INFORMATION RETRIEVAL USING ROBUST NATURAL LANGUAGE PROCESSING

Tomek Strzalkowski

Courant Institute of Mathematical Sciences
New York University
715 Broadway, rm. 704
New York, NY 10003
tomek@cs.nyu.edu

ABSTRACT

We developed a fully automated Information Retrieval System which uses advanced natural language processing techniques to enhance the effectiveness of traditional key-word based document retrieval. In early experiments with the standard CACM-3204 collection of abstracts, the augmented system has displayed capabilities that made it clearly superior to the purely statistical base system.

1. OVERALL DESIGN

Our information retrieval system consists of a traditional statistical backbone (Harman and Candela, 1989) augmented with various natural language processing components that assist the system in database processing (stemming, indexing, word and phrase clustering, selectional restrictions), and translate a user's information request into an effective query. This design is a careful compromise between purely statistical non-linguistic approaches and those requiring rather accomplished (and expensive) semantic analysis of data, often referred to as 'conceptual retrieval'. The conceptual retrieval systems, though quite effective, are not yet mature enough to be considered in serious information retrieval applications, the major problems being their extreme inefficiency and the need for manual encoding of domain knowledge (Mauldin, 1991).

In our system the database text is first processed with a fast syntactic parser. Subsequently certain types of phrases are extracted from the parse trees and used as compound indexing terms in addition to single-word terms. The extracted phrases are statistically analyzed as syntactic contexts in order to discover a variety of similarity links between smaller subphrases and words occurring in them. A further filtering process maps these similarity links onto semantic relations (generalization, specialization, synonymy, etc.) after which they are used to transform user's request into a search query.

The user's natural language request is also parsed, and all indexing terms occurring in them are identified. Next, certain highly ambiguous (usually single-word) terms are dropped, provided that they also occur as elements in some compound terms. For example, "natural" is deleted from a query already containing "natural language" because

"natural" occurs in many unrelated contexts: "natural number", "natural logarithm", "natural approach", etc. At the same time, other terms may be added, namely those which are linked to some query term through admissible similarity relations. For example, "fortran" is added to a query containing the compound term "program language" via a specification link. After the final query is constructed, the database search follows, and a ranked list of documents is returned.

It should be noted that all the processing steps, those performed by the backbone system, and these performed by the natural language processing components, are fully automated, and no human intervention or manual encoding is required.

2. FAST PARSING WITH TTP

TTP (Tagged Text Parser) is based on the Linguistic String Grammar developed by Sager (1981). Written in Quintus Prolog, the parser currently encompasses more than 400 grammar productions. It produces regularized parse tree representations for each sentence that reflect the sentence's logical structure. The parser is equipped with a powerful skip-and-fit recovery mechanism that allows it to operate effectively in the face of ill-formed input or under a severe time pressure. In the recent experiments with approximately 6 million words of English texts,¹ the parser's speed averaged between 0.45 and 0.5 seconds per sentence, or up to 2600 words per minute, on a 21 MIPS SparcStation ELC. Some details of the parser are discussed below.²

TTP is a full grammar parser, and initially, it attempts to generate a complete analysis for each sentence. However, unlike an ordinary parser, it has a built-in timer which regulates the amount of time allowed for parsing any one sentence. If a parse is not returned before the allotted time

¹ These include CACM-3204, MUC-3, and a selection of nearly 6,000 technical articles extracted from Computer Library database (a Ziff Communications Inc. CD-ROM).

² A complete description can be found in (Strzalkowski, 1991).

clapses, the parser enters the skip-and-fit mode in which it will try to "fit" the parse. While in the skip-and-fit mode, the parser will attempt to forcibly reduce incomplete constituents, possibly skipping portions of input in order to restart processing at a next unattempted constituent. In other words, the parser will favor reduction to backtracking while in the skip-and-fit mode. The result of this strategy is an approximate parse, partially fitted using top-down predictions. The fragments skipped in the first pass are not thrown out, instead they are analyzed by a simple phrasal parser that looks for noun phrases and relative clauses and then attaches the recovered material to the main parse structure. As an illustration, consider the following sentence taken from the CACM-3204 corpus:

The method is illustrated by the automatic construction of both recursive and iterative programs operating on natural numbers, lists, and trees, in order to construct a program satisfying certain specifications *a theorem induced by those specifications is proved*, and the desired program is extracted from the proof.

The italicized fragment is likely to cause additional complications in parsing this lengthy string, and the parser may be better off ignoring this fragment altogether. To do so successfully, the parser must close the currently open constituent (i.e., reduce *a program satisfying certain specifications to NP*), and possibly a few of its parent constituents, removing corresponding productions from further consideration, until an appropriate production is reactivated. In this case, TTP may force the following reductions: $SI \rightarrow to\ V\ NP$; $SA \rightarrow SI$; $S \rightarrow NP\ V\ NP\ SA$, until the production $S \rightarrow S\ and\ S$ is reached. Next, the parser skips input to find *and*, and resumes normal processing.

As may be expected, the skip-and-fit strategy will only be effective if the input skipping can be performed with a degree of determinism. This means that most of the lexical level ambiguity must be removed from the input text, prior to parsing. We achieve this using a stochastic parts of speech tagger³ to preprocess the text.

3. WORD SUFFIX TRIMMER

Word stemming has been an effective way of improving document recall since it reduces words to their common morphological root, thus allowing more successful matches. On the other hand, stemming tends to decrease retrieval precision, if care is not taken to prevent situations where otherwise unrelated words are reduced to the same stem. In our system we replaced a traditional morphological stemmer with a conservative dictionary-assisted suffix trimmer.⁴ The suffix trimmer performs essentially two tasks: (1) it reduces inflected word forms to their root forms as specified in the dictionary, and (2) it converts nominalized verb

forms (eg. "implementation", "storage") to the root forms of corresponding verbs (i.e., "implement", "store"). This is accomplished by removing a standard suffix, eg. "stor+age", replacing it with a standard root ending ("+e"), and checking the newly created word against the dictionary, i.e., we check whether the original root ("storage") is defined using the new root ("store"). This allows reducing "diversion" to "diverse" while preventing "version" to be replaced by "verse". Experiments with CACM-3204 collection show an improvement in retrieval precision by 6% to 8% over the base system equipped with a standard morphological stemmer (the SMART stemmer).

4. HEAD-MODIFIER STRUCTURES

Syntactic phrases extracted from TTP parse trees are head-modifier pairs: from simple word pairs to complex nested structures. The head in such a pair is a central element of a phrase (verb, main noun, etc.) while the modifier is one of the adjunct arguments of the head.⁵ For example, the phrase *fast algorithm for parsing context-free languages* yields the following pairs: *algorithm+fast*, *algorithm+parse*, *parse+language*, *language+context free*. The following types of pairs were considered: (1) a head noun and its left adjective or noun adjunct, (2) a head noun and the head of its right adjunct, (3) the main verb of a clause and the head of its object phrase, and (4) the head of the subject phrase and the main verb. These types of pairs account for most of the syntactic variants for relating two words (or simple phrases) into pairs carrying compatible semantic content. For example, the pair [retrieve,information] is extracted from any of the following fragments: *information retrieval system*; *retrieval of information from databases*; and *information that can be retrieved by a user-controlled interactive search process*.⁶ An example is shown in the appendix.⁷

5. TERM CORRELATIONS FROM TEXT

Head-modifier pairs form compound terms used in database indexing. They also serve as occurrence contexts for smaller terms, including single-word terms. In order to determine whether such pairs signify any important association between terms, we calculate the value of the

⁵ In the experiments reported here we extracted head-modifier word pairs only. CACM collection is too small to warrant generation of larger compounds, because of their low frequencies.

⁶ To deal with nominal compounds we use frequency information about the pairs generated from the entire corpus to form preferences in ambiguous situations, such as *natural language processing* vs. *dynamic information processing*.

⁷ Note that working with the parsed text ensures a high degree of precision in capturing the meaningful phrases, which is especially evident when compared with the results usually obtained from either unprocessed or only partially processed text (Lewis and Croft, 1990).

³ Courtesy of Bolt Beranek and Newman.

⁴ We use Oxford Advanced Learner's Dictionary (OALD) MRD.

Informational Contribution (IC) function for each element in a pair. Higher values indicate stronger association, and the element having the largest value is considered semantically dominant. IC function is a derivative of Fano's mutual information formula recently used by Church and Hanks (1990) to compute word co-occurrence patterns in a 44 million word corpus of Associated Press news stories. They noted that while generally satisfactory, the mutual information formula often produces counterintuitive results for low-frequency data. This is particularly worrisome for relatively smaller IR collections since many important indexing terms would be eliminated from consideration. Therefore, following suggestions in Wilks et al. (1990), we adopted a revised formula that displays a more stable behavior even on very low counts. This new formula $IC(x, [x, y])$ is based on (an estimate of) the conditional probability of seeing a word y to the right of the word x , modified with a dispersion parameter for x .

$$IC(x, [x, y]) = \frac{f_{x,y}}{n_x + d_x - 1}$$

where $f_{x,y}$ is the frequency of $[x, y]$ in the corpus, n_x is the number of pairs in which x occurs at the same position as in $[x, y]$, and $d(x)$ is the dispersion parameter understood as the number of distinct words with which x is paired. When $IC(x, [x, y]) = 0$, x and y never occur together (i.e., $f_{x,y} = 0$); when $IC(x, [x, y]) = 1$, x occurs only with y (i.e., $f_{x,y} = n_x$ and $d_x = 1$). Selected examples generated from CACM-3204 corpus are given in Table 2 at the end of the paper. IC values for terms become the basis for calculating term-to-term similarity coefficients. If two terms tend to be modified with a number of common modifiers and otherwise appear in few distinct contexts, we assign them a similarity coefficient, a real number between 0 and 1. The similarity is determined by comparing distribution characteristics for both terms within the corpus: how much information contents do they carry, do their information contribution over contexts vary greatly, are the common contexts in which these terms occur specific enough? In general we will credit high-contents terms appearing in identical contexts, especially if these contexts are not too commonplace.⁸ The relative similarity between two words x_1 and x_2 is obtained using the following formula (α is a large constant):

$$SIM(x_1, x_2) = \log(\alpha \sum_y sim_y(x_1, x_2))$$

where

$$sim_y(x_1, x_2) = MIN(IC(x_1, [x_1, y]), IC(x_2, [x_2, y])) \\ * MIN(IC(y, [x_1, y]), IC(y, [x_2, y]))$$

The similarity function is further normalized with respect to

⁸ It would not be appropriate to predict similarity between *language* and *logarithm* on the basis of their co-occurrence with *natural*.

$SIM(x_1, x_1)$. It may be worth pointing out that the similarities are calculated using term co-occurrences in syntactic rather than in document-size contexts, the latter being the usual practice in non-linguistic clustering (eg. Sparck Jones and Barber, 1971; Crouch, 1988; Lewis and Croft, 1990). Although the two methods of term clustering may be considered mutually complementary in certain situations, we believe that more and stronger associations can be obtained through syntactic-context clustering, given sufficient amount of data and a reasonably accurate syntactic parser.⁹

6. QUERY EXPANSION

Similarity relations are used to expand user queries with new terms, in an attempt to make the final search query more comprehensive (adding synonyms) and/or more pointed (adding specializations).¹⁰ It follows that not all similarity relations will be equally useful in query expansion, for instance, complementary relations like the one between *algol* and *fortran* may actually harm system's performance, since we may end up retrieving many irrelevant documents. Similarly, the effectiveness of a query containing *fortran* is likely to diminish if we add a similar but far more general term such as *language*. On the other hand, database search is likely to miss relevant documents if we overlook the fact that *fortran* is a *programming language*, or that *interpolate* is a specification of *approximate*. We noted that an average set of similarities generated from a text corpus contains about as many "good" relations (synonymy, specialization) as "bad" relations (antonymy, complementation, generalization), as seen from the query expansion viewpoint. Therefore any attempt to separate these two classes and to increase the proportion of "good" relations should result in improved retrieval. This has indeed been confirmed in our experiments where a relatively crude filter has visibly increased retrieval precision.

In order to create an appropriate filter, we expanded the IC function into a global specificity measure called the *cumulative informational contribution function* (ICW). ICW is calculated for each term across all contexts in which it occurs. The general philosophy here is that a more specific word/phrase would have a more limited use, i.e., would appear in fewer *distinct* contexts. ICW is similar to the standard *inverted document frequency* (*idf*) measure except that term frequency is measured over syntactic units rather than

⁹ Non-syntactic contexts cross sentence boundaries with no fuss, which is helpful with short, succinct documents (such as CACM abstracts), but less so with longer texts.

¹⁰ Query expansion (in the sense considered here, though not quite in the same way) has been used in information retrieval research before (eg. Sparck Jones and Tait, 1984; Harman, 1988), usually with mixed results. An alternative is to use term clusters to create new terms, "metaterms", and use them to index the database instead (eg. Crouch, 1988; Lewis and Croft, 1990). We found that the query expansion approach gives the system more flexibility, for instance, by making room for hypertext-style topic exploration via user feedback.

document size units.¹¹ Terms with higher ICW values are generally considered more specific, but the specificity comparison is only meaningful for terms which are already known to be similar. The new function is calculated according to the following formula:¹²

$$ICW(w) = IC_L(w) * IC_R(w)$$

where (with $n_w, d_w > 0$):

$$IC_L(w) = IC([w, _]) = \frac{n_w}{d_w(n_w + d_w - 1)}$$

and analogously for $IC_R(w)$.

For any two terms w_1 and w_2 , and a constant $\delta > 1$, if $ICW(w_2) \geq \delta * ICW(w_1)$ then w_2 is considered more specific than w_1 . In addition, if $SIM_{norm}(w_1, w_2) = \sigma > \theta$, where θ is an empirically established threshold, then w_2 can be added to the query containing term w_1 with weight σ .¹³ In the CACM-3204 collection:

$$\begin{aligned} ICW(algol) &= 0.0020923 \\ ICW(language) &= 0.0000145 \\ ICW(approximate) &= 0.0000218 \\ ICW(interpolate) &= 0.0042410 \end{aligned}$$

Therefore *interpolate* can be used to specialize *approximate*, while *language* cannot be used to expand *algol*. Note that if δ is well chosen (we used $\delta=10$), then the above filter will also help to reject antonymous and complementary relations, such as $SIM_{norm}(pl_i, cobol)=0.685$ with $ICW(pl_i)=0.0175$ and $ICW(cobol)=0.0289$. We continue working to develop more effective filters. Examples of filtered similarity relations obtained from CACM-3204 corpus are given in Table 3.

7. SUMMARY OF RESULTS

The preliminary series of experiments with the CACM-3204 collection of computer science abstracts showed a consistent improvement in performance: the average precision increased from 32.8% to 37.1% (a 13% increase), while the normalized recall went from 74.3% to 84.5% (a 14% increase), in comparison with the statistics of the base system. This improvement is a combined effect of the new stemmer, compound terms, term selection in queries, and query expansion using filtered similarity relations. The choice of similarity relation filter has been found critical in improving retrieval precision through query expansion. It should also be pointed out that only about 1.5% of all

similarity relations originally generated from CACM-3204 were found admissible after filtering, contributing only 1.2 expansion on average per query. It is quite evident significantly larger corpora are required to produce more dramatic results.^{14 15} A detailed summary is given in Table 1 below.

These results, while modest by IR standards, are significant for another reason as well. They were obtained without any manual intervention into the database or queries, and without using any other information about the database except for the text of the documents (i.e., not even the hand generated keyword fields enclosed with most documents were used). Lewis and Croft (1990), and Croft et al. (1991) report results similar to ours but they take advantage of Computer Reviews categories manually assigned to some documents. The purpose of this research is to explore the potential of automated NLP in dealing with large scale IR problems, and not necessarily to obtain the best possible results on any particular data collection. One of our goals is to point a feasible direction for integrating NLP into the traditional IR (Strzalkowski and Vauthey, 1991; Grishman

Tests	org.system	suff.trimmer	query exp.
Recall	Precision		
0.00	0.764	0.775	0.793
0.10	0.674	0.688	0.700
0.20	0.547	0.547	0.573
0.30	0.449	0.479	0.486
0.40	0.387	0.421	0.421
0.50	0.329	0.356	0.372
0.60	0.273	0.280	0.304
0.70	0.198	0.222	0.226
0.80	0.146	0.170	0.174
0.90	0.093	0.112	0.114
1.00	0.079	0.087	0.090
Avg. Prec.	0.328	0.356	0.371
% change		8.3	13.1
Norm Rec.	0.743	0.841	0.842
Queries	50	50	50

Table 1. Recall/precision statistics for CACM-3204

¹¹ We believe that measuring term specificity over document-size contexts (eg. Sparck Jones, 1972) may not be appropriate in this case. In particular, syntax-based contexts allow for processing texts without any internal document structure.

¹² Slightly simplified here.

¹³ The filter was most effective at $\sigma = 0.57$.

¹⁴ KL Kwok (private communication) has suggested that the low percentage of admissible relations might be similar to the phenomenon of 'tight clusters' which while meaningful are so few that their impact is small.

¹⁵ A sufficiently large text corpus is 20 million words or more. This has been partially confirmed by experiments performed at the University of Massachusetts (B. Croft, private communication).

and Strzalkowski, 1991).

ACKNOWLEDGEMENTS

We would like to thank Donna Harman of NIST for making her IR system available to us. We would also like to thank Ralph Weischedel and Marie Meteor of BBN for providing and assisting in the use of the part of speech tagger. KL Kwok has offered many helpful comments on an earlier draft of this paper. In addition, ACM has generously provided us with text data from the Computer Library database distributed by Ziff Communications Inc. This paper is based upon work supported by the Defense Advanced Research Project Agency under Contract N00014-90-J-1851 from the Office of Naval Research, and the National Science Foundation under Grant IRI-89-02304.

REFERENCES

1. Harman, Donna and Gerald Candela. 1989. "Retrieving Records from a Gigabyte of text on a Minicomputer Using Statistical Ranking." *Journal of the American Society for Information Science*, 41(8), pp. 581-589.
2. Mauldin, Michael. 1991. "Retrieval Performance in Ferret: A Conceptual Information Retrieval System." *Proceedings of ACM SIGIR-91*, pp. 347-355.
3. Sager, Naomi. 1981. *Natural Language Information Processing*. Addison-Wesley.
4. Strzalkowski, Tomek. 1991. "TTP: A Fast and Robust Parser for Natural Language." *Proteus Project Memo #43*, Courant Institute of Mathematical Science, New York University.
5. Lewis, David D. and W. Bruce Croft. 1990. "Term Clustering of Syntactic Phrases". *Proceedings of ACM SIGIR-90*, pp. 385-405.
6. Church, Kenneth Ward and Hanks, Patrick. 1990. "Word association norms, mutual information, and lexicography." *Computational Linguistics*, 16(1), MIT Press, pp. 22-29.
7. Wilks, Yorick A., Dan Fass, Cheng-Ming Guo, James E. McDonald, Tony Plate, and Brian M. Sator. 1990. "Providing machine tractable dictionary tools." *Machine Translation*, 5, pp. 99-154.
8. Sparck Jones, K. and E. O. Barber. 1971. "What makes automatic keyword classification effective?" *Journal of the American Society for Information Science*, May-June, pp. 166-175.
9. Crouch, Carolyn J. 1988. "A cluster-based approach to thesaurus construction." *Proceedings of ACM SIGIR-88*, pp. 309-320.
10. Sparck Jones, K. and J. I. Tait. 1984. "Automatic search term variant generation." *Journal of Documentation*, 40(1), pp. 50-66.
11. Harman, Donna. 1988. "Towards interactive query expansion." *Proceedings of ACM SIGIR-88*, pp. 321-331.
12. Sparck Jones, Karen. 1972. "Statistical interpretation of term specificity and its application in retrieval." *Journal of Documentation*, 28(1), pp. 11-20.

13. Croft, W. Bruce, Howard R. Turtle, and David D. Lewis. 1991. "The Use of Phrases and Structured Queries in Information Retrieval." *Proceedings of ACM SIGIR-91*, pp. 32-45.
14. Strzalkowski, Tomek and Barbara Vauthey. 1991. "Fast Text Processing for Information Retrieval." *Proceedings of the 4th DARPA Speech and Natural Language Workshop*, Morgan-Kaufman, pp. 346-351.
15. Strzalkowski, Tomek and Barbara Vauthey. 1991. "Natural Language Processing in Automated Information Retrieval." *Proteus Project Memo #42*, Courant Institute of Mathematical Science, New York University.
16. Grishman, Ralph and Tomek Strzalkowski. 1991. "Information Retrieval and Natural Language Processing." Position paper at the workshop on Future Directions in Natural Language Processing in Information Retrieval, Chicago.

APPENDIX: SAMPLE DATA

DOCUMENT TEXT:

```
*RECORD*
*F* NO
2366
*F* TITLE
Complex gamma function with error control
*F* TEXT
An algorithm to compute the gamma function and
log gamma function of a complex variable is presented.
The standard algorithm is modified in several respects
to insure the continuity of the function value
and to reduce accumulation of round-off errors. In
addition to computation of function values, this
algorithm includes an object-time estimation of round-off
errors. Experimental data with regard to the
effectiveness of this error control are presented.
a fortran program for the algorithm appears in the
algorithms section of this issue.
```

HEAD+MODIFIER PAIRS EXTRACTED:

function+gamma	control+error
present+algorithm	algorithm+compute
compute+function	function+function
function+log	gamma+function
gamma+log	gamma+variable
variable+complex	modify+algorithm,
algorithm+standard	insure+continue,
reduce+accumulate	accumulate+error,
error+round_off	algorithm+include,
include+estimate	estimate+object_time,
estimate+error	error+round_off,
present+data	data+experimental
effective+control	control+error
programme+fortran	section+algorithm
algorithm+issue	

word	head+modifier pair	IC coeff.
<i>theory</i>	<i>theory+mathematical</i>	0.009
<i>mathematical</i>	<i>theory+mathematical</i>	0.016
<i>distribute</i>	<i>distribute+normal</i>	0.040
<i>normal</i>	<i>distribute+normal</i>	0.115
<i>minimum</i>	<i>minimum+relative</i>	0.200
<i>relative</i>	<i>minimum+relative</i>	0.016
<i>retrieve</i>	<i>retrieve+inform</i>	0.086
<i>inform</i>	<i>retrieve+inform</i>	0.004
<i>size</i>	<i>size+medium</i>	0.009
<i>medium</i>	<i>size+medium</i>	0.250
<i>editor</i>	<i>editor+text</i>	0.142
<i>text</i>	<i>editor+text</i>	0.025
<i>system</i>	<i>system+parallel</i>	0.001
<i>parallel</i>	<i>system+parallel</i>	0.014
<i>read</i>	<i>read+character</i>	0.023
<i>character</i>	<i>read+character</i>	0.007
<i>discuss</i>	<i>discuss+panel</i>	0.015
<i>panel</i>	<i>discuss+panel</i>	0.500
<i>implicate</i>	<i>implicate+legal</i>	0.035
<i>legal</i>	<i>implicate+legal</i>	0.083
<i>system</i>	<i>system+distribute</i>	0.002
<i>distribute</i>	<i>system+distribute</i>	0.037
<i>make</i>	<i>make+recommend</i>	0.024
<i>recommend</i>	<i>make+recommend</i>	0.142
<i>infer</i>	<i>infer+deductive</i>	0.095
<i>deductive</i>	<i>infer+deductive</i>	0.142
<i>make</i>	<i>make+arrange</i>	0.006
<i>arrange</i>	<i>make+arrange</i>	0.058
<i>share</i>	<i>share+resource</i>	0.054
<i>resource</i>	<i>share+resource</i>	0.042
<i>comprehend</i>	<i>comprehend+language</i>	0.142
<i>language</i>	<i>comprehend+language</i>	0.002
<i>syntax</i>	<i>syntax+language</i>	0.177
<i>language</i>	<i>syntax+language</i>	0.011
<i>science</i>	<i>science+compute</i>	0.580
<i>compute</i>	<i>science+compute</i>	0.081
<i>maintain</i>	<i>concept+maintain</i>	0.111
<i>cost</i>	<i>cost+maintain</i>	0.005

Table 2. IC coefficients obtained from CACM-3204

word1	word2	SIMnorm
<i>*aim</i>	<i>purpose</i>	0.434
<i>algorithm</i>	<i>technique</i>	0.416
<i>algorithm</i>	<i>method</i>	0.529
<i>acquire</i>	<i>train</i>	0.462
<i>*adjacency</i>	<i>pair</i>	0.499
<i>*algebraic</i>	<i>symbol</i>	0.514
<i>*american</i>	<i>standard</i>	0.719
<i>assert</i>	<i>infer</i>	0.783
<i>back-up</i>	<i>mini-max</i>	0.834
<i>*buddy</i>	<i>time-share</i>	0.622
<i>committee</i>	<i>*symposium</i>	0.469
<i>correct</i>	<i>theorem</i>	0.637
<i>babylonian</i>	<i>old</i>	0.716
<i>critical</i>	<i>final</i>	0.680
<i>best-fit</i>	<i>first-fit</i>	0.871
<i>bound-context</i>	<i>lr</i>	0.948
<i>*duplex</i>	<i>reliable</i>	0.437
<i>deletion</i>	<i>insert</i>	0.411
<i>earlier</i>	<i>previous</i>	0.550
<i>encase</i>	<i>minimum-area</i>	0.991
<i>give</i>	<i>present</i>	0.458
<i>imaginary</i>	<i>real</i>	0.535
<i>incomplete</i>	<i>miss</i>	0.850
<i>input</i>	<i>output</i>	0.401
<i>lead</i>	<i>*trail</i>	0.890
<i>*marriage</i>	<i>stable</i>	0.623
<i>mean</i>	<i>*standard</i>	0.634
<i>method</i>	<i>technique</i>	0.571
<i>memory</i>	<i>storage</i>	0.613
<i>match</i>	<i>recognize</i>	0.563
<i>lower</i>	<i>upper</i>	0.841
<i>minor</i>	<i>*woman</i>	0.689
<i>progress</i>	<i>*trend</i>	0.444
<i>purdue</i>	<i>stanford</i>	1.000
<i>range</i>	<i>variety</i>	0.600
<i>round-off</i>	<i>truncate</i>	0.918
<i>remote</i>	<i>teletype</i>	0.509
<i>pulse</i>	<i>wave</i>	0.591

Table 3. Filtered word similarities (* indicates the more specific term).