

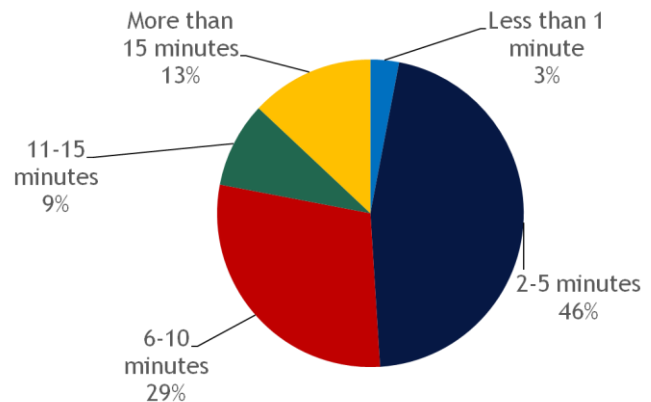


CLINICALLY RELEVANT SEARCH RESULTS IN OVIDMD – A WHITE PAPER

Understanding how a product or service approaches search, its intended use and its target population, is critical to measuring its performance. Google has ambitions to make all of the world’s knowledge searchableⁱ. At OvidMD, *we want to improve patient care by making it easier for physicians to answer clinical questions*. With our broad and trusted base of content, our understanding of clinicians, and our understanding of the challenging context in which they seek answers, we are in a unique position to do just that.

PHYSICIANS ARE LIMITED BY TIME

One of the key constraints in physician-based searching is the time that they are willing or able to allocate to researching any given problem. When asked about how much time they had available when they needed to find an answer to a clinical question, nearly half stated that they had less than 5 minutes. Those 5 minutes include going to the computer, opening the browser, navigating to the site, logging on, performing a search, looking at the results, reading the relevant material and making a decision.




Is it any wonder then, that clinicians may leave half of their own questions unanswered^{ii iii iv}? And, of course, that does not include the fundamentally *unidentified* questions – the questions they didn’t even know to ask.

Physicians repeatedly tell us that they have three primary times when they seek answers to clinical questions:

- 1) When they are with patients (possibly out of the room) and only have a few available minutes.
- 2) At lunch, or on a similar break, when they might have as many as 15 minutes.
- 3) In the evening when they are researching a problem or preparing for upcoming cases.

These are useful guidelines as they suggest that clinicians seek different levels of information depending on their context. A complex document (e.g., a systematic review of the literature) may provide an excellent answer to a clinical question, but due to the time constraint, may be of little or no value when the clinician is in a patient setting. But it is also important not to over-interpret the results. While the time constraint may be





overwhelmingly true for an overly busy resident, it is likely to be overwhelmingly false for a hematologist-oncologist (who may spend hours researching problems on behalf of every one of her patients).

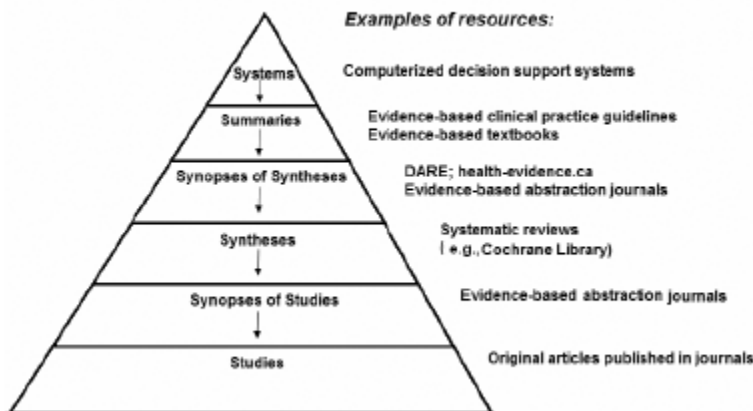
Recognizing this problem, medical content tends to be organized along lines that support clinician’s needs. One model of organization is the “evidence-based pyramid”.

THE EVIDENCE-BASED PYRAMID


In 2001, Brian Haynes MD, PhD, McMaster University described an approach to finding the current best evidence pertaining to clinical questions. The approach described a multi-layer “evidence pyramid” *with original “studies” at the base, “syntheses” (systematic reviews) of evidence just above the base, “synopses” of studies and syntheses next up, and the most evolved evidence-based information “systems” at the top.* Haynes advised searchers that they should begin their search at the top of the pyramid.^v

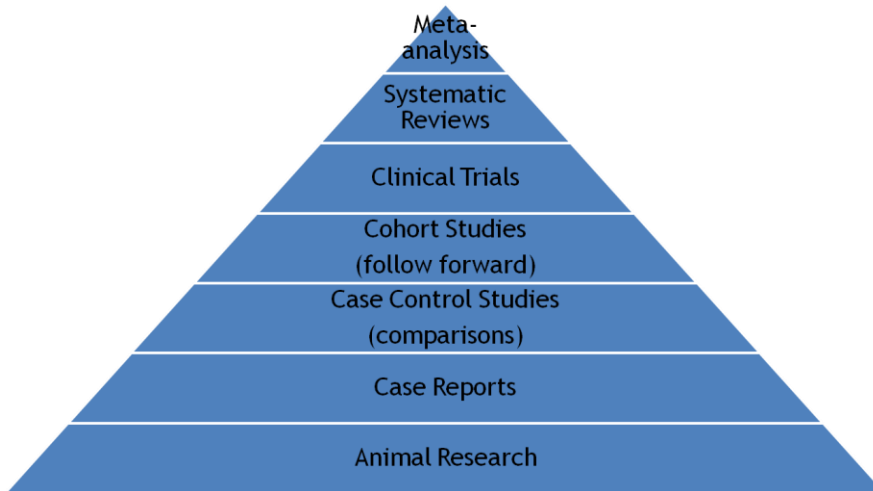
Over the years the model has been refined^{vi} and now includes 6 levels^{vii}. The so-called 6S model is represented below. As the model depicts, each higher level builds on the content below it. Starting with single studies at the bottom and ending with patient-specific, contextually aware content at the top. The articles are short, and worth a read (on OvidMD, no less).

OvidMD offers a potent content set representing each layer in the model.

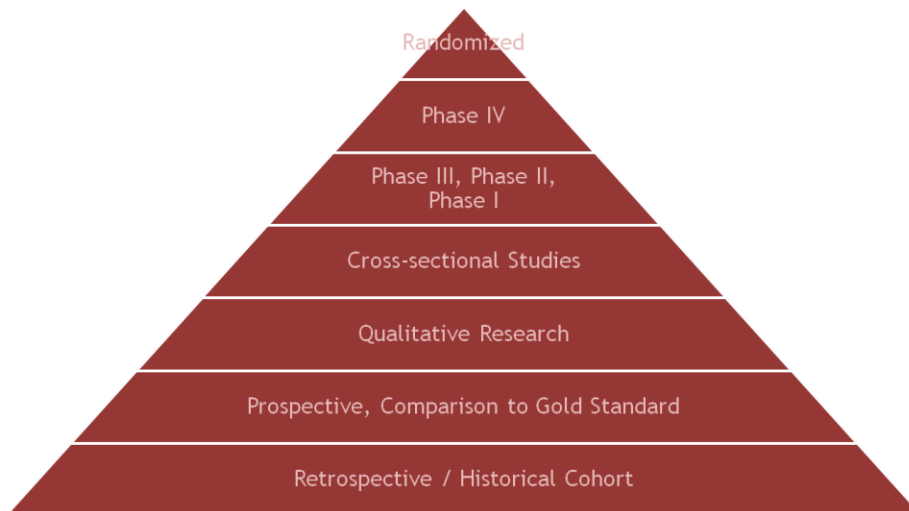


As Haynes points out, every model faces limitations. While OvidMD takes inspiration from the Haynes model, we heavily adapt it to our own context. Other models provide a more granular breakdown using content classifications that may be more familiar. The following model represents MEDLINE publication types.





Within any layer, additional layers of refinement might be considered. For instance, within Clinical Trials the model might look like this:



Our own application of the model is a hybrid of these approaches, combining “system level” content high in the results, and applying fine-levelled granularity to evidence sources, clinical trials, and other sources.

UNDERSTANDING THE EVIDENCE-BASED CONTENT ON OVIDMD

Some elements of OvidMD may not be entirely familiar. The following attempts to offer a brief summary of some selected content.



UPTODATE

- Includes treatment recommendations based on the best medical evidence
- Recommendations are kept current as new studies are released
- Over 4,400 physician authors and editors, all experts in their respective fields, write topic reviews that cover all of the major aspects of a particular condition, including symptoms, tests and diagnosis, and treatment options
- UpToDate is evidence-based and uses a literature-driven updating system; more than 440 journals are monitored by editors and authors^{viii}

EVIDENCE-BASED GUIDELINES

- covering the more than 300 diseases that account for up to 80% of inpatient hospital admissions
- a synthesis of best available evidence on;
 - diagnosis, management, and treatment of conditions that occur in primary care,
 - inpatient and emergency department settings.
- It is written by physicians for physicians, updated quarterly, and recommends a course of action for diagnosis, disease management and drug therapy

COCHRANE DATABASE OF SYSTEMATIC REVIEWS

- 'Gold Standard' for high-quality systematic reviews
- Full-text included in Cochrane Library
- Cochrane Reviews includes complete reviews and protocols (reviews that are still in progress).
- Cochrane Reviews abstracts are in PubMed

ACP JOURNAL CLUB

- American College of Physicians - American Society of Internal Medicine (ACP-ASIM)
- Internal medicine focus
- Abstracts to journal articles selected according to explicit criteria from major peer-reviewed medical journals

GOALS FOR OVIDMD SEARCH

At the outset we said: *we want to improve patient care by making it easier for physicians to answer clinical questions.* It is worth restating that as a more tacit set of objectives for search:

1. **Focus on patient-oriented, rather than disease-oriented content.** Patient-oriented content is concerned with topics like diagnosis, treatment and prognosis. Disease-oriented is concerned with the pathology of diseases. It is the difference between asking about the reduction in mortality associated with a potential hypertension drug, as opposed to the reduction in vasoconstrictors at the epithelial lining of peripheral vessels.
2. **Prioritize content that is highly trusted and highly accessible in the clinical setting.** It often takes only minutes to answers to clinical questions in resources like UpToDate. It can take quite a bit longer to consider a large number of papers describing original research.



3. **Search multiple resources simultaneously.** Remove the burden of looking in multiple resources with different interfaces and different strategies.
4. **Answer a large number of questions rapidly.** While a physician's search may be as simple as "asthma", the number of questions associated with asthma is quite large. The search results themselves are a valuable aid in understanding the clinical landscape.

The balance of this document will focus on the specific steps we take to achieve those goals.

SEARCH 101

Search can seem pretty mysterious and even contradictory, but it doesn't have to be. OvidMD considers dozens of factors and countless variables in performing search. Sometimes these factors are at odds. What follows is a brief description of some of the key relationships. As F. Scott Fitzgerald once famously wrote, *"the test of a first-rate intelligence is the ability to hold two opposed ideas in the mind at the same time, and still retain the ability to function."*^{ix}

CONTROLLED VOCABULARIES, TAXONOMIES AND ONTOLOGIES

An essential resource for any medical search engine is a robust medical vocabulary system.

"Controlled vocabularies", such as the National Library of Medicine's Medical Subject Headings (MeSH) encourage the consistent use of terminology in a given setting. In an ideal world, we might all refer to "stroke" as a "cerebrovascular accident". In fact, we don't.

To ensure that "stroke" returns all relevant documents, OvidMD automatically applies synonyms from Wolters Kluwer's proprietary controlled vocabulary called Lexi. Lexi, in turn, incorporates dozens of other controlled vocabularies. Many of these are compiled and cross-linked in the NLM's Unified Medical Language System.

A few of the vocabularies that are used derive Lexi include:

- MeSH
- Snomed-CT
- ICD-9
- CPT
- DSM

RECALL VS. RELEVANCE



Recall is the ability of a search engine to find every document that matches the query that was submitted. That seems trivial at first, but consider whether or not a search engine should return documents for “myocardial infarction” when a user searches for “heart attack,” or if its abbreviation (MI) happens to be the same as for the state of Michigan.

Relevance is the ability of a search engine to rank documents in a manner that is meaningful to the user. Relevance is subjective, but not arbitrary. By identifying an intended audience, a system can make reasonable assumptions about what a user may be looking for. Whereas a generalized search engine like Google has to satisfy someone looking for vacation destinations as well as someone searching for life saving information, we’re only concerned about the latter. Our relevance model was developed through extensive user research and continues to evolve based on user feedback. We’ll expand on the relevance model later in this document.

INDEX-TIME PERFORMANCE VS. RUNTIME PERFORMANCE


In general, search makes explicit trade-offs in performance. The less work you have to do at runtime (when the query is entered) the faster search performs. Of course, that work still needs to be done. At OvidMD we load as much work as possible into the indexing process.

Indexing is the process of interpreting documents and building a searchable index. (A very simple index would simply be a list of unique words in all documents along with a “pointer” to each occurrence of the word in each document. No matter how much content you include, the index itself would grow to be only roughly the equivalent of a dictionary in terms of size and complexity.) OvidMD parses the various document formats and organizes the index into a number of fields. The fields, and sometimes the content, get normalized across the data. For instance, we have separate fields for titles, subtitles, authors, publication dates, abstracts, indexing keywords, document bodies, image captions, ISBNs, ISSNs, publication types, document types, etc. Not all sources have elements that clearly map to these categories, and not all documents treat the data the same way or even consistently. The syntactic (format) and semantic (meaning) normalization of publication dates, for instance, is a significant challenge.

OvidMD also creates compound fields as a way of grouping elements for weighting when relevance gets calculated. By assigning weights to fields at indexing time, a significant part of relevance ranking can be moved from runtime to index-time, resulting in faster performance for our users.

SIMPLE SEARCH INTERFACE VS COMPLEX OPTIONS

This isn’t as much of an either-or question as other issues in search, but users consistently told us that they wanted a “dead simple” search interface: a Google-like search box, and nothing more. No advanced search, no complicated query language. Experience with similar medical search solutions suggests that fewer than 2% of users searches use “advanced features”.





SIMPLE QUERIES VS. COMPLEX QUERIES

Some systems do well with simple queries, but fall down as queries get more complex. We want users to be able to enter anything from simple keyword searches to complex, natural language questions. OvidMD uses advanced linguistic techniques to analyze queries, to automatically adapt to different word forms, to discard irrelevant expressions (e.g., where can I find information about...), to increase the weight of words based on their uniqueness in the set of all documents, and of course, to apply synonyms.

DO WHAT I MEAN VS. DON'T BE TOO CLEVER


The ideal search solution would understand users and “do what they mean”. OvidMD attempts to do what the user means when it presents results based on the clinical utility of documents, when it provides drug “best bets”, when it automatically expands synonyms to improve recall, and offers faceted navigation to restrict results to certain content subsets or date ranges. These are carefully made choices on behalf of the user. At the same time, we try to guard against being “too clever”, or thinking we know more than we do about what the clinician needs.

OVIDMD SEARCH, STEP BY STEP

RECALL: FINDING ALL THE DOCUMENTS THAT MATCH THE QUERY

OvidMD searches for the exact words as entered by a user, identifying phrases in the input and in the content. It also searches for the morphological variants of the words^x, and for any synonymous terms (e.g., heart attack == myocardial infarction). The synonyms come from LEXI, our own controlled vocabulary. LEXI contains many vocabularies that are part of the UMLS (e.g., MeSH, Snomed).

Some words may be discarded from the user’s query. These are called “anti-phrases” and they refer to the sometimes irrelevant portions of a user’s query. For instance, in the query “where can I find information about hypertension”, it’s pretty clear that the user is not actually looking for the words “where can I find information about”. There are well over 100,000 such phrases that searchers regularly enter. Whether or not the phrase is discarded or diminished is a complex decision, but it accounts for the possibility that the terms actually appear in the corpus (collection of documents). Even if the phrase is not completely discarded, it is likely to be diminished via a calculation called inverse document frequency (see below).





RELEVANCE: PUTTING THE DOCUMENTS IN ORDER

So far, the focus of this discussion has been on “recall” – identifying the documents that match the user’s query, but not doing too much about relevancy ranking. Relevancy is calculated based on a proprietary algorithm in the search engine, and adjusted by our own weights, as described below.

The initial relevance is based on several factors:


- The number of hits in a document.
- The fields where we find hits (e.g., title, abstract, indexing data, body).
- The relative value of any term in describing the document when compared to all documents in the corpus (inverse document frequency).
- The document’s publication date.
- The position of hits relative to each other.

Once the initial ranking has been completed, OvidMD performs another level of ranking:

- **TITLES:** OvidMD gives a sizeable boost to documents with ALL the query terms in the TITLE.
- **KEYWORDS:** OvidMD gives a small boost to keywords (content-specific indexing data such as MeSH).
- **NON-CLINICAL / NON-ACTIONABLE DOCUMENTS:** OvidMD looks for certain characteristics that indicate that the document won’t be of value to a clinician. Characteristics like animal studies, non-medical journals (e.g., theriogenology, nursing), old content (prior to 1996), and things like biographies, portraits and the like are given a **negative** boost. They typically won’t show up very high in the list.
- **HIGH VALUE DOCUMENTS:** OvidMD looks for certain characteristics like “meta-analysis” or “systematic review” that make documents especially relevant.
- **SPECIAL CASES:** The system also does some special case handling. For instance, terms like “pseudotumor cerebri” would normally show up in hypertension searches (intercranial hypertension), even though users probably mean “essential hypertension”. It’s a small number of exceptions, but we give those types of documents a negative boost.
- **PUBLICATION DATE:** We do some special boosting based on the age of the document. New articles get boosted pretty significantly. Old articles get a negative boost. For an old document to show up high in the results, it really has to be a pretty good fit otherwise.



- CONTENT TYPE: Finally, we return to the evidence pyramid presented earlier in this document. **These are the heaviest weights in the system.**
 1. UpToDate.
 2. Evidence Based Guidelines.
 3. Evidence Based Resources (ACPJ, EBMR).
 4. Journals and MEDLINE have a bit of overlap. Within Journals and MEDLINE, we favor journals with a clinical focus, especially those with high impact factor scores (e.g., NEJM). This isn't exact, and it isn't comprehensive, but it does help. Because the content is complete, and because our full text journals tend to be amongst the best available, we have a bias towards showing hits in the full text.
 5. Books, especially where the title of the document exactly matches the user's query.
 6. Drugs get boosted very high, if the user's query is specifically for a drug.
 7. The system does not attempt to do any other ordering for National Guidelines, Patient Handouts, or Drugs.



SPECIAL HANDLING

Some queries express a clear information need. For instance, a search for a drug will bring the drug monograph to the top of the results. Similarly, a search for a book or journal title will bring that item to the top of the results.

ELECTRONIC MEDICAL RECORD AND SEARCH – POSSIBILITIES FOR THE FUTURE

OvidMD does not currently integrate directly with Electronic Medical Records.

CONCLUSION

The stakes are high. A lack of trusted information at the point of care has in the past and present lead to enormous patient losses.

- In the 1840s, before the germ theory of disease, the Austrian physician Ignaz Semmelweis recognized that puerperal (childbirth) fever was contagious. He further found that when physicians would simply by wash their hands in chlorinated lime between seeing patients, mortality rates fell from 12% to 2% in his clinic. But in part because Semmelweis could not offer an explanation, his findings were neither trusted nor



adopted. Semmelweis would later be fired and die in a mental hospital. Thousands of women would needlessly die as a result.

- Some of the exact same issues were revisited in 2005 in Don Berwick's 100,000 Lives Campaign^{xi}. Berwick established a program designed to save 100,000 lives in 18 months by getting some 3,000 hospitals to adopt six changes that, based on *previously published statistical evidence*, would reduce mortality. Amongst the changes --- hand washing to reduce central line infections by more than 90%. The program succeeded, preventing an estimated 122,000 hospital deaths over the 18 month period.
- During World War I, Sir Almroth Wright demonstrated that by cleaning wounds and leaving them open, leukocytes could effectively fight infection. But the information wasn't trusted, busy surgeons were unable to adapt to it, and so procedures from previous conflicts prevailed. Wounds were closed up tight, sealing the fate of the wounded. During The Great War, many more soldiers died of infections from their wounds than from the wounds themselves^{xii}. Again, that despite the fact that an alternative was known, and largely because it was not known to be trustworthy by those treating the patients.
- The examples are endless, and many medical myths still prevail long after contrary published evidence is available^{xiii} (e.g., the slow acceptance of helicobacter pylori as the primary cause of peptic ulcers).
- Indeed, a 2000 study by Balas and Boran at the University of Missouri demonstrated that it took, on average, 17 years for evidence to be implemented for 50% of patients^{xiv}. Put in stark, and heart-breaking terms, although evidence indicated that children should not be placed face-down in cribs as early as 1970, physicians continued to offer this incorrect advice until 1988 (NB: 18 years). Had this evidence been known and trusted, some 60,000 child deaths might have been prevented.^{xv}

It is not an exaggeration to say that you, and OvidMD, are playing an important role in changing those statistics and in improving patient care.

ⁱ Google's mission is to organize the world's information and make it universally accessible and useful.

<http://www.google.com/about/corporate/company/>

ⁱⁱ Gorman PN, Helfand M. Information seeking in primary care: how physicians choose which clinical questions to pursue and which to leave unanswered. *Med Decis Making*. 1995 Apr-Jun;15(2):113-9.

ⁱⁱⁱ A. R. Barrie and A. M. Ward. Questioning behaviour in general practice: a pragmatic study. *BMJ*. 1997 December 6; 315(7121): 1512-1515

^{iv} John W. Ely, MD, MSPH, Jerome A. Osheroff, MD, Saverio M. Maviglia, MD, and Marcy E. Rosenbaum, PhD. Patient-care questions that physicians are unable to answer. *J Am Med Inform Assoc*. 2007 Jul-Aug; 14(4): 407-414.



^v Haynes, R. B. (2001). Of studies, syntheses, synopses, and systems: the "4S" evolution of services for finding current best evidence. *ACP J Club* 134(2): A11-3.

^{vi} Haynes, R. B. (2006). Of studies, syntheses, synopses, summaries, and systems: the "5S" evolution of information services for evidence-based health care decisions. *ACP J Club* 145(3): A8.

^{vii} DiCenso, Alba, Liz Bayley and R. Brian Haynes. "Accessing reappraised evidence: fine-tuning the 5S model into a 6S model," [Editorial]. *ACP Journal Club*, Sept. 15, 2009. 151(3): JC 3-2 – JC 3-3.

^{viii} <http://www.uptodate.com/home/about/journals.html>

^{ix} F. Scott Fitzgerald. *The Crack-up*. 1936

^x Examples of morphological variants include:

- Plural forms of nouns: child, children
- Verb tenses: bleeds, bleeding / cough, coughing
- Adjectives/adverbs: thermal, thermally

^{xi} Overview of the 100,000 Lives Campaign.

<http://www.ihl.org/IHL/Programs/Campaign/100kCampaignOverviewArchive.htm>

^{xii} Thomas Hager. *The Demon Under the Microscope*. Three Rivers Press. 2006.

^{xiii} Ian Ayers. *Super Crunchers. How Should Physicians Treat Evidence-Based Medicine?* Bantam. 2007.

^{xiv} Balas, E.A. and Boran, S.A., *Managing clinical knowledge for healthcare improvement in Yearbook of Medical Informatics*. 2000.

^{xv} Gilbert R, Salanti G, Harden M, See S. Infant sleeping position and the sudden infant death syndrome: systematic review of observational studies and historical review of recommendations from 1940 to 2002. *Int J Epidemiol*. 2005 Aug;34(4):874-87. Epub 2005 Apr 20.

ADDITIONAL RESOURCES

How to ask clinical questions you can answer.

http://www.columbia.edu/itc/hs/medical/clerkships/peds/Student_Information/Reference_Materials/Sackett_1.pdf