Artificial Intelligence in Medicine

Edited by

Peter Szolovits

ROUTLEDGE

Artificial Intelligence in Medicine

AAAS Selected Symposia Series

Artificial Intelligence in Medicine

Edited by Peter Szolovits

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published 1982 by Westview Press

Published 2018 by Routledge
52 Vanderbilt Avenue, New York, NY 10017
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

Library of Congress Catalog Card Number 82- 60046

ISBN 13: 978-0-367-02221-1 (hbk)

About the Book

This book introduces the field of artificial intelligence in medicine, a new research area that combines sophisticated representational and computing techniques with the insights of expert physicians to produce tools for improving health care. An introductory chapter describes the historical and technical foundations of the work and provides an overview of the current state of the art and research directions. The authors then describe four prototype computer programs that tackle difficult clinical problems in a manner similar to that of an expert physician.

The programs presented are INTERNIST, a diagnostic aid that combines a large database of disease/manifestation associations with techniques for problem formulation; EXPERT and the Glaucoma Program which use physiological models for the diagnosis and treatment of eye disease; MYCIN, a rule-based program for diagnosis and therapy selection for infectious diseases; and the Digitalis Therapy Advisor, which aids the physician in prescribing the right dose of the drug digitalis and also explains its actions.

About the Series

The *AAAS Selected Symposia Series* was begun in 1977 to provide a means for more permanently recording and more widely disseminating some of the valuable material which is discussed at the AAAS National Meetings. The volumes in this *Series* are based on symposia held at the Meetings which address topics of current and continuing significance, both within and among the sciences, and in the areas in which science and technology impact on public policy. The *Series* format is designed to provide for rapid dissemination of information, so the papers are reproduced directly from the camera-copy submitted by the authors. The papers are organized and edited by the symposium arrangers who then become the editors of the various volumes. Most papers published in this *Series* are original contributions which have not been previously published, although in some cases additional papers from other sources have been added by an editor to provide a more comprehensive view of a particular topic. Symposia may be reports of new research or reviews of established work, particularly work of an interdisciplinary nature, since the AAAS Annual Meetings typically embrace the full range of the sciences and their societal implications.

WILLIAM D. CAREY
Executive Officer
American Association for
the Advancement of Science

Contents

About the Editor and Authors

Peter Szolovits *is an associate professor in the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. A specialist in artificial intelligence (AI) with an emphasis on medical applications, he is currently concerned with fundamental issues of representation and reasoning, including protocol analysis to discover how clinicians reason about probability and causality, and with programs which model human expert performance in some areas of medical diagnosis and care.*

Randall Davis *is an assistant professor in the Department of Electrical Engineering and Computer Science at Massachusetts Institute of Technology. His research interests are knowledge-based systems, knowledge representation, problem-solving methods, and use of artificial intelligence in medicine. He is the author of* Knowledge-based Systems in Artificial Intelligence *(with D. B. Lenat; McGraw-Hill, 1982).*

Casimir A. Kulikowski *is professor of computer science and associate director of the Laboratory for Computer Science Research at Rutgers University; since 1972 he has also been a senior investigator and associate director of the Rutgers Research Resource on Computers in Biomedicine. His research is in the fields of artificial intelligence and pattern recognition, with emphasis on expert systems and their applications, and he has directed several collaborative projects for the development of expert medical consultation systems.*

William J. Long *is on the research staff at the Laboratory for Computer Science, Massachusetts Institute of Technology. Currently he is working on development and evaluation of the Digitalis Therapy Advisor and on the design of AI programs for heart failure diagnosis and therapy, and he has written on the criteria for computer-generated therapy advice in a clinical domain.*

Ramesh S. Patil, *a research assistant at the Laboratory for Computer Science, Massachusetts Institute of Technology, has done research on the design of a program for the causal understanding and diagnosis of patient illness in electrolyte and acid-base disturbances.*

Harry E. Pople, Jr., *is an associate professor of business and co-director of the Decision Systems Laboratory at the University of Pittsburgh. A computer scientist, he has specialized in artificial intelligence with emphasis on clinical decision making, problem formulation, and problem solving.*

William B. Schwartz *is Vannevar Bush University Professor and professor of medicine at Tufts University and associate member of the Laboratory for Computer Science at Massachusetts Institute of Technology. A physician by training, he has published extensively on renal function and electrolyte and acid-base equilibria in humans. He is also currently involved in development of medical applications of computers, including their use for decision analysis and the simulation of clinical judgment.*

Sholom M. Weiss, *an associate research professor of computer science at Rutgers University, has worked on developing computer-based consultation systems in medicine. His current research interests are the development of general systems for designing expert models and the application of expert systems in medicine, laboratory instrumentation, and oil exploration.*

Preface

This book presents an overview of the major research efforts in the United States in the application of artificial intelligence techniques to medical decision making. It has grown out of a Symposium on Artificial Intelligence in Medicine presented at the 1979 Annual Meeting of the American Association for the Advancement of Science. The chapters in this volume are mainly revisions of presentations at the Symposium, augmented to bring the work more up to date and to address questions arising from the audience and discussions at the meeting.

The Artificial Intelligence in Medicine (AIM) field emerged in the early 1970's in response to several simultaneous needs, opportunities and interests. An increased demand for high-quality medical services coupled with the explosive growth of medical knowledge has led to the suggestion that computer programs could be used to assist physicians and other health care providers in discharging their clinical roles in diagnosis, therapy and prognosis. At the same time, computer science techniques, especially those of the artificial intelligence field, began to reach a maturity with which they could be applied to representing and reasoning about complex, "real world" problems like those arising in medicine. Investigators trained on both the computational and the medical side of these concerns began to develop mutual interests and approaches, and to form coherent collaborative research efforts which have now produced the programs reported on in this volume.

Modern medicine has become technically complex, the standards set for it are very high, conceptual and practical advances are rapid, yet the cognitive capabilities of physicians are pretty much fixed. As more and more data become available to the practicing doctor, as more becomes known about the processes of disease and possible interventions available to alter them, practitioners are called on to know more, to reason better, and to achieve better outcomes for their patients. But how can the physician learn, retain and apply this ever-expanding body of knowledge? Improvements in medical education are a constantly sought but limited means of meeting this challenge. The creation within our medical centers of specialty and sub-specialty services, with experts serving as shared consultants to the rest of the community and training programs to create more experts has been the most important institutional response. However, such a structure also creates major differences among users of the health care system. The general practitioner (or today's "family practice specialist") can no longer embody both a broad and deep expertise in all of medicine. To get the best of care, the patient may need the advice of several specialists over the course of time. Although major medical centers can and do provide such an adequate spectrum of services, those economically disadvantaged or remotely located cannot or do not receive such an intensive concentration of assistance. This problem is especially acute in underdeveloped countries, where even the supply of primary physicians, nurses and technicians is very limited.

If the expertise of consultants can be captured in the form of computer programs which provide advice to less-expert physicians or other health-care providers, then any practitioner could call on that expertise whenever a patient's case suggested the need for careful thought about some aspect of the illness or therapy. The continued astonishing increase in power and decrease in price of computers will put a large (by today's standards) computer within reach of every physician's desk within a few years. The opportunity is there to improve the health-care system by improving each physician's ability to utilize the best available knowledge and the best ways of analyzing medical problems, as encoded in easily-duplicated and updated computer programs.

Implementing this vision is by no means simple. Even technical perfection is no guarantee that a technology will be adopted and used in the ways its proponents envision. What social, ethical, legal, professional, or other reasons may arise to step in the way? And of course technical success is also still more a promise than a reality. The programs described here have made a good start, and have demonstrated that the vision is feasible. Several of them, with needed additional development work, could probably already serve to improve medical care in their areas of competence. The programs have also identified numerous rich problems for artificial intelligence research, and have put their developers at the forefront of research in AI applications.

In this collection of papers, the authors invite you, our readers, to appreciate the excitement of the challenges we face, the accomplishments we have achieved so far, and the progress we plan. The reports included here are from the major research groups in the U.S. No attempt has been made to be exhaustive in coverage, though the best-known programs in the field are represented. Fortunately, the related work of our colleagues and students at our own universities and at other domestic and foreign universities and research centers is too voluminous to include in this work—our field is growing, attracting brilliant young computer scientists and physicians, and finding institutional and financial support.

As editor, I would naturally like to thank all the contributors, whose preparation for the original Symposium and later preparation and revision of these chapters has undoubtedly seemed a never-ending task. I thank Dr. William B. Schwartz, my longtime collaborator, for talking me into editing this volume, and Ms. Joellen Fritsche, who gave editorial advice and bore patient witness for the AAAS during the long birth of this book. I thank my former secretary, Ms. Anne Schmitt, for help with typing and editing the manuscript and drawing illustrations, and my students for helping to read, analyze, and suggest corrections. My wife, Dianne Foster, has been kind enough to match my late nights of editing with almost equally late nights of her own legal studies. All the research reported on here has been touched by the supporting hand of the National Institutes of Health's Division of Research Resources and the National Library of Medicine, two institutions whose leaders have recognized the promise of the field and devoted themselves to its progress.

This book was prepared using computerized text editing, formatting and typesetting systems at MIT's Laboratory for Computer Science. The EMACS text editor, developed by Richard Stallman, has been an excellent tool supporting the incremental editing of countless drafts of the manuscripts. The book was formatted and electronically typeset by Donald Knuth's TEX. Camera-ready copy was produced by our Dover printer, created at Xerox Palo Alto Research Center. During development of the book, as well as in the course of our research, the authors were in contact by the electronic mail facilities of the ARPA network, which allowed us to shuttle manuscript fragments around the country and gave us virtually instant mail. Computer-aided preparation of a book, an experience likely to be shared by growing numbers of authors, is somewhat of a mixed blessing—the authors gain great control over the appearance of the final product, but at the cost of considerably more work that traditionally belongs to the publisher. After this description, it should be more than obvious that flaws in the book belong squarely to the editor.

P. Sz.

Cambridge, Mass.

Taylor & Francis
Taylor & Francis Group
http://taylorandfrancis.com

1. Artificial Intelligence and Medicine

Man strives to augment his abilities by building tools. From the invention of the club to lengthen his reach and strengthen his blow to the refinement of the electron microscope to sharpen his vision, tools have extended his ability to sense and to manipulate the world about him. Today we stand on the threshold of new technical developments which will augment man's reasoning; the computer and the programming methods being devised for it are the new tools to effect this change.

Medicine is a field in which such help is critically needed. Our increasing expectations of the highest quality health care and the rapid growth of ever more detailed medical knowledge leave the physician without adequate time to devote to each case and struggling to keep up with the newest developments in his field. For lack of time, most medical decisions must be based on rapid judgments of the case relying on the physician's unaided memory. Only in rare situations can a literature search or other extended investigation be undertaken to assure the doctor (and the patient) that the latest knowledge is brought to bear on any particular case. Continued training and recertification procedures encourage the physician to keep more of the relevant information constantly in mind, but fundamental limitations of human memory and recall coupled with the growth of knowledge assure that most of what is known cannot be known by most individuals. This is the opportunity for new computer tools: to help organize, store, and retrieve appropriate medical knowledge needed by the practitioner in dealing with each difficult case, and to suggest appropriate diagnostic, prognostic and therapeutic decisions and decision making techniques.

In a 1970 review article, Schwartz speaks of

> the possibility that the computer as an intellectual tool can reshape the present system of health care, fundamentally alter the role of the physician, and profoundly change the nature of medical manpower recruitment and medical education—in short, the possibility that the health-care system by the year 2000 will be basically different from what it is today. [18]

This research was supported (in part) by the National Institutes of Health Grant No. 1 P01 LM 03374 from the National Library of Medicine and Grant No. 1 P41 RR 01096 from the Division of Research Resources.

The key technical developments leading to this reshaping will

> almost certainly involve exploitation of the computer as an 'intellectual,' 'deductive' instrument—a consultant that is built into the very structure of the medical-care system and that augments or replaces many traditional activities of the physician. ... Indeed, it seems probable that in the not too distant future the physician and the computer will engage in frequent dialogue, the computer continuously taking note of history, physical findings, laboratory data, and the like, alerting the physician to the most probable diagnoses and suggesting the appropriate, safest course of action. [18]

This vision is only slowly coming to reality. The techniques needed to implement computer programs to achieve these goals are still elusive, and many other factors influence the acceptability of the programs.

This book is an introduction to the field of *Artificial Intelligence in Medicine,* (abbreviated AIM) which is now taking up the challenge of creating and distributing the tools mentioned above. This introductory chapter defines the problems addressed by the field, gives a short overview of other technical approaches to these problems, introduces some of the fundamental ideas of artificial intelligence, briefly describes the current state of the art of AIM, discusses its technical accomplishments and current problems, and looks at likely future developments. The other four chapters each describe one of the current AIM projects in some detail, pointing out not only the accomplishments of the programs built so far but also what we have learned in the process of creating them.

Definitions

What is "Artificial Intelligence in Medicine?" One introductory textbook defines artificial intelligence (called AI) this way:

> Artificial Intelligence is the study of ideas which enable computers to do the things that make people seem intelligent. ... The central goals of Artificial Intelligence are to make computers more useful and to understand the principles which make intelligence possible. [30]

This is a rather straightforward definition, but it embodies certain assumptions about the idea of intelligence and the relationship between human reasoning and computation which are, in some circles, quite controversial. The coupling of the study of how to make computers useful with the study of the principles which underlie human intelligence clearly implies that the researcher expects the two to be related. Indeed, in the newly-developing field of *cognitive science,* computer models of thought are explicitly used to describe human capabilities.

Historically, researchers in AI have had to defend this linkage against humanist attacks on the reduction of the human intellect to computational steps. The debate has sometimes been heated, as exemplified by the following quote from the introduction to an early collection of AI papers:

> *Is It Possible for Computing Machines to Think?*

> No—if one defines thinking as an activity peculiarly and exclusively *human*. Any such behavior in machines, therefore, would have to be called thinking-like behavior.
>
> No—if one postulates that there is something in the essence of thinking which is *inscrutable, mysterious, mystical*.
>
> Yes—if one admits that the question is to be answered by *experiment and observation*, comparing the behavior of the computer with that behavior of human beings to which the term "thinking" is generally applied.
>
> We regard the two negative views as unscientifically dogmatic. [5, p. 2]

An enlightening review of the history of AI and the bouts between its proponents and adversaries may be found in the recently published *Machines Who Think* [13].

AI in Medicine (AIM) is AI specialized to medical applications. Researchers in AIM need not engage in the controversy introduced above. Although we employ human-like reasoning methods in the programs we write, we may justify that choice either as a commitment to a human/computer equivalence sought by some or as a good engineering technique for capturing the best-understood source of existing expertise on medicine—the practice of human experts. Most researchers adopt the latter view.

The choice to model the behavior of a computer expert in medicine on the expertise of human consultants is by no means logically necessary. If we could understand the functioning in health and in disease of the human body in sufficient depth to model the detailed disease processes which disturb health, then, at least in principle, we could perform diagnosis by fitting our model to the actually observable characteristics of the patient at hand. Further, we could try out possible therapies on the model to select the optimum one to use on the patient. Unfortunately, although biomedical research strives for such a depth of understanding, it has not been achieved in virtually *any* area of medical practice. The AIM methodology does not dogmatically reject the use of non-human modes of expertise in the computer. Indeed, accurate computations of probabilities and solutions of simple differential equations—tasks at which human experts are rather poor without special training—play a role in some of our programs. Nevertheless, most of what we know about the practice of medicine we know from interrogating the best human practitioners; therefore, the techniques we tend to build into our programs mimic those used by our clinician informants.

Relying on the knowledge of human experts to build expert computer programs is actually helpful for several additional reasons: First, the decisions and recommendations of a program can be explained to its users and evaluators in terms which are familiar to the experts. Second, because we hope to duplicate the expertise of human specialists, we can measure the extent to which our goal is achieved by a direct comparison of the program's behavior to that of the experts. Finally, within the collaborative group of computer scientists and physicians engaged in AIM research, basing the logic of the programs on human models supports each of the three somewhat disparate goals that the researchers may hold:

- To develop expert computer programs for clinical use, making possible the inexpensive dissemination of the best medical expertise to geographical regions where that expertise is lacking, and making consultation help available to non-specialists who are not within easy reach of expert human consultants.

- To formalize medical expertise, to enable physicians to understand better what they know and to give them a systematic structure for teaching their expertise to medical students.

- To test AI theories in a "real world" domain and to use that domain to suggest novel problems for further AI research.

History

AIM is certainly not the first use of computers in medicine. Many of the administrative and financial record keeping needs of the hospital, health center, and even small group medical practice have been turned over to computer systems. Such use of computers differs little from similar applications in a wide range of businesses, and few technical developments have been motivated specifically by medical use of what could be called "business computing." Obviously, such use will continue to benefit from the increasing performance of general business-oriented systems; just as computer suppliers now aim for the small retail store as a possible market, they also envision the computerization of even individual doctors' offices, providing billing, scheduling, forms preparation, word processing, and other services.

It appears unlikely, however, that such business uses of computing in medical applications will fulfill the promise to "reshape" medicine. In a recent book on management decision support systems, McCosh and Scott Morton, writing about management information systems (MIS), note that

> despite the tremendous growth in computer-related activities, [MIS] has had little significant impact on management. The kinds of decisions and the ways in which they are made have been very little affected by computers over the last fifteen years. We believe that this can be traced in large part to the lack of proper perspective on the problems involved in augmenting the decision-making ability of management. [12, p. 3]

Similarly, much of the business computing in medicine impacts only on the periphery of the physician's task.

A second, currently much smaller use of computers in medicine is their application to the substance rather than the form of health care. If the computer is a useful manager of billing records, it should also maintain medical records, laboratory data, data from clinical trials, etc. And if the computer is useful to store data, it should also help to analyze, organize, and retrieve it.

Three main approaches to this second type of medical computing have so far been used: the clinical algorithm or flowchart, the matching of cases to large data bases of previous cases, and applications of decision theory. Each of these has had notable successes, but also a more limited applicability than its developers had hoped. All contribute to the development of the AI approaches described here. A good recent review of the state of the art of computer tools for medical decision making can be found in [19] and an accompanying argument for the AI orientation in [25].

Flowcharts

A *flowchart* is conceptually the simplest decision making tool. It encodes, in principle, the sequences of actions a good clinician would perform for any one of some population of patients. We may imagine, for example, recording all sequences of questions asked, answers given, procedures performed, laboratory analyses obtained and eventual diagnoses, treatments and outcomes for a number of patients who present at the emergency room with severe chest pain. If we observe enough patients and allow expert cardiologists to suggest an appropriate retrospective analysis of each case based on their excellent knowledge of the field, we may be able to identify a suitable sequence of actions to take under all possible circumstances. This approach has been successfully applied to the encoding of triage protocols for use by nurses [15], and has also formed the basis for several programs for patient interviewing [20]. A very large flowchart program has also been built for giving therapeutic advice in the acid/base area [1].

The principal deficiency of the flowchart as a general technique for encoding medical decision making knowledge is its lack of compactness and perspicuity. When used in a very large problem domain, the flowchart is likely to become huge, because the number of possible sequences of situations to be considered is enormous.* Furthermore, the flowchart does not include information about its own logical organization: each decision point appears to be independent of the others, no record exists of all logical places where each piece of information is used, and no discipline exists for systematic revision or updating of the program. Therefore, inconsistencies may easily arise due to incomplete updating of knowledge in only some of the appropriate places, the totality of knowledge of the flowchart is difficult to characterize, and the lack of any explicit underlying model makes justification of the program very difficult.

Data Bases

Large *data bases* of clinical histories of patients sharing a common presentation or disease are now being collected in several fields. The growth of data capture and storage facilities and their co-occurring decline in cost make attractive the accumulation of enormous numbers of cases, both for research and clinical uses. Today we are engaged in numerous long-term studies of the health effects of various substances, the eventual outcomes of competing methods of treatment, and the clinical development of diseases. Large databases on significant populations, concentrating on cardiovascular disease, arthritis, cancer and other major medical problems, are now being collected and used to clarify the true incidence of diseases, to identify demographic factors and to measure therapeutic efficacy of drugs and procedures [10, 17, 29].

---

*Its developer has kindly provided the author with a listing of a recent version of the Acid/Base flowchart program mentioned above [1]. The listing, in a variant of the MUMPS language, occupies close to 150 pages.

For clinical purposes, the typical use of large data bases is to select a set of previously known cases which are most similar to the case at hand by some statistical measures of similarity. Then, diagnostic, therapeutic and prognostic conclusions may be drawn by assuming that the current case is drawn from the same sample as members of that set and extrapolating the known outcomes of the past cases to the current one.

The use of collected past records either for research or clinical practice is clearly a data-intensive activity. To sift through the voluminous information at hand, to identify the important generalizations to be found among the thousands of detailed records and to select previous cases likely to shed light on the one under current consideration, numerous statistical techniques have been developed and applied. The literature of medical statistics is large, and will not be reviewed here; a good survey may be found in [26] and accompanying articles.

Although vast collections of data and processing techniques for them are an important advance, the application of this methodology to all of medicine appears unlikely for several reasons. Firstly, the collection and maintenance of the data in a consistent and accessible form is very costly and extremely time consuming. Old data are difficult to reconcile with the new, because continual refinements introduced as medical knowledge deepens introduce distinctions which were absent in previously-collected cases. Rare disorders may be infrequent enough that an insufficient number are seen within the "catchment basin" of any data collection scheme to provide adequate data. Historical and regional differences in nomenclature and interpretation can make the reconciliation of separately-collected data virtually impossible. Thus, it appears likely that only the more common and severe disorders generate enough interest, resources, and clinical cases to make the collection of data practical. Secondly, and equally importantly, the existing expertise of physicians is a highly valuable body of knowledge which cannot be recovered from just the processing of many cases by statistical techniques. A method of diagnosis, prognosis or therapy which relies on the projection of past data without detailed explanations of the causality of the illness under consideration seems unlikely to attract the confidence of physician or patient. People feel the need to explain phenomena in terms of mechanisms they understand, and tend to reject predictions which cannot be understood in such terms. Therefore, clinical judgment based on comparisons with collected data will fill an important but limited role. Other methods of computer use in medicine, relying on the encoding of knowledge held by the expert physician, will be at least as important.

Decision Theory

Decision theory is a mathematical theory of decision making under uncertainty. It assumes that one can quantify the *a priori* and conditional likelihoods of existing states and their manifestations and can similarly determine an evaluation (utility) of all contemplated outcomes. Given these data, decision theory offers a normative, *rational* theory of optimal decision making which is urged by its practitioners as an effective technique for structuring medical decision making problems [16]. Although there is considerable evidence that most human decision makers not specifically trained in decision analysis deviate from this model in their decision making activities [27], the theory is nevertheless appealing as a norm for helping to make explicit the bases of decision making and any existing disagreements among decision makers. Numerous computer programs for decision making in small domains of medicine have employed the decision theoretic formalism [6, 8].

The chief disadvantages of the decision theoretic approach are the difficulties of obtaining reasonable estimates of probabilities and utilities for a particular analysis. Although techniques such as sensitivity analysis help greatly to indicate which potential inaccuracies are unimportant, the lack of adequate data often forces artificial simplifications of the problem and lowers confidence in the outcome of the analysis. Attempts to extend these techniques to large medical domains in which multiple disorders may co-occur, temporal progressions of findings may offer important diagnostic clues, or partial effects of therapy can be used to guide further diagnostic reasoning, have not been successful. The typical language of probability and utility theory is not rich enough to discuss such issues, and its extension within the original spirit leads to untenably large decision problems. For example, one could handle the problem of multiple disorders by considering all possible subsets of the the primitive disorders as mutually competing hypotheses. The number of a priori and conditional probabilities required for such an analysis is, however, exponentially larger than that needed for the original problem, and that is unacceptable.

A second difficulty for decision analysis is the relatively mysterious reasoning of a decision theoretic program—an explanation of the results is to be understood in terms of the numeric manipulations involved in expected value computations, which is not a natural way of thinking for most people. The role of decision theoretic computations is discussed further in [24].

Additional Flexibility

A careful analysis of the shortcomings of any of the above techniques reveals numerous possible improvements. An interesting observation of the AIM community is that the improvements more often involve bringing to bear specific knowledge on selected subproblems of an application than developing a new complete theory for it. For example, in the decision theoretic framework, if most hypotheses are disjoint and most observations are conditionally independent, then it is very helpful to be able to express the few exceptions without resorting to expanding the complete database to give joint probabilities. Flexibility in knowledge representation and problem solving techniques is highly desirable to allow the inclusion of these bits of specific knowledge without needing to magnify greatly the whole program.

The five research projects reported on in this volume all employ AI techniques to represent and reason with their knowledge. In each case, similarities to more traditional forms of program organization will—not surprisingly—be apparent. Each project is pragmatically oriented, with the intent of ultimately producing a clinically significant tool. Although each is based in part on its developers' insights into how expert physicians reason, none is intended as a serious psychological model of human performance in medical reasoning. Thus, aspects of the predetermined clinical flowchart, pattern matching to a data base of known or prototypical cases, and probabilistic reasoning underlie each program where those techniques are appropriate. Of particular interest are the new techniques and their combinations which have been developed for these programs to provide the additional flexibility described above.

Expertise and Common Sense

Encoding human expertise in the computer is amazingly difficult. The difficulty rests both on our lack of understanding of how people know what they know and on technical problems of structuring and accessing large amounts of knowledge in the machine. For an example of a simple human reasoning task that is somewhat beyond the ability of current computer techniques to handle, consider the following dialog, quoted from *The New Yorker* in [7]:

> Mrs. Eloise Dobbs, 38, is married to a feed store owner and she comes to her physician, Dr. Elwood Schmidt, complaining of chest pain. The following dialogue ensues:
>
> "This whole side of my chest hurts, Elwood. It really hurts."
>
> "What about your heart—any irregular beats?"
>
> "I haven't noticed any. Elwood, I just want to feel good again."
>
> "That's a reasonable request. And I think it's very possible you will."
>
> "But what do you think? Is it my heart? Is it my lungs?"
>
> "Now, you won't believe this—but I don't know. I do not know. But I wonder. Are you lifting any sacks down at the store?"
>
> "I lift some. But only fifty pounds or so. And only for the woman customers."
>
> "I think you'd better let your lady customers lift their own sacks. If I know those ladies, they can do it just as well as you can. Maybe better."

The doctor in this story relies not only on his understanding of the physiological basis of pain (that although overexertion can exacerbate some underlying disorder to cause pain, especially in an older person it can cause pain by itself) but also on his knowledge of the patient and her occupation, the common practices of small-town stores, the weight of typical sacks of feed, etc. Therefore, we would not expect even the most sophisticated computer program, charged only with the latest of pathophysiological theory, to arrive at the parsimonious diagnosis of the local doctor.

An optimistic assessment holds that "tricks" like the above do not pose any real difficulty. After all, that reasoning process can be defined in terms of a small set of rules and facts:

1. Try to explain isolated complaints by possible non-pathological causes.
2. Overexertion can cause chest pain.