



Obtaining useful information from expert based sources

David C Slawson, Allen F Shaughnessy

Summary

Clinicians rely heavily on expert based systems—consultation with colleagues, journal reviews and textbooks, and continuing education activities—to obtain new information. The usefulness of sources such as these depends on the relevance and validity of the information and the work it takes to obtain it. Useful information can be distinguished from the useless by asking three questions: Does the information focus on an outcome that my patients care about? Is the issue common to my practice, and is the intervention feasible? If the information is true, will it require me to change my practice? If the answer to all three questions is yes, then the information is a common POEM (patient oriented evidence that matters), capable of improving the lives of your patients and must be evaluated for validity. Conclusions based on results of well designed clinical trials are more likely to be valid than those drawn from observations based on experience in clinical practice. Both members of the team, clinicians and experts, must take responsibility for their respective roles.

Introduction

When educators and academics look for solutions to the problem of managing information in clinical practice they tend to focus on computer resources—hardware and software—for solutions. Clinicians, on the other hand, usually turn to colleagues or other specialists—“wetware”—to find answers. Information from these expert sources can be obtained by direct consultation, via review articles and textbooks, or during continuing education activities.

Expert sources of information are valuable because they are quick, cheap (usually), and easy to use. An expert also provides guidance, support, affirmation, and other psychological benefits that computerised sources cannot provide.¹ While much has been written about the critical assessment of journal articles and reviews, little information is available to guide clinicians in evaluating “expert” sources. We will explore the usefulness of these sources and present a rationale for when to use them and how to evaluate the information.

Determining the usefulness of medical information

When we read journals or textbooks, attend continuing education conferences, or consult with colleagues to

obtain information about treating our patients, our goal is to find the most useful information in the shortest time. The best information has three attributes: it must be relevant to everyday practice, it must be correct, and it should require little work to obtain. These three factors are related as follows:

$$\text{Usefulness of medical information} = \frac{\text{Relevance} \times \text{Validity}}{\text{Work}}$$

Relevance

Relevance is based on the type of information being presented and the frequency of the problem in your practice. The most relevant information will tell you how to help your patients live functional satisfying lives free from pain and symptoms. We call this type of information patient oriented evidence, It is based, as much as possible, on the results of “outcomes based” research rather than “authority based” or “experiential” impressions. It does not stop with surrogate markers² or intermediate level data—what we call disease oriented evidence—but instead evaluates outcomes of importance to people.³

For example, an expert recommending the need for screening for prostate cancer with the prostate specific antigen assay may discuss the accuracy of the test in identifying men with prostate cancer and the survival rates for different stages of prostate cancer (disease oriented evidence). This information does not tell you what you and your patients really want to know: whether they will be better off (live a longer, healthier, happier life) as a result of early identification of the cancer (patient oriented evidence). Only a randomised trial evaluating the overall effect of early detection on mortality and morbidity from prostate cancer will provide this information.

What doctors are truly seeking is “patient oriented evidence that matters” (POEM). This information matters because, if it is valid, it should change what they do in practice. Once a POEM has been identified the frequency of contact with the problem in clinical practice must be considered. When doctors evaluate information sources POEMs are top priority if they provide information that doctors can use to better manage patients with illnesses frequently seen in their clinical practice. These common POEMs will have the greatest impact on patients and therefore have the greatest relevance.

Department of Family Medicine, University of Virginia, Charlottesville, VA 22908, USA
David C Slawson, *associate professor*

Harrisburg Family Practice Residency Program, Harrisburg, PA
Allen F Shaughnessy, *director of research*

Correspondence to: Dr Slawson (dslawson@virginia.edu).

BMJ 1997;314:947-9



Validity

The second factor is the likelihood of the information being true. Conclusions based on results of well designed clinical trials are more likely to be valid than those drawn from observations in clinical practice. But it is not enough to accept evidence at face value simply because it has been published in a well known journal or comes recommended from a specialist.⁴ For example, early research into preventing osteoporosis showed that sodium fluoride increased bone density, and, based on expert recommendation, many patients were subsequently treated.⁵ Further research, however, found that fracture rates were actually increased with the use of fluoride.⁶

Work

This includes factors such as how long it takes to obtain the information, how much it costs, and the amount of mental energy required to track down the answer. Working too hard to establish the validity or relevance of information will lower its usefulness. Too often, however, information sources that require little work also have low validity or relevance and should be used cautiously. Expert based sources are appealing because the work needed to access information is low. The true potential for usefulness of these sources therefore depends on their validity and relevance.

The “lure” of the expert

Perhaps the most important benefit of expert based information is that experience can be used to interpret and apply evidence to the care of difficult patients. Experts can be relied on to perform procedures or to help with diagnosing atypical disease. For example, a paediatric infectious disease specialist may be the best person to confirm the diagnosis of varicella encephalitis in a sick child.

Doctors must be cautious, however, when adopting an expert's anecdotal based treatment advice. This same infectious disease specialist may be too biased by 20 years of experience to endorse the practice of not admitting to the hospital all febrile children under the age of two months. There also is a tendency on the part of the clinicians to develop clinical “rules” out of patient specific recommendations made by consultants. A typical mental shortcut is to think, “The last time I asked, the cardiologist suggested amlodipine. Therefore, I should always use amlodipine.” It may not

have been the expert's intention to provide a wide ranging rule in response to your specific question.

There are several other reasons why expert based information may be inaccurate. Firstly, expert advice may not be based on the most current research. A study of reviews and textbook articles written by experts found that in some cases up to 13 years elapsed between the time that convincing evidence was available to support interventions for managing myocardial infarction and when these interventions were recommended by most experts.⁷ As well as the lag time associated with publication, this delay stems from the proclivity of experts, especially researchers, to favour their beliefs or their prior experiences over evidence derived from outcomes based research, a tendency termed “reverse gullibility.”⁸ For example, at least four randomised clinical trials have shown that patching eyes after a corneal abrasion not only does no good but actually worsens pain and delays healing.⁹⁻¹² In a recent series of letters in a major journal expert ophthalmologists stated that, despite this research, clinicians at their institution had always patched eyes and would continue to do so regardless of information obtained from patient oriented clinical trials.^{13 14}

A second problem with expert information is that there is a tendency for authors of review articles to start by writing their conclusions and then finding the supportive evidence. The potential for unrecognised bias in these articles is high, since only references are used that support these predetermined conclusions. Furthermore, cited references often do not actually support the conclusions. For example, the latest edition of a widely used reference of antimicrobial treatment states that co-trimoxazole should not be used in children with otitis media who do not respond to amoxycillin.¹⁵ However, the reference given to support this recommendation actually states the exact opposite.¹⁶ This flaw is not uncommon: one study of review articles found that 24% of the referenced articles were not correctly summarised.¹⁷ Assertions or information in reviews or texts that contradict what you think is true should be checked for validity. A little raised number at the end of a statement is not an icon of inerrancy.

A third issue is that expert based knowledge is often developed through experience with a highly selected patient population, and this knowledge may not be applicable to the general population.¹⁸ To an endocrinologist who sees a unique population, patients with weight gain, headaches, and fatigue have a hormone secreting tumour until proved otherwise. Similar patients seen by a primary care doctor would probably be evaluated for depression.

Finally, you should not assume that an expert is skilled at evaluating medical research. Techniques for critical appraisal are not well known or easily developed, and experts may not be any better than you are at determining the validity of research findings.

Improving the usefulness of expert based information

Whenever review articles, textbooks, continuing education presentations, or consultants are used to transfer information, both the expert and the clinician are involved in the process. The appropriate use of expert

Benefits and drawbacks of expert based information**Benefits**

- Wisdom gained through experience
- Clinical “pearls” that cannot be derived through the scientific method
- Able to fill in the gaps in current outcomes based knowledge with evidence derived from clinical experience or from their extensive knowledge of the physiological basis of disease

Detriments

- The information may be out of date
- Reverse gullibility, in which clinical experience or disease oriented evidence is favoured over patient oriented evidence
- Knowledge may be based on a highly selected population, and the information may not be applicable to your patients

based sources requires the effort of both groups to assume responsibility for their respective roles.

Clinicians using review articles to keep up to date do not have to read the entire article. It is acceptable to read the article headlines, abstract, or conclusion first. If a statement is not firmly rooted in patient oriented evidence it can be disregarded. If a POEM is found (information with patient oriented outcomes with the potential to change clinical practice) the clinician must then determine its validity. Other sources can be used to confirm validity before any changes in practice are made. Original research and review articles can be evaluated with pre-existing criteria set forth by the Evidence Based Medicine Working Group.¹⁹ A useful source of continually updated reviews based on evidence is the Cochrane Collaboration.²⁰

Clinicians attending continuing education presentations or receiving advice from colleagues should stay alert to glean new information, supported by patient oriented evidence, that would require them to change their practice. Speakers and consultants should clearly identify whether their recommendations are based on patient oriented research outcomes or on other evidence. They should provide some evidence that they have assessed the validity of the information. If not, the listener must take an assertive approach by asking, "What is the evidence to support that recommendation?"

Whether experts are writing reviews or textbook chapters, giving continuing education presentations, or providing consultations, their role remains the same. They should always emphasise POEMs first when forming a recommendation, even when this information conflicts with disease oriented evidence or their own clinical experience. A substantial problem, however, is that POEMs are often unavailable to guide patient care decisions. As a result, experts are often unable to support a particular recommendation with POEMs but have to base their opinion on a knowledge of pathophysiology, studies of surrogate outcomes, experience, or "gut feeling." Thus, when POEMs are not available, experts should use the best available evidence and clearly state that these sources are the basis for their recommendations. By doing so, they can help clinicians to remain more open to changing their practice when patient oriented evidence becomes available, avoiding reverse gullibility.

This mesh of patient oriented evidence and expert opinion provides the most value to the audience and lowers the work required to obtain useful information. Too often lectures are like professional basketball games, in which only the last two minutes are important.

Conclusions

As consumers and providers of new information on the same team, clinicians and experts evaluating information from any source must always keep three questions in mind:

- Will this information have a direct bearing on the health of my patients?
- Is the issue at hand common to my practice?
- If valid, will this information require me to change my current practice?

A negative answer to any of these questions allows you to deflect the information. When the answer to all

Using expert based information

- Determine the level of evidence supporting recommendations—Is the recommendation based on patient oriented or disease oriented evidence?
- Do not make your own clinical rules out of patient specific recommendations made by consultants
- Remember the lag time between fact finding and publication
- Realise that authors of reviews and textbooks and speakers at educational conferences often begin with their predetermined conclusions and then find evidence to support only this point of view
- Wisdom gained through experience helps clinicians diagnose disease and perform procedures. Experience is not adequate to remain proficient in treatments: proficiency also requires a knowledge of the medical literature and the ability to think critically with an open mind

three of these questions is yes, you have found a common POEM capable of improving the lives of your patients. If the information is valid it will require you to change your practice. If you cannot determine validity from the source, you must look elsewhere. If it is not a valid POEM, it is just not necessarily so.

Experts are a valuable source of information for clinicians. Their usefulness, however, depends on the relevance and validity of the information they present. Close attention to these two issues by both experts and clinicians can lead to better information management, and with it, better patient care.

Funding: None.

Conflict of interest: None.

- 1 Smith R. What clinical information do doctors need? *BMJ* 1996; 313:1062-8.
- 2 Fleming TR, DeMets DL. Surrogate end points in clinical trials: are we being misled? *Ann Intern Med* 1996;125:605-13.
- 3 Slawson DC, Shaughnessy AF, Bennett JH. Becoming a medical information master: feeling good about not knowing everything. *J Fam Pract* 1994;38:505-13.
- 4 Shaughnessy AF, Slawson DC, Bennett JH. Becoming an information master: a guidebook to the medical information jungle. *J Fam Pract* 1994;39:489-99.
- 5 Grodberg MG. *Transcripts of testimony before Advisory Committee Meeting, Food and Drug Administration, Rockville, MD, August 18, 1978*. Rockville, MD: Food and Drug Administration, 1978.
- 6 Riggs BL, Hodgson SF, O'Fallon WM, Chao EYS, Wahner HW, Muhs JM, et al. Effect of fluoride treatment on the fracture rate in postmenopausal women with osteoporosis. *N Engl J Med* 1990;322:802-9.
- 7 Antman EM, Lau J, Kupelnick B, Mosteller F, Chalmers TC. A comparison of results of meta-analyses of randomised control trials and recommendations of clinical experts. Treatments for myocardial infarction. *JAMA* 1992;268:240-8.
- 8 Riffenburgh RH. Reverse gullibility and scientific evidence. *Arch Otolaryngol Head Neck Surg* 1996;122:600-1.
- 9 Hulbert MFG. Efficacy of eye pad in corneal healing after corneal foreign body removal. *Lancet* 1991;337:643.
- 10 Kirkpatrick JNP, Hoh HB, Cook SD. No eye pad for corneal abrasion. *Eye* 1993;7:468-71.
- 11 Donnenfeld ED, Selkin BA, Perry HD, Moadel K, Selkin GT, Cohen AJ, et al. Controlled evaluation of a bandage contact lens and a topical non-steroidal anti-inflammatory drug in treating traumatic corneal abrasions. *Ophthalmology* 1995;102:976-84.
- 12 Kaiser PK for the Corneal Abrasion Patching Study Group. A comparison of pressure patching versus no patching for corneal abrasions due to trauma or foreign body removal. *Ophthalmology* 1995;102:1936-42.
- 13 Harris LL, Jampel HD. Treatment of corneal abrasions. *JAMA* 1996;275:837.
- 14 Slawson DC, Shaughnessy AF. Treatment of corneal abrasions. *JAMA* 1996;275:837.
- 15 Sanford JP, Gilbert DN, Sande MA. *The Sanford guide to antimicrobial therapy*. 1996. Dallas: Antimicrobial Therapy, 1996:6.
- 16 Berman S. Otitis media in children. *N Engl J Med* 1995;332:1560-5.
- 17 Neihouse PF, Priske SC. Quotation accuracy in review articles. *DICP Ann Pharmacother* 1989;23:594-6.
- 18 Dunn EV, Norton PG, Stewart M, Tudiver F, Bass MJ, eds. *Disseminating research/changing practice*. Thousand Oaks, CA: Sage Publications, 1994:156-8.
- 19 Oxman AD, Sackett DL, Guyatt GH. Users' guides to the medical literature: I. How to get started. *JAMA* 1993;270:2093-5.
- 20 Bero L, Rennie D. The Cochrane collaboration. Preparing, maintaining, and disseminating systematic reviews of the effects of health care. *JAMA* 1995;274:1935-8.

(Accepted 5 March 1997)

Information retrieval for patient care

Martin Gardner

Information
Retrieval Research
Group, Department
of Computing
Science, University
of Glasgow,
Glasgow G12 8QQ
Martin Gardner,
clinical research fellow
(martin@dcs.gla.ac.uk,
http://www.dcs.
gla.ac.uk/~martin/)

BMJ 1997;314:950-3

Summary

Doctors need clinical information during most consultations with patients, and much of this need could be satisfied by material from online sources. Advances in data communication technologies mean that multimedia information can be transported rapidly to various clinical care locations. However, selecting the few items of information likely to be useful in a particular clinical situation from the mass of information available is a major problem. Current information retrieval systems are designed primarily for use in research rather than clinical care. The design, implementation, and critical evaluation of new information retrieval systems for clinical care should be guided by knowledgeable clinical users.

Introduction

During a consultation with a patient, a doctor must consider two types of information—the patient's medical record and medical knowledge relevant to the present problem. Nowadays, a doctor who chose, as routine practice, to rely on his or her memory of a patient's medical record rather than actually examining the record as part of the consultation might be considered eccentric, complacent, and possibly negligent. There is a huge and rapidly expanding volume of information in journals, textbooks, and other data sources constituting the body of knowledge on which modern medical practice is or should be founded; yet how often is any of this information examined during clinical consultations?

Increasing numbers of doctors recognise the need for such clinical information, and in the near future computer devices and communications networks will be capable of supplying information to the point of care (not only to consulting rooms but to hospital bedsides, operating theatres, patients' homes, road accidents, etc). Unfortunately, a major barrier to progress is that there is as yet no satisfactory solution to the problem of finding those few items of information most likely to be useful in any given situation among the mass of data available.

This general problem of information retrieval is not new or unique to medicine. Over the past 40 years, research into information retrieval has become a large and active discipline with applications in subjects such as law, finance, defence, publishing, research, and entertainment. A consequence of the recent explosive growth of the world wide web is that millions of people now use information retrieval systems on a daily basis.

There are three reasons why doctors should be aware of the principles of information retrieval technology. Firstly, one criterion by which you may judge the maturity of a technology is whether good performance requires technical, in addition to procedural, knowledge. By this criterion, word processing technology, for example, is mature since you can become a highly proficient user without any technical

knowledge about how commands are translated into results. Unfortunately, information retrieval technology is not mature in this respect, and doctors who have some understanding of how information retrieval systems work will get better results than those who do not. Secondly, as a setting for information retrieval, the point of care is very different from other contexts in which information retrieval systems are used. Technology for supplying clinical information to the point of care cannot become mature without the insight, guidance, and commitment of knowledgeable users. Thirdly, information retrieval systems are likely to become important items in healthcare budgets. Purchasing decisions should be influenced by informed clinicians at a local level.

Information needs at the point of care

Many doctors now recognise the need for reference information at the point of care. Indeed, the inaugural article of this section of the *BMJ* addressed just this topic.¹ Richard Smith presented evidence that information needs arise in nearly every consultation between a doctor and a patient, that many of these needs could be satisfied by material in reference sources, and that improved outcomes might accrue. There are problems of both memory registration and memory recall. Thus, no doctor can have read all the information relevant to any particular clinical decision, nor can a doctor expect to have impeccable recall of that sample which he or she has read. In future, healthcare purchasers may expect doctors to justify individual clinical decisions with explicit reference to evidence. More importantly, timely provision of information to the point of care could promote patients' ability both to participate in clinical decision making and to accept responsibility for the outcome.

The information needs of clinicians at the point of care are very different from those of academic researchers in a library or laboratory. Clinicians require access to a much wider range of material (not only journal articles but also passages from textbooks, drug monographs, protocols for patient care, medico-legal information, reference images (of dermatopathology perhaps), etc). They practise in a wide variety of environments (patients' homes and workplaces, wards, clinics, treatment rooms, etc), where standard desktop computers may not be available but information is still required. While researchers require the maximum number of information items with relevance to the topic but do not need rapid browsing interfaces, clinicians require a small representative sample of information items useful for decision making presented in a rapidly browsable manner. The problem of unperceived information needs is much greater for clinicians.



Supplying information to the point of care

The supply of information to the point of care relies on four technologies; information sources in digital form, data communication networks, computer devices at the point of care, and information retrieval systems.

There is already a large volume of medically related information available in digital format. This includes abstracts of journal articles (and the full contents of some), full contents of textbooks, clinical trial repositories, care protocols, critical incident reports, libraries of medical images, medical audio libraries (such as characteristic heart sound recordings of hundreds of cardiac disorders), and video clips of medical procedures (such as endoscopy and fiberoptic intubation). In the near future there will be an explosive increase in the volume of information available in digital format.

Data communication technology is advancing rapidly. Some hospitals already have ATM (asynchronous transfer mode) networks capable of transporting a full size, high resolution chest radiograph in less than a second. A number of general practice surgeries have local area network or ISDN (integrated service digital network) connections. It is now also possible to transmit data at acceptable rates using mobile phone links, or infrared waves over short distances.

Currently available battery powered laptop computers, which are small enough to carry in a briefcase to a patient's home, now have storage devices that can hold the entire contents of many thousands of journal articles or several large textbooks. You can buy palmtop computers which can store several megabits of information, display text and pictures, recognise handwriting, be activated by speech, and be connected to the Internet with a mobile phone.

Advances in these three technologies highlight the need for progress with the fourth—information retrieval.

What is information retrieval?

Most information retrieval researchers would agree that there is no simple definition of information retrieval. I will attempt an implicit definition by citing examples. Systems which involve information retrieval include Index Medicus, commercial interfaces to Medline (such as Ovid and SilverPlatter), and search engines for the world wide web such as AltaVista or Lycos. Systems that do not involve information retrieval include age and sex registers and most electronic medical records—these might rather be termed database systems.

Fundamentally, the differences between information retrieval systems and database systems stem from the fact that information objects in the former tend to be large, complex, heterogeneous, and loosely structured (such as journal articles, book sections, images, audio or video clips, executable programs), whereas in the latter they tend to be small and simple with known value ranges (birth dates, diagnostic codes, lab results, drug prescriptions). As a consequence, information retrieval systems must address the problem that the relevance of any particular information object to a user's need for information is gener-

ally both partial and uncertain. Not only is it difficult for users to create queries that accurately reflect their needs, but their very conception of those needs is often initially vague and can be clarified only by a "dialogue" with the information retrieval system. Thus information retrieval is inherently an interactive process.

Current medical information retrieval systems

There are several theoretical models of information retrieval. Most doctors will have used information retrieval systems based on the boolean model (after George Boole, a 19th century English logician), and some will have used information retrieval systems based on what I shall call, for the purposes of this paper, the ranking model.

Information retrieval based on the boolean model

Within the boolean model, documents (or their surrogates—that is, titles, abstracts, or lists of key words) are considered to be mathematical expressions. In order to find documents of interest, the user in effect creates another expression consisting of terms such as "myasthenia" and "prognosis" and operators with the meaning of *and*, *or*, *not*. (In most systems each term may be specified as a subexpression using "wild card" characters). For each document in the collection, the information retrieval system attempts a process of unification—that is, it attempts to find substitutions that make both expressions the same.

The boolean model has some attractive features. With appropriate indexing structures, boolean systems can run fast on relatively cheap computers. Also, although creating effective boolean expressions can be difficult, the principle is conceptually simple, and it is easy for users to see why documents do or do not match the query.

A minor disadvantage of this model is that the user must learn the syntax for expressing queries (generally different for each system). Two much more serious problems result from the all-or-none nature of the matching. Firstly, these systems tend to exhibit "brittle" responses to modification of a query: a common experience is that a query with three terms returns no matching documents, but removal of any one of the terms produces thousands of matches. Secondly, the

Boolean system of information retrieval

Information need

A patient with epilepsy wishes to know whether it is safe to accept a job which involves prolonged use of a computer screen. What is the evidence?

User input

The doctor constructs a query expression such as "(video *or* monitor *or* screen*) *and* epilepsy"

Note that "screen*" will match any word beginning with "screen"—such as "screens" and, less appropriately, "screened." It is also usually possible to express metaconstraints—for example, that a term must appear in the document title or that the document is written in a particular language or published in a particular range of years

System response

The system finds all documents that match the query expression and the metaconstraints (any document which mentions epilepsy and any one of the other three terms) and presents these as a list

Examples of systems based on boolean model

BIDS Embase; most commercial Medline search systems

Ranking system of information retrieval

Information need

An old man who has been taking low dose digoxin for many years hears that a friend in similar circumstances has been advised to discontinue his treatment. Should he do likewise?

User input

The doctor types a list of words such as "usefulness of low dose digoxin in old person"

System response

Ranking systems typically ignore words such as "of" and "in," augment the query with synonyms from a thesaurus (in this case adding words such as "utility" and "elderly"), and then derive a quantitative measure of similarity between this augmented query and all the documents in the collection

Examples of systems based on ranking model

Knowledge Finder; most general purpose search engines for the world wide web

matching documents cannot be ordered in any useful manner—the most relevant document might be presented as 39th in a list of 122. Though no more than a nuisance for a research user, this is unacceptable for a clinical user at the point of care.

Information retrieval based on the ranking model

Within the ranking model, documents are considered to be objects described by the values of properties related to the words they contain. You can then devise summative measures to assess the similarity between documents. Crudely, a document mentioning "hypertension" three times is considered more similar to one that mentions it four times than to one which does not mention it at all. In order to pose a query, the user constructs a mini-document (essentially just a list of terms without the need for operators or any system specific syntax), and the documents in the collection can be ranked in order of their similarity to it.

Some vendors describe ranking systems as natural language systems. This they most certainly are not. For example, in response to the query "Minoxidil and hypertension but not hair follicle stimulation," all the top 20 documents returned by one such system were about hair growth.

Many similarity metrics have been used, most of which involve weighting terms according to their distribution in the document collection and making corrections for document length. Although these are computationally demanding, advances in computer power mean that systems based on this model are now commercially available, including search engines for the world wide web that analyse millions of pages.

This approach still suffers from the problem of ambiguous terms, particularly when queries are short: for example, "arms" could reference anatomy or weaponry, "tears" could refer to crying or ripping, and "blind loop syndrome" seldom affects vision. Longer queries (5-10 terms or more) give better results, but the vast majority of users construct queries of only one or two terms, perhaps wrongly applying their experience of boolean systems.

Progress in information retrieval

Currently, nearly all fielded medical information retrieval systems suffer from four further limitations.

- The only medium that can be handled is text

- Each information retrieval system can search only one information collection, and each collection can be searched by only one information retrieval system
- Systems cannot adapt their responses to different user circumstances or behaviour
- There is no linkage between information retrieval systems and patient record systems.

However, research in information retrieval is addressing these limitations, and I give a brief overview of progress.

Multimedia information retrieval

Although some workers have addressed retrieval of audio and video information, most research in multimedia information retrieval is currently focused on static images. There are basically two approaches, one based on tagging and the other on image content.

Image tagging—This approach requires that for each image there is an associated piece of text, so that an image can be retrieved by applying existing information retrieval techniques to its textual partner. The associated text might be created by manual annotation, or inferred automatically from the image context (for example, by assuming that words that appear on the same page as a given image are likely to be related to the image's content). This approach is likely to have wide application in medicine because of the ubiquity of pairing image and text information in clinical medical specialties such as radiology, pathology, and microbiology.

Image content—The second approach involves direct matching of image content in terms of relations between the shapes, volumes, colours, and textures that constitute the image. Work on texture matching seems particularly fertile. I can foresee a general practitioner, during a consultation, submitting a digitised photograph of an unusual rash as a query to an information retrieval system, which then returns the best matching images from an annotated reference collection, together with text describing the characteristic features and diagnoses appropriate to each image.

Distributed information retrieval

Clinicians at the point of care require access to a wide variety of information sources, and the world wide web provides a physical infrastructure capable of supporting this. Because of severe time constraints, clinicians should be able to search all appropriate sources with a single interface and a single query.

Prototype systems are now able to translate users' queries into formats acceptable to several different resources and automatically forward these translations in parallel. Further research is needed on methods of predicting which sources are worth searching for any given information need (since the costs of searching all sources would be too high) and on merging the results for presentation to the user (since ranking metrics in different systems are not comparable).

Cognitive dimensions

Ten years ago a canny clinician who needed information would resort not to an inanimate information retrieval system but to a friendly medical librarian, who would clarify the requirements, undertake a library search when there was time, and assess the results in terms of some measure of relevance. In this "mediated mode" of

use, the information retrieval system stands in a relation of cognitive isolation to the clinician, to the clinical context in which the need for information arose, to the clinical task addressed, and to the outcome for the patient. In the future information retrieval will increasingly be undertaken in "immediate mode": that is, by a clinician personally during a consultation in pursuit of answers to a particular clinical problem.

Accordingly, cognitive issues are now a major focus of research in information retrieval. Four important themes in this work are modelling the user, context, and task (so the information retrieval system can adapt its responses to particular circumstances); interactivity (since clarification of the information need is now the responsibility of the clinician); presentation (the amount of time required to assimilate results is crucial to an immediate user); and evaluation (the information retrieval system should be assessed in terms of its impact on the process of consultation as a whole).

Integration with medical records

In principle the integration of information retrieval systems with electronic medical record systems could have three benefits: firstly, saving time, since users would not have to switch between applications and terms could be manually copied (by "cut and paste") or automatically transferred between systems; secondly, improved retrieval effectiveness, since terms from the medical record could be used to enhance the context specificity of a user's query; and, thirdly, support for system initiative—that is, a system might conceivably raise alerts in situations of unperceived need (in which a user was not aware of important new documents and did not initiate a search).

Several research systems have been constructed. Most depend on automated methods for matching terms in free text or problem coded medical records to UMLS (unified medical language system) Metathesaurus concepts, from which queries based on MeSH (medical subject headings) terms can be constructed. Another approach is manually to compile generic queries, each reflecting one type of commonly occurring question, then allow the user to choose

Recommended reading

Information Retrieval: A Health Care Perspective by William R Hersh (Springer-Verlag, 1996. ISBN 0 387 94454 0). Written by a practising doctor, the book has a strong clinical orientation and is detailed, accessible, and up to date (except for the section on the world wide web, which, having been composed more than six months ago, is inevitably obsolescent)

terms from the medical record to be incorporated into the generic query. Integration with medical record systems is likely to be an essential feature for future information retrieval systems working at the point of care, but considerable further developments are needed for convincing clinical benefit to occur.

Conclusions

Many doctors may have concerns about patients' attitudes to information searching as part of a clinical consultation, or may themselves feel uncomfortable with it. If a doctor was to consult a financial adviser, however, would the doctor not have more confidence in one who could present, explain, and evaluate relevant information so as to reach decisions on a basis of shared responsibility, rather than one who simply told the doctor what was best?

Another concern is the practicality of information searching. Even if they are convinced of its benefit, doctors may nevertheless feel that frequent searches are not compatible with an average consultation time of seven minutes. This concern has considerable force with regard to currently available information retrieval technology, and most searches must be undertaken outside clinical hours—that is, it is necessary to adapt the task to suit the technology. A better solution is to develop new information systems specifically for clinical use—that is, to adapt the technology to suit the task.

I believe that the use of information retrieval as a tool in clinical consultations will become as commonplace as the use of a stethoscope is today. New approaches to information retrieval will be required if the potential benefit is to be maximised. There is also a pressing need for new methods of evaluating information retrieval systems for use at the point of care, not simply in terms of whether users consider the information retrieved to be relevant but taking into account the effect of information retrieval on the clinical process as a whole. Such developments will depend on the existence of a large community of users of clinical information with the knowledge and ability to participate in the critical evaluation of new information retrieval systems.

I thank Professor Keith van Rijsbergen, Department of Computing Science, University of Glasgow, for his help in preparing this article.

Funding: My work investigating new architectures for medical information retrieval systems is funded by the *BJA: International Journal of Anaesthesia*.

Conflict of interest: None.

1 Smith R. What clinical information do doctors need? *BMJ* 1996; 313:1062-8.

(Accepted 5 March 1997)

Information retrieval systems in the future

Information need

A middle aged lady with rheumatoid arthritis who is taking indomethacin has made an appointment to see her general practitioner because of shoulder pain

System response

Before the patient arrives, demographic details, coded problem and symptom lists, clinical notes, prescriptions, and consultation histories are automatically transferred from the computerised medical record to the information retrieval system. This searches multiple sources for recent, context specific, information (such as related to differential diagnosis of shoulder pain in an adult with rheumatoid arthritis, clinical protocols for management of rheumatoid arthritis, and complications associated with indomethacin) and arranges these in a simple menu for perusal by the doctor. In addition, the patient may wish to view the information before or after the consultation

Netlines

Virtual café

- Café Herpé (<http://www.cafeherpe.com/>) is a new web site launched by SmithKline Beecham providing medical information on genital herpes in a friendly virtual café environment. Inside Café Herpé you will find a reading lounge, a buffet, an espresso bar, a terrace, and a gallery. You can even summon "the waiter" to help you find information.

Seek and ye shall find

- Filez (<http://www.filez.com/>) is a new web site aimed at helping you find software on the Internet. The site claims that you can search over 75 million files for specific titles. Category headings let you limit your search for titles to Windows, Macintosh, OS/2, and other platforms.
- The Medical World Search site (<http://pride-sun.poly.edu/>) is a useful addition to the list of sites that allow you to search for medical material on the web. The site not only searches its own index of the major medical sites on the web but can pass on your query to other search facilities on the web. The site also features a powerful set of tools to help you refine your query.

Words, words, words

- One of the problems in mastering a new subject is the jargon. Fortunately, there are several dictionaries or glossaries available on line covering a variety of subjects. Julian Dow's Dictionary of Cell Biology is available on <http://www.mblab.gla.ac.uk/~julian/dict.html> (and in book format for the netphobic). You can access the Free On-Line Dictionary of Computing on <http://wombat.doc.ic.ac.uk/>. There is an Epidemiology Dictionary on <http://epidem13.plantsci.cam.ac.uk/~js/glossary/course-gloss96.html>. And if you have dealings with continental Europe, you may find the Multilingual Glossary of medical terms useful on <http://allserv.rug.ac.be/~rvdstich/eugloss/welcome.html>.

Oasis without blur

- The Mayo Health O@sis (<http://healthnet.ivi.com/>) is an online health network and newsletter sponsored by the Mayo Clinic. The site is fully searchable and contains useful information and advice for the general public, some of which is also of interest to biomedical professionals. I particularly liked the article on the psychopathology of life on line (<http://healthnet.ivi.com/ivi/mayo/9702/htm/shyness.htm>).

Clinical chemistry on line

- The Association of Clinical Biochemists' web site (<http://www.leeds.ac.uk/acb/>) houses various items of general medical interest including an assay finder to help researchers find methods or labs to measure a wide variety of hormones, metals, enzymes, and drugs in bodily fluids. There is also free software on offer and some ready made PowerPoint presentations on osteoporosis, inborn errors of metabolism, and other subjects that should prove useful for teaching.

Boil your water

- At the time of writing water supplies in Hertfordshire are contaminated with the protozoal parasite *Cryptosporidium parvum*. For more information on cryptosporidiosis see the New York Department of Health's fact sheet on <http://www.health.state.ny.us/nysdoh/consumer/crypto.htm> or the entry in the Bad Bug Book (<http://vm.cfsan.fda.gov/~mow/chap24.html>). Also available on line are a report on public health threats associated with waterborne cryptosporidiosis (<http://wonder.cdc.gov/WONDER/anon/ANON173/PREV1.00.ex>) and the US Centers for Disease Control's fact sheet (<http://www.cdc.gov/ncidod/publications/brochures/cryptos.htm>). For a selection of reported outbreaks see (<http://www4.ncbi.nlm.nih.gov/htbin-post/Entrez/query?db=m&form=6&uid=8854449&Dopt=m>).

Mailbase

- Mailbase (<http://www.mailbase.ac.uk>) is Britain's major electronic mailing list service for research in higher education. It hosts a range of medical mailing lists (<http://www.mailbase.ac.uk/other/medi-class.html>) and provides hypertext archives of all its lists. Perhaps the most useful for those interested in medicine on the Internet is the medical-it list (<http://www.mailbase.ac.uk/lists/medical-it/>). British academics can start their own lists via the Mailbase service.

Irish Medical Directory

- Irish readers may be interested to visit the Irish Medical Directory's site on <http://ireland.iol.ie/imd/>. The site contains Irish medical news, information on new products and job vacancies, an email directory of Irish doctors, and links to other Irish medical sites.

Brush up on your Netskills

- If you are still a little nervous about using the Internet, then visit the Netskills site (<http://www.netskills.ac.uk/>). Netskills aims to help the British higher education community make effective use of the Internet for teaching, research, and administration. Of particular interest here is the Online Netskills Interactive Tutorial (<http://www.netskills.ac.uk/TONIC/>), which will help you develop your online skills.

You are what you eat?

- If you or your children are regular visitors to McDonalds, you might be interested to visit the McSpotlight site on <http://www.mcspotlight.org/>. The site has grown out of the longest and arguably one of the most important trials in English legal history, the so called "McLibel case." The trial has already lasted over 300 days, with the judge's verdict due some time after Easter. One of the central issues of the case is the nutritional value of fast food. Judge the evidence for yourself (and there's plenty of it on both sides of the case) on <http://www.mcspotlight.org/issues/nutrition/index.html>.

Compiled by Mark Pallen
email m.pallen@qmw.ac.uk
web page <http://www.qmw.ac.uk/~rhm001/mpallen.html>