Knowledge bases in medicine: a review

By Claudia A. Perry, M.L.S.
Head, Reference Department

The New York Academy of Medicine Library
2 East 103d Street
New York, New York 10029

Efforts to represent knowledge effectively have been central to progress in various aspects of medical informatics. These efforts range from relatively simple "electronic textbooks" to fairly sophisticated knowledge-based systems, which function as well as, or even better than, human experts faced with similar problems. Knowledge bases have been developed in many fields, but the relatively limited domains and structured language of medicine, as well as the importance of information in the provision of good medical care, have made research in medical knowledge representation an area of intense activity. This paper reviews representative knowledge bases and knowledge-based systems in medicine: electronic textbooks such as PDQ* and the Hepatitis Knowledge Base (HKB), rule-based systems such as MYCIN, causal models (e.g., CASNET), and hypothesis- or frame-based systems, exemplified by PIP and INTERNIST-1. The paper describes the relationships among divergent approaches and provides a sense of current and future trends. It examines problems in knowledge-based systems, particularly in knowledge representation and acquisition, and the responses to these challenges. The latter include the use of domain-independent software shells for constructing knowledge bases, the adaptation and use of previously existing knowledge bases, and multiple uses of the same knowledge base for different purposes.

The information explosion in science has resulted in multiple efforts to control the increasingly massive continuum of data/information/knowledge. This is particularly true in the field of medicine. The ability of computers to manipulate large sets of symbols has encouraged developments in information retrieval, hospital records management, computer-assisted instruction, computer-assisted decision making, and expert systems. While these systems differ in their purposes and goals, as well as the techniques used to achieve those aims, a common requirement is the need to access a large base of information or knowledge, however knowledge may be defined.

The concept of the knowledge base plays a central role in these diverse efforts to exert control over the information environment. This parallels an important shift in artificial intelligence (AI) research. Early AI efforts focused on the search for procedures that could be combined to make intelligent programs. More recently, Goldstein and Papert asserted that "the fundamental problem of understanding intelligence is not the identification of a few powerful techniques, but rather the question of how to represent large amounts of knowledge in a fashion that permits their effective use and interaction" [1].

Efforts to represent knowledge effectively have been central to progress in various aspects of medical informatics. These efforts range from relatively simple "electronic textbooks" to fairly sophisticated expert systems, which function as well as, or even better than, human experts faced with similar problems. Knowledge bases have been developed in many fields, but the relatively limited domains and structured language of medicine, as well as the importance of knowledge in the provision of good medical care, have made research in medical knowledge representation an area of intense activity. This paper discusses representative knowledge bases in medicine in order to provide an understanding of the nature of the field, the relationships among divergent approaches, and a sense of current and future trends. Particular attention is given to consideration of the problems associated with various knowledge bases and the impact of such problems on the direction of research efforts.

* PDQ is a registered trademark of the National Cancer Institute.

Since the field is extensive and prototypes have proliferated, well-established systems or large-scale projects are emphasized over smaller projects in the initial development stages, except as the latter illustrate new trends and responses to problems. No attempt is made to be comprehensive.

Knowledge bases have been developed in many fields, but the relatively limited domains and structured language of medicine, as well as the importance of knowledge in the provision of good medical care, have made research in medical knowledge representation an area of intense activity.

ELECTRONIC TEXTBOOKS

Despite the successes of automated literature retrieval through systems such as MEDLARS†, it has long been recognized that an important drawback to the effective use of medical information is the need to sift through vast amounts of possibly irrelevant or unreliable material in order to obtain and use needed information. The practice of medicine depends on the application of large stores of knowledge, much of it learned in medical school and in clinical training, but much also absorbed through a process of lifetime learning. Clinicians depend in large part on memory to assist in retrieving knowledge relevant to the case of a particular patient, but their personal knowledge base cannot be static; clinicians must also stay abreast of recent developments in medicine.

The problems associated with the information explosion and the transmission of information were carefully examined by Bernstein et al. in a paper that presented a useful framework for consideration of the potential importance of an artificial knowledge base [2]. The main portion of their paper was a detailed description of a prototype information-transfer system developed by the National Library of Medicine (NLM). The Hepatitis Knowledge Base (HKB) was an effort to assemble a comprehensive collection of current but previously published information reflecting the consensus of experts. This information could be accessed online and updated regularly. Bernstein and his colleagues discussed the construction of the knowledge base through the use of review articles as quality filters; the methods, including computer conferencing, for developing consensus among a panel of experts; and approaches used for updating the system. The knowledge base in this case was intended to synthesize and disseminate existing knowledge and to make information passively available for the practitioner's use "as appropriate for the particular combination of patient, illness, setting, and other circumstances confronted" [3]. While multiple derivative products and services were envisioned for HKB, its use as an active medical consultant was not suggested.

Clinicians depend in large part on memory to assist in retrieving knowledge relevant to the case of a particular patient, but their personal knowledge base cannot be static; clinicians must also stay abreast of recent developments in medicine.

HKB was not extended past the prototype stage, but a very similar product from the National Cancer Institute used many of the same principles in promoting diffusion of information concerning cancer treatment. Like HKB, the Physician Data Query (PDQ) system was assembled by a panel of experts, but continues to be updated regularly, and is available online through a number of vendors [4]. Both systems might be considered to be examples of "electronic textbooks," providing the same sort of authoritative information found in standard medical textbooks, but in a very different and more current format. Although the content may not be radically different from traditional print sources, the electronic format permits timeliness, ongoing modification, and expansion of features and coverage; it also makes possible a very different type of access and use. Unlike HKB, PDQ does not rely solely on previously published material, but provides access to investigational protocols and research in progress. System flexibility has allowed the developers of PDQ to incorporate additional features over time, such as the abstracts of selected cited references and strategies for searching related bibliographic databases.

Very few studies have been published evaluating the use of PDQ by physicians. However, the system has been widely available for direct menu-driven access by physicians through several online vendors (NLM, BRS, MEDIS, TELMED) for a number of years, as well as indirectly through local cancer information services and libraries. It is also available from a number of vendors in CD-ROM format. An organizational meeting of a PDQ User Group at the 1988 meeting of the Medical Library Association (MLA) drew a large crowd of interested attendees, while the mailing list of the newly instituted *PDQ User Group Letter* numbers over 200 [5]. PDQ has clearly advanced far beyond the prototype phase.

Both PDQ and the HKB were designed to collect, screen, and repackage relevant, current, and accurate

† MEDLARS is a registered trademark of the National Library of Medicine.

information for delivery upon request, but the practical difficulties of such a project can be daunting. Bernstein et al. commented that

were additional knowledge bases to be developed, only those disease areas should be selected in which there is high morbidity and mortality, large numbers of physician-patient contacts, high costs of medical care, and large numbers of research publications that make difficult the effective and efficient synthesis of new information for physician-practitioners [6].

It is not surprising that some recent commercially available electronic textbooks (e.g., AIDS Knowledge Base) are also concerned with diseases, like Acquired Immunodeficiency Syndrome, that meet such criteria. Still other electronic textbooks are merely the online equivalents of standard print sources, such as *Merck Index* [7] or Martindale's *Extra Pharmacopeia* [8]. Advances in information technology and resulting changes in publishing practices, however, have made the creation and maintenance of electronic textbooks far more efficient and effective than might have been envisioned even a short time ago. These changes are making the electronic textbook an increasingly common type of knowledge base. Online availability permits more regular updating so that recent research can be incorporated in a manner not readily possible with a printed source. As a result, electronic textbooks can be as selective as their print counterparts, but can provide more current information than would be possible in a standard printed work.

Still other changes are related to medium. The increasing popularity of CD-ROM technology is further expanding the potential of electronic textbooks. A number of knowledge bases available through online vendors, such as BRS or NLM, are now also produced in CD-ROM format, frequently in combination with other products. The quarterly "Compact Library: AIDS"§ produced by the Massachusetts Medical Society's Medical Publishing Group includes the AIDS Knowledge Base, an AIDS subset of the bibliographic data in MEDLINE, and the full text of many of the cited journal references. Lippincott's "OncoDisc"** includes a CD-ROM version of PDQ and NLM's bibliographic file CANCERLIT††, as well as the full text of a number of print oncology textbooks. Still other products sidestep online access altogether, moving directly from print to CD-ROM availability. One such example is Little Brown's "MAXX: Maximum Access to Diagnosis and Therapy," an interactive electronic library of some fifteen handbooks and manuals in clinical medicine. While most such products are updated only quarterly and may not contain any more current information than their print counterparts, the potential for innovative approaches to accessing the information is impressive. Separate knowledge bases may be searched concurrently, linked in a hypertext mode, or displayed simultaneously on screen. Graphics capabilities are also improving dramatically. One anticipates that high-quality illustrations will soon become a common feature in CD-ROM electronic textbook products.

While the aims of electronic textbooks, whether online or CD-ROM-based, are more limited than those of knowledge-base projects dealing with decision support, their importance lies in demonstrating the feasibility of the creation of large-scale, regularly updated knowledge bases:

The knowledge-base program principle is that a centrally coordinated process using outside content experts can perform the difficult, laborious, expensive, technical, and intellectual tasks needed to assemble a knowledge base. Once assembled, it has the potential to be easily, inexpensively, and widely exploited to serve the information needs of very large numbers of users [9].

Clearly, the PDQ knowledge base and similar systems in the commercial sector are convincing demonstrations of the workability of this concept. Knowledge-based systems, a more ambitious effort, also use a knowledge base component, but are explicitly designed with ultimate use in mind.

KNOWLEDGE-BASED SYSTEMS

Knowledge-based or expert systems are computer programs designed to solve problems, generate new information (such as a diagnosis), or provide advice, using a knowledge base and an inference mechanism. Most systems include a user interface and some explanation capability as well. Davis characterized knowledge-based systems as focusing "on the accumulation, representation, and use of knowledge specific to a particular task," but addressed the expanded views of such systems made possible by "the ability to use the same knowledge in several different ways" [10], much as was envisioned by the creators of HKB. The potential and problems associated with multiple users of the same knowledge base are discussed further in a later section.

Duda and Shortliffe differentiated between "knowledge-based systems" and "expert systems," commenting that for many applications there may not be any uniquely qualified human experts, so that it is inappropriate to speak of the development of programs approximating the level of experts. They defined

§ Compact Library: AIDS is a registered trademark of the Massachusetts Medical Society.

** Oncodisc is a registered trademark of J. B. Lippincott Company.

†† CANCERLIT is a registered trademark of the National Library of Medicine.

a knowledge-based system as an AI program whose performance depends more on the explicit presence of a large body of knowledge than on the possession of ingenious computational procedures; by expert system we mean a knowledge-based system whose performance is intended to rival that of human experts [11].

In all cases, knowledge-based systems are considered to be distinct from programs based primarily on mathematical models, statistical techniques, or pattern matching, although some knowledge-based systems may incorporate statistical components along with knowledge bases in an effort to account for uncertainty [12–16]. The knowledge base is clearly the critical and distinguishing element characterizing these decision-support systems.

Knowledge-based or expert systems are computer programs designed to solve problems, generate new information (such as a diagnosis), or provide advice, using a knowledge base and an inference mechanism. Most systems include a user interface and some explanation capability as well.

Knowledge-based systems have a far more ambitious agenda than simply to gather information together for convenient access. In the domain of medicine, programs have been developed to assist with differential diagnosis, suggest particular therapeutic interventions, critique a physician's plan for the treatment of disease, or tutor medical students. Depending on the area of specialization and purpose, different methods of knowledge representation have been deemed appropriate.

Knowledge representation

A variety of different types of knowledge need to be represented in a knowledge-based system. A common distinction is between "what to know" and "what to do," embodied in the knowledge base and the inference mechanism respectively, although in some approaches this line may be blurred [17]. The knowledge base itself can encompass knowledge about facts or relationships drawn from the analysis of data, theoretical concepts, common sense, or world knowledge—in addition to knowledge about solving the specific problem at hand; it may involve heuristics, or "rules-of-thumb" based on experience, as well as expert judgement [18–19]. This expert knowledge has been represented with a variety of schemes: production rules [20], causal-associational networks [21], and "frames" of closely related facts [22–23].

Rule-based systems

One of the best known knowledge-based systems is MYCIN, which employs a rule-based method of knowledge representation to assist in the identification of disease-causing microorganisms and then to advise physician-users concerning appropriate antimicrobial therapy [24–25]. MYCIN encodes its knowledge in approximately 500 rules based on an IF–THEN structure. It permits the user to ask why or how a given decision was determined and provides a mechanism for explanation. At least three evaluation studies of MYCIN have been conducted, resulting in performance approval ratings of just under 75% when compared with experts [26]. More importantly, work on evaluation has provided insight into ways to improve and enhance the system.

A domain-independent version of MYCIN, called EMYCIN, has permitted development of similarly structured knowledge-based systems in other subject areas both inside and outside of medicine. Examples of systems developed using EMYCIN include BLUEBOX and HEADMED, two psychopharmacology advisers [27–28], and PUFF, a program that diagnoses lung diseases [29]. ONCOCIN§§, another MYCIN descendant, assists oncologists with the treatment of cancer patients as part of the routine patient data-management process. ONCOCIN is based on EMYCIN but incorporates changes designed to improve the user interface and speed of interaction [30–31].

Largely due to the success of MYCIN, rule-based systems continue to evolve. The prevalence of rule-based systems reflects Brooks' observation concerning the ease of developing new programs given the existence of a system with related aims.

> If an expert system exists in a similar domain, then the new system can draw on the expertise and skills acquired during the development of the earlier one: the knowledge representation used, methods of inference and control, and overall architecture [32].

Clearly, it is easier to adapt existing frameworks than to begin the process entirely from the beginning. At the same time, many system developers have found a rule-based approach to be suitable. Brooks and Heiser reviewed the advantages of rule-based systems over other systems, including ease of updating through the inclusion of additional rules and ability to respond to the availability of new data [33]. Others, however, have commented that authors of new rules must anticipate the ways in which additions will interact with previously existing rules, and that this may be quite difficult [34]. Brooks and Heiser also

§§ ONCOCIN is a registered trademark of the Board of Trustees of Leland Stanford, Jr., University.

commented on the appropriateness of rule-based systems to domains, such as clinical psychopharmacology, which lack deep theories as to the mechanism of disease. This does not mean, however, that other schemes of knowledge representation have been unsuccessful.

Causal models

A causal-associational network (CASNET) model is considered to be an appropriate vehicle for representing knowledge in areas such as glaucoma, where the causal model of disease is well understood [35-36]. The CASNET model is a particular type of semantic net, a knowledge structure consisting of "*nodes* to represent the concepts, events, characteristics, and values of interest in a system, as well as branches specifying the relationships between nodes" [37]. In the CASNET approach, three levels of description—such as observation, pathophysiological state, and disease category—are used to formulate recommendations for treatment.

A physician hypothesizes as to what may be the nature of a particular illness based on the presence of specific symptoms and prescribes treatment based on what has worked for that diagnosis in the past. The physician then evaluates the success of the diagnosis based on the reaction to therapy.

> As observations are recorded, they are associated with the appropriate states. States are causally related, forming a network that summarizes the mechanisms of disease. Patterns of states in the network are linked to individual disease classifications [38].

Treatment is then recommended based on both the broad disease classification and the specific patient observations.

The CASNET glaucoma consultation model has performed well in evaluations by independent groups of ophthalmologists. Experience with its use showed that

> although we could indeed arrive at a correct set of diagnostic and prognostic conclusions on the basis of strictly causal reasoning, there was need for a more flexible and general scheme of reasoning that would permit the introduction of empirical knowledge where knowledge of disease mechanisms is absent [39].

CASNET has thus been extended and generalized to form the system-building software EXPERT, which can incorporate empirical knowledge along with causal reasoning. AI/RHEUM, a consultation tool developed using EXPERT, has been evaluated extensively for diagnostic accuracy. In a comparison with the diagnoses of Japanese rheumatologists, the program was in full or partial agreement with Japanese decisions in over 90% of the fifty-nine cases evaluated [40]. Despite its CASNET/EXPERT ancestry, AI/RHEUM is essentially a rule-based system and illustrates the increasing tendency of developers of knowledge-based systems to use converging approaches.

Hypothesis-based systems

Still other system developers seek to approximate the reasoning processes physicians engage in when making medical decisions. Bouckaert commented that medical decision making is very similar to scientific hypothesis-making and testing [41]. A physician hypothesizes as to what may be the nature of a particular illness based on the presence of specific symptoms and prescribes treatment based on what has worked for that diagnosis in the past. The physician then evaluates the success of the diagnosis based on the reaction to therapy. Failure to respond requires a reexamination of the initial hypothesis, additional data gathering, and the formation of new hypotheses. The developers of the Present Illness Program (PIP) used just such a conceptual approach in designing a program for diagnosis of edema. Knowledge in PIP is organized into packets of related information called frames. "Within each frame is a rich knowledge structure which includes prototypical findings (signs, symptoms, laboratory data), the time course of the given illness, and rules for judging how closely a given patient might 'match' the disease or state which the frame describes" [42]. Frames are further organized into networks, and include not only information on approximately twenty different diseases, but commonsense knowledge as well. The program combines patient-specific knowledge with the knowledge base in order to develop and test hypotheses in an iterative process. If an hypothesis does not fit, additional information is then sought. Members of the research group have used a similar approach in designing a program to determine appropriate dosages of digitalis [43] and in creating ABEL, a program for the management of acid-base and electrolyte disturbances [44].

Most of the systems described thus far have dealt with discrete subject areas: antimicrobial therapy, the diagnosis of glaucoma, or the management of electrolyte disorders. Given the difficulties of representing "large amounts of knowledge in ways that still allow the effective use of individual facts" [45], most knowledge bases have dealt with limited task domains. INTERNIST-1, however, is an hypothesis-based system using a knowledge base encompassing all of internal medicine. Like PIP and ABEL, INTERNIST-1 represents an effort to model the decision-making behavior of physicians, but its reach is far wider than most knowledge-based systems [46]. As of June 1986, the INTERNIST-1 knowledge base contained 572 diagnoses and recognized over 4,000 patient observations [47].

INTERNIST-1 has been extensively evaluated using published clinicopathological conferences (CPCs—case records of the Massachusetts General Hospital), and has made correct diagnoses for 65% to 75% of the test cases within its scope. The program has been continually modified based on this feedback, with subsequent generations of the system titled INTERNIST-1/CADUCEUS and INTERNIST-1/QUICK MEDICAL REFERENCE (QMR). QMR***, a microcomputer-based version of INTERNIST-1, is designed for practical use by the clinician and corrects many of the perceived deficiencies, such as overlong interactions, which inhibited regular use of prior versions of the system.

PROBLEMS AND RESPONSES TO PROBLEMS IN KNOWLEDGE-BASED SYSTEMS

Problems with knowledge representation schemes

The testing and evaluation of the hypothesis-based INTERNIST-1 in its various forms revealed a number of problems about which its creators have written extensively. They discussed its inadequate representation of pathophysiologic causality, inability to consider the interdependence of manifestations or degrees of severity of findings and diseases, and lack of temporal and anatomical reasoning [48–50]. Of these problems, the most serious is considered to be its occasional inability to attribute findings to their proper causes.

The importance of an explanation capability is assumed by numerous knowledge base developers, who anticipate that physicians will be unlikely to accept blindly decisions they cannot systematically follow and understand [51–53]. This assumption is supported by a survey of physicians' attitudes toward computers, which found that a program's ability to explain diagnostic and treatment decisions to physician users is considered to be its most important attribute [54]. Erdman cautioned that the availability of explanations may lead not only to increased system acceptance, but also to increased confidence in potentially erroneous judgements. His study, which yielded a limited number of significant results, is based on a nonrandom sample of only forty-eight physicians; he makes a useful point in asserting that the benefits of explanation must be demonstrated and not just assumed [55].

User level may be an important factor in deciding how knowledge is to be represented. Research in the adaptation of knowledge bases for pedagogical purposes has, in fact, found that different uses may substantially affect the way a knowledge base must be structured.

Unlike INTERNIST-1, the rule-based MYCIN "was not intended to simulate human problem solving in any formal way" [56], although "it was among the first systems to emphasize the importance of explanation capabilities in medical decision-support tools" [57]. Explanation in MYCIN consists of its ability to respond to a question of why or how a current rule has been invoked or additional information sought, with the premises and action of the rule itself. In this way, MYCIN can reconstruct its reasoning process in arriving at a decision. Fieschi and Joubert commented that, at times, this can yield only a superficial explanation of decisions, since many rules in MYCIN consist largely of heuristics and do not involve a deeper understanding of the disease. In some cases, empirical or associational knowledge may be all that is available; a system cannot embody causal knowledge if none exists [58]. At other times, however, there may be a deeper explanation behind a rule that has not been specifically encoded for the sake of efficiency. Experienced clinicians may not need such detailed explanation, but it may be critical for the understanding of novices. Thus, the user level may be an important factor in deciding how knowledge is to be represented. Research in the adaptation of knowledge bases for pedagogical purposes has, in fact, found that different uses may substantially affect the way a knowledge base must be structured [59].

Another criticism of rule-based systems is that the order in which rules are invoked may be important [60–61]. The ordering of rules may implicitly affect their meaning and interaction, and yet this may not be made clear to a user seeking explanation. In other cases, the order may not be important, but an inflexible system may assert that a different approach is incorrect if the system was programmed to process the rules in a particular order. Rules also may be confusing when examined as part of an explanation, since they are used both to represent knowledge as well as to indicate the way that knowledge is to be used. Lack of consistency in rule structure thus may interfere with understanding of their function. Finally, rules depend to some degree on simplifying assumptions and cannot necessarily handle compli-

*** QMR is a registered trademark of the University of Pittsburgh.

cated manifestations of more than one disease state. This limits their use to specific domains. The problem also exists in other knowledge-based systems such as INTERNIST-1, which "cannot synthesize a general overview in complicated multisystem problems" [62].

While much of this criticism centers around a lack of causal explanatory power, a problem with causal-based systems such as CASNET is that often the pathophysiology of a particular disease is not well understood. In such cases the causal approach is not appropriate. Moreover, Fieschi and Joubert commented that CASNET has been implemented on applications where illness is progressive and where symptoms evolve over time, while early symptoms remain. Such an approach might be less successful with illnesses that do not follow such a characteristic course. It has already been mentioned that EXPERT, CASNET's descendant, has been modified to incorporate empirical findings along with a causal model so that performance and generality can be enhanced. This multifaceted approach has been adopted by researchers throughout the field.

Responses to problems of knowledge representation

Much of medical practice is grounded in empirical knowledge: what works or doesn't work in treating illness. This knowledge may constitute merely an awareness of the significance of an association of particular signs and symptoms, or it may be based on an understanding of the underlying mechanism of disease. Aikins comments that "much of artificial intelligence research has focused on determining the appropriate knowledge representations to use in order to achieve high performance from knowledge-based systems." She believes that there are numerous advantages to using "both framelike structures and rules to solve problems in knowledge-intensive domains" [63]. Aikins described CENTAUR, a descendant of the rule-based system PUFF, which used a combination of frames and rules to overcome problems associated with PUFF. A combination of techniques permitted designers to separate rules concerning expertise from instructions concerned with computation. This encouraged explicit representation of knowledge in data structures called prototypes. CENTAUR is still in the developmental stage, but is an example of efforts to compensate for perceived deficiencies of a strictly rule-based structure. One goal is presumably to enhance understanding of the process by which decisions are made. Such understanding can be as important to the effectiveness of knowledge-based systems as is high performance, as has already been mentioned.

Wallis and Shortliffe reported on efforts to model different types of users and to create explanation functions appropriate to different user levels [64]. They used a semantic network conceptualization of knowledge representation for this prototype, although their previous work had been largely in MYCIN and similar rule-based systems. Like Aikins' incorporation of frames with rules, and EXPERT's use of empirical along with causal knowledge, this work by Wallis and Shortliffe illustrates the trend towards heterogeneous approaches and multiple levels of representation in the development of new knowledge-based systems. Such flexibility in design can be used to correct for specific problems and to increase a system's adaptability. Like human experts, the ability of knowledge-based systems to use one approach in some situations and an alternative, possibly more complex, one in others can improve performance substantially [65].

Much of medical practice is grounded in empirical knowledge: what works or doesn't work in treating illness. This knowledge may constitute merely an awareness of the significance of an association of particular signs and symptoms, or it may be based on an understanding of the underlying mechanism of disease.

Knowledge acquisition

Experts agree that knowledge acquisition and verification are among the most difficult problems connected with the development of knowledge-based systems [66–69]. Duda and Shortliffe asserted that

> the identification and encoding of knowledge is one of the most complex and arduous tasks encountered in the construction of an expert system. The very attempt to build a knowledge base often discloses gaps in our understanding of the subject domain and weaknesses in available representation techniques [70].

Construction of knowledge bases for systems like INTERNIST-1 is measured in person-years and frequently involves large teams of experts. Traditionally, a knowledge engineer, or expert in the technical aspects of a knowledge-based system, works with a human subject expert or several experts. For knowledge transfer and organization to be effective, it helps if each of the two types of experts has some understanding of the other's domain. Spackman and Connelly described four major approaches to knowledge acquisition:

> 1) unstructured interviews with a human expert; 2) observation of the human expert thinking aloud while solving a problem . . . ; 3) interaction with special-purpose com-

puter programs to debug or modify an existing knowledge base; and 4) automated construction of a knowledge base by machine learning programs, which learn from examples of correctly solved problems. These approaches can be used separately or in combination [71].

The acquisition of knowledge from human experts is particularly tedious, while the structuring of the knowledge obtained becomes more and more difficult as the domain of the knowledge base is extended. These difficulties have encouraged efforts to automate the process as much as is feasible.

The acquisition of knowledge from human experts is particularly tedious, while the structuring of the knowledge obtained becomes more and more difficult as the domain of the knowledge base is extended.

The increasing availability of domain-independent "software shells" derived from well-established programs promises to ease the knowledge acquisition process in some applications. "These 'shells' are essentially functional systems stripped of their knowledge base, and they can be used as 'knowledge engineering' tools to develop prototype expert systems" [72]. Mars and Miller referred to similar programs as Knowledge Acquisition (KA) and Knowledge Verification (KV) tools and described a number of examples [73]. HEADMED and BLUEBOX, which psychopharmacology advisers developed using the software shell EMYCIN, demonstrate the potential of such tools to assist in the development of useful knowledge bases by clinicians in relatively short periods of time [74-75]. Similarly, EXPERT, a system-building software program derived from CASNET, aided in the construction of AI/RHEUM [76].

Software shells are most useful for building small knowledge-based systems. The domain of a system such as INTERNIST-1 is so large, however, that it would be impossible for a small group to develop and maintain the knowledge base. Instead, the INTERNIST-1 project developed its own methods of knowledge base construction, which are not unlike the procedures used by the HKB project. These involve literature surveys and consensus building among project team members, followed by extensive testing on actual cases, and regular updating and modification. The process is extremely labor intensive.

Use of existing knowledge bases

Difficulties involved in constructing knowledge bases are such that efforts have been made to use previously existing knowledge bases. Systems using such ready-made knowledge sources include DXplain††† and RECONSIDER (based on *Current Medical Information and Terminology*—CMIT), as well as Pulmonologist, which uses the *Merck Manual* as its knowledge base [77–79]. Research is also being conducted on the integration of CMIT with NLM's MeSH, used in indexing MEDLINE. This should permit the creation of a richer knowledge base than would be available through either knowledge collection alone [80]. Adaptations of thesauri for use in knowledge-based systems show promise in fields where well-structured vocabularies exist. For many applications where no such thesauri are available, however, construction of knowledge bases will remain on unavoidable challenge.

Multiple uses of knowledge bases

Another approach to increasing the return on the effort put into knowledge base construction is to use a knowledge base for more than one purpose. Given the explanation capabilities of many knowledge-based systems designed for decision support, their extended use in teaching students seems like a natural progression. In practice, adapting decision support systems for tutoring is more complicated than it seemed. GUIDON and its successor NEOMYCIN, both based on the MYCIN knowledge base, for example, have required substantially more knowledge than what is contained in MYCIN [81]. MYCIN's rules are frequently associational and do not explicitly incorporate causal reasoning in their heuristics. A system designed to teach students must incorporate such knowledge directly. Moreover, an effective tutoring system must also incorporate teaching expertise. One anticipates that other unplanned uses for a particular knowledge base will also require additional modifications for effective use.

Lower-level uses of a knowledge base do permit direct use without necessary adaptations, however. INTERNIST-1/QMR can be used as an information retrieval tool or "electronic textbook" in much the same manner as PDQ or HKB [82]. Schwartz et al., for example, commented that

> the enormous database of INTERNIST may provide a considerable advantage over textbooks to the physician who is searching for facts about a particular illness. . . . The INTERNIST database has the advantage . . . of including numerical information on the frequency of findings and on the diagnostic importance of each finding [83].

These authors also see value in using knowledge-based systems for less ambitious goals than perhaps

††† DXplain is a registered trademark of the Massachusetts General Hospital.

originally intended, as, for example, to critique diagnostic hypotheses or plans for treatment.

Additional problems with knowledge-based systems

Schwartz et al. see numerous problems with knowledge-based systems in medicine and feel that "we have not reached the point at which artificial-intelligence programs can act as reliable consultants on a wide range of problems" [84]. It is certainly true that these systems have not achieved widespread use, with few exceptions. Shortliffe reported that by 1985 only three medical expert systems, including PUFF and ONCOCIN, were being used routinely in clinical settings [85]. Given the ongoing nature of research on these systems, one wonders how routine their use has been. RECONSIDER has recently been made available for student use at Georgetown University as part of their Integrated Academic Information Management System, but it is yet too early to evaluate the degree of use [86].

A number of reminder systems based on patient record databases have been reasonably successful, but one could argue that these are less knowledge-based systems than they are management information systems. Health Evaluation Through Logical Processing (HELP) has been operational at the LDS Hospital in Salt Lake City for fifteen years, where it provides surveillance alerts for patients with hospital-acquired infections and antibiotic use [87]. A computerized medical record system (CARE) at the Regenstrief Institute for Health Care in Indiana generates reminders to physicians of patient conditions needing attention [88]. COSTAR§§§, a computer-based medical information system at Massachusetts General Hospital, also provides feedback to physicians concerning deficiencies in patient care [89]. It is no doubt relevant to their success that all three of these reminder systems are integrated into the normal routine of hospital operations.

Lack of use of knowledge-based systems is related not only to problems with explanation and knowledge representation, but also to ease of access, the nature of the user interface, and the time-consuming nature of interactions with many systems. Computer literacy, or its lack, is another important factor. A number of authors have argued that unless use of a knowledge base becomes part of regular working procedures, use will be negligible [90–91]. Implementation of several systems on personal computers may help to counteract access problems. Tsuji and Shortliffe anticipated that the development of touch screens and graphic interfaces, such as those popularized by the Macintosh computer, may ease difficulties related to health professionals' lack of typing experience [92]. Younger physicians, increasingly exposed to computer applications in medical school, may also be more accepting of innovations. While all these developments are undoubtedly important, it may also be time to revise the goals for knowledge-based systems.

Schwartz, Miller, and their respective colleagues made a good case for trying to accomplish more with less [93–94]. This may be important for both technical and psychological reasons. Knowledge-based systems have not been nearly as successful, from a technical standpoint, in their role of providing expert advice to health professionals as had been hoped. But psychological factors are also important. Controversy exists over the possible displacement of physicians by expert systems, a controversy reminiscent of similar fears by librarians with the coming of end-user searching. Recent knowledge-based systems attempt to take a less directive role than did early prototypes, while still providing physicians with appropriate decision support. Some, like QMR, are also more flexible in their ability to provide more or less advice to physicians, as desired.

Knowledge-based systems have not been nearly as successful, from a technical standpoint, in their role of providing expert advice to health professionals as had been hoped.

Miller's QMR is a modification of INTERNIST, the aim of which is "to meet the practical needs of clinical medicine" [95]. QMR is implemented using an IBM compatible PC/AT and can function at varying levels of sophistication, as already mentioned. P. L. Miller has developed a system, ESSENTIAL-ATTENDING, that critiques a plan for patient care, rather than advising the user as to what to do [96]. Yet another system, RECONSIDER, provides a list of differential diagnoses for a physician's consideration [97]. It is one of the few programs that has undergone a comparative evaluation with similar systems (PIP and INTERNIST). RECONSIDER is based on the potential of the "library" function of the computer—that is to say, its capacity to remind the user not to forget a particular possibility, much akin to the premise behind CARE, COSTAR, and HELP. Similarly, DXplain, available through the American Medical Association's AMAnet****, suggests reasonable diagnoses that should be considered given a particular set of

§§§ COSTAR is a registered trademark of the Massachusetts General Hospital.

**** AMAnet is a registered trademark of the American Medical Association.

symptoms [98]. All these systems assume a far less active role than was envisioned for PIP or ONCOCIN, but if they are actually used by practicing physicians, in the long run they may have a far greater impact than more ambitious efforts to date.

Contributions of knowledge-base development

Despite undeniable problems and lack of success in the sense of routine use, research in knowledge bases has added immeasurably to understanding of the user interface and of knowledge structure and use. Perhaps the greatest contribution of knowledge-base development has been its role in codifying knowledge and encouraging the development of models of the medical decision-making process. As stated previously, "the very attempt to build a knowledge base often discloses gaps in our understanding of the subject domain and weaknesses in available representation techniques" [99]. Development of a knowledge base makes it possible to present the knowledge of a field formally and systematically, in a way in which it might never have been done before. Efforts to keep such knowledge current also may provide insight into the directions of growth and change in a field.

CONCLUSION

Several major themes arise from an examination of recent progress in the development of medical knowledge bases. One is the increasing tendency of knowledge-base developers to learn from one another and to incorporate more than one approach to knowledge representation in devising knowledge bases. This permits such systems to use different problem-solving approaches as needed in different situations, in a manner not unlike that of human problem solving. The availability of software shells has also reduced the tedium involved in knowledge acquisition, structuring, and verification. This may encourage the future creation and revision of knowledge bases so that developers can focus attention on persistent problems.

A second major point is that by taking a less ambitious course, knowledge bases may be more successful in a practical sense. Knowledge bases that are not used and that remain only as research prototypes are of limited value; only through use will they contribute to our understanding of knowledge structure and use and its role in the decision-making process. Neither will unused databases make possible the provision of improved medical care, which is presumably the ultimate goal. Evaluation studies of such widely available knowledge bases and knowledge-based systems as PDQ, RECONSIDER, and DXplain are needed to determine if these types of knowledge systems do, in fact, make a difference in patient care. The rapid expansion in the availability of electronic textbooks such as "Compact Library: AIDS" signals extensive investment by commercial firms in the development of online and CD-ROM-based knowledge bases. Ideally, this flurry of activity should result in a greater variety of approaches to knowledge representation and access. Clearly, given the difficulties of knowledge-base creation and maintenance, wide-scale implementation and use will be necessary if knowledge bases are to meet the goals of "effective use and interaction" envisioned by Goldstein and Papert [100].

REFERENCES

1. Goldstein I, Papert S. Artificial intelligence, language, and the study of knowledge. Cogn Sci 1977;1(1):84–123.
2. Bernstein LM, Siegel ER, Goldstein CM. The Hepatitis Knowledge Base: a prototype information transfer system. Ann Intern Med 1980 Jul;93(1 part 2):167–81.
3. Ibid., 179.
4. Hubbard SM, Henney JE, Devita VT. A computer data base for information on cancer treatment. N Engl J Med 1987 Feb 5;316(6):315–8.
5. Ricks J. Personal communication. National Cancer Institute. 1988 Dec 19.
6. Bernstein, op. cit., 180.
7. Windholz M, ed. Merck index. 11th ed. Rahway, NJ: Merck, 1989.
8. Reynolds JEF, ed. Martindale: the extra pharmacopeia. 29th ed. London: Pharmaceutical Press, 1989.
9. Bernstein, op. cit., 179.
10. Davis R. Knowledge-based systems. Science 1986 Feb 26;231(4741):957–63.
11. Duda RO, Shortliffe EH. Expert systems research. Science 1983 Apr 15;220(4594):261–8.
12. Ibid.
13. Holman JG, Cookson MJ. Expert systems for medical applications. J Med Eng Technol 1987 Jul/Aug;11(4): 151–9.
14. Shortliffe EH. Medical expert systems—knowledge tools for physicians. West J Med 1986 Dec;145(6):830–9.
15. Shortliffe EH. Computer programs to support clinical decision making. JAMA 1987 Jul 3;258(1):61–6.
16. Torasso P. Knowledge based expert systems for medical diagnosis. Stat Med 1985 Jul/Sep;4(3):317–25.
17. Davis, op. cit.
18. Shortliffe, Medical expert.
19. Brooks HM. Expert systems and intelligent information retrieval. Inf Pr Man 1987;23(4):367–82.
20. Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley, 1984.
21. Weiss S, Kulikowski CA, Amarel S, Safir A. A model-based method for computer-aided medical decision-making. Artif Intell 1978;11(1):145–72.
22. Goldstein, op. cit.
23. Pauker SG, Gorry GA, Kassirer JP, Schwartz WB. Towards the simulation of clinical cognition: taking a present illness by computer. Am J Med 1976 Jun;60(7):981–96.

24. BUCHANAN, op. cit.
25. SHORTLIFFE EH, AXLINE SG, BUCHANAN BG, MERIGAN TC ET AL. An artificial intelligence program to advise physicians regarding antimicrobial therapy. Comput Biomed Res 1973; 6(6):544–60.
26. BUCHANAN, op. cit.
27. BROOKS RE, HEISER JF. Transferability of a rule-based control structure to a new knowledge domain. In: Dunn RA, ed. Proceedings of the Third Annual Symposium on Computer Applications in Medical Care. Washington, DC: IEEE Computer Society Press, 1979:56–63.
28. MORELLI RA, BRONZINO JD, GOETHE JW. Expert systems in psychiatry: a review. J Med Syst 1987 Jun;11(2/3):157–68.
29. BUCHANAN, op. cit.
30. IBID.
31. SHORTLIFFE, Medical expert.
32. BROOKS HM, op. cit., 376.
33. BROOKS RE, op. cit.
34. SCHWARTZ WB, PATIL RS, SZOLOVITS P. Artificial intelligence in medicine: where do we stand? N Engl J Med 1987 Mar 12;16(11):685–8.
35. WEISS, op. cit.
36. WEISS S, KULIKOWSKI CA, SAFIR A. Glaucoma consultation by computer. Comp Biol Med 1978;8(1):25–40.
37. SALTON G, MCGILL MJ. Introduction to information retrieval. New York: McGraw-Hill, 1983.
38. WEISS, A model-based, 145.
39. KULIKOWSKI CA, WEISS SM. Representation of expert knowledge for consultation: the CASNET and EXPERT projects. In: Szolovits P, ed. Artificial intelligence in medicine. Boulder, CO: Westview, 1979:21–55.
40. PORTER JF, KINGSLAND LC, III, LINDBERG DAB, SHAH I ET AL. The AI/RHEUM knowledge-based computer consultant system in rheumatology: performance in the diagnosis of 59 connective tissue disease patients from Japan. Arthritis Rheum 1988 Feb;31(2):219–26.
41. BOUCKAERT A. Medical diagnosis: are expert systems needed? Int J Biomed Comput 1987 Jan;20(1–2):123–33.
42. PAUKER, op. cit., 987.
43. GORRY GA, SILVERMAN H, PAULKER SG. Capturing clinical expertise: a computer program that considers clinical responses to digitalis. Am J Med 1978;64(3):452–60.
44. PARIL RS, SZOLOVITS P, SCHWARTZ WB. Modeling knowledge of the patient in acid-base and electrolyte disorders. In: Szolovits P, ed. Artificial intelligence in medicine. Boulder, CO: Westview, 1979:191–226.
45. GOLDSTEIN, op. cit., 87.
46. MILLER RA, POPLE HE, MYERS JD. INTERNIST-1, an experimental computer-based diagnostic consultant for general internal medicine. N Engl J Med 1982 Aug 19;307:468–76.
47. MILLER RA, MCNEIL MA, CHALLINOR SM, MASARIE FE, JR, ET AL. The INTERNIST-1/QUICK MEDICAL REFERENCE project—status report. West J Med 1986 Dec;145(6):816–22.
48. IBID.
49. MILLER, An experimental.
50. MILLER RA. INTERNIST-1/CADUCEUS: problems facing expert consultant programs. Methods Inf Med 1984; 23(1):9–14.
51. BUCHANAN, op. cit.
52. WEISS, Glaucoma consultation, op. cit.
53. WALLIS JW, SHORTLIFFE EH. Explanatory power for medical expert systems: studies in the representation of causal relationships for clinical consultations. Methods Inf Med 1982;21(3):127–36.
54. TEACH RL, SHORTLIFFE EH. An analysis of physicians' attitudes. In: Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley, 1984:635–52.
55. ERDMAN HP. A computer consultation program for primary care physicians: impact of decisionmaking model and explanation capability. Med Care 1987 Dec;25(suppl. 12): S138–47.
56. BUCHANAN, op. cit., 575.
57. SHORTLIFFE, Computer programs, 64.
58. FIESCHI M, JOUBERT M. Some reflections on the evaluation of expert systems in medicine. Methods Inf Med 1986 Jan; 25(1):15–21.
59. CLANCEY WJ. Use of MYCIN's rules for tutoring. In: Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley, 1984:464–89.
60. HOLMAN, op. cit.
61. FIESCHI, op. cit.
62. MILLER RA, INTERNIST-1, 474.
63. AIKINS JS. A representation scheme using both frames and rules. In: Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford Heuristic Programming Project. Reading, MA: Addison-Wesley, 1984:424–40.
64. WALLIS, op. cit.
65. DUDA, op. cit.
66. SHORTLIFFE, Medical expert.
67. TORASSO, op. cit.
68. MILLER, INTERNIST-1/CADUCEUS.
69. CHARD T. Self-learning for a Bayesian knowledge base: how long does it take for the machine to educate itself? Methods Inf Med 1987;26(4):185–8.
70. DUDA, op. cit., 265.
71. SPACKMAN KA, CONNELLY DP. Knowledge-based systems in laboratory medicine and pathology: a review and survey of the field. Arch Pathol Lab Med 1987 Feb;111(2): 116–9.
72. MORELLI, op. cit., 163.
73. MARS NJI, MILLER PL. Knowledge acquisition and verification tools for medical expert systems. Med Decis Making 1987 Jan/Mar;7(1):6–11.
74. BROOKS RE, op. cit.
75. MORELLI, op. cit.
76. PORTER, op. cit.
77. BLOIS MS, TUTTLE MS, SHERERTZ DD. RECONSIDER: a program for generating differential diagnoses. In: Heffernan HG, ed. Proceedings of the Fifth Annual Symposium on Computer Applications in Medical Care. Washington, DC: IEEE Computer Society Press, 1981:263–8.
78. BARNETT GO, CIMINO JJ, HUPP JA, HOFFER EP. DXplain: an evolving diagnostic decision-support system. JAMA 1987 Jul 3;258(1):67–74.
79. KAR A, MILLER GE, SHEPPARD SV. Pulmonologist: a computer-based diagnosis system of pulmonary diseases. Int J Biomed Comput 1987 Nov;21(3–4):223–35.

80. Lester S, Rada R. A method of medical knowledge base augmentation. Methods Inf Med 1987 Jan;26(1):31-9.
81. Clancey, op. cit.
82. Miller, The INTERNIST-1/QUICK.
83. Schwartz, op. cit., 687.
84. Ibid.
85. Shortliffe, Medical expert, 831.
86. Broering NC, Corn M, Ayers WR, Mistry P. Implementing RECONSIDER, a diagnostic prompting computer system, at the Georgetown University Medical Center. Bull Med Libr Assoc 1988 Apr;76(2):155-9.
87. Evans RS, Larsen RA, Burke JB, Gardner RM et al. Computer surveillance of hospital-acquired infections and antibiotic use. JAMA 1986 Aug 22/29;256(8):1007-11.
88. McDonald CJ, Hui SL, Smith DM, Tierney WM et al. Reminders to physicians from an introspective computer medical record: a two-year randomized trial. Ann Intern Med 1984 Jan;100(1):130-8.
89. Barnett GO, Winickoff R, Dorsey JL, Morgan MM et al. Quality assurance through automated monitoring and concurrent feedback using a computer-based medical information system. Med Care 1978 Nov;16(11):962-70.
90. Shortliffe, Medical expert.
91. Gelernter D, Gelernter J. Expert systems and diagnostic monitors in psychiatry. In: Cohen GS, ed. Proceedings of the Eighth Annual Symposium on Computer Applications in Medical Care. Washington, DC: IEEE Computer Society Press, 1984:45-8.
92. Tsuji S, Shortliffe EH. Graphical access to medical expert systems: I. Design of a knowledge engineer's interface. Methods Inf Med 1986;25(2):62-70.
93. Schwartz, op. cit.
94. Miller, The INTERNIST-1/QUICK.
95. Ibid., 818.
96. Miller PL. Building an expert critiquing system: ESSENTIAL-ATTENDING. Methods Inf Med 1986 Apr;25(2):71-8.
97. Blois, op. cit.
98. Barnett, DXplain.
99. Duda, op. cit., 265.
100. Goldstein, op. cit., 85.

Received July 1989; accepted January 1990

FROM THE *BULLETIN*—50 YEARS AGO

A Bibliography of the Works of Ambroise Paré. Addenda and errata. *By Janet Doe.* 8 p. 1940.

When things are going so and so,
I often think of Janet Doe
And wish that she were round about
To help me on my troubled route.

When *Paré* from the printer came,
Without a flaw or misspelled name,
I thought at last, here is a task,
Perfection reached, as was expected.

And now she lists a long *addenda*,
As useful as a sound suspender
That evens out the load one carries
And enables us to add some *Parés*.

It's nice to know that Waller sought
And finally from a Comtesse bought
The Paris *Peste* of one-five-eight-o,
Thought of as unique in the S.G.O.

To Cushing he sold, as seems fair enough,
Old friends pass on that sought of stuff
To what they call, it's "natural home,"
And New Haven was surely the place for the tome.

Another soon came to the home in Lidköping,
Not very surprising, and no cause for quirking,
For everything goes to the banks of the Lid,
Except a rare item that's here overbid.

Of *errata* a few were bound to creep in,
An 'a' for an 'e'; a 'one' for a 'ten.'
Make changes at once, so your reader may know
That your *Paré*'s complete, by the marvellous Doe.

Bull Med Libr Assoc 1940 Mar;28(3):165