Trusted and Relevant Medical Knowledge: The Promise of Information Retrieval in Biomedicine

Eric Z. Silfen MD, MSHA; and Chintan O. Patel, MS

ABSTRACT

As the world of medicine becomes increasingly digitized, the Web has become a de facto ***resource for physicians to quickly glean pertinent clinical information to carry out diagnostic and therapeutic decisions. At present, physicians face the dual challenge of judging the relevance of the information and trusting its Web source. This paper proposes a trust-relevance framework for conceptualizing computer-accessed medical information resources, a set of criteria for evaluating these information resources, and descriptions of a sample of available online resources. It also presents a usable framework for evaluating information retrieval innovations and explains the different capabilities of representative information retrieval tools and applications. By demystifying the concepts associated with information resources, search engines, and retrieval tools, and presenting a reasonable view of current opportunities as well as future possibilities, the authors hope to provide guidance so physicians can more rapidly adopt innovative computer-assisted search tools for acquiring information that facilitate patient care decision-making.***

KEYWORDS

■ Information retrieval ■ Information extraction ■ Knowledge discovery ■ Search engine ■ Relevance ■ Trust

A recent letter to the editor in the *New England Journal of Medicine*[1] ran the headline: "… And a Diagnostic Test Was Performed." The clinicopathological case under discussion described an infant with diarrhea; an unusual rash (alligator skin); multiple immunologic abnormalities, including low T-cell function; tissue eosinophilia (of the gastric mucosa) as well as peripheral eosinophilia; and an apparent X-linked genetic pattern (several male relatives died in infancy). The Fellow diagnosed a rare syndrome known as IPEX (immunodeficiency, polyendocrinopathy, enteropathy, X-linked). Genetic testing on the baby revealed a mutation in the FOXP3 gene, confirming the diagnosis.

When the visiting professor inquired of the Fellow how the diagnosis was made, the following dialogue ensued:

Visiting Professor: "How did you make that diagnosis?"

Fellow: "I had the skin biopsy report and a chart of the immunologic tests. So I entered the salient features into Google, and it popped right up."

As the world of medicine becomes increasingly digitized, accessing medical information is merely a matter of making a few mouse clicks on a computer screen and initiating a search. As evident from this conversation, the Web has become a *de facto* resource for physicians to quickly glean pertinent clinical information to carry out diagnostic and therapeutic decisions for patients.

The irony and gravity of the situation is that anyone can publish any information on the Web, which raises an important issue when clinicians use the Web to assist with medical practice—how can a physician judge the relevance of the information and trust its source?

Four Challenges

The richness of the clinicopathological conference dialogue emphasizes the numerous challenges associated with weaving new technology into the fabric of medicine. Besides highlighting the issues associated with changing medical practice in general, the myths associated with medical computing emerge as well. To begin, it is worth exploring four challenges posed by a combination of the medical information explosion[2,3] and the role of computers in making this information usable in clinical practice.

The first challenge is diffusing computer innovations into the methods of medicine. Diffusion of innovation among individuals within a social system occurs asynchronously—that is, all individuals do not adopt an innovation at the same pace, but in a well-defined sequence.[4]

However, theory notwithstanding, empirical evidence is most instructive. For example, regarding the diffusion of thrombolytic therapy for the treatment of acute myocardial infarction into mainstream medical practice, the British experience has been extensively reviewed.[5,6] Clearly, and unfortunately, the authors point out that after the initial six-year period of enthusiasm for thrombolytic therapy as a beneficial intervention, an additional eight-year climb was required to reach an acceptable plateau of only 65 percent utilization. After the initial excitement, factors associated with this slowdown in therapeutic acceptance included converting individuals to the new clinical paradigm, concern over possible adverse side-effects, patient-management systems being slow to adapt to the technology, and the threat of negative affects on hospital budgets. For physicians, it is not hard to draw parallels with the diffusion of computer technology into clinical practice.[7,8,9,10,11]

The second challenge concerns assimilating computer technology into the physician's diagnostic and therapeutic armamentarium to catalyze a positive transformation within the art of medicine. Regardless of the power of bits and bytes, the humanness of clinical decision-making must remain tantamount. In addition, physicians must believe that these new computer applications improve the inherently human care of patients. Conversely, if physicians do not shoulder this responsibility, they will fall victim to their worst fear—a degeneration of medical practice into technology in search of patients.

The third and fourth challenges acknowledge the humanness of physicians. For a discipline that prides itself on the rigorous application of the scientific method, admitting that arriving at the correct diagnosis sometimes depends on luck recognizes physician fallibility.[12] Although all physicians strive to know everything that is humanly possible regarding their area of professional expertise, only by recognizing their knowledge gap can physicians appreciate the potential of new computer applications and train themselves to use these tools wisely.

"For a discipline that prides itself on the rigorous application of the scientific method, admitting that arriving at the correct diagnosis sometimes depends on luck recognizes physician fallibility."

Beyond humility, there is the issue of age itself. The adaptability and flexibility of youth to avidly embrace the change associated with computers must be combined with the technological skepticism and wisdom gained through time and experience. Otherwise, all physicians will falter in their quest to rapidly uncover new knowledge, and patients will not benefit from the timely use of the best evidence for medical care.

Given these challenges, what follows is a trust-relevance framework for conceptualizing computer-accessed medical information resources, a set of criteria for evaluating these information resources and descriptions of a sample of available online resources. By reconciling the medical information needs faced by physicians and presenting a reasonable view of current opportunities as well as future possibilities, this paper aims to provide guidance that enables physicians to more rapidly adopt innovative, computerized search tools to retrieve information that facilitates best-evidence medical decision-making for care.

Trust and Relevance are Essential

Trust and relevance are essential characteristics for all medical information that is used for clinical decision-making. The Merriam-Webster online dictionary defines trust as: "Assured reliance on the character, ability, strength, or truth of someone or something; one in which confidence is

placed."[13] For the computing world, trust implies a systematic assessment of the physician's confidence in the information retrieval system to display reliable medical information over the Web. Because trusted medical information is essential to the delivery of high-quality medical care, several trustworthy computing efforts deserve mention.

Founded in 1995, The Health On the Net Foundation is a not-for-profit organization with a mission of helping both laypersons and medical practitioners find useful, credible and reliable medical and health information online. As part of this mission, it has established ethical standards for Web site developers and has issued a code of conduct for medical and health Web sites.[14]

"…the physician faces the daunting challenge of not only searching for trustworthy and relevant medical information and applying that information appropriately for meaningful patient care."

The core of Trustworthy Computing has been the requirement that computing systems be inherently secure, available, and reliable. The National Academies' Committee on Information Systems publication, entitled Trust in Cyberspace,[15] defines a Trustworthy Computing system as one that "does what people expect it to do—and not something else—despite environmental disruption, human user and operator errors, and attacks by hostile parties."

Recently, the National Center for Complementary and Alternative Medicine published 10 criteria to help clinicians evaluate medical Web resources.[16] These criteria reflect a spectrum of issues, ranging from the incentives of the Web service provider for posting the information—in other words, who runs and pays for the site—through the value of the medical information (such as what is the basis of the information, where does it originate, how current is the information, and how is the information chosen) to peer activity, or how the site collects and displays user statistics.

Finally, because Web search engines play a major role in locating and accessing medical information, several online resources can help physicians evaluate Web-enabled information retrieval sources. BIOME is a collection of gateways that provide access to evaluated high-quality Internet resources in the health and life sciences.[17] The Special Advisory Group on Evaluation has published a differentiated six-step guide, "How to Evaluate Internet-based Information Source," that includes documentation regarding each step, such as how to analyze a URL, and how to efficiently answer a medical question of interest.[18]

In addition, the National Library of Medicine, via MedlinePlus, offers a 16-minute tutorial for evaluating health information on the Web that reflects a broader description of the criteria published by the NCCAM.[19] Finally, the Health Summit Working Group of the Health Information Technology Institute of Mitritek Systems published a policy paper that describes seven criteria—credibility, content, disclosure, links, design, interactivity, and caveats—for evaluating the quality of health information provided on the Internet.[20]

Although, these resources are helpful for screening Internet-acquired medical information, they do not provide the necessary scope to help physicians use Web-accessible information retrieval and extraction tools as well as associated clinical decision support applications for best-evidence clinical decision-making.

The Role of Relevance

In computer and library sciences, relevance has a specific meaning when applied to information retrieval. The term is comprised of two complementary metrics. One is recall, or the coverage of the system, such as the number of relevant publications a system returns from the number of relevant publications in the database. The other metric is the precision or the accuracy of the system, which measures the number of relevant publications returned, compared with all publications in the database.

These metrics describe the performance of a search engine in returning the information that users request. In addition, when searching for information, a relevance score is assigned to a search result that represents how accurately the retrieval matches the query. In many cases, that relevance score determines the order in which the information is presented to the user.

For physicians, the term relevance carries a broader meaning. Physicians are not just concerned with the inherent performance of an information search engine. In an epistemological sense, physicians are interested in relevance from a real-world, practical, patient-care perspective. To physicians, the material implications of medical information, as well as the contingency relationships that balance the truthfulness of the information against its contradictions, are most germane. In many ways, the physician's perception of relevance is more akin to that of an economist—that is, relevance as a calculation of the risk associated with use of the information for decision-making for medical purposes.[21]

To reconcile these themes, cognitive scientists have posited that relevance is central to both decision reasoning and communication.[22] By expressing relevance in propositional and goal-directed terms that are applicable to real-world problem solving,[23] physicians can define relevant information as information that supports best-evidence decision-making based on what will work as a matter of

empirical fact, rather than as a state of the clinician's personal beliefs or intuition.[24]

Features of Trusted, Relevant Resources

General themes tend to recur among Web-enabled medical information resources. These themes can be organized into a trust-relevance feature set for evaluating medical resource engines.

Domain adherence describes how well the resource provides requested information when compared with what it proposes to achieve. For example, if the information resource states it retrieves only publications on cardiology, does it fetch some publications on cardiology as well as some publications on endocrinology, while not retrieving additional publications on cardiology?

Domain coverage describes the breadth of the information resource. For example, if the information resource states it retrieves only publications on cardiology, does it fetch only general cardiology publications, or does it also get subspecialty cardiology publications as well, such as pediatric cardiology?

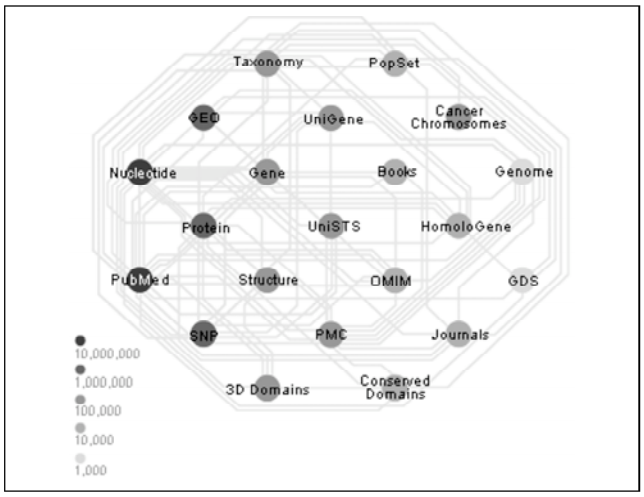

Comprehensibility reflects the physician's ability to understand the information that is retrieved for the purposes of clinical decision-making. Quite simply, how easily can a practicing physician relate the facts presented by the resource to patient care? Do "statistical jugglery" and "techno-jargon" obfuscate information and impede the articulation of a clear plan of care?

Authoritativeness reflects the influence of the medical information. Although an elusive concept, certain questions should be answered when assessing the potential influence of an information resource on medical care. What is the origin of the information? Is the information derived from a primary or secondary source? Does the information resource provide annotated references or citations that will enable physicians to validate the source and facts as presented?

Conciseness of results reflects efficiency by comparing information resources that retrieve an ordered or concise set of articles or summaries with those resources that "dump" a disorganized group of documents or hyperlinks into the physician's lap.

Linkages to related resources enable physicians to expand their medical information horizons. Specifically, ease of connectivity between resources provides complementary or enriching information that could benefit physician inquiry. For example, fueled by discoveries in basic clinical and biological research, the biomedical domain is witnessing an inordinate inflow of information that is constantly updating knowledge with new findings and therapies. Information resources that contend that they are complete should have links to other biomedical knowledge sources that publish new findings (see Figure 1). This connectivity enables physicians to uncover unrecognized previously existing relationships for personal growth and the advancement of medical care.

Cost encompasses direct cash outlays, such as fees, for services; the time-value of money, which assesses how quickly a physician is able to access information during busy office hours; and the opportunity cost of the information search, which measures whether another search method, such as consulting a peer down the hallway, provides the information in a more efficacious manner.

Curation reflects how well the information resource's database is constructed and maintained. Information resources need to perform curation, refinement, or rephrasing of their content to make documents and publications easily accessible to a wider interdisciplinary audience. It is important to know if standards are available that attest to the reliability and validity of the database entries, and if there is published documentation that the information resource is not only useful but has a track record of retrieving appropriate medical information. Finally, it is important to know whether the information resource uses manual or automated curation. With manual curation, knowledgeable experts generate the database content after sifting through existing publications and research. For example, an article in UpToDate is a compilation of reviews by medical domain experts. The automated curation approach employs computer algorithms to sift through millions of knowledge sources to generate database content automatically. Intuitively, the manual approach would tend to produce more trustworthy content. However, sophisticated computer algorithms sometimes produce better results on a larger scale. For example, an article found using Google Scholar[26] is based on the "goodness" of a document, according to the number of other "good" documents that link to it.[27] An ideal approach is to combine the efforts of expert curators and computerized methods, resulting in more reliable and comprehensive as well as speedier database generation. The Entrez/PubMed system applies just such a methodology.[28]

Usability reflects the ease of information retrieval of an

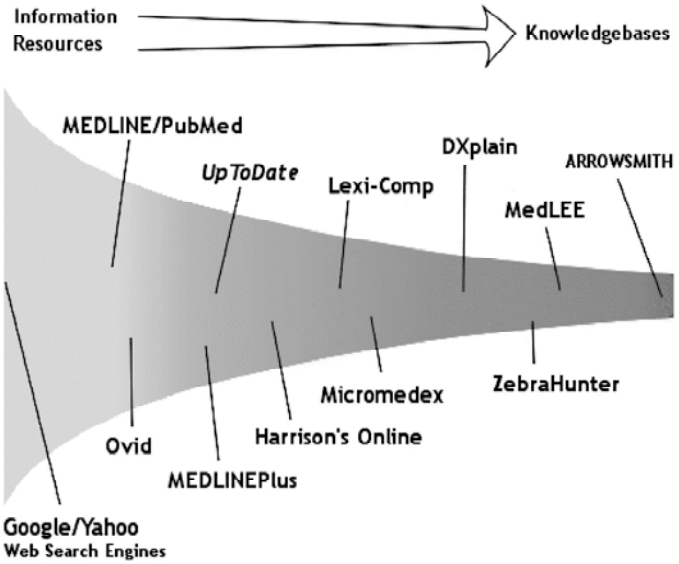

Figure 2. The Spectrum of Information Retrieval Resources in Biomedicine

information resource. Are instructions for use clear and associated with easily accessible help functions that expedite individual learning curves? Do advanced search techniques enable more efficient use of the system? Additionally, from a physician's perspective, an information resource must offer specific, biomedicine search features that are readily accessible. The sociotechnical factors that gauge the feature-accessibility balance are expert-novice configurability and visual noise. Expert-novice configurability implies that a physician should be able to open a Web browser and begin using the information resource without training. In this regard, individuals possess mental models for information searching where they expect to use an empty text box and a search button.29 Unfortunately, the lack of flexibility and the complex queries preferred by physicians require input wizards that disrupt workflow and take significant time before producing the final results. Visual noise represents the amount of irrelevant information presented on the screen along with the actual result. The 79

Web's mass media model enables the display of targeted advertising, pop-ups and malicious programs that degrade users' experiences. For physicians, being bombarded by irrelevant information is a distraction and a deterrent to using the resource.

The Spectrum of Information Retrieval

Based on the trust-relevance feature set, various medical information resources can be described through a spectrum of comparison based on the resource's structure and its information retrieval focus (see Figure 2). At one end of the spectrum, general information resources use text documents or publications to satisfy an information query—for example, if one queries Google or PubMed for information on sodium, documents containing the term Sodium or Na are retrieved. However, at the other end of the spectrum, resources use computer-interpretable knowledgebases. If someone queries for Sodium, the knowledge links in the resource can infer that sodium is related to hypo- or hypernatremia and return information that could be considered more intelligent.

Here are descriptions of and inherent trade-offs between various medical information resources and some examples of how the trust-relevance feature set can be useful for selecting a resource to assist with clinical decision-making.

Search engines (Google and Yahoo). Currently, the World Wide Web has more than 11.5 billion Web pages and is growing each year.[30] Based on the Pew Internet and American Life Project study,[31] 93 million Americans have used the Internet to search for health-related information. The most preferred access points to the Web are search engines[32] such as Google, Yahoo, and MSN, which index much of the content on the Web and provide ranked results for a keyword-based query.

For example, Google[33] is a hypertextual Web search engine designed to crawl the Web for information retrieval. The software employs a proprietary PageRank system[34] that determines an individual page or document's value or ranking,[35] according to the number of other "good" documents to which it is associated or linked. Similarly, most search engines measure authoritativeness by the total number of incoming hyperlinks from other pages, and relevance by the number of keyword matches for the given query.

Unfortunately, the ranking of retrieval results across the different search engines is extremely inconsistent. Therefore, a physician viewing a query result must take additional steps to both verify the relevance and decide whether or not to trust the information displayed on the Web page.

Reference Databases (MEDLINE, Entrez/PubMed, Ovid Medline and MedlinePlus). MEDLINE (Medical Literature Analysis and Retrieval System Online) is the National Library of Medicine's bibliographic database that contains approximately 13 million references to journal articles in the life sciences, with a concentration on biomedicine.[36] MEDLINE indexes journal articles and books using

the Medical Subject Headings of the Unified Medical Language System controlled terminology to create its hierarchy for information retrieval. MEDLINE may be searched via the Entrez Gateway using PubMed, one of a series of databases provided by the National Center for Biotechnology Information, a première Web site for biomedical and bioinformatics research.

Ovid Medline[37] is a proprietary information retrieval application that searches MEDLINE as well as other bibliographic databases (Cochrane EBM Reviews) to provide medical information. Finally, MedlinePlus[38] is a medical information resource that draws from the National Institutes of Health and functions as a medical encyclopedia, a medical dictionary, and a consumer-centric information resource for prescription and nonprescription drugs, consumer health news, and clinical trials information.

"In an epistemological sense, physicians are interested in relevance from a real-world, practical, patient-care perspective."

Reviews and Meta-information (UpToDate and Harrison's Online). The UpToDate[39] information retrieval system gives physicians current, topical medical reviews written by expert physicians. The comprehensive database is fully annotated and undergoes an extensive, peer review process of more than 330 journals per month to ensure that the information and recommendations are accurate and reliable. The system makes available current answers to clinical questions, identifies the clinical manifestations of a wide variety of diseases and disorders, and describes the current options for diagnosis, therapy, and management.

Harrison's Online, the electronic version of Harrison's Principles of Internal Medicine, serves as the flagship for McGraw-Hill's AccessMedicine educational resource for physicians and patients.[40] Via the AccessMedicine home page, users can search Harrison's Principles of Internal Medicine; review current and important updates in medicine through Harrison's Online Updates; scan drug reviews; and explore "Diagnosaurus,"[41] a differential diagnosis decision support tool based on patient symptoms, possible diseases, or organ system involvement.

Medication information (Lexi-Comp and Micromedex). The Lexi-Comp Knowledge Solution and Lexi-Comp Online[42] offer a real-time platform for integrating hospital information systems with specific formularies while enabling unlimited access to the entire site by physicians. Lexi-Comp Online focuses on medication management by providing the most current, clinically relevant content to help physicians make safe point-of-care patient therapeutic decisions. Lexi-Comp offers a comprehensive suite of medication management information resources via managed drug databases, special search functions for therapeutic changes and drug alerts, unique modules such as drug images, and specialized applications that combine the worldwide literature and scientific understanding of drug and natural product interactions to protect against adverse drug events.

Thomson Micromedex[43] provides a broad spectrum of medication management information including the DRUGDEX comprehensive drug reference, the Martindale drug information from Royal Pharmaceutical Society, summarized drug information, the Physician's Desk Reference Online, and the IV INDEX for intravenous medication compatibility.

Decision support (DXplain). Developed at the Laboratory of Computer Science at the Massachusetts General Hospital, the DXplain decision support system serves as both an electronic medical textbook and a medical reference system.[44] In its reference or case analysis mode, DXplain accepts a set of clinical findings (signs, symptoms, and laboratory data) to produce a ranked list of diagnoses that might explain or be associated with the patient's clinical findings. In addition, each finding is assigned an associated disease-independent term importance from 1 to 5 indicating how important it is in explaining the presence of the disease. DXplain provides justification for each differential diagnostic choice, suggests additional clinical information that would help clarify each diagnosis, and lists those clinical manifestations, if any, that would be unusual or atypical for each disease.

As a medical textbook, DXplain can describe more than 2,200 different diseases, emphasizing the etiology, signs and symptoms, pathology, and prognosis of each disease. In addition, DXplain provides a list of diseases that should be considered for more than 4,900 different clinical manifestations as well as many as 10 references for each disease, selected to emphasize clinical reviews when available.

Information extraction and knowledge discovery (MedLEE, ZebraHunter and Arrowsmith). Although the present status of information retrieval is helpful for physicians, it is worthwhile to take a glimpse at what the future holds for retrieving medical information for patient care. Developed at Columbia University, the Medical Language Extraction and Encoding System[45] (MedLEE) is a clinical natural language processing system that uses a combination of linguistic and heuristic domain knowledge to extract, structure, and encode clinical information in textual patient reports so data can be used by physicians for clinical decision-making. Structurally, MedLEE consists of a preprocessor that delineates the sentences of the report; a parser that utilizes the grammar and categories assigned to the sentence phrases to recognize well-formed syntactic and semantic patterns; a phrase regularizer that replaces the

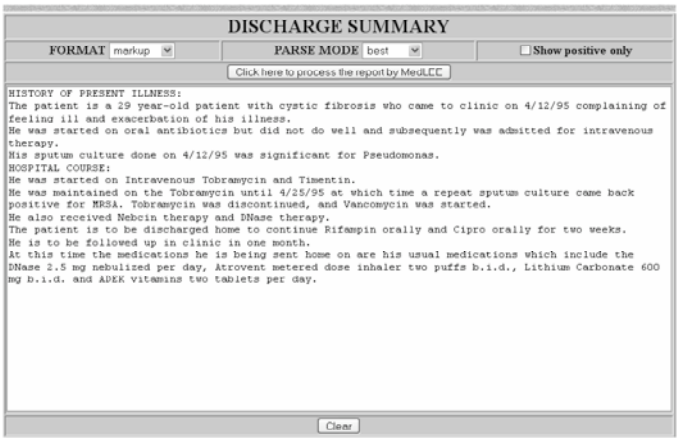
DISCHARGE SUMMARY

FORMAT markup PARSE MODE best Show positive only

Click here to process the report by MedLEE

HISTORY OF PRESENT ILLNESS:
The patient is a 29 year-old patient with cystic fibrosis who came to clinic on 4/12/95 complaining of feeling ill and exacerbation of his illness.
He was started on oral antibiotics but did not do well and subsequently was admitted for intravenous therapy.
His sputum culture done on 4/12/95 was significant for Pseudomonas.
HOSPITAL COURSE:
He was started on Intravenous Tobramycin and Timentin.
He was maintained on the Tobramycin until 4/25/95 at which time a repeat sputum culture came back positive for MRSA. Tobramycin was discontinued, and Vancomycin was started.
He also received Nebcin therapy and DNase therapy.
The patient is to be discharged home to continue Rifampin orally and Cipro orally for two weeks.
He is to be followed up in clinic in one month.
At this time the medications he is being sent home on are his usual medications which include the DNase 2.5 mg nebulized per day, Atrovent metered dose inhaler two puffs b.i.d., Lithium Carbonate 600 mg b.i.d. and ADEK vitamins two tablets per day.

Clear

Figure 3. Dictated Textual Discharge Summary as Prepared by a Typist for Entry into the Medical Records System

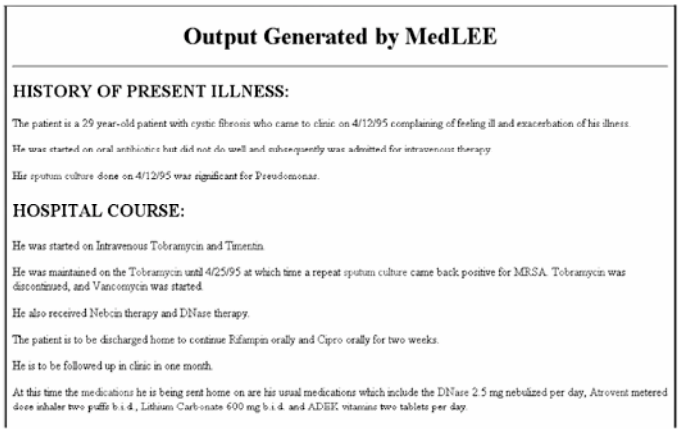
Output Generated by MedLEE

HISTORY OF PRESENT ILLNESS:

The patient is a 29 year-old patient with cystic fibrosis who came to clinic on 4/12/95 complaining of feeling ill and exacerbation of his illness.

He was started on oral antibiotics but did not do well and subsequently was admitted for intravenous therapy.

His sputum culture done on 4/12/95 was significant for Pseudomonas.

HOSPITAL COURSE:

He was started on Intravenous Tobramycin and Timentin.

He was maintained on the Tobramycin until 4/25/95 at which time a repeat sputum culture came back positive for MRSA. Tobramycin was discontinued, and Vancomycin was started.

He also received Nebcin therapy and DNase therapy.

The patient is to be discharged home to continue Rifampin orally and Cipro orally for two weeks.

He is to be followed up in clinic in one month.

At this time the medications he is being sent home on are his usual medications which include the DNase 2.5 mg nebulized per day, Atrovent metered dose inhaler two puffs b.i.d., Lithium Carbonate 600 mg b.i.d. and ADEK vitamins two tablets per day.

Figure 4. A MedLEE Processed Discharge Summary that Marks-up the Diagnoses (red), the Diagnostic Procedures (green), and the Medications (brown), so Physician's Can Expeditiously Evaluate Different Aspects of the Patient's Care

canonic output specified in the lexical definition of the phrase with its associated position in the report; and an encoder that maps the canonic forms into controlled vocabulary terms. The system generates structured output that consists of tagged text in either indented MedLEE, XML, HL7, or clinical markup format that can be used by physicians for categorizing various aspects of patient care (see Figures 3 and 4).[46,47]

Also developed at Columbia University, ZebraHunter[48] is a Web-enabled medical literature information retrieval and extraction application that employs a complex series of wrappers, algorithms, and natural language processing tools linked to a curated database of case reports, clinicopathological conferences, and clinical images. Based on the physician's query, the application retrieves a set of citations and references that are associated with a differential of complex and rare clinical diagnoses based on the patient's particular set of clinical signs, symptoms, and findings.[49]

Finally, supported by the National Library of Medicine and the National Institute of Mental Health's Human Brain-Neuroinformatics Research Project through the University of Illinois at Chicago, Arrowsmith[50] is a medical information retrieval system that queries MEDLINE through the Entrez/PubMed Gateway to identify common items or concepts that are present in two distinct sets of documents or publications for the purpose of knowledge discovery. In addition, the system can identify information that is present in one medical domain that may be relevant to another domain of inquiry. By using the Author-ity tool, the system accepts as input an author's name associated with a specific article in MEDLINE and returns as output a list of all articles with that name ranked by decreasing probability that they are authored by the same individual.[51]

Four Clinical Scenarios

The following scenarios depict four common situations for which information retrieval is essential for patient care. Each section describes the applicable trust and relevance features, the information resources that would satisfy those conditions, and how the resource would be helpful to the physicians under those circumstances.

For biosurveillance, physicians need clinical reports and notifications about new rapidly evolving infectious illnesses. The Center for Disease Control's Morbidity and Mortality Weekly Report[52] used in conjunction with UpToDate are an effective combination of resources that rely on strong domain adherence and coverage as well as authoritativeness for retrieving the most useful information.

In a busy private practice, a physician who is caring for a patient with a rare genetic disorder might wish to look up associated clinical findings. Because work process compatibility, conciseness of results, and usability are important, The Office of Rare Diseases' Term Definitions[53] or Harrison's Online are resources that can quickly retrieve the necessary facts with minimal disruption to the physician's work pattern.

In a high-volume emergency department, when patients come in with various signs and symptoms, the physician is concerned with rapidly retrieving a differential diagnosis to use as a springboard for further evaluation. Under these circumstances, the domain coverage, comprehensibility, and usability of DXplain would be most useful.

When requesting a consultation, a physician might feel fairly certain that a patient is suffering from a rare condition. Because citing corroborating literature is important for supporting this impression, the ZebraHunter bibliographic retrieval is an authoritative and well-curated resource that can increase the physician's diagnostic level of confidence when conferring with a colleague.

Conclusion and Future Directions

In the current medical practice environment, the physician faces the daunting challenge of not only searching for trustworthy and relevant medical information and applying that information appropriately for meaningful patient care. However, in the face of overwhelming odds, current and developing innovations in medical information retrieval will help them.

This paper gives physicians reasonable discussion points regarding these concerns, a usable framework for evaluating information retrieval innovations and explanations of the different capabilities of representative information retrieval tools and applications. By demystifying the concepts associated with information resources, search engines, and retrieval tools, and giving physicians a reasonable view of current opportunities as well as future possibilities, physicians can more rapidly adopt innovative computer-assisted search tools for acquiring information that facilitate patient care decision-making.[54]

To harken back to the start of this article and the use of Google, although serendipity is defined by a combination of sagacity and good fortune,[55] an even mixture of the former and the latter is not tolerable when it comes to the practice of medicine. The Fellow's serendipitous googling of a set of findings to retrieve a correct diagnosis might be very much consistent with Google's exhortation of "I'm Feeling Lucky."[56] However, in medical practice, physicians need the Web to present reliable, consistent, and reproducible "best expert" evidence when they are providing day-to-day care for patients, and Google alone is just not good enough.

Acknowledgment

The authors would like to thank Randolph C. Barrows, Jr., MD; James J. Cimino, MD; Stephen B. Johnson, PhD; John L. Zimmerman, DDS; and Eneida A. Mendonça, MD, PhD, faculty in the Department of Biomedical Informatics at Columbia University, for their critical reading and commentary on this manuscript. This work was supported in part by the National Library of Medicine, Training Grant Number: T-15-NLM007079.

About the Authors

Eric Z. Silfen, MD, MSHA, is a post-doctoral fellow in the Department of Biomedical Informatics and director of the executive education programs.

Chintan O. Patel, MS, is a PhD candidate in the department of biomedical informatics.

References

1. Correspondence. "…And a Diagnostic Test Was Performed". *NEJM*. 2005;353(19):2089-2090.
2. Johnson CM, Johnson TR, Zhang JJ. A User-centered Framework for Redesigning Health Care Interfaces. *Journal of Biomedical Informatics*. 2005;38(1):75-87.
3. Swanson DR. *Medical Literature as a Potential Source of New Knowledge*. Bull Med Libr Assoc. 1990;78(1):29-37.
4. Rogers EM. *Difusion of Innovation*, 5th ed. NewYork:FreePress;2003.
5. Cook A, Packer C, Stevens A. Diffusion of Thrombolysis for Acute Myocardial Infarction from 1981 to 2000 in England: Trend Analysis and Comparison with Need. *International Journal of Technology Assessment in Health Care*. 2004;20(4):531-536.
6. Cook A, Packer C, Stevens A. Influences upon the Diffusion of Thrombolysis for Acute Myocardial Infarction in England: Case Study. *International Journal of Technology Assessment in Health Care*. 2004;20(4):537-544.
7. Morrissey J. Harmonic Divergence. Cedars-Sinai Joins Others in Holding off on CPOE. *Modern Healthcare*. 2004;34(8):16.
8. Kaushal R, Kaveh S, Bates DW. Effects of Computerized Physician Order Entry and Clinical Decision Support Systems on Medication Safety: A Systematic Review. *Arch Intern Med*. 2003;163:1409-1416.
9. Koppel R, et al. Role of Computerized Physician Order Entry Systems in Facilitating Medication Errors. *JAMA*. 2005;293(10):1197-1203.
10. Han YY, et al. Unexpected Increased Mortality after Implementation of a Commercially Sold Computerized Physician Order Entry System. *Pediatrics*. 2005;116(6):1506-12.
11. Audet AM, et al. Information technologies: When Will They Make it into Physicians' Black Bags? *Medscape General Medicine*. 2004;6(4):2.
12. Google: "I'm Feeling Lucky", http://www.google.com/ig?hl=en, accessed on 9/6/2006.
13. Merriam-Webster Online Dictionary, http://www.m-w.com/dictionary/trust, accessed on 9/6/2006.
14. Health on the Net Foundation, http://www.hon.ch/HONcode/Conduct.html, accessed on 9/6/2006.
15. Trust in Cyberspace, http://books.nap.edu/books/0309065585/html/index.html, accessed on 9/6/2006.
16. 10 Things to Know About Evaluating Medical Resources on the Web, http://nccam.nih.gov/health/webresources/, accessed 8/21/2006.
17. BIOME, http://biome.ac.uk/about/, accessed on 8/21/2006.

18. How to Evaluate an Internet-based Information Resource, http://biome.ac.uk/guidelines/eval/howto.html, accessed on 8/21/2006.
19. Evaluating Internet Health Information, http://nlm.nih.gov/medlineplus/webeval/webeval.html, accessed on 8/21/2006.
20. Criteria for Evaluating the Quality of Health Information on the Internet—Policy Paper, http://hitiweb.mitretek.org/docs/policy.htnl#top, accessed on 8/21/2006.
21. Keynes JM. *Treatise on Probability*. London: MacMillan;1921.
22. Sperber D, Wilson D. Précis of Relevance: Communication and Cognition. *Behavioral and Brain Science*. 1987;10:697-754.
23. Gorayska B, Lindsay RO. The Roots of Relevance. *Journal of Pragmatics*. 1993;19:301-323.
24. Lindsay R, Gorayska B. Relevance, Goals and Cognitive Technology. *International Journal of Cognitive Technology*. 2002;1(2):187-232
25. National Center for Biotechnology Information, http://www.ncbi.nih.gov/Database/datamodel/index.html, accessed on 9/6/2006.
26. Google, http://scholar.google.com/, accessed on 9/6/2006.
27. Ridings C, Shishigin M. Page Rank Uncovered. 2002. www.voelspriet2.nl/PageRank.pdf, accessed on 9/6/2006.
28. Aronson AR, et al. The NLM Indexing Initiative. Proc AMIA Ann Symp. 2000:17-21.
29. Calongne, CM. Designing for Web Site Usability. Proceedings of the seventh Annual Consortium for Computing in Small Colleges. 2001:39-45. http://portal.acm.org/citation.cfm, accessed on 9/6/2006.
30. Gulli A, Signorini A. The Indexable Web is More Than 11.5 Billion Pages. In Special interest Tracks and Posters of the 14th International Conference on World Wide Web. ACM Press, New York, NY. 2005:902-903.
31. Pew Internet and American Life Project, http://207.21.232.103/PPF/r/95/report_display.asp, accessed on 9/6/2006.
32. Search Engine Watch, http://searchenginewatch.com/reports/article.php/2156451, accessed on 9/6/2006.
33. Google Technology, http://www.google.com/technology/, accessed on 9/6/2006.
34. Page L. The PageRank Citation: Bringing Order to the Web. January 29, 1998. http://dbpubs.stanford.edu:8090/pub/showDoc.Fulltext?lang=en&doc=1999-66&format=pdf&compression=&name=1999-66.pdf, accessed on 9/6/2006.
35. Henzinger M and Lawrence S. Extracting Knowledge from the World Wide Web. PNAS. 2004;101(Suppl.1):5186-5191 www.pnas.org/cgi/doi/10.1073/pnas.0307528100, accessed on 9/6/2006.
36. MEDLINE, http://medline.cos.com/, accessed on 9/6/2006.
37. Ovid, http://gateway.ut.ovid.com/, accessed on 9/6/2006.
38. MedlinePlus, http://medlineplus.gov/, accessed on 9/6/2006.
39. UpToDate, http://www.uptodate.com/physicians/index.asp, accessed on 9/6/2006.
40. Harrison's Online, http://www.accessmedicine.com/resourceToc.aspx?resourceID=4, accessed on 9/6/2006.
41. Diagnosaurus, http://www.accessmedicine.com/diag.aspx, accessed on 9/6/2006.
42. Lexi-Comp, http://www.lexi.com/web/index.jsp, accessed on 9/6/2006.
43. Micromedex, http://www.micromedex.com/index.html, accessed on 9/6/2006.
44. DXplain, http://www.lcs.mgh.harvard.edu/, accessed on 9/6/2006.
45. MedLEE, http://lucid.cpmc.columbia.edu/medlee/, accessed on 9/6/2006.
46. Friedman C, et al. Automated Encoding of Clinical Documents Based on Natural Language Processing. *J Am Med Inform Assoc*. 2004;11:392-402
47. Friedman C, et al. Representing Information in Patient Reports Using Natural Language Processing and the Extensible Markup Language. *J Am Med Inform Assoc*. 1999;6:76-87
48. ZebraHunter, www.zebrahunter.org, accessed on 9/6/2006.
49. Silfen E, Patel CO, Mendonça E, Friedman, C. ZEBRAHUNTER: Mining Case Reports in MEDLINE to Find Rare Medical Diagnoses and the Relevant Citations. Department of Biomedical Informatics, Columbia University, submitted for publication.
50. Smalheiser NR, Torvik VI, Weeber M, Swanson DR. The Arrowsmith Project: New Tools to assist Biomedical Discovery and Collaboration, presented at the Beckman Institute, University of Illinois at Urbana-Champaign, October 1, 2002. http://arrowsmith.psych.uic.edu/arrowsmith_uic/index.html, accessed on 9/6/2006.
51. Smalheiser, NR. The Arrowsmith Project: 2005 Status Report. Invited address at the 8th International Conference on Discovery Science and the 16th International Conference on Algorithmic Learning Theory, October 8-11, 2005, Singapore. To appear in Lecture Notes in Artificial Intelligence (A. Hoffman, H. Motoda, and T. Scheffer, Editors) 3735:26-43. Lecture Notes in Computer Science Series. Copyright Springer-Verlag, accessed on 9/6/2006.
52. Morbidity and Mortality Weekly Report, http://www.cdc.gov/mmwr/weekcvol.html, accessed on 9/6/2006.
53. Office of Rare Diseases, http://rarediseases.info.nih.gov/, accessed on 9/6/2006.
54. Chen H, Fuller SS, Friedman C, Hersh W. Knowledge Management, Data Mining and Text Mining in Medical Informatics. In: Chen H, Fuller SS, Friedman C, Hersh W, eds. Medical Informatics - Knowledge Management and Data Mining in Biomedicine. 2005
55. Coined by Horace Walpole (1754) after The Three Princes of Serendib (i.e. Ceylon), a Persian fairy tale in which the princes always make discoveries of things they were not in quest of either by accident or sagacity.
56. Google: "I'm Feeling Lucky", http://www.google.com/, accessed on 9/6/2006.