Bodner, R. and Song, F. (1996) Knowledge-based approaches to query expansion in information retrieval. In McCalla, G. (Ed.), *Advances in Artificial Intelligence* (pp. 146-158). New York: Springer.

Knowledge-Based Approaches to Query Expansion in Information Retrieval

Richard C. Bodner and Fei Song
Department of Computing & Information Science
University of Guelph
Guelph, Ontario, Canada N1G 2W1
{rbodner, fsong}@snowhite.cis.uoguelph.ca

Textual information is becoming increasingly available in electronic forms. Users need tools to sift through non-relevant information and retrieve only those pieces relevant to their needs. The traditional methods such as Boolean operators and key terms have somehow reached their limitations. An emerging trend is to combine the traditional information retrieval and artificial intelligence techniques. This paper explores the possibility of extending traditional information retrieval systems with knowledge-based approaches to automatically expand natural language queries. Two types of knowledge-bases, a domain-specific and a general world knowledge, are used in the expansion process. Experiments are also conducted using different search strategies and various combinations of the knowledge-bases. Our results show that an increase in retrieval performance can be obtained using certain knowledge-based approaches.

1 Introduction

The abundance of information available to a user can be overwhelming. Users generally require tools to help them sift through large collections of information and retrieve only those items of interest. The field of information retrieval (IR) is the study of such tools. Although systems have been built with retrieval performance comparable to that of manual methods, further improvement has been slow and difficult [Salton, 1986]. One possible problem is the query size. Many systems do not distinguish between small and large queries. This is a problem since not all queries are created equal. The performance of an IR system is typically proportional to the size of a query [Qui and Frei, 1993]. Long queries are capable of providing enough information for the system to perform reasonably well. Short, especially vague and ill-formed, queries do not provide enough information and are therefore prone to poor retrieval performance. When the results of the two types of queries are averaged, the overall performance of an IR system will be modest at best.

Users may be incapable of formulating long well-defined queries because they do not know much about the problem domain [Brooks, 1987]. Even if they know the domain, they may not want to spend the time entering those long queries. As a result, users may only be able to provide approximate descriptions of their information needs, or simply enter short descriptions, resulting in poor system performance. Thus, there are strong needs for the system to fill in additional information to help users formulate more accurate queries.

This paper explores knowledge-based approaches to query expansion in information retrieval. We assume the conventional vector space model for retrieving information, but queries are first expanded to help improve the retrieval performance.

Both domain-specific and general world knowledge are used and combined in the expansion process. Our goal is to provide an automatic tool so that the user will not be burdened with too much involvement. We further conduct experiments on a standard test collection to investigate the effectiveness of our approaches. Two major results we want to establish are the size of a query that could benefit from such a query expansion process, and the performance characteristics of different search and combination strategies for the two knowledge sources.

2 Information Retrieval and Query Expansion

Many recent IR systems are built on the popular vector space model, which is capable of producing a ranking of the retrieved documents (see Fig. 1, from [Croft, 1993]). The model represents both documents and queries as vectors in a high dimensional space [Turtle and Croft, 1992]. Each dimension corresponds to a feature (or a keyword) of the text. A document vector can be created by removing from text all function words (e.g., and, of, that, etc.) and reducing the remaining terms to their root forms. Based on its frequency in the document, a feature can be assigned a weight, indicating the relative importance of the feature in the vector. A query can be processed roughly in the same manner. A similarity between a document and a query can then be determined, typically by calculating the inner product of the two corresponding vectors. Based on the similarity measure, the retrieved documents are ranked. The vector space model has been shown experimentally to have better performance over the earlier Boolean method [Turtle and Croft, 1992], but the further improvement has been difficult. Only relatively small increases have been obtained over the years of IR research.

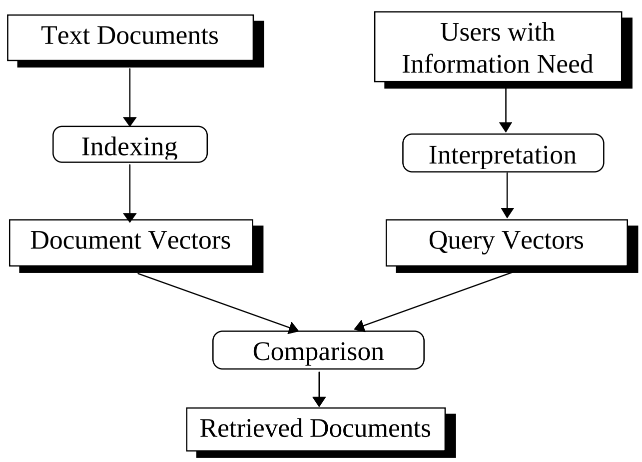

Fig. 1. A General framework for information retrieval

Query expansion methods have been investigated for almost as long as the study of information retrieval. The techniques developed can be classified as user-assisted or automatic. One well-known user-assisted technique is relevance feedback [Salton and Buckley, 1990], which requires a user to iteratively judge the relevance of a set of retrieved documents. The documents identified as relevant are used to refine the original query, and the search process continues until the user is satisfied with the retrieval results. Relevance feedback is a powerful technique, with an improvement of up to 90% for a single search iteration being reported [Salton and Buckley, 1990]. One major advantage of this method is that it lessens the burden on the user to reformulate a query. One major disadvantage is that it does not improve the retrieval performance of the original query. The user has to invest time and effort judging the relevance of retrieved documents before any improvements can be made.

Another user-assisted technique used in most commercial IR systems [Smeaton and van Rijsbergen, 1983] is to incorporate a browsable thesaurus. Given a set of search terms, a user is presented with a list of similar terms which he/she can choose to replace some existing terms or use as additional terms. Once again, this technique relies on the user's knowledge about the problem domain and the user's ability to judge what is an effective search term.

The automatic techniques for query expansion do not rely on users to make relevance judgments, and are often based on language analysis [Sparck-Jones and Tait, 1984] and term co-occurrences [Qui and Frei, 1993]. Language analysis approaches require a deep understanding of queries and documents, usually at higher computational costs. These techniques have also been shown to have only small improvements in retrieval performance.

The co-occurrence approaches can be grouped into four categories [Qui and Frei, 1993]: term classification, document classification, syntactic context, and relevance information. The term classification methods place a term into a class based on its similarity measures with other terms in the class. Expansion is done by matching a search term with a similar class and adding the terms from the class to the query. This process is employed after the indexing of the documents and a query. The document classification is similar to the term classification except that it is used during the indexing phase so that the representations for the query and documents can be enhanced by replacing (or adding) terms from a thesaurus class. The syntactic context methods make use of linguistic knowledge to enhance search terms, and the relevance information methods include the relevance feedback technique.

3 Knowledge-Based Approaches to Query Expansion

We extend the general framework for information retrieval to include a query expansion subsystem. The resulting system is a blending of conventional (statistical) information retrieval and artificial intelligence techniques. It still uses the vector space model to perform the document and query indexing, compare the indexed vectors, and retrieve relevant documents. The artificial intelligence techniques are employed after the queries are interpreted by the conventional methods, and before they are compared with the document collection, as shown in Fig. 2.

The query expansion subsystem contains domain-specific and general world knowledge-bases, both of which are represented as semantic networks. Given a query, we first separate the keywords from the function words. The keywords are then used to search in the semantic networks for additional related words. These additional words are further combined with the original keywords to form an expanded query. Thus, our query expansion process can include additional information that a user is ignorant of and improve the retrieval performance of the system.

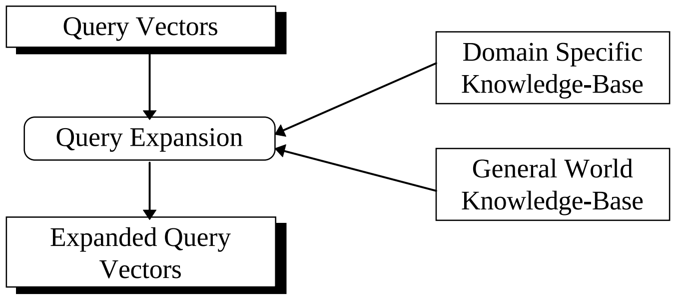

Fig. 2. Query expansion subsystem

3.1 Knowledge-Bases

The main components in our query expansion subsystem are the domain-specific and general world knowledge-bases. Given a document collection, there is often no existing knowledge-base available for the domain, so one has to be created. Since our query expansion process is intended to be used with many document collections, a manual process for creating the domain-specific knowledge-base is not acceptable. Fortunately, there exists a statistical method that can be used to create a domain-specific knowledge-base for a collection of documents [Forsyth and Rada, 1986]. The method works on two assumptions. First, terms with high frequencies have broader meanings and terms with low frequencies have narrower meanings. Second, if the density functions of two terms have the same shape and one term is of a lower frequency, the lower frequency term becomes a descendent of the higher frequency term. Based on these two assumptions, a procedure for creating hierarchies of terms can be given as follows [Frakes and Baeza-Yates, 1992]:

1. All terms in the document collection are grouped into different classes based on their frequencies, with the highest frequency terms at the root and the lowest frequency terms at the leaves.

2. Descendent links (IS-A) are calculated for each class. The similarity between each term in a class is computed with each term in the class above. This allows for multiple parents.

3. All terms without descendants in the class above, are propagated down.

4. Steps 2 and 3 are repeated for all classes.

This procedure results in several hierarchies. These hierarchies are joined together at various classes creating a semantic network. A portion of the semantic net is presented in Fig. 3. A node contains a stemmed term and its frequency in the document collection.

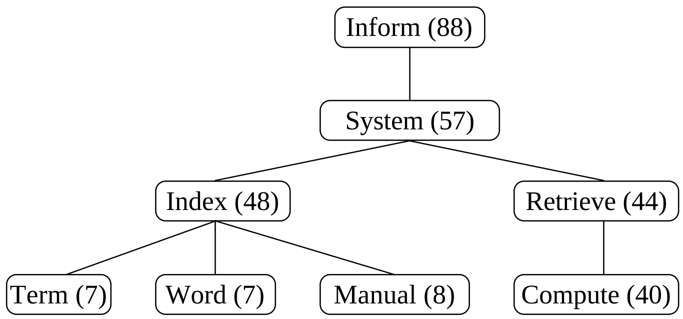

Fig. 3. Portion of the domain-specific knowledge-base

For the general world knowledge-base, we adopt a manually built on-line dictionary, called WordNet [Miller, 1990]. WordNet contains approximately 95,600 word forms, organized into some 70,100 word meanings. The word meanings are joined together to form a semantic network. WordNet differs from standard dictionaries in that the lexicon is divided into four categories: nouns, verbs, adjectives, and adverbs, and there are links between word meanings. Of the four, the noun category is the most robust in terms of the words stored and the links between them. The links are semantic relations, including antonymy, hypernymy/hyponymy (IS-A), and meronym/holonym (PART-OF). It is for this reason that we decide to only use nouns in the general world knowledge-base for our query expansion process. The general world knowledge-base is meant to augment (or “fill in the gaps for”) the domain-specific knowledge-base.

The major difference between the two knowledge-bases is in the organization. The general world knowledge-base is organized around words with similar or related meanings. Due to this organization, the general world knowledge-base contains many more links, besides the IS-A link. The domain-specific knowledge-base uses term co-occurrence as the organization method. This method can distinguish whether two terms are related but it cannot distinguish how [Forsyth and Rada, 1986]. Due to this shortcoming, the knowledge-base may contain extraneous links between terms which are not truly related.

3.2 Search Strategies

Since both knowledge-bases are organized as semantic networks, a constrained spreading activation search can be used. The two constraints imposed on all searches are the distance to travel away from a starting node and the fan-out of a node. The fan-out represents the number of outgoing links from a node. Nodes with a high fan-out are considered to be too general (usually close to the root) and should be avoided.

The search strategies are divided in two categories: isolated and correlated searches. The isolated searches handle the terms separately, ignoring any relationships a term may have with other terms in a user's query. There are three types of isolated searches:

1. Searching only broader terms - following links up the hierarchy towards the root(s). Searching the broader terms should expand a query to allow matching with a greater range of documents. Increasing the scope of a query should enhance the recall.

2. Searching only narrower terms - following links down the hierarchy towards the leaves. Including narrower terms in the expansion should narrow the number of documents that will match the query. The decrease in the scope of the query will not be very large since more terms are being added, but it is thought to be enough so that precision is enhanced.

3. Searching links in both directions. This search is designed to find as many relevant documents as possible by adding all related terms within a specific distance. It is mainly included for completeness.

The correlated search looks at consecutive terms in a query. A technique known as the marker passing is used to search for a path between two terms in a knowledge-base. Such a search can elicit more specific information from the query and thus should be able to produce more meaningful terms for the expansion process. Searching only the narrower or broader terms is not used for the correlated search since it is too restrictive (see Fig. 4). The search can be envisioned as circles radiating from two points. If the circles overlap, a path is found. The possibility for overlapping in a strictly narrower or broader search is greatly reduced.

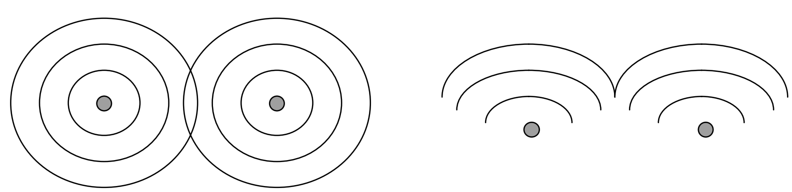

Fig. 4. Circular vs. narrow search for two terms

3.3 Knowledge-Base Combination Strategies

Since there are two knowledge-bases to be used in the query expansion process, we need to decide how they interact with each other. As mentioned above, the general world knowledge-base is intended to assist the domain-specific knowledge-base. This is reflected in the combination strategies we developed, including union, chaining, and mapping.

There are two ways of combining the knowledge-bases using the union strategy: weighted and unweighted. The choice of using a weighted or unweighted union pertains to the intersection portion of the union. For the weighted union, each term in the two expanded sets are associated with a weight (obtained from the original query). If the expanded sets intersect, the weights of the terms are added together. This is done to reflect the assumption that if a term comes from two knowledge sources it must be important. An unweighted union ignores the possibility that some terms are more important than others (similar reasons as the isolated searches). As in the strict sense of a union, the terms are only counted once. The expanded term weights are calculated after the expansion process.

The chaining strategy uses one knowledge-base as the primary source and the other as the secondary source for the expansion. The primary knowledge source is used first to expand the original query, and the result is further expanded through the secondary knowledge source. The primary knowledge source should be the one which is more relevant to the original query. The final expanded query through both knowledge sources is used to retrieve relevant documents.

The mapping strategy is the most complex combination method. Again, one knowledge-base is assigned to be the primary source and is responsible for the expansion of the original query. The secondary knowledge-base is not directly used in the expansion process; instead it is used to assist the primary source by mapping the terms from the original query to those that can be found in the primary knowledge-base. Thus, the mapping process will be useful when the terms in the original query can be found in the secondary knowledge source. In many cases, the primary knowledge source should be the domain-specific knowledge-base since it is more relevant to the problem area and the general world knowledge-base should be used to only assist in the expansion process.

4 Experimental Results and Discussions

To test the effectiveness of our query expansion process, experiments are conducted using the standard ADI document collection. This is a homogeneous collection covering topics in Information Science. Along with the documents are a set of standard queries. Each query is associated with a set of relevant documents. This allows us to compare different information retrieval systems in terms of retrieval performance. Table 1 lists the statistics of the ADI collection (taken from [Crouch and Yang, 1992]).

Table 1. Statistics for the ADI collection

Subject Area	Information Science
Type	Homogeneous
No. of Documents	82
No. of Queries	35
No. of Terms	822
Mean No. of Terms per Document	25.5

Mean No. of Terms per Query	7.1

Based on the standard test collection, the query expansion process is evaluated using the well known performance measures of precision and recall. Precision (P) and recall (R) are defined as [Salton and McGill, 1983]:

$$\text{Recall} = \frac{\text{Number of Items Retrieved and Relevant}}{\text{Number of Total Relevant Items}}$$

$$\text{Precision} = \frac{\text{Number of Items Retrieved and Relevant}}{\text{Number of Total Retrieved Items}}$$

Fig. 5. Equations for precision and recall

Recall measures the ratio of relevant documents retrieved for a user's query to the total number of relevant documents in the collection. Precision measures the ratio of relevant documents retrieved to the total retrieved documents. Retrieval performance can be represented as a precision-recall graph or a single composite measure. An 11-point precision average (at recall levels of 0.0, 0.1, ..., 1.0) is often used to calculate a composite measure for retrieval performance.

4.1 Design of the Experiments

There are four major parameters in our experiments: (1) the best setting for search constraints of fan-out and distance, (2) the sizes of queries which can benefit from the query expansion process, (3) different search strategies, and (4) different knowledge-base combination strategies. The best values for fan-out and distance constraints correspond to the combination that gives the optimal results for the queries we tested. The values are decided through experiments, by first finding the best fan-out and then using this value to obtain the best distance. These constraints need to be identified for each knowledge-base used in the expansion process. Once this information is available, experiments for the search strategies can be started. While conducting the experiments for the search strategies, data are also collected to determine which queries could benefit from such a process. The results of the experiments help us select the best search strategy, which can then be used for testing different knowledge-base combination strategies.

4.2 Best Search Constraint Values

The best values for fan-out and distance are obtained by searching around from the values of 15 for fan-out and 5 for distance [Liu, 1995]. Both isolated and correlated search strategies are used to expand a query to test for different values of fan-out and distance. The expanded set of terms are added to the original query vector using the weighted union method.

The results of our experiments show that a fan-out of 10 and distance of 4 are the best for the domain-specific knowledge-base. For the general world knowledge-base, a fan-out of 15 and distance of 5 are found to be the best. The difference between the fan-out settings can be attributed to the difference in the sizes of the knowledge-bases. The general world knowledge-base is much larger than the domain-specific knowledge-base. The former contains approximately 70,100 nodes for word meanings, as compared to only 822 nodes in the latter. Since the general world knowledge-base is structured as interconnected hierarchies, many nodes near the root will have a high fan-out value.

The size of the knowledge-base does not affect the distance constraint. Indeed, the performance of the general world knowledge-base using a distance of 4 only reduces overall performance by less than one percent. A distance value greater than 5 for either knowledge-base decreases the performance. This is due to the fact that both knowledge-bases represent relationships between words. Words found beyond a distance of 5 links are not strongly related and thus weaken the retrieval performance.

4.3 Search Strategies and Query Sizes

Once the best values for the fan-out and distance are found, the search strategies can be evaluated. Table 2 shows the results of the search strategies for short queries (6-8 terms) and for all queries (3-14 terms). Note that the isolated searches are mainly tested on the domain-specific (DS) knowledge-base, since it produces better results, but the correlated searches are tested on both domain-specific and general world (GW) knowledge-bases for comparisons.

In all the experiments, the results of a conventional IR system based on the vector space model is used as the control. All the isolated search strategies perform poorly. In particular, the isolated narrower search has a performance decrease of 15% for all query sizes from the conventional IR system. Such a decrease is probably due to two factors: query expansion is mainly a recall enhancing technique [Voorhees, 1994] and these strategies treat all words in a query with equal importance. Treating all words the same is not realistic and the problem could be amplified when the query is ill-formed and vague. A user may only understand a portion of his/her domain and may add extraneous information due to this lack of knowledge, resulting in decreased performance.

Table 2. Results for different search strategies

Search Strategy	**11-Point Precision Average**	
	All Queries	Short Queries
Conventional IR System	.41	.39
Isolated search - broader (DS)	.34	.35
Isolated search - narrower (DS)	.35	.39
Isolated search - both (DS)	.31	.30
Correlated - marker passing (DS)	.44	.48

Correlated - marker passing (GW)	.425	.44

The correlated search (or marker passing) strategy has the best performance for all queries, with an increase of near 4% for the general world knowledge-base and over 7% for the domain-specific knowledge-base. This is mainly due to the added constraint that the two terms must be related in the knowledge-base. This constraint addresses the major weakness of the isolated searches. Not all terms in the query are treated equally. If no path exists between two consecutive terms in a query, it is assumed that there is no relationship between the terms and nothing is added to the query. Only terms which have relationships are considered as important and used in the expansion process.

While conducting the experiments on search strategies, we also collected data to identify the sizes of queries which will benefit most from the expansion process. More specifically, we partitioned the queries into three groups: very short queries with fewer than 6 terms, short queries with 6-8 terms, and long queries with more than 8 terms. Note that the terms represented only the keywords, with the function words removed. We found that for the ADI collection, short queries with 6-8 terms or approximately three sentences benefited the most, with the highest increase of performance being 23% for the domain-specific knowledge-base. It is surprising that very short queries with fewer than 6 terms do not benefit much from the expansion process. One reason may be that queries this short do not contain enough information to work from (i.e., no paths can be found between the words). Long queries with more than 8 terms do not benefit much from the expansion process either. In fact some larger queries experience a decrease in the retrieval performance. This can likely be attributed to two reasons. The first and also the more likely reason is that a larger query may already contain enough information to perform adequately. The second reason is the opposite of the first. If a larger query is ill-formed, applying the expansion process may move it even further away from a user's true information need. This is a problem for all queries, but is amplified for larger ones.

4.4 Combination of the Knowledge-Bases

As mentioned previously, the general-world knowledge-base is intended to provide additional information and thus needs to be used in combination with the domain-specific knowledge-base. The results of the five knowledge-base combination strategies are listed in Table 3.

Table 3. Results for different combination strategies

Combination Strategy	**11-point Precision Average**	
	All Queries	Short Queries
Conventional IR System	.41	.39
Union (weighted)	.42	.46
Union (unweighted)	.39	.46
Chaining (DS first)	.43	.48
Chaining (GW first)	.43	.46
Mapping (DS first)	.43	.48

Since the performance of the correlated search strategy surpasses that of the conventional IR system, it is used to search the knowledge-bases and the results are combined using either the union or chaining method. The mapping method uses only one knowledge-base to perform the search, and the domain-specific knowledge-base is used as the primary knowledge-base.

The unweighted union combination (a decrease of 5%) has poorer performance for all queries than the weighted union (an increase of 2%) because the importance of the terms is not considered. For the unweighted union, only one occurrence of a term is used. The fact that a term occurs more than once in the expansion signifies that it is probably more important than other terms which only appear once. The weighted union takes into consideration the possible importance of a term by combining the weights of the duplicate terms. This increases the importance of the term in the expanded vector.

The general world knowledge-base is only intended to augment the domain-specific knowledge-base. The knowledge it contains is too vague to make it a truly effective knowledge source on its own (e.g., for the correlated search strategy; the domain-specific knowledge-base has a much higher performance for short queries). This is shown to be true for the chaining combination strategy. When the general world knowledge-base is used first, performance for short queries is decreased (18% increase in performance, as compared to 23% increase when the domain-specific knowledge-base is used first). The general world knowledge-base adds too many general words and when the domain-specific knowledge-base finishes the query expansion process, it only increases the problem. Using the general world knowledge-base after the domain-specific one proves to be the best. This combination minimizes the problems with the general-world knowledge-base. The mapping combination method performs well because it also emphasizes the use of the domain-specific knowledge-base.

4.5 Comparisons

The composite measure of 11-point precision average is useful in comparing overall performance of an information retrieval system. However, the measure is not useful for identifying characteristics (e.g., favoring high recall) of the search or combination strategies. To see these characteristics it is necessary to view the precision-recall graphs. Of particular interest are the graphs of the correlated search using the domain-specific knowledge-base and the mapping combination method with the domain-specific knowledge-base applied first.

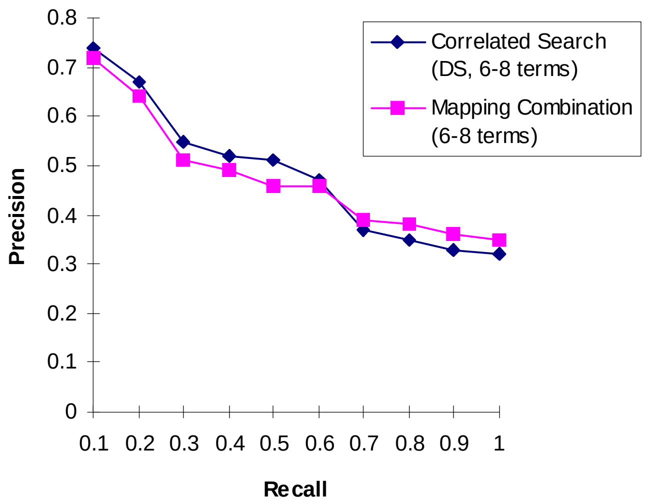

Fig. 6. P-R graph for mapping combination and correlated search

Both strategies have the same performance for short queries (.48), but looking at the graphs it can be seen that the two methods do not have the same performance characteristics. The mapping strategy favors high recall while the correlated search favors low to middle recall. Knowing these characteristics is important for the query expansion process. For example, if the user were a lawyer preparing for a case, he/she would need to know about all relevant material and therefore world require an expansion process which favors a high recall search.

5 Conclusions and Future Directions

Our experiments have shown that an automatically created domain-specific knowledge-base can be an effective tool for query expansion. The general world knowledge-base can also be effective when used with the correlated search (or marker passing) method and either a chaining or mapping combination strategy. Thus, we have demonstrated that shallow artificial intelligence techniques can provide a feasible and useful extension to conventional information retrieval systems. The blending of the two kinds of techniques can increase retrieval performance while still remain efficient in terms of the computational speed. Investigating the retrieval results has revealed that there are different performance characteristics of the various search and combination strategies. These characteristics can be used to provide users with better supports in matching their information needs.

Our query expansion process is intended to expand a user's initial query, and thus can still be used with relevance feedback to further improve the retrieval performance. It has been shown that relevance feedback can produce large gains in retrieval performance [Salton and Buckley, 1990]. However, if no relevant documents are retrieved with the user's initial query, the process would fail. As a result, the user may be forced to enter a different query in order to elicit relevant documents. This can be difficult if the user has limited knowledge in the subject area. Our query expansion process can provide assistance to an inexperienced user by providing additional information that the user is ignorant of. When used with relevance feedback, it will likely be more effective since the user will start with more relevant documents.

Future directions for this research involve testing the query expansion method on other standard collections. Experiments have already begun on the TIME and CACM collections. The CACM is a collection of homogeneous documents about computer science and the TIME is a collection of heterogeneous documents, excerpted from the Time magazine during the 1960's. With the heterogeneous collection, it is hoped that the general world knowledge-base will be of more use, especially for the mapping combination strategy. Also for future work, we intend to add a user modeling component, which would allow us to choose the most appropriate search and/or combination strategies for a user.

Acknowledgments

This research is supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. We gratefully acknowledge this support without which this research would not have been possible.

References

[Brooks, 1987] H. M. Brooks. Expert Systems and Intelligent Information Retrieval. *Information Processing & Management*, 23(4):367-382, 1987.

[Croft, 1993] W. B. Croft. Knowledge-Based and Statistical Approaches to Text Retrieval. *IEEE Expert*, pp. 8-12, April 1993.

[Crouch and Yang, 1992] C. J. Crouch and B. Yang. Experiments in Automatic Statistical Thesaurus Construction. *SIGIR'92, Proceedings of the Fifteenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pp. 77-87, 1992.

[Forsyth and Rada, 1986] R. Forsyth and R. Rada. *Machine Learning, applications in expert systems and information retrieval*. Ellis Horwood Limited, 1986.

[Frakes and Baeza-Yates, 1992] W. B. Frakes and R. Baeza-Yates. *Information Retrieval: Data Structures & Algorithms*. Prentice Hall, New Jersey, 1992

[Liu, 1995] Y. Liu. *Statistical and Knowledge Bases Approaches for Sense Disambiguation in Information Retrieval*. Masters Thesis, Department of Computing and Information Science, University of Guelph, Guelph, Ontario.

[Miller, 1990] G. Miller. Special Issue, WordNet: An on-line lexical database. *International Journal of Lexicography*, 3(4), 1990.

[Qui and Frei, 1993] Y. Qui and H. P. Frei. Concept Base Query Expansion. *SIGIR'93, Proceedings of the Sixteenth Annual International ACM-SIGIR*

Conference on Research and Development in Information Retrieval, pp. 160-169, 1993.

[Salton, 1986] G. Salton. Another Look at Automatic Text-Retreival Systems. *Communications of the ACM*, 29(7):648-656, 1986.

[Salton and Buckley, 1990] G. Salton and C. Buckley. Improving Retrieval Performance by Relevance Feedback. *Journal of the American Society for Information Science*, 41(4):288-297, 1990.

[Salton and McGill, 1983] G. Salton and M. J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, New York, 1983.

[Smeaton and van Rijsbergen, 1983] A. F. Smeaton and C. J. van Rijsbergen. The Retrieval Effects of Query Expansion on a Feedback Document Retrieval System. *The Computer Journal*, 26(3):239-246, 1983.

[Sparck-Jones and Tait, 1984] K. Sparck-Jones and J. I. Tait. Automatic Search Term Variant Generation. *Journal of Documentation*, 40(1):50-66, March 1984.

[Turtle and Croft, 1992] H. R. Turtle and W. B. Croft. A Comparison of Text Retrieval Models. *The Computer Journal*, 35(3):279-290, 1992.

[Voorhees, 1994] E. Voorhees. Query Expansion Using Lexical-Semantic Relations. *SIGIR'94, Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pp. 61-69, 1994.