*Reprinted from: B.C. Vickery (ed),* **Fifty Years of Information Progress.** *Aslib, London, 1994 (pp 118-146).*

# COMPUTER RETRIEVAL

## AS SEEN THROUGH THE PAGES OF JOURNAL OF DOCUMENTATION

S.E. ROBERTSON

Department of Information Science
City University
Northampton Square
London EC1V 0HB, U.K.

## INTRODUCTION

The fifty-year span of the Journal of Documentation coincides very precisely with the lifetime of the modern computer. When the Journal began in 1945, the first recognisably modern, working computer was a year or two away (the Manchester Mark 1 first ran in June 1948), although the foundations had been laid (the von Neumann architecture was defined in 1945).

It did take the Journal a little while to discover the computer as such: it was not, in general, interested in either the machinery itself or the theory of the machinery. It was, however, interested both in practical devices and in theories that might relate to information phenomena or to the handling of information. From the point of view of the practical devices, there is of course a considerable degree of continuity between other forms of mechanization (some already well-established in 1945) and the use of computers. Thus for example early issues of the Journal carried quarterly 'Documentation surveys' -- a classified, annotated bibliography of recent publications. The early heading 'Punched-card techniques' became 'Mechanization in Libraries' in volume 3, September 1947, and thereafter included a number of books and articles on computers in retrieval and other aspects of information work.

In this paper, I will try to pull out and comment upon some of the themes and ideas from the history of the development of computer-based methods in information retrieval, of relevant models and theories, and of the evaluation of retrieval systems, as seen through the pages of the Journal. I attempt to represent, at least in overview, the major concerns revealed in the Journal; nevertheless, the selection of material, and particularly the comments thereon, reflect my own personal biases. Although I cover, at least in general terms, events and discussions taking place outside, references will be entirely to items (papers and reviews) appearing in the Journal. Clearly these were not always the first, nor necessarily the best, examples of their type, but they reflect a particular view of the world. The bibliography is listed strictly in chronological order, rather than in the order of reference in the text; thus as I follow particular themes through their own time-sequences, the reference numbers will indicate roughly the chronological relations with other themes.

## PRECURSOR TECHNOLOGIES

Although it is easy to see much development of computer-based systems as technology driven, there is a strong sense in which the information world was ready and impatiently waiting for the computer. (This comment is not in any way negated by the remark above about the slowness of take-up -- what the information world needed was the practical devices for particular tasks, and it took computers a while to reach that status.) There are several pointers to this state of readiness, but I will emphasise just one: the idea of co-ordinate searching.

Post-co-ordination, or co-ordinate searching, was one of the great inventions of information retrieval. We now recognise the principle simply as Boolean AND, thereby placing its invention well outside the field of IR. But the principle was well established in IR long before it was equated with the Boolean operation, and indeed long before we had computers to do it on. In fact much of the pre-computer mechanization work was stimulated by the desire to allow post-coordination, together with the more mundane twin motivations of reducing clerical tasks and sources of error.[1]

The most important predecessor technology (or group of technologies) was the punched card. The variety of punched-card methods was well represented in the Journal [1] [2] [6] [7]. Fairthorne [10] constructed an informational model of clerical systems based on punched card systems. Garfield [12] gave a practical account of the preparation of indexes using Hollerith cards. Ranganathan and Perry [8] included a discussion of punched cards in a paper on the principles of classification and coding for different purposes. A couple of reviews by Dyson attested to the rate of development of punched-card techniques [11]. In 1956, MacKay reviewed two Russian publications concerning a new system based on punched cards [14]. There was also a curious example of a paper which is pure probability theory, with no citations to any IR literature, about superimposed coding (which was used in some punched-card systems) [15].

There were, of course, several very different types of punched cards. Herman Hollerith’s system was invented in the 1890s, for the analysis of census data (a major component of which was cross-tabulation, which is the statistical equivalent of post-coordination). Machine-sorted Hollerith-type cards were used for information retrieval by a number of large information departments, though they tended to be regarded as ‘elaborate and

---

[1] Anyone who doubts the astonishing, if sometimes misguided, ingenuity of designers of IR mechanisms is referred to the brief description by Roberts [116] (writing in 1984 on the history of the thesaurus in IR) of a system operating at RRE Malvern in 1955. This might be described as a peek-a-boo system, except that the "cards" were 12"x12" metal plates, swinging on a heavy metal post, and "punched" by means of a Black & Decker drill!

expensive' [1]. Other types of cards were optical coincidence or peek-a-boo cards (first invented in the 1910s, and the most obvious response to the need for post-coordination), and edge-notched (hand-sorted) cards. For some purposes in IR, punched paper tape was also useful (see e.g. Wilson [34]).

It is worth noting that the distinction between item cards such as edge-notched (one card per item) and feature cards such as peek-a-boo (one card per index term) transferred directly into computer-based IR systems as the distinction between serial searching and inverted files. Inverted files rapidly won out (so that, by the time Dolby discussed programming languages for IR systems in 1971 [65], he assumed inverted files), but the transfer of the discussion itself is another indication of the readiness for computing technology indicated above.

The last paper on punched cards to be published in the Journal did not appear until 1975 (Jolley [86]). Further, in 1983 Vickers [112] mentions a firm (medium sized, technology based, manufacturing) that is still using peek-a-boo cards.

Another class of mechanized system was microform-based systems. Although these have little connection in general with computer-based systems, Shaw's account of the Rapid Selector [5] included some comments which could transfer directly to computers, on the advantages of electronic selection (speed, space, multiple access points).

## EARLY YEARS

### Perceptions of the possibilities

The first Journal of Documentation paper to include the word 'computer' in its title did not appear until 1964 [36]; in the previous year, a paper described some experiments with a computer program [32]. The first description of an operational system appeared in 1966 [42]. However, readers would have been aware much earlier, from a number of mentions in book reviews and reports of meetings, of the approaching transition from punched cards and other forms of mechanization to computers, despite some apparent scepticism. In 1948, Fairthorne reviewed a book by Vannevar Bush (containing the famous essay, 'As we may think'); the review was affectionate but quite critical of the essay [3]. In the same year, the perception of change was reinforced by the Royal Society Scientific Information Conference, whose proceedings were reviewed in 1949 by Lancaster-Jones [4]; she quotes the

> ...desirability of a full-scale experiment in a scientific library with mechanical methods of indexing and selection.

Another article by Fairthorne discussed various devices that might be useful, including robots to carry books around, and OCR devices [9].

### First steps

In 1957, Anthony reported [17] on (and reviewed [18] the proceedings of) a symposium on 'Systems for Information Retrieval' at Western Reserve University. He made an interesting 3-way classification of systems into (a) manual, (b) machine, and (c) computer systems. Of the last category, he said:

> All those demonstrated were experimental systems which are not likely to produce any practical result for 5-10 years...cost will be high...only national documentation centres...

A Perry and Kent book describing in detail the system at Western Reserve was reviewed by Claridge [23]. Hanson reviewed a National Science Foundation report on current research with the main emphasis on machine methods [22].

A reader of the reviews would also be aware that the idea of machine translation was beginning to emerge [13] [21].[2]

### Theoretical concerns

At the same time, some authors were (according to predisposition) either embracing the possibilities or worrying about whether computers really are going to help. Vickery reviewed another book by Perry and colleagues, which discussed at length their ideas on information retrieval in the new machine age [16]. On the other hand Farradane, in a review of a book by Taube [19], was critical of mechanistic approaches to retrieval; the following year (1959) he presented a paper at a conference (reported by Risk [20]), discussing the problem of appropriate logical structures for machine IR (in the subsequent discussion, Garfield defended machine methods). Farradane published a full paper in the Journal in 1961 [25].

In the same year, Fairthorne reviewed at length a report by Bar-Hillel, and was critical of Bar-Hillel's position that machines could do little for information retrieval [24]; in 1963 he was very critical of a book by Taube, which again suggested that the possibilities for making machines do anything very clever were limited [30].

On the other hand, Vickery reviewed (favourably) a book by Fairthorne [27], in which Fairthorne suggested that the benefits of work on automation might for some years come more from the *ideas* it generates (about the nature of the processes involved) than from actual, practical automation activities. Indeed, already in 1960, a book by Maron (reviewed by Vaswani [28]) suggested the use of statistical methods in indexing. Fairthorne [29] also reviewed a U.S. report describing mathematical research relating to information selection; the report was intended for mathematicians, but the reviewer attempted to interpret it for documentalists. He specifically picked out the idea of Boolean logic for searching; but added the rider that the mathematical papers tend to ignore the problem that 'documentation is inherently imprecise'.

## REAL APPLICATIONS

In the late 60s and early 70s, we began to see real applications of computers to real information retrieval tasks. (By the early 1970s, computers were so commonplace that Green could write a paper discussing the indexing of specialist material, referring throughout to 'data bases' or 'data-bases', without once mentioning whether or not a computer was

---

[2] A very much later review of machine translation (Hutchins [98]) prompts the thought that, whatever doubts one may have about computer-based IR, compared to MT it has been outstandingly successful.

involved [71]!) At the same time, a number of themes emerged. Those that have continued in one form or another until the present day are dealt in later sections, but two that might not seem so familiar now are summarised here.

**Printing and publication**

A major impetus to computerisation in information work lies in the tasks involved in printing and publication. Although there is some sense in which printed services are antithetical to computer-based services such as online, nevertheless the production of printed services was an early, fertile ground for computerisation. Thus Lynch [43] discussed a method for generating printed indexes, and later [57] reviewed a number of methods. The PRECIS system was designed with machine-generated indexes in mind (Austin [81]). Gralewska [54] described a system which included printed index generation; this system reflected what became a universal theme in operational systems, of using a database of machine-readable records, from one or more input processes, to generate a number of different forms of output (in this case both current awareness and retrospective search services, via printed output or machine searching). Similarly, Clough and Bramwell [67] described a package (ASSASSIN) intended for industrial use, which accepted input from tape services and allowed for printed indexes generation as well as machine searching.

Perhaps the most influential development of this period was that done by the major publishers of abstracts journals. Although this was not very well represented in the Journal (one exception is Hyslop's [39] discussion of the system from which Engineering Index was produced), there was an excellent review by Martin [76] of computer typesetting, which gave a very clear summary of why the production of such secondary services was so ripe for this development. Martin did, however, make one comment which, in retrospect, is remarkable for having been so wide of the mark:

> ... authors have suggested that the mere existence of large volumes of text in machine-readable form as a result of the increased use of tape-driven composing machines, with or without computers, could have a significant potential from the point of view of future information systems. In the present writer's view, this potential has been much exaggerated and is unlikely to be exploited...

In fact, it would be hard to exaggerate the effect that such availability did have. The market in the tapes of machine-readable data that were the by-product of this development was just starting in the early 70s (for example, the ASSASSIN system mentioned above was designed to take such input). Very shortly afterwards, two U.S. systems developed as in-house services started offering online search services to subscribers on a commercial basis, based entirely on such tape databases: Lockheed's Dialog and SDC's Orbit. Thus was the modern information industry born.

**SDI**

From the mid-60s until Ronald Reagan hijacked the acronym, SDI meant 'selective dissemination of information', a phrase coined by H.P. Luhn. This was a form of current awareness service (i.e. helping the user to keep up-to-date with current publications), intended to be selective enough for an individual or small group, based on a 'profile' of the interests of the individual or group. Much early work in computerised IR was directed at this end rather than at retrospective retrieval. One reason is probably that before the online era, retrospective searching by computer could be very slow (e.g. overnight or worse); current awareness is more obviously suitable for such hardware.

Corbett [47] ran SDI from a tape service; McCash and Carmichael [61] used an in-house database. Hall *et al.* [58] prepared a number of specialist current awareness bulletins. The ASSASSIN system referred to included an SDI function. Barker *et al.* [72] experimented with methods which might lead to automatic profile construction for SDI (an early form of relevance feedback -- see next section). Leggate [84] reviewed SDI services.

The idea of SDI enjoys periodic revivals of interest -- for example, there was a paper in the Journal in 1990 [132]. More usually now a different name is used -- recent examples have been 'filtering' and 'routing'.

## THE DEVELOPING CONCEPT OF IR

This century's transformation of information retrieval, from card and printed indexes, via various forms of mechanization, to computer-based systems, has been one continuous process of re-evaluation. In some sense, to pick up the story at the point where, more-or-less simultaneously, computers were invented and the Journal of Documentation started, is to begin in the middle. In particular, there is some difficulty with the terminology, which was already in transition at the time computers began to be used; the use of computers then accelerated the process.

Take, for example, the word 'indexing'. Traditionally, it means 'creating an index' or perhaps entering in an (existing) index; the index was a thing which the searcher could see. Nowadays, it is much more likely to mean the assignment of index terms to an item, whatever subsequently happens to them -- and indeed the searcher may have no idea what mechanisms are involved when he or she puts a request to a system. Sometimes the term is used to describe an internal mechanism -- the process of generating an inverted file (in whatever form the input comes). If the system is free-text, then the process of generating an inverted file includes the process of assigning index terms (even if only trivially); if the system involves human indexing, then the processes are mutually exclusive. On the other hand, if we are talking about machine-generated printed indexes, then the distinction is more complex.

Some such terminological problems were becoming apparent at earlier stages of mechanization. Furthermore, as the example shows, the terminological problems reflect more fundamental problems about distinguishing different functions and different parts of processes in IR, given the changing technological context.

### Indexing

*Human indexing for computer retrieval*

The question of whether and what kind of human input is required for a computer retrieval system, beyond the simple bibliographic record, abstract and/or text of the document, is clearly a vexed one. Some authors adapted an existing form of human indexing (or classification) to machine operation (e.g. Caless and Kirk [48] with UDC, or Hines [49] with LC and Dewey codes); others attempted to assess the value of including human indexing (e.g.

Olive, Terry and Datta [78], Barker, Veal and Wyatt [70]). Later, there were some authors who considered quite different kinds of human indexing, on the assumption that they would be used in computer systems (e.g. Kircz [140] with rhetorical structures).

*Automatic indexing*

Whether or not there has been some intellectual input, the question of what machine manipulation of the input record is required at the input stage is another subject for much discussion. Many authors considered automatic indexing processes which might be regarded as equivalent to, or perhaps alternatives to, human indexing. I have already referred to the early book by Maron [28] discussing the possibility of using statistical data as the basis for an automatic indexing process -- an idea that was to be reiterated many times. Another early contribution, by Artandi (in fact the first Journal paper to discuss a computer program in some detail) [32], discussed the detection of proper nouns in text, by means of rules specifying patterns in text indicating the presence of proper nouns. (At a conference I attended in late 1993, a paper discussed the automatic detection in text of citations, with particular reference to author names. *Plus çà change*!) Later, Artandi described a project in which similar rules are used for subject indexing of text [56]; this was interesting for being a long time before the idea of rule-based expert systems become so common.

Dillon and McDonald [111] demonstrated a method of book indexing, using dictionary-based syntactic tagging of words, rule-based identification of multi-word content-bearing units (concepts), and grouping of concepts. Field [87] and Robertson and Harding [115] devised methods by which the system would learn, from a training sample, what index terms a human would assign to a document, given some other information about that document (such as free-text terms). Rada *et al.* [126] discussed augmenting a thesaurus with additional entry terms in order to improve automatic indexing.

Several authors discussed automatic methods which bear some resemblance to indexing, but do not result in conventionally indexed documents with direct subject descriptions. For example, Salton described an experiment with citation indexing [63]; Martyn [38] defined bibliographic coupling (linking items by citations-in-common). Griffiths, Robinson and Willett [114] and Enser [117] derived clusters of documents (classification without labels). Needham and Sparck Jones [33] clustered the terms already used to index document (by some kind of free-text process); again, the clusters themselves do not have labels. Clustering is discussed further in the section on retrieval system theory, below.

*Free text*

The idea that one might do away with anything that looks like an intellectual indexing step altogether, by using free-text records, is an attractive one from an economic point of view, though it only begins to make any kind of sense once the move from punched cards to computers is established. Most early writers assumed that some form of indexing was required; as an intermediate step, Shaw and Rothman [53] proposed that words should be chosen from the text by a human indexer, but not controlled in any way. Gralewska's [54] system included title words as well as controlled-language indexing. But by the early seventies, the possibilities of free-text were being explored, and the free-text versus controlled indexing debate was in full spate. For example, Barker, Veal and Wyatt [70] compared searching of titles and abstracts with index terms; similarly Hersey [66], Olive, Terry and Datta [78]. Although the frequency has declined, there are still some papers on the same lines (for example Cousins [143]).

Actually the distinction between free-text systems and automatic indexing is decidedly fuzzy. Consider the following steps, some of which are present in all free-text systems, and others are sometimes used. Although all might be regarded as elementary or trivial from the point of view of traditional human indexing, they must nevertheless be treated as some sort of indexing operation.

(a) Free-text indexing from part of the item only (e.g. title or abstract). This is clearly making use of someone else's selection of the important words to describe an item.

(b) Word identification. There must be a set of rules for this, dealing not only with word separators such as blank characters and punctuation, but also with upper-lower case, embedded hyphens or hyphens at the end of lines, numbers etc.

(c) Stop-lists. Most systems identify and exclude certain common words (the list is usually manually prepared).

(d) Stemming or suffix stripping.

(e) Dictionary operations such as identification of phrases, acronyms, synonyms.

(f) Inverted file generation. In some sense, this step subsumes all the above. However, it also has its own built-in effects on the later searching stage which might force it to be regarded as a form of indexing in its own right: for example, the traditional inverted index makes it easy to do right-hand-truncation at the searching stage, but much more difficult to allow left-hand-truncation.

It is, of course possible to have a system which does none of these things -- for example, the searching facilities provided in word-processing packages, which involve serial scanning of text. Furthermore, one could probably produce anecdotal arguments against any of them -- for example, although case folding is clearly in general a good thing, the distinction in a medical database between AIDS and (hearing) aids is one that one would like to maintain! But it is generally assumed that even the most minimal retrieval system must include some of these steps. The reasons for this assumption are of two quite different kinds:

(i) Efficiency: There is little hope of providing the kind of speed of search required, on the kind of size databases required, without using inverted files. Some of the above steps are simply necessary for inverted files; some are desirable from the point of view of inverted file size.

(ii) Effectiveness: Many or all of these steps can (to a greater or lesser extent) be justified on the grounds of providing better retrieval performance.

These two reasons have become very firmly intertwined, and it is now difficult to disentangle them. So the possibility (if not now, then in the future) that technological developments such as parallel processing will render the first reason obsolete clouds the issue greatly. Not that it is obvious that the first reason *will* necessarily become obsolete -- although the size and power of computers is increasing at a phenomenal rate, so too is the size of the databases.

Hutchins [50] did propose a system that would work without indexing (the principle was that at search time, all possible variants of search terms as they might appear in text would be generated, and matched against the raw text).[3] However, if one were trying to

[3] Seventeen years later, Sparck Jones and Tait [113] experimented with a very similar procedure, without referencing Hutchins.

implement his system, one would almost certainly use an inverted file for efficiency reasons, and in the process do at least (b) above. Sharp [85] appealed for the term 'natural language' to be used only where the language of the document is not changed at all (though it is not clear whether he would exclude even case conversion).

## Searching

*Boolean logic*

As indicated above, the transition from punched-card systems (which were designed to allow for post-coordination) to Boolean search logic in computer-based systems was a fairly natural one, and from an early stage it was assumed that searching would involve Boolean logic. Fairthorne's discussion of the subject in a review [29] has already been mentioned. Authors have continued to write about the use of Boolean logic (sometimes including those extensions to the logic which text retrieval has induced almost without noticing, such as the implied-OR in truncation and explosion, and term adjacency and proximity operators). For example, Harley, le Minor and Weil [75] discussed a universal search formulation language; Barraclough [93] surveyed work on online retrieval; Dillon and Desper [103] developed a method for the automatic modification of Boolean queries following relevance feedback (see below); Radecki [105] [109] discussed the relation between Boolean and weighted retrieval; Vickery et al. [122] described an expert system which manipulates Boolean search statements. Caless and Kirk [48] extended Boolean logic in a different way when searching on UDC codes, by adding order operators.

*Associative methods*

The alternative to Boolean and similar methods for search statement construction is to use some kind of associative method such as search term weighting. The most obvious difference between the two approaches is that Boolean-type searches are dichotomous: an item is either retrieved or not. Associative methods tend to be used to rank items retrieved, so that the items which match the search statement best are at the top of the ranking. In this respect, associative methods are feasible only in computer systems. Robertson and Belkin [97] discuss the principles of ranking.

Associative methods also tend to make use of obviously statistical information such as term occurrence or co-occurrence or frequency within or between documents etc. However, there is recent interest in making use also of term position information in text, which has usually been seen in Boolean systems (Keen [137]).

Associative methods have been more common in experimental systems or environments than in operational ones: for example, Salton's system [63] [80] or Sparck Jones' experiments [69] [101] [113]. They have also figured more often in theoretical papers (see Robertson [92] for some examples). However, some authors have described work in something approaching an operational environment (e.g. Miller [68]).

*Interactive searching*

In the early seventies, online searching became feasible. Most online systems were based on batch search systems, and they tended to have similar facilities (Barraclough [93]). A few authors began to consider the possibilities of highly interactive retrieval (e.g. Oddy [90]); however, on the whole the most interesting work on interactive IR was happening in the context of manual searching (e.g. Keen [94], Ingwersen [108]). The unwillingness of the major online hosts to make substantial changes in their interfaces lead eventually to consideration of alternative ways to provide interactive help to searchers, such as front-end systems. Two substantial papers in the Journal have surveyed, reviewed and analysed these aids and the principles on which they are based (Efthimiadis [133], Vickery and Vickery [145]). One particular interface was described by Vickery and Vickery [142].

One interactive mechanism which has been the subject of a number of Journal papers is relevance feedback. The principle is that if the user indicates to the system which items (resulting from a first attempt at searching, say) are of interest to her/him, then the system can modify the search statement to match more closely the desired items, and thus find more similar items. In effect, the user's relevance judgements provide indirect evidence as to her/his real need, in addition to the information provided directly in the form of a query. The use of relevance feedback in SDI has already been mentioned (Barker, Veal and Wyatt [72]); other papers on the subject are Sparck Jones [101], Dillon and Desper [103], Wu and Salton [104], Robertson [119] [136], and Hancock-Beaulieu and Walker [144].

Although not all relevance feedback methods are based on this view, the idea fits very well with the probabilistic model for IR (see the section on retrieval system theory, below).

Another aspect of interaction which has received some attention is display. In fact there has been an interesting continuity between the work on printed indexes mentioned above, and later work on the display of records, index entries, concepts and relations on computer screens. Examples of work in this area are Craven [106] [123] [135], Bertrand-Gastaldy and Davidson [120], Sano [138], and Bovey and Brown [124].

### Theories and models

From the point of view of the Journal of Documentation, the major effect of the development of computer-based information retrieval systems (as predicted by Fairthorne) was not so much in the development of specific, practical systems, as in the stimulation of ideas. (The more formally explored and presented ideas, of course, become theories or models.) Experiments or experimental systems are then often used to test the ideas. Fuller discussion of these aspects must wait until after a section on the evaluation of systems.

## INFORMATION RETRIEVAL EXPERIMENT

Although there is now a strong association between computer-based retrieval and evaluation, and although the history of formal evaluation begins almost simultaneously with the history of computer-based methods, the two developed quite independently of each other for some time. Much of the early evaluation work was done on manual systems.

### Cranfield

The remark at the Royal Society conference in 1948 on the desirability of experimentation was mentioned above, but the first major evaluation experiment in information retrieval was the first Cranfield experiment, begun in the late fifties. This experiment was very well represented in the Journal, beginning with a long analytical review by O'Connor of two early reports from the project [26]. This review was interesting for, among other things, suggesting analytical experiments looking at the results for individual queries. (This idea was taken up in the big Medlars experiment in the mid-sixties (Lancaster [62]), but relatively seldom since.) The final Cranfield 1 report also prompted a substantial review (Mote [31]), and then a whole series of longer articles. Kyle [35] and Hyslop [39] both tried to draw conclusions from Cranfield 1 that would be of practical significance in operational systems. Brownson [40] reviewed the state of the art of evaluation, with particular reference to Cranfield 1; Fairthorne [41] discussed some theoretical issues. The idea of testing systems began to spread (e.g. Martyn and Slater [37]; Rolling [42]; Martyn [45]). It is very clear that Cranfield 1 had a major impact on our perception of information retrieval systems, and of the possibility of experimental study of IR.

## Methodology

Brownson [40] also announced the funding of the two major studies of relevance of the sixties, one of which was later reported in the Journal (Cuadra and Katter [51]). Relevance was also discussed by Barhydt [46], and other aspects of testing by Saracevic and Rees [44]. Subsequently, the Journal published many papers on methodological and/or theoretical aspects of evaluation, particularly evaluation measures (Brookes [52]; Robertson [55]; Miller [64]; Brookes [73]; Cleverdon [74]; Heine [77] etc.).

One might argue that, having raised the possibility and hope of being able to treat IR as an experimental discipline, Cranfield 1 dashed it again by simply revealing the extreme difficulty of devising adequate methodologies. Certainly the problems are severe, and although the frequency of methodological papers has declined, this is in no sense because they have been solved. The recent TREC project in the United States (about which more below) has reinforced this point.

## Evaluation experiments

Cleverdon in 1970 reviewed evaluation tests up to that point [59]. He emphasised the way in which the seeking of experimental evidence (rather than philosophical argument or anecdote) had become an accepted method of enquiry in information retrieval. He also particularly excluded from his consideration studies which implement and test only part of a system (for example, inter-indexer consistency studies) -- a point to which I return below.

Many specific tests were reported in the Journal, some on manual and some on computer-based systems (for example, Corbett [47], Shaw and Rothman [53], Searle [60], Salton [63], Miller [68], Sparck Jones [69], Barker, Veal and Wyatt [70] etc.). Some of these tests were intended to provide data for specific decision-making (e.g. about particular features of an operational system, or between competing systems). Some experiments, on the other hand, were intended to inform more generally about information retrieval: to establish general principles of system design or implementation. Relatively rarely, experimental studies were undertaken without actual evaluation, but designed to characterise systems, procedures, methods or databases, in a way which might increase our understanding of evaluation results: an example is van Rijsbergen and Sparck Jones [79].

### Test collections

From Cranfield on, many experiments have been performed on datasets created for earlier experiments (the prime example of this was the Cranfield 2 data, which has been used for an astonishing number and range of experiments since it became available). This led to extensive consideration in the UK in the 70s about the possibility of creating a new, bigger and better test collection of material (documents, requests and relevance judgements). It was referred to as the 'ideal' test collection, and the first published paper about it appeared in the Journal (Sparck Jones and van Rijsbergen [89]).

Unfortunately, the ideal test collection project never got off the ground. However, it eventually inspired a similar project in the United States, which began in 1991: TREC (Text REtrieval Conference). The basis of TREC is that a central organisation builds the test collection, and researchers around the world use it to test their own methods and systems, reporting back to the conference with results presented in a standardised way. The TREC collection is far larger than any previous test collection, which makes the exercise extremely interesting; however, it must also be said that the TREC methodology reflects strongly its origins in the 70s ideal test collection proposal, and (in my view) need substantial development to bring it closer to current concerns, such as highly interactive systems.

## ON RETRIEVAL SYSTEM THEORY

People have been theorizing about information retrieval and retrieval systems since well before the period covered by this review. However, the ideas, theories and models developed during this period, whether or not they make explicit reference to the fact, have been very strongly influenced by both the practice of computer-based retrieval and the understanding and empirical knowledge deriving from evaluation experiments. Two substantial examples of this influence follow.

### The whole system, not the parts

I referred above to the problems of changing and confused terminology in information retrieval. In part this arises because the boundaries between different parts of the process become less clear as we realise the possibilities offered to us by computers. Thus for example one reason for difficulty with the term 'indexing', reinforced by the discussion on free text above, is that some operations can be carried out either at indexing or at search stage. Given this freedom, it is no longer clear what we might theoretically call 'indexing', so the terminological confusion is not surprising. More importantly, in order to understand how a system might behave or perform, we need to have the whole system; it would not make sense, in these circumstances, to even consider trying to construct or to evaluate the indexing stage alone.

This situation strongly suggests a holistic approach to modelling or theorizing about information retrieval. Although not all theoretical papers do so, there is certainly greater awareness of the role that parts play in the whole. To reinforce the terminological point, when Salton [63] or Sparck Jones [83] refer to 'automatic indexing', they both in fact treat the indexing stage as part of the whole retrieval process, and indeed do not strongly distinguish the different parts.

However, holism has severe disadvantages. The reason one would *like* to construct and evaluate (say) the indexing stage alone, is that it would simplify matters greatly. If it were possible to define clearly what the indexing stage is, and what its function is (in relation to the whole), and to measure how well it performs its function, then it would clearly be better to do that (and to build models and theories for that purpose), without being concerned with the other parts of the system. Unfortunately, we seem to be unable to make that separation.

A recent example of this argument that I have come across concerns the evaluation of stemming algorithms. We have some idea what stemming means, and that it contributes a little (not much) to system performance. The problem is, do we evaluate a stemmer by embedding it inside an entire retrieval system, and doing a conventional retrieval test, or do we try to assess it directly? The latter would be much simpler (and potentially much more powerful in a diagnostic sense), but it depends on devising criteria for stemming that we can relate to IR system performance without actually doing the experiment.

**The function of the system**

It is possible to argue (indeed, I have done so on many occasions) that information retrieval systems have been around for at least two-and-a-half millennia. The justification for this argument is that all library classification schemes (as well as various more recent inventions such as card catalogues and printed indexes) are in fact information retrieval systems. I have no difficulty with this statement, but the designers of those systems might not see it that way.

In particular, the purpose or function of a classification system might perhaps have been expressed in terms of the proverb, ‘a place for everything and everything in its place’. The idea of ‘putting a query to’ a classification scheme would seem, on the face of it, absurd.

However, as soon as the concept of an information retrieval system exists, and we begin to try to define what it is for, then it becomes clear that dealing with queries (requests for information), pointing them in the direction of appropriate documents, is precisely what a classification scheme is for. Indeed, a traditional library classification scheme (UDC) was among the four systems tested in Cranfield 1.

This idea has far-reaching implications for theorists of information retrieval, whether they see themselves as addressing library classification or any other possible component of a system. The evaluation experiments actually take a rather narrow and restricted view of the function of the system, which might be expressed in the following way:

> to retrieve in response to a request documents (items) that will be judged by the requester (or end-user) to be relevant to the request (or underlying need, or anomalous state of knowledge).

In effect, the theorist now has the choice of accepting such a definition of function, or of conceptualising the function of the system in a different way. What he or she can no longer do is to ignore the question of function.

An example of the kind of discussion that follows from this observation is given by Robertson and Belkin [97], who address the relation between the question as to whether relevance is binary or multi-valued, and the design of ranking systems.

**Probabilistic models**

The prime example of the influence of the idea of evaluation on theory lies in probabilistic models, which were well represented in the Journal. Essentially a probabilistic model involves a proof that, given certain modelling assumptions, a particular procedure will give optimum performance (Robertson [92]). One of the major links in this proof is the Probability Ranking Principle (Robertson [95]). Miller's paper [68] stimulated some of the work on probabilistic models, and developments were reported in many papers (van Rijsbergen [91], Harper and van Rijsbergen [99], Croft and Harper [102], Radecki [105], Bookstein [110], Thompson [127] etc.).

Relevance feedback, discussed at length above, fits very naturally in the probabilistic framework: documents judged relevant by the user can be taken as providing direct sample evidence concerning the various probabilities of interest in the models. Indeed, in this framework one could see user-provided examples of relevant documents as a *more* natural way to express a query than a verbal description.

Probabilistic models do not necessarily take a holistic view of retrieval -- indeed, one of the problems of probabilistic models of searching is that they take the indexing as given -- but they nevertheless force the integration of some elements that were previously regarded as separate. For example, an associative retrieval technique might involve the assignment of weights to search terms and a match function which measures how similar to the query is any particular document. Some authors treat the two components, the weighting function and the matching function, as separate -- that is, they assume that a decision on a good weighting function is required, and also a decision on a good match function, but do not see any strong connection between the two. Probabilistic models of searching, however, require that the weighting function be regarded as a component of the match function.

**Cognitive models**

Although cognitive models are not well represented in the Journal, one of the early papers on the ASK model of Belkin and others appears here [107], linked to the seminal paper by Oddy [90]. Ingwersen considers manual searching from a cognitive point of view [108]. Daniels reviews cognitive models in IR [121].

Cognitive approaches to IR usually start from the user end. It is, of course, possible to consider authorship as a cognitive activity; however, the considerations above about the function of the system effectively dictate that the user should be central to the cognitive view (retrieval success such as relevance is assessed by the user, not by the author).

Although one can argue that the incorporation of a cognitive model of the user could potentially be of great benefit, the actual use of such models is fraught with difficulties. However, there have been some spinoffs from these concerns: in particular, the idea of using expert-system techniques and/or knowledge bases in the user interface, and concern with user information-seeking behaviour (whether or not s/he actually uses a formal IR system).

Given that much searching was and is undertaken by intermediaries on behalf of end-users, there is an obvious argument for trying to design an interface that has the expertise of the intermediary, and thus makes end-user searching easier. A number of attempts have been made in this direction; the two which are best represented in the Journal are Plexus (Vickery, Brooks, Robinson and Vickery [122]) and its successor TOME Searcher (Vickery and Vickery [145]). Both these systems might be described as knowledge-based: apart from the knowledge and skill of the intermediary, both try to include a knowledge base of a kind with which we have been familiar for over forty years, namely a thesaurus. (The re-emergence of thesauri as knowledge bases reminds one of the remark by M Jourdain in *Le Bourgeois Gentilhomme,* about having spoken prose for over forty years without realising it -- even the time period is right!) Although knowledge engineers from other fields might have difficulty in recognising a thesaurus as a knowledge base, nevertheless it is clear that the description is correct, and indeed would also apply to a more traditional library classification scheme. The big question, though, which is still unresolved, is how to design a computer system to make best use of such knowledge. It is not obvious that it must use traditional expert system techniques such as production rules (like Plexus or TOME); for example, Kim and Kim [131], Rada *et al.* [139], and Lee, Kim and Lee [146] all use thesaurus information in weighted search systems.

Vickery reviews knowledge representation techniques in different areas [118].

The understanding of user information-seeking behaviour has become more and more central to IR system theory and design. A particularly good environment for such study is the library OPAC (online public access catalogue), since huge numbers of searches on OPACs take place daily, without the benefit (or interference) of intermediaries (Hancock-Beaulieu [125] [134] [144], Akeroyd [130]). User behaviour may be understood in a cognitive fashion (Ingwersen [108]), but Ellis [129] argues for a behavioural view which does not require cognitive interpretation. He also [141] distinguishes two paradigms operating in IR research: the physical and the cognitive.

**Mathematical models**

Apart from the probabilistic models discussed above, a number of mathematical models have been enlisted to the service of IR theory. No other class of models, however, seems able to make a direct connection between performance and design, in the way that the probabilistic approach does.

The Swets model, applying signal detection theory to IR systems, does indeed address the question of performance, though less obviously that of design. This model is well represented in the Journal (Brookes [52]; Heine [77] [88]; Bookstein [82]; Hutchinson [96]. Another related model is the Shannon model (Brookes [73]). More generally, however, mathematical models deal with the internals of systems (at least, these seem to be the aspects which are most amenable to mathematical representation).

Perhaps the best-known mathematical model is the vector-space model which is the basis for Salton's SMART system. In fact the first mention of a vector-space model in the Journal is in a review by Vaswani of a report on 'self-organising' files [28]. SMART is represented by a few papers (Salton [63]; Salton and Yang [80]; Wu and Salton [104]); Salton [100] reviews various mathematical models. The vector space model essentially regards the operation of indexing as locating each document as a point in a multi-dimensional space (the axes of the space correspond to the indexing terms available -- thus the space may have thousands of dimensions). Queries are similarly associated with points in the same space. This model does not directly address the question of performance; it does, however, suggest various kinds of mechanisms: for example

(a) associative matching methods generally; more specifically, a match function based on a measure of distance in the space;

(b) relevance feedback: moving the query nearer to the documents judged relevant;
(c) document clustering: bringing together documents which seem to be similar (close);
(d) document space modification: using relevance feedback to adjust the indexing of documents already in the system, by moving them closer to queries to which they have been judged relevant.

Thus the spatial view might be said to encourage, without actually providing any strong justification for, certain kinds of internal mechanism or operation. It is no accident that much work in the vector-space area is strongly empirically based: having suggested a mechanism, the model has nothing to say about whether it might be a good one or not, and leaves this question to experimental resolution.

Clustering in general, whether inspired by the vector space model or not, might be seen in the same light. There are in fact two kinds of clustering in IR: clustering of documents (based on the terms occurring in them), and clustering of terms (based on the documents they occur in). Both methods are represented in the Journal (see above). Similarly, one may use relevance information for the benefit of future searches, either by document space modification as suggested by the vector space model, or by assessing the values of different terms (Biru *et al.* [128]).

## TOWARDS INFORMATION RETRIEVAL

At the beginning of this paper, I said that the information world in 1945 was ready and waiting for the computer. In the ensuing fifty years, the computer has of course become the major tool for the practice of IR. But more importantly, it has been instrumental in changing fundamentally our perception of the nature of information retrieval. Together with the development of theories and models, and of the perception of IR as a field for empirical investigation and experiment, the mere existence of the tool has served to alter our ways of thinking about the problem, almost out of all recognition.

This process is by no means complete. All three reasons remain powerful agents for change, individually and in combination. For example, we have scarcely begun to get to grips with the effective empirical study of highly interactive systems (such as graphic user interfaces) and their use in retrieval. And while there is not now (and may never be) any overall theory of information retrieval, new models and theories continue to contribute to our understanding.

Information retrieval is a real-life problem, and as such has generated many practical papers discussing operational systems and services. It is also an area of exciting intellectual endeavour which is worthy of, and has sometimes attracted, real intellectual power. Reading back over old issues of the Journal, as an active participant myself, I was sometimes inevitably annoyed at what, with hindsight, I perceive as naïvety or worse. I was also, however, often stimulated by the ideas, arguments, ingenuity, and foresight that I found. We have come a long way since 1945, but the journey has been as interesting in retrospect as it was at the time, and as the field is now. I firmly believe that the next fifty years will be just as exciting.

## REFERENCES


1. FULLMAN, B. The planning and equipment of libraries in research organisations and industrial concerns. *Journal of Documentation, 2,* 1946-47, 68-75.

2. FERRIS, L., TAYLOR, K. *and* PERRY, J.W. Bibliography on the uses of punched cards. *Journal of Documentation, 3,* 1947-48, 250-271.

3. FAIRTHORNE, R.A. Review of: BUSH, V. Endless horizons, Public Affairs Press, Washington, 1946. *Journal of Documentation, 4,* 1948, 210-211.

4. LANCASTER-JONES, J. Review of: Royal Society Scientific Information Conference, The Society, London, 1948. *Journal of Documentation, 5,* 1949, 114-118.

5. SHAW, R.R. The Rapid Selector. *Journal of Documentation, 5,* 1949, 164-171.

6. VICKERY, B.C. The structure of a connective index. *Journal of Documentation, 6,* 1950, 140-151.

7. COBLANS, H. Some notes on American practice in documentation. *Journal of Documentation, 6,* 1950, 206-212.

8. RANGANATHAN, S.R. *and* PERRY, J.W. External memory and research. *Journal of Documentation, 7,* 1951, 10-17.

9. FAIRTHORNE, R.A. Automata and information. *Journal of Documentation, 8,* 1952, 164-172.

10. FAIRTHORNE, R.A. Information theory and clerical systems. *Journal of Documentation, 9,* 1953, 101-116.

11. DYSON, G.M. Review of: CASEY, R.S. and PERRY, J.W. Punched Cards, their application to science and industry, Reinhold, New York, 1951. Also of: McGRAW, H.F. Marginal punched cards in college and research libraries, Scarecrow Press, Washington, 1952. *Journal of Documentation, 9,* 1953, 129-130.

12. GARFIELD, E. The preparation of subject-heading lists by automatic punched-card techniques. *Journal of Documentation, 10,* 1954, 1-10.

13. VICKERY, B.C. Review of: LOCKE, W.N. and BOOTH, A.D. Machine translation of languages, Wiley, New York, 1955. *Journal of Documentation, 12,* 1956, 126-127.

14. MACKAY, A. Review of: CHERENIN, V.P. Certain problems of documentation and the mechanisation of information searches; also of: CHERENIN, V.P. and RAKOV, B.M. The experimental information machine of the Institute of Scientific Information of the USSR Academy of Sciences, both published by the USSR Academy of Sciences, 1955. *Journal of Documentation, 12,* 1956, 189-191.

15. OROSZ, G. *and* TAKACS, L. Some probability problems caused by the marking of codes into the superimposition field. *Journal of Documentation, 12,* 1956, 231-234.

16. VICKERY, B.C. Review of: PERRY, J.W., KENT, A. and BERRY, M. Machine literature searching, Interscience, New York, 1956. *Journal of Documentation, 13,* 1957, 164-165.

17. ANTHONY, L.J. Systems for information retrieval. *Journal of Documentation, 13,* 1957, 209-221.

18. ANTHONY, L.J. Review of: SHERA, J.H. *et al.* Documentation in action, Reinhold, New York, 1956. *Journal of Documentation, 13,* 1957, 231-232.

19. FARRADANE, J. Review of: TAUBE, M. *et al.* The mechanization of data retrieval. Documentation Inc., 1957. *Journal of Documentation, 14,* 1958, 160-161.

20. RISK, J.M.S. Proceedings of the International Study Conference on Information Retrieval (summary). *Journal of Documentation, 15,* 1959, 70-80.

21. HILL, E.G. Review of: BOOTH, A.D. *el al.* Mechanical resolution of linguistic problems, Butterworths, 1958. *Journal of Documentation, 15,* 1959, 118-119.

22. HANSON, C.W. Review of: U.S. N.S.F. Office of Scientific Information. Current research and development in scientific documentation. *Journal of Documentation, 15,* 1959, 163.

23. CLARIDGE, P.R.P. Review of: PERRY, J.W. and KENT, A. Tools for machine literature searching, Interscience, 1958. *Journal of Documentation, 15,* 1959, 231-233.

24. FAIRTHORNE, R.A. Review of: BAR-HILLEL, Y. Some theoretical aspects of the mechanization of literature searching, Office of Technical Services, Washington, 1960. *Journal of Documentation, 17,* 1961, 106-110.

25. FARRADANE, J. The challenge of information retrieval. *Journal of Documentation, 17,* 1961, 233-244.

26. O'CONNOR, J. Review of: CLEVERDON, C.W. Aslib Cranfield Research Project: report on the first stage... and interim report on the test programme...[two early reports on the first Cranfield project, 1960]. *Journal of Documentation, 17,* 1961, 252-261.

27. VICKERY, B.C. Review of: FAIRTHORNE, R.A. Towards information retrieval, Butterworths, London, 1961. *Journal of Documentation, 18,* 1962, 91-94.

28. VASWANI, P.K.T. Review of: MARON, M.E. Automatic indexing: an experimental enquiry, Rand Corporation, Santa Monica, 1960. Also of: Advanced Information Systems Company, Report on the organization of large files with self-organising capability. *Journal of Documentation, 18,* 1962, 194-195.

29. FAIRTHORNE, R.A. Review of: Mathematical research related to information selection... U.S. National Bureau of Standards, 1961. *Journal of Documentation, 18,* 1962, 196-198.

30. FAIRTHORNE, R.A. Review of: TAUBE, M. Computers and commonsense: the myth of thinking machines, Columbia University Press, New York, 1961. *Journal of Documentation, 19,* 1963, 32-34.

31. MOTE, L.J.B. Review of: CLEVERDON, C.W. Aslib Cranfield Research Project: report on the testing and evaluation... [final Cranfield 1 report], 1962. *Journal of Documentation, 19,* 1963, 80-81.

32. ARTANDI, S. Mechanical indexing of proper nouns. *Journal of Documentation, 19,* 1963, 187-196.

33. NEEDHAM, R.M. *and* SPARCK JONES, K. Keywords and clumps: recent work on information retrieval at the Cambridge Language Research Unit. *Journal of Documentation, 20,* 1964, 5-15.

34. WILSON, C.W.J. Use of the Friden Flexowriter in the library of the Atomic Energy Research Establishment, Harwell. *Journal of Documentation, 20,* 1964, 16-24.

35. KYLE, B.R.F. Information retrieval and subject indexing: Cranfield and after. *Journal of Documentation, 20,* 1964, 55-69.

36. FREEMAN, R.R. Computers and classification systems. *Journal of Documentation, 20,* 1964, 137-143.

37. MARTYN, J. *and* SLATER, M. Tests on abstracts journals. *Journal of Documentation, 20,* 1964, 212-235.

38. MARTYN, J. Bibliographic coupling. *Journal of Documentation, 20,* 1964, 236.

39. HYSLOP, M.R. The ASM information retrieval system: after Cranfield. *Journal of Documentation, 21,* 1965, 27-42.

40. BROWNSON, H.L. Evaluation of document searching systems and procedures. *Journal of Documentation, 21,* 1965, 261-266.

41. FAIRTHORNE, R.A. Some basic comments on retrieval testing. *Journal of Documentation, 21,* 1965, 267-270.

42. ROLLING, L.N. A computer-aided information service for nuclear science and technology. *Journal of Documentation, 22,* 1966, 93-115.

43. LYNCH, M.F. Subject indexes and automatic document retrieval. *Journal of Documentation, 22,* 1966, 167-185.

44. SARACEVIC, T. *and* REES, A.M. Towards the identification and control of variables in information retrieval experiments. *Journal of Documentation, 23,* 1967, 7-19.

45. MARTYN, J. Tests on abstracts journals: coverage, overlap and indexing. *Journal of Documentation, 23,* 1967, 45-70.

46. BARHYDT, G.C. The effectiveness of non-user relevance assessments. *Journal of Documentation, 23,* 1967, 146-149.

47. CORBETT, L. Chemical Titles -- SDI trial at AWRE. *Journal of Documentation, 23,* 1967, 150-151.

48. CALESS, T.W. *and* KIRK, D.B. An application of UDC to machine searching. *Journal of Documentation, 23,* 1967, 208-215.

49. HINES, T.C. Computer manipulation of classification notations. *Journal of Documentation, 23,* 1967, 216-223.

50. HUTCHINS, W.J. Automatic document selection without indexing. *Journal of Documentation, 23,* 1967, 273-289.

51. CUADRA, C.A. *and* KATTER, R.V. Opening the black box of relevance. *Journal of Documentation, 23,* 1967, 291-303.

52. BROOKES, B.B. The measure of information retrieval effectiveness proposed by Swets. *Journal of Documentation, 24,* 1968, 41-54.

53. SHAW, T.N. *and* ROTHMAN, H. An experiment in indexing by word-choosing. *Journal of Documentation, 24,* 1968, 159-172.

54. GRALEWSKA, A. The design of an information system for the literature of rock mechanics. *Journal of Documentation, 24,* 1968, 197-209.

55. ROBERTSON, S.E. The parametric description of retrieval tests. *Journal of Documentation, 25,* 1969, 1-27 and 93-107.

56. ARTANDI, S. Computer indexing of medical articles -- Project MEDICO. *Journal of Documentation, 25,* 1969, 214-223.

57. LYNCH, M.F. Computer-aided production of printed alphabetical subject indexes. *Journal of Documentation, 25,* 1969, 244-252.

58. HALL, J.L., PALMER, J. *and* POOLE, J.B. An experimental current awareness service in the social sciences. *Journal of Documentation, 26,* 1970, 1-21.

59. CLEVERDON, C. Evaluation tests of information retrieval systems. *Journal of Documentation, 26,* 1970, 55-67.

60. SEARLE, R.H. Human *v.* machine selection for current awareness in mass spectrometry. *Journal of Documentation, 26,* 1970, 221-229.

61. McCASH, W.H. *and* CARMICHAEL, J.J. UDC user profiles as developed for a computer-based SDI service in the iron and steel industry. *Journal of Documentation, 26,* 1970, 295-312.

62. LANCASTER, F.W. Aftermath of an evaluation. *Journal of Documentation, 27,* 1971, 1-10.

63. SALTON, G. Automatic indexing using bibliographic citations. *Journal of Documentation, 27,* 1971, 98-110.

64. MILLER, W.L. The extension of users' literature awareness as a measure of retrieval performance, and its application to MEDLARS. *Journal of Documentation, 27,* 1971, 125-135.

65. DOLBY, J.L. Programming languages in mechanized documentation. *Journal of Documentation, 27,* 1971, 136-155.

66. HERSEY, D.F., FOSTER, W.R., STADLER, E.W. *and* CARLSON, W.T. Free text word retrieval and scientist indexing: performance profiles and cost. *Journal of Documentation, 27,* 1971, 167-183.

67. CLOUGH, C.R. *and* BRAMWELL, K.M. A single computer-based system for both current awareness and retrospective search: operating experience with ASSASSIN. *Journal of Documentation, 27,* 1971, 243-253.

68. MILLER, W.L. A probabilistic search strategy for MEDLARS. *Journal of Documentation, 27,* 1971, 254-266.

69. SPARCK JONES, K. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation, 28,* 1972, 11-21.

70. BARKER, F.H., VEAL, D.C. *and* WYATT, B.K. Comparative efficiency of searching titles, abstracts and index terms in a free-text database. *Journal of Documentation, 28,* 1972, 22-36.

71. GREEN, C.D. Some problems of indexing specialist material drawn from several disciplinary systems. *Journal of Documentation, 28,* 1972, 37-43.

72. BARKER, F.H., VEAL, D.C. *and* WYATT, B.K. Towards automatic profile construction. *Journal of Documentation, 28,* 1972, 44-55.

73. BROOKES, B.C. The Shannon model of IR systems. *Journal of Documentation, 28,* 1972, 160-162.

74. CLEVERDON, C.W. On the inverse relationship of recall and precision. *Journal of Documentation, 28,* 1972, 195-201.

75. HARLEY, A.J., LE MINOR, M. *and* WEIL, A. Towards a universal search formulation language for information retrieval. *Journal of Documentation, 28,* 1972, 202-213.

76. MARTIN, D. Computer typesetting and information. *Journal of Documentation, 28,* 1972, 247-261.

77. HEINE, M.H. The inverse relationship of recall and precision in terms of the Swets model. *Journal of Documentation, 29,* 1973, 81-84.

78. OLIVE, G., TERRY, J.E. *and* DATTA, S. Studies to compare retrieval using titles with that using index terms -- SDI from 'Nuclear Science Abstracts'. *Journal of Documentation, 29,* 1973, 169-191.

79. VAN RIJSBERGEN, C.J. *and* SPARCK JONES, K. A test for the separation of relevant and non-relevant documents in experimental retrieval collections. *Journal of Documentation, 29,* 1973, 251-257.

80. SALTON, G. *and* YANG, C.S. On the specification of term values in automatic indexing. *Journal of Documentation, 29,* 1973, 351-372.

81. AUSTIN, D. The development of PRECIS: a theoretical and technical history. *Journal of Documentation, 30,* 1974, 47-102.

82. BOOKSTEIN, A. The anomalous behaviour of precision in the Swets model, and its resolution. *Journal of Documentation, 30,* 1974, 374-380.

83. SPARCK JONES, K. Automatic indexing. *Journal of Documentation, 30,* 1974, 393-432.

84. LEGGATE, P. Computer-based current awareness services. *Journal of Documentation, 31,* 1975, 93-115.

85. SHARP, J.R. Natural language. *Journal of Documentation, 31*, 1975, 191-194.

86. JOLLEY, J.L. Punched feature cards. *Journal of Documentation, 31,* 1975, 199-215.

87. FIELD, B.J. Towards automatic indexing: automatic assignment of controlled-language indexing and classification from free indexing. *Journal of Documentation, 31,* 1975, 246-265.

88. HEINE, M.H. Measures of language effectiveness and the Swetsian hypothesis. *Journal of Documentation, 31,* 1975, 283-286.

89. SPARCK JONES, K. *and* VAN RIJSBERGEN, C.J. Information retrieval test collections. *Journal of Documentation, 32*, 1976, 59-75.

90. ODDY, R.N. Information retrieval through man-machine dialogue. *Journal of Documentation, 33,* 1977, 1-14.

91. VAN RIJSBERGEN, C.J. A theoretical basis for the use of co-occurrence data in information retrieval. *Journal of Documentation, 33,* 1977, 106-119.

92. ROBERTSON, S.E. Theories and models in information retrieval. *Journal of Documentation, 33,* 1977, 126-148.

93. BARRACLOUGH, E.D. On-line searching in information retrieval. *Journal of Documentation, 33,* 1977, 220-238.

94. KEEN, E.M. On the processing of printed subject index entries during searching. *Journal of Documentation, 33,* 1977, 266-276.

95. ROBERTSON, S.E. The probability ranking principle in IR. *Journal of Documentation, 33,* 1977, 294-304.

96. HUTCHINSON, T.P. An extension of the signal detection model of information retrieval. *Journal of Documentation, 34,* 1978, 51-54.

97. ROBERTSON, S.E. *and* BELKIN, N.J. Ranking in principle. *Journal of Documentation, 34,* 1978, 93-100.

98. HUTCHINS, W.J. Machine translation and machine-aided translation. *Journal of Documentation, 34,* 1978, 119-159.

99. HARPER, D.J. *and* VAN RIJSBERGEN, C.J. An evaluation of feedback in document retrieval using co-occurrence data. *Journal of Documentation, 34,* 1978, 189-216.

100. SALTON, G. Mathematics and information retrieval. *Journal of Documentation, 35,* 1979, 1-29.

101. SPARCK JONES, K. Search term relevance weighting given little relevance information. *Journal of Documentation, 35,* 1979, 30-48.

102. CROFT, W.B. *and* HARPER, D.J. Using probabilistic models of document retrieval without relevance information. *Journal of Documentation, 35,* 1979, 285-295.

103. DILLON, M. *and* DESPER, J. The use of automatic relevance feedback in Boolean retrieval systems. *Journal of Documentation, 36,* 1980, 197-208.

104. WU, H. *and* SALTON, G. The estimation of term relevance weights using relevance feedback. *Journal of Documentation, 37,* 1981, 194-214.

105. RADECKI, T. On a probabilistic approach to determining the similarity between Boolean search request formulations. *Journal of Documentation, 38,* 1982, 14-28.

106. CRAVEN, T.C. Editing of concept sub-networks with a microcomputer-based display system. *Journal of Documentation, 38,* 1982, 29-37.

107. BELKIN, N.J., ODDY, R.N. *and* BROOKS, H.M. ASK for information retrieval. *Journal of Documentation, 38,* 1982, 61-71 and 145-164.

108. INGWERSEN, P. Search procedures in the library -- analysed from the cognitive point of view. *Journal of Documentation, 38,* 1982, 165-191.

109. RADECKI, T. Reducing the perils of merging Boolean and weighted retrieval systems. *Journal of Documentation, 38,* 1982, 207-213.

110. BOOKSTEIN, A. Outline of a general probabilistic retrieval model. *Journal of Documentation, 39,* 1983, 63-72.

111. DILLON, M. *and* McDONALD, L.K. Fully automatic book indexing. *Journal of Documentation, 39,* 1983, 135-154.

112. VICKERS, P.V. Common problems of documentary information transfer, storage and retrieval in industrial organizations. *Journal of Documentation, 39,* 1983, 217-229.

113. SPARCK JONES, K. *and* TAIT, J.I. Automatic search term variant generation. *Journal of Documentation, 40,* 1984, 50-66.

114. GRIFFITHS, A., ROBINSON, L.A. *and* WILLETT, P. Hierarchical agglomerative clustering methods for automatic document classification. *Journal of Documentation, 40,* 1984, 175-205.

115. ROBERTSON, S.E. *and* HARDING, P. Probabilistic automatic indexing by learning from human indexers. *Journal of Documentation, 40,* 1984, 264-270.

116. ROBERTS, N. The pre-history of the information retrieval thesaurus. *Journal of Documentation, 40,* 1984, 271-285.

117. ENSER, P.G.B. Automatic classification of book material represented by back-of-the-book index. *Journal of Documentation, 41,* 1985, 135-155.

118. VICKERY, B.C. Knowledge representation: a brief review. *Journal of Documentation, 42,* 1986, 145-159.

119. ROBERTSON, S.E. On relevance weight estimation and query expansion. *Journal of Documentation, 42,* 1986, 182-188.

120. BERTRAND-GASTALDY, S. *and* DAVIDSON, C.H. Improved design of graphic displays in thesauri -- through technology and ergonomics. *Journal of Documentation, 42,* 1986, 225-251.

121. DANIELS, P.J. Cognitive models in information retrieval -- an evaluative review. *Journal of Documentation, 42,* 1986, 272-304.

122. VICKERY, A., BROOKS, H., ROBINSON, B. *and* VICKERY, B. A reference and referral system using expert system techniques. *Journal of Documentation, 43,* 1987, 1-23.

123. CRAVEN, T.C. An online index entry format based on multiple search terms. *Journal of Documentation, 43,* 1987, 46-53.

124. BOVEY, J.D. *and* BROWN, P.J. Interactive document display and its use in information retrieval. *Journal of Documentation, 43,* 1987, 125-137.

125. HANCOCK, M. Subject searching behaviour at the library catalogue and at the shelves: implications for online interactive catalogues. *Journal of Documentation, 43,* 1987, 303-321.

126. RADA, R., MILI, H., LETOURNEAU, G. *and* JOHNSTON, D. Creating and evaluating entry terms. *Journal of Documentation, 44,* 1988, 19-41.

127. THOMPSON, P. Subjective probability and information retrieval. *Journal of Documentation, 44,* 1988, 119-143.

128. BIRU, T., EL-HAMDOUCHI, A., REES, R.S. *and* WILLETT, P. Inclusion of term relevance in the term discrimination model. *Journal of Documentation, 45,* 1989, 85-109.

129. ELLIS, D. A behavioural approach to information retrieval system design. *Journal of Documentation, 45,* 1989, 171-212.

130. AKEROYD, J. Information seeking in online catalogues. *Journal of Documentation, 46,* 1990, 33-52.

131. KIM, Y.W. *and* KIM, J.H. A model of knowledge-based information retrieval with hierarchical concept graph. *Journal of Documentation, 46,* 1990, 113-136.

132. MONDSCHEIN, L.G. Selective dissemination of information(SDI: relationship to productivity in the corporate R&D environment. *Journal of Documentation, 46,* 1990, 137-145.

133. EFTHIMIADIS, E.N. Online searching aids: a review of front-ends, gateways and other interfaces. *Journal of Documentation, 46,* 1990, 218-262.

134. HANCOCK-BEAULIEU, M. Evaluating the impact of an online library catalogue on subject searching behaviour at the catalogue and at the shelves. *Journal of Documentation, 46,* 1990, 318-338.

135. CRAVEN, T.C. Condensed representation of sentences in graphic displays of text structures. *Journal of Documentation, 46,* 1990, 339-352.

136. ROBERTSON, S.E. On term selection for query expansion. *Journal of Documentation, 46,* 1990, 359-364.

137. KEEN, E.M. The use of term position devices in ranked output experiments. *Journal of Documentation, 47,* 1991, 1-22.

138. SANO, H. Concise display of index entries. *Journal of Documentation, 47,* 1991, 23-35.

139. RADA, R., BARLOW, J., POTHARST, J., ZANSTRA, P. *and* BIJSTRA, D. Document ranking using an enriched thesaurus. *Journal of Documentation, 47,* 1991, 240-253.

140. KIRCZ, J.G. Rhetorical structure of scientific articles: the case for argumentational analysis in information retrieval. *Journal of Documentation, 47,* 1991, 354-372.

141. ELLIS, D. The physical and cognitive paradigms in information retrieval research. *Journal of Documentation, 48,* 1992, 45-64.

142. VICKERY, B.C. *and* VICKERY, A. An application of language processing for a search interface. *Journal of Documentation, 48,* 1992, 255-275.

143. COUSINS, S.A. Enhancing subject access to OPACs: controlled vs. natural language. *Journal of Documentation, 48,* 1992, 291-309.

144. HANCOCK-BEAULIEU, M. *and* WALKER, S. An evaluation of automatic query expansion in an online library catalogue. *Journal of Documentation, 48,* 1992, 406-421.

145. VICKERY, B.C. *and* VICKERY, A. Online search interface design. *Journal of Documentation, 49,* 1993, 103-187.

146. LEE, J.H., KIM, M.H. *and* LEE, Y.J. Information retrieval based on conceptual distance in is-a hierarchies. *Journal of Documentation, 49,* 1993, 188-207.