Using the Internet to Improve Knowledge Diffusion in Medicine+

William M. Detmer, M.D., M.Sc.*
Edward H. Shortliffe, M.D., Ph.D.

Section on Medical Informatics
Stanford University School of Medicine
Stanford, California 94305-5479

Communications of the ACM, August 1997 (in press)

+ This work was supported in part by the National Library of Medicine under a medical informatics training grant (LM-07033), a shared resource grant, CAMIS (LM-05305), and a high performance computing and communications research contract, InterMed (LM-43514).

* Current address: Department of Health Evaluation Sciences, University of Virginia

ABSTRACT

Despite advances in medical diagnosis and therapeutics, diffusion of new medical knowledge to practitioners is typically slow; frequently, it takes years for an important change in medical knowledge to make its way into daily clinical practice. Barriers to knowledge diffusion such as lack of immediate access to up-to-date information resources, ignorance of the availability of relevant information, and limited time for inquiry can be addressed by new network-based retrieval systems that deliver current and relevant information to the point of clinical need. These systems, using terms entered directly by clinicians or automatically by electronic medical-record systems, can rapidly display chunks of relevant information at the time of an information need and provide links to supporting evidence and analysis.

We describe three network-based systems that are in daily use addressing specific clinical information needs: WebMedline, a Web application that facilitates retrieval of the bibliographic and full-text medical literature; CliniWeb, an Internet search engine that retrieves high-quality medical Web sites indexed by a controlled medical vocabulary; and DXplain, a decision-support system that suggests possible diagnoses when given a list of pertinent clinical findings. We then describe MedWeaver, an application that integrates the functionality of these three systems by handling network interactions with each system, translating one system's vocabulary to another, and making available particular services from each system when there is an anticipated information need.

Despite the potential for improving knowledge diffusion in medicine using network-based retrieval systems such as those described here, further research and development is required. Innovations in the areas of content development, information science methods, and technology integration are required and can benefit from the expertise of professionals in medicine, information science, and computer science.

INTRODUCTION

Medical professionals are facing an information crisis. Medical knowledge is expanding and changing at an unprecedented rate, yet practitioners often do not become aware of important advances in a timely manner. While more than 360,000 articles are published in medical journals every year, knowledge diffusion to clinicians is typically slow. For instance, one study found that two years after wide publication, fewer than 50% of general practitioners knew that laser surgery could save the sight of some of their diabetic patients.

Barriers to knowledge diffusion are many [1]. They include clinicians' lack of access to up-to-date information resources, ignorance of the availability of relevant information, and lack of time for inquiry as well as poor organization of available information. Progress could be made if up-to-date information, relevant to clinicians' information needs, were rapidly available in all work settings—office, clinic, hospital ward, library, and home Retrieval systems, using concepts and modifiers entered directly by clinicians or automatically by electronic medical records systems, could rapidly display chunks of relevant summary information and provide links to supporting evidence and analysis.

To realize this vision a combination of content, information science methods, and technology is required. In the area of content, clinicians need access to a wide variety of information including the medical literature, expert summaries as found in textbooks and guidelines, information on medications and diagnostics tests, and procedural knowledge such as health-plan coverage or institutional policies. Much of this content is available today. For instance, MEDLINE, the 9-million-record bibliographic database produced by the National Library of Medicine, contains citations to the last 30 years of medical literature and has been available for three decades [2]. Knowledge-based systems that assist in diagnosis or therapy selection have, after a period of relative obscurity, become more commonplace. And full-text journal articles, textbooks, guidelines, and drug information are increasingly being made available in electronic form by print publishers.

But content alone will not solve the information needs of practicing clinicians. Information science methods are needed to address problems such as how to structure content to achieve optimal retrieval, how to select a resource to best answer a particular question, how to integrate information from several sources into one consistent view, and how to develop search interfaces that are more powerful than generic interfaces because they leverage users' domain knowledge. The field of medical informatics—a multi-disciplinary field that blends biomedicine, decision science, information science, psychology, and computer science [3]—is now focusing on these important areas.

To deliver current, well-structured, high-quality content to the desktops of practicing clinicians, a final component is needed—technology. High-speed data networks, standard protocols, open-systems architectures, and cross-platform applications are necessary technology components. Progress in this area has been accelerated by the emergence of the World Wide Web; universal access to multimedia information via a single application that is available on a variety of computing platforms, once a distant hope, is now closer to reality. Still challenges lie ahead in the areas of high-speed data transmission and data security.

EXAMPLE MEDICAL APPLICATIONS ON THE INTERNET

To exemplify how content, information science methods, and technology can be harnessed to meet clinician's information needs, we present three examples. These systems, each in daily use by clinicians around the country, show how various types of medical knowledge can be delivered to practitioners using Internet technology. A final system, one that integrates functionality from the three example systems, shows how components of these widely distributed systems can be combined in ways that better meet information needs.

WebMedline

WebMedline is a Common Gateway Interface (CGI) application developed at Stanford University [4] that facilitates searching of the medical literature via a standard Web browser. WebMedline uses the Hypertext Markup Language (HTML) to display bibliographic data from MELVYL MEDLINE, the University of California's implementation of the MEDLINE database, as well as hypertext links to corresponding full-text articles on the Internet.

Before WebMedline existed, clinicians accessed MELVYL MEDLINE by initiating a terminal-based Telnet session, logging on to the MELVYL mainframe host, navigating the opening prompts, issuing queries in the MELVYL MEDLINE query language, and displaying or printing results using a specialized language. In contrast, WebMedline users enter queries into HTML forms and view results in hypertext (Figure 1). Users enter text into predefined fields, such as Author, Title, Journal, and Keyword and can choose from pop-up menus the database years, the between-field Boolean operator, the display type, and number of citations to retrieve. Finally, they can choose to constrain the search by standard limiters such as "English only," "Human subjects only," and "Publication Type" (e.g., editorial).

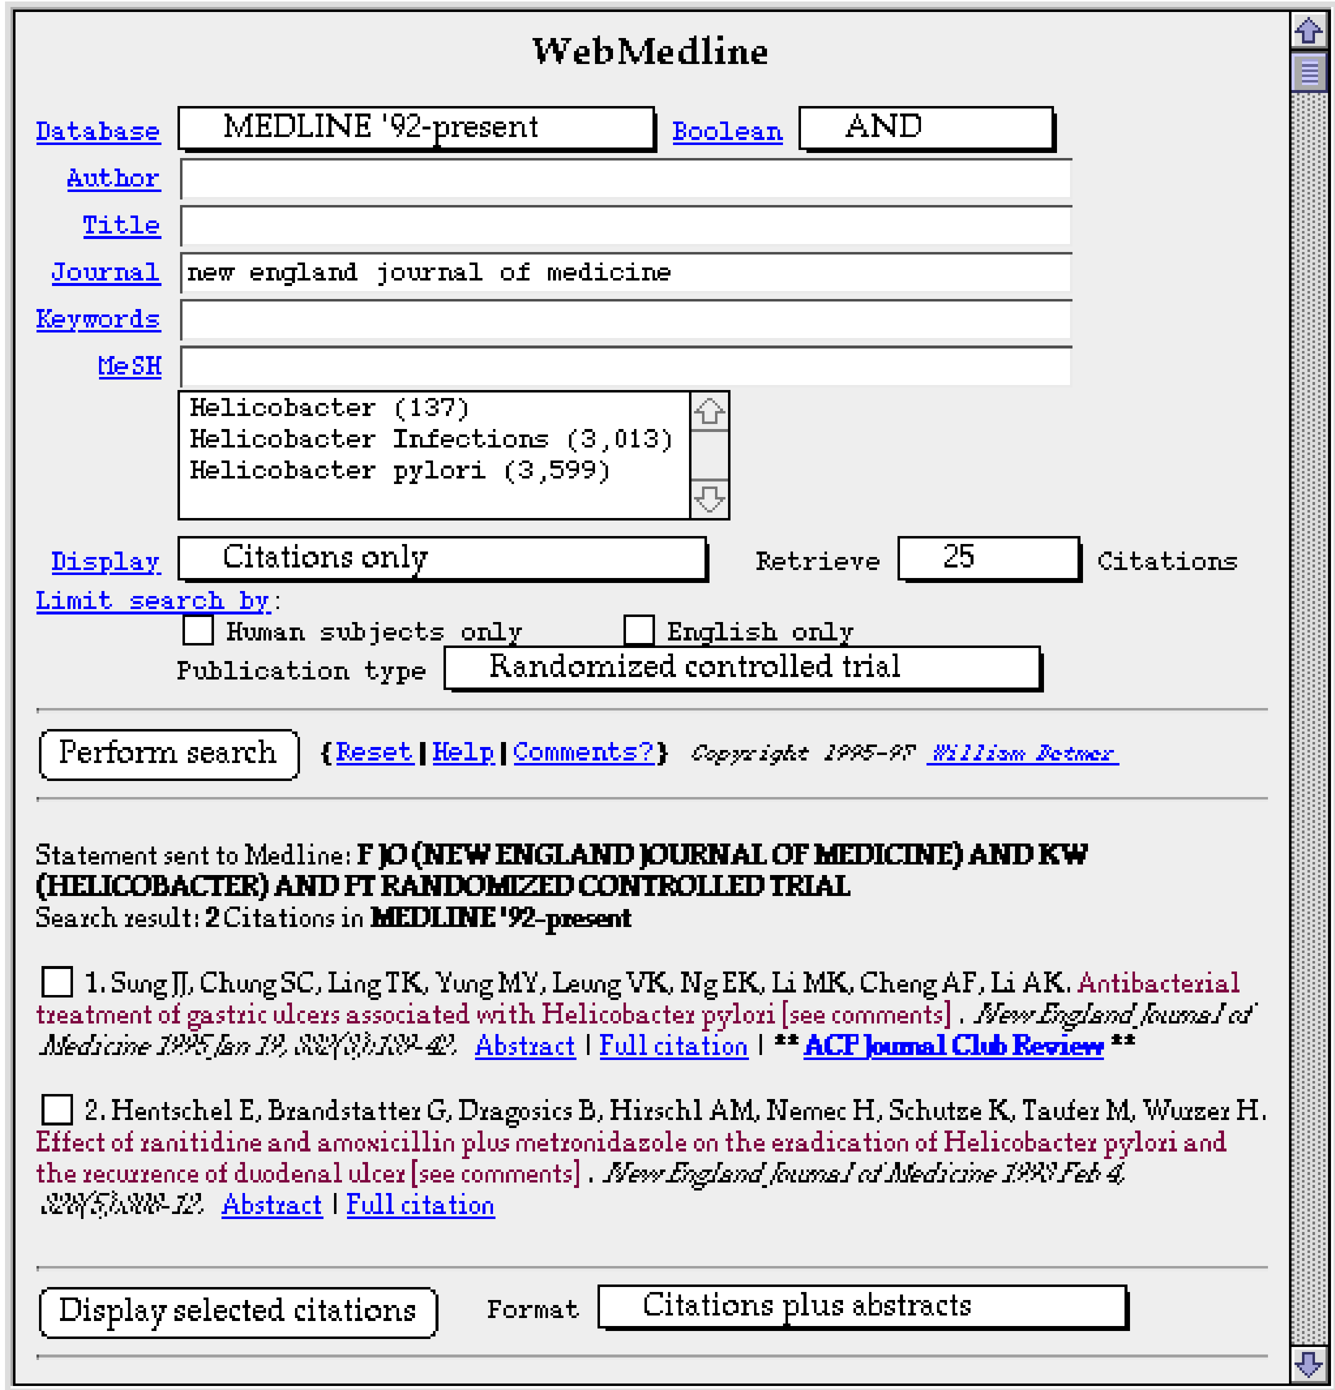

Figure 1. Result of a WebMedline search. These results appear after the user selects the Publication Type "randomized controlled trial" and types "new england journal of medicine" in the Journal box and "helicobacter" in the Keyword box. The top portion of the page shows a new HTML form, which the user can use to further refine the query. For instance, the user can choose one of the MeSH terms that was returned as a by-product of the prior search. The bottom portion of the page contains the results of the search. Note that if a full-text article, such as an *ACP Journal Club* review, exists for a particular citation, a hyperlink is created dynamically. Also note that a user can select a number of citations by checking the box in front of the article and redisplay the selected citations in another form (e.g., with their abstracts).

When a search session is initiated, WebMedline first authenticates the user and then sets up a search session with the MELVYL MEDLINE database. With each query, WebMedline modules transform user input into a legal MELVYL MEDLINE search statement (Query Formulator), retrieve data from MEDLINE (Query Manager), and mark up the results in HTML (Display Manager).

The Query Formulator composes a legal MELVYL MEDLINE search statement by removing stop words from the input, qualifying input terms with appropriate field descriptors, and then joining these qualified input terms with Boolean operators. In addition it attempts to map keywords to controlled indexing terms by first stemming words to their roots and then querying a thesaurus function provided by MELVYL MEDLINE. The resulting controlled indexing terms, called Medical Subject Headings (MeSH), are later shown to the user to assist them in choosing a more precise query term. Next, the Query Manager retrieves data from MEDLINE by connecting to the open search session, issuing the search statement, and requesting output in a desired format. The Display Manager then transforms into HTML the ASCII output of MELVYL MEDLINE, creating hypertext links to full-text documents when it finds a corresponding full-text URL in the WebMedline Link Database.

WebMedline has been in daily use at the University of California, San Francisco and Stanford University since February 1995. In 2 years, over 300,000 sessions have been logged and usage has doubled every 4–6 months. WebMedline provides links to the journals *Science* and the *Journal of Biological Chemistry* , the literature-review publication *ACP Journal Club*, and guidelines from the National Institutes of Health. WebMedline has also been licensed to Ovid Technologies, Inc., an information provider to large biomedical institutions and consortia, and has been incorporated into their Ovid Web Gateway product.

CliniWeb

CliniWeb (http://www.ohsu.edu/cliniweb/) is a retrieval system developed at Oregon Health Sciences University [5] to help health practitioners find useful medical information on the World Wide Web. Human indexers trained in medicine explore the Web, select quality medical resources, assign to them appropriate controlled-vocabulary terms from the Medical Subject Headings (MeSH), and place them in a database. This database enables two access methods: a browsing interface that presents the MeSH hierarchy with all the clinical resources associated with each term and a searching interface that assists users in mapping natural language queries to MeSH terms and then viewing associated Web resources.

Although the WWW contains a vast array of medical information, clinicians are often frustrated by the difficulty in quickly finding authoritative and relevant information. In the vast information space of the Web, clinicians are looking for resources that are practitioner-oriented, produced by reputable sources, and that cover a specific topic in medicine. Without a resource such as CliniWeb, clinicians have to browse or search one of the major Internet search systems. Systems such as Alta Vista or Lycos are poor choices because they are not discriminatory in what they index, and their index methods are word-based and not content-based. Systems such as Yahoo!, which apply some degree of content filtering and assignment to a taxonomy, are an improvement, but still do not provide indexing of individual pages nor indexing at a sufficient level of detail to be useful to a practitioner.

CliniWeb attempts to meet the needs of practitioners by creating a database of human-filtered, practitioner-oriented Web resources that are indexed using MeSH. Medically trained indexers browse web sites of medical schools, governmental health agencies, and medical publishers as well as search databases of Internet sites. When a useful site is identified, information such as page title, institution, and URL are entered into CliniWeb. In addition, words that describe the content of the site are entered into the concept-mapping engine SAPHIRE [6] and a ranked list of possible MeSH matches are displayed. SAPHIRE performs this mapping by matching entered terms against medical concepts and their synonyms as provided by the National Library of Medicine's Unified Medical Language System (UMLS) [7]. From the list of possible MeSH terms, the indexer chooses the most appropriate one and enters it into CliniWeb. Occasionally, MeSH terms need to be manually assigned because of poor automated indexing.

Clinicians can explore CliniWeb in two ways. First, they can browse the MeSH hierarchy by disease or anatomical site. For instance, they can find Internet resources on "heart attack" by traversing the disease hierarchy Cardiovascular Disease/Heart Disease/Myocardial Ischemia/Myocardial Infarction and exploring associated Internet links. Second, they can enter terms into the CliniWeb search interface and retrieve matching MeSH terms ranked by likelihood. As with the indexing step, CliniWeb uses SAPHIRE to map users' free-text entries to MeSH. Users then select the desired MeSH term and receive a list of clinically useful URLs.

In addition to the browsing and searching interfaces provided by CliniWeb's developers, an interface for remote procedure calls exists. This allows application developers from other institutions to integrate CliniWeb services into their own application, as we describe later.

Currently, CliniWeb contains over 10,000 clinically-useful Internet sites and is used approximately 10,000 times per month. Like other indexing services on the Web, future challenges include maintaining existing links, finding and adding new links, storing more meta information that will give users additional context before traversing a link, and improving the mapping from natural language to controlled vocabulary terms.

DXplain

DXplain is a diagnostic decision support system developed at Massachusetts General Hospital [8]. One of the goals of DXplain is to generate a list of possible diagnoses from a group of clinical findings. Users enter clinical terms such as cough, fever, and nightsweats and receive a list of possible diagnoses ranked by likelihood. Additionally, users can explore other information in the knowledge base such as the frequency that a symptom occurs in a particular disease or the chance that a particular disease is present give a certain finding.

Currently, DXplain represents information on approximately 2000 diseases, 4700 clinical findings, and 65,000 interrelationships. A disease profile consists of a set of clinical findings that the developers have decided are relevant. For each finding listed for a disease, three attributes are stored: term frequency, term evoking power, and term importance. *Term frequency* quantifies how often a finding occurs in a disease, *Term* *evoking power* states how strongly the finding supports the diagnosis of the disease and *term importance* measures how consequential is the finding. DXplain also stores for each disease the *disease prevalence* or the baseline likelihood of a disease in the general population and *disease importance*, or how consequential is the disease.

DXplain generates a list of possible diagnoses using a pseudo-probabilistic algorithm. It first evaluates the term importance and term evoking strength of each finding-diagnosis pair and then calculates a summary score for each disease. A disease score is most influenced by positive findings that have high term evoking strength. Findings with intermediate evoking strengths and high term importance contribute moderately to the summary score. After DXplain evaluates each clinical finding it displays the highest ranked diagnoses divided into "common diseases" and "rare or very rare diseases".

For many years, DXplain has been available over the Internet via a Telnet interface. Recently, it was rearchitected as a knowledge base server that provides services for a new Web-based front end (WebDXplain) and other Internet-based applications that require diagnostic services. The new DXplain server accepts UNIX socket connections, recognizes standard function calls, and returns output in a standard format. For instance, the function call *getDDx("password","cough","fever", "nightsweats")* returns a list of likely diagnoses given a password and a list of clinical findings.

DXplain is currently used by more than 200 individuals and institutions around the US. Because of concerns that proper interpretation of its output requires medical knowledge, DXplain is available only to physicians. DXplain is available through a stand-alone Web interface, is integrated with MEDLINE literature searching (see below), and is being evoked using patient data extracted from electronic medical records [9]. Because its diagnostic accuracy is imperfect [10], it is used primarily as a "memory jogger" and an educational tool.

MedWeaver – An Integrated Example

The three stand-alone systems described here can be useful to clinicians, but they are limited because they perform only a specific task—search of the literature, search of the Web, or diagnosis of a patient. The potential of the Internet for healthcare and other professional disciplines is that such separately maintained services can now be integrated and made available in ways that better meet users' information needs.

As an example, we developed MedWeaver, a research prototype that shows how the three systems, maintained at different sites around the country, can be combined (Figure 2). The result is a decision support system that performs assisted searches of the medical literature and directs users to useful Internet sites.

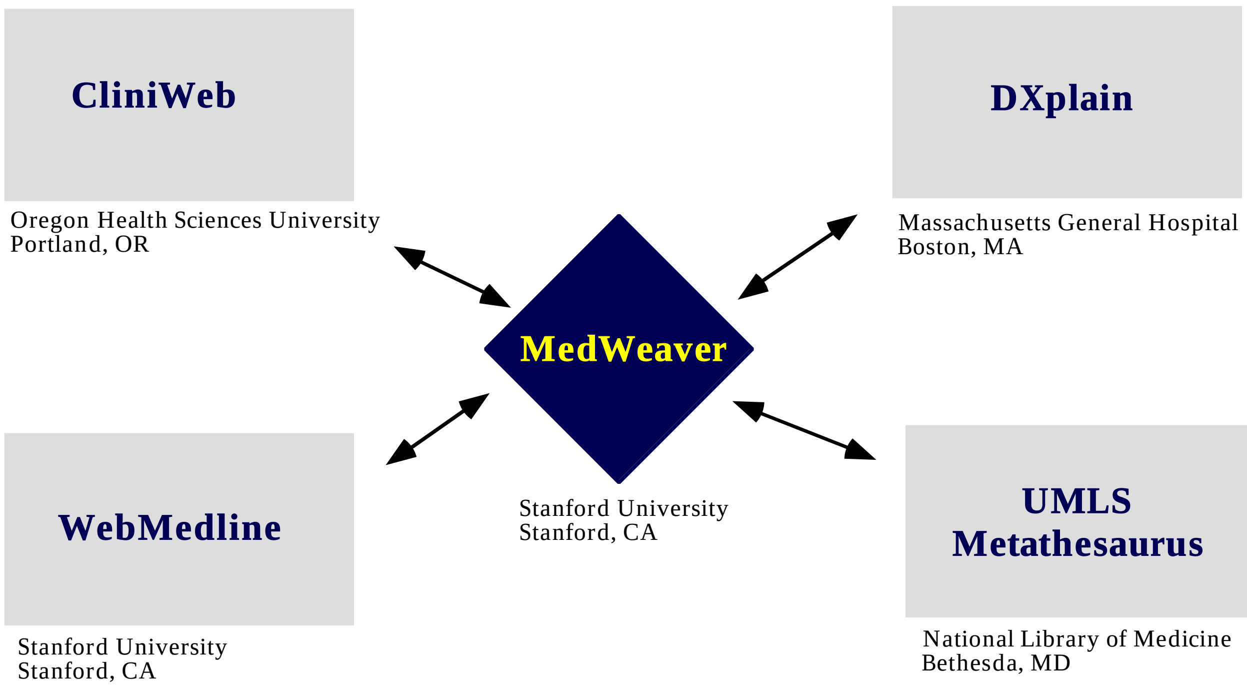

Figure 2. MedWeaver architecture. MedWeaver is a Common Gateway Interface application that integrates diagnostic decision support, literature searching, and retrieval of clinically useful Web sites. MedWeaver contains no content itself, but relies on content provided by the developers of CliniWeb, DXplain, and WebMedline and translation services provided by the UMLS Metathesaurus.

The user begins a MedWeaver session by entering clinical findings. MedWeaver uses DXplain functions to map these findings to controlled vocabulary terms and then to produce a ranked list of possible diagnoses (Figure 3). It then displays for any diagnosis on the list: (1) a summary of that disease (a DXplain function - Figure 4), (2) an explanation as to why the diagnosis appears on list (a DXplain function - Figure 5), (3) an assisted search of the medical literature (a WebMedline function- Figure 6), or (4) a list of clinically useful Internet sites (a CliniWeb function- Figure 7).

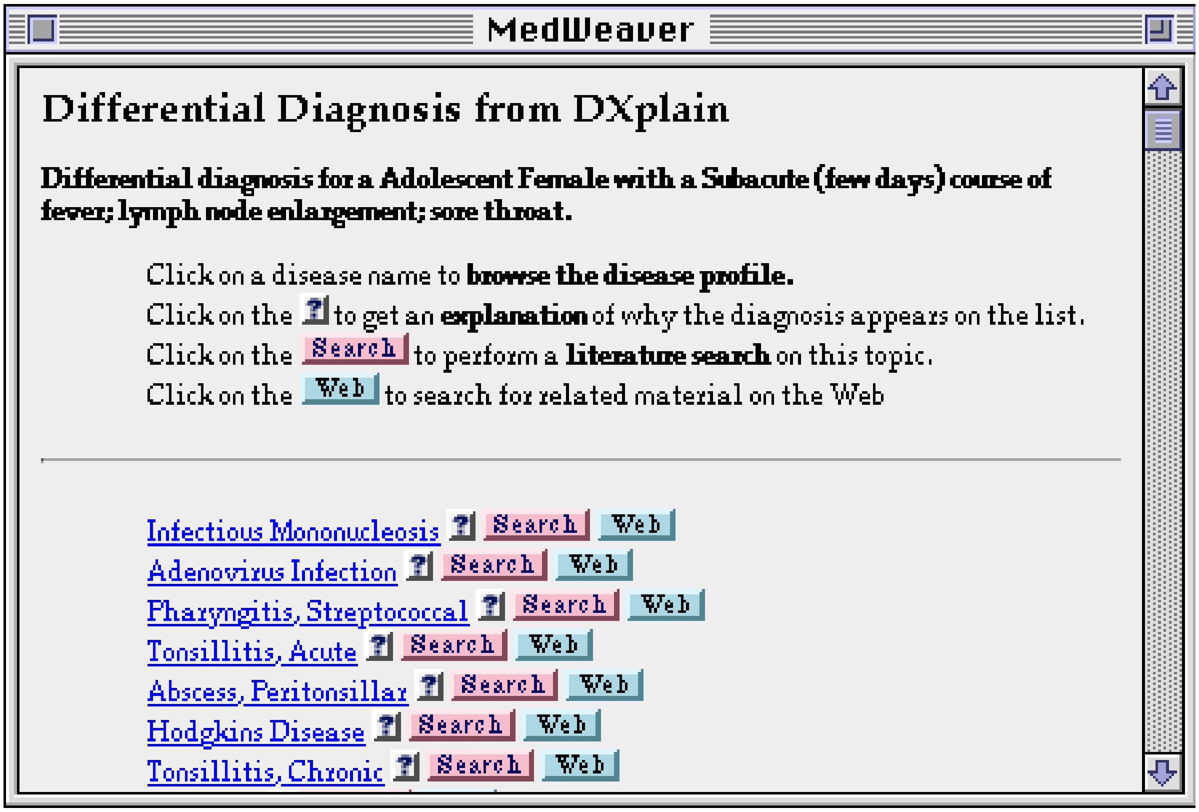

MedWeaver

Differential Diagnosis from DXplain

Differential diagnosis for a Adolescent Female with a Subacute (few days) course of fever; lymph node enlargement; sore throat.

Click on a disease name to **browse the disease profile.**
Click on the ? to get an **explanation** of why the diagnosis appears on the list.
Click on the Search to perform a **literature search** on this topic.
Click on the Web to search for related material on the Web

Infectious Mononucleosis ? Search Web
Adenovirus Infection ? Search Web
Pharyngitis, Streptococcal ? Search Web
Tonsillitis, Acute ? Search Web
Abscess, Peritonsillar ? Search Web
Hodgkins Disease ? Search Web
Tonsillitis, Chronic ? Search Web

Figure 3. Diagnoses generated by DXplain. MedWeaver uses a DXplain function to retrieve a list of possible diagnoses for the clinical findings "adolescent female", "subacute (few days)", "fever", "lymph node enlargement", and "sore throat." From this page, users can retrieve a disease profile, view an explanation, perform an assisted search of the medical literature, or retrieve a list of related Web resources.

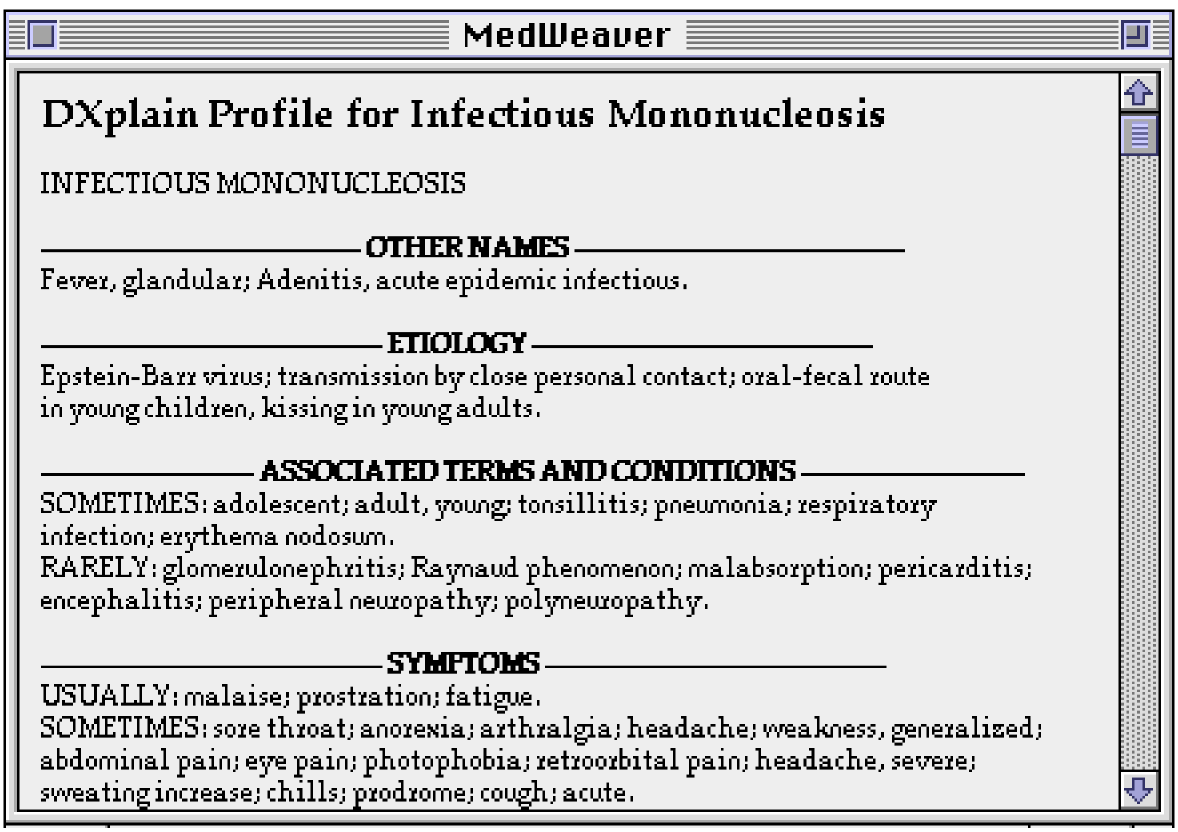

MedWeaver

DXplain Profile for Infectious Mononucleosis

INFECTIOUS MONONUCLEOSIS

OTHER NAMES
Fever, glandular; Adenitis, acute epidemic infectious.

ETIOLOGY
Epstein-Barr virus; transmission by close personal contact; oral-fecal route in young children, kissing in young adults.

ASSOCIATED TERMS AND CONDITIONS
SOMETIMES: adolescent; adult, young; tonsillitis; pneumonia; respiratory infection; erythema nodosum.
RARELY: glomerulonephritis; Raynaud phenomenon; malabsorption; pericarditis; encephalitis; peripheral neuropathy; polyneuropathy.

SYMPTOMS
USUALLY: malaise; prostration; fatigue.
SOMETIMES: sore throat; anorexia; arthralgia; headache; weakness, generalized; abdominal pain; eye pain; photophobia; retroorbital pain; headache, severe; sweating increase; chills; prodrome; cough; acute.

Figure 4. Portions of a disease profile for "infectious mononucleosis" provided by DXplain.

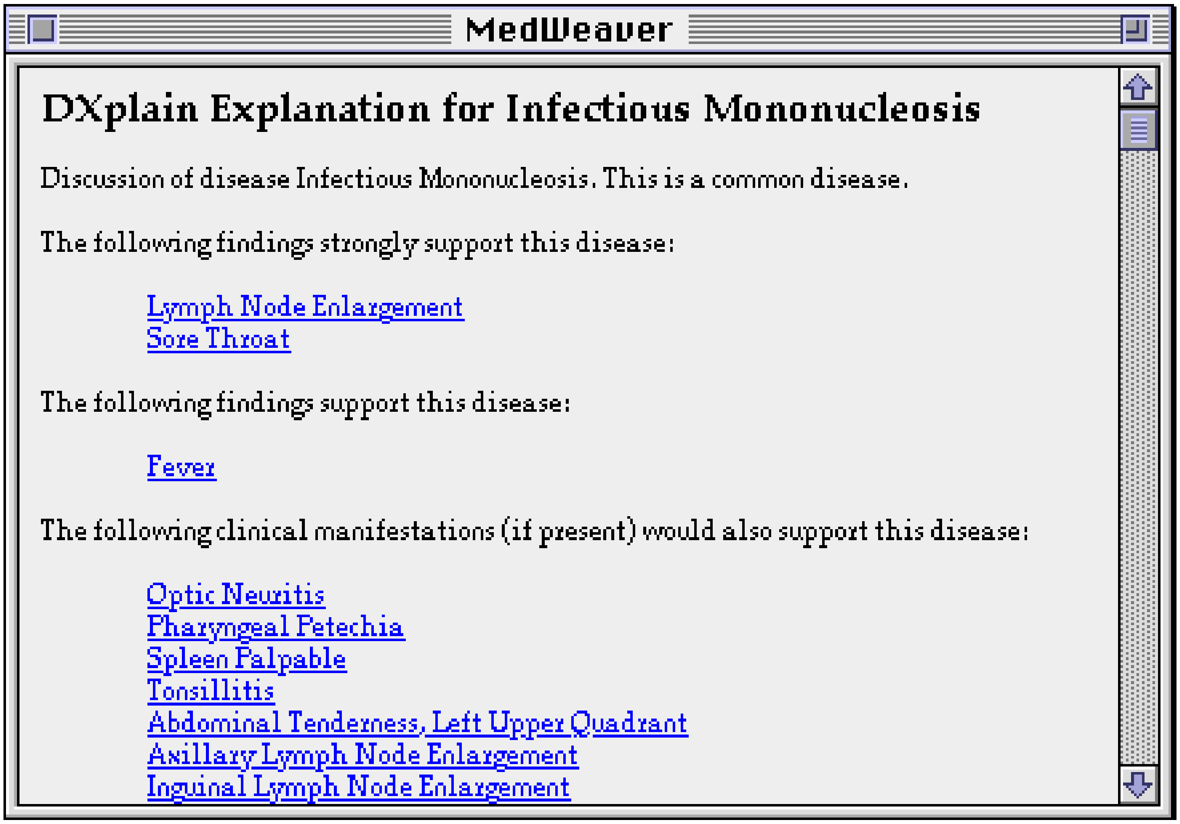

Figure 5. An explanation provided by DXplain for why "infectious mononucleosis" appears on the list of possible diagnoses.

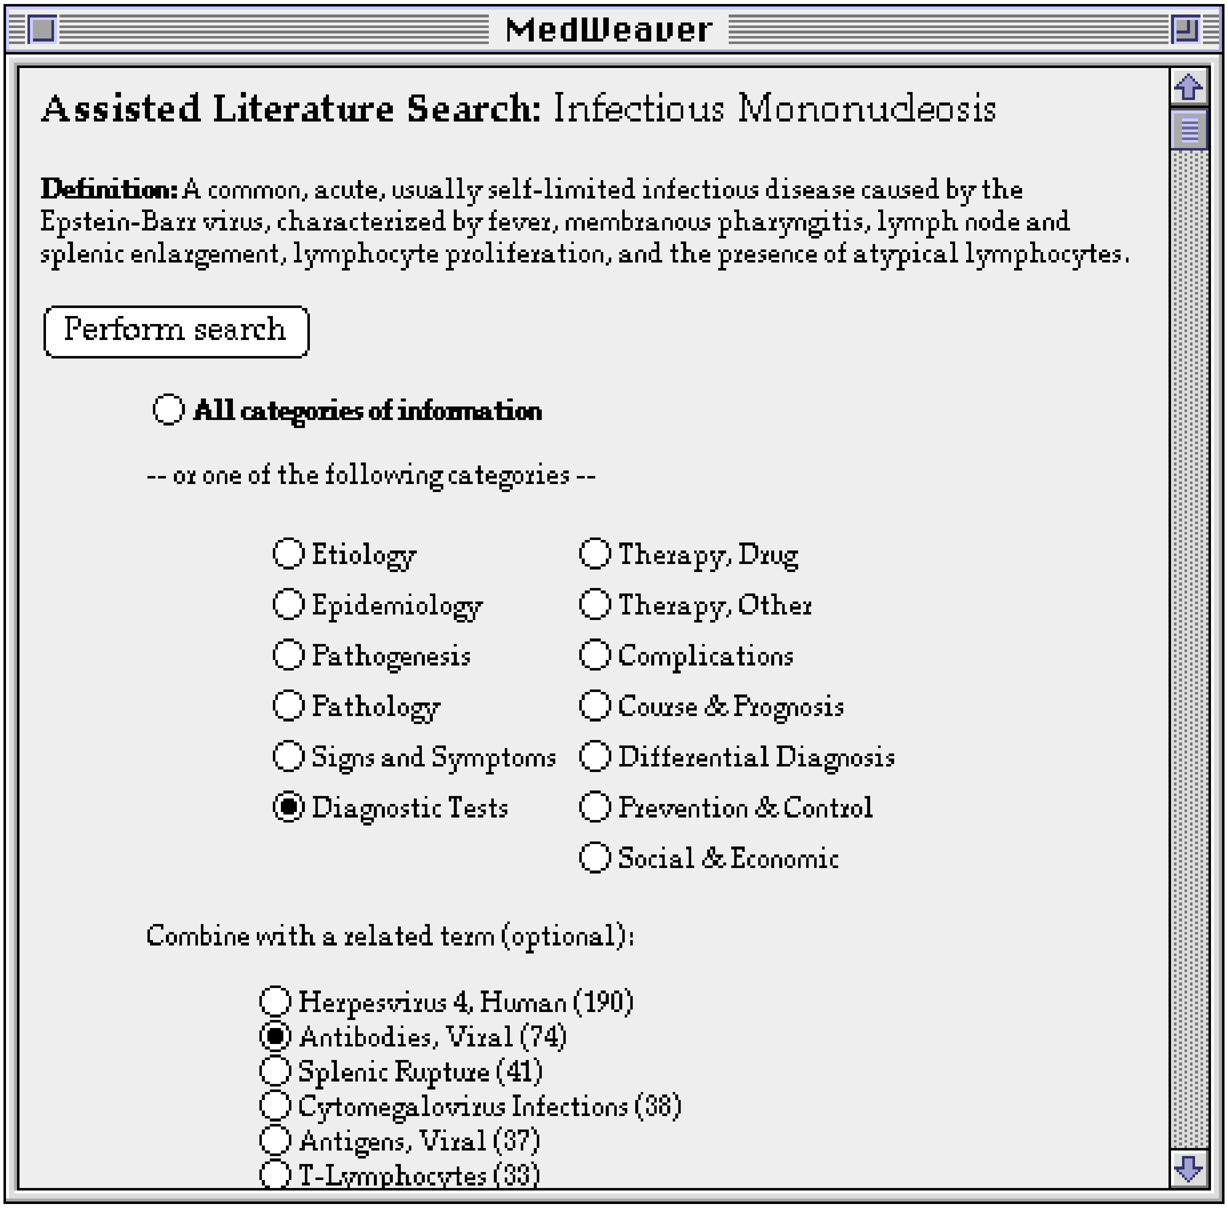

Figure 6. An assisted search of the medical literature . MedWeaver first uses the UMLS Metathesaurus to find the closest MeSH term for the DXplain disease name "infectious mononucleosis." Using this MeSH term, it then retrieves from the Metathesaurus both the term definition and the terms that co-occur in the MEDLINE database. MedWeaver displays this information giving users the ability to limit searches to discrete classes of information such as "diagnostic tests" or "prevention and control". Once a user has specified the desired classes of information and co-occurring terms, MedWeaver generates a sophisticated query of the MEDLINE database and retrieves bibliographic citations and full-text links from WebMedline.

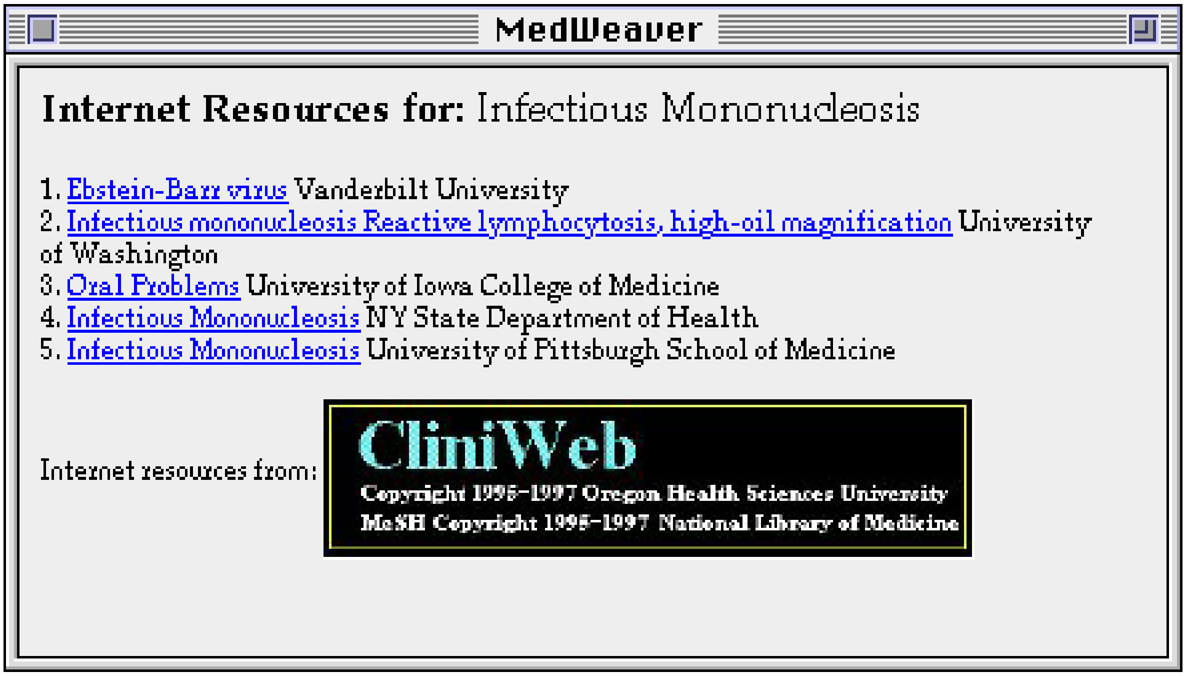

Figure 7. List of clinically useful Internet sites provided by CliniWeb. MedWeaver finds the closest MeSH term for the DXplain disease "infectious mononucleosis," requests from CliniWeb URLs indexed with the specific MeSH term, and then displays the URLs for the user.

To do this, the MedWeaver Common Gateway Interface application takes the output from one information system, transforms it into an intelligent query of another system, manages the interaction with the remote system, and displays the results in a way that anticipates current and future information needs. A major obstacle to this integration is the translation of one system's vocabulary to another. Fortunately, the UMLS Metathesaurus, a collection of medical vocabularies tied together by the concepts they share, was developed by the National Library of Medicine for this purpose [7]. Using common translation functions provided with the Metathesaurus and extended by the developers, MedWeaver can query the Metathesaurus with a DXplain disease name and retrieve the closest matching term from the MeSH vocabulary. MedWeaver then uses this MeSH term to retrieve bibliographic citations from WebMedline or Internet resources from CliniWeb.

DISCUSSION

The above examples show how clinician's information needs can be addressed with a combination of content, information science methods, and Internet technology. Although studies have not yet quantified the benefits of these new types of systems, there are at least four reasons why they may begin to help the knowledge diffusion problem.

First, Internet and Web technology has simplified access to information. Clinicians can now move among office, clinic, hospital ward, library and home and retrieve the information they need through a single, common, cross-platform application. Second, new interfaces shield clinicians from the complexity of retrieval systems. No longer do clinicians need to learn the details of Internet protocols, the physical location of useful resources, or even the query language necessary to retrieve information. Instead, resources can be navigated by following conceptual links or searched with simplified interfaces.

Third, currency of information has been improved. Instead of turning to out-of-date textbooks on the office bookshelf, clinicians can now access the latest fact, summary, or opinion directly from the source. Finally, integration of disparate and widely distributed information is now possible. Clinicians can access information products that combine useful components of previously separate information sources. For instance, clinicians can retrieve journal or textbook information while reviewing patient data from the electronic medical record.

Content Challenges

Despite this potential, many challenges still exist. In the content area, the biggest challenge is how to transform medical publishing from print-based to electronic. For more than 50 years, medical publishers have enjoyed a stable technical and economic environment. Now, with the advent of digital publication, many publishers are finding themselves without the knowledge, vision, or internal processes to guide them in this new era. A crucial decision for many publishers is whether they should remain only content creators or instead become technological innovators in the areas of access, distribution, navigation and search. Many also need to retool their editorial and document management processes to account for both digital and print products. Finally, publishers are struggling to find an economic model for digital publishing that will assure their survival.

Information Science Challenges

The potential of medical knowledge dissemination over networks also presents many exciting information science challenges. Many of these challenges are similar to those in other professional domains such as law, engineering, or computer science [11]. Progress is likely if those with medical and information science expertise work collaboratively on these challenges.

One such challenge is how to structure old and new content so that retrieval is enhanced. Traditional medical retrieval systems use word-based representation methods that can lead to sub-optimal retrieval performance. Newer representation methods such as those that add contextual information to portions of documents may help improve retrieval relevance by focusing retrieval in only relevant semantic regions [12].

Another challenge is how to assist users with resource selection–i.e. which resource should be used to find the answer to a particular question. This assistance is necessary because clinicians are not aware of all the resources available to answer a particular question, nor do they have time to assess which resources is best. To automate the resource-selection process, systems must have knowledge of what questions each resource can answer. Components of such knowledge include the scope, depth, intended audience, currency of information, and reputability of each source. Researchers working with the National Library of Medicine have begun to build such a knowledge resource, called the Information Sources Map (ISM). A component of the Unified Medical Language System, the ISM contains not only descriptive information about the contents of the resource, but also procedural information needed to successfully connect to and retrieve information from a resource [7]. A goal of this work is to create source descriptions that are reusable and sharable among developers.

Another information science challenge is how to organize the output of these systems so they better meet users' information needs. Methods of interest include (1) relevancy ranking based on word characteristics (e.g., part-of-speech, location in the document, or frequency in the document relative to frequency in the collection), (2) relevancy ranking based on user characteristics (e.g., weighting sources and types of information based on a clinician characteristics such as specialty or practice type), and (3) clustering and labeling of similar documents based on statistical correlations between words, phrases, or concepts in the retrieved set. The goal is to present busy clinicians with material only if it is highly relevant.

Yet another challenge is how to develop unified interfaces to disparate resources. Clinicians will want to retrieve information from resources with different structures, yet will prefer to interact with information sources through a single, consistent interface. Because the structure and retrieval paradigms of information sources will vary widely, they need to be unified using some model that encompasses aspects of all the resources and yet is independent of the underlying structure of a particular resource. Fortunately, medical information has conventional categorizations that are shared widely among practitioners and that can be used to compose a commonly understood model of medical information. For instance, a simple model of medical diseases might contain definition, clinical findings, tests, therapies, and complications. Such a medical information model can be used to create a domain-specific query interface and to organize the display of results. To accomplish retrieval and display from disparate sources, resource translators transform queries expressed in the common medical information model to queries in the language of the individual resource.

Technology Challenges

In the technology area, the dissemination of medical knowledge to clinicians has additional challenges. First, connections to high-speed data networks are needed to provide clinicians with access to medical text, images, audio, and video. Many hospitals, clinics, and private offices do not have Internet access or even local area networks because they have not yet found compelling reasons to invest money and personnel in these technologies. However, as the utility of resources such as the electronic medical record and network-based information increases, medical institutions will likely increase investment in the needed network infrastructure.

One of the concerns about giving clinicians access to high-speed networks is how to assure security and confidentiality. While access to information within the institution or across the globe may be beneficial to clinicians, it might also expose clinicians and their institutions to security violations. Already, Web sites are capable of capturing information on the user's location, email address, and query, as well as other information that the user may unwittingly provide. Details of a patient's condition as entered into a retrieval system may result in a breach of patient confidentiality if the data are captured and associated with a particular patient. Advances in security technologies such as firewalls, proxies, and data encryption are essential to keeping medical information systems secure.

CONCLUSION

There is exciting potential to improve the flow of medical information to practitioners by the careful application of digital content, information science methods, and technology. The hope is that improved knowledge diffusion will ultimately boost the quality of medical care delivered to patients. Many challenges exist to make this vision a reality.

Rapid progress is likely if professionals with expertise in medicine, information science, and computer science collaborate to meet these challenges.

ACKNOWLEDGMENTS

We wish to thank William R. Hersh of Oregon Health Sciences University and G. Octo Barnett of Massachusetts General Hospital for providing access to their systems and for contributing valuable feedback during prototype development.

REFERENCES

1. Smith R. What clinical information do doctors need? BMJ 1996;313(7064):1062-1068.
2. Lindberg DA, Siegel ER, Rapp BA, Wallingford KT, Wilson SR. Use of MEDLINE by physicians for clinical problem solving. Jama 1993;269(24):3124-9.
3. Greenes RA, Shortliffe EH. Medical informatics. An emerging academic discipline and institutional priority. Jama 1990;263(8):1114-20.
4. Detmer WM, Shortliffe EH. A model of clinical query management that supports integration of biomedical information over the World Wide Web. Proceedings of the Annual Symposium on Computer Applications in Medical Care 1995:898-902.
5. Hersh WR, Brown KE, Donohoe LC, Campbell EM, Horacek AE. CliniWeb: Managing clinical information on the World Wide Web. J Am Med Inform Assoc 1996;3(4):273-280.
6. Hersh WR, Hickam DH. Information retrieval in medicine: the SAPHIRE experience. J Am Soc Info Sci 1995;46:743-7.
7. Lindberg DA, Humphreys BL, McCray AT. The Unified Medical Language System. Methods Inf Med 1993;32(4):281-91.
8. Barnett GO, Cimino JJ, Hupp JA, Hoffer EP. DXplain. An evolving diagnostic decision-support system. Jama 1987;258(1):67-74.
9. Cimino JJ, Socratous SA, Clayton PD. Internet as clinical information system: application development using the World Wide Web. J Am Med Inform Assoc 1995;2(5):273-84.
10. Berner ES, Webster GD, Shugerman AA, et al. Performance of four computer-based diagnostic systems. N Engl J Med 1994;330(25):1792-6.
11. Special Issue on Digital Libraries. Communications of the ACM 1995;38(4).

12. Purcell GP, Shortliffe EH. Contextual models of clinical publications for enhancing retrieval from full-text databases. Proceedings of the Annual Symposium on Computer Applications in Medical Care 1995:851-7.