Medical Expert Systems—Knowledge Tools for Physicians

EDWARD H. SHORTLIFFE, MD, PhD, *Stanford, California*

Recent advances in the field of artificial intelligence have led to the emergence of expert systems, *computational tools designed to capture and make available the knowledge of experts in a field. Although much of the underlying technology available today is derived from basic research on biomedical advice systems during the 1970s, medical application packages are thus far generally unavailable from the young artificial intelligence industry. Medical expert systems will begin to appear, however, as researchers in medical artificial intelligence continue to make progress in key areas such as knowledge acquisition, model-based reasoning and system integration for clinical environments. It is accordingly important for physicians to understand the current state of such research and the theoretic and logistic barriers that remain before useful systems can be made available. One experimental system, ONCOCIN, provides a glimpse of the kinds of knowledge-based tools that will someday be available to physicians.*

(Shortliffe EH: Medical expert systems—Knowledge tools for physicians, *In* Medical informatics [Special Issue]. West J Med 1986 Dec; 145:830-839)

It was recognized in the early days of the computer age that our modern computational marvels, originally conceived as high-performance mathematical calculators, could be adapted to manipulate text and symbols as well as numbers. Thus computers became viewed as tools for storing and retrieving information, and their use as machines for decision support became a focus of research and development. Accepted tools such as MEDLINE later showed that medicine could benefit significantly from such work.

In the late 1950s scientists first began to suggest that computers might someday play a more active role in helping medical personnel reach decisions about diagnosis and patient management.[1] Instead of viewing computers merely as information sources—rather like electronic textbooks—or as tools for assisting with statistical analysis of large patient data bases, researchers suggested that computers could actually use such information, plus observational data about a patient, to generate individually tailored advice for a specific medical problem. For much of the next 20 years, the focus of such work was mathematical, emphasizing probabilistic reasoning and statistical pattern recognition.[2] In the early 1970s, however, several research groups argued that expert physicians make high-quality decisions without formal numerical analysis and that there must be symbolic or conceptual methods suitable for modeling expert decision making when problems are ill-structured or when formal statistical data are scarce or difficult to acquire. They accordingly turned for research ideas to studies of human problem solving and to the area of computer science with the closest ties to psychology, namely, the field of artificial intelligence (AI).

AI had been born in 1956 at a meeting at Dartmouth College (Hanover, NH) where leading computer scientists first articulated notions of machine intelligence. Alan Turing, the famed British mathematician, had suggested an operational definition for computer-based "intelligent behavior,"[3] but it was at the Dartmouth conference that computer scientists first decided to begin active research in the area. The subsequent early work emphasized the development of general problem-solving techniques and, by 1970, AI was epitomized by analyses and implementation of humanlike reasoning strategies, or heuristics, for focusing attention when solving problems and for maintaining efficient search through a range of possibilities.

Two pieces of work in the late 1960s provided insights

From the Medical Computer Science Group, Knowledge Systems Laboratory, Departments of Medicine and Computer Science, Stanford University School of Medicine, Stanford, California.

Dr Shortliffe is a Henry J. Kaiser Family Foundation Faculty Scholar in General Internal Medicine. Computing facilities were provided through a gift from Xerox Corporation and by the SUMEX-AIM resource under National Institutes of Health grant RR-00785.

Reprint requests to Edward H. Shortliffe, MD, PhD, Medical Computer Science Group, Knowledge Systems Laboratory, Departments of Medicine and Computer Science, Stanford University School of Medicine, Stanford, CA 94305-5479.

ABBREVIATIONS USED IN TEXT
AI = artificial intelligence
PIP = Present Illness Program

that are credited with a paradigm shift in AI research during the following decade. These were MACSYMA,[4] a system to help mathematicians solve complex calculus problems involving symbolic integration, and DENDRAL,[5] a program to help chemists identify unknown compounds from their mass spectral data. These systems differed from earlier AI work in that their power and utility were derived principally from their knowledge of a technical domain—that is, from the fields of expertise of the mathematicians and chemists who helped build them—rather than from a powerful generalized inference technique. The recognition that intelligent behavior by computers would depend on the effective encoding of large amounts of complex knowledge left AI researchers poised, in the early 1970s, to work in fields with an inherent emphasis on the use of expert-level knowledge. Medicine provided a natural focus for such investigations.

What Is an Expert System?

Four experimental systems are generally regarded as having started the research field of artificial intelligence in medicine.[6,7] These were MYCIN, a program to advise physicians on antimicrobial selection for patients with bacteremia or meningitis[8,9]; the Present Illness Program (PIP), a system that gathered data and generated hypotheses about disease processes in patients with renal disease[10]; INTERNIST-1, a large system to assist with diagnosing complex problems in general internal medicine,[11] and CASNET, an ophthalmology advisor designed to assess disease states and to recommend management for patients with glaucoma.[12] All four drew on AI techniques, emphasizing the encoding of large amounts of specialized medical knowledge acquired from the clinical literature and from expert collaborators. None used classical statistical techniques, nor did they base their advice on interpretations of accumulated experience in patient data banks. On the other hand, each was influenced by earlier AI work on general problem-solving techniques, and two of the systems (PIP and INTERNIST-1) explicitly modeled hypotheticodeductive behavior,[13,14] the familiar process by which physicians formulate tentative hypotheses rapidly after obtaining the first few pieces of information about a patient and then let those hypotheses (typically a differential diagnosis) guide further data collection and problem solving.

It was this handful of medical systems, plus a geology advisor known as PROSPECTOR,[15] that led to a growing interest in AI systems that might function as expert consultants. By the late 1970s such systems had become known as "knowledge-based systems" or "expert systems," terms that continue in common use. Thus, the term "expert system" originally implied a computer-based consultation system using AI techniques to emulate the decision-making behavior of an expert in a specialized, knowledge-intensive field.[16] The term has subsequently been broadened as the field has been popularized, so that an expert system's roots in artificial intelligence research can no longer always be presumed. There are arguments for calling *any* decision support system an expert system if it is designed to give expert-level problem-specific advice, even if the underlying programming and analytic techniques differ from the knowledge-based methods developed by AI researchers.

The explosive popularity of the expert systems notion has moved much of the applied work in this field from university research laboratories to the scores of software companies that have emerged in a fledgling industry since 1980. The successes of the technology have been sufficient to attract the interest of many of the major corporations in the country, several of which have invested heavily in start-up AI companies or have sought to develop in-house expertise in the field. The demand for trained AI scientists has skyrocketed, and one financial observer group has predicted total US spending on expert systems at $150 million in 1986, rising to more than $200 million annually within two years (The Yankee Group, quoted in *Expert Systems User,* vol 2, No. 5, August 1986, p 3). Several new journals and trade magazines have appeared, and the American Association for Artificial Intelligence has grown to a membership of more than 12,000 since its inception in 1979.

Most expert systems developed by the young artificial intelligence industry are used internally within the investing companies. The new technology has generally been used to gather experiential knowledge that had previously resisted efforts to encode it using conventional programming techniques. Examples of successful expert system applications in industrial settings include

- A program that assists with the design of a physical layout for computing components assembled within a metal cabinet.
- A system that assists in the interpretation of soil samples recovered from drilling sites during oil exploration.
- A financial management advisor to assist stock brokers.
- Systems to assist in fault diagnosis for items as diverse as electronic circuits, automobiles and steam locomotives.

It is ironic that a field that owes much of its history to research in medical computing should have grown to commercial viability in industrial settings, whereas the AI companies have thus far largely avoided the medical marketplace. By 1985 only three medical expert systems were being used routinely in clinical settings.[17-19] Of these, only one was a commercial product[18] and two functioned without requiring direct use by physicians.[17,18] (A related program is the University of Utah's HELP system,[20] a hospital information system that incorporates decision-support functions by using a representation scheme and reasoning mechanisms with similarities to those in the expert systems cited.)

Although many of the reasons for this phenomenon are logistic or sociologic, most are inherently scientific and have to do with the complex issues that arise when conventional expert system techniques are applied to high-stakes medical problems about which knowledge is still limited, physiologic models are incomplete and the uncertainty in both associations and interpreting data is severe.

Research Perspectives in Medical Artificial Intelligence

For purposes of the following discussion, an expert system will be assumed to be a consultation system that uses AI techniques for encoding knowledge and solving problems with that knowledge. The key AI notions in such systems are as follows:

- Symbolic rather than numeric representations of pertinent knowledge from the domain (field) of application.
- General problem-solving procedures for using knowledge stored in accordance with the prescribed conventions.
- Both general and domain-specific heuristics, or rules of thumb, for limiting search and for dealing with areas in which knowledge is limited or missing.
- Schemes, either numeric or categoric (algorithmic), for dealing with the uncertainty in the domain.

The following sections outline several of the key research and conceptual topics in medical AI research. Because the AI notions listed above may not be familiar to physicians, they are clarified with medical examples at appropriate points within the text.

Knowledge and Heuristics

The distinction between "data" and "knowledge" is generally emphasized in expert systems work. Computer data bases are a familiar concept—that is, collections of individual observations or data points; what, then, is a computer knowledge base? First, factual knowledge tends to be drawn from analyzing data. As such, it is frequently subject to controversy and colored by personal experiences. All physicians are familiar with scientific debates about how a given set of data ought to be interpreted; the 1970s controversy regarding oral hypoglycemic agents after the University Group Diabetes Project reported its results is only one such example. A computer data base might record the observation that "Mr John Jones had a blood pressure of 180/110 mm of mercury on August 3 and a myocardial infarction on September 15." On the other hand, a knowledge base that included information derived from analyzing observational data in a data base might record the *fact* that hypertension is associated with an increased risk of coronary artery disease.

Not all knowledge comprises factual associations of this type, however, and intelligent behavior relies in large part on other categories of knowledge. For example, shared knowledge of the world often provides a background of common sense, definitions and assumptions that hardly seem worth mentioning in normal conversation but which must be explicitly encoded in a computer if its behavior is to be appropriate. Consider, for example, the common-sense knowledge that only women can be pregnant, or the assumption that seriously ill patients are more likely to do well if they are admitted to hospital. We generally do not quote data to support such knowledge; the statements are assumed to be widely accepted as true and hardly need to be defended.

Another category of knowledge is experience-based heuristics, or "rules of thumb." Such knowledge is often rather personal, although it can be taught. For example, there is a common heuristic that it is wise to insert an intravenous line early when evaluating an acutely ill patient in an emergency room. The line can always be removed if not needed, but can be indispensable if a patient's condition deteriorates quickly. Matters of personal and professional style also often have a heuristic character to them—for example, the notion that it is wise to build rapport by conversing with an anxious patient before proceeding to the physical examination.

Fortunately there is also a growing number of medical topics for which we have excellent scientific (mechanistic) models that explain aspects of both normal behavior and disease. Such models can be crucial for expert problem solving because they provide a basis for reasoning from first principles when a patient presents with an unusual diagnostic problem or a difficult management dilemma. Such models can be qualitative ("Insulin controls blood sugar levels by regulating the uptake of glucose by cells") or quantitative ("Bicarbonate may be determined from pH and partial carbon dioxide pressure using the Henderson-Hasselbalch equation"). When an expert system functions in a domain for which good models exist, its performance may be greatly enhanced if such information is encoded in its knowledge base.

Symbolic Representation of Knowledge

Although a number of schemes have been used to represent expert knowledge symbolically in a computer, the rule-based approach is perhaps the most widespread and straightforward. Other representation methods can be studied in any of the several excellent textbooks on artificial intelligence.[21,22] Rules are conditional "if-then" statements that indicate circumstances under which conclusions can be drawn or actions taken. Rules can be categorized in several ways, but for illustrative purposes it is useful to think of four major types:

- *Definitional*—for instance, "*If* a patient is male, *then* the patient is not pregnant or lactating."
- *Cause-to-effect*—for instance, "*If* there is an elevation in the serum parathormone level and the patient has normal renal function, *then* anticipate decreased urinary calcium and increased urinary phosphate levels."
- *Effect-to-cause*—for instance, "*If* a patient has recurrent calcium oxalate kidney stones, *then* consider the diagnosis of primary hyperparathyroidism."
- *Associational*—for instance, "*If* a patient has Gram-negative sepsis and has been seriously burned, *then* the offending organism may be *Pseudomonas aeruginosa.*"

Note that an associational rule tends to describe a relationship for which detailed causal mechanisms are not well understood; much medical knowledge is of this type. In addition, only the first of the four sample rules shown allows a conclusion to be reached with certainty (a so-called categorical rule). The others all lend evidence to a conclusion but do not provide definitive proof. Thus, it is sometimes necessary for the knowledge represented in expert systems to include indications of the degree of certainty associated with the relationship described. AI researchers have experimented with a variety of techniques for dealing with weights of evidence.[23] These vary from classic probabilistic measures and the use of Bayes's theorem, to ad hoc scoring schemes and even to efforts to avoid numeric weights altogether.

Rules such as those just described may be encoded in a computer's memory using any of a number of programming languages. The most commonly used for AI work have been LISP and, more recently, PROLOG. Unlike traditional programming languages like FORTRAN, LISP and PROLOG have been designed more for their ability to express logical relationships and to manipulate symbols than for an ability to compute numbers.

Reasoning With Symbolic Knowledge

Encoding knowledge in a computer serves no purpose if it cannot be retrieved and applied effectively. Much of the work in artificial intelligence has therefore focused on how to access and use knowledge that is stored in large and detailed knowledge bases. As previously mentioned, psychological studies have been influential in this work because of the recognition that human beings have a remarkable ability to select and apply relevant knowledge from the vast amount of information they have acquired over a lifetime. People have developed effective search heuristics for quickly focusing on pertinent facts while avoiding tangential considerations; we obviously do not sequentially consider everything we know until we discover a fact that applies to a given problem.

Reasoning methods in AI systems are frequently placed in one of three categories: goal-directed (also called backward chaining), data directed (also called forward chaining) and hypothesis-directed.[24] In a goal-directed system, rules are selected for consideration because of what they might conclude (if true) that is relevant to a diagnostic or management problem—that is, because of the "then" portion of the rule. For example, if a system needed to identify an organism, it might select for consideration the sample associational rule shown earlier. In a data-directed system, rules are selected for consideration because they use information that has become available about a problem—that is, because of the "if" portion of the rule. For example, if a physician reports that a patient has a calcium oxalate kidney stone, a data-directed system might then invoke the sample effect-to-cause rule and add hyperparathyroidism to the differential diagnosis. A hypothesis-directed system, the type that most closely mimics the hypotheticodeductive behavior mentioned earlier, begins with a data-directed invocation of initial hypotheses but then selects additional rules for consideration based on the set of active hypotheses. This notion is similar to letting a differential diagnosis for a case guide subsequent data collection, which in turn allows the diagnostic hypothesis list to be refined.

In recent years system designers have often sought to avoid the constraints of being committed to any single problem-solving technique. One method for doing such "opportunistic reasoning" has been the notion of a "blackboard model"—a working area in a computer's memory that captures the full scope of a problem and allows the problem-solving mechanism to select whichever reasoning technique and knowledge source are more likely to be useful at a given point. The ONCOCIN system described below uses another type of mixed reasoning strategy. More complete discussions of the alternate approaches can be found in sources such as the three-volume *Handbook of Artificial Intelligence*.[25]

Acquiring and Encoding the Knowledge

Exclusively focusing on representing knowledge and using it within a computer would ignore a key additional issue—how the knowledge for a system is acquired and formulated. It is in this area that much of the expert systems "mystique" has evolved. Persons who work with experts to structure and encode their knowledge of a domain have been dubbed "knowledge engineers." There is little doubt that the process of mapping the ill-structured knowledge of a domain such as a medical subspecialty into a form suitable for machine encoding is among the most difficult and time-consuming parts of the expert system building process. Not only must knowledge engineers be familiar with the technical details of the computational tools available, but they must also be willing to make a major commitment to learning enough about the technical domain of an expert so that discussions of sample problems can be substantive and detailed. It is a special breed of programmer who is able to facilitate this transformation of key elements of expertise from a collaborator's mind to a computer knowledge base.

In recent years the knowledge engineering "bottleneck" in expert systems development has encouraged researchers to develop prototype tools that permit experts to "teach" a computer directly about their specialties. The earliest work of this type was a system named TEIRESIAS that allowed infectious disease experts to update and edit the knowledge base for the MYCIN system by critiquing MYCIN's performance on sample cases and entering in English either new rules or modifications to old ones.[26] More recent work on knowledge acquisition has been exploring the use of graphical techniques for defining knowledge. Diagrams on a computer screen, coupled with manually controlled interactive mechanisms—such as touch screens, light pens or "mouse" pointing devices*—can be made intuitive, permitting experts to outline their knowledge without doing extensive keyboard typing.

Model-Based Reasoning

As was mentioned earlier, medical research is increasingly providing us with mechanistic explanations for phenomena that were once poorly understood. As we gain greater insight into why and how the body works the way it does, experts increasingly use such mechanistic knowledge in reaching decisions, especially when reasoning from first principles to develop creative solutions to unusual or aberrant problems. If computer-based systems lack such knowledge and the ability to use it, their performance is inherently limited. The encoding and use of causal knowledge has accordingly become an area of emphasis in medical AI research.[27,28] Although mathematical modeling has involved medical informatics researchers for several decades, the use of more qualitative models, or models that combine qualitative and quantitative techniques, is a relatively recent development.

When underlying mechanistic models are vague or nonexistent, physicians are forced to reason using empiric associations such as those reported in clinical studies. Much of the knowledge of medicine is of this type. Medical AI researchers have accordingly also begun to look at how expert systems might reason using the kinds of data found in the clinical literature.[29] In this work the "models" are those of statistics and clinical trials rather than mechanistic explanations of observed phenomena.

The Logistics of Human-Computer Interaction

As medical artificial intelligence has begun to move from research laboratories into clinical settings, there has been a growing emphasis on logistic and design issues that will en-

*A "mouse" is a mechanical device named for the taillike cord that connects it to the computer terminal. The user rolls it on the adjacent desktop to control the location of a pointer on the display screen. Items of interest are selected by moving the mouse so that the pointer on the screen is positioned over them and then pressing a button on the top of the mouse.

courage physicians to use computers and assure that they are viewed as a helpful tool rather than as a hindrance to efficiency or as a professional threat. For example, there has been an emphasis on giving advice systems an ability to explain their reasoning, a requirement that seeks to clarify and emphasize the physician's role as the ultimate decision maker in patient management.[30] There has also been a large and growing research effort in the area of human-computer interaction, with studies of alternate interactive techniques (such as keyboards, number pads, touch screens, mouse pointing devices, light pens, speech input) and time-motion studies of clinical settings to provide insight into how best to integrate computational tools in routine patient care environments. It is highly likely that advances in such areas will be as important to the clinical impact of decision-support technologies as will be progress in the fundamental research topics mentioned earlier. A related issue is the inherent resistance to using computers that is frequently mentioned by physicians; until this can be overcome through education and experience, it is likely that medicine will remain among the last areas to benefit from the new decision-support technologies that are beginning to have an impact in other areas of science and society.

An Example of an Expert System

ONCOCIN is an advanced expert system for clinical oncology that has been under development at Stanford University School of Medicine since 1979.* It is designed for use after a diagnosis has been reached, focusing instead on assisting with the management of patients with cancer who are receiving chemotherapy. Because anticancer agents tend to be highly toxic and because their tumor-killing effects are routinely accompanied by damage to normal cells, the rules for monitoring and adjusting treatment in response to a given patient's course over time tend to be complex and difficult to memorize. ONCOCIN integrates a temporal record of a patient's ongoing treatment with an underlying knowledge base

*The author is principal investigator for the ONCOCIN project, an effort by physicians and computer scientists under grant support from the National Library of Medicine and the National Institutes of Health's Division of Research Resources. L. Fagan, MD, PhD, is Project Director, and oncology collaborators include C. Jacobs, MD; R. Carlson, MD; B. Sikic, MD, and R. Lenon, MD.

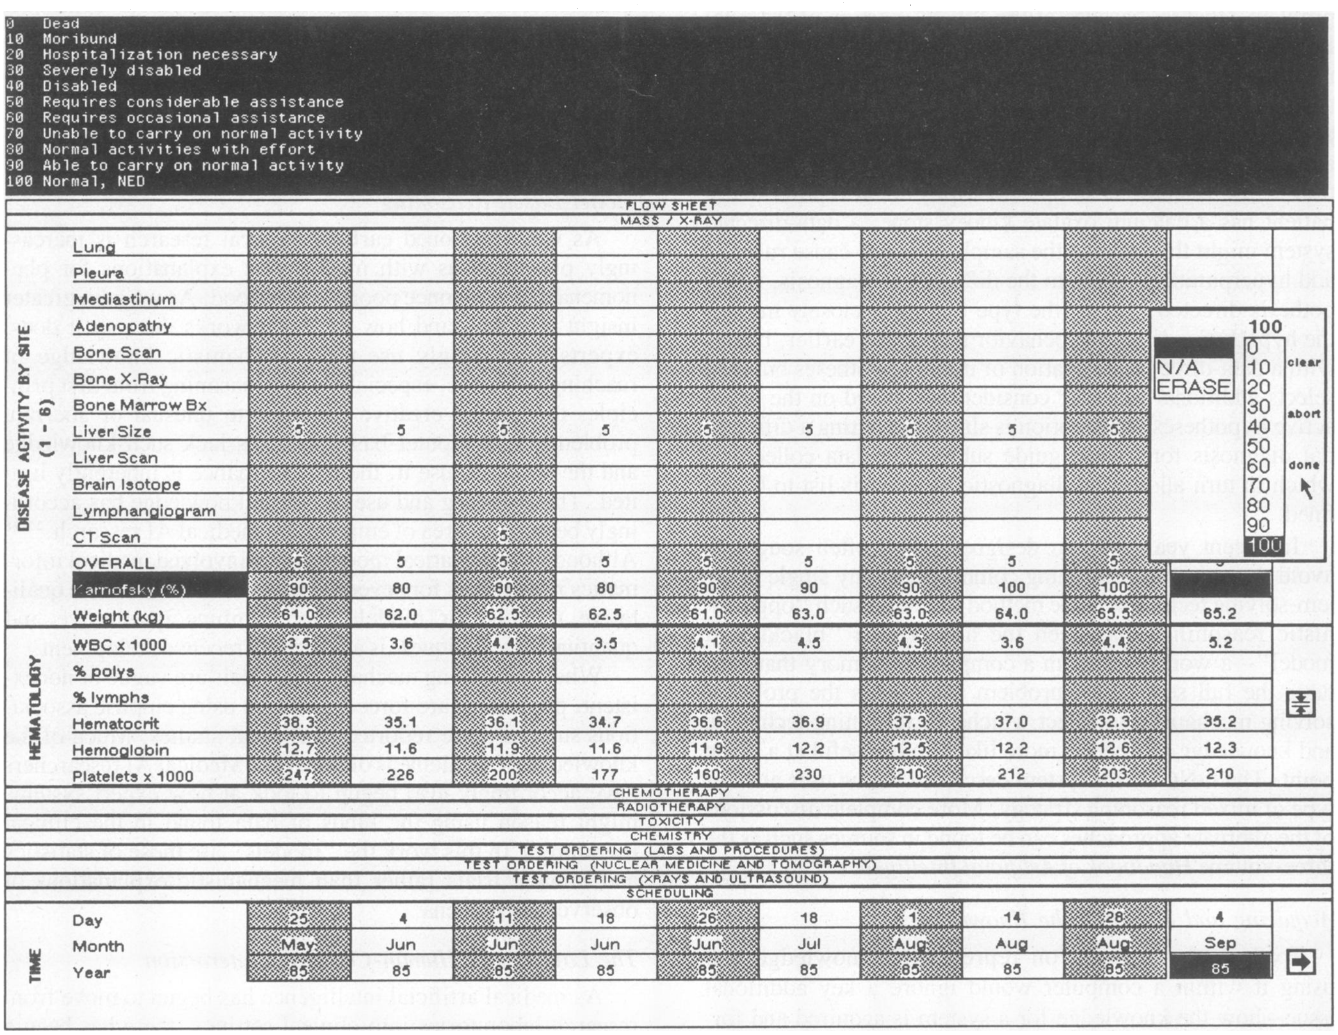

Figure 1.—ONCOCIN's electronic flow sheet: The computer's display screen simulates the appearance of a familiar paper flow sheet. Two sections ("windows") of the form are shown open: "hematology" and "disease activity by site." Closed sections (labeled horizontal bars) may be opened by simply selecting one or more of them with a mouse pointer. In this figure, the physician is logging in the patient's Karnofsky status in the right-hand column of data for the current clinic visit (September 4, 1985). A customized menu register is displayed (see arrow pointer at middle right, controlled by the mouse), and a dark "explanation window" at the top of the screen shows the meaning of the various options for the current data field.

of treatment protocols and rules for adjusting dosage, delaying treatment, aborting cycles, ordering special tests and similar management details. The program uses such knowledge to help physicians with decisions regarding the management of specific patients. This article cannot provide a detailed description of ONCOCIN; other publications have reported the basic system design,[31,32] the results of clinical evaluations[19,33] and the nature of the physicians' interface.[34] Instead, we will briefly describe how this system has addressed the AI research issues outlined in previous sections.

The Logistics of Human-Computer Interaction

A major lesson of past work in clinical computing has been the need to develop methods for integrating a system smoothly into the patient-care environment for which it is intended.[32] In the case of ONCOCIN, the goal has been to provide expert consultative advice as a by-product of the patient data-management process, thereby avoiding the need for physicians to go out of their way to obtain advice. It is intended that oncologists use ONCOCIN routinely for recording and reviewing patient data on the computer's screen, regardless of whether they feel they need decision-making assistance. This process replaces the conventional recording of data on a paper flow sheet and thus seeks to avoid being perceived as an additive task. In accordance with its knowledge of a patient's chemotherapy protocol, ONCOCIN then provides assistance by suggesting appropriate therapy at the time that a day's treatment is to be recorded on the flow sheet. Physicians maintain control of the decision, however, and can override the computer's recommendation if they wish. ONCOCIN also indicates the appropriate interval until a patient's next treatment and reminds the physician of radiologic and laboratory studies required by the treatment protocol.

Although the original ONCOCIN prototype was developed on a large time-shared computer,[19,31] it was clear that disseminating the technology to clinics and private offices would require transferring the system to smaller, less expensive machines. ONCOCIN has thus been rewritten to run on single-user LISP machines, computers optimized to run the LISP programming language mentioned earlier. (We use Xerox 1100 series workstations, but similar LISP machines are available from several other manufacturers.) In addition to large memories and enough power to handle systems the size of ONCOCIN, such machines provide high-quality graphics screens and mouse pointing devices. Thus, ONCOCIN's interface for physicians has been redesigned to take advantage of these advanced graphics capabilities.[34]

The interface (Figure 1) uses multiple movable "windows" to simulate a traditional paper flow sheet used in an oncology clinic.* Each window represents a section of the flowsheet. Because not all the sections can be viewed simultaneously—as they would more than fill the screen—individual windows can be "closed" or compressed when not in use. A physician enters data on ONCOCIN's flow sheet using the mouse pointer and specially developed software input devices called "registers." Such registers are displayed on the screen and manipulated using the mouse and its selection button. For example, registers let the physician select an item from choices (displayed on a "menu") or enter numbers (using a graphic calculator keypad). Registers have been built for almost every kind of data input on the flow sheet; the computer's keyboard is used to enter only textual information such as names and addresses. Because such text information is generally entered once by a data manager when the patient is first seen in the clinic, registers and the mouse pointer allow physicians to avoid using a keyboard.

*The interface system described here has been largely the work of C. Lane and C. Wulfman, MS.

	To determine the dose of methotrexate administered in VAM chemotherapy in protocols 20-83-1 and 2091:
If:	the serum creatinine level (in mg per dl) exceeds 1.5
Then:	do not give methotrexate

Figure 2.—A simple ONCOCIN rule: Rules in ONCOCIN are conditional statements that indicate circumstances under which a given conclusion can be reached about a patient's management. The rule shown is one of many that could simultaneously be applied in considering the case of a patient undergoing protocol-directed therapy for small-cell carcinoma of the lung. VAM = VP-16-213 (etoposide), Adriamycin (doxorubicin) hydrochloride and methotrexate

Although ONCOCIN at first appears to be designed as a data-management tool for following events that occur over time, there is an underlying reasoning element that uses the data as they are entered, along with a patient's historical information, to determine the recommended therapy for that clinic visit. The knowledge for making such decisions is derived from chemotherapy protocols and from the experience of collaborating oncologists who have assisted in ONCOCIN's development. When the "chemotherapy" section of the graphic form is opened at the end of the session (see Figure 1), recommended drug doses are automatically filled in on the flow sheet by the computer, and the explanation window provides the reasons for any dosage attenuation, delays in treatment or aborted cycles of chemotherapy that have been recommended. The physician may override such advice by simply entering a different drug dose on the flow sheet in lieu of ONCOCIN's suggestion.

Symbolic Representation of Knowledge

The internal knowledge representation for ONCOCIN is based on the notion of decision rules such as those described earlier.* There are hundreds of such rules in ONCOCIN, some of which are specific to particular treatment protocols and others of which apply generally across all cancer management strategies. Many encode management heuristics provided by oncologists and never explicitly stated in protocol documents. A typical simple ONCOCIN rule is shown in Figure 2. Such rules are internally arranged in a data structure that reflects the hierarchic organization of the field of cancer chemotherapy (Figure 3). By storing rules in the computer in accordance with such hierarchic conventions, it can be assured that knowledge is considered and applied only when the context is appropriate. For example, the rule shown in Figure 2 is explicitly associated with the drug methotrexate, chemotherapy VAM (VP-16-213 [etoposide], Adriamycin [doxorubicin] hydrochloride and methotrexate) and protocol 20-83-1 or protocol 2091 (see Figure 3); it would never be considered if methotrexate were being administered in some other situation.

*The knowledge representation scheme used on ONCOCIN has been largely the work of S. Tu, MS; M. Kahn, MD; M. Musen, MD, and J. Ferguson.

Reasoning With Symbolic Knowledge

The reasoning strategy used by ONCOCIN is guided by the hierarchic structure of the knowledge base described earlier. The focus of attention within the hierarchy at any given time determines the reasoning context and thus the rules that are likely to be pertinent. Because a physician may enter information on the flow sheet in whatever order seems most natural for a given case, however, the system's reasoning

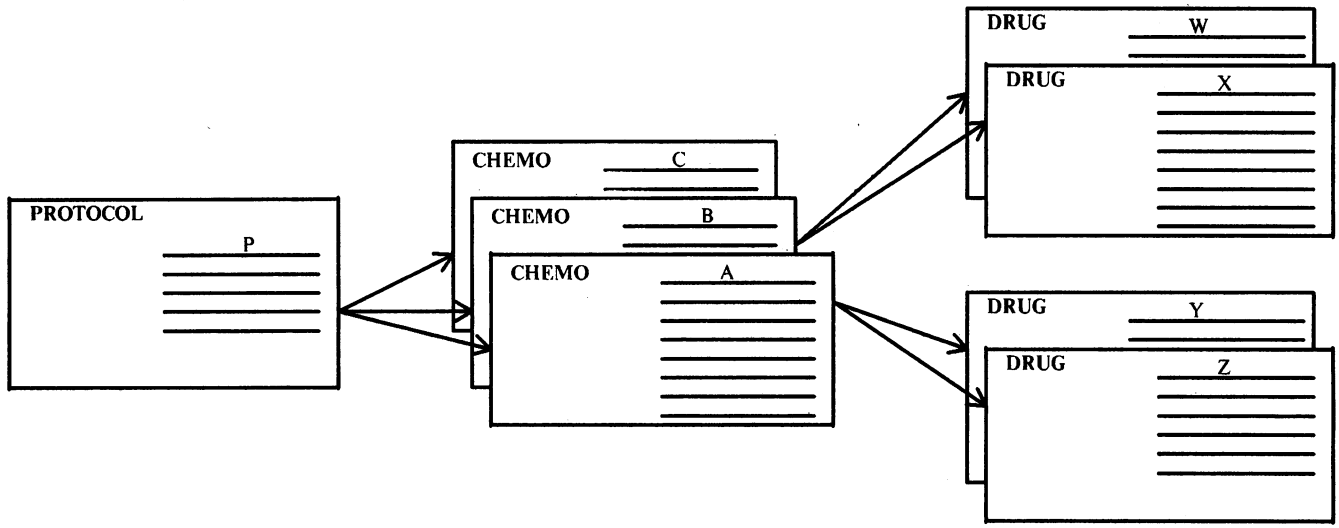

Figure 3.—A portion of ONCOCIN's domain hierarchy: Relevant concepts in ONCOCIN's knowledge base are naturally organized according to the hierarchy shown here. In this example, a protocol *P* involves treatment with chemotherapies *A, B* or *C.* Chemotherapy *A* requires administration of drugs *Y* and *Z,* whereas chemotherapy *B* uses drugs *W* and *X.*

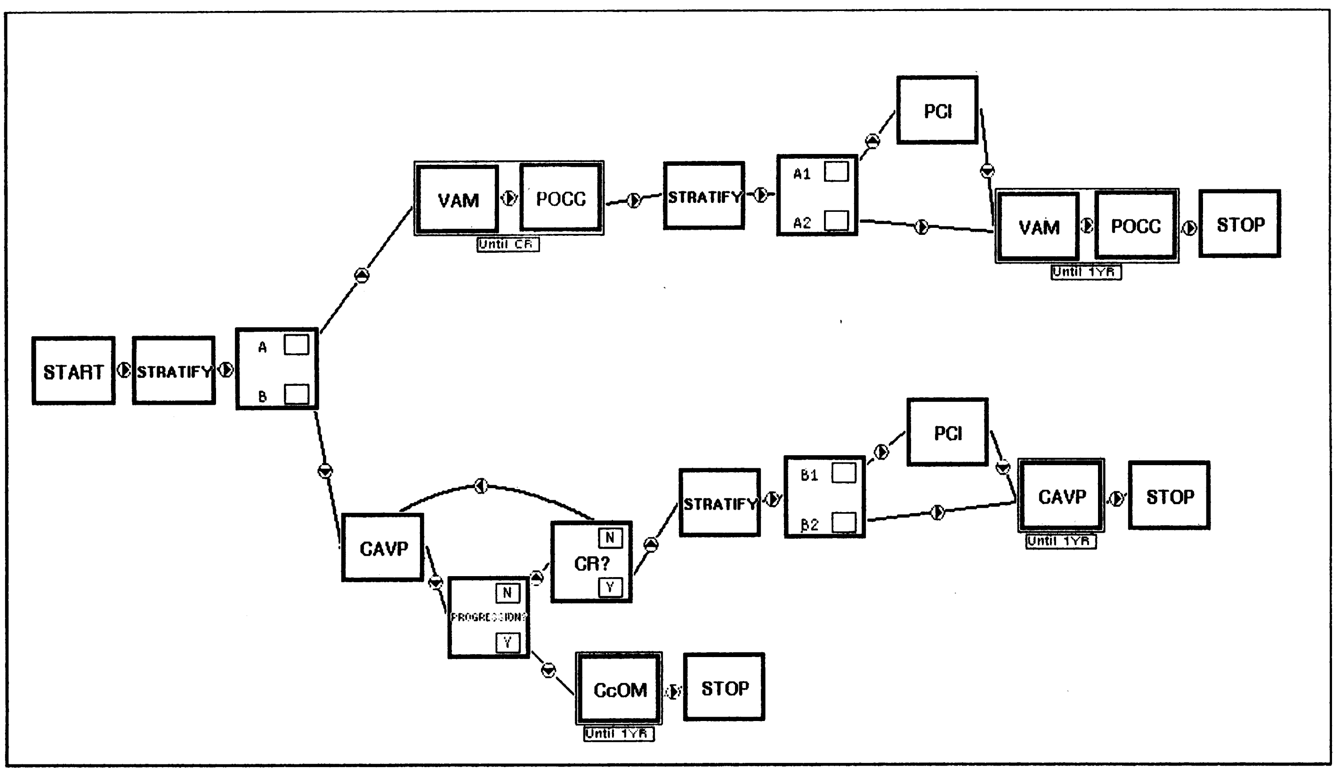

Figure 4.—The schematic description of a protocol entered using OPAL: Flow charts of this sort are entered by manipulating boxes on the display screen and using the mouse pointer to indicate connections between them. The procedural descriptions correspond to flow diagrams included at the beginning of most protocol documents. Individual boxes stand for events such as chemotherapy—for instance, VAM (VP-16-213 [etoposide], Adriamycin [doxorubicin hydrochloride] and methotrexate), POCC (procarbazine hydrochloride, Oncovin [vincristine sulfate], CCNU [lomustine] and cyclophosphamide) or CAVP (cyclophosphamide, Adriamycin and VP-16-213)—or radiation therapy—for instance, prophylactic cranial irradiation (PCI). The diagram shown here represents the schema for a protocol for small-cell carcinoma of the lung in use at our institution. Note that cases are stratified and patients randomly assigned to one of two arms (*A* and *B*) designed to compare different chemotherapies (using VAM and POCC combined versus using CAVP), with subsequent randomization to assess the effect of PCI. CcOM = CCNU, Oncovin and methotrexate; CR = complete remission

strategy must also respond to patient data as they are entered. Schemes have accordingly been devised so that data entered by the physician are remembered but not actually used until they are needed by the hierarchically guided reasoning system. To minimize conflicts between the data the physician chooses to enter and the information needed by the reasoning system, the interactive screen environment suggests pathways through the flow sheet that will result in the most efficient consideration of a given case (ONCOCIN highlights those regions of the screen where it would next like to see data entered). But the ultimate organization of the data-entry process is intentionally left to the physician. Thus, ONCOCIN's reasoning strategy is determined both by the nature of the interaction between a physician and the computer-based flow sheet *and* by the hierarchical organization of its knowledge. It accordingly does not fit cleanly into any single reasoning category but has both data-directed and goal-directed features.

Acquiring and Encoding the Knowledge

Early knowledge-base development for ONCOCIN was arduous and time-consuming. It was difficult for project knowledge engineers to study chemotherapy protocol documents and to encode their contents in accordance with ONCOCIN's knowledge representation scheme. In addition, there were frequent gaps in protocol logic that required extensive sessions with oncology collaborators who helped interpret the several areas in which protocols were ambiguous, incomplete or internally inconsistent. As the structure of the oncology domain became clearer (see, for example, the simple hierarchy in Figure 3 that emerged from the knowledge-engineering process), we developed the notion of a knowledge acquisition system that would permit expert clinicians to enter protocols directly onto ONCOCIN's knowledge base. The resulting system, known as OPAL, has been used experimentally by the collaborating oncologists on the

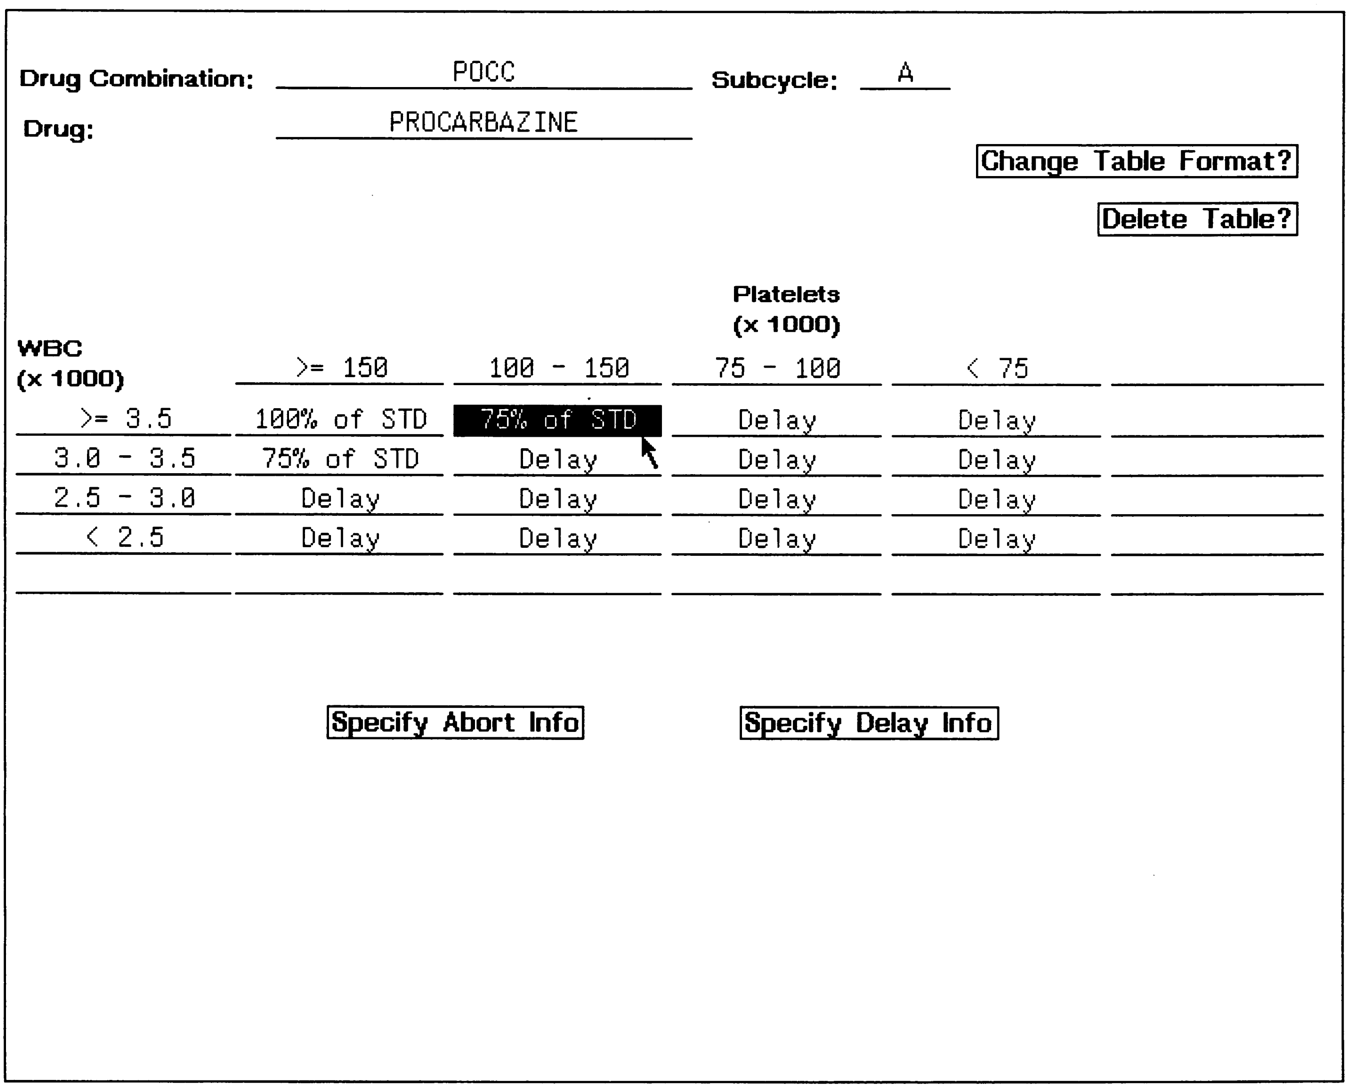

Drug Combination: POCC **Subcycle:** A

Drug: PROCARBAZINE

Change Table Format?

Delete Table?

Platelets (x 1000)

WBC (x 1000)	>= 150	100 - 150	75 - 100	< 75	
>= 3.5	100% of STD	75% of STD	Delay	Delay	
3.0 - 3.5	75% of STD	Delay	Delay	Delay	
2.5 - 3.0	Delay	Delay	Delay	Delay	
< 2.5	Delay	Delay	Delay	Delay	

Specify Abort Info

Specify Delay Info

Figure 5.—A sample OPAL form: Much of an interactive session with OPAL requires completing forms such as the one shown here. As in the ONCOCIN interface (Figure 1), most entries may be completed by pointing at areas of interest and selecting from a customized register that appears on the screen. Here an oncologist has entered a dosage-attenuation schedule for procarbazine hydrochloride when it is given in the *A* subcycle of POCC (procarbazine hydrochloride, Oncovin [vincristine sulfate], CCNU [lomustine] and cyclophosphamide) chemotherapy. The highlighted choice indicates that when a patient's leukocyte count (WBC) is greater than or equal to 3,500 per μl and the platelet count is between 100,000 and 150,000 per μl, 75% of the standard dose (STD) should be administered.

project.* Early experience suggests it will greatly facilitate the process of protocol entry and modification for ONCOCIN.

OPAL is a graphical program for use by an oncologist who wishes to enter a new chemotherapy protocol for use by ONCOCIN or to edit an existing protocol.[35] Although the system is designed for use by oncologists who have been trained in its use, it does not require an understanding of the internal representations or reasoning strategies used by ONCOCIN. The system may be used in two interactive modes, depending on the type of knowledge to be entered. The first permits the entry of a graphical description of the overall flow of the therapy process (Figure 4). The oncologist manipulates boxes on the screen that stand for various steps in the protocol. The resulting diagram is then translated by OPAL into computer code for use by ONCOCIN. Thus, by drawing a flow chart that describes the protocol schematically, the physician is effectively programming the computer to carry out the procedure appropriately when ONCOCIN is later used to guide the management of a patient enrolled in that protocol.

OPAL's second interactive mode permits an oncologist to describe the details of the individual events specified in the graphical description. For example, the rules for administering a given chemotherapy will vary greatly depending on a patient's response to earlier doses, intercurrent illnesses and toxicities, hematologic status and so forth. Figure 5 shows one of the forms provided by OPAL for this type of specification. It permits the entry of an attenuation schedule for an agent based on a patient's leukocyte count and platelet count at the time of treatment. Tables such as this are generally found in the written version of chemotherapy protocols. Thus, OPAL permits oncologists to enter information using familiar forms displayed on a computer screen. The contents of such forms are subsequently translated into rules and other knowledge structures for use by ONCOCIN.

Model-Based Reasoning

Although the knowledge of cancer chemotherapy is rich and complex, protocols seldom refer directly to underlying models of drug action. The guidelines in a protocol are, rather, high-level composite descriptions of expert advice, based on the study designers' experience and biologic models of the therapeutic agents and their mechanisms of action. We have observed, however, that when protocols fail to cover a complex clinical situation that arises for a given patient, expert oncologists will turn to underlying mechanistic models and use them to assist in the decision-making process. ONCOCIN has no such knowledge; it must therefore occasionally decline to make a recommendation and instead refers a physician to the study overseer for a decision about how to manage a particular complex problem. It is accordingly a long-range goal to add model-based expert-level reasoning to ONCOCIN's performance.

Our research in oncologic model-based reasoning is embodied in a program known as ONYX.† This system is based on the observation that creative planning strategies in the oncology domain (and many other fields) appear to involve a three-step process[36]: (1) heuristically generating a small number of plans, that is, plausible responses to the problem at hand, (2) mentally simulating (also called "envisioning") how a patient would respond over time if each of those plans were carried out and (3) selecting a preferred plan based on the likelihood of the various possible outcomes and the value placed on those outcomes by the patient and physician. Step 2 in this process involves patient-specific simulation of tumor pathophysiology and drug action but it also depends on recognizing that the outcomes of interventions cannot be predicted with certainty and that probabilistic predictions are more realistic. Thus, model-based probabilistic simulations in ONYX are coupled to a decision analytic module that assists with the third step in the process. Although the work outlined here is preliminary, the proposed planning architecture appears to be a fruitful area for basic research in medical artificial intelligence.

*OPAL is largely the work of D. Combs, M. Musen, MD, and J. Walton.

†ONYX is largely the work of C. Langlotz, MS; H. Rappaport, MD; S. Tu, MS; M. Kahn, MD, and R. Bhatia.

Conclusion

As the ONCOCIN project shows, applied work in medical expert systems typically touches on a broad range of research issues ranging from logistic (such as system integration) to psychological (such as designing human-computer interfaces) to theoretic (such as causal model-based reasoning). It is this melding of real-world needs with promising technologies and challenging theoretic concerns that has made the field of medical artificial intelligence particularly appealing to the physicians and computer scientists working in the area. The long-term challenges are well recognized—such as the need for mechanisms that will assure completeness and currency of shared knowledge bases, or the need for integrated computing systems that will allow physicians to access a variety of expert systems and other computational tools from a single workstation in their office or on the hospital wards. Yet, it is only through enhanced educational opportunities for health personnel that we will see the emergence of physicians who can identify the quality systems from among those that are available and accept the computer's role as a knowledge-enhancing tool rather than as a replacement for a physician's own thoughtful assessments. It is ultimately an informed professional community that will determine the clinical impact of medical advice systems.

REFERENCES

1. Ledley RS, Lusted LB: Reasoning foundations of medical diagnosis. Science 1959; 130:9-21
2. Shortliffe EH, Buchanan BG, Feigenbaum EA: Knowledge engineering for medical decision making: A review of computer-based clinical decision aids. Proc IEEE 1979; 67:1207-1224
3. Turing AM: Computing machinery and intelligence. Mind 1950; 59:433-460
4. Moses J: Symbolic integration: The stormy decade. Commun ACM [Association for Computing Machinery] 1971; 8:548-560
5. Lindsay RK, Buchanan BG, Feigenbaum EA, et al (Eds): Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project. New York, McGraw-Hill, 1980
6. Szolovits P (Ed): Artificial Intelligence in Medicine. Boulder, Colo, Westview Press (AAAS Symposium Series), 1982
7. Clancey WJ, Shortliffe EH (Eds): Readings in Medical Artificial Intelligence: The First Decade. Reading, Mass. Addison-Wesley, 1984
8. Yu VL, Fagan LM, Wraith SM, et al: Antimicrobial selection by a computer: A blinded evaluation by infectious disease experts. JAMA 1979; 242:1279-1282
9. Buchanan BG, Shortliffe EH (Eds): Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project. Reading, Mass, Addison-Wesley, 1984
10. Pauker SG, Gorry GA, Kassirer JP, et al: Toward the simulation of clinical cognition: Taking a present illness by computer. Am J Med 1976; 60:981-995
11. Miller RA, Pople HE, Myers JD: INTERNIST-1: An experimental computer-based diagnostic consultant for general internal medicine. N Engl J Med 1982; 307:468-476

12. Weiss SM, Kulikowski CA, Amarel S, et al: A model-based method for computer-aided medical decision making. Artif Intell 1978; 11:145-172

13. Elstein AS, Shulman LS, Sprafka SA (Eds): Medical Problem Solving: An Analysis of Clinical Reasoning. Cambridge, Mass, Harvard University Press, 1978

14. Kassirer JP, Gorry GA: Clinical problem solving: A behavioral analysis. Ann Intern Med 1978; 89:245-255

15. Campbell AN, Hollister VF, Duda RO, et al: Recognition of a hidden mineral deposit by an artificial intelligence program. Science 1982; 217:927-929

16. Duda RO, Shortliffe EH: Expert systems research. Science 1983; 220:261-268

17. Aikins JS, Kunz JC, Shortliffe EH, et al: PUFF: An expert system for interpretation of pulmonary function data. Comput Biomed Res 1983; 16:199-208

18. Weiss SM, Kulikowski CA, Galen RS: Developing microprocessor-based expert models for instrument interpretation, *In* Proceedings of the 7th International Joint Conference on Artificial Intelligence. Los Altos, Calif, Morgan-Kaufmann, 1981, pp 853-855

19. Hickam DH, Shortliffe EH, Carlson RW, et al: The treatment advice of a computer-based cancer chemotherapy protocol advisor. Ann Intern Med 1985; 103:928-936

20. Pryor RA, Garner RM, Clayton PD, et al: The HELP system. J Med Syst 1983; 7:87-102

21. Winston PH: Artificial Intelligence, 2nd Ed. Reading, Mass, Addison-Wesley, 1984

22. Charniak E, McDermott D: Introduction to Artificial Intelligence. Reading, Mass, Addison-Wesley, 1985

23. Kanal LN, Lemmer JF (Eds): Uncertainty in Artificial Intelligence. Amsterdam, North-Holland, 1986

24. Shortliffe EH: Reasoning methods in medical consultation systems: Artificial intelligence approaches. Comput Programs Biomed 1984; 18:5-14

25. Barr A, Cohen P, Feigenbaum EA (Eds): The Handbook of Artificial Intelligence. Los Altos, Calif, Morgan-Kaufmann, 1981 (vol 1), 1982 (vols 2 & 3)

26. Davis R: Interactive transfer of expertise: Acquisition of new inference rules. Artif Intell 1979; 12:121-158

27. Patil RS, Szolovits P, Schwartz WB: Causal understanding of patient illness in medical diagnosis, *In* Proceedings of 7th International Joint Conference on Artificial Intelligence. Los Altos, Calif, Morgan-Kaufmann, 1981, pp 893-899

28. Cooper GF: A diagnostic method that uses causal knowledge and linear programming in the application of Bayes' formula. Comput Methods Programs Biomed 1986; 22:223-237

29. Rennels GD, Shortliffe EH, Stockdale FE, et al: Reasoning from the clinical literature: The Roundsman system, *In* Proceedings of MEDINFO-86. Amsterdam, North-Holland, 1986, pp 771-775

30. Teach RL, Shortliffe EH: An analysis of physician attitudes regarding computer-based clinical consultation systems. Comput Biomed Res 1981; 14:542-558

31. Shortliffe EH, Scott AC, Bischoff AC, et al: ONCOCIN: An expert system for oncology protocol management, *In* Proceedings of 7th International Joint Conference on Artificial Intelligence. Los Altos, Calif, Morgan-Kaufmann, 1981, pp 876-881

32. Bischoff MB, Shortliffe EH, Scott AC, et al: Integration of a computer-based consultant into the clinical setting, *In* Proceedings of the 7th Symposium on Computer Applications in Medical Care. Washington, DC, IEEE Computer Society Press, 1983, pp 149-152

33. Kent DL, Shortliffe EH, Bischoff MB, et al: Improvements in data collection through physician use of a computer-based chemotherapy treatment consultant. J Clin Oncol 1985; 3:1409-1417

34. Lane CD, Walton JD, Shortliffe EH: Graphical access to a medical expert system—Part II: Design of an interface for physicians. Methods Inf Med 1986; 25:143-150

35. Musen MA, Fagan LM, Combs DM, et al: Facilitating knowledge entry for an oncology therapy advisor using a model of the application area, *In* Proceedings of MEDINFO-86. Amsterdam, North-Holland, 1986, pp 46-50

36. Langlotz CP, Fagan LM, Tu SW, et al: Combining artificial intelligence and decision analysis for automated therapy planning assistance, *In* Proceedings of MEDINFO-86. Amsterdam, North-Holland, 1986, pp 794-798