Chapter 3

Transaction log analysis

3.1 Introduction

3.1.1 Source of the data

This chapter presents, mainly in tabular form, the results of statistical analysis of the transaction logs gathered from use of the Okapi system described in Appendix A. The logs record sessions on the City University library catalogue and on a section of the INSPEC database from May 1990 to the end of June 1991. They are treated as three separate subsets:

- *CAT1*: library users identified by their library card numbers
- *CAT2*: library users who chose not to identify themselves
- *INSP*: registered INSPEC users

The logs forming CAT2 were gathered from a single terminal placed in the university library among the terminals to the university's integrated library system CLSI. CAT1 comprises logs of searches at the library terminal conducted by users who identified themselves and also logs of catalogue searches by people who had registered to use the system from anywhere on the campus network. These users also had access to INSPEC. It is not known how many of the searches in CAT1 were carried out in the library and how many from other terminals. The INSPEC searches were all done by registered users from terminals somewhere on the campus network (other than the library).

The sets comprise almost all the logs generated between May 1990 and end June 1991. Logs of searches performed by the experimenters have been removed as far as possible, although a few may remain in CAT2.

3.1.2 Sessions and searches

For the purposes of machine analysis sessions and searches are defined as follows.

- A *session* or user session starts when a user invokes a search input screen (Figure A.4) and finishes when the QUIT key is pressed or the system times out or becomes disconnected.
- A *search* starts when the system displays any of the search input screens (such as those shown in Figures A.4 and A.20) and finishes when the session finishes or when the user requests a new or edited search.
  - A search is *aborted* if the user pressed the QUIT key or allowed the system to time out or chose the "new" or "edit" options before the system displayed a "result" (Figure A.6) for the search. This also covers the case where none of the terms in the search was found during the lookup and the user chose not to replace them (Figure A.7) so that the search was effectively empty.
  - A search *failed* if it was not aborted but retrieved no records.

The above definitions of sessions and searches are artificial, and are used simply because they enable sessions and searches to be algorithmically identified. Machine analysis has the advantage of objectivity, but sessions as defined are often far from being a natural unit of search activity. A more meaningful definition would be something along the lines of "a session is a sequence of searches by a single user for items on a single topic or closely related topics". But looking at logs shows that it is not at all unusual for a user to do one or more searches on a subject, then some more on another subject which is apparently not closely related, followed by reversion to the original topic. Nor is it possible in general to decide on the degree of relationship between two searches, or even what we mean by "closely related".

Recognition of session boundaries is another problem. Although Okapi users had to go through a kind of login procedure it was not unusual for a user to logout and reconnect almost immediately, rather than use the "new search" or "edit search" commands. Probably we should have amalgamated "sessions" known to be by the same user when there was no more than a few minutes time lapse between them.

3.1.3 Transaction logs

Virtually all sessions were automatically logged (it is of course possible to run the system without logging). Some details of log layout changed during the course of the experiments, but the general style and layout have remained

unchanged since 1984 [7]. The logs record all user keystrokes, all record displays more or less as the user saw them, and enough coded information about the state of the system to enable the interaction to be completely reconstructed. An annotated example is shown in Appendix B. There is one log for each user session.

The logs are designed rather more for human readability than for automatic analysis and it is quite difficult to write reliable analysis programs. These programs are likely to fail on logs which are defective, usually because of a data communication breakdown or because the host machine crashed. The programs used for the analysis in this chapter were written in the unix *awk* language, and were verified by hand on a reasonable sample of logs. Analysis took place in two stages. The first stage digested a log into a few summary lines. Figure 3.1 shows an example.

Figure 3.1: Transaction log summary

```
User: 098920007
File: userlog.0002.Z
Database: city
Start_time: 900903142242
Search 1 (first)
accounting management control<RET>
Lookup: account (1622), management (6036), control (2841),
Terms: 3
Result orig: OK 535 60 26
Result more: OK 157 157 0
Orig: Bfs 90 Fls 3 Chos 2 Band1 2 Band2 0 Band3 0 Mores 1 Clss 0
Mores: Bfs 54 Fls 1 Chos 1 Band1 0 Band2 0 Band3 1
End of session. Time 395
```

This user typed in the search "accounting management control". The system stemmed "accounting" to "account". The next line gives the number of postings for each of the three terms. The system reported "26 books match your search well — (535 books found altogether)" (cf. Figure A.6). The user looked at the first 90 records in brief (10 screens), three in full, and chose two as relevant. Both the chosen records were among the first 26, those which matched the search well. The user did one iteration of query expansion, displaying 54 brief records and looking at and choosing one, which was not from the first two screens of records.

These summaries identify time, user and database, search statements, term lookup results and overall search result. For each search the summary shows the number of brief and full records which the user saw, followed by the number chosen relevant. There is some information about the source of the records chosen as relevant. In the case of an "original" search (one entered by the user), chosen records are divided into those which contain all the terms of the search, those which match the search fairly well but not perfectly and those which do not match the search closely. In the case of a

query expansion search ("More") records are divided into those chosen from the first screen, the second screen and subsequent screens.

3.2 Sessions and searches

Table 3.1 gives summary information about the numbers of users, sessions and searches in each of the three datasets.

Table 3.1: Sessions and searches

	CAT1		CAT2		INSP	
Users	705		N/K		86	
Sessions	2212		1401		562	
Searches	4952		2574		1763	
– aborted	462	(9.3%)	211	(8.2%)	126	(7.1%)
– failed	201	(4.1%)	117	(4.5%)	36	(2.0%)
– no recs disp	458	(9.2%)	204	(7.9%)	160	(9.1%)
– records disp	3830	(77.3%)	2042	(79.3%)	1441	(81.7%)
Mean srches/session	2.03		1.69		2.91	

Searches/session is calculated on non-aborted searches only.
No display means that the search found some records but the user did not choose to see any of them.

There are some noticeable differences between the three sets. In particular, the number of *searches per session* differs quite markedly (significantly at 0.1% on a t-test) between each pair of datasets, with INSPEC having the highest number and unidentified catalogue users (CAT2) the lowest. It will be seen later that CAT2 searches also had a much lower usage of query expansion than the other two sets (Tables 3.7 and 3.8). This is discussed further 3.6.

3.2.1 Aborted searches

The proportion of aborted searches looks very high, but this is mainly due to our definition. In up to three-quarters of aborted searches the user had not in fact entered a search at all but had pressed the QUIT key or left the system to time out from a search input screen. In most other cases the aborted search would count as a failed search on systems which do not compel the user to negotiate terms which it cannot find (Figure A.7).

3.2.2 Failed searches

The proportion of (non-aborted) searches which failed to retrieve any records is very low indeed (3.9% overall), and significantly lower for INSPEC than for the catalogue sets. Some keyword-type online catalogues have returned failure rates of up to 40%. Some of this difference is accounted for by our definitions of *failed* and *aborted*, but in the main it is due to

- Okapi's "best match" search
- stemming and automatic cross-referencing.

These features were not under investigation in the present project — see [16] for an account of the relevant experiments. However, Table 3.2 suggests that the overall failure rate might well have been over 30% without these devices.

Table 3.2: Distribution of number of records retrieved

Search result	CAT1	CAT2	INSP
Exact match:			
"high" (more than 180 recs)	12.3	13.6	16.7
"mid" (10–180 records)	30.9	30.7	40.1
"low" (1–9 records)	34.2	32.7	28.6
Subtotal	77.4	77.0	85.4
Fairly good match	8.6	8.2	9.0
Poor match	14.1	14.8	5.6
Total searches	4288	2246	1601

The figures are percentages of the number of searches which retrieved some records.
Exact match means records which would have been retrieved by a boolean AND — they contain all the terms in the search.
Fairly good match generally means records attaining about 2/3 of the weight they would have if they contained all the terms of the search.
Poor match generally means records attaining between 1/2 and 2/3 of the weight they would have had if they contained all the terms of the search.

3.3 Records retrieved, displayed and chosen

From here on failed and aborted searches are excluded from the datasets.

3.3.1 Numbers of records retrieved

Table 3.2 shows the distribution of the number of records retrieved for each of the three datasets. The two catalogue datasets give very similar figures, but INSPEC searches have a marked tendency to retrieve more records.

3.3.2 Effect of number of records reported on record display

Overall, in about 10% of "successful" searches (i.e. searches which retrieved some records), the user chose not to look at any records, even in a brief display (this figure is derived from Table 3.1). The Okapi system did not automatically go into record display on retrieval as it was felt that the user should have a picture of the overall results of the search (Figure A.6). One would expect some searchers to be put off if the system reports that a large number of records match.[1]

Table 3.3 shows the proportion of searches in which the user did not display any records. It shows that a "high" result is moderately discouraging.

Table 3.3: Percent of non-failed searches in which no records were displayed

Search result	CAT1 (4288 srches)	CAT2 (2246 srches)	INSP (1601 srches)
"high"	37.6	24.6	35.0
"mid"	9.1	9.2	5.0
"low"	3.1	3.3	3.3
Fairly good match	4.4	1.1	3.5
Poor match	12.8	12.0	15.7
Overall	10.7	9.1	10.0

The figures are percentages of the number of searches in each category.
(The actual numbers of searches are implicit in Table 3.2.)

There is little difference between the datasets. The system itself does not discourage users from looking at what might be thought excessive results. The Okapi system used over the local ethernet usually gave very rapid updating of screen displays and it did not seem tedious to the experimenters to look at a great many screens of brief records. It appears that experienced users are more likely than inexperienced to display records from large sets: for the CAT1 dataset 56.1% of "high" searches by first time users led to a display, but this increased to 70% of "high" searches by users with eight or more sessions (the overall CAT1 figure is 62.1%).

Further analysis shows that the proportion of searches in which no records are displayed increases smoothly to reach about 43% when the system shows that more than 1000 records match the search.

[1] On the other hand it was not unusual to see a case where a user seemed content to browse 40 or more screens of brief records. In particular this was observed several times in catalogue searches for "accounting" or "accountancy"; the user gave the impression of trying to get an overall picture of the library's holdings.

Action following a search where the user displayed no records

A limited further investigation of these cases was carried out by visual examination of the log summaries in 118 cases taken successively from the CAT1 set. They were classified as shown in Table 3.4.

Table 3.4: Action following search with no records displayed (CAT1 set)

Action	Cases		Examples
Narrow	61	(51.7%)	"PhD" → "How to get a PhD"
Abandon(1)	17	(14.4%)	"civil engineering" (579 recs)
Abandon(2)	13	(11.0%)	"o", "ew", "ESC O P"
Different search	11	(9.3%)	"endocrine disorders" → "diabetes"
Rephrase	8	(6.8%)	"aids" → "hiv aids"
Broaden	4	(3.4%)	"pediatric larynx" → "larynx"
Repeat	4	(3.4%)	
Total	118	(100.0%)	

Abandon(1) means abandon following a "sensible" (but usually highly posted) search. *Abandon(2)* means abandon following accidental or foolish input. This was sometimes caused by program bugs or data communication problems.

3.3.3 Full records

From now on the datasets will be limited to those searches in which the user actually looked at some brief records. This eliminates about 10% of the searches (Table 3.3). Table 3.5 gives the distributions of the number of full records displayed, and the number of records chosen relevant, for each of the reduced datasets. In about a third of these searches no full records were displayed. Table 3.6 gives the corresponding figures restricted to searches in which at least one full record *was* seen.

As Table 3.5 shows, in about a third of the searches no records at all were displayed in full. In the case of the catalogue database the brief records did not show the location of the item (Figure A.8) and for INSPEC they did not show the source. Hence searches in which no full records are displayed cannot be called successful if the object of a bibliographic search is to locate bibliographic items. However, examination of the sessions in which these searches occur, and sometimes of other neighbouring sessions by the same user, reveals that in a large proportion of cases the search is either preceded or followed by an identical or similar search in which full records *are* displayed.

INSPEC users have a marked tendency to display more full records than catalogue users. Part of this difference may accounted for by the fact that an appreciable proportion of catalogue searches, but few INSPEC searches,

Table 3.5: Full records displayed and chosen: searches in which some brief records were displayed

	CAT1		CAT2		INSP	
Records	Disp	Chos	Disp	Chos	Disp	Chos
None	34.2	49.7	33.5	50.9	31.7	57.7
1	28.0	23.7	28.1	25.2	21.2	14.1
2–3	16.9	13.2	18.7	13.1	18.7	11.9
4–7	11.9	8.7	11.7	6.8	13.5	8.3
8–	9.0	5.0	8.0	3.9	14.9	8.0
Total searches	3830		2042		1441	
Mean records per search	2.9	1.7	2.5	1.5	4.1	2.2

The figures are percentages of all searches in which some brief records were displayed.

Table 3.6: Full records displayed and chosen: searches in which some full records were displayed

	CAT1		CAT2		INSP	
Records	Disp	Chos	Disp	Chos	Disp	Chos
None	0.0	23.5	0.0	26.1	0.0	38.1
1	42.5	36.0	42.2	37.9	31.1	20.6
2–3	25.7	20.0	28.0	19.6	27.5	17.5
4–7	18.2	13.1	17.7	10.9	19.9	12.2
8–	13.7	7.4	13.3	6.1	21.5	11.5
Total searches	2520		1358		984	
Mean records per search	4.3	2.7	3.7	2.2	5.9	3.2

The figures are percentages of all searches in which some full records were displayed.

are attempts to locate a specific work, despite the notice on the library terminal emphasising that the system is for subject searching only and the information on the search input screen (Figure A.4).

3.3.4 Relevance judgements

On displaying a record in full the user is compelled to answer a form of relevance question ("Is this the sort of thing you are looking for?"). The object, for the system, is to obtain relevance information for use in query expansion. The motivation, for the user, is less obvious. There is at this point in the interaction some information available on request which informs users that if they choose some books the system will look for similar ones if they type MORE. We have not tried to find out what proportion of users looked at this information — certainly some did. Many users will also have learnt that the system allowed them to review or to print out (or have

emailed to them) a list of the items they had chosen (Figure A.18). The fact that in more than a quarter of the searches in which full records were displayed the user chose no records at all (Table 3.6) is not in itself an indication that the user did not find any useful records.

3.4 Query expansion

One of the objects of these experiments was to evaluate the automatic query expansion feature of the Okapi system (2.1.1). Before query expansion can be done it is necessary for the user to provide some feedback in the form of records chosen as relevant. The interaction is shown in Figures A.11 and A.12. Query expansion is not available unless at least one record has been displayed in full, and the user has given an affirmative answer to the relevance question.[2]

From here on searches in which no full records are displayed are excluded from the analysis.

3.4.1 Takeup of query expansion option

Table 3.7 shows the proportion of searches in which query expansion was used, broken down by the number of records reported as retrieved (cf Table 3.2), given that at least one record had been chosen relevant.

The table shows, perhaps not surprisingly, that query expansion is much more likely to be used in searches where few records were retrieved or where none of the retrieved records contained all the terms of the search.

The overall figures for the takeup of query expansion are not significantly different between CAT1 and INSPEC, but CAT2 shows a very significantly lower takeup than the other two. We have been unable to account for this difference. There is further discussion of this point in 3.6.

Table 3.8 shows the proportion of searches in which query expansion was used by the number of records already chosen relevant by the user.

Here, we find that query expansion is less likely to be used when the user has only chosen one record. Many of these cases will be those where the user is not seeking an exhaustive search. It is also less likely to be used when the user has already chosen a large number of records.

[2]Query expansion was not always offered unless two or three records had been chosen relevant, although it was almost always available if the user knew the command. Whether it was offered depended on the number of records which matched the search well and on how many had been chosen.

Table 3.7: Takeup of query expansion by search result

Search result	CAT1	CAT2	INSP
	(% of srches using QE)		
Exact match:			
"high" (more than 180 recs)	20.0	22.7	22.6
"mid" (10–180 records)	26.2	20.9	30.5
"low" (1–9 records)	34.8	26.6	40.8
Fairly good match	50.7	48.7	44.7
Poor match	53.7	48.0	71.4
Overall	33.5	27.5	35.3
Number of searches	1925	1002	609

The figures are percentages of the number of searches with the given result.

Table 3.8: Takeup of query expansion by number of records chosen

Number of recs chosen	CAT1	CAT2	INSP
	(% of srches using QE)		
1	21.8	16.3	22.8
2–3	46.2	39.6	41.1
4–7	49.7	46.0	50.5
8–	30.1	30.0	34.5
Overall	33.5	27.5	35.3
Number of searches	1928	1003	609

The figures are percentages of the number of searches in which the given number of records had been chosen relevant before query expansion.

3.4.2 Results of query expansion searches

In some of the searches in which query expansion was used it was used more than once, with a mean of about 1.3 expansion searches per original search (*original search*: search typed in by a user). To arrive at the figures in Table 3.9 all the expansion searches in an original search have been aggregated.

About 95% of query expansion searches were successful in the sense that they led to the system finding some records. There does not seem to be an obvious way of measuring the "nearness" of records retrieved by query expansion to the records used for relevance feedback. Hence the system was given a fairly arbitrary, and low, cutoff for the minimum weight to be achieved by records in an expansion search. However, as Table 3.9 shows, in about 40% of cases no full records were displayed and in about 55% no records were chosen relevant.

Table 3.9 gives the distributions of the number of full records displayed and the number of records chosen from record sets retrieved by query ex-

pansion. This should be compared with Table 3.5, which gives the same distributions for the totality of records retrieved in all the searches in which the user displayed some brief records.

Table 3.9: Records displayed and chosen from query expansion

Number of records	CAT1		CAT2		INSP	
	Disp	Chos	Disp	Chos	Disp	Chos
	(Percent of searches)					
0	42.6	58.4	40.9	54.3	32.6	50.2
1	17.6	15.5	16.3	13.4	17.7	17.7
2–3	13.9	11.3	15.6	13.4	20.0	13.5
4–7	13.2	9.1	14.1	13.0	14.0	12.6
8–	12.7	5.7	13.0	5.8	15.8	6.0
Total searches	646		276		215	
Mean records per search	3.3	1.8	3.3	2.1	3.8	2.2

The proportions of searches in which no records were displayed in full, and in which no records were chosen are somewhat higher than the overall figures shown in Table 3.5. However, the means are also higher. This suggests that query expansion searches are more likely than original searches to produce no useful records, but that in the 40%–50% of cases where they *do* produce useful records the query expansion record lists are a richer source than the original lists.

3.4.3 Source of records displayed and chosen

Tables 3.10 and 3.11 compare searches with and without query expansion with respect to the sources of displayed and chosen records and the proportion of records chosen. *Bands* 1, 2 and 3 refer to the source of chosen records. For an original search, a record is from Band 1 if it contains all the terms of the search, from Band 2 if it matches fairly well and Band 3 otherwise. For query expansion searches Bands 1 and 2 refer to the first and second screens of brief records and Band 3 to any subsequent screen.

3.4.4 How useful was query expansion?

(In this discussion the datasets are restricted to searches in which at least one record was chosen relevant.) Query expansion made a substantial contribution to the choice of relevant records. It was used in 33.5% of CAT1 searches and in 35.3% of INSPEC searches Table 3.7. For reasons unknown, it was less frequently used (27.5% of searches) in CAT2. Its use was more frequent in searches where few or no records matched the original search exactly. In CAT1 QE was responsible for 37.5% of the records chosen as

Table 3.10: Sources of displayed and chosen records: CAT1

	QE not used	QE used			Overall		
		Orig	QE	Tot	Orig	QE	Tot
	(means per search)						
Brief recs	34.8	30.6	58.6	89.2	33.4	19.8	53.2
Full recs	3.8	4.3	3.3	7.6	3.9	1.1	5.1
Chosen: (total)	2.8	3.0	1.8	4.8	2.9	0.6	3.5
-band 1	2.2	2.1	0.8	2.9	2.2	0.3	2.5
-band 2	0.3	0.3	0.3	0.6	0.3	0.1	0.4
-band 3	0.3	0.6	0.7	1.3	0.4	0.2	0.6
% of bfs chosen	8.0	9.8	3.0	5.4	8.6	3.1	6.5
% of fulls chosen	74.3	70.9	53.4	63.2	73.1	53.5	68.7
Searches	1270			637			1917

The searches are the subset of CAT1 consisting of all those where at least one record was chosen.

relevant in searches in which it was used, and for 17.1% over all searches in which any records were chosen (Table 3.10). The corresponding figures for INSPEC are 31.9% and 15.4% (Table 3.11).

QE was in some sense less efficient that the user searches in finding relevant records. In QE searches more than three times as many brief records were scanned for each record chosen than in original (user-entered) searches (Tables 3.10 and 3.11). A portion of this difference must be put down to the fact that by the time searchers use query expansion they have already chosen some records, and hence their need for more records must in general be less than it was in the early stages of the search — they can afford to be more choosy. There is some rather tenuous evidence for this in the fact that in QE searches users also chose a somewhat lower proportion of displayed records than in original searches, suggesting that they were less ready to answer "yes" to the relevance question.[3]

Note on ranking

Records displayed following query expansion come out in decreasing weight order, where the weight of a record is the sum of the weights of the terms by which it has been retrieved. It is worth noting that for the CAT1 set 44% of the records chosen relevant are from the first screen (nine records), 17% from the second screen and remaining 39% from subsequent screens

[3] But another hypothesis might be that these people were so desperate for records that they exercised *less* discrimination in selecting brief records for display, and so were more likely to reject the full record. Or they might just have been uncertain about what they were looking for.

Table 3.11: Sources of displayed and chosen records: INSPEC

	QE not used	QE used			Overall		
		Orig	QE	Tot	Orig	QE	Tot
	(means per search)						
Brief recs	51.5	47.3	80.3	127.6	50.0	28.3	78.3
Full recs	6.5	7.6	3.8	11.4	6.9	1.3	8.2
Chosen: (total)	4.3	4.7	2.2	6.9	4.4	0.8	5.2
-band 1	3.6	3.8	1.0	4.8	3.7	0.3	4.0
-band 2	0.4	0.5	0.3	0.8	0.4	0.1	0.6
-band 3	0.3	0.4	0.9	1.3	0.3	0.3	0.6
% of bfs chosen	8.3	10.0	2.7	5.4	8.9	2.7	6.6
% of fulls chosen	65.6	62.2	56.9	60.4	64.3	56.9	63.1
Searches	394			1214			608

The searches are the subset of INSP consisting of all searches where at least one record was chosen.

(Table 3.10). The corresponding figures for INSPEC are 45%, 14% and 41% respectively.

3.5 System use patterns

Tables 3.12 and 3.13 give the distribution of the number of sessions by identified users over the four month period Jan–Apr 1991. Table 3.12 includes everyone who searched the given databases during the period and Table 3.13 limits the figures to those who had searched the databases at least once prior to Jan 1 1991 and therefore excludes new users who arrived during the period under observation. The latter table also gives the number of users who did *not* use the system again during the period. The tables show that most catalogue users only used the system once or twice, but there was a small number of people who used it regularly, or at least intensively. Among INSPEC users, who had all been to the trouble to register for use of the system, there are fewer occasional users, and a nucleus of regular users who provided the data for the work described in Chapter 5.

3.6 Differences between classes of users

In previous sections some significant differences show up between catalogue users who chose to remain anonymous and other users. In particular, the former made fewer searches per session and were less likely to use query expansion. One would expect to find differences between people using a

Table 3.12: System use Jan–Apr 1991: all identified users

Catalogue		INSPEC	
Sessions	Users	Sessions	Users
1	137	1	14
2–3	68	2–3	18
4–7	34	4–7	9
8–15	12	8–15	5
16–	6	16–	3
717	257	220	81

Table 3.13: System use Jan–Apr 1991: people who had used system before Jan 1 1991

Catalogue		INSPEC	
Sessions	Users	Sessions	Users
0	400	0	32
1	33	1	2
2–3	19	2–3	11
4–7	16	4–7	8
8–15	7	8–15	4
16–	5	16–	3
373	490	176	60

system from an open terminal and those using it from their own desks. In the rather small open access library of City University it is easy to get up from the terminal and go and look at the shelves, and this alone might be expected to lead to shorter and less exhaustive sessions. However, a comparison between catalogue use by identified users at the terminal in the library and at their own terminals showed no significant difference. One would also expect considerable differences between INSPEC and catalogue searching. However, the latter is less pronounced than the difference between identified and unidentified users at the library terminal. There does seem to be some "real" difference between anonymous and identified catalogue users, whether the latter are using the system from their own terminals or in the library.

3.7 Best match searching

These experiments were not intended to evaluate best match, ranked output searching. This was done several years ago in the experiments described in

[16]. It is, however, an essential feature in any system which has an automatic query expansion facility. AQE can only be done on such a system[4].

Table 3.2 shows that of searches which retrieved some records, 23% of catalogue searches and 14.6% of INSPEC searches would have failed completely on a system which required all the search terms to be present. Further, Table 3.10 shows that about 24% of the total records chosen from original searches (i.e. prior to query expansion) were retrieved by only some of the terms in the search (CAT1 set). The corresponding figure for INSPEC, where searches tended to retrieve more records (Table 3.2), is 16%. The proportions when highly posted searches are excluded are very much higher.

[4]In theory, although not always in practice, "best match" can be provided in a front end to a conventional boolean system — see, for example, [12]