Think local, search global? Comparing search engines for searching geographically specific information.

Alastair G Smith
School of Information Management
Victoria University of Wellington

Abstract

This study evaluates the retrieval of New Zealand information using three local New Zealand search engines, four major global search engines and three metasearch engines. Searches for NZ topics were carried out on all the search engines, and the relative recall calculated. The local search engines did not achieve higher recall than the global search engines or metasearch engines, but no search engine achieved more than 45% recall. Despite the theoretical advantage of searching the databases of several individual search engines, metasearch engines did not achieve higher recall. 36% of relevant pages for the queries were outside the .nz domain. Implications for searching for geographically specific information, and for evaluation of search engines, are discussed.

Keywords: search engines, metasearch engines, evaluation, recall, overlap, geographic areas, New Zealand

---

Introduction

Search engines have become the paramount method of searching for information on the Internet (Kehoe et al. 1999). Users who are looking for information about a specific geographic area, such as New Zealand, can choose between using a local search engine which covers their specific area, or a global search engine which attempts to include the entire Web. A search for information about New Zealand, for instance, can be carried out on global search engines such as AllTheWeb or Google, or on a local search engine such as SearchNZ (http://www.searchnz.co.nz/), SearchNow, (http://www.searchnow.co.nz/), or NZ Explorer (http://nzexplorer.co.nz/). This article reports a study of the performance of both global and local search engines for a set of queries relating to New Zealand. As well as providing guidance for searchers for geographically specific information, the implications for the design and operation of search engines are examined. While this study relates to New Zealand information, it has implications for searching for information about other geographical areas. For example Canada, Britain, Malaysia, and Australia have their own geographically specific search engines.

There have been a number of studies of search engine performance, reviewed by Oppenheim *et al* (Oppenheim et al. 2000). In general, studies have found that recall (the proportion of available relevant sites found by a search) is low; and that the overlap between search sets recovered by different search engines is minimal ((Bar-Ilan 1998),

(Clarke & Willett 1997), (Gordon & Pathak 1999), (Hsieh-Yee 1998), (Leighton & Srivastava 1999), (Dresel et al. 2001), (Olvera Lobo 2000)).

A traditional measure of search engine effectiveness has been recall. Using recall as an evaluation measure for Internet search engines poses two difficulties:

1. The number of available relevant documents is impossible to measure, short of visiting every page on the Web.
2. Search engines retrieve large sets of documents, but present them in a ranked order. Users tend to consult only the highest ranked documents, so the question is not whether the search engine retrieves relevant documents, but whether they are ranked high enough to be seen by the user.

The first problem can be overcome to a degree by the use of "relative recall" (Clarke & Willett 1997) where the relevant documents retrieved by a number of search engines are pooled, and the recall for any one search engine is regarded as the proportion of the pool retrieved by that particular search engine. While there are limitations to this approach (Fricke 1998), it provides a practical way of comparing the recall of a number of search engines, particularly if a significant number of search engines is compared.

The second problem can be addressed by considering only documents ranked in the first 10 or 20 of the retrieved set. Clearly this doesn't reflect all the relevant documents retrieved by the search engine, but is a realistic measure of those documents likely to be found by a user.

Precision (the proportion of relevant documents in a retrieved set) is also a traditional measure of search effectiveness. However if the relevant documents in the first *n* relevancy ranked documents is used to compare different search engines, precision and recall are proportional (except for the rare event that less than *n* documents are retrieved), and using both measures is not necessary. (For example, if search engine A retrieves *a* relevant documents in the top 20, and there are c possible relevant documents to be found, then precision is a/20 and recall is a/c. If search engine B retrieves *b* relevant documents in the top 20, for search engine B precision is b/20 and recall is b/c, and the ratio of both precision and recall for A to B is *a/b*).

For the purposes of this study, "search engines" are defined as search tools that use an automated crawler to index words on web pages, thus allowing full text searching of the Web. These differ from directories such as Yahoo! where human intervention is used to categorise websites. Metasearch engines, which forward search terms to a number of different search engines and collate the results, are attractive due to the lack of overlap between single search engines; by using a metasearch engine a single search can in theory cover a greater part of the Web (Tomaiuolo 1999), (Jacso 2001). The current study included three metasearch engines, to see how well they handled New Zealand queries, and how they compared with the single search engines.

In many cases a global search engine is available at a local URL. For example Google is available at the URL http://www.google.co.nz/ as well as at http://www.google.com/. The

local URL usually searches the same database as the international one, but provides an option for restricting the search to local sites.

Methodology

In September 2002, a set of 10 queries relating to NZ topics were searched on the search engines under study. The search engines studied were:

Four global search engines:

- AlltheWeb/FAST http://www.alltheweb.com/
- Google http://www.google.co.nz/
- HotBot http://hotbot.lycos.com/
- Altavista http://nz.altavista.com/

Three local search engines:

- SearchNZ http://www.searchnz.co.nz/
- SearchNow http://www.searchnow.co.nz
- NZExplorer http://nzexplorer.co.nz/

Three metasearch engines:

- Excite http://www.excite.com/
- Vivisimo http://vivisimo.com/
- Surfwax http://www.surfwax.com/

The global and metasearch engines were chosen because they are widely used and performed comparatively well in pilot searches. The metasearch engines are ones that provide a single combined list of results: it would be difficult to identify the "top 20 hits" using a metasearch engine such as Dogpile which presents separate lists from each search engine. The three local search engines were the only ones available at the time of the research.

The queries used for the study were derived from queries submitted to an Australian search tool, OzSearch (http://www.ozsearch.com.au/), which at the time provided a webpage listing recent queries. These queries were adapted to NZ topics, since no equivalent service was available for New Zealand.

The queries used are listed in appendix 1.

For each search, the URLs of relevant web pages in the first 20 hits were noted. In carrying out searches, the best search was carried out that could be done with the functionality of the specific search engine. Searches were generally not restricted to the .nz domain except when this gave the best results, for instance when it was likely that relevant sites would be in the .nz domain, and defining "New Zealand" was otherwise difficult (for example in the queries for floristry schools and truck rental companies). Searching and judgement of relevance was carried out by the researcher. While there is an argument for using real queries and searchers, and judgement of relevance by actual requestors, in this case it was the relative performance of the different search engines that was being investigated, rather than an absolute measure of performance.

An issue in judging relevance in the web environment is whether a page is relevant if it links to a page with the desired information. In the current study, the information had to be on the indexed page to be judged relevant, but it is arguable that the relevant page might be found by following links from the pages produced by the search engines. Exploring the possible links, however, would be very time consuming, and lead to difficult relevance decisions.

In some cases (for example the *Flags of the World* page for the Maori Flag, http://fotw.vexillum.com/flags/nz_mao.html) a relevant page might have a number of mirror sites. In these cases only one page was counted as relevant. Similarly, a site may have a number of interlinked pages containing the relevant information (for instance a search engine might show a number of different pages from the same truck rental company); only one page from a site was counted in these cases.

One factor that clearly influences retrieval is the size of the search engine database. Do global search engines index as many pages in the New Zealand domain as local search engines do? To test the relative coverage by the search engines of the .nz domain, 40 searches on single words were carried out on each of the search engines, restricted to the .nz domain. Half of the words were general (e.g. hubble, cheats, sulphur…) and half were NZ specific (e.g. *Waitangi* – the location where an important treaty was signed between Maori and British government in 1840; *Tangi* – a Maori funeral; *sharemilker* – a dairy farmer who runs cows on land owned by another). The total hits across all the words for each search engine gave an indication of the coverage of database for the New Zealand domain. This is based on a method used by Notess (Notess 2002). This exercise was not attempted for the metasearch engines, since they do not maintain their own databases.

Results

Overall recall was calculated by taking the total number of relevant pages found by each search engine for all the queries, and dividing it by the total number of relevant pages found by all search engines for all queries. Figure 1 shows the overall recall for the search engines (the same data is presented as a table in appendix 2).

Figure 1: Overall recall for search engines

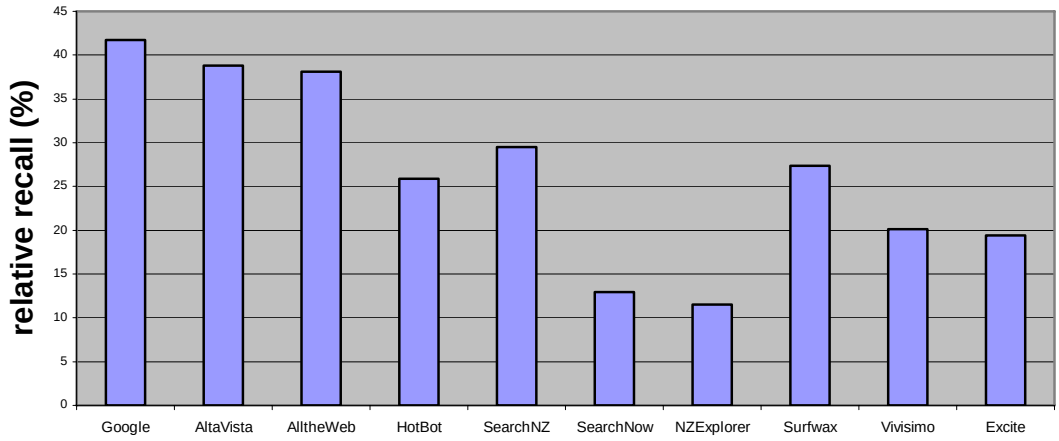

This shows that three of the four global search engines (Google, AltaVista, and AlltheWeb) had a higher recall of NZ information than the best of the local NZ search engines (SearchNZ). The other two local search engines performed poorly. This indicates that at least for New Zealand, geographically specific search engines do not offer great advantages over global search engines. A possible reason for this is the greater sophistication of the global search engines. It was notable that the two poorly performing local search engines had limited search features: NZ Explorer for example had a limit of three words per search, and did not use Boolean operators. The conclusion that global search engines tend to have higher recall than specialised search engines is broadly supported by Glander-Höbel's study of searching for chemical information [Glander-Höbel, 2001 #694].

In the course of carrying out the study, it became clear that several search features affected recall.

- OR operator: this enabled a search for alternative terms that might describe a concept. NZ Explorer did not have the OR operator, and this affected retrieval.
- Brackets to control the order of operations. This was not available in SearchNow, and was not documented for several other databases.
- Ability to set ranking terms independently of the search terms. This was only available in AltaVista's advanced mode. An example of the utility of being able to rank independently is in the *Bone People* search: here it was useful to include terms such as "novel" "book" "fiction" etc to discriminate against web sites featuring a music group called "The Bone People", however if these terms were used in ranking, sites were favoured that were primarily about writing in general rather than the specific novel.
- Limitation on the size of search statements. SearchNZ had a limit of 60 characters, which affected the ability to include alternative terminology.

- Poor documentation of search features. While the major search engines such as AltaVista and Google have extensive help pages documenting the search features, others had little guidance.

While no attempt was made to compare the extent of old and duplicate links, this was clearly a problem with some of the search engines. While the metasearch engines include algorithms to eliminate duplicates across search engines, it was noticeable that their results included more duplicates than the single engine results.

An interesting limitation of the search engine approach to searching was the search for the website of the New Zealand Lottery Grants Board. Many hits for a search on the phrase "new zealand lottery grants board" are to organisations which acknowledge funding from the board on their website. This reduced the probability that the Board's website would appear in the top ranked hits. Only half the search engines ranked the Board's website in the top 20, although it is likely that the site was indexed by more of the search engines.

It is also notable that none of the search engines achieved more than 45% recall: so on average a searcher relying on only one search engine will find less than half the relevant sites.

Interestingly, despite the theoretical advantage of covering a wider range of web pages, the metasearch engines did not achieve a better result than the other search engines. One possible reason for this are that metasearch engines do not allow full use of the search features of individual search engines. Metasearch engines appear to differ widely in the extent to which the search command is adapted to use the features of the individual search engines, and few provide complete documentation of what commands are available. Also, some major search engines such as Google block metasearch engines, which reduces their coverage.

Another significant point is that a metasearch engine selects the highest ranking hits from each search engine, and combines these into a list. However unless the user looks further down the hit list than for a single search engine, they will be seeing fewer results from each individual search engine than they would if they searched the engines separately, and won't gain an advantage from the wider coverage of the metasearch engine. For example if a metasearch engine combines the top 20 hits from five search engines, and presents a combined list of 100 items, a user who examines the top 20 hits in the metasearch list will be only seeing an average of four hits from each search engine. In practice ranking algorithms are not so good that relevant pages appear regularly in the top four hits, and the user may not achieve a better result than by looking at the top 20 hits for one of the single search engines. This example is a simplification (for example, metasearch engines also remove duplicates when combining the search engine lists) but illustrates that metasearch engines do not wholly overcome the problem of limited search engine coverage.

To what extent was database coverage of the .nz domain important in achieving recall? To investigate this, the relative coverage of the .nz domain by the different search engines was approximated by totaling the hits for 40 single word terms, as described in the methodology section above. The recall of the different search engines is plotted against the relative coverage in figure 2.

Figure 2: Recall vs coverage of NZ domain

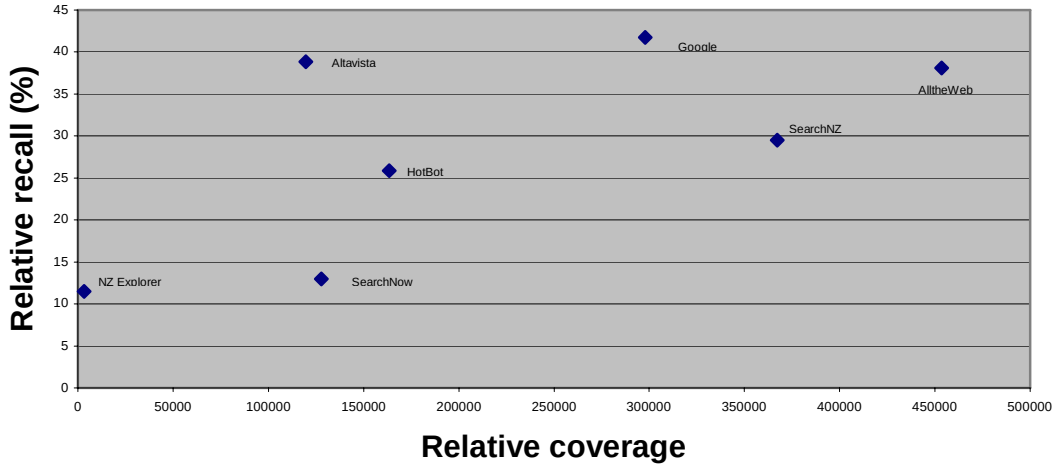

This indicates that AltaVista has achieved a larger number of relevant hits for its coverage of the .nz domain. A possible reason for this is the flexibility of searching in AltaVista, which in its advanced mode allows full Boolean logic, and the ability to set ranking terms independently of the search terms. On the other hand, although AlltheWeb and SearchNZ had the highest coverage of the .nz domain, this did not result in a proportionally greater number of relevant hits. This indicates that although the number of pages in the search engine database is important, it is not the only factor in achieving high recall. Figures for the relative coverage are shown in appendix 2.

The study also examined overlap of search engine results: to what extent were relevant web pages found by multiple search engines? Figure 3 shows a histogram of the number of search engines to find the same relevant web page (the same data is presented as a table in appendix 3).

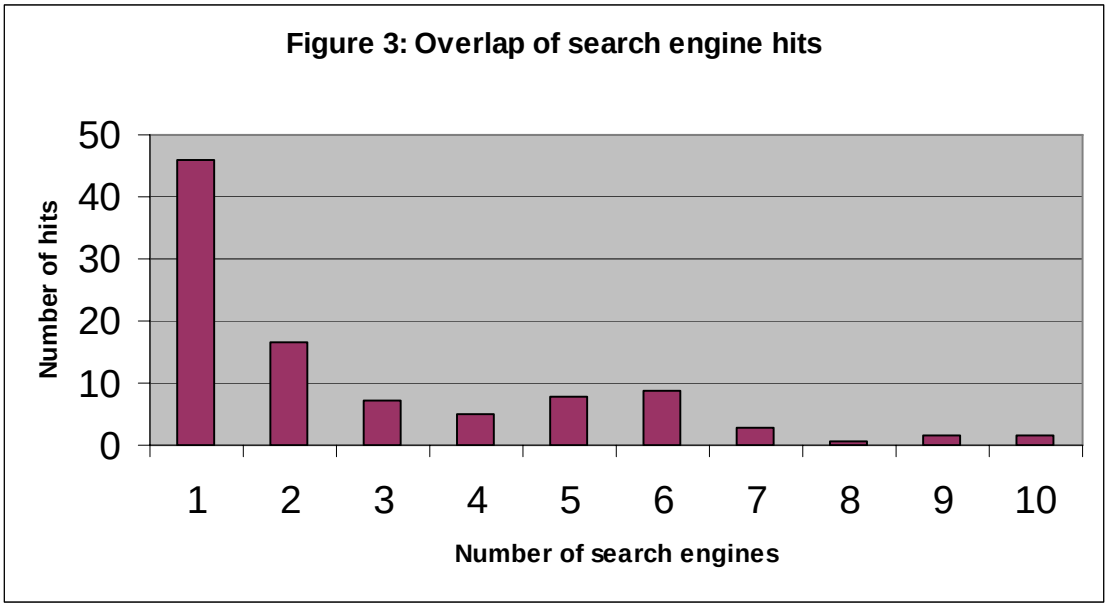

This shows that of the relevant web pages found across all the queries and search engines, 46% were found by only one search engine. On the other hand, only 1% of relevant pages were found by all 10 search engines. This lack of overlap is despite the inclusion of the three metasearch engines, which because they are searching multiple search engine databases, could be expected to increase the degree of overlap. Figure 3 reinforces the view that searching several search engines is necessary to achieve a complete search. Interestingly, metasearch engines do not appear to be a substitute for searching several search engines.

Another potential limitation of a local search engine is that it only searches the geographic domain. In this study, 36% of the relevant pages were outside the .nz domain. This occurred for a number of different reasons. For instance relevant information could be on a page maintained in New Zealand, but on a server registered in a generic top level domain such as .com or .org. A non-NZ server can also contain information relevant to NZ. For instance the Flags of the World web site, maintained internationally with servers in Sweden, Switzerland and elsewhere, had a good page of information about the Maori flag. The implication is that searches that are restricted to a local geographic domain may miss significant information related to the area. In practice some of the local search engines did include some sites outside the .nz domain, for instance the nz.com website that resides in the US, but contains NZ information. However it is impossible for a local search engine to include all possible sites that might be relevant to the geographic area.

Conclusion

The result of this study shows that, at least for New Zealand information, local search engines do not offer a significant advantage over global search engines. In practice, global search engines achieve similar or better coverage of the local domain, and tend to have more sophisticated search features that enable a higher level of recall. However, there is a low level of overlap between the results of different search engines, and no

search engine achieved more than 45% recall. This means that for a search to be as comprehensive as possible, it must include several search engines.

While the size of the search engine database is important, search features are also significant. For example, AltaVista's sophisticated search features apparently contributed to its high recall, despite a relatively low coverage of the .nz domain.

Another issue with using local search engines, or restricting searches to a local geographic domain, is that many relevant pages appear outside the local domain. In this study, 36% of relevant pages were not in the .nz domain.

A surprising result was that metasearch engines did not achieve significantly higher recall than the single search engines, despite the theoretical advantage of searching the databases of several different search engines. This could be due to a number of factors, including less efficient use of search features, blocking by major search engines, and the fact that metasearch engines only present a limited number of hits from each search engine. It would be valuable to carry out a more extensive comparison of the search performance of metasearch engines compared with single search engines.

References

(URLs checked 23 September 2002)

Bar-Ilan, J. 1998, 'On the overlap, the precision and estimated recall of search engines. A case study of the query "ERDOS"', *Scientometrics,* vol. 42, no. 2, pp. 207-28.

Clarke, S. J. & Willett, P. 1997, 'Estimating the recall performance of Web search engines', *Aslib Proceedings,* vol. 49, no. 7, pp. 184-9.

Dresel, R., Hornig, D., Kaluza, H., Peter, A., Rossmann, N. & Sieber, W. 2001, 'Evaluation deutscher Web-Suchwerkzeuge: ein vergleichender Retrievaltest. Eval uation of German Web search tools: a comparative retrieval test', *nfd Information: Wissenschaft und Praxis,* vol. 52, no. 7, pp. 381-4, 386-92.

Fricke, M. 1998, 'Measuring recall', *Journal of Information Science,* vol. 24, no. 6, pp. 409-417.

Gordon, M. & Pathak, P. 1999, 'Finding information on the World Wide Web: the retrieval effectiveness of search engines', *Information Processing and Management,* vol. 35, no. 2, pp. 141-80.

Hsieh-Yee, I. 1998, 'The retrieval power of selected search engines: how well do they address general reference questions and subject questions?', *The Reference Librarian,* vol. 60, pp. 24-47.

Jacso, P. 2001, 'Rating the metasearch engines', *Information Today,* vol. 18, no. 11, pp. 28, 33.

Kehoe, C., Pitkow, J., Sutton, K. & Aggarwal, G. 1999, *GVU's 10th WWW User Survey*, Georgia Institute of Technology: Graphic, Visualization, & Usability Center. Available: http://www.cc.gatech.edu/gvu/user_surveys/survey-1998-10/.

Leighton, H. V. & Srivastava, J. 1999, 'Precision among World Wide Web Search Services (Search Engines): Alta Vista, Excite, Hotbot, Infoseek, Lycos', *Journal of the American Society for Information Science*, vol. 50, no. 10, pp. 870-81.

Notess, G. R. 2002, *Search Engine Statistics: Relative Size Showdown*, Available: http://www.searchengineshowdown.com/stats/size.shtml.

Olvera Lobo, M. D. 2000, 'Rendimiento de los sistemas de recuperacion de informacion en al World Wide Web: revision metodologica. Methodology for evaluating information retriev al in WWW search engines', *Revista Espanola de Documentacion Cientifica*, vol. 23, no. 1, pp. 63-77.

Oppenheim, C., Morris, A., McKnight, C. & Lowley, S. 2000, 'The evaluation of WWW search engines', *Journal of Documentation*, vol. 56, no. 2, pp. 190-211.

Tomaiuolo, N. 1999, 'Are metasearches better searches?', *Searcher*, vol. 7, no. 1.

Appendices

1. Search queries

- Location of floristry schools in New Zealand.
- A description and image of the Maori flag.
- Who were the scriptwriters for the film *Scarfies*?
- Information on the novel *The Bone People*.
- Information about the Otago Central Rail Trail.
- A listing of New Zealand electronic journals.
- New Zealand companies that rent trucks.
- Information on the payment of British pensions in New Zealand
- The website of the New Zealand Lottery Grants Board
- Legal firms specialising in family/matrimonial law in NZ

2. Hits, Recall, and relative database coverage of the .nz domain, for the search engines studied

Search engine	total hits	recall %	Relative coverage
Google	58	42	297897

AltaVista	54	39	119691
AlltheWeb	53	38	453401
HotBot	36	26	163659
SearchNZ	41	29	367230
SearchNow	18	13	127748
NZExplorer	16	12	3416
Surfwax	38	27	N/a
Vivisimo	28	20	N/a
Excite	27	19	N/a

Notes:

- A total of 139 relevant web pages were retrieved across all queries and all search engines.
- "Relative coverage" is the total number of hits found in the search engine database for 40 single word queries on the .nz domain, and is a measure of the search engine's coverage of the .nz domain.

3. Overlap of search engine hits

Number of search engines	Number of hits	Percentage of total hits
1	64	46
2	23	17
3	10	7
4	7	5
5	11	8
6	12	9
7	4	3
8	1	1
9	2	1
10	2	1