The History of Information Retrieval Research

This paper provides some perspective on the history of information retrieval that had its beginnings long before the creation of the Internet and provides some enlightened predictions on possible future directions of the field.

By Mark Sanderson and W. Bruce Croft

ABSTRACT | This paper describes a brief history of the research and development of information retrieval systems starting with the creation of electromechanical searching devices, through to the early adoption of computers to search for items that are relevant to a user's query. The advances achieved by information retrieval researchers from the 1950s through to the present day are detailed next, focusing on the process of locating relevant information. The paper closes with speculation on where the future of information retrieval lies.

KEYWORDS | History; information retrieval; ranking algorithms

I. INTRODUCTION

The long history of information retrieval does not begin with the Internet. It is only in the last decade and a half of the IEEE's 100 years that web search engines have become pervasive and search has become integrated into the fabric of desktop and mobile operating systems. Prior to the broad public day-to-day use of search engines, information retrieval (IR) systems were found in commercial and intelligence applications as long ago as the 1960s. The earliest computer-based searching systems were built in the late 1940s and were inspired by pioneering innovation in the first half of the 20th century. As with many computer technologies, the capabilities of retrieval systems grew with increases in processor speed and storage capacity. The development of such systems also reflects a rapid progression away from manual library-based approaches of

acquiring, indexing, and searching information to increasingly automated methods.

An IR system locates information that is relevant to a user's query. An IR system typically searches in collections of unstructured or semistructured data (e.g., web pages, documents, images, video, etc.). The need for an IR system occurs when a collection reaches a size where traditional cataloguing techniques can no longer cope. Similar to Moore's law of continual processor speed increase, there has been a consistent doubling in digital storage capacity every two years. The number of bits of information packed into a square inch of hard drive surface grew from 2000 bits in 1956 to 100 billion bits in 2005 [1]. With the growth of digitized unstructured information and, via high-speed networks, rapid global access to enormous quantities of that information, the only viable solution to finding relevant items from these large text databases was search, and IR systems became ubiquitous.

This brief review of past work focuses on the algorithms that take a user's query and retrieve a set of relevant documents. This paper opens with a review of the early developments of electromechanical and computational devices that searched manually generated catalogs. This is followed by a description of how IR moved to automatic indexing of the words in text and how complex Boolean query languages gave way to simple text queries. The automatic techniques and theories that supported them have continued to be developed for more than 40 years, and provided the framework for successful web search engines. This review finishes with a perspective on the future challenges for IR.

M. Sanderson is with the School of Computer Science and Information Technology, RMIT University, Melbourne, Vic. 3001, Australia (e-mail: mark.sanderson@rmit.edu.au).
**W. B. Croft** is with the Department of Computer Science, University of Massachusetts, Amherst, MA 01003-9264 USA (e-mail: croft@cs.umass.edu).

II. PREHISTORY—MECHANICAL AND ELECTROMECHANICAL DEVICES

Conventional approaches to managing large collections of information originate from the discipline of librarianship.

Fig. 1. *Diagrams from the U.S. patent granted to Goldberg in 1931.*

Commonly, items such as books or papers were indexed using cataloging schemes. Eliot and Rose claim this approach to be millennia old: declaring Callimachus, a third century BC Greek poet as the first person known to create a library catalog [2]. Facilitating faster search of these physical records was long researched, for example, Rudolph filed a U.S. patent in 1891 for a machine composed of catalog cards linked together, which could be wound past a viewing window enabling rapid manual scanning of the catalog. Soper filed a patent for a device in 1918 [3], where catalog cards with holes, related to categories, were aligned in front of each other to determine if there were entries in a collection with a particular combination of categories. If light could be seen through the arrangement of cards, a match was found.

Mechanical devices that searched a catalog for a particular entry were also devised. The first person to build such a system appears to be Emanuel Goldberg who tackled this problem in the 1920s and 1930s. A series of patents were granted to Goldberg for a machine that searched for a pattern of dots or letters across catalog entries stored on a roll of microfilm. Goldberg's original patents can be found on the websites of the German and U.S. patent offices. Part of the U.S. version filed in 1928 [4] is shown in Fig. 1. Here it can be seen that catalog entries were stored on a roll of film (no. 1 of the figure). A query (2) was also on film showing a negative image of the part of the catalog being searched for; in this case, the first and sixth entries on the roll. A light source (7) was shone through the catalog roll and query film, focused onto a photocell (6). If an exact match was found, all light was blocked to the cell causing a relay to move a counter forward (12) and for an image of the match to be shown via a half silvered mirror (3), reflecting the match onto a screen or photographic plate (4 and 5).

According to a biography of Goldberg by Buckland [5], three prototypes of the machine were built, one of which was said to be integrated into Goldberg's desk. Buckland quotes a colleague of Goldberg as saying: "He was telling us that he was the only person in the world as far as he knew who had on his desk a document retrieval capability... He would dial a number, press a button and after three seconds [a microfilmed copy of] the document would be projected."

A number of other researchers produced devices based on similar principles: Mooers [6] described investigations by Davis and Draeger in 1935 on searching with microfilm. This work, according to Mooers, was taken up by Vannevar Bush in the late 1930s, who, with his students, built a film-based searching prototype. The work undoubtedly informed Bush's well-known proposal of the Memex system in 1945 [7]. The culmination of this approach appears to be Shaw's rapid selector [8], which was reported to search through a 2000-ft reel of film. Each half of the film's frames had a different purpose: one half for "frames of material," and the other for "index entries." It is stated that 72 000 frames were stored on the film, which in total were indexed by 430 000 entries. Shaw reported that the selector was able to search at the rate of 78 000 entries per minute.

Other mechanical technologies were examined. Luhn, for example, made a selector using punch cards, light, and photocells. Prototypes of this system were completed in 1950 and demonstrated in 1951. A key feature of this system was that a consecutive sequence of characters could be matched within a larger string [9]. The system searched at the rate of 600 cards per minute. At this time, the term "information retrieval" was first used. Presenting a paper at a conference in March 1950, Calvin Mooers wrote: "The problem under discussion here is machine searching and

retrieval of information from storage according to a specification by subject... It should not be necessary to dwell upon the importance of information retrieval before a scientific group such as this for all of us have known frustration from the operation of our libraries—all libraries, without exception" [10]. Mooers' paper described IR systems using punch cards. As reported by Jahoda [11], these mechanical systems continued to be developed and used until the advent of computers when this approach to IR was surpassed.

III. EARLY USE OF COMPUTERS FOR IR

To discuss means of dealing with a perceived explosion in the amounts of scientific information available, a specially convened conference was held by the U.K.'s Royal Society in 1948. At it, Holmstrom described a "machine called the Univac" capable of searching for text references associated with a subject code. The code and text were stored on a magnetic steel tape [12]. Holmstrom stated that the machine could process "at the rate of 120 words per minute."[1] It appears that this is the first reference to a computer being used to search for content.

Mitchell [13] described a project to model the use of a Univac computer to search 1 000 000 records indexed by up to six subject codes; it was estimated that it would take 15 h to search that many records. Nanus [14] detailed a number of computer-based IR projects run in the 1950s, including one system from General Electric that searched over 30 000 document abstracts; see also [15] for another review of implementations of computerized IR systems in that decade, including mention of IR work conducted in the Soviet Union in the 1950s. The impact of computers in IR is highlighted when Hollywood drew public attention to the innovation with the comedy *Desk Set*, which came out in 1957. It centered on a group of reference librarians who were about to be replaced by a computer.

IR as a research discipline was starting to emerge at this time with two important developments: how to index documents and how to retrieve them.

A. Indexing—The Move Toward Words

In the field of librarianship, the way that items were organized in a collection was a topic that was regularly debated. The classic approach was to use a hierarchical subject classification scheme, such as the Dewey Decimal system, which assigned numerical codes to collection items. However, alternatives were proposed, most notably Taube *et al.*'s Uniterm system [16], which was essentially a proposal to index items by a list of keywords. As simple an idea as this seems today, this was at the time a radical step. A few years later, Cleverdon conducted a detailed comparison of retrieval effectiveness using Uniterms and the

more classic classification techniques [17]. His conclusion that Uniterms were as good as and possibly better than other approaches caused much surprise and his work came under extensive scrutiny [18]. However, Cleverdon's experimental results were found to be correct and as a result the use of words to index the documents of an IR system became established. Many aspects of Cleverdon's *test collection* approach to evaluation are still used in both academic research and commercial search testing today.

B. Ranked Retrieval

The style of search used by both the electromechanical and computer-based IR systems was so-called Boolean retrieval. A query was a logical combination of *terms* (a synonym of word in IR literature), which resulted in a set of those documents that exactly matched the query. Luhn [19] proposed and Maron *et al.* [20] tested an alternative approach, where each document in the collection was assigned a score indicating its relevance to a given query. The documents were then sorted and those at the top ranks were returned to the user. The researchers manually assigned keywords to a collection of 200 documents, weighting those assignments based on the importance of the keyword to the document. The scores assigned to the documents were based on a probabilistic approach. The researchers hand tested their ranked retrieval method, showing that it outperformed Boolean search on this test collection with 39 queries. In the same year as Maron *et al.*'s work, Luhn suggested "that the frequency of word occurrence in an article furnishes a useful measurement of word significance" [21]; his approach later became known as *term frequency* weighting.

This ranked retrieval approach to search was taken up by IR researchers, who over the following decades refined and revised the means by which documents were sorted in relation to a query. The superior effectiveness of this approach over Boolean search was demonstrated in many experiments over those years; see [22, p. 237] for a list of these experiments. Work in the 1950s established computers as the definitive tool for search. What followed was the growth of a commercial search sector and the consolidation of IR as an increasingly important research area.

IV. THE 1960s

The 1960s saw a wide range of activities reflecting the move from simply asking if IR was possible on computers to determining means of improving IR systems. One of the major figures to emerge in this period was Gerard Salton, who formed and led a large IR group, first at Harvard University (Cambridge, MA), and then at Cornell University (Ithaca, NY). The group produced numerous technical reports (the ISR reports), establishing ideas and concepts that are still major areas of investigation today.

One of these areas is the formalization of algorithms to rank documents relative to a query. Of particular note

---

[1]Note that the UNIVAC is not generally thought to have come into existence until 1951, the date when the first machine was sold; Holmstrom presumably saw or was told about a preproduction version.

was an approach where documents and queries were viewed as vectors within an N-dimensional space (N being the number of unique terms in the collection being searched). This was first proposed by Switzer [23], and later, the similarity between a document and query vector was suggested by Salton to be measured as the cosine of the angle between the vectors using the cosine coefficient [24, p. 236].

Another significant innovation at this time was the introduction of relevance feedback [25]. This was a process to support iterative search, where documents previously retrieved could be marked as relevant in an IR system. A user's query was automatically adjusted using information extracted from the relevant documents. Versions of this process are used in modern search engines, such as the "related articles" link on Google Scholar. Relevance feedback was also the first (but not the last) use of machine learning in IR.

Other IR enhancements examined in this period included the clustering of documents with similar content; the statistical association of terms with similar semantic meaning, increasing the number of documents matched with a query by expanding the query with lexical variations (so-called stems), or with semantically associated words. For coverage of this past research, see [24], [26], and [27].

In this decade, commercial search companies emerged out of the development of bespoke systems built for large companies or government organizations. See [28] for a description of one of these early systems that was searching tens of thousands of items. Bjørner states that one of the first companies dedicated to providing search was Dialog formed in 1966 from the creation of an IR system for the National Aeronautics and Space Administration (NASA) [29].

A striking aspect of this time was the low level of interaction between the commercial and IR research communities. Despite researchers' consistent demonstration that ranked retrieval was a superior technique, almost all commercial searching systems used Boolean search. This situation did not change until the early to mid-1990s with systems such as WESTLAW's WIN system [30] and the growth of web search engines.

V. THE 1970s

One of the key developments of this period was that Luhn's term frequency (tf) weights (based on the occurrence of words within a document) were complemented with Spärck Jones's work on word occurrence across the documents of a collection. Her paper on inverse document frequency (idf) introduced the idea that the frequency of occurrence of a word in a document collection was inversely proportional to its significance in retrieval: less common words tended to refer to more specific concepts, which were more important in retrieval [31]. The idea of

combining these two weights (tf ● idf) was quickly adopted; see [32] for an early exploration of such ideas.

A number of researchers worked to formalize the retrieval process. Salton synthesized the outputs of his group's work on vectors to produce the vector space model [33]. This approach to describing the retrieval process underpinned many research retrieval systems and much research for the coming two decades. Today, the ranking formulas proposed by Salton are rarely used, however, viewing documents and queries as vectors in a large dimensional space is still common.

An alternative means of modeling IR systems involved extending Maron et al.'s idea of using probability theory. Robertson defined the probability ranking principle [34], which determined how to optimally rank documents based on probabilistic measures with respect to defined evaluation measures. A further paper from Robertson and Spärck Jones [35] along with a derivation of the probabilistic model in Van Rijsbergen's book [27] stimulated much research on this form of modeling. Van Rijsbergen showed that the basic probabilistic model assumed that words in a document occurred independently of each other, which is a somewhat unrealistic assumption. Incorporating term dependency into ranked retrieval started to be examined, which led to a wide range of research in later years.

VI. THE 1980s TO THE MID-1990s

Building on the developments of the 1970s, variations of tf ● idf weighting schemes were produced (Salton and Buckley [36] reviewed an extensive range) and the formal models of retrieval were extended. The original probabilistic model did not include tf weights and a number of researchers worked to incorporate them in an effective and principled way. Among other achievements, this work ultimately led to the ranking function BM25 [37], which, although not as principled an approach as some researchers would have liked, has proven to be a highly effective ranking function and is still commonly used.

Advances on the basic vector space model were also developed and probably the most well known is latent semantic indexing (LSI), where the dimensionality of the vector space of a document collection was reduced though singular-value decomposition [38]. Queries were mapped into the reduced space. Deerwester et al. claimed the reduction caused words with common semantic meaning to be merged resulting in queries matching a wider range of relevant documents. Tests in the original paper were described as only "modestly encouraging"; nevertheless, the paper has been highly influential.

Unlike the purely numerical approach of LSI for extending the range of documents a query could match, others explored computational linguistics approaches considering the syntax of words, their semantics; addressing anaphora, ambiguity, and named entities. A great deal of

this work led to little or no improvement in the effectiveness of retrieval systems. One technique that was found to produce a level of improvement was stemming, the process of matching words to their lexical variants. Although stemming algorithms date back to the 1960s, Porter in the late 1970s developed a compact set of English language stemming rules; his Porter stemmer [39] continues to influence stemming design today.

A. Text REtrieval Conference (TREC)

One concern in the academic community in the late 1980s and early 1990s was that the size of document collections being used for testing was small compared to the collections that some commercial search companies were working with at the time. Voorhees and Harman formed Text REtrieval Conference (TREC), an annual exercise where a large number of international research groups collaborated to build test collections several orders of magnitude larger than had been in existence before [40]. Working with these new data sets showed that the existing weighting and ranking functions were not ideally suited for these different collections. It was also becoming clear that different collections required different ranking and weighting approaches. This realization was to be further confirmed as web search engines started to be developed in the late 1990s.

B. Learning to Rank

Up to this point, the ranking functions used in search engines were manually devised and tuned by hand through experimentation. Fuhr [41] described work where the retrieval function was learned based on relevant documents identified for an existing set of queries. Whereas Rocchio's relevance feedback tuned the query for a particular search, Fuhr's idea was to tune the ranking function for all queries for a particular document collection. The idea was followed up soon after [42], [43], but only became truly effective when more training data became available in web query logs in the 2000s, along with better learning methods that are able to handle large numbers of features.

VII. THE MID-1990s TO THE PRESENT

Although Berners-Lee created the World Wide Web in late 1990, the number of websites and quantity of pages was relatively small until 1993. In those initial years, conventional manual cataloging of content sufficed. In the middle of 1993, as recorded by Gray's survey,[2] there were around 100 websites; six months later, there were over four times that number, and six months after that, the number had increased fourfold again. Web search engines started to emerge in late 1993 to cope with this growth. The arrival of the web initiated the study of new problems in IR. This point also marked a time when the interaction between the commercial and research-oriented IR communities was much stronger than it had been before. Ideas developed in earlier years were pushed further and implemented in the commercial search sector.

A. Web Search

Until the rise of the web, the collections that people searched were selected and edited from authoritative sources by the search service providers. On the web, however, the situation was different. Search engine developers quickly realized that they could use the links between web pages to construct a crawler or robot to traverse and gather most web pages on the Internet; thereby automating acquisition of content.[3] However, this approach did nothing to ensure a crawled collection contained only authoritative material. Unscrupulous authors discovered that by manipulating the content of a page, they could alter its rank on a search engine. Methods to combat such manipulation (i.e., the various types of spam) and to also identify the best pages on the web were needed.

Two important developments to achieving these goals were link analysis and searching of anchor text, i.e., searching both the content of a web page and the text of links pointing (anchoring) to that page. Both developments were related to earlier work on the use of citation data for bibliometric analysis and search, and using "spreading activation" search in hypertext networks. The anchor text, almost always a brief summary of the page, was recognized early on as a valuable source of information (e.g., McBryan's work in 1994 [44]). The anchor texts were generally written by a number of people, making manipulation of that text harder to achieve. Using anchor text was a key feature of the Google search engine from its early development [45], along with the more famous use of link analysis methods: PageRank developed by the creators of Google and HITS which was developed at the same time by Kleinberg [46].

Adding link analysis and multiple text representations of documents to existing document ranking functions meant that the internal algorithm of an IR system was becoming complex. Correctly setting parameters for these different features was a challenge, which caused a revisiting of the learning to rank approaches started by Fuhr. He was hampered by a lack of training data, however, as search engines became popular, it was realized that logs of user interactions could be exploited for this purpose. The log data are very noisy, but solutions for extracting valuable information were found; see, for example, Joachims's use of logs to train a rank function [47].

---

[2]http://www.mit.edu/~mkgray/net/web-growth-summary.html.

[3]The first crawler for web search was developed for the JumpStation search engine built by Fletcher in late 1993. The search engine offered only basic search of part of the web pages gathered. The first full text search engine using a crawler was WebCrawler released in 1994.

B. Exploiting Query Logs

Automated exploitation of information extracted from the logs of search engines was also examined at that time. Although logs were stored and examined for many years [48], the most they had been used for was to inform subsequent manual adjustment of a searching system [49]. The true potential of extracting valuable information from these logs was only realized when large volumes of people started to use web search engines. Examining users' queries, click patterns, and reformulations of queries enabled researchers to develop more effective query processing techniques based on understanding the user's "intent," such as automated spell correction [50]; automated query expansion [51], and more accurate stemming [52].

C. Other Advances

In the same way that query log analysis had long been known about, but was only researched in detail more recently, it had long been recognized that different users with different information needs might search for that need using the same query [53]. IR systems should be able to serve such diverse needs by finding "differently relevant" documents to rank. Only since the late 1990s has there been a concerted effort to tackle this problem. Carbonell and Goldstein's [54] description of their maximal marginal relevance (MMR) diversity system was a key paper in generating interest in this area.

The retrieval models that are the basis of the core ranking function of IR systems continued to be developed in this period. Of particular note was the introduction of a probabilistic approach using language models, described by Ponte and Croft [55] and by Hiemstra [56]. By taking a new view of the matching process between documents and queries, the language model approach provided new understanding of a wide range of IR processes, such as relevance feedback, forming clusters of documents, and term dependence. Metzler and Croft [57] showed, for example, that incorporating term dependence in the form of proximity operators in a ranking function significantly outperforms term-independence models.

D. New Areas of Search

The applications of search and the field of IR continue to evolve as the computing environment changes. The most obvious recent example of this type of change is the rapid growth of mobile devices and social media. One response from the IR community has been the development of *social search*, which deals with search involving communities of users and informal information exchange. New research in a variety of areas such as user tagging, conversation retrieval, filtering and recommendation, and collaborative search is starting to provide effective new tools for managing personal and social information. An important early paper in this area dealt with desktop search [58], which has many similar characteristics to current search applications in the mobile world.

Much research on web IR has focused on short queries, which have little linguistic structure (typically a single noun compound). Another development has been supporting users who issue longer, more natural questions. Much of this work started with the question answering task in TREC [59] that dealt with finding simple answers in text to a limited range of questions (such as the "wh" questions "who" and "when"). This then progressed into the more detailed questions found in large community-based question-answering archives. Researchers have also been developing techniques that provide more focused answers for more detailed questions. The success of applications such as Apple's Siri, IBM's Watson, and Yahoo! Answers is in part due to this research.

VIII. FURTHER READING

Beyond ranking functions, a wide range of IR research was extensively studied in areas such as information seeking behavior, interface design, implementation of search engines, evaluation, and specializations for particular collection types (e.g., social media, multimedia, etc.). Books from Manning *et al.* [60], Hearst [61], Croft *et al.* [62], and Baeza-Yates and Ribeiro-Neto [63] provide excellent coverage of these other research areas. Recent detailed reviews of specific research areas in IR can be found in the journal *Foundations and Trends in Information Retrieval* and the Morgan Claypool *Synthesis Lectures on Information Concepts, Retrieval, and Services*.

IX. CONCLUSION AND FUTURE DIRECTIONS

The 20th and early 21st centuries were transformational in the way people accessed information. In 1912, a person with an information need would probably go to a local library and, using a card catalog, locate books or documents that hopefully answered that need. Because of the relative inconvenience of accessing information in that way, that person would most likely only seek to answer a small number of questions. The scope of information available to people would be limited by the size of their library; for a small number of very important needs, a loan across libraries might have been arranged. Because of the ubiquity of web-based search, it need hardly be said what the current state of the art is: for those with an Internet connection, one can instantaneously access hundreds of terabytes of web pages, video clips, news, images, social media, scanned books, academic papers, music, television programs, and films; almost always through search engines. In the last few years, the access has been also possible from a mobile phone. Just about the only thing in common between the situation today and 100 years ago is that both services are generally free at the point of use.

Because the systems that are accessible today are so easy to use, it is tempting to think the technology behind

them is similarly straightforward to build. This review has shown that the route to creating successful IR systems required much innovation and thought over a long period of time.

When considering possible future directions, Apple's 1987 Knowledge Navigator vision of IR is still a strong exemplar of how search systems might develop. The short film showed a college professor pulling together a lecture presentation at the last minute. The professor used a form of a tablet computer running an IR system presented as an agent capable of impeccable speech recognition, natural dialog management, a high level of semantic understanding of the searcher's information needs, as well as unbounded access to documents and federated databases.

The Knowledge Navigator identified and connected the professor to a colleague who helped him with the lecture. The broader implications of finding people (rather than documents) to aid with information needs that we see facilitated in the vast growth of social media was not really addressed in the Apple vision. What it also did not encompass was the portability of computer devices opening the possibility of serving information needs pertinent to the particular local context of location, location type, route, the company one is in, or a combination of all these factors.

Today's web search engines seem a simple tool compared to such visions of the future: there are still many opportunities to improve. ∎

REFERENCES

[1] C. Walter, *Insights: Kryder's Law*. Singapore: Scientific American, 2005.

[2] S. Eliot and J. Rose, *A Companion to the History of the Book*. New York: Wiley, 2009.

[3] H. E. Soper, "Means for compiling tabular and statistical data," U.S. Patent US00 135 169 231-1920, 1920.

[4] E. Goldberg, "Statistical machine," U.S. Patent 183 838 929-1931, 1931.

[5] M. K. Buckland, *Emanuel Goldberg and His Knowledge Machine: Information, Invention, and Political Forces*. Westport, CT: Greenwood, 2006.

[6] C. N. Mooers, "The next twenty years in information retrieval: Some goals and predictions," *Proc. Western Joint Comput. Conf.*, 1959, pp. 81–86.

[7] V. Bush, "As we may think," *The Atlantic Monthly*, vol. 176, no. 1, pp. 101–108, 1945.

[8] R. R. Shaw, "The rapid selector," *J. Documentation*, vol. 5, no. 3, pp. 164–171, 1949.

[9] "New tools for the resurrection of knowledge," *Chem. Eng. News*, vol. 32, no. 9, pp. 866–869, 891, 1954.

[10] C. N. Mooers, "The theory of digital handling of non-numerical information and its implications to machine economics," in *Proc. Assoc. Comput. Mach. Conf.*, 1950.

[11] G. Jahoda, "Electronic searching," *The State of the Library Art*, vol. 4, pp. 139–320, 1961.

[12] J. E. Holmstrom, "Section III. Opening plenary session," in *Proc. Roy. Soc. Sci. Inf. Conf.*, London, U.K., pp. 77–94, 1948.

[13] H. F. Mitchell, "The use of the univ AC FAC-tronic system in the library reference field," *Amer. Documentation*, vol. 4, no. 1, pp. 16–17, 1953.

[14] B. Nanus, "The use of electronic computers for information retrieval," *Bull. Med. Library Assoc.*, vol. 48, no. 3, pp. 278–291, Jul. 1960.

[15] H. L. Brownson, "Research on handling scientific information," *Science*, vol. 132, pp. 1922–1931, 1960.

[16] M. Taube, C. D. Gull, and I. S. Wachtel, "Unit terms in coordinate indexing," *Amer. Documentation*, vol. 3, no. 4, pp. 213–218, Jan. 1952.

[17] C. W. Cleverdon, "The evaluation of systems used in information retrieval (1958: Washington)," in *Proc. Int. Conf. Sci. Inf.*, 1959, pp. 687–698.

[18] C. W. Cleverdon, "The significance of the Cranfield tests on index languages," in *Proc. 14th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, Chicago, IL, 1991, pp. 3–12.

[19] H. P. Luhn, "A statistical approach to mechanized encoding and searching of literary information," *IBM J. Res. Develop.*, vol. 1, no. 4, pp. 309–317, 1957.

[20] M. E. Maron, J. L. Kuhns, and L. C. Ray, "Probabilistic indexing: A statistical technique for document identification and retrieval," Thompson Ramo Wooldridge Inc., Los Angeles, CA, Data Systems Project Office, Technical Memo. 3, Jun. 1959.

[21] H. P. Luhn, "The automatic creation of literature abstracts," *IBM J. Res. Develop.*, vol. 2, no. 2, pp. 159–165, 1958.

[22] K. Spärck Jones, Ed., *Information Retrieval Experiment*. Oxford, U.K.: Butterworth-Heinemann, 1981.

[23] P. Switzer, "Vector images in document retrieval," Harvard Univ., Cambridge, MA, ISR-4, 1963.

[24] G. Salton, *Automatic Information Organization and Retrieval*. New York: McGraw-Hill, 1968.

[25] J. J. Rocchio, "Relevance feedback in information retrieval," Harvard Univ., Cambridge, MA, ISR-9, 1965.

[26] M. E. Stevens, V. E. Giuliano, and L. B. Heilprin, *Proc. Symp. Stat. Assoc. Methods Mechanized Documentation*, Washington, DC, 1964.

[27] C. J. van Rijsbergen, *Information Retrieval*. Oxford, U.K.: Butterworth-Heinemann, 1979.

[28] B. K. Dennis, J. J. Brady, and J. A. Dovel, "Five operational years of index manipulation and abstract retrieval by computer," *J. Chem. Documentation*, vol. 2, no. 4, pp. 234–242, Oct. 1962.

[29] S. Bjørner and S. C. Ardito, "Online before the internet—Part 1: Early pioneers tell their stories," *Searcher: The Magazine for Database Professionals*, vol. 11, no. 6, pp. 36–46, Jun. 2003.

[30] H. Turtle, "Natural language vs. Boolean query evaluation: A comparison of retrieval performance," in *Proc. 17th Annu. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 1994, pp. 212–220.

[31] K. Spärck Jones, "A statistical interpretation of term specificity and its application in retrieval," *J. Documentation*, vol. 28, no. 1, pp. 11–21, 1972.

[32] G. Salton and C. S. Yang, "On the specification of term values in automatic indexing," Dept. Comput. Sci., Cornell Univ., Ithaca, NY, TR 73-173, 1973.

[33] G. Salton, A. Wong, and C. S. Yang, "A vector space model for automatic indexing," *Commun. ACM*, vol. 18, no. 11, pp. 613–620, 1975.

[34] S. E. Robertson, "The probability ranking principle in IR," *J. Documentation*, vol. 33, no. 4, pp. 294–304, 1977.

[35] S. E. Robertson and K. Spärck Jones, "Relevance weighting of search terms," *J. Amer. Soc. Inf. Sci.*, vol. 27, no. 3, pp. 129–146, 1976.

[36] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," *Inf. Process. Manage.*, vol. 24, no. 5, pp. 513–523, 1988.

[37] S. E. Robertson, S. Walker, S. Jones, M. M. Hancock-Beaulieu, and M. Gatford, "Okapi at TREC-3," *NIST Special Publication SP*. Gaithersburg, MD: National Institute of Standards and Technology, pp. 109–126, 1990.

[38] S. Deerwester, S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman, "Indexing by latent semantic analysis," *J. Amer. Soc. Inf. Sci.*, vol. 41, no. 6, pp. 391–407, 1990.

[39] M. F. Porter, "An algorithm for suffix stripping," *Program, Electron. Library Inf. Syst.*, vol. 14, no. 3, pp. 130–137, 1980.

[40] E. M. Voorhees and D. K. Harman, *TREC: Experiment and Evaluation in Information Retrieval*. Cambridge, MA: The MIT Press, 2005.

[41] N. Fuhr, "Optimum polynomial retrieval functions based on the probability ranking principle," *ACM Trans. Inf. Syst.*, vol. 7, no. 3, pp. 183–204, 1989.

[42] N. Fuhr and C. Buckley, "A probabilistic learning approach for document indexing," *ACM Trans. Inf. Syst.*, vol. 9, no. 3, pp. 223–248, 1991.

[43] W. S. Cooper, F. C. Gey, and D. P. Dabney, "Probabilistic retrieval based on staged logistic regression," in *Proc. 15th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 1992, pp. 198–210.

[44] O. A. McBryan, "GENVL and WWWW: Tools for taming the Web," in *Proc. 1st Int. World Wide Web Conf.*, 1994. [Online]. Available: http://www94.web.cern.ch/WWW94/PrelimProcs.html

[45] S. Brin and L. Page, "The anatomy of a large-scale hypertextual web search engine," *Comput. Netw. ISDN Syst.*, vol. 30, no. 1–7, pp. 107–117, 1998.

[46] J. M. Kleinberg, "Authoritative sources in a hyperlinked environment," *J. ACM*, vol. 46, no. 5, pp. 604–632, 1999.

[47] T. Joachims, "Optimizing search engines using clickthrough data," in *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Disc. Data Mining*, 2002, pp. 133–142.

[48] D. Meister and D. J. Sullivan, "Evaluation of user reactions to a prototype on-line information retrieval system," NASA, Canoga Park, CA, NASA CR-918, Prepared under Contract NASw-1369 by Bunker-RAM0 Corp., Oct. 1967.

[49] S. E. Robertson and M. M. Hancock-Beaulieu, "On the evaluation of IR systems," *Inf. Process. Manage.*, vol. 28, no. 4, pp. 457–466, 1992.

[50] S. Cucerzan and E. Brill, "Spelling correction as an iterative process that exploits the collective knowledge of web users," in *Proc. EMNLP*, 2004, pp. 293–300.

[51] F. Radlinski and T. Joachims, "Query chains: Learning to rank from implicit feedback," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Disc. Data Mining*, 2005, pp. 239–248.

[52] F. Peng, N. Ahmed, X. Li, and Y. Lu, "Context sensitive stemming for web search," in *Proc. 30th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2007, pp. 639–646.

[53] J. Verhoeff, W. Goffman, and J. Belzer, "Inefficiency of the use of Boolean functions for information retrieval systems," *Commun. ACM*, vol. 4, no. 12, pp. 557–558, 1961.

[54] J. Carbonell and J. Goldstein, "The use of MMR, diversity-based reranking for reordering documents and producing summaries," in *Proc. 21st Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 1998, pp. 335–336.

[55] J. M. Ponte and W. B. Croft, "A language modeling approach to information retrieval," in *Proc. 21st Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 1998, pp. 275–281.

[56] D. Hiemstra, "A linguistically motivated probabilistic model of information retrieval," *Res. Adv. Technol. Digital Libraries*, vol. 1513, *Lecture Notes in Computer Science*, Berlin, Germany: Springer-Verlag, pp. 569–584, 1998, DOI: 10.1007/3-540-49653-X_34.

[57] D. Metzler and W. B. Croft, "A Markov random field model for term dependencies," in *Proc. 28th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2005, pp. 472–479.

[58] S. Dumais, E. Cutrell, J. J. Cadiz, G. Jancke, R. Sarin, and D. C. Robbins, "Stuff I've seen: A system for personal information retrieval and re-use," in *Proc. 26th Annu. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2003, pp. 72–79.

[59] E. M. Voorhees, "The TREC question answering track," *Natural Lang. Eng.*, vol. 7, no. 4, pp. 361–378, 2001.

[60] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge, U.K.: Cambridge Univ. Press, 2008.

[61] M. A. Hearst, *Search User Interfaces*, 1st ed. Cambridge, U.K.: Cambridge Univ. Press, 2009.

[62] W. B. Croft, D. Metzler, and T. Strohman, *Search Engines: Information Retrieval in Practice*, 1st ed. Reading, MA: Addison-Wesley, 2009.

[63] R. Baeza-Yates and B. Ribeiro-Neto, *Modern Information Retrieval: The Concepts and Technology Behind Search*, 2nd ed. Reading, MA: Addison-Wesley, 2011.

ABOUT THE AUTHORS

Mark Sanderson received the B.Sc. (honors) and Ph.D. degrees in computer science from the University of Glasgow, Glasgow, U.K., in 1988 and 1997, respectively.

From 1998 to 1999, he was a Postdoctoral Researcher at the University of Massachusetts Amherst, Amherts. Then, he worked as a faculty member at the Information School, the University of Sheffield, Sheffield, U.K., until 2010. He is now a Professor at the School of Computer Science and Information Technology, RMIT University, Melbourne, Vic., Australia.

Prof. Sanderson is an Associate Editor of *Information Processing and Management* and *ACM Transactions on the Web*. He is Co Program Chair of the 2012 Annual International ACM SIGIR Conference.

W. Bruce Croft received the B.Sc. (honors) degree in applied math and computer science and the M.Sc. degree in computer science from Monash University, Melbourne, Vic., Australia, in 1973 and 1974, respectively, and the Ph.D. degree in computer science from the University of Cambridge, Cambridge, U.K., in 1979.

In 1979, he joined the faculty of the Department of Computer Science, University of Massachusetts Amherst, Amherst, where he currently holds the position of Distinguished Professor.

Prof. Croft was a member of the National Research Council Computer Science and Telecommunications Board (2000–2003) and Editor-in-Chief of the *ACM Transactions on Information Systems* (1995–2002). He was elected a Fellow of the Association for Computing Machinery (ACM) in 1997. He received the Research Award from the American Society for Information Science and Technology in 2000 and the Gerard Salton Award from the ACM Special Interest Group in Information Retrieval (SIGIR) in 2003.