

# Complementing WordNet with Roget's and Corpus-based Thesauri for Information Retrieval

Rila Mandala, Takenobu Tokunaga and Hozumi Tanaka

Department of Computer Science  
Tokyo Institute of Technology  
2-12-1 Oookayama Meguro-Ku  
Tokyo 152-8522 Japan  
{rila,take,tanaka}@cs.titech.ac.jp

## Abstract

This paper proposes a method to overcome the drawbacks of WordNet when applied to information retrieval by complementing it with Roget's thesaurus and corpus-derived thesauri. Words and relations which are not included in WordNet can be found in the corpus-derived thesauri. Effects of polysemy can be minimized with weighting method considering all query terms and all of the thesauri. Experimental results show that our method enhances information retrieval performance significantly.

## 1 Introduction

Information retrieval (IR) systems can be viewed basically as a form of comparison between documents and queries. In traditional IR methods, this comparison is done based on the use of common index terms in the document and the query (Salton and McGill, 1983). The drawback of such methods is that if semantically relevant documents do not contain the same terms as the query, then they will be judged irrelevant by the IR system. This occurs because the vocabulary that the user uses is often not the same as the one used in documents (Blair and Maron, 1985).

To avoid the above problem, several researchers have suggested the addition of terms which have similar or related meaning to the query, increasing the chances of matching words in relevant documents. This method is called query expansion. A thesaurus contains information pertaining to paradigmatic semantic relations such as term synonymy, hypernymy, and hyponymy (Aitchison and Gilchrist, 1987). It is thus natural to use a thesaurus as a source for query expansion.

Many researchers have used WordNet (Miller, 1990) in information retrieval as a tool for query

expansion (Voorhees, 1994; Smeaton and Berrut, 1995), computing lexical cohesion (Stairmand, 1997), word sense disambiguation (Voorhees, 1993), and so on, but the results have not been very successful.

Previously, we conducted query expansion experiments using WordNet (Mandala et al., to appear 1999) and found limitations, which can be summarized as follows :

- Interrelated words may have different parts of speech.
- Most domain-specific relationships between words are not found in WordNet.
- Some kinds of words are not included in WordNet, such as proper names.

To overcome all the above problems, we propose a method to enrich WordNet with Roget's Thesaurus and corpus-based thesauri. The idea underlying this method is that the automatically constructed thesauri can counter all the above drawbacks of WordNet. For example, as we stated earlier, proper names and their interrelations are not found in WordNet, but if proper names bear some strong relationship with other terms, they often cooccur in documents, as can be modelled by a corpus-based thesaurus.

Polysemous words degrade the precision of information retrieval since all senses of the original query term are considered for expansion. To overcome the problem of polysemous words, we apply a restriction in that queries are expanded by adding those terms that are most similar to the entirety of the query, rather than selecting terms that are similar to a single term in the query.

In the next section we describe the details of our method.

## 2 Thesauri

### 2.1 WordNet

In WordNet, words are organized into taxonomies where each node is a set of synonyms (a synset) representing a single sense. There are 4 different taxonomies based on distinct parts of speech and many relationships defined within each. In this paper we use only noun taxonomy with hyponymy/hypernymy (or is-a) relations, which relates more general and more specific senses (Miller, 1988). Figure 1 shows a fragment of the WordNet taxonomy.

The similarity between word  $w_1$  and  $w_2$  is defined as the shortest path from each sense of  $w_1$  to each sense of  $w_2$ , as below (Leacock and Chodorow, 1988; Resnik, 1995)

$$sim(w_1, w_2) = \max[-\log(\frac{N_p}{2D})]$$

where  $N_p$  is the number of nodes in path  $p$  from  $w_1$  to  $w_2$  and  $D$  is the maximum depth of the taxonomy.

### 2.2 Roget's Thesaurus

In Roget's Thesaurus (Chapman, 1977), words are classified according to the ideas they express, and these categories of ideas are numbered in sequence. The terms within a category are further organized by part of speech (nouns, verbs, adjectives, adverbs, prepositions, conjunctions, and interjections). Figure 2 shows a fragment of Roget's category.

In this case, our similarity measure treat all the words in Roget as features. A word  $w$  possesses the feature  $f$  if  $f$  and  $w$  belong to the same Roget category. The similarity between two words is then defined as the Dice coefficient of the two feature vectors (Lin, 1998).

$$sim(w_1, w_2) = \frac{2|R(w_1) \cap R(w_2)|}{|R(w_1)| + |R(w_2)|}$$

where  $R(w)$  is the set of words that belong to the same Roget category as  $w$ .

### 2.3 Corpus-based Thesaurus

#### 2.3.1 Co-occurrence-based Thesaurus

This method is based on the assumption that a pair of words that frequently occur together in the same document are related to the same subject. Therefore word co-occurrence information can be used to identify semantic relationships between words (Schutze and Pederson, 1997; Schutze and Pederson, 1994). We use mutual information as a

tool for computing similarity between words. Mutual information compares the probability of the co-occurrence of words  $a$  and  $b$  with the independent probabilities of occurrence of  $a$  and  $b$  (Church and Hanks, 1990).

$$I(a, b) = \log \frac{P(a, b)}{P(a)P(b)}$$

where the probabilities of  $P(a)$  and  $P(b)$  are estimated by counting the number of occurrences of  $a$  and  $b$  in documents and normalizing over the size of vocabulary in the documents. The joint probability is estimated by counting the number of times that word  $a$  co-occurs with  $b$  and is also normalized over the size of the vocabulary.

#### 2.3.2 Syntactically-based Thesaurus

In contrast to the previous section, this method attempts to gather term relations on the basis of linguistic relations and not document co-occurrence statistics. Words appearing in similar grammatical contexts are assumed to be similar, and therefore classified into the same class (Lin, 1998; Grefenstette, 1994; Grefenstette, 1992; Ruge, 1992; Hindle, 1990).

First, all the documents are parsed using the Apple Pie Parser. The Apple Pie Parser is a natural language syntactic analyzer developed by Satoshi Sekine at New York University (Sekine and Grishman, 1995). The parser is a bottom-up probabilistic chart parser which finds the parse tree with the best score by way of the best-first search algorithm. Its grammar is a semi-context sensitive grammar with two non-terminals and was automatically extracted from Penn Tree Bank syntactically tagged corpus developed at the University of Pennsylvania. The parser generates a syntactic tree in the manner of a Penn Tree Bank bracketing. Figure 3 shows a parse tree produced by this parser.

The main technique used by the parser is the best-first search. Because the grammar is probabilistic, it is enough to find only one parse tree with highest possibility. During the parsing process, the parser keeps the unexpanded active nodes in a heap, and always expands the active node with the best probability.

Unknown words are treated in a special manner. If the tagging phase of the parser finds an unknown word, it uses a list of parts-of-speech defined in the parameter file. This information has been collected from the Wall Street Journal corpus and uses part of the corpus for training and the rest for testing. Also, it has separate lists for such information as special suffices like -ly, -y, -ed, -d, and -s. The accuracy of this parser is reported

Synonyms/Hypernyms (Ordered by Frequency) of noun correlation  
2 senses of correlation

Sense 1

correlation, correlativity  
=> reciprocity, reciprocity  
=> relation  
=> abstraction

Figure 1: An Example WordNet entry

9. Relation. -- N. relation, bearing, reference, connection, concern, . cognation ; correlation c. 12; analogy; similarity c. 17; affinity, homology, alliance, homogeneity, association; approximation c. (nearness) 197; filiation c. (consanguinity) 11[obs3]; interest; relevancy c. 23; dependency, relationship, relative position.  
comparison c. 464; ratio, proportion.  
link, tie, bond of union.

Figure 2: A fragment of a Roget's Thesaurus entry

as parseval recall 77.45 % and parseval precision 75.58 %.

Using the above parser, the following syntactic structures are extracted :

- Subject-Verb  
a noun is the subject of a verb.
- Verb-Object  
a noun is the object of a verb.
- Adjective-Noun  
an adjective modifies a noun.
- Noun-Noun  
a noun modifies a noun.

Each noun has a set of verbs, adjectives, and nouns that it co-occurs with, and for each such relationship, a mutual information value is calculated.

- $I_{sub}(v_i, n_j) = \log \frac{f_{sub}(n_j, v_i)/N_{sub}}{(f_{sub}(n_j)/N_{sub})(f(v_i)/N_{sub})}$   
where  $f_{sub}(v_i, n_j)$  is the frequency of noun  $n_j$  occurring as the subject of verb  $v_i$ ,  $f_{sub}(n_j)$  is the frequency of the noun  $n_j$  occurring as subject of any verb,  $f(v_i)$  is the frequency of the verb  $v_i$ , and  $N_{sub}$  is the number of subject clauses.
- $I_{obj}(v_i, n_j) = \log \frac{f_{obj}(n_j, v_i)/N_{obj}}{(f_{obj}(n_j)/N_{obj})(f(v_i)/N_{obj})}$   
where  $f_{obj}(v_i, n_j)$  is the frequency of noun  $n_j$  occurring as the object of verb  $v_i$ ,  $f_{obj}(n_j)$  is the frequency of the noun  $n_j$  occurring as object of any verb,  $f(v_i)$  is the frequency of

the verb  $v_i$ , and  $N_{sub}$  is the number of object clauses.

- $I_{adj}(a_i, n_j) = \log \frac{f_{adj}(n_j, a_i)/N_{adj}}{(f_{adj}(n_j)/N_{adj})(f(a_i)/N_{adj})}$   
where  $f(a_i, n_j)$  is the frequency of noun  $n_j$  occurring as the argument of adjective  $a_i$ ,  $f_{adj}(n_j)$  is the frequency of the noun  $n_j$  occurring as the argument of any adjective,  $f(a_i)$  is the frequency of the adjective  $a_i$ , and  $N_{adj}$  is the number of adjective clauses.
- $I_{noun}(n_i, n_j) = \log \frac{f_{noun}(n_j, n_i)/N_{noun}}{(f_{noun}(n_j)/N_{noun})(f(n_i)/N_{noun})}$  where  
 $f(a_i, n_j)$  is the frequency of noun  $n_j$  occurring as the argument of noun  $n_i$ ,  $f_{noun}(n_j)$  is the frequency of the noun  $n_j$  occurring as the argument of any noun,  $f(n_i)$  is the frequency of the noun  $n_i$ , and  $N_{noun}$  is the number of noun clauses.

The similarity  $sim(w, w_2)$  between two words  $w_1$  and  $w_2$  can be computed as follows :

$$\frac{\sum_{(r,w) \in T(w_1) \cap T(w_2)} (I_r(w_1, w) + I_r(w_2, w))}{\sum_{(r,w) \in T(w_1)} I_r(w_1, w) + \sum_{(r,w) \in T(w_2)} I_r(w_2, w)}$$

Where  $r$  is the syntactic relation type, and  $w$  is

- a verb, if  $r$  is the subject-verb or object-verb relation.
- an adjective, if  $r$  is the adjective-noun relation.

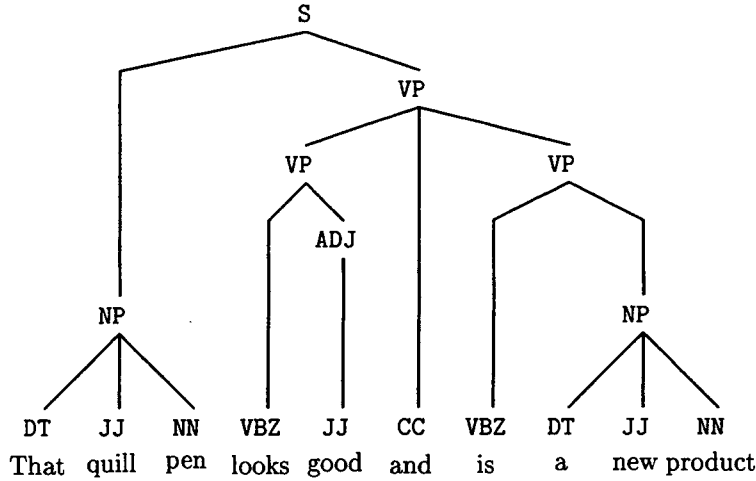


Figure 3: An example parse tree

- a noun, if  $r$  is the noun-noun relation.

and  $T(w)$  is the set of pairs  $(r, w')$  such that  $I_r(w, w')$  is positive.

### 3 Combination and Term Expansion Method

A query  $q$  is represented by the vector  $\vec{q} = (q_1, q_2, \dots, q_n)$ , where each  $q_i$  is the weight of each search term  $t_i$  contained in query  $q$ . We used SMART version 11.0 (Salton, 1971) to obtain the initial query weight using the formula  $lrc$  as follows :

$$\frac{(\log(tf_{ik}) + 1.0) * \log(N/n_k)}{\sqrt{\sum_{j=1}^n [(\log(tf_{ij} + 1.0) * \log(N/n_j))^2]}}$$

where  $tf_{ik}$  is the occurrence frequency of term  $t_k$  in query  $q_i$ ,  $N$  is the total number of documents in the collection, and  $n_k$  is the number of documents to which term  $t_k$  is assigned.

Using the above weighting method, the weight of initial query terms lies between 0 and 1. On the other hand, the similarity in each type of thesaurus does not have a fixed range. Hence, we apply the following normalization strategy to each type of thesaurus to bring the similarity value into the range  $[0, 1]$ .

$$sim_{new} = \frac{sim_{old} - sim_{min}}{sim_{max} - sim_{min}}$$

The similarity value between two terms in the combined thesauri is defined as the average of their similarity value over all types of thesaurus.

The similarity between a query  $q$  and a term  $t_j$  can be defined as follows :

$$simqt(q, t_j) = \sum_{t_i \in q} q_i * sim(t_i, t_j)$$

where the value of  $sim(t_i, t_j)$  is taken from the combined thesauri as described above.

With respect to the query  $q$ , all the terms in the collection can now be ranked according to their  $simqt$ . Expansion terms are terms  $t_j$  with high  $simqt(q, t_j)$ .

The  $weight(q, t_j)$  of an expansion term  $t_j$  is defined as a function of  $simqt(q, t_j)$ :

$$weight(q, t_j) = \frac{simqt(q, t_j)}{\sum_{t_i \in q} q_i}$$

where  $0 \leq weight(q, t_j) \leq 1$ .

The weight of an expansion term depends both on all terms appearing in a query and on the similarity between the terms, and ranges from 0 to 1. The weight of an expansion term depends both on the entire query and on the similarity between the terms. The weight of an expansion term can be interpreted mathematically as the weighted mean of the similarities between the term  $t_j$  and all the query terms. The weight of the original query terms are the weighting factors of those similarities (Qiu and Frei, 1993).

Therefore the query  $q$  is expanded by adding the following query :

$$\vec{q}_e = (a_1, a_2, \dots, a_r)$$

where  $a_j$  is equal to  $weight(q, t_j)$  if  $t_j$  belongs to the top  $r$  ranked terms. Otherwise  $a_j$  is equal to 0.

The resulting expanded query is :

$$q_{expanded} = q \circ q_e$$

where the  $\circ$  is defined as the concatenation operator.

The method above can accommodate polysemy, because an expansion term which is taken from a different sense to the original query term is given a very low weight.

## 4 Experiments

Experiments were carried out on the TREC-7 Collection, which consists of 528,155 documents and 50 topics (Voorhees and Harman, to appear 1999). TREC is currently de facto standard test collection in information retrieval community.

Table 1 shows topic-length statistics, Table 2 shows document statistics, and Figure 4 shows an example topic.

We use the title, description, and combined title+description+narrative of these topics. Note that in the TREC-7 collection the description contains all terms in the title section.

For our baseline, we used SMART version 11.0 (Salton, 1971) as information retrieval engine with the *lnc.ltc* weighting method. SMART is an information retrieval engine based on the vector space model in which term weights are calculated based on term frequency, inverse document frequency and document length normalization.

Automatic indexing of a text in SMART system involves the following steps :

- **Tokenization** : The text is first tokenized into individual words and other tokens.
- **Stop word removal** : Common function words (like *the*, *of*, *an*, etc.) also called stop words, are removed from this list of tokens. The SMART system uses a predefined list of 571 stop words.
- **Stemming**: Various morphological variants of a word are normalized to the same stem. SMART system uses the variant of Lovin method to apply simple rules for suffix stripping.
- **Weighting** : The term (word and phrase) vector thus created for a text, is weighted using *tf*, *idf*, and length normalization considerations.

Table 3 gives the average of non-interpolated precision using SMART without expansion (baseline), expansion using only WordNet, expansion

using only the corpus-based syntactic-relation-based thesaurus, expansion using only the corpus-based co-occurrence-based thesaurus, and expansion using combined thesauri. For each method we also give the relative improvement over the baseline. We can see that the combined method outperform the isolated use of each type of thesaurus significantly.

Table 1: TREC-7 Topic length statistics

Topic Section	Min	Max	Mean
Title	1	3	2.5
Description	5	34	14.3
Narrative	14	92	40.8
All	31	114	57.6

## 5 Discussion

In this section we discuss why our method using WordNet is able to improve information retrieval performance. The three types of thesaurus we used have different characteristics. Automatically constructed thesauri add not only new terms but also new relationships not found in WordNet. If two terms often co-occur in a document then those two terms are likely to bear some relationship.

The reason why we should use not only automatically constructed thesauri is that some relationships may be missing in them. For example, consider the words *colour* and *color*. These words certainly share the same context, but would never appear in the same document, at least not with a frequency recognized by a co-occurrence-based method. In general, different words used to describe similar concepts may never be used in the same document, and are thus missed by co-occurrence methods. However their relationship may be found in WordNet, Roget's, and the syntactically-based thesaurus.

One may ask why we included Roget's Thesaurus here which is almost identical in nature to WordNet. The reason is to provide more evidence in the final weighting method. Including Roget's as part of the combined thesaurus is better than not including it, although the improvement is not significant (4% for title, 2% for description and 0.9% for all terms in the query). One reason is that the coverage of Roget's is very limited.

A second point is our weighting method. The advantages of our weighting method can be summarized as follows:

- the weight of each expansion term considers the similarity of that term to all terms in the

Table 2: TREC-7 Document statistics

Source	Size (Mb)	# Docs	Median # Words/Doc	Mean # Words/Doc
<b>Disk 4</b>				
FT	564	210,158	316	412.7
FR94	395	55,630	588	644.7
<b>Disk 5</b>				
FBIS	470	130,471	322	543.6
LA Times	475	131,896	351	526.5

Title :  
ocean remote sensing

Description:  
Identify documents discussing the development and application of spaceborne ocean remote sensing.

Narrative:  
Documents discussing the development and application of spaceborne ocean remote sensing in oceanography, seabed prospecting and mining, or any marine-science activity are relevant. Documents that discuss the application of satellite remote sensing in geography, agriculture, forestry, mining and mineral prospecting or any land-bound science are not relevant, nor are references to international marketing or promotional advertizing of any remote-sensing technology. Synthetic aperture radar (SAR) employed in ocean remote sensing is relevant.

Figure 4: Topics Example

original query, rather than to just one query term.

- the weight of an expansion term also depends on its similarity within all types of thesaurus.

Our method can accommodate polysemy, because an expansion term taken from a different sense to the original query term sense is given very low weight. The reason for this is that the weighting method depends on all query terms and all of the thesauri. For example, the word *bank* has many senses in WordNet. Two such senses are the financial institution and river edge senses. In a document collection relating to financial banks, the river sense of *bank* will generally not be found in the cooccurrence-based thesaurus because of a lack of articles talking about rivers. Even though (with small possibility) there may be some documents in the collection talking about rivers, if the query contained the finance sense of *bank* then the other terms in the query would also tend to be concerned with finance and not rivers. Thus rivers would only have a relationship with the *bank* term and there would be no relations with other terms in the original query, resulting in a low weight.

Since our weighting method depends on both the query in its entirety and similarity over the three thesauri, wrong sense expansion terms are given very low weight.

## 6 Related Research

Smeaton (1995) and Voorhees (1994; 1988) proposed an expansion method using WordNet. Our method differs from theirs in that we enrich the coverage of WordNet using two methods of automatic thesaurus construction, and we weight the expansion term appropriately so that it can accommodate polysemy.

Although Stairmand (1997) and Richardson (1995) proposed the use of WordNet in information retrieval, they did not use WordNet in the query expansion framework.

Our syntactic-relation-based thesaurus is based on the method proposed by Hindle (1990), although Hindle did not apply it to information retrieval. Hindle only extracted subject-verb and object-verb relations, while we also extract adjective-noun and noun-noun relations, in the manner of Grefenstette (1994), who applied his

Table 3: Average non-interpolated precision for expansion using single or combined thesauri.

Topic Type	Base	Expanded with				
		WordNet only	Roget only	Syntac only	Cooccur only	Combined method
Title	0.1175	0.1276 (+8.6%)	0.1236 (+5.2 %)	0.1386 (+17.9%)	0.1457 (+24.0%)	0.2314 (+96.9%)
Description	0.1428	0.1509 (+5.7%)	0.1477 (+3.4 %)	0.1648 (+15.4%)	0.1693 (+18.5%)	0.2645 (+85.2%)
All	0.1976	0.2010 (+1.7%)	0.1999 (+1.2%)	0.2131 (+7.8%)	0.2191 (+10.8%)	0.2724 (+37.8%)

syntactically-based thesaurus to information retrieval with mixed results. Our system improves on Grefenstette's results since we factor in thesauri which contain hierarchical information absent from his automatically derived thesaurus.

Our weighting method follows the Qiu and Frei (1993) method, except that Qiu used it to expand terms from a single automatically constructed thesaurus and did not consider the use of more than one thesaurus.

This paper is an extension of our previous work (Mandala et al., to appear 1999) in which we did not consider the effects of using Roget's Thesaurus as one piece of evidence for expansion and used the Tanimoto coefficient as similarity coefficient instead of mutual information.

## 7 Conclusions

We have proposed the use of different types of thesaurus for query expansion. The basic idea underlying this method is that each type of thesaurus has different characteristics and combining them provides a valuable resource to expand the query. Wrong expansion terms can be avoided by designing a weighting term method in which the weight of expansion terms not only depends on all query terms, but also depends on their similarity values in all type of thesaurus.

Future research will include the use of a parser with better performance and the use of more recent term weighting methods for indexing.

## 8 Acknowledgements

The authors would like to thank Mr. Timothy Baldwin (TIT, Japan) and three anonymous referees for useful comments on the earlier version of this paper. We also thank Dr. Chris Buckley (SabIR Research) for support with SMART, and Dr. Satoshi Sekine (New York University) for providing the Apple Pie Parser program. This

research is partially supported by JSPS project number JSPS-RFTF96P00502.

## References

- J. Aitchison and A. Gilchrist. 1987. *Thesaurus Construction: A Practical Manual*. Aslib.
- D.C. Blair and M.E. Maron. 1985. An evaluation of retrieval effectiveness. *Communications of the ACM*, 28:289-299.
- Robert L. Chapman. 1977. *Roget's International Thesaurus (Forth Edition)*. Harper and Row, New York.
- Kenneth Ward Church and Patrick Hanks. 1990. Word association norms, mutual information and lexicography. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pages 76-83.
- Gregory Grefenstette. 1992. Use of syntactic context to produce term association lists for text retrieval. In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 89-97.
- Gregory Grefenstette. 1994. *Explorations in Automatic Thesaurus Discovery*. Kluwer Academic Publisher.
- Donald Hindle. 1990. Noun classification from predicate-argument structures. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistic*, pages 268-275.
- Claudia Leacock and Martin Chodorow. 1988. Combining local context and WordNet similarity for word sense identification. In Christiane Fellbaum, editor, *WordNet, An Electronic Lexical Database*, pages 265-283. MIT Press.
- Dekang Lin. 1998. Automatic retrieval and clustering of similar words. In *Proceedings of the COLING-ACL'98*, pages 768-773.

- Rila Mandala, Takenobu Tokunaga, and Hozumi Tanaka. to appear, 1999. Combining general hand-made and automatically constructed thesauri for information retrieval. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*.
- George A. Miller. 1988. Nouns in WordNet. In Christiane Fellbaum, editor, *WordNet, An Electronic Lexical Database*, pages 23–46. MIT Press.
- George A. Miller. 1990. Special issue, WordNet: An on-line lexical database. *International Journal of Lexicography*, 3(4).
- Yonggang Qiu and Hans-Peter Frei. 1993. Concept based query expansion. In *Proceedings of the 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 160–169.
- Philip Resnik. 1995. Using information content to evaluate semantic similarity in a taxonomy. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 448–453.
- R. Richardson and Alan F. Smeaton. 1995. Using WordNet in a knowledge-based approach to information retrieval. Technical Report CA-0395, School of Computer Applications, Dublin City University.
- Gerda Ruge. 1992. Experiments on linguistically-based term associations. *Information Processing and Management*, 28(3):317–332.
- Gerard Salton and M McGill. 1983. *An Introduction to Modern Information Retrieval*. McGraw-Hill.
- Gerard Salton. 1971. *The SMART Retrieval System: Experiments in Automatic Document Processing*. Prentice-Hall.
- Hinrich Schutze and Jan O. Pederson. 1994. A cooccurrence-based thesaurus and two applications to information retrieval. In *Proceedings of the RIAO 94 Conference*.
- Hinrich Schutze and Jan O. Pederson. 1997. A cooccurrence-based thesaurus and two applications to information retrieval. *Information Processing and Management*, 33(3):307–318.
- Satoshi Sekine and Ralph Grishman. 1995. A corpus-based probabilistic grammar with only two non-terminals. In *Proceedings of the International Workshop on Parsing Technologies*.
- Alan F. Smeaton and C. Berrut. 1995. Running TREC-4 experiments: A chronological report of query expansion experiments carried out as part of TREC-4. In *Proceedings of The Fourth Text REtrieval Conference (TREC-4)*. NIST special publication.
- Mark A. Stairmand. 1997. Textual context analysis for information retrieval. In *Proceedings of the 20th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pages 140–147.
- Ellen M. Voorhees and Donna Harman. to appear, 1999. Overview of the Seventh Text REtrieval Conference (TREC-7). In *Proceedings of the Seventh Text REtrieval Conference*. NIST Special Publication.
- Ellen M. Voorhees. 1988. Using WordNet for text retrieval. In Christiane Fellbaum, editor, *WordNet, An Electronic Lexical Database*, pages 285–303. MIT Press.
- Ellen M. Voorhees. 1993. Using wordnet to disambiguate word senses for text retrieval. In *Proceedings of the 16th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pages 171–180.
- Ellen M. Voorhees. 1994. Query expansion using lexical-semantic relations. In *Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pages 61–69.