COMPUTERS AND BIOMEDICAL RESEARCH 22, 113–135 (1989)

Knowledge Management as a Decision Support Method: A Diagnostic Workup Strategy Application

ROBERT A. GREENES,* DAVID B. TARABAR, MADELINE KRAUSS, GINA ANDERSON, WALTER J. WOLNIK, LAWRENCE COPE, ERIC SLOSSER, AND WILLIAM HERSH

Harvard Medical School, Department of Radiology, Brigham and Women's Hospital, Boston, Massachusetts 02115

Received August 8, 1988

We have explored the potential of a computer-based approach called "knowledge management" to aid in clinical problem solving and education. The major features of the approach are its ability to support flexible and immediate access by a user to relevant knowledge and annotation and organization of the knowledge for personal use and subsequent retrieval. We illustrate this approach with its application to diagnostic workup strategy problems. In this application, knowledge may be in the form of static narrative text, diagrams, pictures, graphs, tables, flow charts, or bibliographic citations. Other more dynamic forms of knowledge may be the result of simulations, "what if" analyses or modeling, quantitative mathematical or statistical calculation, or heuristic inference. User assessment has demonstrated the system's ease of use and user perception of its desirability, but underscores the need for a "critical mass" of knowledge before such an approach will be widely utilized.

INTRODUCTION

Computers now offer the potential of playing a much greater role than ever before in aiding the physician in his or her professional activities. A computer-based approach we believe to be among the most promising is referred to as *knowledge management* (*1*). Knowledge management is a methodology for facilitating the problem-oriented access of a user to knowledge and expertise that is useful to a specific decision-making task. The knowledge itself may be in the form of static information derived from textbooks, access to archival literature from a bibliographic data base, or more dynamic invocation of quantitative and artificial intelligence-based analyses and inference procedures, as appropriate to the decision problem. As contrasted with *model-based* decision support methodologies, such as rule-based expert systems or probability-based prediction models, knowledge management does not rely on a

* To whom correspondence and reprint requests should be addressed at Brigham and Women's Hospital, Department of Radiology, 75 Francis Street, Boston, MA 02115.

0010-4809/89 $3.00

particular decision support paradigm; rather, it is a *problem-based* approach, which provides a framework for access to whatever kind of knowledge is available and pertinent to the current problem. A knowledge management system ideally provides a comprehensive framework that supports keyword-based information retrieval, traversal of "hypertext" (nonlinear text) and "hypermedia" knowledge bases (*2*), and invocation of simulations or dynamic inference procedures, as appropriate. It should provide navigational aids for exploring a knowledge base to locate desired information or to examine its structure, and it should permit tailoring of pathways through the knowledge base and annotation of knowledge for personal use.

Knowledge management is now attractive as an approach to providing decision support in clinical practice for several reasons: First, high resolution displays can be used, which provide easy-to-read text, including multiple forms, sizes, and styles, as well as use of graphics for portraying detailed pictures, and charts. Displays can be provided in multiple "windows," giving the ability to make easily available a larger quantity of information than actually fits on the screen of the computer. The use of "mouse" pointing devices and "pull-down" and "pop-up" menus facilitating rapid selection of windows or choices within windows makes interaction with the computer intuitive and obvious. High capacity storage is now possible for personal computer workstations, including CD-ROM devices capable of inexpensively storing 500 million characters of data on single 5-inch digital optical disks. Networking technology also allows the personal workstation to interact with hospital computer systems, clinical laboratory systems, or ambulatory record systems to retrieve patient data; with electronic mail systems, to communicate with colleagues; and with bibliographic search services and other information utilities, to access data and knowledge that are not available locally. Workstations with these capabilities are increasingly becoming feasible for physicians to have on their own desktops.

The approach of knowledge management is also highly desirable for educational purposes, particularly in case-oriented problem-solving settings, or when specific items of information, or experience with particular topics, needs to be obtained. It is especially relevant as attention in medical curricula is shifting away from emphasis on memorization and factual recall toward provision of resources that can be used for information retrieval and problem solving—modes that parallel the needs of clinical medicine (*3*).

A focus on diagnostic workup strategy. Knowledge management is of interest in terms of its ability to provide access to a wide variety of knowledge, e.g., relating to differential diagnosis, pathophysiologic basis for diseases, diagnostic workup strategy, or issues of patient management for specific conditions. In this paper we shall illustrate the approach by its application to the support for diagnostic workup strategy selection involving imaging procedures, as a prototype for broader application of the approach. The work described here involved the development of a knowledge base for diagnostic workup strategy support, CASPER, to be used in conjunction with a knowledge management system which was concurrently developed, known as Explorer-1 (*1*). Motivations for the particular application to diagnostic workup strategy involving imaging procedures include the following considerations:

(1) Technological developments in radiology. The field of radiology has been experiencing a period of rapid technological innovation and development which has been responsible to date for a succession of new diagnostic procedures and therapeutic methods—with the promise of continued expansion of "high technology" capabilities into the future. As a consequence, multiple modalities have become available for workup of many diagnostic problems. The wide spectrum of radiologic procedures that are reasonable to consider for any specific clinical condition, with their overlapping indications, differing risks, preparation requirements, and quality of performance, makes the choice of procedure for a patient difficult under the best of conditions.

(2) Cost constraints. Deployment and use of radiologic imaging technology are currently being constrained by several factors, foremost among them a new climate of cost consciousness in medicine, in which both diagnostic and therapeutic procedures are regarded more thoroughly with respect to their cost than previously. A challenge facing the physician is to minimize costs without compromising quality of medical care.

(3) The nature of the decision problem. Physicians are poorly equipped to make cost–benefit trade-offs in evaluating alternative diagnostic workup strategies. Data bases, analytic tools, and expert judgment, providing relevant information about the appropriate workup for particular clinical problems, are generally unavailable to the physician when faced with these problems. The problem is not simply one of choosing the optimal test from among competing alternatives. Patient management often involves sequential testing. The results of one test must be interpreted in terms of patient management objectives, to determine whether the degree of uncertainty has been sufficiently resolved to permit a change in treatment. If not, then the process of diagnostic strategy selection continues in an iterative fashion. Thus the diagnostic algorithm appropriate to a specific clinical problem may have a large number of branches, reflecting the particular characteristics of the patient, and the results obtained through testing up to that point. Because of the degree of specialized knowledge that must go into constructing suitable strategies for the variety of clinical problems likely to be encountered, it is not surprising that most clinicians would like to have available to them a way of obtaining consultation on the selection of appropriate diagnostic strategy.

(4) Difficulties in supporting this decision-making task. Radiology departments have attempted to meet the need for providing expert advice in several ways, through books, continuing education courses, daily radiology conferences, and consultation with referring physicians. However, none of these methods has been carried out entirely satisfactorily, leading Heilman, in 1982, to write an editorial lamenting the inadequacy of the interface between the radiologist and the referring physician (*4*). The difficulty is often one of logistics or immediacy of availability of the consultation when the physician has the

particular question. Also subspecialization within radiology makes it difficult even for a particular radiologist to always be able to provide appropriate consultation for the range of problems encountered.

(5) Potential role for the computer. As indicated above, one way to aid in the selection of diagnostic strategy which is being explored is the use of the computer as a decision support tool. Among the capabilities a computer can provide is access to a data base of tests that could be considered for each clinical problem, including information about the performance characteristics of the tests. Through use of the computer, analytic procedures can be used for assessing the appropriateness of the test in relation to specific patient management objectives (*5–7*). The computer can also provide educational material about the nature and interpretation of these analytic methods. Finally, the computer can be used to retrieve expert advice about the general approach and algorithm to be followed for various clinical problems, including the features of specific tests, as well as pertinent references from the literature.

Goals of this project. This application is designed as a knowledge resource for the clinician considering an imaging workup for a variety of common clinical problems, as well as an aid for the radiologist asked to consult on imaging workup strategy. The application is intended for use on a personal workstation, by physicians in a variety of settings. It is thus not tied to a particular hospital information system (HIS), although it can operate, on a multi-tasking basis with current microcomputer operating system software, along with terminal emulation programs that are used to access the HIS. It was developed as a prototype of a potentially expanding library of knowledge resources which could be made available to physicians through desktop workstations.

The principal aim of development was to identify the requirements involved in preparing knowledge management applications, to identify knowledge management system features that would be useful, authoring tool characteristics, as well as difficulties in adapting content material originally prepared for other purposes. Another goal was to carry out a user assessment to determine whether perception of ease of use and desirability of the approach were sufficiently high to warrant the much larger efforts of maintaining the knowledge base on an ongoing basis, and scaling up the approach to include many more domains of knowledge.

Methods

Design of CASPER. The design objectives described in the previous section have been addressed in a knowledge management application we have named CASPER (computer-aided selection of procedures and evaluation of results). It is aimed specifically at providing the information and evaluations needed for the determination of appropriate diagnostic workup strategy for clinical problems that involve radiologic imaging procedures. We initially implemented CASPER as a specific application program for physician decision support (*8*).

Another adaptation of our original approach by Kahn *et al.* (*9*) has emphasized the encoding of clinical algorithms, as contained in CASPER, in the form of a set of rules. However, we became convinced during further development that a more generalized framework, employing a wide variety of knowledge management capabilities, would have the greatest likelihood of being useful to physicians. Not only is the power of CASPER considerably enhanced, but the ability of a user to integrate information from other sources as well as CASPER is thereby facilitated. We describe CASPER therefore as one of the first full-size applications of the role of knowledge management as a decision support and educational methodology.

The central feature of CASPER is a computer "knowledge base" dealing with the radiologic imaging workup of 62 clinical problems. The knowledge base includes, for each clinical problem, information on the general characteristics of the patients being evaluated, a description of the approach to the diagnostic workup, a list of the diagnostic conditions being considered, a flowchart of the clinical algorithm, descriptions of the diagnostic tests available, along with the quantitative data about efficacy, and pertinent references from the literature. For each test, the data available include the indications, preparation requirements, advantages and drawbacks, success rate, sensitivity and specificity, and reasons for false positives and false negatives, as well as other additional comments.

The CASPER knowledge base is composed of a collection of "frames" or modular units of knowledge. A total of 1232 frames are included in the knowledge base. An individual frame is an arbitrary collection of text and/or a diagram or chart or set of references on a particular specific topic which is retrieved as a unit; examples are a frame describing the general characteristics of a patient who might be considered as having the clinical problem of biliary obstruction, or a frame discussing the use, drawbacks, sensitivity, and specificity of ultrasound for this problem.

The CASPER knowledge base application is accessed through a knowledge management system known as Explorer-1 (*1*), which we have implemented on the Macintosh computer (Apple Computer, Inc.). It utilizes the features of the Macintosh, including multiple-window access to knowledge, so that each frame can be presented in its own window, scrolling through knowledge content in individual window frames, pull-down menus, and interaction with a simple mouse pointing device.

The CASPER knowledge content is based in large part on material compiled as a department-wide effort of the Brigham and Women's Hospital Department of Radiology, and available in handbook form (*10*), as well as augmented by other sources. This material was derived by extensive literature review, and by peer review of each chapter by colleagues and by the editors of the handbook. For use in a knowledge management environment, the material was reorganized for the frame-based mode of retrieval and access which is provided by the Explorer-1 system. An authoring system was developed for use with Explorer-1 which permitted narrative content material to be created, or imported from a word processing program and restructured, and facilitated the incorporation of hypertext links and links to other decision support modules. Each content frame was given a title, and a set of keywords was defined for it, to permit keyword-based retrieval. A thesaurus was constructed from the keywords that had been defined.

The computer implementation has permitted a variety of dynamic capabilities to be added, in addition to hypertext-oriented linking and cross-referencing of content, and keyword retrieval. For concepts such as sensitivity, specificity, or prevalence, frames containing educational material can be retrieved which explains these concepts and related ones. To further illustrate such didactic material, pictorial simulation frames are available (*11*, *12*) which allow the user to dynamically view the consequences of having two overlapping result distributions, for nondiseased and diseased patients, in terms of their impact on sensitivity and specificity, or visualize the effect of pretest probability of disease on predictive value. Manipulations such as changes in the pretest probability, the mean or standard deviation of the distribution for nondiseased or diseased populations, or the cutoff threshold, can be made and the consequences viewed pictorially.

In addition, performance characteristics of the tests can be related to patient-specific factors, enabling the user to assess the potential impact of the test in a particular patient. By entering estimates of pretest probability of disease and the desired patient management threshold probability for either ruling-in or ruling-out disease in order to change treatment, a user is able to assess the potential value of the test in facilitating that objective. For example, such quantities can be determined as the post-test probability of disease given a positive or negative result, the chance of a positive or negative result, or the chance that the test will permit either the rule-in or the rule-out threshold to be exceeded (the assignment potential of the test). These quantities can be portrayed in either tabular or graphic form. As with the simulations discussed above, input quantities such as pretest probability estimate of disease, or rule-in or rule-out threshold, can be modified and the effect on the computed quantities viewed dynamically. This permits a kind of "what if . . ." analysis that lets the physician assess whether under any of the conditions considered reasonable in terms of the input quantities entered the test would be useful in the patient. This is a means for compensating for the known imprecision of judgmental probability estimates, and for subjecting conclusions or interpretations to a "sensitivity analysis" to determine how stable they are.

Selection of knowledge content to be retrieved can be done either (a) by selection of topics from menus or (b) by means of keyword search for a pertinent term or phrase. The frame is then displayed in a window on the screen. A user is able to easily branch from one content window to another to access related material by means of hypertext links, in which mouse-based selection is made of "hot spots" or designated choices of words, phrases, or picture elements that are highlighted to indicate that the additional material is available. Explorer-1 itself is a user environment for traversing a knowledge base, and CASPER is one of any number of possible knowledge domains it could be used to access. The Explorer-1 system also provides facilities for the author to create knowledge frames, index them, and link them into the knowledge base.

In addition to multiple-window access to knowledge, other means are provided for facilitating navigation through the knowledge base, and browsing through material. This capability is provided because it is clear that with multiple windows on a "desktop" image on a computer, it quickly becomes difficult for the user to keep track of what is there, or to identify other potential avenues to pursue. An "overview mode" is therefore available, which allows the user to inspect the path he or she has traversed through the knowledge base, using symbols or brief titles to represent the contents of each window. This permits the user to determine quickly what other frames of knowledge are potentially accessible from each point, and to open other windows as desired, in order to explore them.

In overview mode, the user can also edit the diagram of opened windows, deleting those of little interest, and retaining those that may be of particular interest or that would be helpful subsequently in accessing a particular part of the knowledge base. This edited state of the desktop can be saved as a personalized knowledge path, under a name supplied by the user, and restored at a later time, as a basis for further knowledge retrieval.

User assessment. We assessed users to determine whether the design of the knowledge management environment fulfilled its objectives of being easy to use to retrieve specific knowledge pertaining to diagnostic workup strategy, and the extent to which its environment facilitated browsing and the pursuit of curiosity. This was carried out by having 12 physicians (5 residents in radiology and 7 attending physicians) use the system to obtain the answers to three specific questions of their choice, by monitoring their use of the system, and by assessing their reactions to it. The subjects were chosen to be those with prior general experience with the use of the Macintosh computer system for unrelated applications (e.g., word processing or data base access), so that their use of this application could be assessed without that variable. Prior to using CASPER, subjects were given a brief explanation of the kinds of knowledge encompassed by CASPER and demonstration of its use (for no more than 5 min) but no written documentation. Two of the subjects had prior experience with development of small portions of the knowledge content of the Brigham handbook on which CASPER was based.

RESULTS

Operation of CASPER. CASPER is able to be used in an entirely stand-alone mode, for example, whenever reference to its content or use of it for decision support is desired. Yet the microcomputer on which it is implemented can also be used to interact with a hospital information system, through one of many available terminal emulation programs. When a problem requiring consultation a

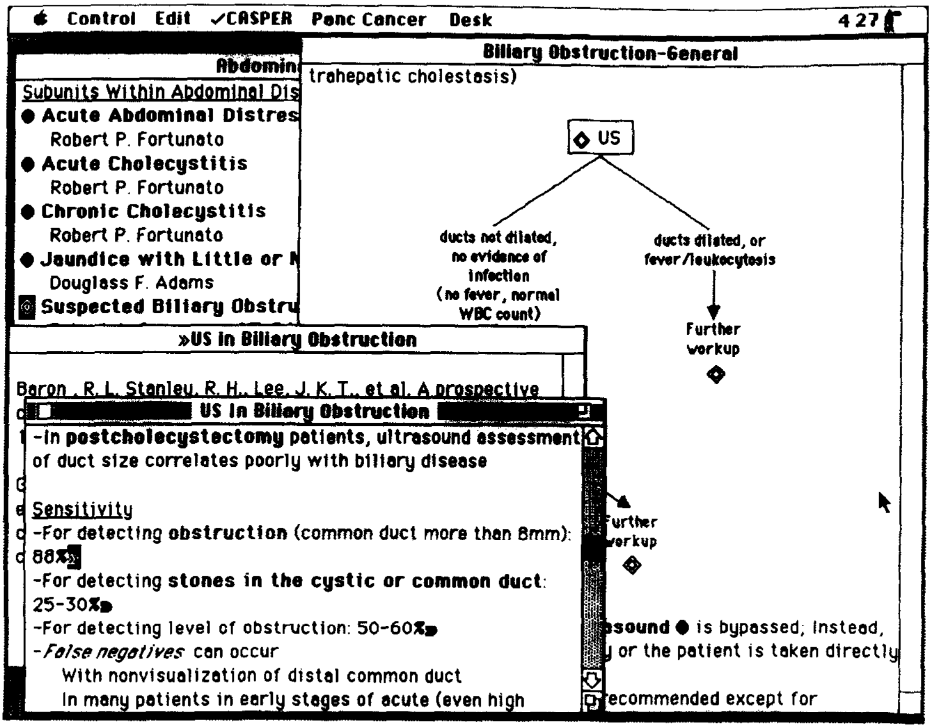
Control Edit ✓CASPER Panc Cancer Desk
4 27
Biliary Obstruction-General
trahepatic cholestasis)
US
ducts not dilated, no evidence of infection (no fever, normal WBC count)
ducts dilated, or fever/leukocytosis
Further workup
Abdomin
Subunits Within Abdominal Dis
Acute Abdominal Distres
Robert P. Fortunato
Acute Cholecystitis
Robert P. Fortunato
Chronic Cholecystitis
Robert P. Fortunato
Jaundice with Little or
Douglass F. Adams
Suspected Biliary Obstru
»US in Biliary Obstruction
Baron, R. L. Stanley, R. H., Lee, J. K. T., et al. A prospective
US in Biliary Obstruction
-In postcholecystectomy patients, ultrasound assessment of duct size correlates poorly with biliary disease
Sensitivity
-For detecting obstruction (common duct more than 8mm): 88%
-For detecting stones in the cystic or common duct: 25-30%
-For detecting level of obstruction: 50-60%
-False negatives can occur
With nonvisualization of distal common duct
In many patients in early stages of acute (even high
urther workup
sound ● is bypassed; instead,
or the patient is taken directly
ecommended except for

b

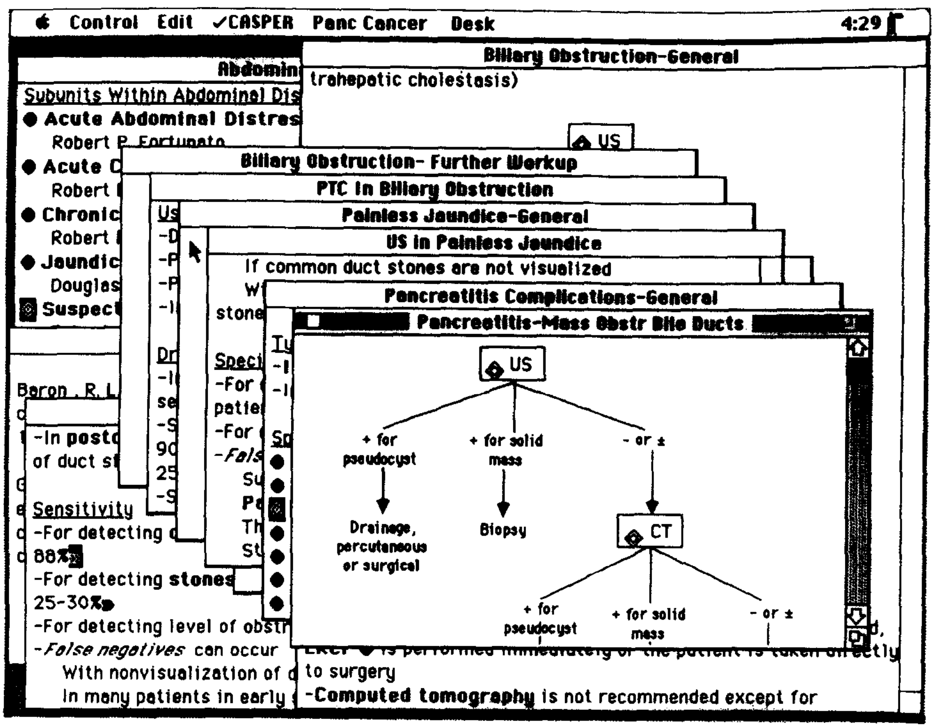
Control Edit ✓CASPER Panc Cancer Desk
4:29
Biliary Obstruction-General
trahepatic cholestasis)
Abdomin
Subunits Within Abdominal Dis
Acute Abdominal Distres
Biliary Obstruction- Further Workup
PTC in Biliary Obstruction
Painless Jaundice-General
US in Painless Jaundice
If common duct stones are not visualized
Pancreatitis Complications-General
Pancreatitis-Mass Obstr Bile Ducts
US
+ for pseudocyst
+ for solid mass
- or ±
Drainage, percutaneous or surgical
Biopsy
CT
+ for pseudocyst
+ for solid mass
- or ±
Sensitivity
-For detecting
88%
-For detecting stones
25-30%
-For detecting level of obst
-False negatives can occur
With nonvisualization of
In many patients in early
to surgery
-Computed tomography is not recommended except for

or reference to CASPER is encountered, the program can then be accessed through the multi-tasking facilities of the microcomputer's operating system.

Use of the system typically begins with identification of the specific problem or diagnostic test of interest. This is done by specifying the topic via selection either from a menu or by keyword or phrase. Either a diagnostic problem or a particular procedure may be indicated as a menu choice. For each diagnostic problem area, access is provided to information about typical patient characteristics, the approach to the imaging workup, a recommended clinical algorithm, the imaging tests available, and pertinent references. Figure 1 shows use of Explorer-1 to access the CASPER knowledge base, in order to find information on the workup of suspected biliary obstruction, and specifically on the role of ultrasound, followed by exploration of a variety of related topic areas.

Note that, besides menu selection of pertinent topics, keyword lookup provides an alternative method of retrieval allowing specific frames to be accessed without selecting a topic from a menu. This is analogous to using the index in a book to find a topic of interest rather than using the table of contents. For example, the user may request information on "ultrasound for cancer of the pancreas." The keyword lookup procedure ignores connecting words such as "for," "of," and "the" and does a logical AND of the keywords given to find those frames from an indexed list that pertain to the given words. A limited number of synonyms are built into the keyword directory, and the lookup procedure treats the words entered as if they are stem words, so that a word in the index matches the input if it contains the input word as its leading characters. Thus "ultraso canc panc" would retrieve the relevant frames, and would also be likely to retrieve frames where index terms such as "ultrasonic" and "pancreatic" were used instead of "ultrasound" and "pancreas." When a frame is selected through this process, it is displayed in a window on the desktop. If the keyword phrase does not match a unique frame, available choices are presented to the user, from which a particular option may be selected.

FIG. 1. Two representative screen displays (a,b) during interaction with CASPER and the Explorer-1 knowledge management system, showing multiple "windows" of information retrieved by interaction between the user and the system. The user has first specified which knowledge bases are of interest (shown in the menu bar on the top line, in which CASPER and a medical textbook chapter on pancreatic cancer are indicated, referring to content relating to diagnostic workup strategy and to pancreatic cancer pathophysiology and clinical management, respectively). The check in front of CASPER indicates that the currently topmost window comes from that knowledge base. The user has selected particular windows by identifying topics of interest through a combination of either keyword phrase lookup, menu selection, graphical navigation (see Fig. 3), or the following of hypertext links or "hot spots." The latter involves indicating a topic of interest in a displayed window, by selecting it with the mouse pointing device, as a result of which a new window is displayed on top of the previous windows, with the requested information. Note that windows can contain a combination of text and graphics, including clinical algorithms or flow charts, as shown in the large window on the right in (a). Choosing the Desk Menu (at right of menu bar on top line) displays a list of all the windows on the screen, including those currently obscured by other overlapping windows, providing the ability to the user to quickly review what knowledge content has been accessed in Explorer-1's windows, and to bring any of these to the front.

In this manner, multiple access paths are provided to the relevant knowledge. For example, if one wished to approach retrieval, not by considering a problem first, but by considering a particular test, and only then a specific clinical problem from those offered for which the test was potentially useful, this could be done by simply entering a keyword phrase indicating the test of interest. As a result, a list of frames describing use of that test for various clinical problems would be presented. Information about the test for any particular indication could then be viewed.

From within the discussion of a test for a particular indication, a "what if . . ." analysis can be invoked to dynamically assess the potential value of the test for a specific patient (Fig. 2). By either "hot spot" selection or keyword lookup, didactic material and interactive simulations can also be accessed that provide information about various quantitative and probabilistic performance measures and analytic methods. These performance measures and analytic methods relate to the evaluation and interpretation of diagnostic tests and include distributions of test results, sensitivity, specificity, ROC curves, predictive value, Bayes theorem, thresholds, assignment potential, and assignment strength.

A "frame" of content is presented in a CASPER window. This may be a sequence of narrative paragraphs about some particular topic, a diagram or picture illustrating the topic, or list of options available to the user. The windows in which CASPER frames are presented may be moved around on the computer desktop, shrunk or enlarged, or the content may be "scrolled" or "paged" up or down within the window. Frame content may be larger than the window itself, but typically not very much larger. Rather, additional content material is reached through selection of "hotspots"—highlighted content material within the window that indicates that other material can be obtained on that topic, to expand, illustrate, define, or otherwise delineate it, or that a procedure is available that will perform some specific function at that point. Examples of the latter might be invocation of analytic procedures to relate the

FIG. 2. *Dynamic procedural capabilities.* Some windows present displays that are produced dynamically as a result of data determined through interaction with the user or as the result of calculations and heuristic inferences, e.g., a differential diagnosis or a probabilistic analysis or simulation. Shown here are (a) selection of "what if . . ." analysis in description of the use of ultrasound for detecting biliary obstruction; (b) display of a spreadsheet in which the test characteristics and various other input parameters are used to calculate several measures of test efficacy; (c) modification of the spreadsheet through entry of different estimates for input parameters, in this case pretest probability of disease, and recalculation of efficacy measures; (d) graphical view of the first two calculated measures in the spreadsheet; (e) graphical view of the last two calculated measures in the spreadsheet; and (f) a display of the graphical relationship between the result distributions for disease and nondiseased patients as a basis for the definitions of sensitivity, specificity, and predictive value. In all these displays the input parameters can be modified resulting in dynamic alteration of the display. In (b)–(e) this permits a graphical "sensitivity analysis" to be performed on the calculated efficacy measures. In the graphical simulation in (f) these same capabilities can be used to dynamically modify the relationships of the parameters as an aid to understanding the concepts involved.

a

b

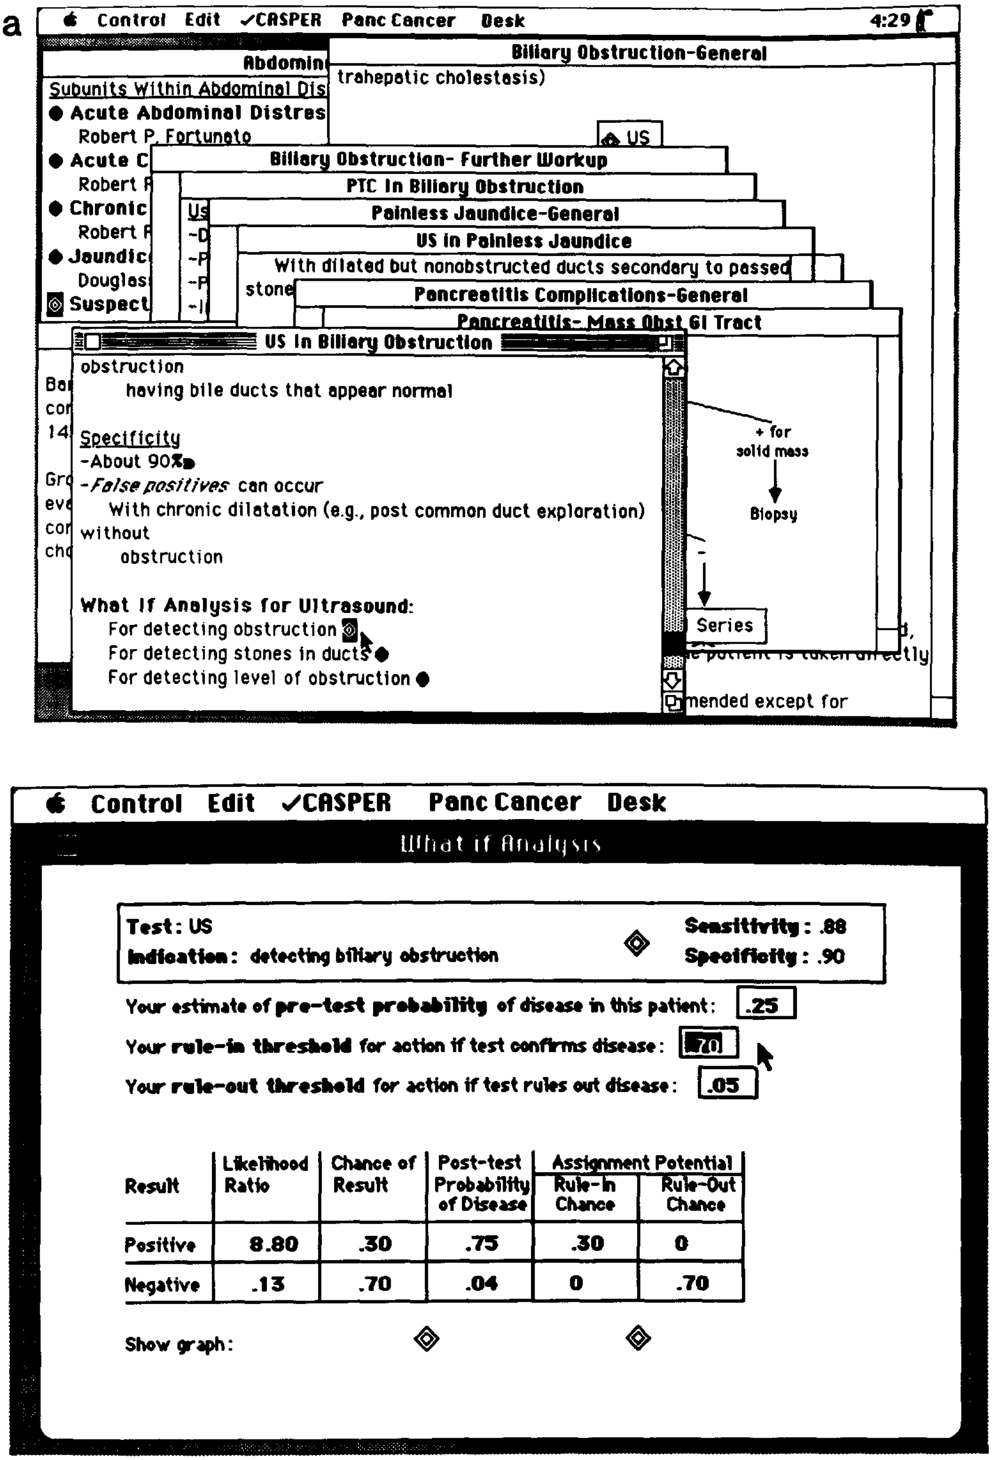
Control Edit ✓CASPER Panc Cancer Desk
4:29
Biliary Obstruction-General
trahepatic cholestasis)
Subunits Within Abdominal Dis
Acute Abdominal Distres
Robert P. Fortunato
Biliary Obstruction- Further Workup
PTC In Biliary Obstruction
Painless Jaundice-General
US in Painless Jaundice
With dilated but nonobstructed ducts secondary to passed
Pancreatitis Complications-General
Pancreatitis- Mass Obst GI Tract
US In Biliary Obstruction
obstruction
having bile ducts that appear normal
Specificity
-About 90%
-False positives can occur
With chronic dilatation (e.g., post common duct exploration)
without
obstruction
What If Analysis for Ultrasound:
For detecting obstruction
For detecting stones in ducts
For detecting level of obstruction
+ for
solid mass
Biopsy
Series
mended except for
Control Edit ✓CASPER Panc Cancer Desk
What if Analysis
Test: US
Indication: detecting biliary obstruction
Sensitivity: .88
Specificity: .90
Your estimate of pre-test probability of disease in this patient: .25
Your rule-in threshold for action if test confirms disease: .70
Your rule-out threshold for action if test rules out disease: .05
Result
Likelihood Ratio
Chance of Result
Post-test Probability of Disease
Assignment Potential
Rule-In Chance
Rule-Out Chance
Positive 8.80 .30 .75 .30 0
Negative .13 .70 .04 0 .70
Show graph:

c

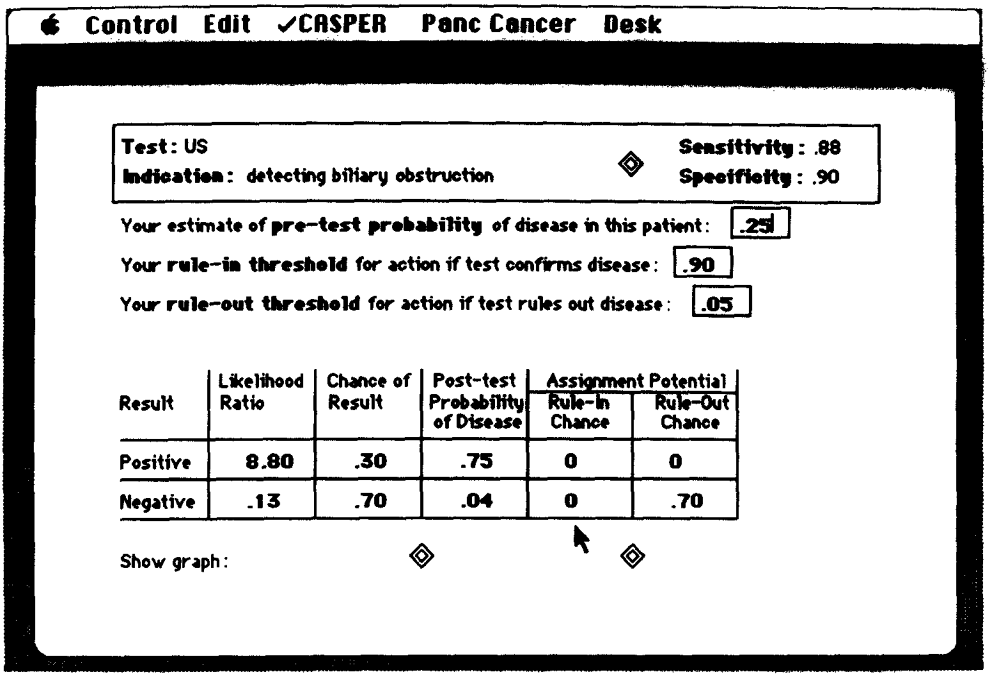
Control Edit ✓CASPER Panc Cancer Desk
Test: US
Indication: detecting biliary obstruction
Sensitivity: .88
Specificity: .90
Your estimate of pre-test probability of disease in this patient: .25
Your rule-in threshold for action if test confirms disease: .90
Your rule-out threshold for action if test rules out disease: .05
Result
Likelihood Ratio
Chance of Result
Post-test Probability of Disease
Assignment Potential
Rule-In Chance
Rule-Out Chance
Positive 8.80 .30 .75 0 0
Negative .13 .70 .04 0 .70
Show graph:

d

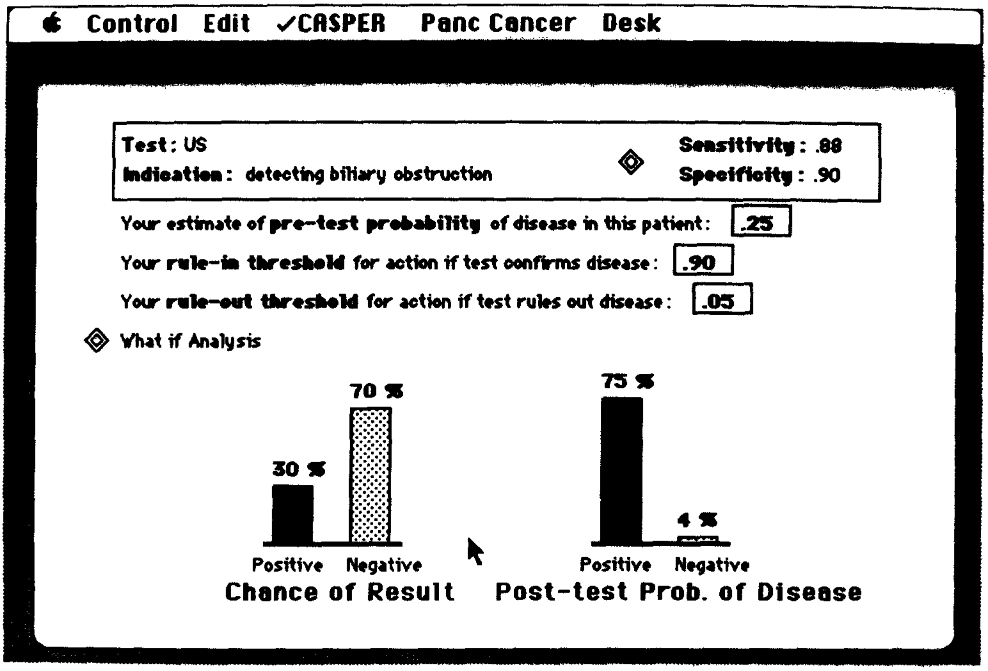
Control Edit ✓CASPER Panc Cancer Desk
Test: US
Indication: detecting biliary obstruction
Sensitivity: .88
Specificity: .90
Your estimate of pre-test probability of disease in this patient: .25
Your rule-in threshold for action if test confirms disease: .90
Your rule-out threshold for action if test rules out disease: .05
What if Analysis
30 %
70 %
Positive Negative
Chance of Result
75 %
4 %
Positive Negative
Post-test Prob. of Disease

e

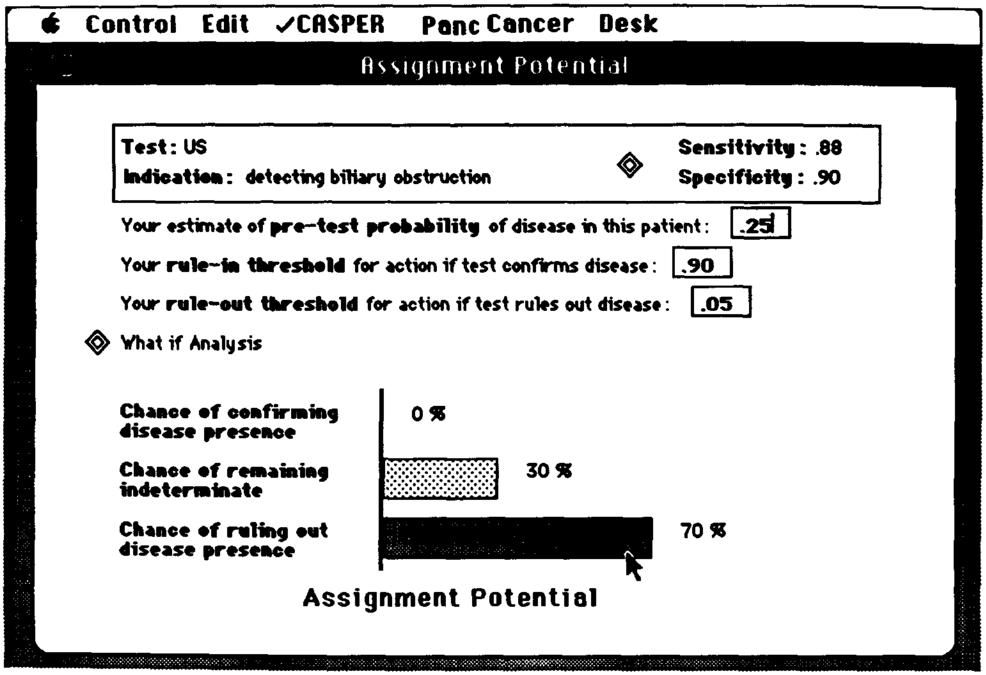
Control
Edit
✓CASPER
Panc Cancer
Desk
Assignment Potential
Test: US
Indication: detecting biliary obstruction
Sensitivity: .88
Specificity: .90
Your estimate of pre-test probability of disease in this patient: .25
Your rule-in threshold for action if test confirms disease: .90
Your rule-out threshold for action if test rules out disease: .05
What if Analysis
Chance of confirming disease presence
0 %
Chance of remaining indeterminate
30 %
Chance of ruling out disease presence
70 %
Assignment Potential

f

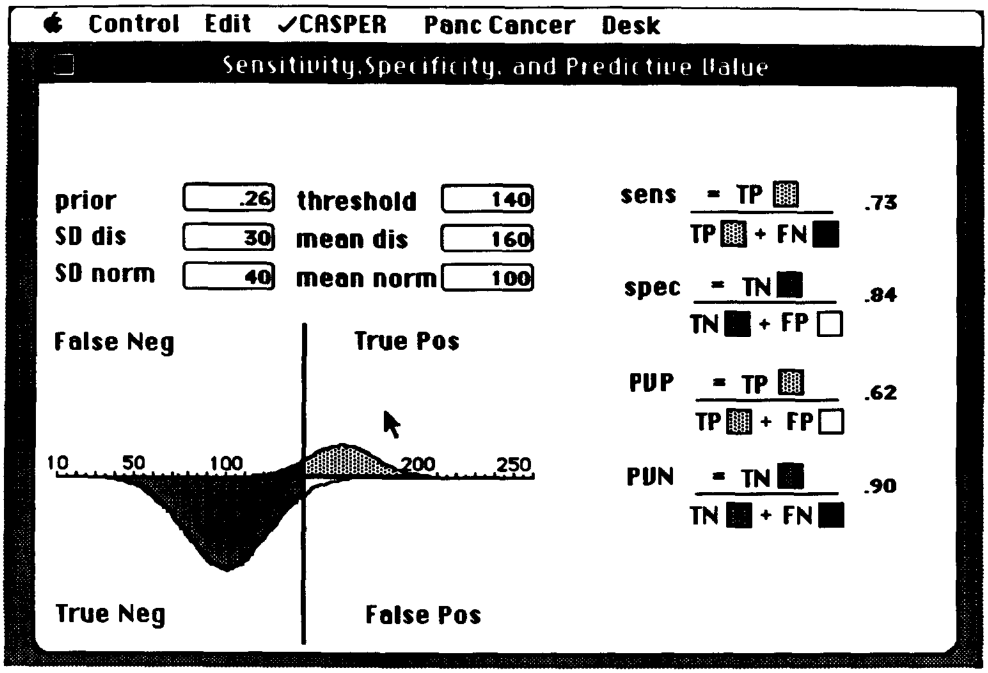
Control
Edit
✓CASPER
Panc Cancer
Desk
Sensitivity, Specificity, and Predictive Value
prior .26
SD dis 30
SD norm 40
threshold 140
mean dis 160
mean norm 100
False Neg
True Pos
10
50
100
200
250
True Neg
False Pos
sens = TP / TP + FN .73
spec = TN / TN + FP .84
PVP = TP / TP + FP .62
PVN = TN / TN + FN .90

test to patient-specific characteristics (the "what if . . ." analysis described above) and to plot the relationship graphically, or simulation procedures designed to illustrate a concept or relationship and to allow manipulation of it to increase understanding of it.

Figure 3 shows the use of Overview Mode to portray a map of the pathway a user has taken through the knowledge base during a particular session. Each box represents a content window on the desktop, and each link, a branch from one window that was taken to another window. Windows can be inspected through a "zoom" command. Potential links not taken can be examined. Windows and links no longer of interest can be excised from the map, and the display redrawn. The edited map can be saved, to be later restored as the starting point for a future interaction with CASPER.

User assessment. Assessment of the use of CASPER by the 12 subjects to answer three queries is shown in Table 1. Queries to CASPER were assessed in terms of total time required to find the relevant information, number of frames of content viewed per query, which system features were utilized, and the extent to which browsing or pursuit of curiosity occurred versus targeted information retrieval. All subjects were able to obtain answers from the system to the queries they sought, provided that the questions were actually answerable; in other words, lack of answer of a question was in no case attributable to inability to find the relevant knowledge because of the system design or organization of the material, but only because of intrinsic limitations of the content knowledge of CASPER.

TABLE 1

EXPERIENCE OF 12 RADIOLOGISTS (5 RESIDENTS, 7 STAFF) IN USING CASPER TO OBTAIN THE ANSWERS TO THREE QUERIES EACH, CONCERNING DIAGNOSTIC WORKUP STRATEGY OR THE CHARACTERISTICS AND APPLICABILITY OF SPECIFIC DIAGNOSTIC TESTS

	Time per query (min)	No. of frames per query	Time per query, last query only (min)	No. of frames per query, last query only
Queries	36	36	12	12
Mean	3.9	4.9	2.8	3.8
SD	1.9	2.1	1.0	1.5
max	10	10	5	6
min	2	1	2	1

	% Queries using zoom	% Queries using "what if"	% Queries using overview	% Frames due to browsing	% Queries begun via keyword
Queries	36	36	36	36	36
Mean	86.0	42.0	17.0	19.7	14.9
SD	35.0	50.0	38.0	16.5	23.1

Note. Subjects had prior experience with the Macintosh computer but not with this application and were given no more than 5 min of specific training and no written documentation prior to use.

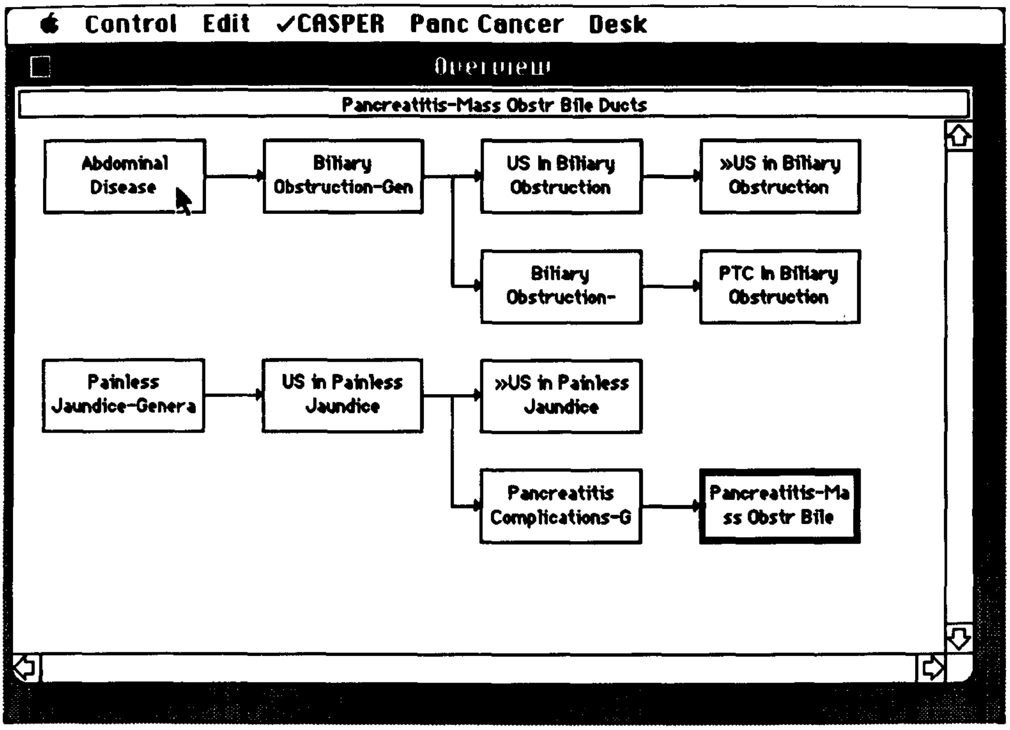

FIG. 3. *Overview mode.* A menu choice in Explorer-1 causes a graphical map of the sequence of interactions the user has taken to be displayed. Each box is an "icon" representing a window that has been activated, and the arrows indicate the connectivity between windows that has been followed by the user. Since the boxes are small, the full title of a currently selected window is displayed just below the menu bar. From this display the user is able to "zoom" in on the actual contents of any window, or can view a list of branches available to be followed from the window or a list of other windows which could have led to the given window, from which any of these other windows can be selected and added to the map. The map can also be edited in other ways, deleting windows that are not of principal interest (perhaps having been used only in a pathway toward other windows containing the information actually sought). The edited state of the map can then be saved in the computer's disk under a name assigned by the user; such saved state maps can form a basis for a personalized knowledge directory, and can be restored as desired as the starting point for subsequent knowledge access in the future.

Average time per query was 3.9 ± 1.9 min, considering all queries. To identify any effect of increasing experience, we separately determined the average time per query considering the last of the three queries per subject only; average time per query for this subset of queries dropped to 2.8 ± 1.0 min, thus demonstrating a moderate learning effect. Number of frames examined per query dropped from 4.9 ± 2.1 for all queries to 3.8 ± 1.5 when considering the last of the three queries per subject only. The two subjects who had prior involvement with the development of portions of the content material for the Brigham handbook on which CASPER was based did not retrieve

TABLE 2

SUBJECTIVE ASSESSMENTS OF 12 SUBJECTS REGARDING EASE OF OPERATION OF SYSTEM, POTENTIAL USEFULNESS OF THIS APPROACH ON A WIDE SCALE, AND DESIRABILITY AS AN ALTERNATIVE TO TEXTBOOK KNOWLEDGE

	Ease of operation	Usefulness on a wide scale	Desirability versus textbook
Subjects	12	12	12
Mean	5.0	4.8	4.1
SD	0.0	0.5	1.0
max	5	5	5
min	5	4	3

Note. Assessments are on a 5-point rating scale, 5 highest, 1 lowest. For the comparison with textbook, 5 favors system, 1 favors textbook, and 3 represents equal rating.

information any faster or with fewer queries than the inexperienced subjects, demonstrating no appreciable biasing effect of this prior knowledge.

Subjects used keyword search (analogous to using the index in a book) as the initial access to a problem in approximately 15% of queries, whereas they used menu selection (analogous to using the table of contents in a book) as the initial access to a problem in 85% of queries. Approximately 80% of frames accessed were judged to be directly related to the initial purpose of the query and the remaining 20% were regarded as browsing or pursuit of curiosity. "What if . . ." analysis was invoked in 42% of queries and was commented upon as a very useful tool. The "zoom" capability to permit visualization of a window's content in full screen size was used for at least one window in 86% of queries. Overview mode as a means of orientation of users was used in only 17% of queries.

The relatively low frequency of use of overview mode appeared to be related to the relatively small number of frames needed to answer each of the queries, as a result of which loss of orientation of the user was not much of a problem. This is borne out by informal evaluation we have conducted of browsing behavior of users examining the content of two knowledge bases we adapted for use with Explorer-1 from medical textbook chapters on angina and on lung cancer. The typical number of frames retrieved by a user exceeded 10 when exploring either of these chapters, and overview mode was invoked in almost all cases, which was felt by users to be important in order to avoid the feeling of "getting lost" in the knowledge base.

Evaluation of user attitudes, as indicated in Table 2, unequivocally indicated that the system was regarded as easy to use. All 12 subjects rated it "very easy" on a subjective 5-point rating scale. Once familiar with the Macintosh computer, users had no difficulty determining how to use CASPER to access specific knowledge of interest to them. Reaction to its modes of presentation and user opinion of ease of access to relevant material were uniformly positive.

When asked about potential usefulness of this approach to knowledge retrieval on a widespread basis, subjects considered it moderately to highly useful, with an average of 4.8 on a 5-point scale. Those who did not rate the approach as highly useful consistently indicated that the only negative aspects were related to current limitations of content, in situations in which the subjects wished to either learn more about the underlying pathophysiology of a disease process, its differential diagnosis, and the actual radiologic imaging appearance, or other information that was beyond the scope of CASPER's knowledge base. The need for updating and revising the knowledge base in fast-changing areas such as diagnostic imaging was also evidenced by analysis of 3 of 36 queries (8%) in which the information in CASPER had already become outdated.

When asked to compare this approach to knowledge retrieval to that of a textbook, the computer-based approach received an average rating of 4.1 on a 5-point scale (where 5 favored the computer exclusively, 3 represented equal value of computer and textbook, and 1 favored the textbook exclusively). Almost all subjects (83%) pointed out that the computer and the textbook each had advantages in specific situations.

Discussion

The application of knowledge management to clinical decision support and education in medicine is promising, but the rapidity of development in this area and its impact on practice depend on a number of factors.

"Critical mass" issues. We expect that use of computer systems for the task of knowledge management will be only sporadic until sufficient content knowledge is available to enable computers to function as consultants in a wide variety of topics—not just diagnostic workup strategy, but also differential diagnosis, treatment, and other areas. This is likely because a physician's willingness to use a computer system for a particular task must overcome the inertia of learning to use the program and must be based on a reasonable expectation of benefit for the current problem task. In the work reported here, our assessment of CASPER was aimed only at determining ease of use of the approach to retrieving problem-specific knowledge. Other kinds of evaluation that were specifically not carried out were those aimed at determining the extent to which the system would be utilized routinely in practice setting, or at assessing the impact of CASPER on actual diagnostic test ordering practices.

Computers are increasingly being used by physicians for *data* management, for both clinical and other data. Yet in the area of *knowledge* management, capabilities are still quite limited. The knowledge resources (as opposed to data resources) available to most physicians currently are largely only those for bibliographic retrieval. Until a "critical mass" of knowledge resources becomes available that can aid physicians with a wide variety of decision support capabilities tailored to specific problem-solving needs, we would not expect that physicians would normally turn to computers as their primary knowledge source. While this critical mass is not available yet, recent software and hardware developments have stimulated considerable developer interest in this area, and we may expect this situation to change rapidly.

Coupling of knowledge and data. Ideally, a knowledge resource such as CASPER should be coupled with a clinical information system so that its advice and knowledge could be made available as physicians were reviewing patient test results and planning their subsequent workup strategy. Such knowledge access could be triggered either automatically, through system rules, as provided by the HELP system at Latter Day Saints Hospital in Utah (*13*), or by the Regenstrief Information System in Indiana (*14*), or it can be invoked by specific user query. In both the HELP system and the Regenstreif System, logic is provided for detecting clinical situations which require decision support, although coupling to external knowledge support resources such as CASPER or to the bibliographic literature is not yet provided. The Brigham and Women's Hospital–Beth Israel Hospital Clinical Information System (*15*) provides online access to PaperChase (*16*), a MEDLINE bibliographic data base implementation, although this is not coupled to specific clinical activities of the physician. As new generation HIS system designs increasingly focus on use of microcomputer-based physician workstations, with system functions of the HIS being implemented to operate as servers in a workstation-based client-server network, such integration of patient data transactions and supporting knowledge resources will become increasingly feasible. It is our view that the workstation provides the most suitable locus for coupling data and knowledge, for monitoring system alert messages and detecting other alerts directly, for identification of the particular knowledge resources best suited to the problem, and for invocation of those resources. The knowledge resources themselves may be either local, at the workstation, or also accessed via network server requests.

Knowledge base maintenance and support. A problem that is raised by the prospect of electronic knowledge access is the expectation that such information be continually up-to-date and current. This requires a commitment to maintenance, updating, and support of the knowledge base, as well as of the software, that can be expected to extend long beyond the original development period. In a fast-changing area such as diagnostic radiology, a portion of the knowledge included in the CASPER knowledge base is already out of date. Clearly, to make support of such applications viable on an ongoing basis, either a wide user base is necessary in order to provide economic viability as a product or commitment to maintaining the content over time must exist from a professional organization or other body.

User interface. CASPER is intended as a prototype knowledge-based decision support tool, for the user who requires access to specific knowledge in the area of diagnostic workup strategy selection. However, additional developments are focusing on how to incorporate knowledge of other types, and to make this knowledge available as well. Examples include pathophysiologic knowledge, laboratory tests available for the same clinical problems, an expanded array of clinical problems, differential diagnostic capabilities, and therapeutic options. With respect to differential diagnosis, we have adapted the Internist-1/QMR knowledge base and inference procedures (*17*) to permit them to work in concert with CASPER and Explorer-1, as well as independently, on the same computer system. (This adaptation has been possible because of the generous cooperation of the developers of QMR at the University of Pittsburgh.)

Figure 4 depicts schematically the kinds of interactions that can be supported through use of a knowledge management framework such as represented by Explorer-1. The general philosophy that underlies our development is that knowledge is usually sought by a user in a particular problem context, i.e., when the user is faced with a "need to know" in order to make a particular decision. The information requirements of physicians in practice have been well documented by Covell *et al.* (*18*) and may deal with diagnosis, the nature of specific diseases, the characteristics of specific tests, or the properties of therapeutic modalities. Medical educators now believe that learning, in fact, occurs best in such settings, because of the motivation of the user and the relevance of the information gained to a real problem (*3*).

Often, however, the inquiry prompts additional curiosity, as suggested by the material retrieved, and such pursuit of curiosity should also be supported. The desire to pursue curiosity was demonstrated in our evaluation study, but to a lesser extent than we believe to actually be present, because of the subjects' prior knowledge of the limitations of the knowledge domain. Specific analysis assistance in differential diagnosis, in diagnostic strategy selection, or in determining a specific treatment regimen may be desired. In addition to learning specific facts in a problem context, such as the drug of choice, it may be of interest to the user to follow other relationships or pursue other specific kinds of information, e.g., of what pharmacological family the drug is a member, its mechanism of action, or the way in which it is metabolized. This is accomplished in the CASPER application by the ability at any time to select additional frames by either new menu choices, selection of "hot spots" from currently displayed frames, or keyword lookup. In addition, the use of overview mode allows the user to quickly inspect the pathway that has been taken through the knowledge base, and to examine, for any frame, the list of potential frames that could be accessed from that frame, or through which this frame could have been accessed. Overview mode thus gives the user the capability to navigate through the knowledge base in a variety of ways and to explore the relations among content material, as dictated by interest or curiosity. As we have indicated, the relative infrequency of use of this mode in our evaluation study was attributable, we believe, to the relatively small number of content frames the users needed to traverse to obtain answers to their queries. For more involved queries, or with a wider range of knowledge bases available, the likelihood of getting lost and needing such navigational aids would be higher. The ability to edit the state of the desktop by either selecting specific frames to be retained or including "note frames" or annotations further augments these capabilities. These features give the user the means to tailor the knowledge base and develop personalized maps of the access paths to relevant information for future use.

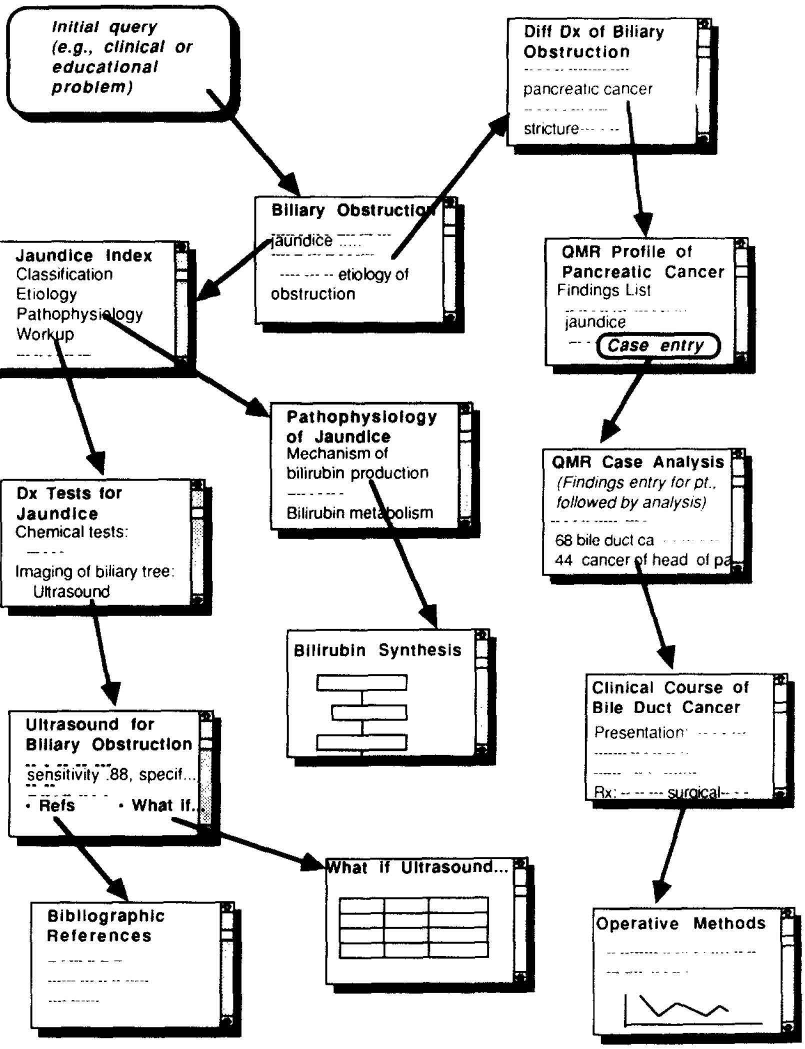

FIG. 4. *A schematic view of the use of knowledge management system to support medical decision making and education.* The use of a system such as Explorer-1 is extrapolated in this diagram to show a potential pathway a user might take through the knowledge available, both for targeted information retrieval and for pursuit of curiosity. It can be expected that the system would typically be initially accessed to provide information specific to a particular problem, or even triggered automatically by a patient care information system detecting the need for a specific alert message to the physician. The initial knowledge retrieved may consist of pointers to other more specific details, or it may suggest other related topics. The user may alternatively think of topics in which the relationship to the current topic is not explicitly recognized in the knowledge base, but which are nonetheless of interest to the user to pursue. Some of the knowledge may be static, in the form of text, images, charts, and tables. Other knowledge may be in the form of dynamic procedures, e.g., for heuristic inference, probabilistic estimation, and simulation. In all of these circumstances, a goal of the knowledge management environment is to facilitate access in a convenient and consistent fashion.

Research issues. Methods for establishing and maintaining links among knowledge units, and across different knowledge types, are particularly important. The need to formalize the structure of the knowledge, and support access to it as well as navigation through it, will be especially acute as knowledge bases expand in size and number through the incorporation of different knowledge domains and involve contributions not only by a single author but also via communal authorship and iterative refinement of knowledge. Much of our current research, as well as that of other workers, involves exploration of these complex issues. Problems that are relatively easy to solve when the knowledge base consists of several hundred frames of information become much different in scope and complexity when one deals with thousands of frames.

Near-term prospects. Knowledge management has immediate potential for aiding the practice of medicine by supporting the information-seeking needs of the clinician as problems are encountered in the process of patient care. The computer workstation used to support knowledge management can readily be used, in a clinical setting, to concurrently emulate a computer terminal for accessing patient data in a hospital or ambulatory information system; we expect that such interaction will frequently directly precipitate the need for knowledge access. Much research and development has occurred in the past decade or more in the attempt to create decision support systems to aid the decision-making tasks of physicians (*19*). These have taken the forms of both analytic and quantitative, probabilistically based procedures to aid in diagnostic test assessment (*20*), differential diagnosis, and evaluation of costs and benefits of alternative strategies, as well as less quantitative "artifical intelligence" approaches based on codified expert judgments and heuristic reasoning process (*21*). However, these various approaches to decision support have generally received only limited acceptance and have made relatively little impact on the practice of medicine despite their promise and potential. This stems from several factors, including the extent of the expertise that must be available to the computer for it to truly function as an expert, the human interface problems involved in conveying sufficient patient-specific data to the program for it to be useful, and "locus of control" issues with respect to how physicians would like to utilize such decision support programs. Our knowledge management approach overcomes many of the limitations of these decision support systems by providing access to a wide variety of knowledge under user control, while retaining the ability to invoke specific expert or analytic modules in particular contexts for which they are particularly beneficial.

The use of computers for knowledge management and selective knowledge retrieval in medicine is a relatively new pursuit, made practical by the high resolution displays, improved human interface, and mass storage capabilities of personal workstations, as well as the evolution of networking technology allowing these systems to operate as front ends to other patient care information systems. New hypertext softward and other development tools for microcomputers will stimulate this kind of development. An example is the recently introduced Hypercard program (Apple Computer, Inc.) which makes it relatively easy for a user to implement connections among related content material. The Explorer-1 system with the CASPER application has been an opportunity to explore the use of the approach of knowledge management in an arena in which the need for knowledge access in evaluating a number of complex alternatives is particularly apparent. However, we expect that the potential of the technology will be realized only as knowledge pertaining to a wide variety of topics becomes available in an integrated fashion, and as physicians increasingly rely on the computer as a knowledge source in their practice. We predict that this capability will evolve rapidly in the next few years.

ACKNOWLEDGMENTS

This work was supported in part by Grants LM 04572 and LM 07037 from the National Library of Medicine, DHHS, and by a grant from the CAMDAT Foundation, Inc., Farmington, Connecticut. We sincerely thank the staff and resident physicians of Brigham and Women's Hospital Department of Radiology, who willingly served as subjects in the evaluation of CASPER and offered helpful criticisms and suggestions.

REFERENCES

1. GREENES, R. A. Knowledge management as an aid to medical decision making and education: The Explorer-1 system. "Proc. MEDINFO 86," pp. 895–899. Elsevier, Amsterdam/The Netherlands, 1986.
2. CONKLIN, J. Hypertext: An introduction and survey. *IEEE Comput.* **20,** 17 (1985).
3. Physicians for the Twenty-First Century: The GPEP Report. Report of the Panel on the General Professional Education of the Physician and College Preparation for Medicine. Amer. Assoc. of Med. Colleges, Washington, DC 1984.
4. HEILMAN, R. S. What's wrong with radiology? *N. Engl. J. Med.* **306,** 477 (1982).
5. GREENES, R. A., CAIN, K. C., AND BEGG, C. B. Performance measures of diagnostic tests. 1. Tools for prospective test assessment. *Med. Decis. Making* **4**(1), 7 (1984).
6. GREENES, R. A., BEGG, C. B., CAIN, K. C., SWETS, J. A., FEEHRER, C. E., AND MCNEIL, B. J. Performance measures of diagnostic tests. 2. Assignment potential and assignment strength. *Med Decis. Making* **4**(1), 17 (1984).
7. CAIN, K. C., GREENES, R. A., AND BEGG, C. B. Performance measures of diagnostic tests. 3. U-factor. *Med. Decis. Making* **4**(1), 33 (1984).
8. GREENES, R. A. Toward more effective radiologic consultation: Design of a desktop workstation to aid in the selection and interpretation of diagnostic procedures. "Proceedings, 8th Conference on Computer Applications in Radiology, St. Louis, May 1984," pp. 553–562. Amer. Coll. of Radiol., Chicago, IL, 1984.
9. KAHN, C. E., MESSERSMITH, R. N., AND JOKICH, M. D. PHOENIX: An expert system for selecting diagnosis imaging procedures. *Invest Radiol.* **22**(12), 978 (1987).
10. MCNEIL, B. J., AND ABRAMS, H. L. (Eds.) "Brigham and Women's Hospital Handbook of Diagnostic Imaging." Little Brown, Boston, MA, 1986.
11. GREENES, R. A. Interactive microcomputer-based graphical tools for physician decision support: Aids to test selection and interpretation and use of Bayes theorem. *Med. Decis. Making* **3**(1), 15 (1983).
12. POLLACK, M. A., AND GREENES, R. A. A pictorial simulation construction kit for enhancing knowledge-based learning. "Proc. MEDINFO 86," pp. 887–890. Elsevier, Amsterdam/The Netherlands, 1986.

13. Pryor, T. A., Gardner, R. M., Clayton, P. D., and Warner, H. R. The HELP system. *J. Med. Syst.* **7**(2), 87 (1983).

14. McDonald, C. J., Hui, S. L., Smith, D. M., Tierny, W. M., Cohen, S. J., Weinberger, M., and McCabe, G. P. Reminders to physicians from an introspective computer medical record. A two-year randomized trial. *Ann. Int. Med.* **100**(1), 130 (1984).

15. Bleich, H. L., Beckley, R. F., Horowitz, G. L., Jackson, J. D., Moody, E. S., *et al.* Clinical computing in a teaching hospital. *N. Engl. J. Med.* **312**(12), 756 (1985).

16. Horowitz, G. L., Jackson, J. D., and Bleich, H. L. PaperChase: Self-service bibliographic retrieval. *J. Amer. Med. Assoc.* **250**(18), 2494 (1983).

17. Miller, R. A., McNeil, M. A., Challinor, S. M., Masarie, F. E. and Myers, J. D. The INTERNIST-1/QUICK MEDICAL REFERENCE project—Status report. *West. J. Med.* **145**, 816 (1986).

18. Covell, D. G., Uman, G. C., and Manning, P. R. Information needs in office practice: Are they being met? *Ann. Int. Med.* **103**, 596 (1985).

19. Reggia, J. A., and Tuhrim, S. (Eds.) "Computer-Assisted Medical Decision Making" (two volumes). Springer-Verlag, New York, 1985.

20. Greenes, R. A. Computer-aided diagnostic strategy selection. *Radiat. Clinics N. Amer.* **24** (No. 1), 105–120 (1986).

21. Szolovits, P., Patil, R. S., and Schwartz, W. B. Artificial intelligence in medical diagnosis. *Ann. Int. Med.* **208**, 80 (1988).