

# An Analysis of Ontology-based Query Expansion Strategies

Roberto Navigli and Paola Velardi  
Dipartimento di Informatica  
Università di Roma “La Sapienza”  
{navigli,velardi}@di.uniroma1.it

## Abstract

Sense based query expansion never proved its effectiveness except for the so-called “open domain question answering” task. The present work is still inconclusive at this regard, due to some experimental limitations, but we provide interesting evidence suggesting new guidelines for future research. Word sense disambiguation is in fact only one of the problems involved with sense based query expansion. The second is how to use sense information (and ontologies in general) to expand the query. We show that expanding with synonyms or hyperonyms has a limited effect on web information retrieval performance, while other types of semantic information derivable from an ontology are much more effective at improving search results.

## 1 Introduction

Despite the growing effort of the Semantic Web community to demonstrate that ontologies, and, in general, semantic knowledge may indeed improve the accessibility of web documents by humans and machines, no strong experimental results are yet available to support this convincement. The most important Information Retrieval (IR) conferences (SIGIR, TREC)<sup>1</sup> show the predominance of standard keyword-based techniques, improved through the use of additional mechanisms such as document structure analysis (Cutker et al. 1999) and query expansion using statistical methods (Carpineto et al. 2002) and query logs (Cui et al 2002). High-performant search engines rely also on the exploitation of the hypertextual relations in a document, using anchor analysis (Eiron and McCuley 2003) and link analysis (Borodin et al 2001). The effectiveness of the various techniques depends on the task, e.g. subject finding vs. site retrieval, as well as on the dimension of the query (short vs. long). Query expansion seems

particularly useful for short queries of two-three words, that represent the standard case for users of web search engines.

Except for the so-called “open domain question answering” task (Moldovan et al. 2002), the use of knowledge bases in state-of-art web retrieval systems is almost absent. Published results on sense based query expansions are not very recent (Voorhees,1993) (Sanderson, 1994). A more recent work (Gonzalo et al., 1998) analyzes the effect of expanding a query with WordNet synsets, in a “canned” experiment where all words are manually disambiguated. Gonzalo and his colleagues show that a substantial increase in performance is obtained only with less than 10% errors in the word sense disambiguation (WSD) task. Since WSD is known as one of the hardest problems in Artificial Intelligence, this study left us with little hope that sense based query expansion might indeed rival with statistical methods.

We believe that the complexity of WSD is only one of the problems, the second being *how to use sense information to effectively expand the query*. In the literature, sense based query expansion is performed replacing senses with taxonomic information, e.g. synonyms or hyperonyms. However, the most successful query expansion methods seem to suggest that the best way to expand a query is by adding words that often co-occur with the words of the query, i.e. words that, on a probabilistic ground, are believed to pertain to the same *semantic domain* (e.g. *car* and *driver*). Query expansion terms are extracted either from an initial set of top retrieved documents (Carpineto et al. 2002) or from query logs, i.e. associations between a query and the documents downloaded by the user (Cui et al 2002). This latter source of co-occurrence information is obviously more precise, but proprietary.

In this study, we experiment the possibility of using ontological information to extract the *semantic domain* of a word. Rather than using taxonomic relations for sense based query expansion (e.g. synonyms and hyperonyms) we expanded with the words in a sense *definition*. In our experiment, we use the Google retrieval engine, the

<sup>1</sup> <http://www.informatik.uni-trier.de/~ley/db/conf/sigir/>  
<http://trec.nist.gov/>

search topics of the TREC 2001 web track<sup>2</sup> to query the web, and WordNet 1.6<sup>3</sup> to extract word senses and sense-related information. Our results are preliminary, both because of the limited size of the experiment, and because of some limitations imposed by the use of Google. Still, we show a systematic improvement over the unexpanded query, especially when query expansion terms are chosen from the sense definitions of the query words. Interestingly, the improvement is considerable even when word sense disambiguation performance is lower than the 90% suggested by Gonzalo and his colleagues.

The paper is organized as follows: section 2 describes the disambiguation algorithm, which relies on the WordNet lexical knowledge base. Section 3 describes the different query expansion methods adopted in the experiment. Section 4 presents a discussion of the results. Conclusions are drawn in section 5.

## 2 Word Sense Disambiguation

Word sense disambiguation is known as one of the most complex tasks in the area of artificial intelligence. We do not even attempt here a survey of the field, but we refer the interested reader to the Senseval home page (<http://www.senseval.org/>) for a collection of state of art sense disambiguation methods, and the results of public competitions in this area. During the most recent Senseval evaluation, the best system in the English all-words task (Mihalcea and Moldovan, 2001) reached a 69% precision and recall, a performance that (Gonzalo et al., 1998) claim to be well below the threshold that produces improvements in a text retrieval task.

However, for a query expansion task it is not necessary to pursue high recall, but rather high precision. As we show in sections 3 and 4, even expanding only monosemous words in a query may produce a significant improvement over the unexpanded query.

Therefore we developed an algorithm that may be tuned to produce high precision, possibly at the price of low recall. The algorithm belongs to the class of structural pattern recognition methods (Pavlidis, 1977). Structural pattern recognition is particularly useful when instances have an inherent, identifiable organization, which is not captured by feature vectors. In our work we use a graph representation to describe instances (word senses). Shortly, the algorithm is as follows:

Let  $Q = \{w_1, w_2, \dots, w_n\}$  be the initial query (stop words are pruned as usual)

Let

$$S(w_k) = \{S_j^k \mid S_j^k \in \text{Synset}_{\text{WordNet}}(w_k), w_k \in Q\}$$

<sup>2</sup> <http://trec.nist.gov/pubs/trec10/papers/web2001.ps.gz>

<sup>3</sup> <http://www.cogsci.princeton.edu/~wn/>

be the WordNet synonym sets (synsets) of  $w_k$ ,  $k=1, \dots, n$ . Let further

$$C_x = (S_{x_1}^1, S_{x_2}^2, \dots, S_{x_n}^n)$$

be a possible configuration of senses for  $Q$  ( $x_k$  is a sense index between 1 and the number of possible senses for  $w_k$ ).

For each configuration  $C_x$ , do the following:

1. Create semantic networks for each sense;
2. Intersect semantic networks;
3. Assign a score the configuration.

Finally select  $C_{best} = \arg \max_x (Score(C_x))$ .

In the next sections the three steps will be described in detail.

### 2.1 Creation of semantic networks

For every  $w_k \in Q$  and every synset  $S_j^k$  of  $w_k$  (where  $S_j^k$  is the  $j$ -th sense of  $w_k$  in WordNet) we create a *semantic net*.

Semantic nets are automatically built using the following semantic relations: *hyperonymy* (car *is-a* vehicle, denoted with  $\rightarrow^{\text{@}}$ ), *hyponymy* (its inverse,  $\rightarrow^{\sim}$ ), *meronymy* (room *has-a* wall,  $\rightarrow^{\#}$ ), *holonymy* (its inverse,  $\rightarrow^{\%}$ ), *pertainymy* (dental *pertains-to* tooth  $\rightarrow^{\vee}$ ), *attribute* (dry *value-of* wetness,  $\rightarrow^{\text{=}}$ ), *similarity* (beautiful *similar-to* pretty,  $\rightarrow^{\&}$ ), *gloss* ( $\rightarrow^{\text{gloss}}$ ), *topic* ( $\rightarrow^{\text{topic}}$ ), *domain* ( $\rightarrow^{\text{dl}}$ ).

Every relation is directly extracted from WordNet, except for *gloss*, *topic* and *domain*.

The *topic* and the *gloss* relations are obtained parsing with a NL processor respectively the SemCor<sup>4</sup> sentences including a given synset  $S_j^k$  and WordNet concept definitions (called *glosses*). *SemCor* is an annotated corpus where each word in a sentence is assigned a sense selected from the WordNet sense inventory for that word; an example is the following:

*Movement#7 itself was#7 the chief#1 and often#1 the only#1 attraction#4 of the primitive#1 movies#1 of the nineties#1.*

The *topic* relations extracted from Semcor identify semantic co-occurrences between two related nodes of the semantic network (e.g. *chief#1*  $\rightarrow^{\text{topic}}$  *attraction#4*).

As far as the gloss relation is concerned, it is worth noticing that words in glosses do not have sense tags in WordNet, therefore we use an algorithm for gloss disambiguation that is a variation of the WSD algorithm described in this section. For example, for sense #1 of

<sup>4</sup> <http://engr.smu.edu/~rada/semcor/>

*bus* “a vehicle carrying many passengers; ...” the following relations are created:

$$\text{bus}\#1 \rightarrow^{\text{gloss}} \text{vehicle}\#1, \text{bus}\#1 \rightarrow^{\text{gloss}} \text{passenger}\#1$$

Finally, the *domain* relation is extracted from the set of domain labels (e.g. *tourism*, *chemistry*, *economy*..) assigned to WordNet synsets by a semiautomatic methodology described in (Magnini and Cavaglia, 2000). To reduce the dimension of a SN, we consider only concepts at a distance not greater than 3 relations from  $S_j^k$  (the SN center). The dimension of the SN has been experimentally tuned.

Figure 1 is an example of SN generated for sense #1 of *bus*.

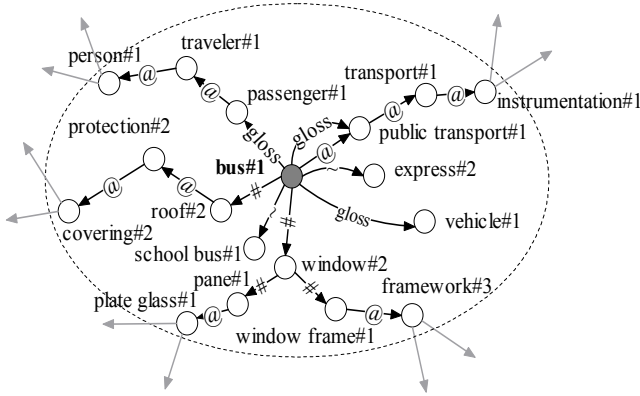


Figure 1. The semantic net for sense 1 of *bus*.

## 2.2 Intersecting semantic networks and scoring configurations

Let then  $SN(S_j^k)$  be the semantic network for sense  $j$  of word  $w_k$ . Given a configuration of senses  $C_x$ , for a query  $Q$ , semantic networks are intersected pair-wise, and the number of common nodes are counted. Let  $SN(S_j^k) \cap SN(S_m^l)$  be one such intersection. Common nodes  $S$  are those that can be reached from both SN centers through directed paths, e.g.:  $S_j^k \rightarrow^* S^* \leftarrow S_m^l$  where  $\rightarrow^*$  denotes a sequence of nodes and arcs of any type.

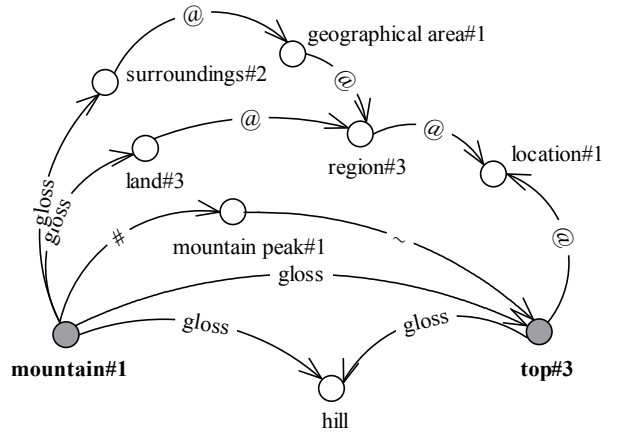


Figure 2. The patterns between *mountain#1* and *top#3*.

Figure 2 shows an example of intersection between the SN of sense 1 of *mountain* and the SN of sense 3 of *top*. There are 2 common nodes (*location#1* and *hill#1*), plus the direct gloss relation between the two central senses (therefore also the SN center *top#3* is common, according to our definition).

For each sense configuration, the score is computed as the total number of common nodes (e.g. 3 in the previous example):

$$\text{Score}(C_x) = \sum_{S', S'' \in C_x, S' \neq S''} |SN(S') \cap SN(S'')|$$

Furthermore, common nodes are ordered according to the inverse of the length of intersecting paths in which they participate. Let then  $[S]^x$  be the ordered list of shared nodes in a configuration  $C_x$ .

## 3. The experiment

The objective of the experiment described in this paper is only in part the evaluation of the WSD algorithm described in previous section, that is still under refinement. Rather, our purpose is to obtain a better insight on the *use* of sense information for improving web search.

To this end, we used five sense-based expansion methods, and two strategies to choose expansible words. The following expansion methods are explored:

1. Synset expansion: “expansible” words are replaced by their synsets, retrieved by the algorithm of previous section.
2. Hyperonym expansion: “expansible” words are augmented by their WN direct hyperonyms.
3. Gloss synset expansion: “expansible” words are augmented by the synsets of its glosses

(disambiguated by an ad-hoc version of our WSD algorithm).

4. Gloss words expansion: “expansible” words are augmented with the words in their glosses.
5. Common nodes expansion: “expansible” words are augmented with the words whose synsets are in  $\{S_j\}^x$ .

According to the first strategy, expansible words are only monosemous words. In the second, we expand words whose synset, selected according to the WSD algorithm of section 2, has at least  $k$  ( $k > 0$ ) nodes in common with other synsets of the query. The first strategy ensures maximum sense disambiguation precision, while the second allows it to tune the best precision-recall trade off, through the parameter  $k$ .

We queried the web with the first 24 of the 50 queries used in the TREC2001 web track. The queries (called “topics”) include the actual query (*title*) but also text to explain the query (*description*) and describe precisely the type of documents that should be considered relevant (*narrative*).

For example:

```
<top>
<num> Number: 518
<title> how we use statistics to aid our decision making?
<desc> Description:
Find documents that reference the use of statistical
data in decision-making.
<narr> Narrative:
A relevant document will describe a specific statistical
method that is used to assist decision-making.
</top>
```

Clearly, *description* and *narrative* texts cannot be used to expand the query, but only to manually verify the correctness of retrieved documents, as we did. To query the web, we used Google, which revealed not to be the best choice to exploit our algorithm, due to the limitation of 10 words per query. Therefore, for longer queries we are faced with the problem of choosing only a fragment of the candidate expansion words.

However, we felt that our results could be stronger if we show an improvement in performance using the most popular search engine.

For each query, we retrieved the first 10 top ranked pages without query expansion, and then we repeated the search with each of the sense based expansion methods outlined above. When expansion terms are synsets, terms of a synset are put in OR. Whenever plain query terms + expansion terms exceed the threshold imposed by Google, we simply choose the first words of the list, a strategy that is optimized only for method 5, since the list  $[S]^x$  is ordered according to the strength of the intersection of each synset  $S$ .

The results are shown in Table 1a and b.

#### 4. Discussion

Table 1a shows the results of the first four different expansion methods, when only monosemous words are expanded. The method five is not tested, since in each query rarely more than one word is monosemous. Consequently, intersecting paths between synsets are found only for one or two queries, which makes the evaluation not relevant.

Expanding only monosemous words is a very conservative policy, still, Table 1a shows interesting results. Every strategy produces an improvement. In particular, expanding with gloss words produces a 26,83% improvement over the plain query words. Interesting enough, the increase in performance (or at least a non-decrease) is systematic. The only critical cases are those where the query includes a named entity (e.g. topic 527: “can you info on booker t. washington?”). Since, regrettably, we did not use in this experiment any additional tool for the treatment of named entities, and since in the “TREC topics“ names are often not capitalized, “booker” is interpreted as the monosemous concept *booker, booking agent*, with obvious consequences on retrieval performance.

Table 1b is the same experiment as for in Table 1a, but now the policy is to expand all words whose synset has at least one node ( $k=1$ ) in common with some other synset of the chosen configuration. The sense configuration for ambiguous words is chosen according to our WSD algorithm of section 2.

Table 1b shows the results of six query types (unexpanded plus our five methods). The WSD algorithm attempted to disambiguate 52 words included in the 24 queries. The precision was 82,97% (39/47) and the recall 90,38% (47/52). This is a good result, though not comparable with WSD literature, given the limited size of the experiment. In agreement with the high precision requirement in (Gonzalo et al. 2000), the WSD precision is not enough to systematically improve over the monosemous words experiment, however the results are very close, in some case slightly better (synonyms) in some case slightly worst (gloss words). The gloss words expansion strategy achieves an improvement that is still very high, confirming that a better expansion strategy may overcome the problem of imprecise WSD.

This is a very interesting result, and even though the size of the experiment should be increased, this behavior seems very consistent across the various queries. Only in few cases some gloss word causes a decrease in performance. One interesting case is the topic “uniforms in public schools”. The gloss for *public school#1* is “a free school supported by taxes and controlled by a school board”. Here the word *tax* causes unwanted hits during web search. We would expect this problem be reduced by the “common nodes expansion” technique (the last column in Table 1b), but unfortunately this technique works very badly. In 11 cases no common nodes were

found<sup>5</sup>, and in the other cases the results are mixed, leading eventually to a decrease in performance with respect to the plain query.

Looking at the data, it appears that often there are interesting common nodes, but our node weighting method, though intuitive, does not capture them.

Furthermore, many common nodes are related to the query word synsets by a hyperonymy relation, a relation that was already confirmed as a bad expansion strategy. Same examples are useful here.

First, we provide an example of retrieved sense-based expansion words for the five methods, relative to the TREC topic 501: “deduction and induction in English?”

The first two lines show respectively, the senses chosen for each word by the WSD algorithm, and by the manual annotators (in some case more than one sense seemed appropriate). Then, the WordNet concepts extracted for each of the expansion policies are shown.

**WSD:** {deduction#3, induction#3, English#1}  
**WSD manual:** {deduction#4|3, induction#3, English#1}  
**synsets:** {English}, {generalization, induction, inductive\_reasoning}, {deduction, entailment, implication}  
**hyperonyms:** {reasoning, logical\_thinking, abstract\_thought}, {inference, illation}  
**gloss synsets:** {England}, {detailed, elaborate, elaborated}  
**gloss words:** relating, culture, England, characteristic, detailed, facts, general, principles, reasoning, implied, deduced, is\_inferred, entailed  
**common nodes (with weights):** {reasoning, logical\_thinking, abstract\_thought}:0,33, {syllogism}:0.16, {argumentation, local\_argument, line\_of\_reasoning, line}:0.2, {thinking, thought, cerebration, intellection, mentation}:0.2, {deduction, deductive\_reasoning, synthesis}:0.2, {analysis, analytic\_thinking}:0.2, {conjecture}:0.2, {reasoning, logical\_thinking, abstract\_thought}:0,33

It is interesting also to provide examples of common patterns between semantic networks of word senses. In the previous query, several common patterns are found between *deduction#3* and *induction#3* (figure 3).




Figure 3. A path between *deduction#3* and *induction#3*.

where  $[S_i]$  is the common node according to our definition.

*Induction* and *deduction* are semantically close words, and in fact the common node technique work well here.

Often common paths are found also between less intuitively related words, e.g. in TREC topics “information

<sup>5</sup> in these cases (marked with \*) the plain query word strategy is used, to allow a comparison with the other columns of Table 1b.

about what manatees eat” and “hair loss is a symptom of what diseases”. In the first query, the common nodes technique find patterns like that in figure 4.

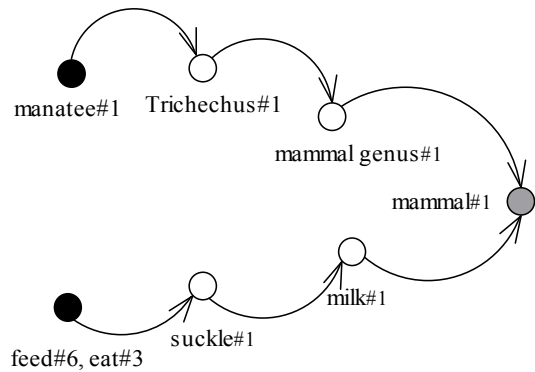


Figure 4. A path between *manatee#1* and *eat#3*.

The final set of chosen common nodes for this query is: *animal*: 0.25, *mammal*:0.2, *placental*:0.16, *animal\_order*:0.16. These nodes are overly general, and cause noise during expansion. As shown in Table 1b, the number of relevant pages drops from 9 (plain query words) to 6 when using the common nodes technique for this query.

For the second query, the final set of common nodes is: *medical*:0,33 and *patient*:0,33. These nodes have the same level of generality as the query words. Contrary to “manatee eat”, common nodes improve “hair loss” search from 7 to 9 relevant pages.

In many cases, named entities are the cause of problems, since their related synsets are often not useful to expand the query (figure 5 shows such a common path for the topic history on cambodia?).




Figure 5. A path connecting *history#1* and *Cambodia#1*.

The problem here is that named entities are *instances*, not concepts, and should not be expanded at all. In WordNet, instead, there is an ontological confusion between instances and concepts, which are treated in the same way. As we already remarked, we plan to have an ad-hoc treatment of named entities in our future experiments.

To summarize, each of the retrieved sense-based information types is in principle useful for query expansion, and even expanding only monosemous words in a query may provide a significant improvement in retrieval. Ideally, the sense-based expansion algorithm should be able to exploit and combine at best each of the available strategies, but this is matter of future research. Analyzing in detail the data, we found that words in the same *semantic domain* (and *same level of generality*) of the query words appear as the best candidates for expansion. Expanding with gloss words gives the best results by large, while hyperonyms expansion is a less performant strategy. The common nodes technique produces improvement (almost) only when the selected common nodes are related through non taxonomic relations (e.g. the *medical care* example), while it works quite badly if the selected common nodes are hyperonyms of the query word senses.

These findings seem very consistent throughout our data, however, in order to declare our results conclusive, it is indeed necessary to experiment on much larger corpora, for example, over the TREC web corpus, an experiment that we plan to conduct in the near future.

## 5. Conclusions

In this paper we presented a word sense disambiguation method based on structured pattern recognition, and we used this method to explore several sense-based strategies for web search query expansion. By means of a small but interesting set of experiments, we could draw interesting conclusions on the type of sense-related information that appears more useful for web search, and obtain evidence on possible developments of the work.

## Acknowledgments

This work has been in part funded by the MURST-CNR Web Learning national project.

## References

(Borodin et al 2001) Borodin a., Roberts G., Rosenthal J., Tsaparas P. : Finding Authorities and Hubs from Link Structures on the WWW” WWW10, may 1-5. 2001 Hong-Kong.

(Carpineto et al. 2002) Claudio Carpineto, Giovanni Romano, Vittorio Giannini: Improving retrieval feedback with multiple term-ranking function combination. TOIS 20(3): 259-290 (2002).

(Craswell and Hawking, 2002) Overview of the TREC-2002 Web Track,  
[http://trec.nist.gov/pubs/trec11/t11\\_proceedings.html](http://trec.nist.gov/pubs/trec11/t11_proceedings.html)

(Cui et al. 2002) Cui H. Wen J. Nie J. Ma W. Probabilistic Query expansion using query logs” WWW202, may 7-11, Hawaii, USA, ACM 1-58113-449-5/02/0005.

(Cutler et al. 1999) Cutler M. Deng H. Maniccom S. Meng W. “A new study on using HTML structure to improve retrieval” 11<sup>th</sup> IEEE Conf. on Tools with AI, 1999.

(Eiron and McCuley 2003) Eiron N. and McCuley K. Analysis of Anchor Text for Web Search, SIGIR 2003, Toronto, Canada.

(Gonzalo et al. 1998) Julio Gonzalo Felisa Verdejo Irina Chugur Juan Cigarr'an ”Indexing with WordNet synsets can improve text retrieval” Proceedings of the COLING/ACL '98 Workshop on Usage of WordNet for NLP.

(Magnini and Cavaglia, 2000) Magnini, B. and Cavaglia, G.: Integrating Subject Field Codes into WordNet. Proceedings of the 2nd International Conference on Language resources and Evaluation, LREC2000, Atenas.

(Mahesh et al, 1999) Mahesh K.,Kud J. Dixon P. Oracle at TREC 8: A lexical Approach , Proc. of TREC 8, NIST, 1999.

(Mihalcea and Moldovan, 2001) Rada Mihalcea, Dan I. Moldovan: A Highly Accurate Bootstrapping Algorithm for Word Sense Disambiguation. International Journal on Artificial Intelligence Tools 10(1-2): 5-21 (2001).

(Moldovan et al. 2002) Moldovan, D. Harabagiu S., Girju R., Morarescu P., Llacatusu F., Novischi A., Badulescu A., Bolohan O.: LCC Tools for Question Answering, [http://trec.nist.gov/pubs/trec11/t11\\_proceedings.html](http://trec.nist.gov/pubs/trec11/t11_proceedings.html)

(Pavlidis 1977) T. Pavlidis Structured Pattern Recognition, Springer-Verlag Berlin 1977,

(Sanderson, 1994) Sanderson M. Word Sense Disambiguation and Information Retrieval 17<sup>th</sup> Int. Conf. on Research and Development in Information Retrieval, 1994.

(Voorhees 1993) Voorhees E. Using WordNet to disambiguate Word Senses for Text retrieval, ACM-SIGIR Pittsburgh, PA, 1993.

**Table 1 a) Retrieved correct pages when using sense information only from monosemous words**

<b>Monosemous words only</b>	<b>Plain query words</b>	<b>+ synsets</b>	<b>+hyperonyms</b>	<b>+gloss hyperonyms</b>	<b>+ gloss words</b>
Deduction induction	5	5	5	4	6
Prime factor	4	4	5	3	9
Vikings Scotland	4	4	7	9	7
Manatee eat	9	10	10	9	9
History skateboarding	7	7	7	9	9
Hair loss	7	7	9	4	8
Oppenheimer	3	3	6	9	9
Diseases smoking	9	9	10	9	8
Tornadoes	7	7	5	8	8
Earthquakes	9	9	7	4	10
Bell	5	5	5	7	9
Halloween	5	4	0	4	8
Titanic	1	1	1	3	3
Decision making	2	3	2	0	6
Black plague	4	1	4	5	5
Mojave	4	4	4	5	7
Booker Washington	3	3	1	0	0
Hygrometer	2	5	3	5	3
Cambodia	8	8	8	5	6
Hypnosis	7	7	10	8	8
School uniforms	7	7	10	8	8
Artists 1700	1	1	0	1	3
Canadian b. Codes	9	9	8	6	7
FHA	1	3	2	0	2
<b>Avg. correct pages (over first 10)</b>	5,125	5,25	5,208333	5,291667	6,5
<b>% variation with respect plain query words</b>		<b>2,44%</b>	<b>1,63%</b>	<b>3,25%</b>	<b>26,83%</b>

**Table 1 b) Retrieved correct pages when using sense information for all disambiguated words**

<b>WSD with k=1</b>	<b>Plain Query words</b>	<b>+Synonyms</b>	<b>+Hyperonyms</b>	<b>+Gloss synsets</b>	<b>+Gloss words</b>	<b>+ Common nodes</b>
Deduction induction	5	6	7	2	6	6
Prime factor	4	4	5	3	9	*4
Vikings Scotland	4	4	7	9	10	5
Manatee eat	9	10	10	10	10	6
History skateboarding	7	7	7	9	9	*7
Hair loss	7	7	8	6	5	9
Oppenheimer	3	3	6	9	9	*3
Diseases smoking	9	10	9	8	7	7
Tornadoes	7	8	5	5	7	2
Earthquakes	9	9	7	4	10	*9
Bell	5	5	5	7	9	*5
Halloween	5	4	0	4	8	*5
Titanic	1	1	1	3	3	*1
Decision making	2	1	3	4	5	2
Black plague	4	2	4	5	6	*4
Mojave	4	4	1	3	3	3
Booker Washington	3	3	2	6	3	1
Hygrometer	2	5	3	4	5	4
Cambodia	8	7	9	7	6	2
Hypnosis	7	7	10	8	8	*7
School uniforms	7	7	6	4	3	3
Artists 1700	1	1	0	1	2	*1
Canadian b. Codes	9	9	8	6	6	*9
FHA	1	3	2	0	2	3
<b>Avg. correct pages over first 10</b>	5,125	5,291667	5,125	5,291667	6,291667	4,5
<b>%Variation with respect to plain query words</b>		<b>3,25%</b>	<b>1,63%</b>	<b>3,25%</b>	<b>22,76%</b>	<b>-12,20%</b>