



Search Engines for the World Wide Web: A Comparative Study and Evaluation Methodology

Heting Chu

Palmer School of Library & Information Science, Long Island University
Brookville, New York

Marilyn Rosenthal

Library Reference Department, Long Island University
Brookville, New York

ABSTRACT

Three Web search engines, namely, Alta Vista, Excite, and Lycos, were compared and evaluated in terms of their search capabilities (e.g., Boolean logic, truncation, field search, word and phrase search) and retrieval performances (i.e., precision and response time) using sample queries drawn from real reference questions. Recall, the other evaluation criterion of information retrieval, is deliberately omitted from this study because it is impossible to assume how many relevant items there are for a particular query in the huge and ever changing Web system. The authors of this study found that Alta Vista outperformed Excite and Lycos in both search facilities and retrieval performance although Lycos had the largest coverage of Web resources among the three Web search engines examined. As a result of this research, we also proposed a methodology for evaluating other Web search engines not included in the current study.

INTRODUCTION

Though a latecomer in the Internet family, the World Wide Web (WWW or the Web) has rapidly gained popularity and become the second most widely used application of the Internet [1]. The publicity WWW has gained is so great that many people naively equate WWW with the Internet. The friendly user interface and the hypermedia features of WWW have been attracting a significant number of users as well as information providers. As a result, the web has become a sea of all kinds of data, making any query into the huge information reservoir extremely difficult.

In order to overcome this difficulty in retrieving information from WWW, more than two dozen companies and institutions quickly developed various search aids [2] such

as Lycos and Excite. However, since there are usually only one or two search aids for other Internet applications (e.g., Archie for FTP, and Veronica for Gopher), why have at least two dozen search engines been developed for the Web so far? The sheer number invites research. For instance, what features do various Web search engines offer? How do they differ from one another in performance? Is there a single Web search engine that out-performs all others in information retrieval? The current study attempts to seek answers to those questions.

RELATED STUDIES

Web search engines did not come into existence until 1994. The literature covering them has an even shorter time span. In fact, a survey of the literature indicates that the number of evaluation studies done on Web search engines is small, and the majority of those publications (e.g., Shirky, 1995; Taubes, 1995; Wildstrom, 1995) are descriptive in nature.

Eventually, people went a step further by starting to evaluate Web search engines in addition to describing them. Notess examined Lycos, WebCrawler, World-Wide Web Worm, Harvest Broker, CUI, and CUSI in one article (1995a) and InfoSeek in another (1995b). Based on online documentation provided by those Web search engines and personal usage, Notess recommended that "for single keyword searches of a large database, use Lycos". "For multiword searches with an AND, try WebCrawler". "For a time-consuming comprehensive search, use CUSI". In addition, Notess also compared InfoSeek with Lycos and WebCrawler in terms of coverage, precision, and currency.

In a more recent publication, Courtois, Baer, and Stark (1995) evaluated the performances of about 10 different Web search aids including CUI, Harvest, Lycos, Open Text, World-Wide Web Worm, and Yahoo. Using 3 sample search questions along with other information available about the search engines, the authors concluded that, among other things, Open Text was the best at the time of their study "with its flexible, powerful search interface and quick response". They also concluded that "For novices, WebCrawler offers the easiest interface". In a different study, Scoville (1996) surveyed a wide range of Web search engines, and suggested that Excite, InfoSeek, and Lycos should be added to one's list of favorites because they can retrieve "accurate results from easy-to-use interfaces".

Leighton (1995) did a study of Web search engines for course work, actually employing the evaluation criterion of precision. The findings were not submitted to a journal for publication because of the fast changing nature of the search engines. Leighton evaluated InfoSeek, Lycos, WebCrawler and World-Wide Web Worm using 8 reference questions from a university library as search queries. The author found that "Lycos and the free part of InfoSeek have about the same precision with Lycos just a nose ahead" while WebCrawler gave "surprisingly bad precision". "WWWorm was good enough that usually retrieved at least one or two hits" for

the given queries with high precision.

Kimmel (1996) examined World Wide Web Worm, Lycos, WebCrawler, Open Text, Jumpstation II, AliWeb, and Harvest based on documentation provided by the search engines along with a couple of single word test searches (e.g., pollution, ebola). The author's focus was, like many other publications, on describing the features of these various search engines even though the number of hits produced by test searches were also listed. The author, in summary, indicated that "Of the robot-generated databases presented here, Lycos appears to be the strongest system overall".

c|net, a company specialized in evaluating online products and services, distributed a comparative study of 19 Web search engines on its Web site (Leonard, 1996). The search engines were tested on their accuracy of results, ease of use and provision of advanced options using 15 queries specifically composed for the evaluation. Most of the queries resemble reference questions asked in public libraries. According to the two feature tables generated by the evaluation, Alta Vista seems to be the best choice among individual search engines, while All-in-One Search Page and the Internet Sleuth achieved the highest ranking for meta- or unified search engines.

The reported findings obviously do not appear to agree with one another. The methodologies and evaluation criteria used by those studies differed as well. Can a feasible methodology be developed to help Web users select a search engine, out of the great number of choices, that is most appropriate to their specific search needs? The authors of this study are trying to do so by first evaluating the searching capabilities and performance of selected Web search engines currently available.

SCOPE AND OBJECTIVE OF THE STUDY

As indicated previously, Web search aids are variously referred to as catalogs, directories, indexes, search engines, or Web databases (Courtois, Baer, & Stark, 1995). Since the current study focuses on the search capability and performance of Web search aids, we decide to use the phrase "search engine" as the formal expression. On the other hand, according to our understanding, a search engine should at least allow users to compose their own search queries rather than simply follow pre-specified search paths or hierarchy as in the case of certain catalogs. Thus, due to its origin and its browsing component, Yahoo is not included in our study despite the fact that it is one of the most widely used search aids for Web resources.

We are also aware that many search engines index not only Web information but also resources stored on other Internet applications such as discussion groups and Gopher. But, we chose to consider only Web databases to be consistent with the objective of our study. Moreover, we did not cover unified Web search engines such as CUSI (Configurable Unified Search Index, <http://pubweb.nexor.co.uk/public/cusi/doc/list.html>) since search tools of that kind do not provide

anything new except putting together existing individual ones. Although some of them (e.g., MetaCrawler) have added such new features as removing duplicates, their searching mechanism remains the same.

In addition, most of the Web search engines are available to users free of charge. It seems that these free services will continue to be available to the Internet community in the foreseeable future. Given the fact that users will naturally choose search engines that can be accessed at no cost to them, our study excludes fee-based Web search services such as InfoSeek even though we understand that it may indeed perform well in retrieving Web resources.

During the process of selecting Web search engines to be evaluated, we paid particular attention to covering those representing diversity so that our choices would comprise different types of Web search engines. We applied the same criterion in choosing sample queries for our performance evaluation. Sample search queries were drawn from real reference questions.

With selected search engines, we compared their search capabilities such as Boolean logic, truncation, field searching, and word/phrase searching. Furthermore, we also evaluated the performance of the selected search engines with respect to precision and response time. Recall, the other commonly used evaluation criterion for information retrieval performance, was deliberately omitted from this study because it is impossible to determine how many relevant items there are for a particular query in the huge and ever-changing Web system. The ultimate goal of this study is, as stated earlier, to develop a feasible methodology for evaluating all Web search engines.

SELECTED SEARCH ENGINES AND THEIR FEATURES

Three search engines, namely Alta Vista, Excite [3], and Lycos, were examined based on the selection criteria discussed above. Out of the three selected search engines, Lycos has the longest history, while Alta Vista the shortest. The following summary information is mainly derived from online documentation for the three search engines, plus some publications and personal experience in using them.

Alta Vista (<http://www.altavista.digital.com>)

Alta Vista began to be developed in the Summer of 1995 at Digital's Research Laboratories in Palo Alto, California, and was formally delivered to the Web on December 15, 1995.

It indexes the full text of over 16,000,000 Web pages (by January 1996) with unspecified update frequencies. According to its documentation, Alta Vista can fetch 2.5 million pages a day following the Robots Exclusion Standard, and index 1 GB of

text per hour. Alta Vista supports Boolean searching, term as well as phrase searching (i.e., proximity searching with the NEAR operator), field searching (eg, title:steelhead; url:home.html), right-hand truncation with some restriction, and case-sensitive searching if only the first letter of a word is capitalized.

Alta Vista provides three display options: compact, standard, and detailed although the latter two are the same. The display order or relevancy ranking of search results is determined by the location (e.g., in title or the body of text) of matching words, occurrence frequencies of matching words, and distance (i.e., how many words apart) between the matching words. However, only the first few words of a document found are displayed, which may limit users' ability to judge its relevancy without referring to the full version of the document. In addition, general search terms such as "computer" and "analysis" are automatically ignored in Alta Vista.

Excite (<http://www.excite.com>)

Excite was developed by Architext Software, a company initially based in a garage. It claims 1.5 million fully indexed Web pages (Scoville, 1996), and its index is updated approximately once a week.

Excite allows keyword searching as well as concept searching since it is able to determine related concepts from document collections, eliminating the need for external manually-defined representations such as thesauri. An example of concept searching given by Excite is that a search query about "intellectual property rights" will retrieve all documents about the topic even if terms such as "software piracy" or "copyright law" rather than the actual matching words appear in the document. In other words, the search engine itself handles synonyms and related terms, taking the burden of vocabulary control off users' shoulders. As for keyword search, query terms are both AND'ed and OR'ed in each search, but a higher weight is given to results with terms AND'ed. However, Excite does not support at present other advanced search options than those being described already.

Equipped with automatic abstracting capability, Excite is able to generate an abstract for each of the Web pages it indexes, which is a very unique and fine feature that many of its counterparts do not have. But, there are no different formats for displaying search results. In addition, its online documentation appears somewhat unorganized.

Lycos (<http://www.lycos.com>)

Lycos, representing the first 5 letters of the Latin name for wolf spider, was originally designed at Carnegie Mellon University. It was later sold to America Online and became Lycos, Inc. at which Michael Mauldin, the person who has overseen Lycos' development is still a full-time employee. Although commercialized, Lycos continues to provide free services to the Internet community.

By the end of January 1996, Lycos has indexed over 95% (ca. 19 million unique URLs including FTP and Gopher) of Web resources, making it the largest Web search engine in its family. Nevertheless, it does not index the full text of a Web page. Rather, it only extracts the title and a portion of a document (e.g., the smaller of the first 20 lines or 20% of the document). This practice has been singled out by Lycos' competitors as its most salient weakness. Around 50,000 documents are added, deleted, or updated in the Lycos index everyday.

Lycos supports Boolean logic, and furthermore, it incorporates that feature in such a way that the users do not have to type the Boolean operators when conducting a search. For example, one only needs to select the search option "Match all terms (AND)" to use the AND operator. Another search feature Lycos provides is to match query terms against Web documents at 5 different levels, namely, Loose match, Fair match, Good match, Close match, Strong match. Nevertheless, no specific explanation is given as to how the different levels of match are determined. Truncation is automatically done in Lycos during a search, which may result in some unwanted search outcome. Phrase search is not supported by Lycos so any queries with phrases cannot be appropriately executed.

On the other hand, Lycos implements a wide variety of display options. Users are given the choices of viewing 10, 20, 30, or 40 research results a time. In addition, each search result can be displayed using the summary, standard, or detailed format. The detailed format corresponds with the long abstracts Lycos prepares, which include URL, title, outline, keys, abstract, description, date, and other related information. The summary format contains what Lycos' short abstracts have: URL and descriptions. In terms of coverage, the standard format lies somewhere between the summary and detailed formats. The online documentation available at Lycos' Web site describes the composition of each output segment (e.g., outline and keys) in detail.

In summary, the three different Web search engines show diversity in their search capabilities, user interface, and quality of documentation. The next section of this paper will discuss the performance evaluation of the selected Web search engines.

SAMPLE QUERIES AND THE TEST ENVIRONMENT

Nine out of the following ten search queries were extracted from real reference questions handled by the librarians at Long Island University. These questions are intended to be used for testing various features each search engine claims to have, as well as to represent different levels of searching complexity.

Reference Questions

1. volunteerism in society
2. classical Greek philosophy
3. memory and neurobiology
4. sexual differences and mathematical ability
5. psychological analysis of contemporary British artist Francis Bacon
6. violence among athletes
7. computers and learning disabilities
8. NAFTA
9. plagiarism
10. Long Island University

As can be seen, some of the questions consist of single words (e.g., #8 & #9), and some constitute phrases (e.g., #4 & #7). Some queries require the use of Boolean logic (e.g., #1 & #6) while others do not (e.g., #10). Most of the questions are about general themes (e.g., #1), but some deal with specific topics (e.g., #5). In addition, truncation and case-sensitivity can be tested in several cases (e.g., #2 & #4). Since Excite allows concept searching, questions such as #4 and #6 are used to test, for example, whether Web documents using synonyms for "athletes" (e.g., sportsman) would be retrieved while searching the term "athletes". Question #10 was the only query composed by the authors in order to test the field search capability (e.g., title search) of the selected search engines.

Search Queries

According to the specific syntax of each search engine selected, three separate search queries were constructed for every reference question listed above. The terms and characters listed after the name of each search engine are the queries actually typed in during the searches. It became obvious from these queries that Excite and Lycos have very similar syntaxes.

- #1 Alta Vista: volunteerism +society
Excite: volunteerism society
Lycos: volunteerism society
- #2 Alta Vista: "classical Greek philosophy"
Excite: classical Greek philosophy
Lycos: classical Greek philosophy
- #3 Alta Vista: memory +neurobiology
Excite: memory neurobiology
Lycos: memory neurobiology
- #4 Alta Vista: "sexual difference*" + "mathematical ability"

- Excite: sexual differences mathematical ability
Lycos: sexual differences mathematical ability
- #5 Alta Vista: "psychological analysis" +"British artist" +"Francis Bacon"
Excite: British artist Francis Bacon
Lycos: British artist Francis Bacon
- #6 Alta Vista: violence +athlete*
Excite: violence athletes
Lycos: violence athletes
- #7 Alta Vista: computers +"learning disabilit*"
Excite: computers learning disabilities
Lycos: computer learning disabilities
- #8 Alta Vista: NAFTA
Excite: NAFTA
Lycos: NAFTA
- #9 Alta Vista: plagiarism
Excite: plagiarism
Lycos: plagiarism
- #10 Alta Vista: title:"Long Island University"
Excite: Long Island University
Lycos: Long Island University

The Test Environment

Both Netscape and Lynx were used as Web browsers for the searches since Netscape supports the full capability of a Web browser while Lynx can generate search results for downloading without HTML tags for easier reading.

Whenever there are different search options available (e.g., Alta Vista's simple search & advanced search), the simple mode is used in order to relate the findings of this study to those with little searching background. In the case of Lycos, the "Loose match" and "Match all terms (AND)" options were selected for all the queries. The latter decision was based on the rationale that none of the ten questions entails the use of other listed choices such as "Match any term (OR)" and "Match 2 terms". As for the display options, the most detailed one available is always favored since this option would provide us with more information for evaluation.

Due to the time factor, we only examined up to 10 Web records [4] for each query. As all the selected search engines display results in descending order of relevance calculated one way or another, we believe that this should not critically affect the validity of our study.

PERFORMANCE EVALUATION

Evaluation Criteria

Lancaster and Fayen (1973) once listed 6 criteria for assessing the performance of information retrieval systems. They are: 1) Coverage, 2) Recall, 3) Precision, 4) Response time, 5) User effort, and 6) Form of output. Although the criteria were set up more than two decades ago and a great deal has been done to reduce user effort (e.g., design friendly user interface) in using the system, they still seem quite applicable to evaluating information retrieval systems today. In addition, the Web is essentially an information storage and retrieval system characterized by its enormous size, hypermedia structure, and distributed architecture. In the previous section, we already examined to some extent the coverage, user effort and output form of the three selected Web search engines. As planned, we also intentionally excluded the recall criterion from our evaluation. Therefore, what remains to be measured are their precision and response time in information retrieval.

Response Time

Surprisingly, the response time for the three Web search engines did not vary greatly, and the average waiting time (i.e., between issuing a search command and displaying the first batch of search results on the screen) for every search was in the range of 1-5 seconds. Also, no significant difference was found in response time between peak or off hours of Web usage, as Netscape was used for all the searches on a Sunday afternoon whereas searches using Lynx were done during a Monday morning. Overall, Alta Vista appeared to be the quickest in responding, followed by Lycos and Excite.

Precision of Search Results

Relevance of retrieved Web records was determined separately by both authors on the basis of the up to 10 Web records we downloaded for each query [5]. We did not try to read the full text Web documents by following the links provided because of time considerations and reliability of the Web linkages. In order to delineate the overall performance of each Web engine we examined, we not only computed precision scores for each individual query, but also calculated average precision among all 10 searches for every search engine included in the study, as was done in the TREC series (Harman, 1995). In addition, we tabulated the mean precision for each sample query so that some light can be shed on the suitability of using Web search engines for certain questions. (See Table 1)

While Alta Vista and Excite always retrieved at least 10 Web records for each query, Lycos sometimes could not find anything at all on some of the topics (e.g., #4 & #5). We have omitted listing the total number of Web records retrieved by the three search engines since Excite, unlike Alta Vista and Lycos, does not provide that figure up front, and one must trace down along the path of "Next Documents" in order to get the final count.

Table 1 Precision (P) Chart for Three Web Search Engines

Sample Query	Alta Vista		Excite		Lycos		Mean P
	Sum P/#	P	Sum P/#	P	Sum P/#	P	
#1	8.5/10	0.85	8.0/10	0.8	2.0/3	0.67	0.77
#2	8.0/10	0.8	6.0/10	0.6	4.0/4	1.0	0.8
#3	10.0/10	1.0	7.5/10	0.75	9.0/10	0.9	0.88
#4	3.5/10	0.35	1.0/10	0.1	0	0	0.2
#5	1.0/10	0.1	3.0/10	0.3	0	0	0.13
#6	9.0/10	0.9	0.5/10	0.05	4.5/10/	0.45	0.47
#7	7.5/10	0.75	3.0/10	0.3	0	0	0.35
#8	10.0/10	1.0	10.0/10	1.0	1.0/10	1.0	1.0
#9	10.0/10	1.0	6.0/10	0.6	8.0/8 [6]	1.0	0.87
#10	10.0/10	1.0	0/10	0	1.0/2	0.5	0.5
Mean P	N/A	0.78	N/A	0.45	N/A	0.55	0.59

In comparison, Alta Vista obtained the highest precision score (0.78) among the three Web search engines while Lycos had 0.55 and Excite got the lowest one (0.45). There are however very few duplicates in some 250 Web records we perused, which suggests that the constructions (e.g., spiders and indexes) of each Web search engine are diversified enough that they consequently represent different portions of the entire magnificent Web system.

On the other hand, it is evident from Table 1 that some search queries (e.g., #4 & #5) are not really suitable for searching on the Web because of the complicated nature of those questions and the fact that sophisticated facilities such as proximity searching are yet to be developed for Web search engines.

Evaluation Specific to Alta Vista, Excite, & Lycos

Search capabilities

Of the three search engines evaluated, only Alta Vista truly supports phrase (the NEAR proximity) searching, an important feature that any search tools should have. The other two experienced great difficulty (e.g., no hits or irrelevant hits) in retrieving queries containing phrases. The lack of phrase searching capability directly contributed to the relatively low precision scores Excite and Lycos obtained.

While Excite does not support truncation, Lycos automatically truncates every possible query term (e.g., "violence" was truncated as "violence", "violenced", "violenceo", and "violences"), which inevitably brings a lot of noise into search

results. Although case-sensitivity is not a significant feature for search tools, only Alta Vista offers it on a limited scale.

We deliberately tested the concept searching capability of Excite, and the results were satisfactory. On the other hand, the field searching feature provided by Alta Vista is quite unique among Web search engines, which partially explains why it did the best in Query #10 -- Long Island University, while Excite could only retrieve results about Long Island and nothing about Long Island University in its first 10 records displayed.

Output option

As described previously, both Alta Vista and Lycos offered different output options for Web records whose contents are exclusively "extracted" from original Web documents. So it is not unusual to encounter cases where a so-called abstract or summary suddenly ends in the middle of a sentence. In comparison, Excite presents readable as well as meaningful summaries for the Web documents retrieved thanks to its automatic abstracting technique.

Among the three search engines compared, Lycos presented the greatest amount of information in a Web record including URL, outline, keys, abstract, description, and other related data. Nevertheless, a closer look at the contents reveals that a lot of the information displayed is either redundant or of little practical value. For instance, an outline always repeats the title of a Web document, and an abstract consists of the first 20 lines or 20% of the document, whichever is smaller.

Documentation and interface

Interface and documentation are the two factors that affect users' efforts in learning and using a Web search engine. Generally speaking, Lycos has good interface and documentation while the interface for Alta Vista (especially the one for displaying results) and the documentation for Excite need to be improved.

In a word, Alta Vista should be the first choice, among the three Web search engines evaluated, for users who expect search results of high precision. Otherwise, the selection of a search engine for Web navigation can be based on personal preference for the documentation, interface, or other features specific to it.

AN EVALUATION METHODOLOGY FOR WEB SEARCH ENGINES

Based on our knowledge and experience gained from the current study, we believe that one needs to consider the following aspects when evaluating a Web search engine.

Composition of Web Indexes

Whenever a Web search request is issued, it is the web index generated by Web robots or spiders, not the web pages themselves, that has been used for retrieving information. Therefore, the composition of Web indexes affects the performance of a Web search engine. There are three components that the authors would like to inspect regarding the makeup of a Web index, namely, coverage, update frequency and the portions of Web pages indexed (e.g., titles plus the first several lines, or the entire Web page). We understand that the magnitude of all three components depends largely on the power and sophistication of the hardware and software that make the Web index or database. On the other hand, larger coverage, frequent updates and fulltext indexing do not necessarily mean better Web search engines in other measurements.

Search Capability

A competent Web search engine must include the fundamental search facilities that Internet users are familiar with, which include Boolean logic, phrase searching, truncation, and limiting facilities (e.g., limit by field). Because the searching capabilities of a Web search engine ultimately determine its performance, absence of these basic functions will severely handicap the search tool, as we have seen in the case of Excite and Lycos which do not support phrase searching.

Retrieval Performance

Retrieval performance is traditionally evaluated on three parameters: precision, recall and response time. While the three variables can all be quantitatively measured, extra caution should be exercised when one judges the relevance of retrieved items and estimates the total number of documents relevant to a specific topic in the Web system.

Output Option

This evaluation component should be examined from two perspectives. One is the number of output options a Web search engine offers, whereas the other deals with the actual content of the output. Sometimes, one search engine may appear quite impressive in one aspect, but in reality it cannot satisfy its users because of its weakness in the other facet of this evaluation criterion. The output content, to a certain degree, is decided by the way a search engine is constructed. That is, does it purely extract or actually index the Web system when building its database? The answer should be apparent.

User Effort

User effort refers to documentation and interface in this study. Well-prepared

documentation and a user-friendly interface play a notable role in users' selection of Web search engines. Since there are more than two dozen of them available, the attractiveness of each Web search engine is expressed, to its users, mainly in its documentation and interface. In other words, users will not use a search engine unless they are comfortable with its interface, and able to read and comprehend its documentation when consulted. Table 2 is a representation of the evaluation methodology we have proposed, using the three Web search tools chosen for the study as examples. We hope that it would provide practical guidelines in evaluating a Web search engine.

Table 2 An Evaluation Methodology of Web Search Engines

Evaluation Criteria	Web Search Engine			
Choice	Alta Vista	Excite	Lycos	Your Choice
Web Indexes				
Coverage [7]	16 Mil.	1.5 Mil.	19 Mil.	
Update	Unspecified	Weekly	Weekly	
Web page indexed	Partial	Full	Partial	
Search Capability				
Boolean search	Yes	Yes	Yes	
Proximity search	Yes	No	No	
Truncation	Limited	No	Automatic	
Field search	Yes	No	No	
Case-sensitivity	Limited	No	No	
Concept search	No	Yes	No	
Retrieval Performance				
Precision	0.78	0.45	0.55	
Recall	Not Tested	Not Tested	Not Tested	
Response time	< 3 sec.	3-5 sec.	About 3 sec.	
Output				
# of formats	2	1	3	
Content	Extract	Abstract	Extract	
User Effort				
Documentation	Good	Poor	Very good	
Interface	Fair	Good	Good	

CONCLUDING REMARKS

Web search engines no doubt are different, in various aspects, from the well-established search tools developed by services such as Dialog and SilverPlatter. Therefore, they require a different evaluation methodology, and we made a preliminary attempt with three search engines and ten sample queries.

In the future, we plan to apply the proposed methodology to a wider scope with the hope that our research findings will truly enable Web users to select a search engine appropriate to their specific search needs, and help Web search engine developers design even better ones for the Internet community.

REFERENCES

Courtois, Martin P., Baer, William M., and Stark, Marcella. (November/December 1995). Cool tools for searching the Web: A performance evaluation. *Online*, 19(6), 14-32.

Harman, Donna. (1995). Overview of the second Text REtrieval Conference (TREC-2). *Information Processing & Management*, 31(3), 271-289.

Lancaster, F.W., and Fayen, E.G. (1973). *Information Retrieval On-Line* Los Angeles, CA: Melville Publishing Co. Chapter 6.

Leighton, H. Vernon. (1995). Performance of four World Wide Web (WWW) index services: InfoSeek, Lycos, WebCrawler, and WWWorm. <http://www.winona.msus.edu/services-f/library-f/webind.htm>.

Leonard, Andrew J. (1996). Where to find anything on the net. <http://www.cnet.com/Content/Reviews/Search/>.

Notess, Greg R. (July/August 1995a). Searching the World-Wide Web: Lycos, WebCrawler and More. *Online*, 19(4), 48-53.

Notess, Greg R. (August/September 1995b). The InfoSeek Databases. *Database*, 85-87.

Scoville, Richard. (January 1996). Special report: Find it on the net! *PC World*, also available at <http://www.lycos.com>.

Shirky, Clay. (October, 1995). Finding needles in haystacks. *Netguide*, 87-90.

Taubes, Gary. (September 8, 1995). Indexing the Internet. *Science*, 269, 1354-1356.

Wildstrom, Stephen H. (September 11, 1995). Feeling your web around the Web.

Business Week, 22.

ENDNOTES

1. Electronic mail is currently the most widely used Internet application.
 2. Such Web search aids take assorted names. Catalogs, indexes, directories, and search engines are some of the examples.
 3. We noticed that the first letter of "Excite" should not be capitalized according to the practice of its developers. But we have intentionally altered the spelling to avoid any confusion that may arise with the actual word "excite".
 4. We define a Web record as all the information displayed for a retrieved Web document. We understand that the contents of Web records differ from one search engine to another.
 5. Precision scores were assigned to each retrieved item using a three-level (1 for relevant, 0.5 for somewhat relevant, and 0 for irrelevant) scoring method. The differences between the two judgements were averaged whenever needed.
 6. Out of the 10 downloaded Web records, there are 2 duplicates.
 7. All the figures were obtained in January 1996.
-

© 1996, American Society for Information Science. Permission to copy and distribute this document is hereby granted provided that this copyright notice is retained on all copies and that copies are not altered.