expense and errors of transcription.

(3) Analysis of data can be complicated by missing values that frequently occur, outlying values, and poor reproducibility of data across time and among physicians. Unfortunately, neither statistical nor data analysis methods provides a basis for acquiring, representing and manipulating the domain-specific knowledge of experts about how to handle problems in data.

(4) There is no formal correlation between the way expert physicians approach patient management decisions and the way that statistical programs arrive at recommendations. Psychological studies have now shown repeatedly that human clinical decision-making depends upon skilled use of large amounts of symbolic knowledge of medicine [13], [12]. Lacking expert clinical knowledge, statistical analysis systems have a severely limited ability to explain either the clinical significance or the clinical or the physiological basis for their results.

Thus, knowledge of medicine will be the most natural basis for explanation of the clinical significance and decisions of clinical medicine which is necessary to explain the clinical significance and the clinical or physiological basis for their results.

Lacking this correlation,

# 4 Mathematical Models of Physical Processes

## 4.1 Overview

Pathophysiologic processes can be well-described by mathematical formulae in a limited number of clinical problem areas. Selected data from a single patient is the typical input to a CADM system based on a mathematical model. Large amounts of raw data are often summarized using mathematical techniques: computing an average is a simple example. The processing is mathematical, using either direct analytical or simulation solution methods. The model itself is usually a mathematical representation of a general law of physics as it may apply in a particular situation.

Programs to assist with digitalis dosing have progressed to the inclusion of broader medical knowledge over the last ten years. The earliest work was Jelliffe's [24] and was based upon his considerable experience studying the pharmacokinetics of the cardiac glycosides. His computer program used mathematical formulations based on parameters such as therapeutic goals (e.g., desired predicted blood levels), body weight, renal function, and route of administration. In one study he showed that computer recommendations reduced the frequency of adverse digitalis reactions from 35% to 12% [25]. Later, another group revised the Jelliffe model to permit a feedback loop in which the digitalis blood levels obtained with initial doses of the drug were considered in subsequent therapy recommendations [40]. More recently, a third group in Boston, noting the insensitivity of the first two approaches to the kinds of nonnumeric observations that experts tend to use in modifying digitalis therapy, augmented the pharmacokinetic model with a patient-specific model of clinical status [21]. Running their system in a monitoring mode, in parallel with actual clinical practice on a cardiology service, they found that each patient in the trial in whom toxicity developed had received more digitalis than would have been recommended by their program.

## 4.2 Discussion of the Methodology

A general problem with mathematical approaches is that they depend on appropriate selection of the model to use, on assignment of default values to parameters whose measured values are unknown or unknowable, and ultimately on interpreting the clinical significance of the results of an analysis. Interpretation of results must be done both in terms of properties of the data used in the analysis itself and in terms of the general situation, including factors such as history, current therapies and clinical course of the patient. As with clinical algorithms and statistical approaches to CADM, the techniques of mathematical analysis themselves provide no basis for representing and manipulating the knowledge necessary to set up the model and to interpret its results.

# 5 Decision Theoretical Approaches

## 5.1 Overview

Bayes' Theorem is only one of several techniques used in the larger field of decision analysis, and there has recently been increasing interest in the ways in which decision theory might be applied to medicine and adapted for automation. Several excellent medically-oriented reviews of the field are available in textbooks [55], journal articles (such as an entire issue of the New England Journal of Medicine which was devoted to papers on this methodology [23]), and recently in a series of "Clinical Decision Conferences" in the journal Medical Decision Making [22]. The latter series describes cases referred to the Clinical Decision Analysis Service at the New England Medical Center in Boston.

In general terms, decision analysis can be seen as any attempt to consider values associated with choices, as well as probabilities, in order to analyze the processes by which decisions are made or should be made. Schwartz identifies the calculation of "expected value" as central to formal decision analysis [39]. Ginsberg contrasts medical classification problems (e.g., diagnosis) with broader decision problems (e.g., "What should I do for this patient?"), and asserts that most important medical decisions fall in the latter category and are best approached through decision analysis [17]. The following topics are among the central issues in the field.

(1) Decision Trees. The decision making process can be seen as a sequence of steps in which the clinician selects a "path" through a network of plausible events and actions. Nodes in this tree-shaped network are of two kinds: decision nodes, where the clinician must choose from a set of actions, and chance nodes, where the outcome is not directly controlled by the clinician but is a probabilistic response of the patient to some action taken.

(2) Expected Values. In actual practice physicians make sequential decisions based on more than the probabilities associated with the chance node that follows. For example,

the best possible outcome is not necessarily sought if the costs associated with that path far outweigh those along alternate pathways.

(3) Eliciting Values. Obtaining from physicians and patients the cost and values they associate with various tests and outcomes can be a formidable problem, particularly since formal analysis requires expressing the various costs in standardized units. Lotteries are a useful new approach [36]. Inferences regarding values can be made by identifying the odds, in a hypothetical lottery, at which the physician or patient is indifferent regarding taking a course of action with certain outcome and betting on a course with preferable outcome but with a finite chance of significant negative costs if the "bet" is lost.

(4) Test Evaluation. Since the tests which lie at decision nodes are central to clinical decision analysis, it is crucial to know the predictive value of tests that are available. This leads to consideration of test sensitivity, specificity, receiver operator characteristic curves, and sensitivity analysis. Such issues are discussed in [26], [30] and [55].

## 5.2 Discussion of the Methodology

In 1973, Gorry reported results of using a program based on decision analysis for management of acute renal failure [19]. Gorry's program used subjective probability estimates from experts, rather than probabilities derived from formal statistical studies. The excellent performance of his program suggests the importance of the clinical analysis that underlies the decision theoretical approach. The reasoning steps in managing clinical cases have been dissected in such detail that small errors in the probability estimates are apparently much less important than they were shown to be for deDombal's purely Bayesian approach [28]. Gorry suggests this may be simply because the decisions made by the program are based on the combination of large aggregates of such numbers, but this argument should apply equally for a Bayesian system. It seems to us more likely that distillation of the clinical domain in a formal decision tree gives the program so much more knowledge of the clinical problem that the quantitative details become somewhat less critical to overall system operation. The explicit decision network is a powerful knowledge structure; the "knowledge" in deDombal's system lies in conditional probabilities alone and there is no larger scheme to override the propagation of error as these probabilities are mathematically manipulated by the Bayesian routines.

The decision theory approach is not without problems, however. Perhaps the most difficult problem is that considerable expertise in both decision analysis and in the domain area are needed to construct and use the decision tree itself. From the perspective of the user, the difficulty is similar to the issue of context, (discussed in Section 4.2). Knowledge of the domain-independent method of decision analysis does not, in itself, provide a basis for representing a decision tree which pertains to a particular issue. In addition, if a generic tree is produced by one investigator, in one institution, then there is inevitable and appropriate question about its use in a different institution which uses somewhat different procedures and for which sensitivity and specificity data may be somewhat different. Overlapping or coincident diseases are also not well-managed, unless specifically included in the analysis, and the Bayesian foundation for many of the calculations still assumes mutually exclusive and exhaustive disease categories. Problems of symptom conditional dependence still remain, and there has been no easy way to include knowledge regarding the time course of diseases. Gorry points out that his program was also incapable of recognizing circumstances in which two or more actions should be carried out concurrently. Furthermore decision theory per se does not provide the kind of focusing mechanisms that clinicians tend to use when they assume an initial diagnostic hypothesis in dealing with a patient and discard it only if subsequent data make that hypothesis no longer tenable. An additional difficult problem is assigning numerical values (e.g., dollars) to a human life or a day of health, etc. Some critics feel this is a major limitation to the methodology [52].

## 6 Symbolic Reasoning Approaches

## 6.1 Overview

In the early 1970's researchers at several institutions simultaneously began to investigate the potential applications of artificial intelligence (AI) to clinical decision making. A recent series of books [3], [9] reviews the general field of AI, and two other recent books review the applications of AI in medicine [48], [8]. The term "artificial intelligence" is generally accepted to include those computer applications which perform symbolic reasoning rather than numeric calculation. Examples include programs that reason about mineral exploration, organic chemistry, or molecular biology; programs that converse in English and understand spoken sentences; and programs that generate theories from observations.

The basis for symbolic reasoning programs, and the source of their power, is the use of qualitative judgments - codified as "heuristics", or experiential "good guess" knowledge. In contrast, the basis for numerical calculation programs is the use of analytical equations or statistical techniques. Heuristics encode knowledge about a problem area, such as medical diagnosis. They are used to focus the attention of the reasoning program on parts of the problem that seem most critical, and on parts of the knowledge base that seem most relevant. They directly address the issue of context discussed above in Section 4.2 by both focusing the reasoning process and by deleting irrelevant items from consideration. The result is that these programs pursue a line of reasoning as opposed to following a sequence of steps in a calculation.

Gorry's landmark 1973 paper first critically analyzed conventional approaches to computer-based clinical decision making and outlined his motivation for turning to newer symbolic techniques [20]. He used his acute renal failure program [19] to illustrate the problems that can arise when decision analysis is used alone. In particular, he analyzed some of the cases on which the renal failure program had failed but the physicians considering the cases had performed well. His conclusions from these observations include the following four points:

(1) Clinical judgment is based less on detailed knowledge of pathophysiology than it is on gross chunks of knowledge and detailed experience from which rules of thumb are derived.

(2) Clinicians know facts, of course, but their knowledge is also largely judgmental. The rules they learn allow them to focus attention and generate hypotheses quickly. Such heuristics permit them to avoid detailed search through the entire problem space.

(3) Clinicians recognize levels of belief or certainty associated with many of the rules they use, but they do not routinely quantitate or utilize these certainty concepts in any formal statistical manner.

(4) It is easier for experts to state their rules in response to perceived misconceptions in others than it is for them to generate such decision criteria a priori.

Gorry's group at MIT and Tufts developed new approaches to examining the renal failure problem in light of these observations [35].

Due to the limitations of the older data-intensive techniques, it was perhaps inevitable that some medical researchers would turn to the AI field for new methodologies. Major AI research areas include knowledge representation, heuristic search, natural language understanding and generation, and models of thought processes -- all topics clearly pertinent to the problems we have been discussing. Furthermore, AI researchers were beginning to look for applications to which they could apply some of the techniques they had developed in theoretical domains. This community of researchers has grown in recent years, and a complete issue of the journal Artificial Intelligence was devoted to applications of AI to biology, medicine and chemistry [45](1).

Using symbolic reasoning techniques, CADM systems have been developed during the last decade in several areas of scientific, technical and commercial interest. These programs are called expert systems because they exploit knowledge acquired from experts and because they address problems normally thought to require human specialists for their solution. Expert systems research is reviewed in a recent article [11], and one of the first such expert systems, MYCIN, is discussed in some detail. The clinical goal of MYCIN is to make

-----------------------------------------

(1) Many of the AI-based systems described in this this article were developed on the SUMEX-AIM computing resource, a nationally shared system devoted entirely to applications of AI to the biomedical sciences. SUMEX-AIM computers are physically located at Stanford and Rutgers Universities, but they are used by researchers nationwide via connections to a computer network. The resource is funded by the Division of Research Resources, Biotechnology Resources Program, National Institutes of Health.

diagnoses and to recommend therapy suggestions for patients suffering from infectious diseases. The system goals and methodologies are discussed in detail in [42]. MYCIN illustrates the approach to design of expert systems. The input to MYCIN includes clinical facts, such as information about the patient history and laboratory test results. MYCIN makes decisions by processing rules which describe knowledge about clinical medicine. Rules are acquired from infectious disease experts, and one of their key features is that both the using clinician and the computer can understand their meaning. The output of MYCIN is a suggested diagnosis and a therapy suggestion. Following user request, MYCIN identifies the rules it used in making its decisions and thereby explains the clinical basis of its suggestions to the user.

Pople, Myers and Miller have developed a system called INTERNIST-1 whose goal is to assist with diagnosis and test selection for all diseases in internal medicine [32]. The program utilizes a hierarchic disease categorization, an ad hoc scoring system for quantifying symptom-disease relationships, plus some clever heuristics for focusing attention, discriminating between competing hypotheses, and diagnosing concurrent diseases. The system designers characterize INTERNIST-1 as a research project, not yet ready for routine clinical use. For them, the most troublesome system limitations are its inability to reason anatomically or temporally and its limited representation and use of pathophysiological knowledge. Pople discusses these problems and approaches for their solution in detail [37]. In addition, the system does not yet have the well developed human engineering features which are necessary for routine use.

Miller's recent article summarizes the results of using INTERNIST-1 to analyze 19 trial cases reported in New England Journal of Medicine clinicopathological exercises [32]. These cases have the advantage of being well documented and diagnostically very challenging, but they are not representative of the majority of cases seen in routine medical practice because of their complexity. In this complex set of cases, 43 anatomically verified diagnoses were present. INTERNIST-1 correctly made 25 diagnoses; the attending physicians made 28, and the expert physicians made 35. No other diagnostic computer system could analyze such a large number of complex cases drawn from the large field of internal medicine. In his editorial, Barnett suggests that the contribution of

INTERNIST-1 should not be measured simply by its limitations or failures [2]. Rather, experiments such as INTERNIST provide insight into both the human diagnostic process and the way that computers can help the practitioner. Barnett continues: "The issue is whether such artificial intelligence models can reach conclusions similar to those of a competent clinician and can then justify those conclusions in a rational and clinically acceptable fashion."

Most AI programs have been developed on relatively large mainframe computers. Recently, some systems have been developed to work on personal "work stations". For example, the ONCOCIN program is now used routinely in the Stanford oncology outpatient clinic [44]. By concentrating on state-of-the-art AI techniques developed during the 70's, selecting both hardware and software to provide a user-friendly interface, and carefully selecting a task domain (managing patients enrolled in chemotherapy protocols), ONCOCIN has succeeded in appealing to user physicians, and it plays a gradually increasing role in their clinical operations. In addition, in recent years, several medical diagnosis programs based on AI techniques have been made to work on small and relatively inexpensive computers. For example, a small expert system designed for operation on a chip has been reported [56]. The PUFF program, which interprets pulmonary function test results, runs routinely on a personal computer in a hospital laboratory [1]. Taking measured data directly from a laboratory computer, PUFF identifies the presence and severity of one or more of obstructive lung disease, restrictive disease and diffusion defect. The program produces a report for each patient which explains the reasons for its diagnosis, considers the measured and potential effect of drug therapy, and discusses the potential of further testing when appropriate to confirm a diagnosis. The PUFF system report is reviewed by a physician, edited if necessary (about 90% of the reports need no editing), signed, and entered into the patient record. The program has interpreted over 4000 cases since it first went into routine use in 1978 at Pacific Medical Center in San Francisco.

### 6.2 Discussion of the Methodology

Symbolic reasoning techniques differ from the other methodologies mentioned in this article in that the computer techniques themselves are as yet experimental and rapidly changing. Whereas the computations involved in Bayes' Theorem, for example, involve straightforward application of computing techniques already well-developed, basic researchers in computer science continue to develop new methodologies for knowledge representation, language understanding, heuristic search, and the other symbolic reasoning problems we have mentioned. Thus the AI programs tend to be developed in highly experimental environments where short term practical results are often unlikely to be found. Prototype programs typically require large amounts of space and tend to be slow, particularly in time-sharing environments. As has been true for most of the methodologies discussed, AI researchers have still not developed adequate methods for handling concurrent diseases, assessing the time course of disease, nor acquiring adequate structured knowledge from experts. Furthermore, inexact reasoning techniques tend to be developed and justified largely on intuitive grounds.

Szolovits and Pauker reviewed some applications of AI to medicine and have attempted to weigh the successes of this young field against the very real problems that lie ahead [47]. They identify several serious deficiencies of current systems. For example, termination criteria are still poorly understood. In addition, although some of these programs now cleverly mimic some of the reasoning styles observed in experts [13],[12], it is less clear how to keep the systems from abandoning one hypothesis and turning to another one as soon as new information suggests another possibility. Programs that operate this way appear to digress from one topic to another -- a characteristic that decidedly alienates a user regardless of the validity of the final diagnosis or advice.

Despite these significant limitations, AI techniques do provide a way to respond to many of Gorry's observations regarding the inadequacies of prior methodologies [20]. There are now several programs responsive to his criticisms. INTERNIST-I has received close attention within both the computer science and the medical communities because of its ability to make multiple and complex diagnoses within a very broad problem domain. The psychological research into the nature of diagnosis shows the complexity of the process followed by humans and the richness of the knowledge used in following that process. INTERNIST-I is the best documented example of the power of the symbolic processing methods to make diagnostic decisions at a high level of competence and to justify its conclusions to human decision-makers.

## 7 Integration of multiple techniques

Sections 3-7 of this paper discussed the approach, the strengths and the limitations of five basic methods for computer-based medical decision making. No one technique is best for all applications, and some techniques have strengths in the same areas that others have limitations. Thus, it is not surprising that some investigators have attempted to develop systems which combine multiple techniques.

Integrated systems show power and promise. Warner's group in Salt Lake City, for example, has developed a computer-based system which interfaces with a hospital information system and implements a variety of decision-making approaches [53]. The HELP system includes many simple clinical algorithms, a large data bank, some mathematical models, and programs which use Bayesian and decision-analysis methods to generate advice, to identify expected values of decisions, and to perform sensitivity analyses. The system is in routine use in the Latter Day Saints Hospital, Salt Lake City, and is remarkable both for its successful focus on decision-making and for its integration of CADM techniques into the routine of patient care.

Bleich's system combines techniques of algorithms and mathematical modeling [4], while the Digitalis advisor, previously mentioned in section 5.1, was designed to recommend digitalis dose rates by combining a symbolic model of therapeutic goals with a mathematical model of digitalis pharmacokinetics [21]. This program uses heuristics to develop a description, called a "patient specific model", of an individual patient. The associated mathematical model uses the results of heuristic analysis as input, and it produces quantitative results which are added to the patient-specific model.

Another example, the AI/MM system, represents and manipulates a physiological model of fluid and electrolyte balance. The program includes symbolic knowledge of anatomy and physiological function, and it uses simple mathematical descriptions of the relevant first principles of physics and physiology [27]. Thus, AI/MM incorporates Yamamoto's suggestion that mathematical models might be incorporated into AI-based systems [58]. Using its symbolic knowledge base, AI/MM analyzes the physiological model to infer the behavior of the renal system. It explains the basis for behavior in terms either of first principles or of high-level heuristics. This symbolic analysis is used to motivate, to initialize and to evaluate the of quantitative analyses; numerical results in turn are considered in symbolic analysis of physiological behavior.

Blum has done work on the discovery of causal relations in a clinical data base [5]. Blum's program (RX) integrates techniques of artificial intelligence and statistics. The principal objective of RX is to automate the process of generating and testing hypotheses about causal relations in large data banks. RX uses knowledge about clinical medicine for generating and testing hypotheses about new causal relations. It uses AI techniques to identify relevant knowledge and to choose how to employ it in particular situations, and it uses statistical techniques to test the validity of hypothesized relations. Although RX is not a CADM system per se, one can envision CADM systems which use the insight that AI and statistics can be combined.

## 7.1 Discussion

The systems discussed in this section suggest some of the potential benefits of integrating different decision making techniques into a unified system. This integration may have important benefits for the clinical user. The HELP system, for example, is an integrated system. It provides current patient data and data about similar kinds of patients who have been seen in the institution. Accordingly, it can suggest common implications of the patient data, as conceived by thoughtful experts. It can use algorithmic, mathematical, statistical, or decision-analysis techniques for interpreting the data of an individual patient. In addition, an integrated system offers great power to the expert decision-maker. For example, the HELP system provides a collection of tools which the decision-maker can use, in the context of problem-oriented patient management, in specifying criteria for giving routine diagnostic and therapeutic advice and for giving warnings about abnormal situations. Thus, as the HELP system illustrates, an integrated system can allow the clinical user to obtain both data and its analysis by any of the several specialized scientific disciplines which might be appropriate in a specific case.

## 8 Conclusions

We have claimed that there are two recurring issues to confront in considering the field of computer-based clinical decision making:

(1) How can we design systems that help physicians to reach better, more reliable decisions in a broad range of applications, and

(2) How can we more effectively encourage the use of such systems by physicians or other intended users?

We shall summarize by reviewing these points separately.

### Performance Issues

Central to assuring a program's adequate performance is a matching of the most appropriate technique with the problem domain. We have seen that the structured logic of clinical algorithms can be effectively applied to triage functions and other primary care problems, but they would be less naturally matched with complex tasks such as the diagnosis and management of acute renal failure. Good statistical data may support an effective Bayesian program in settings where diagnostic categories are small in number, non-overlapping, and well-defined, but the lack of higher level domain knowledge limits the effectiveness of the Bayesian approach in more complex patient management or diagnostic environments. A mathematical approach may support decision making in certain well-described fields in which observations are typically quantified, and related by functional expressions. These examples, and others, demonstrate the the need for thoughtful consideration of the technique most appropriate for managing a clinical problem. In general the simplest effective methodology is to be preferred, but acceptability issues must also be considered as discussed below.

Recent work suggests the potential of combining two or more decision-making techniques in a CADM system. The systems discussed in Section 8 suggest that there is more power in a combination of two or more processing techniques than there is in any of the techniques individually used in these systems. However, whether this power will prove to have practical clinical importance will be determined only if such integrated systems are further developed and are used in different settings. In addition to the appeal of combining techniques in general, there seems to be a special opportunity for combining AI with the quantitative techniques of statistical and mathematical analysis and, potentially, with decision analysis. AI offers a formalism for representing and manipulating the heuristic knowledge of expert human problem solvers. This knowledge can also be used to address the issue of recognizing or establishing an appropriate context for using some processing technique, and heuristic knowledge can be used to assess the qualitative clinical significance of computed results. AI/MM uses AI and simple mathematical models; it uses quantitative and qualitative problem solving to support each other, and it shows the epistemological basis of inference steps, distinguishing between inferences based on physical law and on heuristic associations. RX uses AI and statistics; it generates and interprets statistical results in a more powerful way than would be possible with either technique alone.

It is always appropriate to ask whether computer-based approaches are needed at all for a given decision making task. Some clinical algorithm developers, for example, have discarded the machine [41], [50], and Schwartz et al. pointed out that a useful decision analysis can often be accomplished in a qualitative manner using paper and pencil [39]. Similarly, it is appropriate to ask whether potentially complex and expensive computer processing is needed for a task. Sometimes, as in the work of Mesel, the needs of users can be met effectively by a simple computer-based approach which can be developed relatively quickly and run economically [31].

Finally, it is important to consider the extent to which a program's "understanding" of its task domain will heighten its performance, particularly in settings where knowledge of the field tends to be highly judgmental and poorly quantified. We use the term "understanding" here to refer to the degree of judgmental or structural knowledge (as opposed to data) that is contained in the program. Analyses of human clinical decision making [13], [12] suggest that as decisions move from simple to complex, a physician's reasoning style becomes less algorithmic and more heuristic, with qualitative judgmental knowledge and the conditions for invoking it coming increasingly into play. accordingly, it is likely that medical computing researchers will have to become "knowledge engineers" as they look for effective ways to match the knowledge structures that they use to the complexity of the tasks they are undertaking.

### Acceptability Issues

A recurring observation as one reviews the literature of computer-based medical decision making is that essentially no system has been effectively utilized outside of the environment in which it was developed, even when its performance has been shown to be excellent! This suggests that it may be an error to concentrate our research effort primarily on improving the decision making performance of computers when there is evidently much more required before these systems will have substantial clinical impact. An editorial critique suggests that issues of terminal design and interface between user and computer must be addressed effectively before computers substantially effect clinical medicine [16]. A heightened awareness of "human engineering" issues among medical computing researchers may also help to improve acceptance of computers by physicians. Fox reviewed this field in detail [15]. The issues range from the mechanics of interaction at a computer terminal to program characteristics designed to make the system appear as a tool for the physician rather than a dogmatic advice-giving machine.

Adequate attention must also be given to the severe time constraints perceived by physicians. Ideally they would like programs to take no more time than they currently spend when accomplishing the same task on their own. Time and schedule pressures are similarly likely to explain the greater resistance to automation among interns and residents than among medical students or practicing physicians in Startsman's study [46].

Finally it must be noted that acceptability issues should generally be considered from the outset in a system's design because they may dictate the choice of methodology as much as the task domain itself does. The role of formal knowledge structures to facilitate explanation capabilities, for example, may argue in favor of using symbolic reasoning techniques even when a somewhat less complex methodology might have been adequate for the decision task.

It may be tempting to conclude that the biases of medical personnel against computers are so strong that systems will inevitably be rejected, regardless of performance, and in fact there are some data to support this view [46]. However, recent studies have shown relatively positive attitudes of medical personnel toward computers [49], and there are now examples of applications in which initial resistance to automated techniques has gradually been overcome through the incorporation of adequate system benefits [54].

In spite of their goals and achievements, the systems described in this paper have not been used outside of the institutions in which they were developed. Commercial neglect of CADM systems probably stems from at least three issues: lack of demand from the hospital and office-practice communities for computer-assisted decision-making; unwillingness of vendors to develop products for which there is limited apparent demand; and lack of convincing evidence that computer decision-support provides more effective or more economical care.

Perhaps one of the most revealing lessons about acceptability is an observation regarding the system of Mesel et al. [31]. The physicians in Mesel's study accepted the guidance of protocols for the management of chemotherapy in their cancer patients. It is likely that the key to acceptance in this instance is the fact that these physicians had previously had no choice but to refer their patients with cancer to a tertiary care center where all complex chemotherapy was administered. The introduction of the protocols permitted these physicians to undertake tasks that they had previously been unable to do, and it simultaneously allowed maintenance of close doctor-patient relationships and helped the patients avoid frequent long trips to the center. The motivation for the physician to use the system is clear in this case. It is reminiscent of Rosati's assertion that physicians will first welcome computer decision aids when they become aware that colleagues who are using the machine have a clear advantage in their practice [38].

In summary, the trend toward increased use of medical knowledge in clinical decision programs has been in response to desires for both improved performance and improved acceptance of such systems. As greater experience is gained with these techniques, and as they become better known throughout the medical computing community, it is likely that we will see increasingly powerful unions between symbolic reasoning and the alternate methodologies we have discussed. One lesson to be drawn lies in the recognition that there is basic computer science research to be done in medical computing, and that the field is more than the application of established computing techniques in medical domains.

## References


1. Aikins, J.S., Kunz, J.C., Shortliffe, E.H., Fallat, R.J., "PUFF: an expert system for interpretation of pulmonary function data" Comp. Biomed. Res. (1983, in press).

2. Barnett, G.O. "The computer and clinical judgement." (editorial) N. Eng. J. Med. 307, pp. 493-494 (1982).

3. Barr, A., Feigenbaum, E.A. Handbook of Artificial Intelligence Vol. 1-2, Los Altos, CA: William Kaufman Inc., (1982).

4. Bleich, H.L. "Computer-based consultation: Electrolyte and acid-base disorders", Am. J. Med. 53, PP. 285-291 (1972).

5. Blum, R.L. "Discovery, confirmation, and incorporation of causal relationships from a large time-oriented clinical database: The RX project." Comp. Biomed. Res. 15, 2 pp. 164-187 (1982).

6. Buchanan, B.G. and Feigenbaum, E.A. "Dendral and Meta-Dendral: their applications dimension." Artificial Intelligence 11, pp. 5-24 (1978).

7. Cannon, S.R., Gardner, R.M. "Experience with a computerized interactive protocol system using HELP." Comp. Biomed. Res. 13, PP. 399-409 (1980).

8. Clancey, W.J. and Shortliffe, E.H. Readings in Medical Artificial Intelligence: the First Decade. Reading, MA: Addison-Wesley (1983).

9. Cohen, P.R., Feigenbaum, E.A. Handbook of Artificial Intelligence Vol. 3, Los Altos, CA: William Kaufman Inc., (1982).

10. deDombal, F.T., Leaper, D.J., Horrocks, J.C., et al. "Human and computer-aided diagnosis of abdominal pain: further report with emphasis on performance of clinicians." Brit. Med. J. 1, pp. 376-380 (1974).

11. Duda, R.O. and Shortliffe, E.H. "Expert Systems Research." Science 220, pp. 261-268 (1983).

12. Eddy, D.M., Clanton, C.H. "The Art of Diagnosis." New Eng. J. Med., 306, pp. 1263-1268 (1982).

13. Elstein, A.S., Shulman, L.S., and Sprafka, S.A. Medical Problem Solving: An Analysis of Clinical Reasoning. Cambridge, Ma.: Harvard Univ. Press, (1978).

14. Feigenbaum, E.A. "The Art of Artificial Intelligence: Themes and case studies of

knowledge engineering." AFIPS Conference Proc., NCC (1978). Vol. 47. Montvale, N.J.: AFIPS Press, (1978), p.227.

15. Fox, J. "Medical computing and the user." Int. J. Man-Machine Studies 9, pp. 669-686 (1977).

16. Friedman, R.B. and Gustafson, D.H. "Computers in clinical medicine: a critical review."(editorial) Comp. Biomed. Res. 8, pp. 199-204 (1977).

17. Ginsberg, A.S. "The diagnostic process viewed as a decision problem." In Computer Diagnosis and Diagnostic Methods, (J.A. Jacquez, ed.), Springfield, IL: Charles C. Thomas, (1972).

18. Goldman, L., Weinberg, M. et al. "A computer-derived protocol to aid in the diagnosis of emergency room patients with acute chest pain." New Eng. J. Med. 307, pp. 588-596 (1982).

19. Gorry, G.A., Kassirer, J.P., Essig, A., and Schwartz, W.B. "Decision analysis as the basis for computer-aided management of acute renal failure." Amer. J. Med 55, pp. 473-484 (1973).

20. Gorry, G.A. "Computer-assisted clinical decision making." Meth. Inform. Med. 12, pp. 45-51 (1973).

21. Gorry, G.A., Silverman, H., and Pauker, S.G. "Capturing clinical expertise: a computer program that considers clinical responses to digitalis." Amer. J. Med 64, pp. 452-460 (1978).

22. Gottleib, J.E., Pauker, S.G. "Whether or not to administer amphotericin to an immunosuppressed patient with hematologic malignancy and undiagnosed fever." Medical Decision Making 1 pp.75-93 (1981).

23. Inglefinger, F.J. "Decision in medicine" (editorial). N. Eng. J. Med. 293, pp. 254-255 (1975).

24. Jelliffe, R.W., Buell, J., Kalaba, R., et al. "A computer program for digitalis dosage regimens." Math. Biosci. 9, pp. 179-193 (1970).

25. Jelliffe, R.W., Buell, J., and Kalaba, R. "Reduction of digitalis toxicity by computer-assisted glycoside dosage regimens." Anns. Int. Med. 77, pp. 891-906 (1972).

26. Komaroff, A.L. "The variability and inaccuracy of medical data." Proc IEEE 67, pp. 1196-1207 (1979).

27. Kunz, John C. "Use of artificial intelligence and simple mathematics to analyze a physiological model", Ph.D. thesis, Stanford University (1983).

28. Leaper, D.J., Horrocks, J.C., Staniland, J.R., and deDombal, F.T. "Computer-assisted diagnosis of abdominal pain using 'estimates' provided by clinicians." Brit. Med. J. 4, pp. 350-354 (1972).

29. McDonald, C.J., Wilson, G.A., McCabe, G.P. "Physician response to computer reminders" J. Am. Med. Assoc. 244, pp. 1579-1581 (1980).

30. McNeil, B.J. and Adelstein, S.J. "Determining the value of diagnostic and screening tests." J. Nucl. Med. 17, pp. 439-448 (1977).

31. Mesel, E., Wirtschafter, D.D., Carpenter, J.T., et al. "Clinical algorithms for cancer chemotherapy - systems for community-based consultant-extenders and oncology centers." Meth. Inform. Med. 15, pp. 168-173 (1976).

32. Miller, R.A., Pople, H.E., and Myers, J.D. "INTERNIST-1, an experimental computer-based diagnostic consultant for general internal medicine." N. Eng. J. Med. 307, pp. 468-476 (1982).

33. Multiple Risk Factor Intervention Trial Research Group. "Multiple risk factor intervention trial" JAMA 248, pp. 1465-1477 (1982).

34. Patrick, E.A. "Pattern Recognition in Medicine," Systems, Man and Cybernetics Review, 6, pp. 4 (1977).

35. Pauker, S.G., Gorry, G.A., Kassirer, J.P., and Schwartz, W.B. "Towards the simulation of clinical cognition: taking a present illness by computer." Amer. J. Med. 60:981-996 (1976).

36. Pauker, S.P. and Pauker, S.G. "Prenatal diagnosis: a directive approach to genetic counseling using decision analysis." Yale J. Biol. Med. 50:275-289 (1977).

37. Pople, H. E. "Heuristic Methods for imposing structure on ill-structured problems: the structuring of medical diagnostics." in Artificial Intelligence in Medicine, edited by P. Szolovits, AAAS Symposium 51, Westview Press (1981).

38. Rosati, R.A., Wallace, A.G., and Stead, E.A. "The way of the future." Arch. Intern. Med. 131, pp. 285-287 (1973).

39. Schwartz, W.B., Gorry, G.A., Kassirer, J.P., and Essig, A. "Decision analysis and clinical judgment." Amer. J. Med 55, pp. 459-472 (1973).

40. Sheiner, L.B., Halkin, H., Peck, C., et al. "Improved computer-assisted digoxin therapy." Anns. Int. Med. 82, pp. 619-627 (1975).

41. Sherman, H., Reiffen, B., and Komaroff, A.L. "Ambulatory care systems." In Problem-Directed and Medical Information Systems (M.F. Driggs, ed.), New York: Intercontinental Medical Book Corporation, 1973, pp. 143-171.

42. Shortliffe, E.H. Computer-Based Medical Consultations: MYCIN, New York: Elsevier/North Holland, (1976).

43. Shortliffe, E.H., Buchanan, B.G., Feigenbaum, E.A. "Knowledge engineering for medical decision making: a review of computer-based clinical decision aids." Proc. of IEEE 67, pp. 1207-1224 (1979).

44. Shortliffe, E.H., Scott, A.C., Bischoff, M.B., et al "ONCOCIN: an expert system for oncology protocol management." Proc. of Am. Assoc. Art. Intelligence pp. 876-880 (1981).

45. Sridharan, N.S. Guest editorial. Artificial Intelligence 11, pp. 1-4 (1978).

46. Startsman, T.S., and Robinson, R.E. "The attitudes of medical and paramedical personnel towards computers." Comp. Biomed. Res. 5, pp. 218-227 (1972).

47. Szolovits, P. and Pauker, S.G. "Categorical and probabilistic reasoning in medical diagnosis." Artificial Intelligence 11, pp. 115-144 (1978).

48. Szolovits, P., Artificial Intelligence in Medicine. AAAS Selected Symposium 51, Westview Press (1982).

49. Teach, R.L. and Shortliffe, E.H. "An analysis of physician attitudes regarding computer-based consultation systems." Comp. Biomed. Res. 14, 6 pp. 542-558 (1981).

50. Vickery, D.M. "Computer support of paramedical personnel: the question of quality control." Medinfo 74, pp. 281-287, North-Holland: Amsterdam, 1974.

51. Wagner, G., Tautu, P., and Wolber, U. "Problems of medical diagnosis: a bibliography." Meth. Info. Med. 17, pp. 55-74 (1978).

52. Warner, H.R. Knowledge sectors for logical processing of patient data in the HELP system." Proc. of 2nd. Ann. Symp. on Computer Applications in Medical Care, IEEE, Wash. D.C.,(1978), pp. 401-404.

53. Warner, H.R. Computer-Assisted Medical Decision Making. New York: Academic Press, (1980).

54. Watson, R.J. "Medical staff response to a medical information system with direct physician-computer interface." MEDINFO 74, p. 299-302, Amsterdam: North-Holland Publishing Company, (1974).

55. Weinstein, M.C., Fineberg, H.V., et al. Clinical Decision Analysis. Philadelphia: W.B. Saunders Co., (1980).

56. Weiss, S.M. and Kulikowski, C.A. "Developing microprocessor-based expert models for instrument interpretation." Proceedings IJCAI-81,pp. 853-855, Vancouver, B.C., (1981).

57. Wirtschafter, D.D., Scalise, M., Henke, C., Gams, R.A. "Do information systems improve the quality of clinical research? Results of a randomized trial in a cooperative multi-institutional cancer group." Comput. Biomed. Res. 14, PP. 78-90 (1981).

58. Yamamoto, W.S., and Walton, E.S. "On the evolution of the physiological model" in Annual Review of Biophysics and Bioengineering, 4, p. 81-102 (1975).

59. Ziporyn, T., "Computer-assisted medical decision-making: interest growing." J. Am. Med. Assoc. 248, pp. 913-918 (1982).

60. Zoltie, N., Horrocks, J.C., and deDombal, F.T. "Computer-assisted diagnosis of dyspepsia - report on transferability of a system, with emphasis on early diagnosis of gastric cancer." Meth. Inform. Med. 16, pp. 89-92 (1977).