
Chapter One

Introduction to Web Search Engines

1.1 A SHORT HISTORY OF INFORMATION RETRIEVAL

Today we have museums for everything—the museum of baseball, of baseball players, of crazed fans of baseball players, museums for world wars, national battles, legal fights, and family feuds. While there’s no shortage of museums, we have yet to find a museum dedicated to this book’s field, a museum of information retrieval and its history. Of course, there are related museums, such as the Library Museum in Borås, Sweden, but none concentrating on information retrieval. **Information retrieval**¹ is the process of searching within a document collection for a particular information need (called a **query**). Although dominated by recent events following the invention of the computer, information retrieval actually has a long and glorious tradition. To honor that tradition, we propose the creation of a museum dedicated to its history. Like all museums, our museum of information retrieval contains some very interesting artifacts. Join us for a brief tour.

The earliest document collections were recorded on the painted walls of caves. A cave dweller interested in searching a collection of cave paintings to answer a particular information query had to travel by foot, and stand, staring in front of each painting. Unfortunately, it’s hard to collect an artifact without being gruesome, so let’s fast forward a bit.

Before the invention of paper, ancient Romans and Greeks recorded information on papyrus rolls. Some papyrus artifacts from ancient Rome had tags attached to the rolls. These tags were an ancient form of today’s Post-it Note, and make an excellent addition to our museum. A tag contained a short summary of the rolled document, and was attached in order to save readers from unnecessarily unraveling a long irrelevant document. These abstracts also appeared in oral form. At the start of Greek plays in the fifth century B.C., the chorus recited an abstract of the ensuing action. While no actual classification scheme has survived from the artifacts of Greek and Roman libraries, we do know that another elementary information retrieval tool, the table of contents, first appeared in Greek scrolls from the second century B.C. Books were not invented until centuries later, when necessity required an alternative writing material. As the story goes, the Library of Pergamum (in what is now Turkey) threatened to overtake the celebrated Library of Alexandria as the best library in the world, claiming the largest collection of papyrus rolls. As a result, the Egyptians ceased the supply of papyrus to Pergamum, so the Pergamenians invented an alternative writing material, parchment, which is made from thin layers of animal skin. (In fact, the root of the word *parchment* comes from the word *Pergamum*.) Unlike papyrus,

¹The boldface terms that appear throughout the book are also listed and defined in the Glossary, which begins on page 201.

parchment did not roll easily, so scribes folded several sheets of parchment and sewed them into books. These books outlasted scrolls and were easier to use. Parchment books soon replaced the papyrus rolls.

The heights of writing, knowledge, and documentation of the Greek and Roman periods were contrasted with their lack during the Dark and Middle Ages. Precious few documents were produced during this time. Instead, most information was recorded orally. Document collections were recorded in the memory of a village's best storyteller. Oral traditions carried in poems, songs, and prayers were passed from one generation to the next. One of the most legendary and lengthy tales is *Beowulf*, an epic about the adventures of a sixth-century Scandinavian warrior. The tale is believed to have originated in the seventh century and been passed from generation to generation through song. Minstrels often took poetic license, altering and adding verses as the centuries passed. An inquisitive child wishing to hear stories about the monster Grendel waited patiently while the master storyteller searched his memory to find just the right part of the story. Thus, the result of the child's search for information was biased by the wisdom and judgement of the intermediary storyteller. Fortunately, the invention of paper, the best writing medium yet, superior to even parchment, brought renewed acceleration to the written record of information and collections of documents. In fact, *Beowulf* passed from oral to written form around A.D. 1000, a date over which scholars still debate. Later, monks, the possessors of treasured reading and writing skills, sat in scriptoriums working as scribes from sunrise to sunset. The scribes' works were placed in medieval libraries, which initially were so small that they had no need for classification systems. Eventually the collections grew, and it became common practice to divide the holdings into three groups: theological works, classical authors of antiquity, and contemporary authors on the seven arts. Lists of holdings and tables of contents from classical books make nice museum artifacts from the medieval period.

Other document collections sprung up in a variety of fields. This dramatically accelerated with the re-invention of the printing press by Johann Gutenberg in 1450. The wealthy proudly boasted of their private libraries, and public libraries were instituted in America in the 1700s at the prompting of Benjamin Franklin. As library collections grew and became publicly accessible, the desire for focused search became more acute. Hierarchical classification systems were used to group documents on like subjects together. The first use of a hierarchical organization system is attributed to the Roman author Valerius Maximus, who used it in A.D. 30 to organize the topics in his book, *Factorum ac dictorum memorabilium libri IX* (Nine Books of Memorable Deeds and Sayings). Despite these rudimentary organization systems, word of mouth and the advice of a librarian were the best means of obtaining accurate quality information for a search. Of course, document collections and their organization expanded beyond the limits of even the best librarian's memory. More orderly ways of maintaining records of a collection's holdings were devised. Notable artifacts that belong in our information retrieval museum are a few lists of individual library holdings, sorted by title and also author, as well as examples of the Dewey decimal system (1872), the card catalog (early 1900s), microfilm (1930s), and the MARC (MACHINE Readable Cataloging) system (1960s).

These inventions were progress, yet still search was not completely in the hands of the information seeker. It took the invention of the digital computer (1940s and 1950s) and the subsequent inventions of computerized search systems to move toward that goal. The

first computerized search systems used special syntax to automatically retrieve book and article information related to a user's query. Unfortunately, the cumbersome syntax kept search largely in the domain of librarians trained on the systems. An early representative of computerized search such as the Cornell SMART system (1960s) [146] deserves a place in our museum of information retrieval.

In 1989 the storage, access, and searching of document collections was revolutionized by an invention named the World Wide Web by its founder Tim Berners-Lee [79]. Of course, our museum must include artifacts from this revolution such as a webpage, some HTML, and a hyperlink or two. The invention of linked document collections was truly original at this time, despite the fact that Vannevar Bush, once Director of the Office of Scientific Research and Development, foreshadowed its coming in his famous 1945 essay, "As We May Think" [43]. In that essay, he describes the memex, a futuristic machine (with shocking similarity to today's PC and Web) that mirrors the cognitive processes of humans by leaving "trails of association" throughout document collections. Four decades of progress later, remnants of Bush's memex formed the skeleton of Berners-Lee's Web. A drawing of the memex (Figure 1.1) by a graphic artist and approved by Bush was included in *LIFE* magazine's 1945 publishing of Bush's prophetic article.

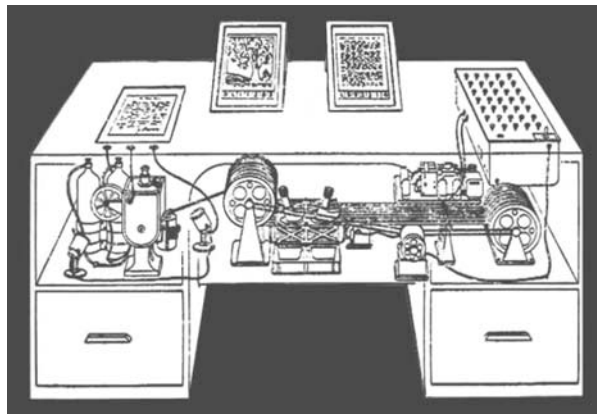


Figure 1.1 Drawing of Vannevar Bush's memex appearing in *LIFE*. Original caption read: "Memex in the form of a desk would instantly bring files and material on any subject to the operator's fingertips. Slanting translucent screens supermicrofilm filed by code numbers. At left is a mechanism which automatically photographs longhand notes, pictures, and letters, then files them in the desk for future reference."

The World Wide Web became the ultimate signal of the dominance of the Information Age and the death of the Industrial Age. Yet despite the revolution in information storage and access ushered in by the Web, users initiating web searches found themselves floundering. They were looking for the proverbial needle in an enormous, ever-growing information haystack. In fact, users felt much like the men in Jorge Luis Borges' 1941 short story [35], "The Library of Babel", which describes an imaginary, infinite library.

When it was proclaimed that the Library contained all books, the first impression was one of extravagant happiness. All men felt themselves to be the masters of an intact and secret treasure. There was no personal or world problem whose eloquent solution did not exist in some hexagon.

. . . As was natural, this inordinate hope was followed by an excessive depression. The certitude that some shelf in some hexagon held precious books and that these precious books were inaccessible seemed almost intolerable.

Much of the information in the Library of the Web, like that in the fictitious Library of Babel, remained inaccessible. In fact, early web search engines did little to ease user frustration; search could be conducted by sorting through hierarchies of topics on Yahoo, or by sifting through the many (often thousands of) webpages returned by the search engine, clicking on pages to personally determine which were most relevant to the query. Some users resorted to the earliest search techniques used by ancient queriers—word of mouth and expert advice. They learned about valuable websites from friends and linked to sites recommended by colleagues who had already put in hours of search effort.

All this changed in 1998 when **link analysis** hit the information retrieval scene [40, 106]. The most successful search engines began using link analysis, a technique that exploited the additional information inherent in the hyperlink structure of the Web, to improve the quality of search results. Web search improved dramatically, and web searchers religiously used and promoted their favorite engines like Google and AltaVista. In fact, in 2004 many web surfers freely admit their obsession with, dependence on, and addiction to today's search engines. Below we include the comments [117] of a few Google fans to convey the joy caused by the increased accessibility of the Library of the Web made possible by the link analysis engines. Incidentally, in May 2004 Google held the largest share of the search market with 37% of searchers using Google, followed by 27% using the Yahoo conglomerate, which includes AltaVista, AlltheWeb, and Overture.²

- “It’s not my homepage, but it might as well be. I use it to ego-surf. I use it to read the news. Anytime I want to find out anything, I use it.”—Matt Groening, creator and executive producer, *The Simpsons*
- “I can’t imagine life without Google News. Thousands of sources from around the world ensure anyone with an Internet connection can stay informed. The diversity of viewpoints available is staggering.”—Michael Powell, chair, Federal Communications Commission
- “Google is my rapid-response research assistant. On the run-up to a deadline, I may use it to check the spelling of a foreign name, to acquire an image of a particular piece of military hardware, to find the exact quote of a public figure, check a stat, translate a phrase, or research the background of a particular corporation. It’s the Swiss Army knife of information retrieval.”—Garry Trudeau, cartoonist and creator, *Doonesbury*

Nearly all major search engines now combine link analysis scores, similar to those used by Google, with more traditional information retrieval scores. In this book, we record the history of one aspect of *web* information retrieval. That aspect is the link analysis or *ranking* algorithms underlying several of today’s most popular and successful search

²These market share statistics were compiled by comScore, a company that counted the number of searches done by U.S. surfers in May 2004 using the major search engines. See the article at <http://searchenginewatch.com/reports/article.php/2156431>.

engines, including Google and Teoma. Incidentally, we'll add the **PageRank** link analysis algorithm [40] used by Google (see Chapters 4-10) and the **HITS** algorithm [106] used by Teoma (see Chapter 11) to our museum of information retrieval.

1.2 AN OVERVIEW OF TRADITIONAL INFORMATION RETRIEVAL

To set the stage for the exciting developments in link analysis to come in later chapters, we begin our story by distinguishing **web information retrieval** from **traditional information retrieval**. Web information retrieval is search within the world's largest and linked document collection, whereas traditional information retrieval is search within smaller, more controlled, nonlinked collections. The traditional nonlinked collections existed before the birth of the Web and still exist today. Searching within a university library's collection of books or within a professor's reserve of slides for an art history course—these are examples of traditional information retrieval.

These document collections are nonlinked, mostly static, and are organized and categorized by specialists such as librarians and journal editors. These documents are stored in physical form as books, journals, and artwork as well as electronically on microfiche, CDs, and webpages. However, the mechanisms for searching for items in the collections are now almost all computerized. These computerized mechanisms are referred to as search engines, virtual machines created by software that enables them to sort through virtual file folders to find relevant documents. There are three basic computer-aided techniques for searching traditional information retrieval collections: Boolean models, vector space models, and probabilistic models [14]. These search models, which were developed in the 1960s, have had decades to grow, mesh, and morph into new search models. In fact, as of June 2000, there were at least 3,500 different search engines (including the newer web engines) [37], which means that there are possibly 3,500 different search techniques. Nevertheless, since most search engines rely on one or more of the three basic models, we describe these in turn.

1.2.1 Boolean Search Engines

The **Boolean model** of information retrieval, one of the earliest and simplest retrieval methods, uses the notion of exact matching to match documents to a user query. Its more refined descendants are still used by most libraries. The adjective Boolean refers to the use of Boolean algebra, whereby words are logically combined with the Boolean operators AND, OR, and NOT. For example, the Boolean AND of two logical statements x and y means that both x AND y must be satisfied, while the Boolean OR of these two statements means that at least one of these statements must be satisfied. Any number of logical statements can be combined using the three Boolean operators. The Boolean model of information retrieval operates by considering which keywords are present or absent in a document. Thus, a document is judged as relevant or irrelevant; there is no concept of a partial match between documents and queries. This can lead to poor performance [14]. More advanced fuzzy set theoretic techniques try to remedy this black-white Boolean logic by introducing shades of gray. For example, a title search for `car AND maintenance` on a Boolean engine causes the virtual machine to return all documents that use both words in the title. A relevant document entitled "Automobile Maintenance" will not be returned. Fuzzy Boolean engines use fuzzy logic to categorize this document as somewhat relevant and return it to the user.

The car maintenance query example introduces the main drawbacks of Boolean search engines; they fall prey to two of the most common information retrieval problems, **synonymy** and **polysemy**. Synonymy refers to multiple words having the same meaning, such as *car* and *automobile*. A standard Boolean engine cannot return semantically related documents whose keywords were not included in the original query. Polysemy refers to words with multiple meanings. For example, when a user types *bank* as their query, does he or she mean a financial center, a slope on a hill, a shot in pool, or a collection of objects [24]? The problem of polysemy can cause many documents that are irrelevant to the user's actual intended query meaning to be retrieved. Many Boolean search engines also require that the user be familiar with Boolean operators and the engine's specialized syntax. For example, to find information about the phrase *iron curtain*, many engines require quotation marks around the phrase, which tell the search engine that the entire phrase should be searched as if it were just one keyword. A user who forgets this syntax requirement would be surprised to find retrieved documents about interior decorating and mining for iron ore.

Nevertheless, variants of the Boolean model do form the basis for many search engines. There are several reasons for their prevalence. First, creating and programming a Boolean engine is straightforward. Second, queries can be processed quickly; a quick scan through the keyword files for the documents can be executed in parallel. Third, Boolean models scale well to very large document collections. Accommodating a growing collection is easy. The programming remains simple; merely the storage and parallel processing capabilities need to grow. References [14, 75, 107] all contain chapters with excellent introductions to the Boolean model and its extensions.

1.2.2 Vector Space Model Search Engines

Another information retrieval technique uses the **vector space model** [147], developed by Gerard Salton in the early 1960s, to sidestep some of the information retrieval problems mentioned above. Vector space models transform textual data into *numeric vectors* and *matrices*, then employ *matrix analysis*³ techniques to discover key features and connections in the document collection. Some advanced vector space models address the common text analysis problems of synonymy and polysemy. Advanced vector space models, such as LSI [64] (Latent Semantic Indexing), can access the hidden semantic structure in a document collection. For example, an LSI engine processing the query *car* will return documents whose keywords are related semantically (in meaning), e.g., *automobile*. This ability to reveal hidden semantic meanings makes vector space models, such as LSI, very powerful information retrieval tools.

Two additional advantages of the vector space model are **relevance scoring** and **relevance feedback**. The vector space model allows documents to partially match a query by assigning each document a number between 0 and 1, which can be interpreted as the likelihood of relevance to the query. The group of retrieved documents can then be sorted by degree of relevancy, a luxury not possible with the simple Boolean model. Thus, vector space models return documents in an ordered list, sorted according to a relevance score. The first document returned is judged to be most relevant to the user's query.

³Mathematical terms are defined in Chapter 15, the Mathematics Chapter, and are italicized throughout.

Some vector space search engines report the relevance score as a relevancy percentage. For example, a 97% next to a document means that the document is judged as 97% relevant to the user's query. (See the Federal Communications Commission's search engine, <http://www.fcc.gov/searchtools.html>, which is powered by *Inktomi*, once known to use the vector space model. Enter a query such as `taxes` and notice the relevancy score reported on the right side.) Relevance feedback, the other advantage of the vector space model, is an information retrieval tuning technique that is a natural addition to the vector space model. Relevance feedback allows the user to select a subset of the retrieved documents that are useful. The query is then resubmitted with this additional relevance feedback information, and a revised set of generally more useful documents is retrieved.

A drawback of the vector space model is its computational expense. At query time, distance measures (also known as similarity measures) must be computed between each document and the query. And advanced models, such as LSI, require an expensive singular value decomposition [82, 127] of a large matrix that numerically represents the entire document collection. As the collection grows, the expense of this matrix decomposition becomes prohibitive. This computational expense also exposes another drawback—vector space models do not scale well. Their success is limited to small document collections.

Understanding Search Engines

The informative little book by Michael Berry and Murray Browne, *Understanding Search Engines: Mathematical Modeling and Text Retrieval* [23], provides an excellent explanation of vector space models, especially LSI, and contains several examples and sample code. Our mathematical readers will enjoy this book and its application of linear algebra algorithms in the context of traditional information retrieval.

1.2.3 Probabilistic Model Search Engines

Probabilistic models attempt to estimate the probability that the user will find a particular document relevant. Retrieved documents are ranked by their odds of relevance (the ratio of the probability that the document is relevant to the query divided by the probability that the document is not relevant to the query). The probabilistic model operates recursively and requires that the underlying algorithm guess at initial parameters then iteratively tries to improve this initial guess to obtain a final ranking of relevancy probabilities.

Unfortunately, probabilistic models can be very hard to build and program. Their complexity grows quickly, deterring many researchers and limiting their scalability. Probabilistic models also require several unrealistic simplifying assumptions, such as independence between terms as well as documents. Of course, the independence assumption is restrictive in most cases. For instance, in this document the most likely word to follow `information` is the word `retrieval`, but the independence assumption judges each word as equally likely to follow the word `information`. On the other hand, the probabilistic framework can naturally accommodate a priori preferences, and thus, these models do offer promise of tailoring search results to the preferences of individual users. For example, a

user's query history can be incorporated into the probabilistic model's initial guess, which generates better query results than a democratic guess.

1.2.4 Meta-search Engines

There's actually a fourth model for traditional search engines, meta-search engines, which combines the three classic models. **Meta-search engines** are based on the principle that while one search engine is good, two (or more) are better. One search engine may be great at a certain task, while a second search engine is better at another task. Thus, meta-search engines such as *Copernic* (www.copernic.com) and *SurfWax* (www.surfwax.com) were created to simultaneously exploit the best features of many individual search engines. Meta-search engines send the query to several search engines at once and return the results from all of the search engines in one long unified list. Some meta-search engines also include subject-specific search engines, which can be helpful when searching within one particular discipline. For example, *Monster* (www.monster.com) is an employment search engine.

1.2.5 Comparing Search Engines

Annual information retrieval conferences, such as TREC [3], SIGIR, CIR [22] (for traditional information retrieval), and WWW [4] (for web information retrieval), are used to compare the various information retrieval models underlying search engines and help the field progress toward better, more efficient search engines. The two most common ratings used to differentiate the various search techniques are precision and recall. **Precision** is the ratio of the number of relevant documents retrieved to the total number of documents retrieved. **Recall** is the ratio of the number of relevant documents retrieved to the total number of relevant documents in the collection. The higher the precision and recall, the better the search engine is. Of course, search engines are tested on document collections with known parameters. For example, the commonly used test collection Medlars [6], containing 5,831 keywords and 1,033 documents, has been examined so often that its properties are well known. For instance, there are exactly 24 documents relevant to the phrase *neoplasm immunology*. Thus, the denominator of the recall ratio for a user query on *neoplasm immunology* is 24. If only 10 documents were retrieved by a search engine for this query, then a recall of $10/24 = .41\bar{6}$ is reported. Recall and precision are information retrieval-specific performance measures, but, of course, when evaluating any computer system, time and space are always performance issues. All else held constant, quick, memory-efficient search engines are preferred to slower, memory-inefficient engines. A search engine with fabulous recall and precision is useless if it requires 30 minutes to perform one query or stores the data on 75 supercomputers. Some other performance measures take a user-centered viewpoint and are aimed at assessing user satisfaction and frustration with the information system. A book by Robert Korfhage, *Information Storage and Retrieval* [107], discusses these and several other measures for comparing search engines. Excellent texts for information retrieval are [14, 75, 163].

1.3 WEB INFORMATION RETRIEVAL

1.3.1 The Challenges of Web Search

Tim Berners-Lee and his World Wide Web entered the information retrieval world in 1989 [79]. This event caused a branch that focused specifically on search within this new document collection to break away from traditional information retrieval. This branch is called **web information retrieval**. Many web search engines are built on the techniques of traditional search engines, but they differ in many important ways. We list the properties that make the Web such a unique document collection. The Web is:

- huge,
- dynamic,
- self-organized, and
- hyperlinked.

The Web is indeed huge! In fact, it's so big that it's hard to get an accurate count of its size. By January 2004, it was estimated that the Web contained over 10 billion pages, with an average page size of 500KB [5]. With a world population of about 6.4 billion, that's almost 2 pages for each inhabitant. The early exponential growth of the Web has slowed recently, but it is still the largest document collection in existence. The Berkeley information project, "How Much Information," estimates that the amount of information on the Web is about 20 times the size of the entire Library of Congress print collection [5]. Bigger still, a company called BrightPlanet sells access to the so-called Deep Web, which they estimate to contain over 92,000TB of data spread over 550 billion pages [1]. BrightPlanet defines the Deep Web as the hundreds of thousands of publicly accessible databases that create a collection over 500 times larger than the Surface Web. Deep webpages can not be found by casual, routine surfing. Surfers must request information from a particular database, at which point, the relevant pages are served to the user dynamically within a matter of seconds. As a result, search engines cannot easily find these dynamic pages since they do not exist before or after the query. However, Yahoo appears to be the first search engine aiming to index parts of the Deep Web.

The Web is dynamic! Contrast this with traditional document collections which can be considered static in two senses. First, once a document is added to a traditional collection, it does not change. The books sitting on a bookshelf are well behaved. They don't change their content by themselves, but webpages do, very frequently. A study by Junghoo Cho and Hector Garcia-Molina [52] in 2000 reported that 40% of all webpages in their dataset changed within a week, and 23% of the .com pages changed daily. In a much more extensive and recent study, the results of Fetterly et al. [74] concur. About 35% of all webpages changed over the course of their study, and also pages that were larger in size changed more often and more extensively than their smaller counterparts. Second, for the most part, the size of a traditional document collection is relatively static. It is true that abstracts are added to MEDLINE each year, but how many? Hundreds, maybe thousands. These are minuscule additions by Web proportions. Billions of pages are added to the Web each year. The dynamics of the Web make it tough to compute relevancy scores for queries when the collection is a moving, evolving target.

The Web is self-organized! Traditional document collections are usually collected and categorized by trained (and often highly paid) specialists. However, on the Web, anyone can post a webpage and link away at will. There are no standards and no gatekeepers policing content, structure, and format. The data are volatile; there are rapid updates, broken links, and file disappearances. One 2002 U.S. study reporting on “link rot” suggested that up to 50% of URLs cited in articles in two information technology journals were inaccessible within four years [1]. The data is heterogeneous, existing in multiple formats, languages, and alphabets. And often this volatile, heterogeneous data is posted multiple times. In addition, there is no editorial review process, which means errors, falsehoods, and invalid statements abound. Further, this self-organization opens the door for sneaky **spammers** who capitalize on the mercantile potential offered by the Web. *Spammers* was the name originally given to those who send mass advertising emails. With one click of the send button, spammers can send their advertising message to thousands of potential customers in a matter of seconds. With web search and online retailing, this name was broadened to include those using deceptive webpage creation techniques to rank highly in web search listings for particular queries. Spammers resorted to using minuscule text font, hidden text (white on a white background), and misleading **metatag** descriptions to fool early web search engines (like those using the Boolean technique of traditional information retrieval). The self-organization of the Web also means that webpages are created for a variety of different purposes. Some pages are aimed at surfers who are shopping, others at surfers who are researching. In fact, search engines must be able to answer many types of queries, such as transactional queries, navigational queries, and informational queries. All these features of the Web combine to make the job for web search engines Herculean.

Ah, but the Web is hyperlinked! This linking feature, the foundation of Vannevar Bush’s memex, is the saving grace for web search engines. Hyperlinks make the new national pastime of surfing possible. But much more importantly, they make focused, effective searching a reality. This book is about ways that web search engines exploit the additional information available in the Web’s sprawling link structure to improve the quality of their search results. Consequently, we focus on just one aspect of the web information retrieval process, but one we believe is the most exciting and important. However, the advantages resulting from the link structure of the Web did not come without negative side effects. The most interesting side effects concern those sneaky spammers. Spammers soon caught wind of the link analysis employed by major search engines, and immediately set to work on link spamming. Link spammers carefully craft hyperlinking strategies in the hope of increasing traffic to their pages. This has created an entertaining game of cat and mouse between the search engines and the spammers, which many, the authors included, enjoy spectating. See the asides on pages 43 and 52.

An additional information retrieval challenge for any document collection, but especially pertinent to the Web, concerns precision. Although the amount of accessible information continues to grow, a user’s ability to look at documents does not. Users rarely look beyond the first 10 or 20 documents retrieved [94]. This user impatience means that search engine precision must increase just as rapidly as the number of documents is increasing. Another dilemma unique to web search engines concerns their performance measurements and comparison. While traditional search engines are compared by running tests on familiar, well studied, controlled collections, this is not realistic for web engines. Even small web collections are too large for researchers to catalog, count, and create estimates of the precision and recall numerators and denominators for dozens of queries. Comparing two

search engines is usually done with user satisfaction studies and market share measures in addition to the baseline comparison measures of speed and storage requirements.

1.3.2 Elements of the Web Search Process

This last section of the introductory chapter describes the basic elements of the web information retrieval process. Their relationship to one another is shown in Figure 1.2. Our purpose in describing the many elements of the search process is twofold: first, it helps emphasize the focus of this book, which is the ranking part of the search process, and second, it shows how the ranking process fits into the grand scheme of search. Chapters 3-12 are devoted to the shaded parts of Figure 1.2, while all other parts are discussed briefly in Chapter 2.

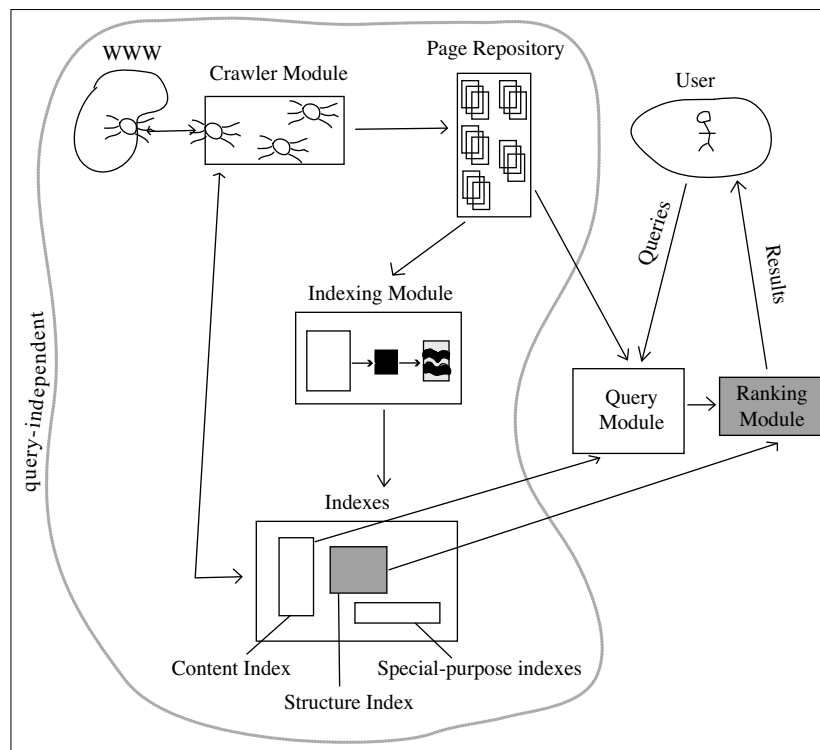


Figure 1.2 Elements of a search engine

- **Crawler Module.** The Web's self-organization means that, in contrast to traditional document collections, there is no central collection and categorization organization. Traditional document collections live in physical warehouses, such as the college's library or the local art museum, where they are categorized and filed. On the other hand, the web document collection lives in a cyber warehouse, a virtual entity that is not limited by geographical constraints and can grow without limit. However, this geographic freedom brings one unfortunate side effect. Search engines must

do the data collection and categorization tasks on their own. As a result, all web search engines have a crawler module. This module contains the software that collects and categorizes the web's documents. The crawling software creates virtual robots, called **spiders**, that constantly scour the Web gathering new information and webpages and returning to store them in a central repository.

- **Page Repository.** The spiders return with new webpages, which are temporarily stored as full, complete webpages in the page repository. The new pages remain in the repository until they are sent to the indexing module, where their vital information is stripped to create a compressed version of the page. Popular pages that are repeatedly used to serve queries are stored here longer, perhaps indefinitely.
- **Indexing Module.** The indexing module takes each new uncompressed page and extracts only the vital descriptors, creating a compressed description of the page that is stored in various indexes. The indexing module is like a black box function that takes the uncompressed page as input and outputs a "Cliffnotes" version of the page. The uncompressed page is then tossed out or, if deemed popular, returned to the page repository.
- **Indexes.** The indexes hold the valuable compressed information for each webpage. This book describes three types of indexes. The first is called the **content index**. Here the content, such as keyword, title, and anchor text for each webpage, is stored in a compressed form using an **inverted file** structure. Chapter 2 describes the inverted file in detail. Further valuable information regarding the hyperlink structure of pages in the search engine's index is gleaned during the indexing phase. This link information is stored in compressed form in the **structure index**. The crawler module sometimes accesses the structure index to find uncrawled pages. **Special-purpose indexes** are the final type of index. For example, indexes such as the image index and pdf index hold information that is useful for particular query tasks.

The four modules above (crawler, page repository, indexers, indexes) and their corresponding data files exist and operate independent of users and their queries. Spiders are constantly crawling the Web, bringing back new and updated pages to be indexed and stored. In Figure 1.2 these modules are circled and labeled as **query-independent**. Unlike the preceding modules, the query module is **query-dependent** and is initiated when a user enters a query, to which the search engine must respond in **real-time**.

- **Query Module.** The query module converts a user's natural language query into a language that the search system can understand (usually numbers), and consults the various indexes in order to answer the query. For example, the query module consults the content index and its inverted file to find which pages use the query terms. These pages are called the relevant pages. Then the query module passes the set of relevant pages to the ranking module.
- **Ranking Module.** The ranking module takes the set of relevant pages and ranks them according to some criterion. The outcome is an ordered list of webpages such that the pages near the top of the list are most likely to be what the user desires. The ranking module is perhaps the most important component of the search process because the output of the query module often results in too many (thousands

of) relevant pages that the user must sort through. The ordered list filters the less relevant pages to the bottom, making the list of pages more manageable for the user. (In contrast, the similarity measures of traditional information retrieval often do not filter out enough irrelevant pages.) Actually, this ranking which carries valuable, discriminatory power is arrived at by combining two scores, the **content score** and the **popularity score**. Many rules are used to give each relevant page a relevancy or content score. For example, many web engines give pages using the query word in the title or description a higher content score than pages using the query word in the body of the page [39]. The popularity score, which is the focus of this book, is determined from an analysis of the Web's hyperlink structure. The content score is combined with the popularity score to determine an **overall score** for each relevant page [30]. The set of relevant pages resulting from the query module is then presented to the user in order of their overall scores.

Chapter 2 gives an introduction to all components of the web search process, except the ranking component. The ranking component, specifically the popularity score, is the subject of this book. Chapters 3 through 12 provide a comprehensive treatment of the ranking problem and its suggested solutions. Each chapter progresses in depth and mathematical content.