

Query Expansion Using Local and Global Document Analysis

Jinxi Xu and W. Bruce Croft
Center for Intelligent Information Retrieval
Computer Science Department
University of Massachusetts, Amherst
Amherst, MA 01003-4610, USA
xu@cs.umass.edu croft@cs.umass.edu

Abstract

Automatic query expansion has long been suggested as a technique for dealing with the fundamental issue of word mismatch in information retrieval. A number of approaches to expansion have been studied and, more recently, attention has focused on techniques that analyze the corpus to discover word relationships (global techniques) and those that analyze documents retrieved by the initial query (local feedback). In this paper, we compare the effectiveness of these approaches and show that, although global analysis has some advantages, local analysis is generally more effective. We also show that using global analysis techniques, such as word context and phrase structure, on the local set of documents produces results that are both more effective and more predictable than simple local feedback.

1 Introduction

The problem of word mismatch is fundamental to information retrieval. Simply stated, it means that people often use different words to describe concepts in their queries than authors use to describe the same concepts in their documents. The severity of the problem tends to decrease as queries get longer, since there is more chance of some important words co-occurring in the query and relevant documents. In many applications, however, the queries are very short. For example, applications that provide searching across the World-Wide Web typically record average query lengths of two words [Croft et al., 1995]. Although this may be one extreme in terms of IR applications, it does indicate that most IR queries are not long and that techniques for dealing with word mismatch are needed.

An obvious approach to solving this problem is query expansion. The query is expanded using words or phrases with similar meaning to those in the query and the chances of matching words in relevant documents are therefore increased. This is the basic idea behind the use of a thesaurus

in query formulation. There is, however, little evidence that a general thesaurus is of any use in improving the effectiveness of the search, even if words are selected by the searchers [Voorhees, 1994]. Instead, it has been proposed that by automatically analyzing the text of the corpus being searched, a more effective thesaurus or query expansion technique could be produced.

One of the earliest studies of this type was carried out by Sparck Jones [Sparck Jones, 1971] who clustered words based on co-occurrence in documents and used those clusters to expand the queries. A number of similar studies followed but it was not until recently that consistently positive results have been obtained. The techniques that have been used recently can be described as being based on either global or local analysis of the documents in the corpus being searched. The global techniques examine word occurrences and relationships in the corpus as a whole, and use this information to expand any particular query. Given their focus on analyzing the corpus, these techniques are extensions of Sparck Jones' original approach.

Local analysis, on the other hand, involves only the top ranked documents retrieved by the original query. We have called it *local* because the techniques are variations of the original work on local feedback [Attar & Fraenkel, 1977, Croft & Harper, 1979]. This work treated local feedback as a special case of relevance feedback where the top ranked documents were assumed to be relevant. Queries were both reweighted and expanded based on this information.

Both global and local analysis have the advantage of expanding the query based on all the words in the query. This is in contrast to a thesaurus-based approach where individual words and phrases in the query are expanded and word ambiguity is a problem. Global analysis is inherently more expensive than local analysis. On the other hand, global analysis provides a thesaurus-like resource that can be used for browsing without searching, and retrieval results with local feedback on small test collections were not promising.

More recent results with the TREC collection, however, indicate that local feedback approaches can be effective and, in some cases, outperform global analysis techniques. In this paper, we compare these approaches using different query sets and corpora. In addition, we propose and evaluate a new technique which borrows ideas from global analysis, such as the use of context and phrase structure, but applies them to the *local* document set. We call the new technique local context analysis to distinguish it from local feedback.

In the next section, we describe the global analysis procedure used in these experiments, which is the *Phrasefinder* component of the INQUERY retrieval system [Jing & Croft,

1994]. Section 3 covers the local analysis procedures. The local feedback technique is based on the most successful approaches from the recent TREC conference [Harman, 1996]. Local context analysis is described in detail.

The experiments and results are presented in section 4. Both the TREC [Harman, 1995] and WEST [Turtle, 1994] collections are used in order to compare results in different domains. A number of experiments with local context analysis are reported to show the effect of parameter variations on this new technique. The other techniques are run using established parameter settings. In the comparison of global and local techniques, both recall/precision averages and query-by-query results are used. The latter evaluation is particularly useful to determine the robustness of the techniques, in terms of how many queries perform substantially worse after expansion. In the final section, we summarize the results and suggest future work.

2 Global Analysis

The global analysis technique we describe here has been used in the INQUERY system in TREC evaluations and other applications [Jing & Croft, 1994, Callan et al., 1995], and was one of the first techniques to produce consistent effectiveness improvements through automatic expansion. Other researchers have developed similar approaches [Qiu & Frei, 1993, Schütze & Pedersen, 1994] and have also reported good results.

The basic idea in global analysis is that the global context of a *concept* can be used to determine similarities between *concepts*. Context can be defined in a number of ways, as can concepts. The simplest definitions are that all words are concepts (except perhaps stop words) and that the context for a word is all the words that co-occur in documents with that word. This is the approach used by [Qiu & Frei, 1993], and the analysis produced is related to the representations generated by other dimensionality-reduction techniques [Deerwester et al., 1990, Caid et al., 1993]. The essential difference is that global analysis is only used for query expansion and does not replace the original word-based document representations. Reducing dimensions in the document representation leads to problems with precision. Another related approach uses clustering to determine the context for document analysis [Crouch & Yang, 1992].

In the Phrasefinder technique used with INQUERY, the basic definition for a concept is a noun group, and the context is defined as the collection of fixed length *windows* surrounding the concepts. A noun group (phrase) is either a single noun, two adjacent nouns or three adjacent nouns. Typical effective window sizes are from 1 to 3 sentences. One way of visualizing the technique, although not the most efficient way of implementing it, is to consider every concept (noun group) to be associated with a *pseudo-document*. The contents of the pseudo-document for a concept are the words that occur in every window for that concept in the corpus. For example, the concept *airline pilot* might have the words *pay*, *strike*, *safety*, *air*, *traffic* and *FAA* occurring frequently in the corresponding pseudo-document, depending on the corpus being analyzed. An INQUERY database is built from these pseudo-documents, creating a concept database. A filtering step is used to remove words that are too frequent or too rare, in order to control the size of the database.

To expand a query, it is run against the concept database using INQUERY, which will generate a ranked list of phrasal concepts as output, instead of the usual list of document names. Document and collection-based weighting of match-

ing words are used to determine the concept ranking, in a similar way to document ranking. Some of the top-ranking phrases from the list are then added to the query and weighted appropriately. In the Phrasefinder queries used in this paper, 30 phrases are added into each query and are downweighted in proportion to their rank position. Phrases containing only terms in the original query are weighted more heavily than those containing terms not in the original query.

Figure 1 shows the top 30 concepts retrieved by Phrasefinder for the TREC4 query 214 "What are the different techniques used to create self induced hypnosis". While some of the concepts are reasonable, others are difficult to understand. This is due to a number of spurious matches with noncontent words in the query.

The main advantages of a global analysis approach like the one used in INQUERY is that it is relatively robust in that average performance of queries tends to improve using this type of expansion, and it provides a thesaurus-like resource that can be used for browsing or other types of concept search. The disadvantages of this approach is that it can be expensive in terms of disk space and computer time to do the global context analysis and build the searchable database, and individual queries can be significantly degraded by expansion.

3 Local Analysis

3.1 Local Feedback

The general concept of local feedback dates back at least to a 1977 paper by Attar and Fraenkel [Attar & Fraenkel, 1977]. In this paper, the top ranked documents for a query were proposed as a source of information for building an automatic thesaurus. Terms in these documents were clustered and treated as quasi-synonyms. In [Croft & Harper, 1979], information from the top ranked documents is used to re-estimate the probabilities of term occurrence in the relevant set for a query. In other words, the weights of query terms would be modified but new terms were not added. This experiment produced effectiveness improvements, but was only carried out on a small test collection.

Experiments carried out with other standard small collections did not give promising results. Since the simple version of this technique consists of adding common words from the top-ranked documents to the original query, the effectiveness of the technique is obviously highly influenced by the proportion of relevant documents in the high ranks. Queries that perform poorly and retrieve few relevant documents would seem likely to perform even worse after local feedback, since most words added to the query would come from non-relevant documents.

In recent TREC conferences, however, simple local feedback techniques appear to have performed quite well. In this paper, we expand using a procedure similar to that used by the Cornell group in TREC 4 & 3 [Buckley et al., 1996]. The most frequent 50 terms and 10 phrases (pairs of adjacent non stop words) from the top ranked documents are added to the query. The terms in the query are reweighted using the Rocchio formula with $\alpha : \beta : \gamma = 1 : 1 : 0$.

Figure 2 shows terms and phrases added by local feedback to the same query used in the previous section. In this case, the terms in the query are stemmed.

One advantage of local feedback is that it can be relatively efficient to do expansion based on high ranking documents. It may be slightly slower at run-time than, for

hypnosis	meditation	practitioners
dentists	antibodies	disorders
psychiatry	immunodeficiency-virus	anesthesia
susceptibility	therapists	dearth
atoms	van-dyke	self
confession	stare	proteins
katie	johns-hopkins-university	growing-acceptance
reflexes	voltage	ad-hoc
correlation	conde-nast	dynamics
ike	illnesses	hoffman

Figure 1: Phrasefinder concepts for TREC4 query 214

hypnot	hypnotiz	19960500
psychosomat	psychiatr	immun
mesmer	franz	suscept
austrian	dyck	psychiatrist
shesaid	tranc	professor
hallucin	18th	centur
hilgard	11th	unaccept
19820902	syndrom	exper
physician	told	patient
hemophiliac	strang	cortic
ol	defic	muncie
spiegel	diseas	imagin
suggest	dyke	feburar
immunoglobulin	reseach	fresco
person	numb	katie
psoriasis	treatment	medicin
17150000	ms	franz-mesmer
austrian-physician	psychosomat-medicin	
hypnot-state	fight-immun	intern-congress
late-18th	diseas-fight	hypnotiz-peopl
ms-ol		

Figure 2: Local feedback terms and phrases for TREC4 query 214

example, Phrasefinder, but needs no thesaurus construction phase. Local feedback requires an extra search and access to document information. If document information is stored only for this purpose, then this should be counted as a space overhead for the technique, but it likely to be significantly less than a concept database. A disadvantage currently is that it is not clear how well this technique will work with queries that retrieve few relevant documents.

3.2 Local Context Analysis

Local context analysis is a new technique which combines global analysis and local feedback. Like Phrasefinder, noun groups are used as concepts and concepts are selected based on co-occurrence with query terms. Concepts are chosen from the top ranked documents, similar to local feedback, but the best passages are used instead of whole documents. The standard INQUERY ranking is not used in this technique.

Below are the steps to use local context analysis to expand a query Q on a collection.

1. Use a standard IR system (INQUERY) to retrieve the top n ranked passages. A passage is a text window of fixed size (300 words in these experiments [Callan, 1994]).

There are two reasons that we use passages rather than documents. Since documents can be very long and

about multiple topics, a co-occurrence of a concept at the beginning and a term at the end of a long document may mean nothing. It is also more efficient to use passages because we can eliminate the cost of processing the unnecessary parts of the documents.

2. Concepts (noun phrases) in the top n passages are ranked according to the formula

$$bel(Q, c) = \prod_{t_i \in Q} (\delta + \log(af(c, t_i)) idf_c / \log(n))^{idf_i}$$

Where

$$\begin{aligned} af(c, t_i) &= \sum_{j=1}^{j=n} ft_{ij} fc_j \\ idf_i &= \max(1.0, \log_{10}(N/N_i)/5.0) \\ idf_c &= \max(1.0, \log_{10}(N/N_c)/5.0) \end{aligned}$$

- c is a concept
- ft_{ij} is the number of occurrences of t_i in p_j
- fc_j is the number of occurrences of c in p_j
- N is the number of passages in the collection
- N_i is the number of passages containing t_i
- N_c is the number of passages containing c
- δ is 0.1 in this paper to avoid zero bel value

The above formula is a variant of the $tf\ idf$ measure used by most IR systems. In the formula, the af part

rewards concepts co-occurring frequently with query terms, the *idf_c* part penalizes concepts occurring frequently in the collection, the *idf_q* part emphasizes infrequent query terms. Multiplication is used to emphasize co-occurrence with all query terms.

3. Add m top ranked concepts to Q using the following formula:

$$\begin{aligned} Q_{new} &= \#WSUM(1.0 \ 1.0 \ Q \ w \ Q') \\ Q' &= \#WSUM(1.0 \ w_1 \ c_1 \ w_2 \ c_2 \ \dots \ w_m \ c_m) \end{aligned}$$

In our experiments, m is set to 70 and w_i is set to $1.0 - 0.9 * i / 70$. Unless specified otherwise, w is set to 2.0. We call Q' the auxiliary query. $\#WSUM$ is an INQUERY query operator which computes a weighted average of its components.

Figure 3 shows the top 30 concepts added by local context analysis to TREC4 query 214.

Local context analysis has several advantages. It is computationally practical. For each collection, we only need a single pass to collect the collection frequencies for the terms and noun phrases. This pass takes about 3 hours on an Alpha workstation for the TREC4 collection. The major overhead to expand a query is an extra search to retrieve the top ranked passages. On a modern computer system, this overhead is reasonably small. Once the top ranked passages are available, query expansion is fast: when 100 passages are used, our current implementation requires only several seconds of CPU time to expand a TREC4 query. So local context analysis is practical even for interactive applications. For queries containing proximity constraints (e.g. phrases), Phrasefinder may add concepts which co-occur with all query terms but do not satisfy proximity constraints. Local context analysis does not have such a problem because the top ranked passages are retrieved using the original query. Because it does not filter out frequent concepts, local context analysis also has the advantage of using frequent but potentially good expansion concepts. A disadvantage of local context analysis is that it may require more time to expand a query than Phrasefinder.

4 Experiments

4.1 Collections and Query Sets

Experiments are carried out on 3 collections: TREC3 that comprises Tipster 1 and 2 datasets with 50 queries (topics 151-200), TREC4 that comprises Tipster 2 and 3 datasets with 49 queries (topics 202-250) and WEST with 34 queries. TREC3 and TREC4 (about 2 GBs each) are much larger and more heterogeneous than WEST. The average document length of the TREC documents is only 1/7 of that of the WEST documents. The average number of relevant documents per query with the TREC collections is much larger than that of WEST. Table 1 lists some statistics about the collections and the query sets. Stop words are not included.

4.2 Local Context Analysis

Table 2 shows the performance of local context analysis on the three collections. 70 concepts are added into each query using the expansion formula in section 3.2.

Local text analysis performs very well on TREC3 and TREC4. All runs produce significant improvements over the baseline on the TREC collections. The best run on

TREC4 (100 passages) is 23.5% better than the baseline. The best run on TREC3 (200 passages) is 24.4% better than the baseline. On WEST, the improvements over the baseline are not as good as on TREC3 and TREC4. With too many passages, the performance is even worse than the baseline.

The high baseline of the WEST collection (53.8% average precision) suggests that the original queries are of very good quality and we should give them more emphasis. So we downweight the expansion concepts by 50% by reducing the weight of auxiliary query Q' from 2.0 to 1.0. Table 3 shows that downweighting the expansion concepts does improve performance.

It is interesting to see how the number of passages used affects retrieval performance. To see it more clearly, we plot the performance curve on TREC4 in figure 4. Initially, increasing the number of passages quickly improves performance. The performance peaks at a certain point. After staying relatively flat for a period, the performance curves drop slowly when more passages are used. For TREC3 and TREC4, the optimal number of passages is around 100, while on WEST, the optimal number of passages is around 20. This is not surprising because the first two collections are an order of magnitude larger than WEST. Currently we do not know how to automatically determine the optimal number of passages to use. Fortunately, local context analysis is relatively insensitive to the number of the passages used, especially for large collections like the TREC collections. On the TREC collections, between 30 and 300 passages produces very good retrieval performance.

5 Local Text Analysis vs Global Analysis

In this section we compare Phrasefinder and local context analysis in term of retrieval performance. Tables 4-5 compare the retrieval performance of the two techniques on the TREC collections. On both collections, local context analysis is much better than Phrasefinder. On TREC3, Phrasefinder is 7.8% better than the baseline while local context analysis using the top ranked 100 passages is 23.3% better than the baseline. On TREC4, Phrasefinder is only 3.4% better than the baseline while local context analysis using the top ranked 100 passages is 23.5% than the baseline. In fact, all local context analysis runs in table 2 are better than Phrasefinder on TREC3 and TREC4. On both collections, Phrasefinder hurts the high-precision end while local context analysis helps improve precision. The results show that local context analysis is a better query expansion technique than Phrasefinder.

We examine two TREC4 queries to show why Phrasefinder is not as good as local context analysis. For one example, "China" and "Iraq" are very good concepts for TREC4 query "Status of nuclear proliferation treaties - violations and monitoring". They are added into the query by local context analysis but not by Phrasefinder. It appears that they are filtered out by Phrasefinder because they are frequent concepts. For the other example, Phrasefinder added the concept "oil spill" to TREC4 query "As a result of DNA testing, are more defendants being absolved or convicted of crimes". This seems to be strange. It appears that Phrasefinder did this because "oil spill" co-occurs with many of the terms in the query, e.g., "result", "test", "defendant", "absolve" and "crime". But "oil spill" does not co-occur with "DNA", which is a key element of the query. While it is very hard to automatically determine which terms are key elements of a query, the product function used by local context analysis for selecting expansion concepts should be

hypnosis	brain-wave	ms.-burns
technique	pulse	reed
brain	ms.-olness	trance
hallucination	process	circuit
van-dyck	behavior	suggestion
case	spiegel	finding
hypnotizables	subject	van-dyke
patient	memory	application
katie	muncie	approach
study	point	contrast

Figure 3: Local Context Analysis concepts for query 214

collection	WEST	TREC3	TREC4
Number of queries	34	50	49
Raw text size in gigabytes	0.26	2.2	2.07
Number of documents	11,953	741,856	567,529
Mean words per document	1,970	260	299
Mean relevant documents per query	29	196	133
Number of words in a collection	23,516,042	192,684,738	169,682,351

Table 1: Statistics on text corpora

collection	Number of passages										
	10	20	30	40	50	100	200	300	500	1000	2000
TREC4	29.5	29.9	30.2	30.3	30.4	31.1	31.0	30.7	29.9	29.0	27.9
	+17	+18.6	+19.8	+20.3	+20.6	+23.5	+23.0	+21.8	+18.6	+15	+10.7
TREC3	36.6	37.5	38.7	39.0	38.9	38.9	39.3	39.1	38.3	37.6	36.6
	+16	+18.9	+22.6	+23.6	+23.2	23.3	+24.4	+23.7	+21.3	+19	+16.0
WEST	54.8	55.4	54.5	54.6	54.2	54.2	53.1	52.7	52.1	51.7	51.7
	+1.9	+3.0	+1.3	+1.6	+0.7	+0.8	-1.3	-2.0	-3.2	-3.9	-3.9

Table 2: Performance of local context analysis using 11 point average precision

collection	Number of passages										
	10	20	30	40	50	100	200	300	500	1000	2000
WEST	55.9	56.5	55.6	55.7	55.8	55.6	54.6	54.4	53.6	53.7	53.7
	+3.8	+5.0	+3.4	+3.6	+3.7	+3.3	+1.6	+1.2	-0.4	-0.1	-0.1

Table 3: Downweight expansion concepts of local context analysis on WEST. The weight of the auxiliary query is reduced to 1.0

better than the sum function used by Phrasefinder because with the product function it is harder for some query terms to dominate other query terms.

6 Local Text Analysis vs Local Feedback

In this section we compare the retrieval performances of local feedback and local context analysis. Table 7 shows the retrieval performance of local feedback.

Table 8 shows the result of downweighting the expansion concepts by 50% on WEST. The reason for this is to make a fair comparison with local context analysis. Remember that we also downweighted the expansion concepts of local context analysis by 50% on WEST.

Local feedback does very well on TREC3. The best run produces a 20.5% improvement over the baseline, close to the 24.4% of the best run of local context analysis. It is also relatively insensitive to the number of documents used for feedback on TREC3. Increasing the number of documents from 10 to 50 does not affect performance much.

It also does well on TREC4. The best run produces a 14.0% improvement over the baseline, very significant, but lower than the 23.5% of the best run of local context analysis.

It is very sensitive to the number of documents used for feedback on TREC4. Increasing the number of documents from 5 to 20 results in a big performance loss. In contrast, local context analysis is relatively insensitive to the number of passages on all three collections.

On WEST, local feedback does not work at all. Without downweighting the expansion concepts, it results in a significant performance loss over all runs. Downweighting the expansion concepts only reduces the amount of loss. It is also sensitive to the number of documents used for feedback. Increasing the number of feedback documents results in significantly more performance loss.

It seems that the performance of local feedback and its sensitivity to the number of documents used for feedback depend on the number of relevant documents in the collection for the query. From table 1 we know that average number of relevant documents per query on TREC3 is 196, larger than 133 of TREC4, which is in turn larger than 29 of WEST. This corresponds to the relative performance of local feedback on the collections.

Tables 4-6 show a side by side comparison between local feedback and local context analysis at different recall levels on the three collections. Top 10 documents are used for local

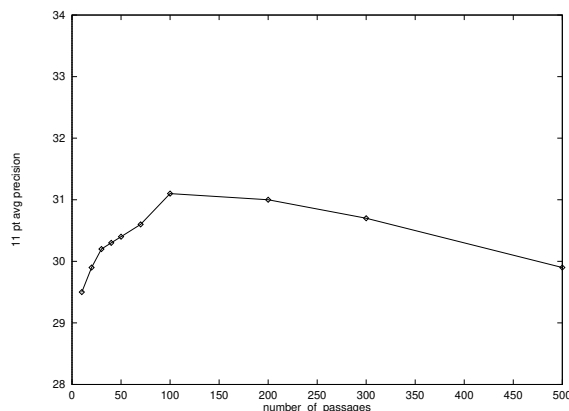


Figure 4: Performance curve of local context analysis on TREC4

Recall	Precision (% change) - 49 queries			
	base	Phrasefinder	lf-10doc	lca-100p
0	71.0	68.6 (-3.3)	68.4 (-3.6)	73.2 (+3.2)
10	49.3	48.6 (-1.6)	52.8 (+7.0)	57.1 (+15.7)
20	40.4	40.0 (-1.0)	43.2 (+7.0)	46.8 (+16.0)
30	33.3	33.9 (+1.8)	36.0 (+8.0)	39.9 (+19.8)
40	27.3	28.0 (+2.5)	29.8 (+9.2)	35.3 (+29.1)
50	21.6	23.9 (+10.3)	24.5 (+13.2)	29.9 (+38.4)
60	14.8	18.8 (+27.1)	19.7 (+33.4)	23.6 (+59.8)
70	9.5	11.8 (+24.7)	14.8 (+56.9)	17.9 (+89.1)
80	6.2	8.1 (+31.0)	10.8 (+74.7)	11.8 (+91.0)
90	3.1	4.2 (+33.6)	6.4 (+104.6)	5.7 (+80.2)
100	0.4	0.6 (+24.0)	0.9 (+93.3)	0.8 (+88.2)
average	25.2	26.0 (+3.4)	27.9 (+11.0)	31.1 (+23.5)

Table 4: A comparison of baseline, Phrasefinder, local feedback and local context analysis on TREC4. 10 documents for local feedback (lf-10doc). 100 passages for local context analysis (lca-100p)

feedback and top 100 passages are used for local context analysis in these tables. In table 6 for WEST, the expansion concepts are downweighted by 50% for both local feedback and local context analysis.

We also made a query-by-query comparison of the best run of local feedback and the best run of local context analysis on TREC4. Of 49 queries, local feedback hurts 21 and improves 28, while local context analysis hurts 11 and improves 38. Of the queries hurt by local feedback, 5 queries have a more than 5% percent loss in average precision. The worst case is query 232, whose average precision is reduced from 24.8% to 4.3%. Of those hurt by local context analysis, only one has a more 5% percent loss in average precision. Local feedback also tends to hurt queries with poor performance. Of 9 queries with baseline average precision less than 5%, local feedback hurts 8 and improves 1. In contrast, local context analysis hurts 4 and improves 5. Its tendency to hurt “bad” queries and queries with few relevant documents (such as the WEST queries) suggests that local feedback is very sensitive to the number of relevant documents in the top ranked documents. In comparison, local context analysis is not so sensitive.

It is interesting to note that although both local context analysis and local feedback find concepts from top ranked passages/documents, the overlap of the concepts chosen by them is very small. On TREC4, the average number of unique terms in the expansion concepts per query is 58 by local feedback and 78 by local context analysis. The average overlap per query is only 17.6 terms. This means local

context analysis and local feedback are two quite different query expansion techniques. Some queries expanded quite differently are improved by both methods. For example, the expansion overlap for query 214 of TREC4 (“What are the different techniques used to create self-induced hypnosis”) is 19 terms, yet both methods improve the query significantly.

7 Conclusion and Future Work

This paper compares the retrieval effectiveness of three automatic query expansion techniques: global analysis, local feedback and local context analysis. Experimental results on three collections show that local document analysis (local feedback and local context analysis) is more effective than global document analysis. The results also show that local context analysis, which uses some global analysis techniques on the local document set outperforms simple local feedback in terms of retrieval effectiveness and predictability.

We will continue our work in these aspects:

1. local context analysis: automatically determine how many passages to use, how many concepts to add to the query and how to assign the weights to them on a query by query basis. Currently the parameter values are decided experimentally and fixed for all queries.
2. Phrasefinder: a new metric for selecting concepts. Currently Phrasefinder uses Inquiry’s belief function, which is not designed to select concepts. We

Recall	Precision (% change) – 50 queries						
	base	Phrasefinder		lf-10doc		lca-100p	
0	82.2	79.4	(-3.3)	82.5	(+0.4)	87.0	(+5.9)
10	57.3	60.1	(+4.8)	64.9	(+13.3)	65.5	(+14.3)
20	46.2	50.4	(+9.1)	56.1	(+21.5)	57.2	(+23.8)
30	39.1	43.3	(+10.7)	48.3	(+23.5)	48.4	(+23.8)
40	32.7	36.9	(+12.8)	41.6	(+26.9)	42.7	(+30.4)
50	27.5	31.8	(+15.9)	36.8	(+34.1)	37.9	(+38.0)
60	22.6	26.1	(+15.1)	30.9	(+36.7)	31.5	(+39.3)
70	18.0	20.6	(+14.0)	25.2	(+40.0)	25.6	(+42.1)
80	13.3	15.8	(+18.6)	19.4	(+45.7)	19.4	(+45.7)
90	7.9	9.4	(+18.7)	11.5	(+44.3)	11.7	(+47.3)
100	0.5	0.8	(+60.9)	1.2	(+143.5)	1.4	(+177.0)
average	31.6	34.1	(+7.8)	38.0	(+20.5)	38.9	(+23.3)

Table 5: A comparison of baseline, Phrasefinder, local feedback and local text analysis on TREC3. 10 documents for local feedback (lf-10doc). 100 passages for local context analysis (lca-100p)

Recall	Precision (% change) – 34 queries				
	base	lf-10doc-dw0.5	lca-100p-w1.0		
0	88.0	81.9	(-7.0)	92.1	(+4.7)
10	80.0	76.9	(-4.0)	84.3	(+5.4)
20	77.5	71.4	(-7.8)	78.5	(+1.3)
30	74.1	68.2	(-7.9)	73.9	(-0.1)
40	62.9	60.8	(-3.3)	61.8	(-1.7)
50	57.5	56.8	(-1.2)	56.8	(-1.2)
60	49.7	50.1	(+0.8)	50.7	(+2.2)
70	41.5	42.1	(+1.3)	44.2	(+6.4)
80	32.7	33.1	(+1.1)	36.4	(+11.2)
90	19.3	21.8	(+13.0)	22.6	(+17.1)
100	8.6	9.3	(+7.8)	10.0	(+15.3)
average	53.8	52.0	(-3.3)	55.6	(+3.3)

Table 6: A comparison of baseline, local feedback and local text analysis on WEST. 10 documents for local feedback with weights for expansion units downweighted by 50% (lf-10doc-dw0.5). 100 passages for local context analysis with weight for auxiliary query set to 1.0 (lca-100p-w1.0).

hope a better metric will improve the performance of Phrasefinder.

8 Acknowledgements

We thank Dan Nachbar and James Allan for their help during this research. This research is supported in part by the NSF Center for Intelligent Information Retrieval at University of Massachusetts, Amherst.

This material is based on work supported in part by NRD Contract Number N66001-94-D-6054. Any opinions, findings and conclusions or recommendations expressed in this material are the author(s) and do not necessarily reflect those of the sponsor.

References

- [Attar & Fraenkel, 1977] Attar, R., & Fraenkel, A. S. (1977). Local Feedback in Full-Text Retrieval Systems. *Journal of the Association for Computing Machinery*, 24(3), 397–417.
- [Buckley et al., 1996] Buckley, C., Singhal, A., Mitra, M., & Salton, G. (1996). New Retrieval Approaches Using SMART : TREC 4. In Harman, D., editor, *Proceedings of the TREC 4 Conference*. National Institute of Standards and Technology Special Publication. to appear.
- [Caid et al., 1993] Caid, B., Gallant, S., Carleton, J., & Sudbeck, D. (1993). HNC Tipster Phase I Final Report. In *Proceedings of Tipster Text Program (Phase I)*, pp. 69–92.
- [Callan et al., 1995] Callan, J., Croft, W. B., & Broglio, J. (1995). TREC and TIPSTER experiments with IN-QUERY. *Information Processing and Management*, pp. 327–343.
- [Callan, 1994] Callan, J. P. (1994). Passage-level evidence in document retrieval. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 302–310.
- [Croft et al., 1995] Croft, W. B., Cook, R., & Wilder, D. (1995). Providing Government Information on The Internet: Experiences with THOMAS. In *Digital Libraries Conference DL'95*, pp. 19–24.
- [Croft & Harper, 1979] Croft, W. B., & Harper, D. J. (1979). Using probabilistic models of document retrieval without relevance information. *Journal of Documentation*, 35, 285–295.
- [Crouch & Yang, 1992] Crouch, C. J., & Yang, B. (1992). Experiments in automatic statistical thesaurus construction. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 77–88.

collection	number of documents used					
	5	10	20	30	50	100
TREC4	28.7	27.9	26.9	27.2	26.7	26.1
	+14.0	+11.0	+6.8	+8.2	+6.2	+3.5
TREC3	36.6	38.0	37.6	37.7	37.7	36.6
	+16.0	+20.5	+19.1	+19.4	+19.3	+15.8
WEST	49.6	49.8	46.2	44.1	40.0	35.1
	-7.8	-7.5	-14.2	-18.0	-25.6	-34.7

Table 7: Performance of local feedback using 11 point average precision.

collection	number of documents used					
	5	10	20	30	50	100
WEST	52.6	52.0	48.7	47.5	44.5	40.0
	-2.2	-3.3	-9.5	-11.6	-17.2	-25.7

Table 8: Downweight expansion concepts of local feedback by 50% on WEST

[Deerwester et al., 1990] Deerwester, S., Dumais, S., Furnas, G., Landauer, T., & Harshman, R. (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science*, 41,391–407.

[Harman, 1995] Harman, D. (1995). Overview of the Third Text REtrieval Conference (TREC-3). In Harman, D., editor, *Proceedings of the Third Text REtrieval Conference (TREC-3)*, pp. 1–20. NIST Special Publication 500-225.

[Harman, 1996] Harman, D., editor (1996). *Proceedings of the TREC 4 Conference*. National Institute of Standards and Technology Special Publication. to appear.

[Jing & Croft, 1994] Jing, Y., & Croft, W. B. (1994). An association thesaurus for information retrieval. In *Proceedings of RIAO 94*, pp. 146–160.

[Qiu & Frei, 1993] Qiu, Y., & Frei, H. P. (1993). Concept based query expansion. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 160–169.

[Schütze & Pedersen, 1994] Schütze, H., & Pedersen, J. O. (1994). A cooccurrence-based thesaurus and two applications to information retrieval. In *Proceedings of RIAO 94*, pp. 266–274.

[Sparck Jones, 1971] Sparck Jones, K. (1971). *Automatic Keyword Classification for Information Retrieval*. Butterworth, London.

[Turtle, 1994] Turtle, H. (1994). Natural language vs. Boolean query evaluation: A comparison of retrieval performance. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 212–220.

[Voorhees, 1994] Voorhees, E. (1994). Query expansion using lexical-semantic relations. In *Proceedings of ACM SIGIR International Conference on Research and Development in Information Retrieval*, pp. 61–69.