City Research Online

City St George's, University of London

Citation: Bhogal, Jagdev & MacFarlane, A. (2013). Ontology Based Query Expansion with a Probabilistic Retrieval Model. Paper presented at the 6th Information Retrieval Facility Conference, IRFC 2013, 07-10-2013 - 09-10-2013, Limassol, Cyprus.

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: https://openaccess.city.ac.uk/id/eprint/5267/

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to City Research Online policy.

City Research Online: http://openaccess.city.ac.uk/ publications@citystgeorges.ac.uk

Ontology Based Query Expansion with a Probabilistic Retrieval Model

Jagdev Bhogal
Birmingham City University
Faculty of TEE
School of CTN
Millennium Point
Curzon Street
Birmingham
B7 4XG

Andrew Macfarlane
City University London
Centre for Interactive Systems Research
School of Informatics
Northampton Square
London
EC1V 0HB

Abstract.

This paper outlines the problems of traditional information retrieval systems and examines the use of ontologies for defining query context. The information retrieval system used is based on the probabilistic retrieval model. We extend the use of relevance feedback (RFB) and pseudo-relevance feedback (PF) query expansion techniques using information from a news domain ontology. The aim is to assess the impact of the ontology on the query expansion results with respect to recall and precision. We also tested the results for varying the relevance feedback parameters (number of terms or number of documents). The factors which influence the success of ontology based query expansion are outlined. Our findings show that ontology based query expansion has had mixed success. The use of the ontology has vastly increased the number of relevant documents retrieved, however, we conclude that for both types of query expansion, the PF results are better than the RFB results.

Keywords: Ontology, Query Expansion, Probabilistic Retrieval Model, Okapi, relevance feedback, pseudo-relevance feedback

adfa, p. 1, 2011.
© Springer-Verlag Berlin Heidelberg 2011

1 Introduction

In traditional information retrieval (IR) systems, the search process was iterative. Relevance feedback information was taken from the user so the retrieval process could be repeated using the additional relevance information. However since the users might be reluctant to provide feedback, researchers started focusing on contextual IR (Bhogal, MacFarlane & Smith 2007). Contextual IR integrates the user context into the retrieval process. Context can be inferred in many different ways. An ontological model can effectively disambiguate meanings of words from free text sentences (Buckland 2003). Ontologies can be used to infer context for ambiguous queries. The concepts in the ontology can be used for word sense disambiguation and subsequent query expansion.

A collection independent ontology is used for our experiments and ontology based query expansion is applied to the news domain. The ontological approach is suitable for the information intensive news domain. News is the communication of information on current events which is presented by print, broadcast, internet or word of mouth to a third party or mass audience. An ontology is a collective body of knowledge which is usually created and shared by users who are experts in that domain. A news ontology is usually created and shared by a group of specialists in the news field such as journalists, editors and Press standards organisations. Domain specific ontologies are used to model specialised vocabulary from that field such as medical terms. The news domain doesn't have a specific vocabulary as such it just uses plain English language in an accepted journalistic style. However what is important within this domain is the structure of news items. The structure of a news item includes: Headline, subheading, date, author, event description. News ontologies can be used to assist in different tasks such as news categorisation/classification, reasoning; searching; news annotation; updating, news summarization and news alerts. The chosen ontology has been derived from news articles so it is appropriate to use it for the searching task on the TREC document collection.

2 Methodology

This paper seeks to address questions such as whether the use of query expansion increases recall, precision or both and secondly how ontology based query expansion compares with relevance feedback/pseudo-relevance feedback techniques. This is the first time this particular TREC document collection and news ontology are being used in conjunction with each other so the results produced can provide useful baseline statistics for other researchers who want to carry out retrieval experiments using this particular combination of document collection and ontology.

The paper attempts to combine both approaches of relevance feedback query expansion and ontology based query expansion. A detailed investigation is carried out into the area of query expansion using a news ontology in a probabilistic retrieval

environment. Since we are interested in the news domain, an appropriate document collection and domain-specific ontology is selected. The Okapi system uses pseudo relevance and relevance feedback techniques (Robertson et al 1997) and the relevance feedback information can be based on pre-stored relevance judgments which indicate for each document whether it is relevant to the topic query or not. The techniques have proved to be successful to a certain extent - the revised retrieval model will build on the existing retrieval model and incorporate the use of the ontology information into the query expansion process.

The two main parameters of relevance feedback are: selection of terms and the sample size of relevant documents. In the Okapi system traditionally these have been 100 terms and 20 documents. Billerbeck and Zobel (2004) state that the choice of query expansion parameters used can affect the retrieval performance. As part of this research we experiment in varying these relevance feedback parameters and analyse the impact on the results. Another question that is addressed is whether to use all expanded terms or select the top 3 query expansion terms.

Short queries are better candidates for query expansion because they have insufficient terms to describe the information need and tend to be more ambiguous (Navigli and Velardi 2003). Therefore the query files are based on the topic titles (defined in section 2.2) only because they form shorter queries compared to queries based on the topic description.

We expand all queries and do not attempt to assess their ambiguity. With regards to term selection, all query terms are used for query expansion. Each query term in the Okapi Index is searched to provide new query terms, however in addition to this, the parent-child database is searched to provide ontology based query expansion terms.

Query Expansion is effective in increasing recall, it is less successful than relevance feedback (RFB) and may be as good as pseudo-relevance feedback (PF) (Billerbeck and Zobel 2003). Our research experiments can be used to test these claims.

2.1 Description of the Ontology

The WNO ontology was created by Kallipolitis et al (2007) studied a large number of international news articles from news agency websites and as a result based the ontology on 11 subjects which they felt were sufficiently representative in the domain of world news. The WNO ontology adheres to a tree structure with a maximum depth of 2 levels with class information provided in a top down fashion. The size of the WNO is 29.Kb making it easier to navigate and process programmatically.

WNO is written in XML in News Industry Text Format (NITF) which is published by IPTC and is designed to standardize the content and structure of text-based news articles; xml enhances system portability; WNO is relatively easy to process and is

based on the industry standards news codes taxonomy produced by ITPC (NewsML, 2008). NewsML provides a set terms for the news domain. This set of terms also known as Newscodes includes a hierarchy of terms and concepts that can be used to describe news in any field of interest. This hierarchical structure or taxonomy shown in Figure 1 consists of three levels:

Subject: topics at this level provide a description of the editorial content of news at a high level
Subjectmatter: a Subjectmatter provides a more precise description
Subjectdetail: provides the most specific description compared to the higher levels.

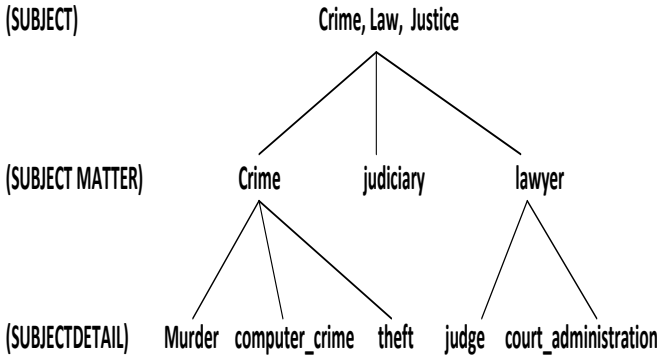

Fig. 1. Example of WNO instantiation

2.2 Chosen Test Collection

We selected the TREC newswire document collection from TREC (Disk2) because it is a reasonable size (over 231,000 documents) and even though it is not as large as other document collections it has associated topics/queries and also the relevance judgements were readily available (Harman 1993). Therefore it is ideal to use as a test collection for information retrieval evaluation. TREC document collections are widely accepted by the information retrieval research community. We used the adhoc task and topics 51-300 used (250 topics in total). A TREC topic is a natural language statement of information need written by real users of retrieval systems. Topics are distinct from queries because they contain more detail than the latter. Queries are constructed from topics and submitted to the retrieval system.

Disk2 is a smaller collection size in comparison to other document collections but the advantage of Disk 2 is that it contains a wider range of topics. The document collection contains news articles and non-news based articles. News based articles were not separated out and the entire collection was used because the aim was to use the as many relevance assessments in the document collection as possible. The non-news articles in the collection introduced "noise" to discover whether the news ontology

ranked news articles higher than non-news articles. There is no strong evidence to suggest that the news ontology favours news articles over non-news articles possibly because we are not putting any emphasis on the structure of the news articles. Only key terms are being used for the search thus we are treating all articles news or non-news in the same manner. If any structural feature of news articles are incorporated in the search process then it is likely that news articles would appear higher up in the ranked set of results.

2.3 Experiment Design

The document collection is indexed in Okapi which uses the probabilistic retrieval model (Sparck-Jones et al 2000). The document collection is indexed on the TREC document id (DOCNO), heading (HEAD) and description (TEXT) fields. Additionally, the News ontology is searched and hierarchical node relationship information is recorded in the parent-child database. The new system employs RFB and PF techniques but in expands the query further by making use of the parent-child information obtained from the ontology. The parent node(s) of a query term will broaden the query and the child node(s) of a query term will make a query more specific. The two main parameters of relevance feedback are: selection of terms and the sample size of relevant documents. We investigate the effect of varying the number of terms/documents relevance feedback parameters (Table 1).

Purpose of experiment	Experiment Number
Test ontology based query expansion compared to original system	Experiment 1 uses standard relevance feedback parameter values of 20 documents and 20 terms
Test the effect of varying the number of terms relevance feedback parameter	Experiments 2, 3, 4, 5, 6 use term relevance feedback parameters of 5, 10, 15, 100 and 200 respectively
Test the effect of varying the number of documents relevance feedback parameter	Experiments 7, 8, 9, 10, 11 use document relevance feedback parameters of 5, 10, 15, 100 and 200 respectively
Test the effect of selecting a subset of the expanded terms	Experiment 12 uses standard relevance feedback parameter values of 20 documents and 20 terms but only selects the top 3 expansion terms

Table 1. Summary of Experiments

In the Okapi system traditionally these have been 20 terms and 20 documents. Billerbeck and Zobel (2004) state that the choice of query expansion parameters used can affect the retrieval performance. There is no real consensus on the optimum num-

ber of documents to use for Query expansion. Sparck-Jones (1979) used 20, Robertson et al (1995) used 1000 (too much effort for very little return). Search routines were developed which used relevance feedback for query expansion and the resulting set of expanded terms were expanded even further by using associated broader and narrower ontological terms. Experimental results were evaluated using retrieval effectiveness metrics.

2.4 Metrics

Different types of evaluation metrics are required to evaluate the performance of each retrieval model and conduct a comparison. Recall and Precision are single-value metrics which evaluate the quality of an unordered set of documents returned by the system. For systems that return a ranked sequence of documents, it is desirable to also consider the order in which the returned documents are presented. Three other metrics used are mean average precision (MAP), Bpref and precision-recall curves. T-tests are carried out on Document level averages, Precision-Recall , MAP, and Average Recall statistical data to measure the statistical significance of these results. These measures are commonly used by other information retrieval systems thus making it easier to compare our results against those of other systems.

3 System

Okapi is an experimental IR system, written to examine various aspects of interactive IR research, including such tasks as bibliographic search and full-text search (Macfarlane et al 2010). The system uses the Probabilistic Retrieval Model and BM25 weighting functions are used to rank the documents (Sparck-Jones et al 2000). BM25 is a best match operator which retrieves ore relevant documents higher up the rank (MacFarlane and Tuson 2009). The probabilistic retrieval model is a highly effective retrieval model that makes explicit distinctions between occurrences of terms in relevant and non-relevant documents (Sparck-Jones et al 2000). It calculates the probability of a document being relevant if it contains certain terms. Figure 2 shows an overview of all the components which make up the final system. A single processor Sun SS10 with 64MB of core and about 12GB of disk was used as the main development machine and file server.

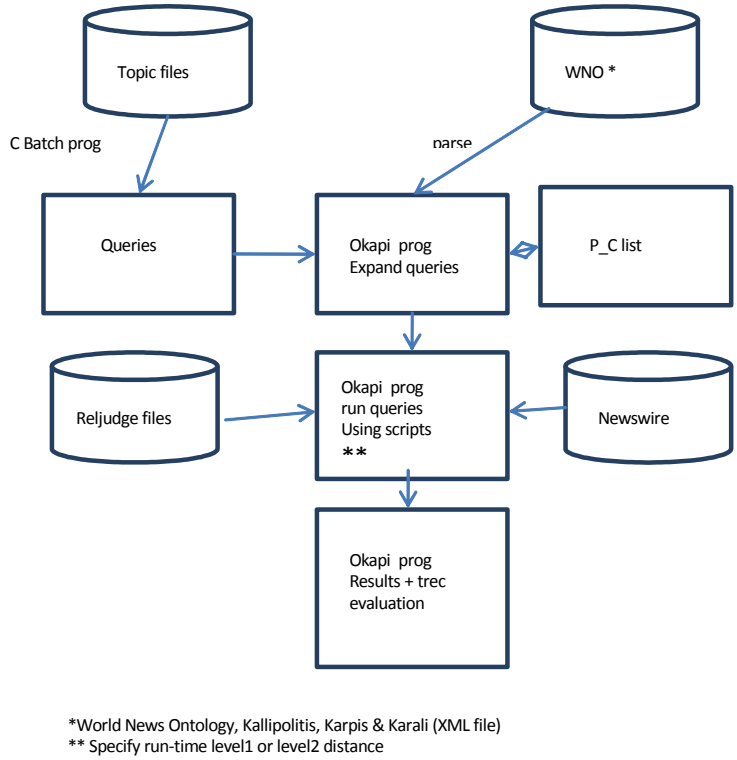

Fig. 2. System Overview

We build a separate database containing semantic information such as parent-child relationships between ontology nodes. This was required so we could transfer the ontology knowledge in an appropriate format and make it accessible to the Okapi software. This information is used to supply additional terms for expanding the original query terms. The information is stored in memory using a list structure which consists of: (childnode; parentnode; original term; weighting, r; nwords; levelno).

The level information shows whether the ontology parent-child term is obtained at a distance of 1 node (level1) or 2 nodes (level2) from the current node. This was done to find out the optimum level of ontology processing required to improve the retrieval results. Childnode is the descendant of the original terms, parentnode is the ancestor of the original term and original term is either a term from the query topic or from the refined wordlist resulting from the relevance feedback process. Weighting is the weighting value; r is the number of relevant documents for the term; nwords is the number of words in the list and levelno is distance from the original term to the parent/child term.

So for example :

If we searched for the term Crime then there would be the following list entries:
(Police; Crime_Law_Justice; Crime; 0.0, r; nwords;1);
(Law_Enforcement; Crime_Law_Justice; Crime; 0.0; r; nwords; 2);

The weighting values are updated during the three types of search routines.

4 Experiment Results

Use of the ontology has vastly increased the number of relevant documents retrieved. The ontology improves results for topics considered to be hard and non-hard topics.

The ontology has a better effect higher up the rank for the PF runs of Document Level Averages metric and the RFB runs starts improving from the lower end of the ranked set of documents which implies that the PF runs have more to gain from varying the relevance feedback parameters and do benefit from the use of the ontology. With the RFB runs, use of the ontology based terms for query expansion distorts the retrieval of relevant documents and is only useful at the lower end of the ranked list. In our opinion the RFB is harder to improve on because the top N documents used for RFB are already judged to be relevant so RFB without the use of the ontology produces good results which are hard to improve on. For the PF runs, the top N documents are **assumed** to be relevant because they are ranked highly by the system. These documents might not contain as many relevant query expansion terms as the RFB documents so any relevant additional ontology based query expansion terms will result in an improvement.

The use of ontology based query expansion has achieved high Recall results. This is possibly because query topics have a higher number of hits in the ontology for broader searches and for each hit, few ontology terms are retrieved but a higher proportion of the terms retrieved are relevant compared to ontology terms retrieved for narrower searches. An explanation for this is that quite a large number of query topic

terms are being found in the ontology and even though each of these only has one parent node associated, the use of these parent nodes is retrieving more relevant documents. Sometimes when searching for parent nodes, the ontology produces relevant terms. In other cases the ontology produces non-relevant terms which have a negative effect on precision and recall as shown in the example below:

TOPIC NUMBER = 90 (“data proven reserves oil natural gas producers”)
current word is oil
--> economy_business_finance

In this case the query topic is more to do with oil as an energy source and has nothing to do with economy_business_finance.

The use of ontology based query expansion has only increased mean average precision for a few cases but overall the precision is usually identical to the baseline or sometimes even below the baseline. In “narrower searches”, fewer query topic terms are matched with the ontology terms. Where a match occurs, the ontology term tends to have many more child terms associated with it but the precision-recall depends on the number of child terms that are relevant to the query topic and the number of relevant documents that contain the child term. In a few cases a larger number of relevant results are produced by the ontology which results in improved precision-recall. However in most other cases, just because an ontology term has lots of associated child terms, does not necessarily mean that the number of relevant documents retrieved will increase vastly. For example in narrower searches, where a term is quite general, many child nodes are retrieved of which only one or two might be relevant. Alternatively, the term produced is so general it does not improve the precision results at all because it retrieves a large number of documents which contain the general term and many of these documents are not relevant to the query topic. Another example to illustrate lack of improvement in performance retrieval is where many of the child terms are relevant to the search term but not relevant to the query. Therefore the use of ontology based query expansion has only increased mean average precision for a few cases but overall the precision is usually identical to the baseline or sometimes even below the baseline. The reason for this is that more ontology child terms are retrieved but a smaller proportion of these are actually relevant, thus having minimum impact on precision.

Retrieval results have improved with the use of the ontology but there is no clear trend that increasing the number of terms/documents results in improved retrieval. The number of terms parameter for relevance feedback benefits the PF and PF with ontology results but the number of documents parameter also has an effect on the RFB results. For example the graphs for document level averages (Figure 2) show more of an improvement compared to the Precision-Recall graphs (Figure 3).The reason for this is that it is easier to achieve improvements in precision in the top 5 or top 10 documents compared to achieving improvements in precision at recall .10 es-

pecially if the document collection is large. For example if the document collection is 20,000 documents, 0.10 recall calculates to 2000 documents.

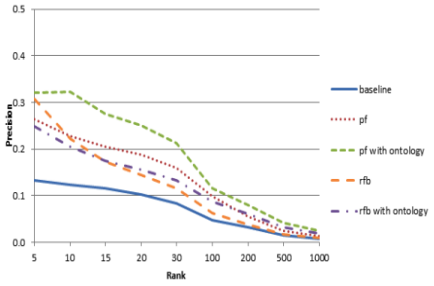

Fig. 2. Topic 51-100 Results Document Level Averages

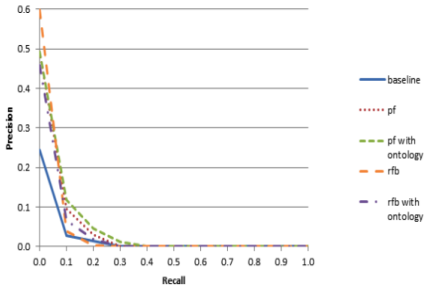

Fig. 3. Topic 51-100 Results Precision-Recall

Improved retrieval results depend on the ontology coverage of the topic in breadth and depth; the similarity of terms between the ontology and the document collection; and finally the document collection coverage of the ontology terms. The ontology could have a lot of terms relevant to the topic but these terms might not be contained in many documents thus there is minimum impact on performance.

The ontology results in improvements for some topic-sets but not for others. First of all, when searching an ontology using query topic terms, we need to find at least one hit in the ontology for any improvements to take place. Some topics have more ontology hits than others. The second success factor relies not on just the number of hits in the ontology but on the retrieved ontology terms being relevant to the query topic.

We use stemmed keywords when searching the ontology, so its possible that the actual ontology hits are irrelevant and/or the retrieved ontology terms are irrelevant.

Topic no = 223 (responsible great emergence Microsoft computer industry)
ORIG WORD IS emergenc
--> explosion

Emergence has been stemmed to emergenc, and ontology picks up non-relevant term explosion which is more related to emergency than emergence.

Suppose we find a good set of ontology terms to expand the query with, then the next factor in improving retrieval relates to finding enough documents in the document collection that contain the ontology term and are relevant to the query topic. If the match between the ontology and the document collection is poor, then even though the ontology terms are relevant to the query topic, because there aren't enough documents containing that term, query expansion has minimal effect on recall/precision. Alternatively, if the parent/child term obtained from the ontology is too

general, then many documents are retrieved but very few of these are relevant to the query topic.

A topic hardness measure is calculated as the average over a given set of runs of precision for each topic after all relevant documents have been retrieved OR after 100 documents have been retrieved if more than 100 documents are relevant. The measure is oriented towards high-recall performance and how well systems do at finding all relevant documents. If no system does well on a query then it can be called a hard query. According to TREC hardness measure given in Buckley et al (1996) the performance for TREC 4 (topic 201-250) and TREC 5 (251-300) drops from 0.676 to 0.672 and 0.556 respectively. These are seen to be difficult topics because they are progressively shorter in length and higher level in nature. This trend is mirrored in the SMART experiments. For example in TREC1 the precision is 0.2431 and in TREC2 the best precision has improved to 0.2594 but in TREC 4 and TREC 5 the precision dropped to 0.1507 and 0.1038 respectively.

For each run, we can compared across the various metrics to see which topics occur the most. Again, topics51-100 and topics251-300 have the highest frequency across the various metrics. According to Buckley (1996), topic251-300 is considered to be a hard topic set. So the ontology seems to have improved the retrieval performance for a hard topic as well those considered not to be hard.

We can analyse table 2 for statistical significance (*=significant **=very significant). For the Document level averages, twice as many PF results are significant/very significant compared to the RFB results. For the Precision-Recall curves, only the PF results are significant/very significant. Recall is the metric with the highest number of statistically significant results.

So we have high recall at expense of precision. This is good for the some domains because professional searchers such as investigative journalists prefer to obtain as much information about a given news story as possible. Lawyers need to look at all case statutes in order to produce a strong argument otherwise missed case articles will weaken their evidence. In the same way investigative journalists need to ensure they have accessed all relevant articles in order to produce a thorough report on the subject they are investigating otherwise they will open to criticism if gaps in the research are found. Also the results analysis shows the document level average results are better than recall –precision and the document level averages (PF runs) are benefitting from the ontology higher up the rank. Again this would indicate that the ranking algorithm is working and searchers tend to concentrate on the documents occurring higher up in a ranked set of results. The documents for PF are “assumed” to be relevant because they appear high up in the system ranking, whereas the documents for RFB are judged by human assessors as actually being relevant. It would be difficult to improve retrieval performance on the RFB relevant documents, however the PF runs have more to gain from these other factors than RFB.

Experiment	Doc level Averages		Precision-Recall		Recall		MAP		BPref	
	PF	RFB	PF	RFB	PF	RFB	PF	RFB	PF	RFB
5 terms	51-100 **(0.0001)**	251-300	251-300	51-100	51-100 **(0.001)**	51-100 **(0.002)**	51-100	51-100	51-100 (0.005)	51-100
10 terms	251-300 (0.013)	251-300	251-300	51-100	151-200 **(0.004)**	51-100 (0.002)	251-300	51-100	251-300 **(0.000)**	51-100
15 terms	51-100 **(0.003)**	51-100	251-300	51-100	151-200 **(0.006)**	201-250 **(0.003)**	251-300	51-100	251-300 **(0.001)**	51-100
100 terms	251-300 **(0.002)**	51-100 (0.034)	251-300	251-300	251-300 (0.044)	201-250 (0.010)	251-300	51-100 (0.022)	251-300 (0.045)	51-100
200 terms	51-100 **(0.000)**	151-200	251-300	151-200	151-200 **(0.001)**	151-200 **(0.004)**	51-100 **(0.008)**	151-200	101-150 **(0.004)**	101-150
5 docs	251-300 (0.021)	151-200	251-300	51-100	151-200	151-200	251-300	151-200	251-300 **(0.003)**	51-100 **(0.008)**
10 docs	101-150	251-300	251-300	51-100	51-100	51-100	251-300	51-100	51-100	51-100
15 docs	251-300	251-300	251-300	51-100	151-200	51-100	251-300	51-100	251-300	101-150
100 docs	51-100 **(0.002)**	201-250 **(0.008);** 251-300 **(0.009)**	251-300	51-100	51-100	151-200	51-100	251-300 (0.11)	51-100 (0.029)	51-100
200 docs	251-300 **(0.000)**	151-200 **(0.007)**	251-300	251-300	251-300	251-300	101-150	251-300 **(0.008)**	101-150 **(0.003)**	51-100 (0.011)
20 terms/docs	51-100 (0.023)	51-100	251-300	51-100	201-250 (0.013)	201-250 (0.049)	251-300 (0.013)	51-100 **(0.008)**	251-300 **(0.001)**	51-100 **(0.009)**
Top 3 expansion terms	51-100 **(0.001)**	151-200	251-300	51-100	51-100 **(0.000)**	51-100 **(0.001)**	51-100 **(0.001)**	51-100 **(0.000)**	51-100 **(0.000)**	51-100 **(0.001)**

Table 2. Overall Results Summary and Statistical Significance

Robin and Ramalho (2003) used disk2 of the TREC collection and the WordNet ontology to expand query words with some of their synonyms and hypernyms. For comparison purposes, the document collection is the same but we have used a news based ontology to obtain synonyms and hypernyms instead. The other difference is that Robin and Ramalho used the F-measure metric instead of BPref. Finally they used bounds of 10, 15, 20, 30 and 50 documents, we used 5, 10, 15, 20, 100 and 200 terms/documents. They found that all expansion strategies improve overall effectiveness by improving recall more than they worsen precision (in relative terms). Their results show that recall can be boosted up by as much as 72.4% relative to the no expansion case. They also expand to the first-level in the ontology. Their best query expansion strategy yields only a 2.51% improvement reaching 9.3% and only 11% of all relevant documents together with 77.5% irrelevant ones. For bounded precision for the top 20, 30, 40 and 50 documents, precision respectively improved by 1%, 12%, 17% and 37%. In comparison to Robin and Ramalho's work, our results are just as good if not better for recall and precision. Even though there was in improvement in our results for some topics across different runs, unlike Robin and Ramalho we did

not discover any linear trend resulting from increasing the number of terms/documents.

5 Conclusion

It is important to compare our findings with those of other related research. The use of ontologies for query expansion has had mixed success (Gonzalo et al 1998) because they are effective in increasing recall and less successful than RFB but as good as PF (Billerbeck and Zobel 2004). Our findings support these statements. Our attempts at ontology based query expansion have had mixed success. Use of the ontology has vastly increased the number of relevant documents retrieved. We can conclude that for both types of query expansion, the PF results are better than the RFB results. Our findings are similar to that of Billerbeck and Zobel (2003) in that ontology based query expansion enhances recall, and produces bigger improvements for PF compared to RFB. The ontology has a better effect higher up the rank for the PF runs of Document Level Averages metric and the RFB runs starts improving from the lower end of the ranked set of documents which implies that the PF runs have more to gain from varying the relevance feedback parameters and do benefit from the use of the ontology. Query expansion seems to be more successful only on relevant documents (Ogawa and Mano 2001, Billerbeck and Zobel 2003). In support of this statement, use of the ontology based terms for query expansion in RFB runs is distorting the retrieval of relevant documents and is only useful at the lower end of the ranked list.

Our work can be improved by conducting further research on better term selection. Selecting query expansion terms based on relatedness to the whole query is more effective (Qiu and Frei (1993); Mandala, Tokunaga and Tanaka (1998)) In TREC 8 (Robertson and Walker 1999), a term selection measure is used for selective expansion to measure the statistical significance of any given term's association with relevance. The research indicates that the choice of level (5%, 1% or 0.1%) is largely arbitrary and recommends setting the criterion in relation to the size of the vocabulary 1/Vec (c is a constant, positive or negative).

To increase intelligence, the system should recognise synonyms and utilise homography - a spelling method that represents every sound by a character. Our system does not at present have these features. Compound words add complexity to the query expansion process however, further research is needed on the effective deployment of compound words in query expansion.

Finally we could apply our query expansion algorithms to different ontologies to see what difference each ontology makes to the query expansion process and the reasons why one ontology is inherently better than another. For example the NEWS ontology is larger in size which indicates it has more coverage of the news domain. It also has a more complex lattice structure and deeper levels of nodes than the ontology we used. It would be more complex to process but could produce enriched results.

Query expansion has been successful to a certain extent but there is still scope to improve the techniques for selecting and designing algorithms for optimum parameter choice and only expanding queries which would benefit from the query expansion process.

References

1. Bhogal, J., MacFarlane, A., Smith, P. (2007). "A review of Ontology based query expansion." Information Processing and Management 43(4): 866-886.
2. Billerbeck, B., Zobel, J. (2004). Questioning Query Expansion: An Examination of Behaviour and Parameters. Proceedings of the Fifteenth Australasian Database Conference, ADC 2004, Dunedin, New Zealand, Australian Computer Society.
3. Buckland, M. (2003). Translingual information management using domain ontologies. [online] available at http://metadata.sims.berkeley.edu/GrantSupported/tides.html [Accessed 10th September 2003]
4. Buckley, C. trec_eval [online] available at http://trec.nist.gov/trec_eval/ [Accessed 14th July 2009]
5. Gonzalo, J., Verdejo, F., Chugur, I., Cigarran, J. (1998). Indexing with WordNet synsets can improve text retrieval. Coling-ACL 98, Canada.
6. Harman D., (1993). Overview of the First Text REtrieval Conference, 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Pittsburgh, USA, June 27 - July 1, 1993. pp. 36-48, Proceedings. ACM 1993.
7. Kallipolitis, L., Karpis, V., Karali, I. (2007). World News Finder: How we Cope without the Semantic Web. Artificial Intelligence and Applications (AIA), IASTED/ACTA Press.
8. MacFarlane, A., Secker, A., May, P., Timmis, J. (2010). "An experimental comparison of a genetic algorithm and a hill-climber for term selection." Journal of Documentation 66(4): 513-531.
9. MacFarlane, A., Tuson, A. (2009). "Local Search: A guide for the Information Retrieval Practitioner." Information Processing and Management 45(1): 159-174.
10. Mandala, R., Tokunaga, T., Tanaka, H (1999). "Combining Multiple Evidence from Different Types of Thesaurus for Query Expansion." Research and Development in Information Retrieval.
11. Navigli, R., Velardi, P (2003). An analysis of Ontology-based query expansion strategies. Workshop on Adaptive Text Extraction and Mining (ATEM 2003), in the 14th European Conference on Machine Learning (ECML 2003),September 22-26th,, Cavtat-Dubrovnik, Croatia.
12. NewsML (2008). Metadata Taxonomies for news [online] available at http://www.iptc.org/site/NewsCodes/ [Accessed 12th October 2008]
13. Ogawa, Y., Mano, H. (2001). Selecting expansion terms in automatic query expansion. ACM-SIGIR Annual Int. Conf. on Research and Development in Information Retrieval.
14. Robertson, S., Walker, S . (1999). Okapi/Keenbow at TREC-8. TREC.
15. Robertson, S., Walker, S., HancockBeaulieu, M., Gatford, M., Payne, A. (1995). Okapi at TREC-4. The Fourth Text REtrieval Conference (TREC-4).
16. Robin, J., Ramalho, F., 2003. Can ontologies improve web search engine effectiveness before the advent of the semantic web? In SBBD 2003 (pp. 157–169).