BEME Guide No 3: Systematic searching for evidence in medical education—Part 1: Sources of information*

ALEX HAIG[1] & MARSHALL DOZIER[2]
[1]*NHS Education for Scotland, Edinburgh, UK;* [2]*University of Edinburgh, Erskine Medical Library, Edinburgh, UK*

SUMMARY *Searching for evidence to inform best practice in medical education is a complex undertaking. With very few information sources dedicated to medical education itself, one is forced to consult a wide range of often enormous sources—and these are dedicated to either medicine or education, making a medical education search all the more challenging. This guide provides a comprehensive overview of relevant information sources and methods (including bibliographic databases, grey literature, hand searching and the Internet) and describes when they should be consulted. The process of constructing a search is explained: identifying and combining core concepts, using Boolean algebra and search syntax, limiting results sets, and making best use of databases' controlled vocabularies. This process is illustrated with images from search screens and is followed by numerous examples designed to reinforce skills and concepts covered. The guide has been developed from the ongoing experience gained from the systematic searches conducted for the Best Evidence Medical Education Collaboration, and concludes by looking ahead to initiatives that will shape future searching for medical education evidence.*

Introduction

> Knowledge is of two kinds. We know a subject ourselves, or we know where we can find information upon it. (Dr Samuel Johnson, 18 April 1775, in James Boswell, *Life of Johnson*)

Like other teachers, trainers and researchers, medical education professionals are increasingly expected to base their practice on best evidence—yet this evidence is located in a bewildering number of diverse and incompatible sources. Those who set out to search for evidence are frequently not convinced that they have been looking in the best sources, much less that they have found the most relevant and useful papers. The explosion of information over the last decade has produced an enormous challenge for those seeking to navigate the evidence base quickly and effectively—but it has also produced enormous opportunities for those who do learn how to search successfully.

Evidence to support best practice is found in many formats and sources. *Database*s are nearly always the most abundant resource, so this guide concentrates on this medium. However, this guide also covers accessing the *grey literature*, searching subject gateways and the web at large, as well as other search methods and sources to ensure that a search is as comprehensive and cohesive as required.

There are many reasons to consult the evidence base: to answer a specific question, to identify experts to consult or work with, to determine what evidence already exists before starting research, or to keep current with what is being practised. Recent developments in medical education, such as the focus on student-centred learning and the increasing attention given to research and scientific method (General Medical Council, 2003), have increased the importance of empowering the individual to find the evidence he/she requires by him/herself.

Individual queries range from simple to complex. This guide covers the wide range of skills and sources required to answer simple individual questions quickly and appropriately, but its larger aim is to provide guidance for systematic and comprehensive retrievals of evidence. The Best Evidence Medical Education (BEME) Collaboration (BEME Collaboration, 2003) is an international group producing systematic reviews for medical education. It is through the experience gained by BEME's efforts that this guide has been produced to help the researcher overcome the challenges that are faced when searching for evidence in medical education.

The challenges

> It is a very sad thing that nowadays there is so little useless information. (Oscar Wilde, 1854–1900)

The foremost challenge in searching for evidence in medical education is that there are very few comprehensive sources dedicated to the profession. For most queries, *bibliographic database*s are the medium likely to contain the majority of evidence, yet there is no indexed database for medical education. Whilst keyword databases do exist, these lack the relational *subject heading*s that enable accurate and consistent searching. Instead, one must turn to either medical (e.g. Medline) or educational (e.g. ERIC) databases—neither of which adequately collects or indexes medical education content.

*Database*s index publications with descriptors (or *subject heading*s) that describe the concepts central to each of the individual documents. Often *bibliographic database*s are focused on either medicine or education but they are frequently inadequate for retrieving citations in medical education. Even when adequate subject headings for medical education do exist they are incorrectly and/or incompletely used because these databases' emphasis is (understandably) focused on meeting their own objectives—namely describing medicine **or** education.

Correspondence: Alex Haig, NHS Education for Scotland, 22 Queen Street, Edinburgh EH2 1NT, UK. Tel: +44 (0)131 247 6618; fax: +44 (0)131 225 9970; email: alex.haig@nes.scot.nhs.uk
*Part 2: Constructing searches, will appear in the September issue.

A BEME case study

> It is in fact nothing short of a miracle that the modern methods of instruction have not yet entirely strangled the holy curious of inquiry. It is a very grave mistake to think that the enjoyment of seeing and searching can be promoted by means of coercion and a sense of duty. (Albert Einstein, 1879–1955)

The BEME Group FEENASS (*fee*dback *i*n *ass*essment), which is researching the use of feedback in medical education, provides an interesting illustration of the challenges in retrieving medical education information. Searches designed to *scope* their topic, 'feedback in assessment', were conducted in spring 2001 to determine the extent of the evidence available across the major *database*s. These were measured for search *sensitivity* and *specificity*.

Sensitivity (or *recall*) measures what percentage of the total number of known citations on a topic was actually retrieved by the electronic search (the total, or gold standard, is determined by *hand searching* the journals). Sensitivity ranged from 6.5% to 19.6%, depending on the database. This means that even in the database with the highest sensitivity, four-fifths of relevant citations were not appearing. As atrocious as these results are, those experienced in medical education searching will hardly find them extraordinary.

Specificity (or *precision*) measures what percentage of the search results was actually relevant to the query (i.e. the positive predictive value). For the FEENASS group specificity across the databases came in at 17.5%, about average for BEME pilot groups, which ranged from 6% to 34% (Haig, 2001). Clearly, all groups were finding that most results were irrelevant, but the negative impact for the BEME pilot groups' time and other resources was exacerbated, given that they had to look at the collective results of many databases—often totalling over 10,000 results.

It is trying enough for an individual conducting a quick search to find that his/her best strategy still produces a majority of false *hit*s, but for a systematic search that nets thousands of results the experience can be frustrating and forbidding.

Search strategies can be written to improve *specificity*, but not without sacrificing *sensitivity*. A systematic review group cannot afford to miss results, and with the growing emphasis on evidence-based practice, other groups and individuals are becoming less able to do so. Before we examine how to improve search skills to reduce these difficulties, it is important to understand how they arise.

The failings of Subject Headings examined

A closer look at how the world's largest medical *database*, Medline, indexes the concept of feedback clearly illustrates the problems with searching *subject heading*s. The FEENASS Group's concept of feedback is common with that of the medical education profession. When searched as a subject in Medline, feedback maps to these different subject headings: *feedback*; *feedback, biochemical*; and *feedback, psychological*. The scope notes define them as follows:

- *Feedback*: A mechanism of communication within a system in that the input signal generates an output response which returns to influence the continued activity or productivity of that system.
- *Feedback, Biochemical*: A mechanism of communication among life processes to coordinate development, reproduction, and homeostasis. In humans, feedback loops are especially important for communication between organs that are spatially separated. Virtually all hormones from the nervous and endocrine systems are under feedback control: by peripheral hormones, cations, metabolites, osmolarity or extracellular fluid volume.
- *Feedback, Psychological*: A mechanism of information stimulus and response that may control subsequent behaviour, cognition, perception, or performance. (From APA *Thesaurus of Psychological Index Terms*, 8th edn: National Library of Medicine, 2003).

Feedback, Biochemical is obviously not relevant as it is used to index records concerning physiology, metabolism, immunology etc. Of the other two possibilities, *Feedback* would appear to relate to mechanical systems and processes, and *Feedback, Psychological* to the learning. However, when *Feedback* is combined with the subject heading *Education, Medical* it finds 367 results while *Feedback, Psychological* locates only four (Ovid Technologies, 2002).

Unfortunately Medline's most suitable descriptor, *Knowledge of Results* (*Psychology*) does not appear when feedback is mapped to the *subject heading*s:

- A principle that learning is facilitated when the learner receives immediate evaluation of learning performance. The concept also hypothesises that learning is facilitated when the learner is promptly informed whether a response is correct, and, if incorrect, of the direction of error.

Despite the suitability of this term, there are only seven citations found when it is searched with *Education, Medical*. The problems encountered can therefore be described as a mixture of failure of the system to map to the most appropriate term, lexical ambiguity (or overlap) between possible subject headings, and *indexing* with the least appropriate of possible subject headings.

The situation is made more problematic by the fact that the concept of feedback is often not indexed at all. If feedback is searched as a *free-text* word and combined with the *subject heading* Education, Medical (with duplicates using the feedback subject headings removed) over 700 citations are found. The vast majority of these results are relevant to the topic to varying degrees—but often match the topic as well as, or better than, those documents that have been indexed with feedback descriptors.

Conclusions from the case study

The problems illustrated by the feedback example are representative of the challenges facing searching for medical education evidence in general. Searches have low *sensitivity* (i.e. many relevant citations are missed) and low *specificity* (i.e. many results that are returned are irrelevant). Nobody who embarks on a search wants to miss key results (least of all for a systematic search), nor does anyone wish to drown in a sea of irrelevant results. It is the purpose of this guide to describe how to improve sensitivity while not compromising specificity.

Covering multiple databases

> What information consumes is rather obvious: it consumes the attention of its recipients. Hence, a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it. (Herbert Simon, Economist)

It has been well documented that medical searches must cover multiple *database*s if they are not to risk missing substantial amounts of significant evidence (Brettle *et al.*, 1998; Minozzi *et al.*, 2000; Avenell *et al.*, 2001). In the case of medical education, consulting multiple sources is even more important—for a number of reasons.

As mentioned above, the evidence is not contained in sources dedicated to medical education itself and therefore medical, educational and other specialist sources need to be identified. It is critical that search strategies for each source be tailored to the databases' individual specifications.

Researchers will discover that database coverage is patchy. Medline, the world's largest medical database, indexes less than a quarter of existing medical journals. While no similar measure has been taken for medical education, it is reasonable to assume that the figure would be similar, if not worse. For example, neither *Education for Primary Care* nor *Education for Health* is indexed in Medline; however, both are indexed in Embase.

In addition, *indexing* is also incomplete. Medline and Embase both index the journals *Medical Teacher* and *Medical Education*, but until recently both journals were partially indexed with only selected articles appearing in the databases. Although both journals are now fully indexed, past evidence from these key journals will be missed unless they are indexed retrospectively.

Finally, as the feedback in assessment example proved, even when journals are fully indexed the *subject heading*s often fail to capture the subject requirements of medical education. Therefore every effort must be made to perfect search techniques to overcome these problems.

There is an overlap between and within medical, educational and other databases but despite the potential issue of duplication these problems make it far more advantageous to search all possible relevant sources (duplicates can be eliminated in seconds by bibliographic software anyway).

Review of information sources

> Lord Ronald said nothing; he flung himself from the room, flung himself upon his horse and rode madly off in all directions. (Stephen Leacock, 1869–1944, *Nonsense Novels*, 1911)

The number of potential sources relevant to a search in medical education is vast and confusing. There exists, however, a principle core of databases that should be consulted for any comprehensive search, as they are almost certain to contain essential evidence; secondary databases should be employed according to the nature of the search topic, and the time and other resources allocated to the search.

Core bibliographic databases

Medline. The National Library of Medicine in Maryland, USA produces Medline (Index Medicus). It contains over 73,000 citations indexed as medical education and over 300,000 additional citations that are considered educationally relevant. It is available through commercial vendors and freely available through numerous websites, notably PubMed.

Embase. This is the second largest medical database and is owned by Elsevier Science in The Netherlands. Access is by subscription only. Embase contains over 43,000 citations indexed as medical education and more than 100,000 that are related to education in a health environment. Embase tends to index more European journals, while Medline is more focused on North American research. The database has very strong pharmaceutical coverage and indexes several key titles in medical education not covered elsewhere.

CINAHL. The Cumulative Index to Nursing and Allied Health Literature is the world's largest database for nursing and the professions allied to medicine. Although there are fewer than 1500 citations indexed as medical education there are over 100,000 educationally relevant ones that could inform a medical education query. Access is by subscription.

ERIC. The Education Resource Information Centre is the world's largest education database (1,000,000 + records) and is freely available on the web or by subscription. Although the emphasis is on primary and secondary education, there are over 17,000 citations related to medical education. Many more will be relevant in a supporting context, providing evidence in education that could be applied to medical education (e.g. use of multiple-choice questions).

BEI. As it has a British focus, the British Education Index is much smaller than ERIC but like its larger American counterpart contains citations that are directly or indirectly relevant to answering queries in medical education. Access is by subscription, but partial free access is available on the web.

PsycINFO. While PsycINFO obviously concentrates on psychiatric and psychological content, there are over 4000 records indexed as medical education with well over 100,000 concerning education in a broader context, as well as teaching and learning. Access is by subscription.

Note on database access. Most databases are accessed by subscription, with organizations or individuals paying to access the raw data through search software vendors such as Ovid, Dialog or SilverPlatter. Several key ones (notably Medline and ERIC) are freely available on the web. Paid subscription access, with the possible exception of PubMed, does provide more rigorous and comprehensive searching.

Additional databases

Additional databases should be consulted when one needs to be as comprehensive as possible (such as for a systematic review), or when the database compliments the search by the nature of the subjects it indexes.

Additional keyword databases. There are two keyword databases that collect references in medical education, and both are freely available on the web. Although they lack a professional indexing system specific to medical education, instead relying on users attempting to match their search terms with keyword lists, this does compromise accuracy. These databases contain thousands of records not found elsewhere and so they cannot be ignored.

(1) *Research and Development Resource Base* (*RDRB*): *www.cme.utoronto.ca/rdrb*. RDRB collects information to assist study of physician performance, programme evaluation, change and healthcare outcomes. It collates literature from a broad range of continuing education topics from databases such as Medline, Embase, ERIC and CINAHL. It also contains conference abstracts from sources such as the Society for Academic Continuing Medical Education and the Alliance for Continuing Medical Education. The database is hosted at the University of Toronto.

(2) *Topics in Medical Education* (*TIMELIT*): *www.timelit.org*. TIMELIT contains nearly 50,000 records of direct relevance to medical education covering a broad range of subjects. The database combines references extracted from the major databases with specially selected citations from a comprehensive range of sources. TIMELIT's strength is these latter citations—many of which are not to be easily found by any other means. The database is hosted at the University of Dundee.

Additional indexed databases The following databases will not necessarily be relevant to every search but should be considered for most searches. Medical education evidence is widely dispersed so this is not an exhaustive list, as less typical searches will require the researcher to consult additional sources. In these instances consulting an information professional is advisable in these cases.

(1) *AMED (Allied and Complementary Medicine Database)*. The researcher should consider AMED for educational searches in complementary medicine, palliative care and the professions allied to medicine (including podiatry, physiotherapy, occupational therapy and rehabilitation).

(2) *ASSIA (Applied Social Sciences Index and Abstracts)*. ASSIA is a social science database of more than 650 English-language titles. Although not directly relevant to many medical education searches it does index relevant sociological content not found elsewhere.

(3) *BNI (British Nursing Index)*. The British Nursing Index is much smaller than CINAHL, but does contain educational material (obviously largely British) not found in CINAHL and the other large databases. An essential source for comprehensive searches in nursing and the professions allied to medicine.

(4) *EBM Databases—Cochrane databases*, *ACP Journal Club and DARE*. The various evidence-based medicine databases contain full reviews, appraisal of reviews, protocols for research, trials and commentary. Educational material forms a small but growing proportion of this evidence, though much of this is already contained within Medline, which permits more robust searching (use of *controlled vocabulary* and *subheading*s).

Though the EBM databases frequently provide no unique content to an educational search they occasionally can provide invaluable material—particularly in the form of commentary/appraisal of research, contacts and bibliographies (often containing *grey literature—see below*).

(5) *HMIC (Health Management Information Consortium)*. HMIC is actually the combined files of three separate databases collating health management information. HMIC indexes a surprising amount of medical education material where it meets management. The database is good for *grey literature* and has strong European coverage.

(6) *SOCIOFILE*. The premier database for sociology, social planning and policy, SOCIOFILE can provide supporting and contextual evidence for medical education searches that have a significant societal dimension.

Other methods of searching

SCI (Science Citation Index). SCI is actually another *database*, but one that uses a rather different method of searching. SCI indexes abstracts and bibliographic information from nearly 4000 journals—but its true strength is that it makes possible *cited reference searching* (i.e. retrieving all citations that reference a particular citation, author, author within a journal, etc.). This cited reference searching, along with SCI's broad interdisciplinary *scope*, reveals relationships between research and disciplines that would otherwise remain undiscovered. Although it is an expensive database (and therefore not always available) and does not permit robust subject searches, SCI should be searched when possible, particularly if your search topic would benefit from evidence across professions.

The ISI Web of Science/Web of Knowledge provides unified access to SCI in addition to the Social Science and Arts and Humanities Index.

Ancestry searching. Ancestry searching is the process of searching the bibliographies of relevant papers to discover references missed by other methods. The Science Citation Index has thankfully reduced this time-consuming task to a few keystrokes for many journals, but there still will be crucial papers not included in the SCI. In these cases bibliographies should be matched with the search results you are compiling.

Hand searching Hand searching is literally the searching of print (or electronic) journals volume by volume, issue by issue, and article by article. This obviously requires large amounts of time and resources and is an extremely tedious process. However, for a comprehensive search such as is required for a systematic review, hand searching may be necessary.

Medical hand searches have most notably been conducted by the Cochrane Collaboration to identify controlled trials. This work has demonstrated that without hand searches not only were large numbers of papers missed (Hopewell *et al.*, 2000), but specialized non-English journals should be included as well if a search is ever to be comprehensive (Bereczki *et al.*, 2000).

Hand searches conducted by BEME to date suggest a similar situation exists for medical education (Haig & Coupar, 2001; Harden *et al.*, 2002). While these were subject-based

hand searches and not searches for particular study designs, the results demonstrated that hand searches picked up citations not found by *database*s. This was partly due to the inadequate subject *indexing* (mentioned above) by databases, as well as the fact that many crucial medical education titles (notably *Medical Education* and *Medical Teacher*) have only recently started to have every article indexed.

Ideally all possibly relevant journals would be hand searched for all queries. However, in practical terms the extent of *hand searching* comes down to the availability of resources (time, labour, funds and access to the desired journals themselves) to conduct the hand search. Hand searches should be planned well in advance, taking into consideration what resources can be invested, how to integrate results, and which journals to prioritize for searching.

Experts in the field. Finally, for a search to be truly comprehensive, efforts should be made to contact leading figures with professional experience in the subject. It can be difficult to identify and contact experts, though the *Internet* has made the process much easier. There is no guarantee that experts will have the time or inclination to reply, but they can potentially be a very rich source of information that is not easily found by other means.

Grey literature.   Grey literature is best defined as:

> That which is produced on all levels of government, academics, business and industry in print and electronic formats, but which is not controlled by commercial publishers. (4th International Conference on Grey Literature, 1997)

From this definition it is obvious that a considerable number of publication formats constitute grey literature, including:

- academic papers;
- census data;
- committee reports;
- conference papers;
- corporate documents;
- discussion papers;
- dissertations;
- government reports;
- house journals;
- market surveys;
- newsletters;
- ongoing research;
- preprints;
- proceedings;
- research reports;
- standards;
- technical reports;
- theses;
- trade literature;
- translations;
- working papers.

(1) Challenges.   The variety of publication types obviously does not make the task of systematic searching any easier. It is essential to remember that 'grey' is not a comment on quality, but refers to the medium of (non-commercial) delivery.

Grey literature presents a significant challenge when searching systematically for evidence. While it is a primary source of evidence, it is notoriously difficult to locate, retrieve and manage.

Not only is grey literature diverse in format but it exists in a fluid environment. The large changes that the commercial publishing industry has experienced has also affected grey literature, most significantly the *Internet* and networked personal computers.

The creators of grey literature are often universities, research institutes, industry or government; because these organizations rarely have the widespread dissemination of information as a primary objective, retrieving the items becomes more challenging. Indeed, grey literature can often be intended for a partially restricted, or even confidential, audience.

(2) Advantages.   Grey literature does enjoy many key advantages over commercially produced literature. Most grey literature can be created and distributed comparatively quickly as the process tends to be less structured and formalized. Without lengthy procedural delays, this rapid dissemination allows the evidence to reach the level of practice much more effectively, and grey literature is therefore often a type of evidence that is received and used to inform opinion. While this in itself does not influence systematic searching, it should be noted that grey literature can have a strong effect on policy makers. Grey literature is often excellent for providing context for findings contained in commercial sources. Grey literature can link discoveries and provide insight to the decision making and the environments in which it occurred. Finally, because this material is frequently in-process, policy oriented, informal and/or discipline-specific, it is also pivotal for intra-professional communication, because research/scientific policy, protocols and journalism are nearly always grey.

Grey literature is often valued for being succinct; with very focused content it is often easy to isolate relevant portions of evidence quickly. At the same time the material contained in grey literature is likely to have been thoroughly researched. This is particularly the case with technical reports and government documents; a grey version of a document may contain far more detail than what is made available when the item becomes commercially published.

Another aspect of grey literature that makes its consideration critical is the sheer growth in volume of information contained in grey literature. It has been estimated that the volume of grey literature is currently growing at three to four times the rate of commercially available literature. In 1992 the British Library Document Supply Centre (BLDSC) at Boston Spa held some 3,000,000 items that were considered to be grey literature and these items had been collected over the previous 30 years. However, by 2000 the collection had grown to over 17 million items (Helmer, 2002).

Grey searching can reveal large amounts of evidence not found by traditional searching. One study looked at grey searches across a variety of medical topics and found that 26.1–41.6% of the total evidence base was found by grey searches. A wide variety of types of study design was found, and interestingly one grey search found 33.6% (73) randomized controlled trials—not only the gold standard for most systematic medical studies, but the type of study

design one might assume would be comprehensively covered by conventional means (Helmer, 1999).

Perhaps the most import reason for including grey literature in a systematic search is the impact is has when considered against other types of evidence. A published study (McAuley *et al.*, 2000) found 33% of meta-analyses they examined contained grey literature. In these studies grey literature accounted for between 4.5% and 75% of the total number of studies. Critically, they found that the commercially published studies averaged a 15% larger estimation of effect than the grey literature. Commercially published literature has long been suspected of being more likely to report positive findings, and it is therefore crucial to include grey literature to give a balanced representation.

The *Internet* has had an enormous effect on the growth and diversification of grey literature. There are many electronic forms that grey literature appears in, including personal webpages, listservs, *usenet*, *blogs*, digital libraries and various other types of electronic files. The variety of electronic sources can make identifying and using these sources effectively a time- and resource-intensive activity, but electronic sources are certainly quicker to access than paper methods.

This diversity of format also occurs in the print form and makes the acquisition, storage and retrieval of grey literature very difficult for libraries. It can be both difficult and costly for libraries to acquire items, given the varying methods by which grey literature is distributed. As a large amount of grey literature is produced by and for conferences it is often necessary to be aware of the event in advance or perhaps even attend it. The distribution of grey literature tends to be limited to a small number of organizations that have been identified previously as being of interest. In terms of the physical storage of grey literature, libraries have difficulty in managing materials that have no standard format or size. Finally, grey literature very often requires original and time-consuming *indexing*/cataloguing because the descriptive or bibliographic information required is not available for purchase off the shelf from commercial cataloguing services.

To balance these difficulties, there are a number of factors that have made grey literature more attractive to researchers and practitioners than commercially published literature. The cost of commercially published literature has been increasing at a much higher rate than inflation for many years, with some individual publishers becoming increasingly dominant through mergers and takeovers within the industry. It has also been argued that publishers have been reticent in adopting new technologies. Grey literature, by contrast, has readily adapted to technological change and appeared in new formats and by new means. Such flexibility is certainly a strength in a rapidly developing market.

Ideally there would be a centralized authority that would store and distribute grey literature in medical education. Unfortunately the prospect of such an authority undertaking this role in the near future is unlikely. One organization, Grey Net, which looked as if could become a generalized central authority by providing network services, hosting conferences and producing an international journal for grey literature, has recently been shut down. The web does compensate for this to some extent as sites have started to collect and supply grey information more comprehensively. To date, however, this process has tended to be very unevenly distributed with the UK and America making some considerable progress while other countries and regions are hardly represented.

As with all types of evidence, grey literature needs to undergo a rigorous evaluation of quality. Organizations produce copious amounts of grey literature with much of it being of limited utility; the immediate task one faces when using grey literature for research purposes is in discovering the signal from the noise. When evaluating grey literature it can be complicated because grey literature is frequently not in a standard format or design; for example, foreign-language items may need to be translated. Finally, because little grey literature is peer reviewed, quality varies considerably and a thorough evaluation process is essential.

When evaluating grey literature one should consider assessing the statement of findings, testing any hypotheses presented, checking the authors' details and experience, and looking at the source of the item (where was it found?; what are the institutional affiliations?). In addition, evaluating the research methodology and data-collection procedures is as necessary for grey literature as it is for conventionally available studies.

(3) Sources including grey literature. Many of the databases in Information Sources list types of grey literature in varying degrees. In addition to this, consider the following sources:

- Adobe PDF: *searchpdf.adobe.com*—a *search engine* exclusively for Adobe documents—a format in which grey literature often appears.
- CHID: *chid.nih.gov*—Combined Health Information Database for the American federal government.
- ClinicalTrials.gov: *www.clinicaltrials.gov*—an index of clinical trials, ongoing and complete. While educational trials are understandably few, there will be some of relevance to particular topics.
- CRISP database: *www.commons.cit.nih.gov/crisp/*—Computer Retrieval of Information on Scientific Projects collates US federally funded biomedical research. Its educational content has an emphasis on patient education.
- DARE (Database of Abstracts of Reviews of Effectiveness): *nhscrd.york.ac.uk*—a collection of appraised reviews, economic assessments and health technology research hosted by the University of York. Educational content is small, but of high quality.
- Dissertation Abstracts: *www.lib.uni.com/dissertations/gateway*—the most recent two years of dissertations and previews. This free version is provided by Digital Dissertations, but the full database and full text require a subscription.
- National Guidelines Clearinghouse: *www.guideline.gov/index.asp*—educational guidelines here are almost exclusively public health education. Similar sites exist for other countries e.g. NICE: www.nice.org.uk for England and Wales.
- National Research Register: *www.doh.gov.uk/research/nrr.htm*—a *database* of current and newly completed research projects, the majority of which are funded by Britain's National Health Service.
- National Technical Information Service: *www.ntis.gov/search*—provided to search for American government scientific, technical and business information.

RIGHTS LINK

- OMNI: *omni.ac.uk*—an example of a high-quality gateway to *Internet* pages in health and medicine.
- TRIP: *www.tripdatabase.com*—Turning Research Into Practice simultaneously searches over 75 high-quality medical sites.

(a) *Library catalogues*. The following national library catalogues all index *grey literature*. These can be searched in English, but obviously non-English equivalents can be consulted when needed.

- British Library: *www.bl.uk*
- COPAC: *www.copac.ac.uk* (s combined catalogue of Britain and Ireland's largest universities)
- Library of Congress: *catalog.loc.gov*.
- National Library of Australia: *www.nla.gov.au/catalogue/*
- National Library of Canada: *www.nlc-bnc.ca/7/2*
- National Library of Medicine: www.locatorplus.gov
- National Library of the Netherlands: *www.kb.nl/kb/resources/frameset_catalogi-en.html*

(b) *Government sites and professional organizations*. Government and professional websites can be an excellent source for grey literature evidence. The following sites are examples relevant to Scotland/Britain, but obviously each relevant country's equivalents should be consulted.

- British Medical Association: *www.bma.org.uk*
- DOH: *www.doh.gov.uk/index.html*
- General Medical Council: *www.gmc-uk.org*
- Royal Colleges: *www.rcplongon.ac.uk/general/gen_ links.htm*
- Scottish Executive: *www.scotland.gov.uk/pages/default. aspx*
- Scottish Parliament: *www.scottish.parliament.uk*

(c) *Subscription or paid sources*

- Conference Papers Index—a database that includes the Cambridge Scientific Abstracts.
- Health Management Information Consortium—a combined database of health management and healthcare from the UK Department of Health and Stationery Office.
- Index to Theses: *www.theses.com*—a complete list of theses and abstracts accepted by British and Irish universities.
- Northern Light: *www.northernlight.com*—formerly a free search engine, this site is now of potential interest for its Special Collection of paid access documents. Searching is free.
- SIGLE: *www.cas.org/ONLINE/DBSS/sigless.html*—the System for Information on Grey literature in Europe is a large grey literature database produced by the European Association for Grey literature in Europe.
- Web of Science Proceedings Monthly—updates from nearly 5000 international conferences.
- World Cat, Proceedings First: www.oclc.org/home—the world's largest unified library catalogue.
- ZETOC: *zetoc.mimas.ac.uk*—a British (British Library) example of a database and alerting service from the tables of contents of journals, as well as conference proceedings.

(d) *Archives*

- Biomed Central: *www.biomedcentral.com*—an independent medical web publishing house with a variety of free peer-reviewed journals, including medical education (BMC Medical Education).

- Netprints (BMJ): *clinmed.netprints.org/home.dtl*—"A repository of non-peer reviewed original research" provided by the BMJ Publishing Group.
- PubMed Central: *www.pubmedcentral.com*—unified free access to a number of top peer-reviewed journals. Papers will inevitably appear in the databases above, but this site allows free access to the full text.

Web searching Searching the web is an essential part of any systematic search. The benefits of using web searching to identify trials (ongoing or unpublished) have been well documented in medicine (Eysenbach *et al.*, 2001). Because evidence relevant to medical education can be found in such disparate and numerous locations, a web search is essential if the search is to be considered comprehensive. A web search can be considered an integral part of the *grey literature* search—it can find evidence that is not located anywhere else.

Web searching presents its own challenges and rewards. The most immediate is the sheer size of the web itself. At present there is no accurate measure on the number of pages the web contains. The largest *search engine*s access over three billion pages, though the total number of public and private pages is thought to be many times larger still. The sheer volume of pages is a definite obstacle in locating relevant pages from false *hit*s, but the enormous size of the web also means there is relevant material available. And because no search engine covers the entire web (nor is any expected to in the immediate or medium future) one can never assume a search, even across multiple engines, is exhaustive.

The accuracy of search engines has been improving consistently, however, and within the last 12 months it has become routine for engines to search files other than only web (.html) pages. Microsoft Word and Adobe Acrobat files are perhaps the two most important file types to be included alongside the usual web pages, as both these formats add a huge wealth of potential evidence stored on academic and other servers.

(1) Searching the web is different from database searching. A *database* search using *controlled vocabulary* is pre-coordinated, as searchers select *subject heading*s that have already been created and assigned to the *database* records. A web search is post-coordinated, as searchers must select their own terms and leave it to the *search engine* to match the terms to web pages and rank them by relevance. Theoretically a post-coordinated search is more flexible and ultimately more accurate; however, as anyone who as ever searched the web can attest to, this has yet to occur in practice. Web searches produce enormous amounts of irrelevant *hit*s.

While at present the web is too murky and unwieldy to provide clear and succinct searches, ongoing initiatives may make systematic web searches more of a possibility in the future. If *metadata* standards such as Dublin Core and IMS are widely implemented, the web in effect would have a more powerful *controlled vocabulary* than most *bibliographic database*s. In the meantime, relevancy ranking of *search engine*s continues to improve, as anyone who has been searching the web over the last few years can informally testify.

'I must've seen it in a USENET posting'; that's sort of like hearsay evidence from Richard Nixon. (Blair Houghton)

From the beginning questions have been raised about the accuracy and value of much of the evidence found on the web. Evidence found on the web is almost certain to lack traditional quality controls like peer review. While the quality information on health websites is thought to be improving (Pandolfini *et al.*, 2002), it would be wrong to assume this was necessarily the case with evidence relevant to medical education.

The medical world is also further along in establishing models for rating website quality. A good example of a guide for judging the quality of medical sites is Wilson (Wilson, 2002). The author outlines a range of tools that can be employed, along with respective costs and benefits. These include: Codes on conduct, Self-applied code of conduct or quality label, User guidance system, Filtering tools, and Quality and Accreditation labels awarded by third parties. While these are only partially applicable to medical education, the broad criteria for appraising websites is interdisciplinary:

(i) *Authority*—are the authors clearly stated and reputable?
(ii) *Accuracy*—does the information appear accurate?
(iii) *Currency*—is the material up to date?
(iv) *Scope*—does the subject area match your research question?
(v) *Objectivity*—can you detect bias?

Instruments to assess the validity of medical information on the web are not uncommon. One prominent one is the DISCERN Tool (www.discern.org.uk/) created by the University of Oxford and the British Library. While it would have to be adapted to consider the educational context, the principles remain the same.

This now moves from searching to critical appraisal, which is not covered in this guide. It is hoped that the ongoing work of the BEME Collaboration will produce common criteria to appraise websites alongside the array of study designs being examined.

(2) Search engines and sites. The one constant in web searching is change. *Search engine*s rise and fall, as do the techniques they employ. Some points to bear in mind before undertaking a web search:

- No single search engine covers the entire web.
- Web searches are not 'live' but are searching the stored results of the last sweeping search of the web the search engine has made—only when you click on the results are you making a live connection.
- Most search engines offer *Boolean algebra*, though this frequently is not immediately apparent and often is only seen in their advanced search options.
- Some search engines automatically link words (without Boolean commands) with AND, while others default to OR.
- Results are ranked with what are judged to be the most relevant ones first; every search engine has a different way of doing this, looking at the documents' titles, text, links and associations, and hidden *metadata*.

- Some search engines initially offer only paid results, i.e. pages that companies have paid to appear when certain keywords are searched.
- There are probably fewer search engines now than there were a couple years ago; fierce competition has led to closures and alliances where one search engine merely retrieves results from another (competition has also noticeably improved results).
- The recommendations below will change over time...

(3) Search engines—recommended

Google

- *www.google.co.uk*—From obscurity five years ago, Google has arguably become the world's most popular *search engine*—with good reason. The first to use an algorithm that analyses what other sites a website links to in addition to its text and titles, Google provides accuracy that very few can rival.

The basic search page is simple and free from clutter. It defaults to a Web search, but can also retrieve Images, News, postings from Usenet and Directory results.

Usenet is a massive collection of news group forums on nearly every imaginable topic; obviously the vast majority will be irrelevant, but with a *focus*ed query (e.g. 'summative assessment') one can quickly cut to the significant *hit*s. Google's Directory is a human-created hierarchical collection of subject categories—good for broad subject searches, but not for finding obscure papers.

The Advanced Search page offers (limited) Boolean commands (through text boxes), limits (such as language, file type and date), pages that link to a page, and pages that are conceptually similar to a page. In addition you can focus a search to specific universities' web pages (at the time of writing this includes American and Canadian universities).

Google also has cached pages—or stored pages from its previous searches of the web. This store allows users to view web pages that may no longer be accessible.

alltheweb
find it all

- *www.alltheweb.com*—Rivalling Google in size, Alltheweb (also known as FAST) is nowhere nearly as well known but can also rival Google's accuracy. Alltheweb also defaults to a simple Web search, but also covers Images, Audio, Video, News, FTP (file transfer) sites.

The Advanced search allows slightly more powerful Boolean searching than Google, as well as the standard range of limits. In addition there are Word Filters which, most usefully, allow a user to specify what words should occur in the documents. For example, if one was interested in the educational uses of portfolios in general, but specifically in medical education, one might search 'portfolios education' in the main box while entering 'medical' in the Should Word Filter.

Unfortunately, Alltheweb places paid listings (sponsored pages) at the top of its results list. At the time of writing, the

company looked to as if it was to be acquired by another, so changes in the future are perhaps more likely than usual.

TEOMA

- *www.teoma.com*—A relatively new *search engine*, Teoma is not nearly as large as Google or Alltheweb. Nevertheless, it is well worth using and watching because it has created a powerful link analysis algorithm that produces extremely accurate results.

Teoma's results fall into three categories Results (standard web pages ranked for relevancy), Refine (suggestions on how to narrow large searches) and Resources ('expert' subject similar pages). While the last two additions are not the primary reason for using this engine, they can often bring surprisingly applicable content to a web search.

Teoma also has an Advanced Search that is comparable to two mentioned above.

(4) Other choices

WiseNut

- *www.wisenut.com*—A new *search engine* that has shown considerable promise, Wisenut is currently being redeveloped. The engine has a simple interface and largely bases its results on links analysis. With a successful upgrade this engine could become a very significant player.

HOTBOT

- *www.hotbot.com*—In years past Hotbot was an impressive engine in itself. Today, however, it searches four separate engines: Google, Alltheweb, Teoma and Inktomi. While the first three are engines in their own right, Inktomi is not but is a separate massive collection of web pages that can be accessed by a number of engines. Hotbot is probably the best of these and has a decent array of options in its advanced search page.

altavista

- *uk.altavista.com*—While other engines now outperform AltaVista, the engine still provides accurate results and the range of features one would expect of a high-quality search engine. Past research has found AltaVista was very good at *indexing* academic websites.

Meta-Search engines. Meta-search engines do not search the web themselves. Instead, they collect the results from any number of regular search engines and (hopefully) eliminate duplicates and rank them relevantly. In the past these tools were of limited value, particularly because they could not take advantage of individual engines' features. The best meta-engines have now overcome this and have noticeably improved their relevancy ranking.

Recommended:

Pro Fusion

- www.profusion.com—Profusion is perhaps the best *meta-search engine*. Not only does it provide very good results, it has a considerable number of features and options including Boolean commands, vertical searching through subject categories and the ability to select from which engines it collates its results.

ixquick

- *Ixquick.com/uk*—Fast and accurate, Ixquick is a good choice, offering relevance scores and the ability to select and search most of the web's top engines.

(1) Other choices

Vivísimo

- *Vivisimo.com*—Vivisimo is a clustering engine. It works by retrieving results from the major *search engine*s and uses an algorithm to collate them conceptually and hierarchically. Results are post-coordinated, as search terms are not fitted into existing categories but grouped together after each individual search. Claiming this creates heuristic searches, it does allow for a distinct way for exploring results. Not recommended for specific focused searches, but do consider it for broad conceptual searches.

- *www.kartoo.com*—Kartoo provides good results but is also worth considering for the way it displays them. Rather than textual lists and/or categories, Kartoo collates results visually on a map that the searcher is then free to navigate. It also searches in French, Spanish, Portuguese and German.

★★★

Law of Diminishing Returns: Law stating that if one factor of production is increased while the others remain constant, the overall returns will relatively decrease after a certain point. (*The Columbia Encyclopaedia*, 6th edn, 2001)

Ultimately one does have to decide when to stop searching. This is particularly applicable to web searching, as you are searching what, in practice if not reality, is an infinite resource.

Notes on contributors

ALEX HAIG is the Information Scientist for NHS Education for Scotland and has been involved with the BEME Collaboration for the last three years. His experience with

systematic reviews has been focused on comprehensive evidence retrieval, data abstraction and critical appraisal. His other interests include knowledge management and information theory, evaluating healthcare information systems, and the organizational use of IT.

MARSHALL DOZIER is Reader Services Librarian (medicine) at the University of Edinburgh Library. Her current activities are largely in information skills training at all levels, from first-year medical students to practising health professionals. She also works with researchers carrying out systematic literature searches.

References

AVENELL, A. *et al.* (2001) Lessons for search strategies from a systematic review, in The Cochrane Library of nutritional supplementation trials in patients after hip fracture, *American Journal of Clinical Nutrition*, 73(3), pp. 505–509.

BARTLEBY [available at: http://www.bartleby.com/, accessed 16 June 2003].

BERECZKI, D. *et al.* (2000) A Hungarian example for handsearching specialized national health care journals of small countries for controlled trials. Is it worth the trouble?, *Health Libraries Review*, 17(3), pp. 144–147.

BEST EVIDENCE MEDICAL EDUCATION COLLABORATION [available at: http://www.bemecollaboration.org, accessed 4 April 2003].

BRETTLE, A. *et al.* (1998) Searching for information on outcomes: do you need to be comprehensive?, *Quality in Health Care*, 7(3), pp. 163–167.

EYSENBACH, G. *et al.* (2001) Evaluation of the usefulness of *Internet* searches to identify unpublished clinical trials for systematic reviews, *Medical Informatics*, 26(3), pp. 203–218.

GENERAL MEDICAL COUNCIL [available at: http://www.gmc-uk.org/med_ed/, accessed 4 April 2003].

HAIG, D.A.J. (2001) Evidence retrieval in medical education: obstructions and opportunities, paper presented at Association of Medical Education in Europe, Berlin, Germany.

HAIG, D.A.J. & COUPAR, H.A. (2001) Evidence retrieval in medical education: obstructions and opportunities, paper presented at Association of Medical Education in Europe, Berlin, Germany.

HARDEN, V., HARDEN, R.M., HAIG, A., MCMANUS, N.K. & LILLEY, P.M. (2002) *The Objective Structured Clinical Examination* (*OSCE*) *Review Project* (Dundee, Scotland, NHS Education for Scotland).

HELMER, D. (2002) Etext on Health Technology Assessment (HTA) Information Resources, National Library of Medicine, Bethesda, MD, USA.

HELMER, D. (1999) How do various fugitive literature searching methods impact the comprehensiveness of literature uncovered for systematic review?, paper presented at 4th International Conference on *Grey Literature*, Washington, USA, 4–5 October.

HOPEWELL, S. *et al.* (2000) A comparison of *handsearching* versus MEDLINE searching to identify reports of randomized controlled trials, *Statistics in Medicine*, 21(11), pp. 1625–1634.

MCAULEY, L. *et al.* (2000) Does the inclusion of *grey literature* influence estimates of intervention effectiveness reported in meta-analyses?, *Lancet*, 356(9237), pp. 1228–1231.

MINOZZI, S. *et al.* (2000), Searching for rehabilitation articles on Medline and EMBASE: an example with cross-over design, *Archives of Physical Medicine Rehabilitation*, 81(June), pp. 720–722.

MIT USENET QUOTES [available at: http://www.mit.edu/afs/sipb/user/marthag/text/h, accessed 16 June 2003].

NATIONAL LIBRARY OF MEDICINE [available at: http://www.nlm.nih.gov/MeSH/Mbrowser.html, accessed 4 April 2003].

OVID TECHNOLOGIES [available at: http://gateway.ovid.com, accessed 26 November 2002].

PANDOLFINI, C. *et al.* (2002) Follow up of quality of public oriented health information on the *world wide web*: systematic re-evaluation, *British Medical Journal*, 324(3773), pp. 582–583.

QUOTE WORLD [availabe at: http://www.quoteword.org/, accessed 16 June 2003].

THE QUOTATIONS PAGE [available at http://www.quotationspage.com/, accessed 16 June 2003].

WEBLUST QUOTE WORLD [available at: http://www.weblust.com/, accessed 16 June 2003].

WILSON, P. (2002) How to find the good and avoid the bad or ugly: a short guide to tools for rating quality of health information on the *internet*, *British Medical Journal*, 324(3773), pp. 598–602.

Glossary

Ancestry searching	The process of searching the bibliographies of relevant papers to discover references missed by other methods.
Bibliographic database	A database containing bibliographic information about publications, such as title, author and so on, but not usually the full text of publications.
Blogs	'We*b logs*'. Individuals' chronological personal writings and collected links.
Boolean algebra	AND, OR and NOT are combining commands that you can use to combine search terms and therefore refine your results. Boolean algebra is also known as 'Boolean logic' and the commands as 'logical operators'.
Broader	In the context of Subject headings, means a Subject heading that is more general in scope or at a higher level in the hierarchy of the Controlled thesaurus.
Concepts, search	The subject(s) of a search. Search concepts can be ideas, theories, outcomes, interventions or populations.
Controlled thesaurus [or index or vocabulary]	A controlled thesaurus is a list of standard subject terms from which indexers select subject headings to describe the content of articles or other publications in a consistent manner.
Database	A structured electronic information file, maintained to facilitate the retrieval of information.
Database fields	A logical defined unit of data, e.g. author, subject heading, year
Exhaustivity	The level of subject detail to which publications are described by indexers— ideally, a Controlled thesaurus should be comprehensive enough to allow to a well-rounded and precise description of publications.
Explode	Exploding a subject heading retrieves all its smaller subordinate terms. In general, it is always a good idea to Explode to ensure a search is comprehensive.
Focus	When you search a subject heading with Focus, that subject heading will be an essential concept to all results retrieved. This reduces the number of the results and may result in missed items, but increases the relevance.
Free-text search	A search that will look for a term or phrase in all specified or available fields, regardless of contextual meaning.

Granularity	The degree of description or detail; related to the scope of terms, particularly with regard to the level of detail indicated by Subject headings, where high granularity is equivalent to a very narrow subject heading and low granularity is equivalent to a broad subject heading.
Grey literature	That which is produced on all levels of government, academics, business and industry in print and electronic formats, but which is not controlled by commercial publishers (4th International Conference on Grey Literature, 1997).
Hand searching	Hand searching is literally the searching of print (or electronic) journals volume by volume, issue by issue, article by article.
Hit	A search result—referring to the records matching your search term.
Html	Hypertext mark-up language—most web pages are written with this.
Hyperlink	A link that takes your web browser to another location.
Inclusive	Encompassing all the concepts described. Synonymous with sensitive.
Indexing	The process of classifying an item (for example an article) with subject descriptors to make it easier to retrieve.
Internet	The worldwide network of computer networks.
MeSH	Medical Subject heading—the terms in the Controlled thesaurus created by the National Library of Medicine and used by indexers to describe the content of articles indexed in Index Medicus (and therefore PubMed and Medline).
Metadata	Descriptive information, such as bibliographic details, but also subject headings, publication type, language, etc. A bibliographic database like Medline is full of metadata, but not usually the full-text articles (that is, the actual data).
Methodology filter	A 'ready-made' search of terms that will retrieve specific types of reports, e.g. cohort studies, controlled trials, diagnostic use, etc. Filters are not usually subject specific—they are meant to be applicable to any subject search.
Narrower	In the context of Subject headings, means a Subject heading that is more specific in scope or at a subordinate level in the hierarchy of the Controlled thesaurus.
Polysemy	Polysemy occurs when a search term has multiple meanings, e.g. simulation, which could be about computer simulations, heuristics, the Markov Processes, role-playing, or patient simulators. Polysemy is one cause of irrelevant results from free-text searches.
Precision	In database searching, means the proportion of hits that are relevant. A search with high precision has few irrelevant results, and is also described as having high 'specificity'. A danger with highly precise search techniques is that some relevant information will be missed because it was not well described by the author or indexer.
Proximity operator	Proximity (or adjacency) operators are used when one wants to find two or more terms within a certain distance of one another and can usually only be employed in free-text searching. The operator varies between search interfaces e.g. 'adjx' (Ovid), 'nearx' (Silverplatter), where 'x' refers to the number of words apart the search terms may appear; or 'same' (Web of Science) which requires the words to be in the same sentence.
Qualitative research	Research involving detailed verbal descriptions of characteristics, cases, and settings. Qualitative research typically uses observation, interviewing and document review to collect data.
Quantitative research	Research that examines phenomena through the numerical representation of observations and statistical analysis.
Recall	In database searching, the proportion of relevant hits retrieved by a search—a search with high recall is called 'sensitive'.
Scope	The extent or range of a subject.
Search engine	An enormous database of Internet sites, usually compiled by robots.
Search filter	A search filter is a series of search commands designed to retrieve a particular type of result. Filters might be created to locate a particular type of study (e.g. controlled trials) or query specific (participants, educational aspects, outcomes) studies (e.g. undergraduate students, OSCEs, or competences).
Search interface	The 'front end' to a database—what you see when you are searching. This phrase can also imply the tools and facilities presented by the programming behind the interface.
Search set	A single search statement—for example, in a search history like the one built up in the sample Medline search in this guide, each line in the search history is a 'set'. Building up a history of search sets as in the sample is called 'set searching'.
Search strategy	Sometimes used interchangeably with 'search history' but can also refer to a larger systematic strategy for retrieving information that also includes handsearching journals and discovering unpublished research.

Selectivity — The choice of Subject headings made by the indexer to anticipate the concepts of greatest utility to the database user; also the degree to which a search is specified or narrowed.

Sensitivity — When referring to a literature search, means inclusive, so that you get more hits, and may get some irrelevant ones. Synonymous with high 'recall' in that a sensitive search aims to retrieve the highest proportion possible of relevant database records.

Specificity — When referring to a literature search, means exclusive, so that you get fewer hits to sift through, but may miss some relevant information. Synonymous with 'precision'.

Subheading — In the context of Subject headings, represents a popular facet of study related to a Subject heading

Subject heading — A term used to describe the content of a publication—usually derived from a Controlled thesaurus.

Synonymy — Where concepts can be named using different terms—this a problem with free-text searching, which requires you to think of as many terms as possible to capture the various words or phrases that authors could use in writing about the same concept. In theory, Subject headings should mitigate this issue.

Thesaurus — See Controlled thesaurus.

Truncation — Truncation (or wildcard searching) is the substitution of a character to retrieve variations in spelling and word ending. It cannot be used with the set terms of a controlled vocabulary, but is a powerful aid in improving the sensitivity of free-text searches. The truncation symbol varies from one search interface to another (see help or search tips to find out). Ovid uses the $: portfolio$—finds the singular and plural (portfolio or portfolios). evaluat$—finds multiple variations, evaluate, evaluation, evaluations, evaluative, evaluator, etc. Internal truncation is also possible. With Ovid, the # can be used to find spelling variations between versions of English: an#esthesia or "standardi#ed patient$".

URL — Uniform Resource Locator—a website's address: www.bmj.com.

Usenet — Archived electronic discussion lists on nearly any topic imaginable; first appeared on the Internet in 1981.

Web browser — A software application that lets you view web documents written in html—like Netscape or Internet Explorer

World wide web — One part of the Internet—other parts include email and discussion groups.