

Using Argumentation to Retrieve Articles with Similar Citations from MEDLINE

Imad Tbahriti^{1,2}, Christine Chichester¹, Frédérique Lisacek^{1,3}, Patrick Ruch⁴

¹Geneva Bioinformatics (GeneBio) SA, 25, avenue de Champel, Geneva

²Computer Science Dept., University of Geneva

³Swiss Institute of Bioinformatics, Geneva

⁴SIM, University Hospital of Geneva

imad.tbahriti;christine.chichester;frederique.lisacek@genebio.com - patrick.ruch@epfl.ch

Abstract

The aim of this study is to investigate the relationships between citations and the scientific argumentation found in the abstract. We extracted citation lists from a set of 3200 full-text papers originating from a narrow domain. In parallel, we recovered the corresponding MEDLINE records for analysis of the argumentative moves. Our argumentative model is founded on four classes: PURPOSE, METHODS, RESULTS, and CONCLUSION. A Bayesian classifier trained on explicitly structured MEDLINE abstracts generates these argumentative categories. The categories are used to generate four different argumentative indexes. A fifth index contains the complete abstract, together with the title and the list of Medical Subject Headings (MeSH) terms. To appraise the relationship of the moves to the citations, the citation lists were used as the criteria for determining relatedness of articles, establishing a benchmark. Our results show that the average precision of queries with the PURPOSE and CONCLUSION features is the highest, while the precision of the RESULTS and METHODS features was relatively low. A linear weighting combination of the moves is proposed, which significantly improves retrieval of related articles.

1 Introduction

Numerous techniques help researchers locate relevant documents in an ever-growing mountain of scientific publications. Among these techniques is the analysis of bibliographic information, which can identify conceptual connections between large numbers of articles. Although helpful, most of these systems deliver masses of documents to the researcher for analysis, which contain various degrees of similarity. This paper introduces a method to determine the similarity of a bibliographic co-citation list, that is the list of

citations that are shared between articles, and the argumentative moves of an abstract in an effort to define novel similarity searches.

Authors of biological papers develop arguments and present the justification for their experiments based on previously documented results. These results are represented as citations to earlier scientific literature and establish the links between old and new findings. The assumption is that the majority of scientific papers employing the same citations depict related viewpoints. The method described here is applied to improve retrieval of similar articles based on co-citations, but other applications are foreseen. Documents that should be conceptually correlated due to bibliographic relatedness but which propose different or novel arguments are often not easily located in the majority of bibliographically correlated articles. Our system can be tuned to identify these documents. Conversely, such a system could also be used as a platform to aid authors by means of automatic assembly or refinement of their bibliographies through the suggestion of citations coming from documents containing similar arguments.

The rest of this paper is structured as follows: section 2 describes the background related to experiments using citations or argumentation that compare aspects connected to the logical content of publications. Section 3 details the method and the generation of the different indexes used in our analyses, e.g. the citation index, the four argumentative indexes and the abstract index (abstract, title and keywords). Section 4 presents the results of the evaluations we performed. Section 5 closes with a summary of the contribution of this work, limitations and future work.

2 Background

Digital libraries aim at structuring their records to facilitate user navigation. Interfaces visualizing

overlapping relationships of the standard library fields such as author and title in document collections are usually the most accessible to the user. Beyond these well-known targets, researchers (see de Bruijn and Martin, 2002, or Hirschman and al. 2002, for a survey) interested in information extraction and retrieval for biomedical applications have mostly focused on studying specific biological interactions (Stapley and Benoit, 2000; Nédellec et al., 2002; Dobrokhotov et al., 2003) and related entities (Collier et al., 2000; Humphreys et al., 2000; Yu et al., 2002; Yamamoto et al., 2003; Albert et al., 2003) or using terms in biomedical vocabularies (Nazarenko et al., 2001; Ruch et al., 2004; Srinivasan and Hristovski, 2004). The use of bibliographical and argumentative information (McKnight and Srinivasan 2003) has been less well studied by researchers interested in applying natural language processing to biomedical texts.

2.1 Citations

Originating from bibliometrics, citation analysis (White, 2003) has been used to visualize a field via a representative slice of its literature. Co-citation techniques make it possible to cluster documents by scientific paradigm or hypothesis (Noyons et al., 1999). Braam et al., (1991) have investigated co-citation as a tool to map subject-matter specialties. They found that the combination of keyword analysis and co-citation analysis was useful in revealing the cognitive content of publications. Peters et al., (1995) further explored the citation relationships and the cognitive resemblance in scientific articles. Word profile similarities of publications that were bibliographically coupled by a single, highly cited article were compared with publications that were not bibliographically coupled to that specific article. A statistically significant relationship has been established between the content of articles and their shared citations. This result will serve as basis to establish our benchmark without relevance judgments (Wu and Crestani, 2003; Soboroff et al., 2001).

2.2 Argumentation in biomedical abstracts

Scientific research is often described as a problem solving activity. In full text scientific articles this problem-solution structure has been crystallized in a fixed presentation known as *Introduction, Methods, Results* and *Conclusion*. This structure is often presented in a much-compacted version in the abstract and it has been clearly demonstrated by Schuemie et al., (2004) that abstracts contain a higher information density than full text. Correspondingly, the 4-move problem-solving structure (standardized according to ISO/ANSI

guidelines) has been found quite stable in scientific reports (Orasan, 2001). Although the argumentative structure of an article is not always explicitly labeled, or can be labeled using slightly different markers (as seen in Figure 1), a similar implicit structure is common in most biomedical abstracts (Swales, 1990). Therefore, to find the most relevant argumentative status that describes the content of the article, we employed a classification method to separate the content dense sentences of the abstracts into the argumentative moves.

INTRODUCTION: Chromophobe renal cell carcinoma (CCRC) comprises 5% of neoplasms of renal tubular epithelium. CCRC may have a slightly better prognosis than clear cell carcinoma, but outcome data are limited. **PURPOSE:** In this study, we analyzed 250 renal cell carcinomas to a) determine frequency of CCRC at our Hospital and b) analyze clinical and pathologic features of CCRCs. **METHODS:** A total of 250 renal carcinomas were analyzed between March 1990 and March 1999. Tumors were classified according to well-established histologic criteria to determine stage of disease; the system proposed by Robson was used. **RESULTS:** Of 250 renal cell carcinomas analyzed, 36 were classified as chromophobe renal cell carcinoma, representing 14% of the group studied. The tumors had an average diameter of 14 cm. Robson staging was possible in all cases, and 10 patients were stage 1) 11 stage II; 10 stage III, and five stage IV. The average follow-up period was 4 years and 18 (53%) patients were alive without disease. **CONCLUSION:** The highly favorable pathologic stage (RI-RII, 58%) and the fact that the majority of patients were alive and disease-free suggested a more favorable prognosis for this type of renal cell carcinoma.

Figure 1: Example of an explicitly structured abstract in MEDLINE. The 4-class argumentation model is sometimes split into classes that may carry slightly different names, as illustrated in this example by the INTRODUCTION marker.

3 Methods

We established a benchmark based on citation analysis to evaluate the impact of using argumentation to find related articles. In information retrieval, benchmarks are developed from three resources: a document collection, a query collection and a set of relevance rankings that relates each query to the set of documents. Existing information retrieval collections normally contain user queries composed of only a few words. These short queries are not suitable for evaluating a system tailored to retrieve articles with similar citations. Therefore, we have created the collection and tuned the system to accept long queries such as abstracts (Figure 2).

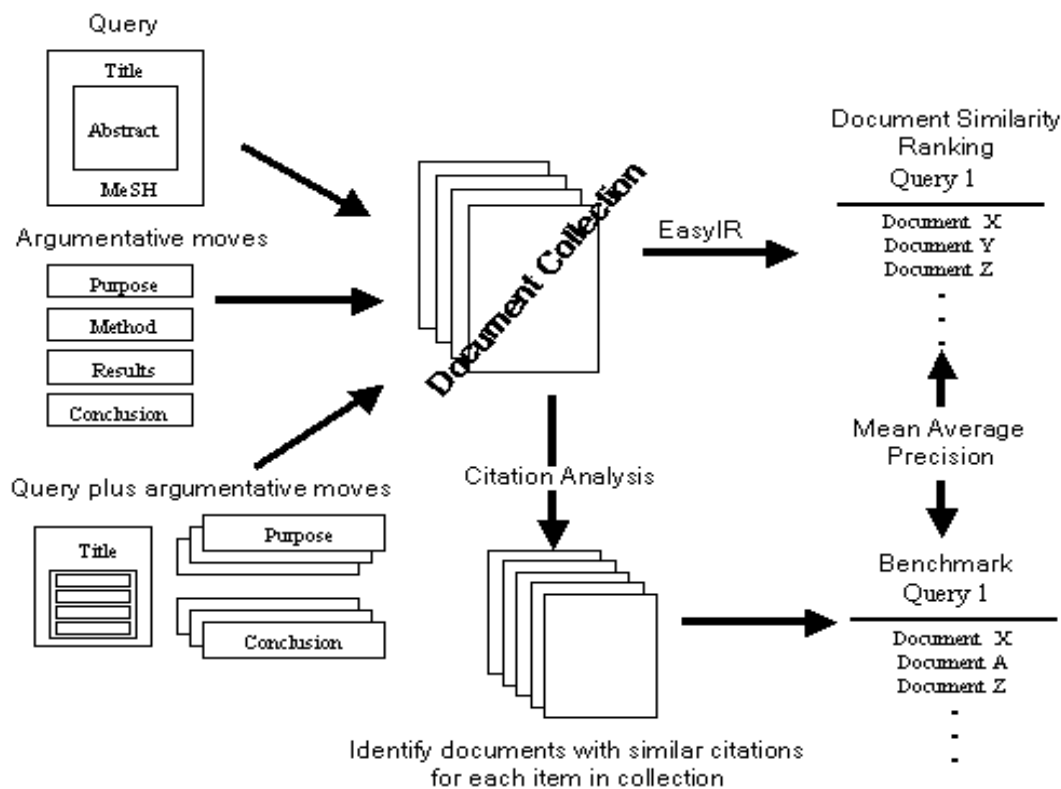


Figure 2: Flowchart for the chain of experimental procedures. The benchmark was assembled from citations shared between documents and compared to the document similarity ranking of EasyIR.

3.1 Data acquisition and citation indexing

All the data used in these experiments were acquired from MEDLINE using the PubMed interface.

Document Collection. The document set was obtained from PubMed by executing a set of Boolean queries to recover articles related to small active peptides from many animal species excluding humans. These peptides hold the promise of becoming novel therapeutics. The set consisted of 12500 documents, which were comprised of abstract, title and MeSH terms. For 3200 of these documents we were able to recover the full text including the references for citation extraction and analysis.

Queries. Following statistical analysis confirmed by Buckley and Voorhees (2000), four sets of 25 articles were selected from the 3200 full text articles. The title, abstract and MeSH terms fields were used to construct the queries. For testing the influence the argumentative move, the specific sentences were extracted and tested either alone or in combination with the queries that contained the title, abstract and MeSH terms.

Citation analysis. Citation lists were automatically extracted from 3200 full-text articles that were correspondingly represented within the document set. This automatic parsing of citations was manually validated. Each citation was represented as a unique ID for comparison

purposes. Citation analysis of the entire collection demonstrated that the full-text articles possessed a mean citation count of 28.30 ± 24.15 (mean \pm S.D.) with a 95% CI = 27.47 – 29.13. Within these records the mean co-citation count was 7.79 ± 6.99 (mean \pm S.D.) with a 95% CI = 7.55 – 8.03. As would be expected in a document set which contains a variety of document types (reviews, journal articles, editorials), the standard deviations of these values are quite large.

Citation benchmark. For each set of queries, a benchmark was generated from the 10 cited articles that contained the greatest number of co-citations in common with the query. For the benchmark, the average number of cited articles that have more than 9 co-citations was 15.70 ± 6.58 (mean \pm S.D.). Query sets were checked to confirm that at least one sentence in each abstract was classified per argumentative class.

3.2 Metrics

The main measure for assessing information retrieval engines is mean average precision (MAP). MAP is the standard metric although it may tend to hide minor differences in ranking (Mittendorf and Schäuble, 1996).

3.3 Text indexing

For indexing, we used the easyIR system¹, which implements standard vector space IR schemes (Salton et al., 1983). The term-weighting schema

¹ <http://lithwww.epfl.ch/~ruch/softs/softs.html>.

composed of combinations of term frequency, inverse document frequency and length normalization was varied to determine the most relevant output ranking. Table 1 gives the most common term weighting factors (atc.atn, ltc.atn); the first letter triplet applies to the document, the second letter triplet applies to the query (Ruch, 2002).

Term Frequency	
First Letter	$f(tf)$
n (natural)	tf
l (logarithmic)	$1 + \log(tf)$
a (augmented)	$0.5 + 0.5 \times (tf/\max(tf))$
Inverse Document Frequency	
Second Letter	$f(1/df)$
n(no)	1
t(full)	$\log(N/df)$
Normalization	
Third Letter	$f(\text{length})$
n(no)	1
c(cosine)	$\sqrt{\rho_1^2 + \rho_2^2 + \dots + \rho_n^2}$

Table 1. Weighting parameters, following SMART conventions.

3.4 Argumentative classification

The classifier segmented the abstracts into 4 argumentative moves: PURPOSE, METHODS, RESULTS, and CONCLUSION.

CONCLUSION [00160116]	The highly favorable pathologic stage (RI-RII, 58%) and the fact that the majority of patients were alive and disease-free suggested a more favorable prognosis for this type of renal cell carcinoma.
METHODS [00160119]	Tumors were classified according to well-established histologic criteria to determine stage of disease; the system proposed by Robson was used.
METHODS [00162303]	Of 250 renal cell carcinomas analyzed, 36 were classified as chromophobe renal cell carcinoma, representing 14% of the group studied.
PURPOSE [00156456]	In this study, we analyzed 250 renal cell carcinomas to a) determine frequency of CCRC at our Hospital and b) analyze clinical and pathologic features of CCRCs.
PURPOSE [00167817]	Chromophobe renal cell carcinoma (CCRC) comprises 5% of neoplasms of renal tubular epithelium. CCRC may have a slightly better prognosis than clear cell carcinoma, but outcome data are limited.
RESULTS [00155338]	Robson staging was possible in all cases, and 10 patients were stage 1) 11 stage II; 10 stage III, and five stage IV.

Figure 3: The classification results for the abstract shown in Figure 1. In each box, the attributed class is first, followed by the score for the class, followed by the extracted text segment. In this example, one of RESULTS sentences is misclassified as METHODS

The classification unit is the sentence which means that abstracts are preprocessed using an ad hoc sentence splitter. The confusion matrix for the four argumentative moves generated by the classifier is given in Table 2. This evaluation used explicitly structured abstracts; therefore, the argumentative

markers were removed prior to the evaluation. Figure 3 shows the output from the classifier, when applied to the abstract shown in Figure 1. After extraction, each of the four types of argumentative moves was then used for indexing, retrieval and comparison tasks.

	PURP	METH	RESU	CONC
PURP	93.55%	0%	3.23%	3%
METH	8%	81%	8%	3%
RESU	7.43%	5.31%	74.25%	13.01%
CONC	2.27%	0%	2.27%	95.45%

Table 2. Confusion matrices for each argumentative class.

3.5 Argumentative combination

We adjusted the weight of the four argumentative moves, based on their location and then combined them to improve retrieval effectiveness. The query weights were recomputed as indicated in equation (1).

$$W_{\text{new}} = W_{\text{old}} * S_c * k_c \quad (1)$$

$c \in \{\text{PURPOSE; METHODS; RESULTS; CONCLUSION}\}$

W_{old} : the feature weight as given by the query weighting (ltc)

S : the *normalized* score attributed by the argumentative classifier to each sentence in the abstract. This score is attributed to each feature appearing in the considered segment

k : a constant for each value of c . The value is set empirically using the tuning set (TS). The initial value of k for each category is given by the distribution observed in Table 4 (i.e., 0.625, 0.164, 0.176, 0.560 for the classes, PURPOSE, METHODS, RESULTS and CONCLUSION respectively), and then an increment step (positive and negative) is varied to get the most optimal combination.

This equation combines the score (S_c) attributed by the original weighting (ltc) for each feature (W_{old}) found in the query with a boosting factor (k_c). The boosting factor was derived from the score provided by the argumentative classifier for each classified sentence. For these experiments, the parameters were determined with a tuning set (TS), one of the four query sets, and the final evaluation was done using the remaining three sets, the validation sets (VS). The document feature factor (atn) remained unchanged.

4 Results

In this section, we described the generation of the baseline measure and the effects of different conditions on this baseline.

4.1 Comparison of text index parameters

The use of a domain specific thesaurus tends to improve the MAP when compared to the citation benchmark, 0.1528 vs. 0.1517 for ltc.atn and 0.1452 vs. 0.1433 for atc.atn (Table 3). The ltc.atn weighting schema in combination with the thesaurus produced the best results, therefore these parameters were more likely to retrieve abstracts found in the citation index and thus were used for all subsequent experiments.

	Set 1	Set 2	Set 3	Set 4	Average
atc.atn	0.1402	0.1417	0.1438	0.1476	0.1433
atc.atn + T	0.1440	0.1431	0.1477	0.1465	0.1452
ltc.atn	0.1505	0.1528	0.1506	0.1529	0.1517
ltc.atn + T	0.1524	0.1534	0.1530	0.1539	0.1532

Table 3. Mean average precision (MAP) for each query set (1,2,3, and 4) with different term weighting schemas. The last column gives the average MAP. T represents the thesaurus

4.2 Argumentation-based retrieval

For demonstrating that argumentative features can improve document retrieval, we first determined which argumentative class was the most content bearing. Subsequently, we combined the four argumentative classes to again improve document retrieval.

	PURP	METH	RESU	CONC	ltc.atn + T
MAP	0.0958 (62.5%)	0.0251 (16.4%)	0.0270 (17.6%)	0.0858 (56.0%)	0.1532

Table 4. MAP results from querying the collection using only the argumentative move.

To determine the value of each argumentative move in the retrieval, the argumentative categorizer first parses each query abstract, generating four groups each representing a unique argumentative class. The document collection was separately queried with each group. Table 4 gives the MAP measures for each type of argumentation. Table 4 shows the sentences classified as PURPOSE provide the most useful content to retrieve similar documents. Baseline precision of 62.5% is achieved when using only this section of the abstract. The CONCLUSION move is the second most valuable at 56% of the baseline. The METHODS and RESULTS sections appear less content bearing for retrieving similar documents, 16.4% and 17.6%, respectively, of the baseline. Each argumentative set represents roughly a quarter of the textual content of the original abstract. Querying with the PURPOSE section, (25% of the available textual material) realizes almost 2/3 of the average precision and for the CONCLUSION section, it is more than 50% of the baseline precision. In information retrieval queries and documents are often seen as symmetrical

elements. This fact may imply the possible use of the argumentative moves as a technique to reduce the size of the indexed document collection or to help indexing pruning in large repositories (Carmel and al. 2001).

4.3 Argumentative overweighting

As implied in Table 4, Table 5 confirms that overweighting the features of PURPOSE and CONCLUSION sentences results in a gain in average precision (respectively +3.39% and +3.98 for CONCLUSION and PURPOSE) as measured by citation similarity. More specifically, Table 5 demonstrates the use of PURPOSE and CONCLUSION as follows:

- PURPOSE applies a boosting coefficient to features classified as PURPOSE by the argumentative classifier;
- CONCLUSION applies a boosting coefficient to features classified as CONCLUSION by the argumentative classifier;
- COMBINATION applies two different boosting coefficients to features classified as CONCLUSION and PURPOSE by the argumentative classifier.

The results, in Table 5, from boosting PURPOSE and CONCLUSION features are given alongside the MAP and show an improvement of precision at the 5 and 10 document level. At the 5-document level the advantage is with the PURPOSE features, but at the 10-document level boosting the CONCLUSION features is more effective. While the improvement brought by boosting PURPOSE and CONCLUSION features, when measured by MAP is modest (3-4%), the improvement observed by their optimal combination reached a significant improvement: + 5.48%. The various combinations of RESULTS and METHODS sections did not lead to any improvement.

Argumentation has typically been studied in relation to summarization (Teufel and Moens, 2002). Its impact on information retrieval is more difficult to establish although recent experiments (Ruch et al., 2003) tend to confirm that argumentation is useful for information extraction, as demonstrated by the extraction of gene functions for LocusLink curation. Similarly, using the argumentative structure of scientific articles has been proposed to reduce noise (Camon et al., 2004) in the assignment of Gene Ontology codes as investigated in the BioCreative challenge. In particular, it was seen that the use of ‘Material and Methods’ sentences should be avoided. A fact which is confirmed by our results with the METHOD argumentative move.

	MAP	Precision at 5	Precision at 10
lrc.atn + T	0.1532	0.2080	0.1840
PURPOSE	0.1593 (+3.98%)	0.2240	0.1760
CONCLUSION	0.1584 (+3.39%)	0.2160	0.1920
COMBINATION	0.1616 (+5.48%)	0.2320	0.1960

Table 5. Retrieval results for the argumentative classes PURPOSE and CONCLUSION, and the combination of both classes.

5 Conclusion and Future work

We have reported on the construction of an information retrieval engine tailored to search for documents with similar citations in MEDLINE collections. The tool retrieves similar documents by giving more weight to features located in PURPOSE and CONCLUSION segments. The RESULTS and METHODS argumentative moves are reported here as less useful for such a retrieval task. Evaluated on a citation benchmark, the system significantly improves retrieval effectiveness of a standard vector-space engine. In this context, it would be interesting to investigate how argumentation can be beneficial to perform *ad hoc* retrieval tasks in MEDLINE (Kayaalp et al., 2003).

Evidently using citation information to build our benchmark raises some questions. Authors may refer to other work in many ways to benefit the tone of their argument. Specifically, there are two major citation contexts, one where an article is cited negatively or contrastively and one where an article is cited positively, or the authors state that their own work originates from the cited work. In this study we have not made a distinction between these contexts but we consider this as an avenue for building better representations of the cited articles in future work. Finally, we are now exploring the use of the tool to detect inconsistencies between articles. We hope to use citation and content analysis to identify articles containing novel views so as to expose differences in the consensus of the research area's intellectual focus. The idea is to retrieve documents having key citation similarity but show some dissimilarity regarding a given argumentative category.

Finally, we have observed that citation networks in digital libraries are analogous to hyperlinks in web repositories. Consequently using web-inspired similarity measures may be beneficial for our purposes. Of particular interest in relation to

argumentation, is the fact that citations networks, like web pages, are hierarchically nested graph with argumentative moves introducing intermediate levels (Bharat et al., 2001).

Acknowledgements

We would like to thank Patrick Brechbühl for his assistance in organizing the computing environment and Ron Appel for his support.

References

- S. Albert, S. Gaudan, H. Knigge, A. Raetsch, A. Delgado, B. Huhse, H. Kirsch, M. Albers, D. Rebholz-Schuhmann and M. Koegele. 2003. Computer-assisted generation of a protein-interaction database for nuclear receptors. *Journal of Molecular Endocrinology*, 17(8): 1555-1567.
- K. Bharat, B. Chang, M. Rauch Henzinger, M. Ruhl: Who Links to Whom: Mining Linkage between Web Sites. *ICDM 2001*: 51-58
- R. R. Braam, H.F. Moed, and A.F.J. van Raan. 1991 Mapping of science by combined co-citation and word analysis, I: Structural Aspects, *Journal of the American Society for Information Science*, 42 (4): 233-251.
- C. Buckley and E. M. Voorhees. 2000. Evaluating evaluation measure stability, *ACM SIGIR*, p. 33-40.
- E. Camon et al. Personal communication on BioCreative Task 2 Evaluation. *BMC Bioinformatics Special Issue on BioCreative.2004*. To be submitted.
- D. Carmel, E. Amitay, M. Herscovici, Y. Maarek, Y. Petruschka and A. Soffer: Juru at TREC 10 - Experiments with Index Pruning. *TREC 2001*
- N. Collier, C. Nobata and J.I. Tsujii. 2000. *Extracting the Names of Genes and Gene Products with a Hidden Markov Model*. *COLING 2000*. 201-207.
- B. de Bruijn and J. Martin. *Getting to the (c)ore of knowledge: mining biomedical literature*. 2002. In *International Journal of Medical Informatics*, P Ruch and R Baud, eds., pages 7-18, Volume 67, Issues 1-3, 4, p. 7-18
- P. B. Dobrokhotov, C. Goutte, A. L. Veuthey, and É. Gaussier: *Combining NLP and probabilistic categorisation for document and term selection for Swiss-Prot medical annotation*. 2003. *ISMB 2003*, 91-94.
- W. Hersh, S. Moy, D. Kraemer, L. Sacherek and D. Olson. 2003. More Statistical Power Needed: The OHSU TREC 2002 Interactive Track Experiments, *TREC 2002*.

- L Hirschman, JC Park, JI Tsujii, L Wong, C Wu: Accomplishments and challenges in literature data mining for biology. *Bioinformatics* 18(12): 1553-1561 (2002)
- K. Humphreys, G. Demetriou and R. Gaizauskas. 2000. *Two Applications of Information Extraction to Biological Science Journal Articles: Enzyme Interactions and Protein Structures* In Proceedings of the Workshop on Natural Language Processing for Biology, held at the Pacific Symposium on Biocomputing (PSB2000).
- M. Kayaalp, A.R. Aronson, S.M. Humphrey, N.C. Ide, L.K. Tanabe, L.H. Smith, D. Demner, R.R. Loane, J.G. Mork, and O. Bodenrieder. 2003. *Methods for accurate retrieval of MEDLINE citations in functional genomics*. In Notebook of the TREC-2003, pages 175-184, Gaithersburg, MD.
- L. McKnight and P. Srinivasan. 2003. *Categorization of Sentence Types in Medical Abstracts*. Proceedings of the 2003 AMIA conference.
- H. Mima, S. Ananiadou, G. Nenadic, and J. Tsujii. *A methodology for terminology-based knowledge acquisition and integration*, 2002. COLING. Morgan Kaufmann.
- E Mittendorf and P Schäuble. 1996. Measuring the effects of data corruption on information retrieval. SDAIR Proceedings.
- A. Nazarenko, P. Zweigenbaum, B. Habert and J. Bouaud. 2001. *Corpus-based Extension of a Terminological Semantic Lexicon*, Recent Advances in Computational Terminology. John Benjamins, 2001.
- C. Nédellec, M. Vetah and P. Bessières. 2001. *Sentence filtering for information extraction in genomics, a classification problem*. In Proceedings PKDD, pages 326-237, Springer-Verlag, Berlin.
- E.C.M. Noyons, H.F. Moed, and M. Luwel. 1999. A Bibliometric Study Combining Mapping and Citation Analysis for Evaluative Bibliometric Purposes. *Journal of the American Society for Information Science*, 50(2):115-131.
- C. Orasan. 2001. *Patterns in scientific abstracts*. In Proceedings of Corpus Linguistics, 433-445.
- H. P. F. Peters, R.R. Braam, and A.F.J. van Raan. 1995. Cognitive Resemblance and Citation Relations in Chemical Engineering Publications. *Journal of the American Society for Information Science*, 46 (1): 9-21.
- J.C. Reynar and A. Ratnaparkhi.. 1997. A maximum entropy approach to identifying sentence boundaries. In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, 16-19.
- P. Ruch. 2002. *Using Contextual Spelling Correction to Improve Retrieval Effectiveness in Degraded Text Collections*. COLING 2002. Morgan Kaufmann.
- P. Ruch, R. Baud and A. Geissbühler. 2003. Learning-free Text Categorization, *AIME*, M Dojat, E Keravnou and P Barahona (Eds.). 199-208, LNAI 2780. Springer.
- P. Ruch, C Chichester, G Cohen, G Coray, F Ehrler, H Ghorbel, H Müller, and V Pallotta. 2004. *Report on the TREC 2003 Experiment: Genomic Track*, TREC.
- M.J. Schuemie, M. Weeber, B.J.A Schijvenaars, E.M. van Mulligen, C.C. van der Eijk, R. Jeliert B. Mons, and J. A. Kors. 2004. Distribution of information in biomedical abstracts and full text publications. *Bioinformatics*. Submitted.
- I Soborrof, C. Nicholas and P. Cahan. 2001. Ranking Retrieval Systems without Relevance Judgments. SIGIR 2001: 66-73
- P. Srinivasan and D. Hristovski. 2004. *Distilling Conceptual Connections from MeSH Co-Occurrences*. MEDINFO 2004. Submitted.
- B Stapley and G Benoir. 2000. BioBibliometrics: information retrieval and visualisation from co-occurrences of gene names in MEDLINE abstracts. *Pac. Symp. Biocomp.* 5:526-537.
- J. Swales. 1990 *Genre analysis: English in academic and research settings*. Cambridge University Press, UK.
- S. Teufel and M. Moens: Summarizing Scientific Articles: Experiments with Relevance and Rhetorical Status. *Computational Linguistics* 28(4): 409-445 (2002)
- H. White. 2003. Pathfinder networks and author cocitation analysis: a remapping of paradigmatic information scientists. *J. Am. Soc. Inf. Sci. Technol* 54(5) 423-434.
- S. Wu and F. Crestani. 2003. Methods for Ranking Information Retrieval Systems Without Relevance Judgments. SAC 2003: 811-816. ACM.
- K. Yamamoto, T. Kudo, A. Konagaya and Y. Matsumoto. 2003. *Protein name tagging for biomedical annotation in text*. ACL Workshop on Natural Language Processing in Biomedicine, pp. 65-72, July 2003.
- H. Yu, V. Hatzivassiloglou, C. Friedman, I.H. Iossifov, A. Rzhetsky and W.J. Wilbur. 2002. *A rule-based approach for automatically identifying gene and protein names in MEDLINE abstracts: A proposal*. ISMB 2002.