

Chapter 2

Query expansion in a library catalogue

2.1 Introduction

This Chapter presents results on the use and non-use of automatic query expansion (AQE) by users in City University library. Further statistical analysis of some larger datasets is given in Chapter 3.

Automatic query expansion was first used in an Okapi system developed for the project reported by [17]. In that project automatic query expansion was evaluated in a controlled environment. The perceived helpfulness and general acceptability of the facility in this initial project provided a sound basis for further evaluation. The current project thus aimed to carry out a more extensive evaluation of query expansion in an operational library setting, using real users with real information needs.

2.1.1 Automatic query expansion

The idea behind AQE is that features of records which a user has judged relevant are likely to be useful in retrieving additional relevant records. The Okapi implementation uses the assumption that the higher the odds on a feature appearing in a record known to be relevant the higher the odds that the same feature will occur in relevant records which have not yet been seen by the user.

The present implementation is almost identical to that described in [17, pp 27–29]. It is summarized here. When the user answers “yes” to the relevance question (Figure A.10) the system extracts terms from certain

fields of the record and adds them to a pool of terms which may be used in a query expansion search.¹

When the user chooses the query expansion option the system calculates a weight for each term in the query expansion pool, using the following formula based on the $F4'$ formula given by Robertson and Sparck Jones in [11]:

$$\log((r + 0.5/R - r + 0.5)/(n - r + 0.5/N - n - R + r + 0.5)) \quad (2.1)$$

where N is the number of indexed records in the database

n is the number of postings for the term

R is the number of records chosen relevant

and r is the number of chosen records which contain the term.

The pool of terms is sorted in decreasing weight order, and the system starts at the top of the list and selects any term for which the user has not already seen all the records, until the term weight reaches zero or enough² terms have been selected or all the available terms have been selected. The selected terms are then used in a search, and the retrieved records ranked in descending weight order (the weight of a record is the sum of the weights of the terms by which it has been retrieved). The interaction is shown in Figures A.12 and A.13. In what follows a searcher's use of AQE is often referred to as the "MORE" option, because that is how it was described to users ("*Type More to look for books similar to the ones you have chosen*").

2.1.2 Installation of catalogue system

An Okapi version of the library catalogue was installed in the library at City. The system was also accessible remotely on the University's Ethernet local area network. Usage of the system in the library and over the network was monitored and tested using a combination of data gathering methods.

2.2 Background

2.2.1 Library database

The City University library database contained some 155,000 records at the time of this study. It is estimated that two-thirds are UK MARC records with LCSH and PRECIS subject fields and that the remaining third, which have been generated locally, have no subject headings.

¹The fields used are determined by parameters. During the experiment these were for the catalogue database *title*, *subject headings*, *classification code* and for INSPEC and LISA *title*, *feature headings*, *descriptors*.

²*enough* is a parameter, which was set to 24 throughout these experiments

2.2.2 Library setting

The Okapi system was installed in the library at the end of May 1990. A terminal was located alongside the CLSI OPAC terminals and was identified by a notice indicating that it was an experimental subject searching system available to all users. Unlike the CLSI system, Okapi did not provide loan information and catered for subject searching only. Further explanatory details about the project were included in leaflets placed next to the terminal and also online through the "information" option (Figure A.3). In keeping with the design philosophy that Okapi is usable at sight and does not require training, no specific instructions on its use were offered.

To distinguish between user sessions, searchers were invited to press a black key when they had completed their session. A message "Press Black key to quit" was displayed on all screens. Only a small minority of users left the terminal without using the black key. If users failed to exit, the system would automatically time out and return to the opening screen.

2.2.3 Network access

The Okapi library catalogue was made available over the University network to registered users. The system also provided access to the INSPEC and LISA databases. Registration forms were sent to selected departments (including Computer Science and Information Science), and were available in the library. Staff and postgraduate students accounted for the majority of registered users.

2.3 Evaluation methodology

The evaluation of the library catalogue system focused on the usability and retrieval effectiveness of the query expansion facility. A combined methodology was adopted whereby quantitative data was first collected through transaction log analysis, followed by more qualitative data provided by the replay of searches, user questionnaires and interviews.

2.3.1 Transaction logs

All search sessions on the library catalogue, which were carried out in the library and over the network, were logged automatically. The data from the transaction logs was analyzed using software developed for the project. In the six month period from June to November 1990 a total of 391 searchers used the Okapi catalogue system in the library. Eight hundred and fifty-eight user sessions consisting of 1,876 searches were analyzed.

The log analysis provided an initial overview of usage. The information derived from the logs also formed a basis for the design of the user questionnaires and interview schedules. In addition printed versions of the transaction logs were used to encourage users to comment on their search in post-search interviews.

2.3.2 Search replays

The logs were analysed to identify which searches had not used the query expansion facility. The searches which met the appropriate criteria were replayed to determine whether query expansion could have been beneficial. Searches were considered to be a reliable source of data for testing query expansion if they satisfied the following conditions:

1. the "MORE" option had been offered in the original search but not used,
2. the search topic was unambiguous, that is the search statement and books selected in the original search were obviously related, and
3. the experimenter could make relevance judgements with some degree of confidence.

A sample of 53 searches extracted from the transactions logs were replayed by the experimenter to the point at which the original search ended. Searches were then continued by using the "MORE" option based on the records actually selected by the searchers. Query expansion was considered to be useful only if at least one further item, judged to be relevant, appeared in the first screen of references (nine references) and had not been previously displayed.

Search replays were also used as part of post-search interviews for 12 searchers who had not used the query expansion facility when available. In that situation however, relevance judgements were made by the searchers themselves.

2.3.3 Pre-search and post-search questionnaires

Data was also collected directly from users in the library in the final phase of the project, from September to November 1990, using both pre-search and post-search questionnaires. The sample for this evaluation exercise consisted of 120 users who were asked to participate as they approached the Okapi terminal. The pre-search questionnaire established users' previous experience with computerized catalogues. The questionnaire was extended later in the test period (34 cases) to include questions about users' intentions or expectations in relation to the number of books they hoped to find.

In the course of the pre-search questionnaire, the experimenter asked users whether they would be willing to participate in a structured interview at the end of their search session. A brief post-search questionnaire was administered to the 75 cases who declined. The post-search questionnaire focused on the ease of use of the system and the searcher's satisfaction with the search results, including the use of the query expansion facility.

2.3.4 Post-search interviews

The 45 users who were willing to discuss their search session in greater depth were interviewed in a separate room away from the catalogue area where a second terminal and a printer were available. Participants were offered a printout of the references they had selected.

The interviews were divided into two categories according to whether or not the searcher had used the query expansion facility.

Searches with query expansion

If volunteers had used query expansion (the "MORE" option), the aim of the interview was to have the user talk about the level of satisfaction with the system in general and with the query expansion facility in particular. A printed transaction log of the search session was available to review the search and to encourage participants to comment on their search as fully as possible. The interviewer asked a set of questions for each subsequent search within a user session (Appendix).

Searches without query expansion

For searches where the "MORE" option was not used, the search was replayed with the user and query expansion was carried when made available. Users were asked why they had decided not to use the facility and were then asked for relevance judgements on additional references found.

2.3.5 Structured interviews of network catalogue users

Fifteen regular users of the Okapi system over the network were interviewed. The predominant use of Okapi over the network was to search databases other than the library catalogue (see Chapter 3). The catalogue appears to have been used primarily to check if the library had specific items users had found on one of the other databases.

2.4 Findings

2.4.1 Overall usage of query expansion facility

The query expansion facility was offered by the system only when searchers had chosen at least two items as relevant and so the “MORE” option did not appear on the screen until this condition had been satisfied. Table 2.1 sets out the overall usage of the query expansion option as shown in the transaction logs.

Table 2.1: Availability and usage of query expansion

QE option	Searches	%
Available and used	250	13
Available and not used	567	30
Not available	1059	57
Total	1876	100

The “MORE” option was thus available in 817 searches. It was utilized in 31% of the searches where it was offered.

In order to ascertain whether or not experience with the system had any influence on the use of the “MORE” option, a further analysis was undertaken correlating the number of search sessions carried out by users with the use of the query expansion facility (Table 2.2). It would seem that first time users were just as likely to use the query expansion facility as were more frequent users of the system. The option would thus appear to have been sufficiently self-explanatory and easy to use. At the same time greater experience with the system does not seem to have promoted the use of the query expansion facility.

Table 2.2: Usage of the “MORE” option and search sessions

No. of sessions	1	2-3	4-7	8-	Total
No. of users	217	117	45	12	391
No. of searches MORE avail.	190	280	205	142	817
No. of searches MORE used	70	71	55	54	250
% searches MORE used	37	25	27	38	31

2.4.2 Effectiveness of query expansion

The effectiveness of the query expansion option was measured in two ways. Firstly the number of items selected as relevant in searches with and without

query expansion were compared (Table 2.3). Secondly the number of items selected as relevant before and after query expansion were also compared (Table 2.4).

Table 2.3: Items selected with and without query expansion

	Searches	Items	Items/Search
QE not used	1626	1236	0.8
QE used	250	1302	5.2

Table 2.4: Items displayed and selected before and after query expansion

	No. displayed	No. selected	% selected
Before QE	2332	1236	53
After QE	1062	517	49

It would thus appear that searches where the query had been expanded led to a significantly higher number of items being selected as relevant overall. The items being selected were more or less equally distributed before and after the “MORE” option has been used. Table 2.5 however shows that 53% (133) of searches where the “MORE” option was used resulted in no further items being selected. In 29% of the searches two and more further items were selected.

Table 2.5: Items selected per search after query expansion

Selected	Searches	
0	133	(53%)
1	44	(18%)
2-3	29	(12%)
4-7	28	(11%)
8+	16	(6%)
	250	(100%)

2.4.3 Perceived usefulness of query expansion facility

The 120 users who participated in the post-search questionnaire and structured interviews, accounted for 127 queries. The “MORE” option was used in 58 queries. In 52 of the 58 additional useful items were found and 6 found no further useful items. 41 out of the 45 who did not use the “MORE” option when available, did not do so because they had found all the books they wanted already. Some of these searchers also indicated that they did

not know what the “MORE” option would do. The remaining four users felt that it would not have been helpful since there would seem to be no point in looking for more items similar to those already found if the items found were not what was required.

2.4.4 Search intentions

Only 34 of the subjects were asked how many items they hoped to find in searching the catalogue. Users didn’t intend to carry out exhaustive searches but half expected a bit more than the minimum of one or two books (Table 2.6). As discussed above, if users met their expectations in the system’s initial output, then there may have been no incentive to use the query expansion facility.

Table 2.6: Search intentions

Intention	No. of users
Find one or two books	12
Find a reasonable selection	14
Find as many books as possible	3
Find one particular item	4
None of the above	2
Total	34

2.4.5 Replay of searches with “MORE” option not used

In 30% of searches the “MORE” option was available but not used (Table 2.1). A number of these searches were replayed to ascertain whether or not query expansion would have been useful. 53 searches were selected from transaction logs for replay and relevance judgements were made by the experimenter. Another 12 searches were replayed as part of the post-search interview with the user providing the relevance judgements. Table 2.7 presents the result of the searches replayed to include query expansion. The “MORE” option appeared to find at least one potentially useful record in 53% of the searches, whereas 43% did not appear to find any extra useful records that had not already been seen by the user in the original search.

Eleven of the twelve replayed searches where the user provided the relevance judgements led to at least one useful item being found. Hence the query expansion facility was not used to its full potential.

Table 2.7: Query expansion in replayed searches

Items found	Searches	%
At least one useful	39	60
None useful	23	35
None	3	5
Total	65	100

2.4.6 Previous experience with online catalogues

The pre-search questionnaire established users' previous experience with online catalogues. 102 out of 120 users (85%) were first time users of Okapi, and only 18 (15%) had used it before. However 88% had used another online catalogue previously, mostly the library's CLSI system (84%), whilst 30% had used online catalogues in other libraries (Table 2.8). The fact that it was the beginning of term and the Okapi system was only one terminal available amongst a bank of eight CLSI terminals could account for the high percentage of first time users of Okapi.

Table 2.8: Users' experience of online catalogues

System	Users	%
CLSI	69	58
CLSI + other	32	27
Other	4	3
None	15	12
Total	120	100

2.4.7 Ease of use

Users were asked to rate the ease of use of the system. This was set on a scale of one to five with one being very difficult and five very easy (Table 2.9).

Table 2.9: Ease of use of Okapi

	1	2	3	4	5	Total
No. of users	0	0	13	64	43	120
%	0	0	11	53	36	

The 101 users who had also used the CLSI system (Table 2.8) were asked

to state which they found easier to use. Table 2.10 reveals that Okapi was rated as easier to use by 91%.

Table 2.10: Ease of use of Okapi compared to CLSI

Rating	No. of users	%
Equally easy	2	2
CLSI easier	4	4
OKAPI easier	92	91
Don't know	3	3
Total	101	100

Two of the four users who thought CLSI was easier to use qualified their choice by saying they were regular users of CLSI and more familiar with it. Another of the four said that Okapi was less easy to use because it did not cater for author searching.

2.4.8 Help facility

Okapi is designed to be used without any help facility. The 75 users who answered the post-search questionnaire were asked if they had required help. 70 (93%) were of the opinion that help was not necessary. Only four (5%) felt that help would have been useful and one user did not know. Problems encountered included one user who could not find how to do further searches without quitting and starting again. Another user said he wanted to use the "MORE" option but could not see it listed. The relevance question ("is this the sort of book you are looking for?") was not understood by one searcher. The labelling of the Return key with an arrow symbol on the keyboard caused some difficulty in one other case.

2.4.9 User satisfaction with search outcome

Searchers were asked to indicate their satisfaction with the search outcome on a scale of one to five with five being very satisfied (Table 2.11). 80% of users gave a positive response, 16% had some reservations and 4% were negative.

Table 2.11: User satisfaction

Satisfaction	1	2	3	4	5	Total
No. of users	3	3	19	51	44	120
Percent of users	2	2	16	43	37	100

2.4.10 Browsing references

Since Okapi displays search results in a ranked order there was some interest to find out to what extent users needed to browse through references in order to find useful items. Table 2.12 shows how searchers in the post-search questionnaire perceived their effort in finding useful references. Of the 75 users, two failed to find any useful references and one user carried out two searches, which accounts for the total of 74 searches.

Table 2.12: Perceived number of references looked at

Refs. looked at	Searches	%
“a few”	54	73
“quite a few”	5	7
“a lot”	15	20
Total	74	100

It would appear that most users did not feel that they had to browse through a large number of references. In some instances however there were some very high postings particularly for broad single term searches such as ‘accounting’. Although users actually browsed long lists of references they did not perceive that they had.