

Legal Texts Summarization by Exploration of the Thematic Structures and Argumentative Roles

Atefeh Farzindar and Guy Lapalme

RALI, Département d'Informatique et Recherche Opérationnelle
Université de Montréal, Québec, Canada, H3C 3J7
{farzinda, lapalme}@iro.umontreal.ca

Abstract

In this paper we describe our method for the summarization of legal documents helping a legal expert determine the key ideas of a judgment. Our approach is based on the exploration of the document's architecture and its thematic structures in order to build a table style summary for improving coherency and readability of the text. We present the components of a system, called LetSum, built with this approach, its implementation and some preliminary evaluation results.

1 Introduction

The goal of a summary is to give the reader an accurate and complete idea of the contents of the source (Mani, 2001). In this research, we focused on a problem referred to as *legal text summarization*. As ever larger amounts of legal documents become available electronically, interest in automatic summarization has continued to grow in recent years. In this paper, we present our method for producing a very short text from a long legal document (a record of the proceedings of federal courts in Canada) and present it as a table style summary. The goal of this project is to develop a system to create a summary for the needs of lawyers, judges and experts in the legal domain. Our approach investigates the extraction of the most important units based on the identification of thematic structures of the document and the determination of semantic roles of the textual units in the judgment (Farzindar, 2004). The remainder of the paper is organized as follows. Section 2 introduces the motivation of the research and the context of the work. Section 3 reports on the results of our analysis of a corpus of legal abstracts written by professional abstractors. Section 4 describes our method for the exploration of document architecture and the components of the system that we have developed to produce a summary. Section 5 presents some related work in this domain. Section 6 concludes the paper and presents some preliminary evaluation results for the components of our

system.

2 Context of the Work

In Canada, the Canadian Legal Information Institute project (CANLII) aims at gathering legislative and judicial texts, as well as legal commentaries, from federal, provincial and territorial jurisdictions in order to make primary sources of Canadian law accessible for free on the Internet (<http://www.canlii.org>). The large volume of legal information in electronic form creates a need for the creation and production of powerful computational tools in order to extract relevant information in a condensed form.

But why are we interested in the processing of previous legal decisions and in their summaries? First, because a court order generally gives a solution to a legal problem between two or several parties. The decision also contains the reasons which justify the solution and constitute a *law jurisprudence precedent* from which it is possible to extract a legal rule that can be applied to similar cases. To find a solution to a legal problem not directly indicated in the law, lawyers look for *precedents* of similar cases. For a single query in a data base of law reports, we often receive hundreds of documents that are very long to study for which legal experts and law students request summaries. In Quebec REJB (Répertoire électronique de jurisprudence du Barreau) and SOQUIJ (Société québécoise d'information juridique) are two organizations which provide manual summaries for legal resources, but the human time and expertise required makes their services very expensive. For example the price of only one summary with its full text, provided by SOQUIJ is 7.50 \$ can. Some legal information systems have been developed by private companies like QuickLaw in Canada and WESTLAW and LEXIS in the United States, however no existing system completely satisfies the specific requirements of this field.

One reason for the difficulty of this work is the complexity of the domain: specific vocabularies of

Between:

JASPER NATIONAL PARK Applicants and THE ATTORNEY GENERAL OF CANADA Respondent,
Docket: T-1557-98

Judgment	Professional abstract	Role
[1] This application for judicial review arises out of a decision (the Decision) announced on or about the 30th of June 1998 by the Minister of Canadian Heritage (the Minister) to close the Maligne River (the River) in Jasper National Park to all boating activity, beginning in 1999.	Judicial review of Minister of Canadian Heritage's decision to close Maligne River in Jasper National Park to all boating activity beginning in 1999 to protect habitat of harlequin ducks.	INTRODUCTION
[7] The applicants offer commercial rafting trips to Park visitors in this area each year from mid-June to sometime in September.	Applicants offer commercial rafting trips on River.	CONTEXT
[10] Consequently, a further environmental assessment regarding commercial rafting on the Maligne River was prepared in 1991. The assessment indicated that rafting activity had expanded since 1986, with an adverse impact on Harlequin ducks along the Maligne River.	1991 environmental assessment indicating rafting having adverse impact on harlequin ducks along river.	CONTEXT

Table 1: Alignment of the units of the original judgment with the professional abstract

the legal domain and legal interpretations of expressions produce many ambiguities. For example, the word *sentence* can have two very different meanings: one is a sequence of words and the other is a more particular meaning in law, the decision as to what punishment is to be imposed. Similarly *disposition* which means nature, effort, mental attitude or property but in legal terms it means the final part of a judgement indicating the nature of a decision: acceptance of a inquiry or dismissal.

Most previous systems of automatic summarization are limited to newspaper articles and scientific articles (Saggion and Lapalme, 2002). There are important differences between news style and the legal language: statistics of words, probability of selection of textual units, position of paragraphs and sentences, words of title and lexical chains relations between words of the title and the key ideas of the text, relations between sentences and paragraphs and structures of the text.

For judgments, we show that we can identify discursive structures for the different parts of the decision and assign some argumentative roles to them. Newspapers articles often repeat the most important message but, in law, important information may appear only once. The processing of a legal document requires detailed attention and it is not straight forward to adapt the techniques developed for other types of document to the legal domain.

3 Observations from a Corpus

3.1 Composition

Our corpus contains 3500 judgments of the Federal Court of Canada, which are available in HTML on <http://www.canlii.org/ca/cas/fct/>. We analyzed manually 50 judgments in English and 15 judgments in French as well as their summaries written by professional legal abstractors. The average size of the documents that are input to our system are judgments between 500 and 4000 words long (2 to 8 pages), which form 80% of all 3500 judgments; 10% of the documents having less than 500 words (about one page) and so they do not need a summary. Only 10% of the decisions have more than 4000 words. Contrary to some existing systems (Moens et al., 1999) that focus only on limited types of judgments, such as criminal cases, our research deals with many categories of texts such as: *Access to information, Administrative law, Air law, Broadcasting, Competition, Constitutional law, Copyright, Customs and Excise - Customs Act, Environment, Evidence, Human rights, Maritime law, Official languages, Penitentiaries, Unemployment insurance and etc.*

3.2 Structure of Legal Judgments

During our corpus analysis, we compared model summaries written by humans with the texts of the original judgments. We have identified the organisational architecture of a typical judgment.

Thematic structures	Content	Judgment	Summary
DECISION DATA	Name of the jurisdiction, place of the hearing, date of the decision, identity of the author, names of parties, title of proceeding and Authority and doctrine		
INTRODUCTION	Who? did what? to whom?	5 %	12 %
CONTEXT	Facts in chronological order or by description	24 %	20 %
JURIDICAL ANALYSIS	Comments by the judge, finding of facts and application of the law	67 %	60 %
CONCLUSION	Final decision of the court	4 %	8 %

Table 2: Table of summary shows the thematic structures in a judgement and percentage of the contribution of each thematic structure in source judgment and its human made summary

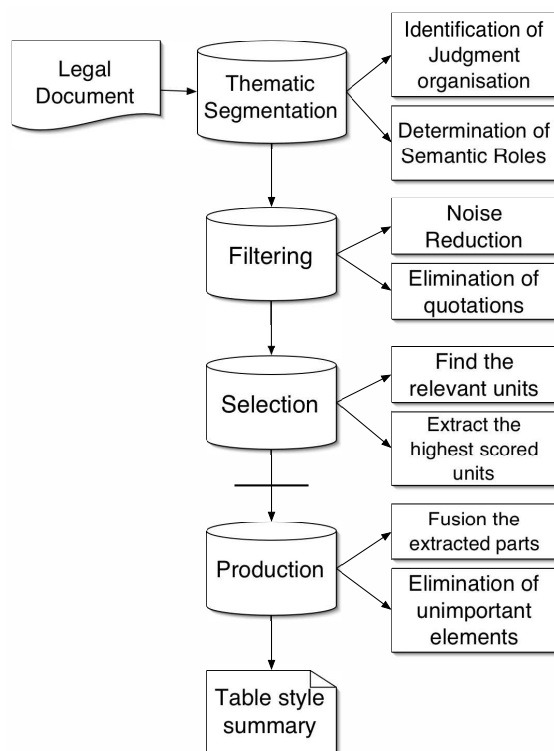


Figure 1: The procedural steps for generating of table style summary

The paragraphs that address the same subject are grouped as members of a block. We annotated the blocks with a label describing their semantic roles. We also manually annotated *citations* which are textual units (sentence or paragraph) quoted by the judge as reference, for example an article of law or other jurisprudence. The citations account for a large part of the text of the judgment, but they are not considered relevant for the summary, therefore these segments will be eliminated during the infor-

mation filtering stage.

The textual units considered as important by the professional abstractors were aligned manually with one or more elements of the source text. Table 1 shows an example of an alignment between a human summary and the original judgment. We look for a match between the information considered important in the professional abstract and the information in the source documents. Our observation shows that, for producing a summary, a professional ab-

stractor mainly relies on the manual extraction of important units while conforming to general guidelines. The collection of these selected units forms a summary.

During this analysis, we observed that texts of jurisprudence are organized according to a macrostructure and contain various levels of information, independently of the category of judgment. Proposed guidelines by Judge Mailhot of the Court of Appeal of Quebec (Mailhot, 1998) and (Branting et al., 1997) on legal judgments support this idea that it is possible to define organisational structures for decisions. Jurisprudence is organized by the discourse itself, which makes it possible to segment the texts thematically.

Textual units dealing with the same subject form a thematic segment set. In this context, we distinguish the layered thematic segments, which divide the legal decisions into different discursive structures. The identification of these structures separates the key ideas from the details of a judgment and improves readability and coherency in the summary. We will present the argumentative roles of each level of discourse, and their importance in the judgment from the point of view of the key and principal ideas. Table 2 shows the structure of a jurisprudence and its different discourse levels. Therefore, in the presentation of a final summary, we propose to preserve this organization of the structures of the text in order to build a table style summary with five themes:

DECISION DATA contains the name of the jurisdiction, the place of the hearing, the date of the decision, the identity of the author, names of parties, title of proceeding, authority and doctrine. It groups all the basic preliminary information which is needed for planning the decision.

INTRODUCTION describes the situation before the court and answers these questions: who are the parties? what did they do to whom?

CONTEXT explains the facts in chronological order, or by description. It recomposes the story from the facts and events between the parties and findings of credibility on the disputed facts.

JURIDICAL ANALYSIS describes the comments of the judge and finding of facts, and the application of the law to the facts as found. For the legal expert this section of judgment is the most important part because it gives a solution to the problem of the parties and leads the judgment to a conclusion.

CONCLUSION expresses the *disposition* which is the final part of a decision containing the information about what is decided by the court. For example, it specifies if the person is discharged or not or the cost for a party.

During our corpus analysis, we computed the distribution of the information (number of words shown in Table 2) in each level of thematic structure of the judgment. The average length of a judgment is 3500 words and 350 words for its summary i.e. a compression rate of about 10%.

4 Method for Producing Table Style Summary

Our approach for producing the summary first identifies thematic structures and argumentative roles in the document. We extract the relevant sentences and present them as a table style summary. Showing the information considered important which could help the user read and navigate easily between the summary and the source judgment. For each sentence of the summary, the user can determine the theme by looking at its rhetorical role. If a sentence seems more important for a user and more information is needed about this topic, the complete thematic segment containing the selected sentence could be presented. The summary is built in four phases (Figure 1): thematic segmentation, filtering of less important units such as citations of law articles, selection of relevant textual units and production of the summary within the size limit of the abstract.

The implementation of our approach is a system called LetSum (Legal text Summarizer), which has been developed in Java and Perl. Input to the system is a legal judgment in English. To determine the Part-of-Speech tags, the tagger described by (Hepple, 2000) is used. The semantic grammars and rules are developed in JAPE language (Java Annotations Pattern Engine) and executed by a GATE transducer (Cunningham et al., 2002).

4.1 Components of LetSum

Thematic segmentation for which we performed some experiments with two statistic segmenters: one described by Hearst for the TextTiling system (Hearst, 1994) and the C99 segmenter described by Choi (Choi, 2000), both of which apply a clustering function on a document to find classes divided by theme. But because the results of these numerical segmenters were not satisfactory enough to find the thematic structures of the legal judgments, we decided to develop a segmentation process based on the specific knowledge of the legal field.

Category of section title	Linguistic markers	Examples of section title
Begin of the judgment	<i>decision, judgment, reason, order</i>	<i>Reasons for order, Reasons for judgment and order</i>
INTRODUCTION	<i>introduction, summary</i>	<i>Introduction, Summary</i>
CONTEXT	<i>facts, background</i>	<i>The factual background, Agreed statement of facts</i>
JURIDICAL ANALYSIS	<i>analysis, decision, discussion</i>	<i>Analysis and Decision of the court</i>
CONCLUSION	<i>conclusion, disposition, cost</i>	<i>Conclusion and Costs</i>

Table 3: The linguistic markers in section titles

Each thematical segment can be associated with an argumentative role in the judgment based on the following information: the presence of significant section titles (Table 3 shows categories and features of the section titles), the absolute and relative positions of a segment, the identification of direct or narrative style (as the border of CONTEXT and JURIDICAL ANALYSIS segments), certain linguistic markers.

The linguistic markers used for each thematic segment are organized as follows:

CONTEXT introduces the parties with the verb to be (eg. *the application is company X*), describes the application request like: *advise, indicate, request* and explains the situation in the past tense and narration form.

In JURIDICAL ANALYSIS, the judge gives his explanation on the subject thus the style of expression is direct such as: *I, we, this court*, the cue phrases (Paice, 1981) like: *In reviewing the sections No. of the Act, Pursuant to section No., As I have stated, In the present case, The case at bar is.*

In CONCLUSION the classes of verbs are: *note, accept, summarise, scrutinize, think, say, satisfy, discuss, conclude, find, believe, reach, persuade, agree, indicate, review*, the concepts such as: *opinion, conclusion, summary, because, cost, action*, the cue phrases: *in the case at bar, for all the above reasons, in my view, my review of, in view of the evidence, finally, thus, consequently, in the result.* This segment contains the final result of court decision using phrases such as: *The motion is dismissed, the application must be granted.* The important verbs are: *allow, deny, dismiss, grant, refuse.*

Filtering identifies parts of the text which can be eliminated, without losing relevant information for the summary. In a judgment, the citation units (sentence or paragraph) occupy a large volume in the text, up to 30%, of the judgment, whereas their contents are less important for the summary. This is why we remove citations inside blocks of thematic segments. We thus filter two categories of segments:

submissions and arguments that report the points of view of the parties in the litigation and citations related for previous issues or references to applicable legislation. In the case of eliminating a citation of a legislation (eg. law's article), we save the reference of the citation in DECISION DATA in the field of authority and doctrine.

The identification of citations is based on two types of markers: direct and indirect. A direct marker is one of the linguistic indicators that we classified into three classes: verbs, concepts (noun, adverb, adjective) and complementary indications. Examples of verbs of citation are: *conclude, define, indicate, provide, read, reference, refer, say, state, summarize.* Examples of the concepts are: *following, section, subsection, page, paragraph, pursuant.* Complementary indications include numbers, certain preposition, relative clauses and typographic marks (colon, quotation marks).

The indirect citations are the neighboring units of a quoted phrase. For example, in Table 4 a citation is shown. For detecting CITATION segment units such as ***paragraph 78(1)***, which ***reads as follows:*** are identified using direct markers (shown here in bold) but surrounding textual units with numbers are also quotations. We thus developed a linear integration identification mechanism for sentences following a quoted sentence for determining a group of citations.

Selection builds a list of the best candidate units for each structural level of the summary. LetSum computes a score for each sentence in the judgment based on heuristic functions related to the following information: position of the paragraphs in the document, position of the paragraphs in the thematic segment, position of the sentences in the paragraph, distribution of the words in document and corpus ($tf \cdot idf$). Depending on the given information in each layered segment, we have identified some cue words and linguistic markers. The thematic segment can change the value of linguistic indicators. For example, the phrase *application is dismissed* that can

be considered as a important feature in the CONCLUSION might not have the same value in CONTEXT segment. At the end of this stage, the passages with the highest resulting scores are sorted to determine the most relevant ones.

Production of the final summary in which the selected sentences are normalized and displayed in tabular format. The final summary is about 10% of source document. The elimination of the unimportant sentences takes into account length statistics presented in Table 2. In the INTRODUCTION segment, units with the highest score are kept within 10% of the size of summary. In the CONTEXT segment, the selected units occupy 24% of the summary length. The contribution of the JURIDICAL ANALYSIS segment is 60% and the units with the role CONCLUSION occupy 6% of the summary.

4.2 Current State of LetSum

Table 4 shows an example of the output after the execution of the **Selection** module of LetSum (modules of Figure 1 up to the horizontal line) applied on a judgment of Federal Court of Canada (2468 words). **Thematic segmentation** module has divided the text into structural blocks according to the rhetorical roles (given to the left of braces in Table 4). The **Filtering** module removes citation blocks and its enumerated quoted paragraphs (e.g. paragraph (15) in tablet). **Selection** module chooses total relevant textual units (shown in bold in Table 4) in each thematic segment. The units are selected according to their argumentative role in the judgement. Here the length of all extracted units is 313 words.

Preliminary evaluations of components of LetSum are very promising; we obtained 0.90 F-measure for thematic segmentation and 0.97 F-measure for filtering stage (detection of 57 quoted segment correctly on 60).

From this information, the **Production** module (currently being implemented) could concatenate textual units with some grammatical modification to produce a short summary.

5 Related Research

LetSum is the one of the few systems developed specifically for the summarization of legal documents. All of these approaches attest the importance of the exploration of thematic structures in legal documents.

The FLEXICON project (Smith and Deedman, 1987) generates a summary of legal cases by using information retrieval based on location heuristics, occurrence frequency of index terms and the

use of indicator phrases. A term extraction module that recognizes concepts, case citations, statute citations and fact phrases leads to a document profile. This project was developed for the decision reports of Canadian courts, which are similar to our corpus.

SALOMON (Moens et al., 1999) automatically extracts informative paragraphs of text from Belgian legal cases. In this project a double methodology was used. First, the case category, the case structure and irrelevant text units are identified based on a knowledge base represented as a text grammar. Consequently, general data and legal foundations concerning the essence of the case are extracted. Secondly, the system extracts informative text units of the alleged offences and of the opinion of the court based on the selection of representative objects.

More recently, SUM (Grover et al., 2003) examined the use of rhetorical and discourse structure in level of the sentence of legal cases for finding the main verbes. The methodology is based on (Teufel and Moens, 2002) where sentences are classified according to their argumentative role.

These studies have shown the interest of summarization in a specialized domain such as legal texts but none of these systems was implemented in an environment such as CANLII which has to deal with thousands of texts and produce summaries for each.

6 Conclusion

In this paper, we have presented our approach for dealing with automatic summarization techniques. This work refers to the problem of processing of a huge volume of electronic documents in the legal field which becomes more and more difficult to access. Our method is based on the extraction of relevant units in the source judgment by identifying the discourse structures and determining the semantic roles of thematic segments in the document. The presentation of the summary is in a tabular form divided by the following thematic structures: DECISION DATA, INTRODUCTION, CONTEXT, JURIDICAL ANALYSIS and CONCLUSION. The generation of summary is done in four steps: thematic segmentation to detect the document structures, filtering to eliminate unimportant quotations and noises, selection of the candidate units and production of table style summary. The system is currently being finalized and preliminary evaluation results are very promising.

7 Acknowledgements

We would like to thanks LexUM group of legal information-processing laboratory of the Public

DECISION DATA	<p>Name of the jurisdiction: Federal Court of Canada, Place of the hearing: Ottawa Date of the decision: 31/12/97, Identity of the author: J.E. Dubé Names of parties: Commissioner of official languages of Canada, Applicant - and - Air Canada, Respondent Title of proceeding: Official languages, Docket number: T-1989-96 Authority and doctrine : Official Languages Act, R.S.C., 1985 (4th Supp.), c. 31</p>
INTRODUCTION	<p>(1) An order was made by this Court on February 4, 1997 authorizing the respondent (Air Canada) to raise preliminary objections to the notice of an originating motion filed by the applicant (the Commissioner). As a result, this motion filed by Air Canada on March 18, 1997 raises six alternative preliminary objections asking the Court to strike out in part the motion made by the Commissioner on September 6, 1996 under section 78 of the Official Languages Act.</p>
CONTEXT	<p>1. Facts (2) The Commissioner's originating motion, which was filed with the consent of the complainant Paul Comeau, concerns Air Canada's failure to provide ground services in the French language at the Halifax airport. The Commissioner asks this Court to declare that there is a significant demand for services in French in Air Canada's office at the Halifax airport and that Air Canada is failing to discharge its duties under Part IV of the Act. Part IV establishes language-related duties for communications with and services to the public, including the travelling public, where there is significant demand. (3) The Commissioner's motion is filed by the complainant Paul Comeau. ...</p>
CITATION	<p>(15) The point of departure is paragraph 78(1), which reads as follows: 78. (1) The Commissioner may (a) within the time limits prescribed by paragraph 77(2)(a) or (b), apply to the Court for a remedy under this Part in relation to a complaint investigated by the Commissioner if the Commissioner has the consent of the complainant. (b) appear before the Court on behalf of any person who has applied under section 77 for a remedy under this Part; or (c) with leave of the Court, appear as a party to any proceedings under this Part.</p>
ANALYSIS	<p>(16) Air Canada's position is therefore that the Commissioner may only apply for a remedy limited to facts relating to a specific complaint, the investigation of that complaint and the resulting reports and recommendations. In my view, this interpretation is too narrow and is inconsistent with the general objectives of the Act and its remedial and quasi-constitutional nature. ...</p>
CONCLUSION	<p>7. Conclusion (29) Thus, to ensure that the judge presiding at the hearing on the merits can correctly assess the situation in light of all the material evidence, no reference or evidence filed by the Commissioner in the three affidavits mentioned above should be struck out. (30) This motion to strike by Air Canada with respect to the preliminary objections must accordingly be dismissed.</p>

Table 4: Output produced by the LetSum's modules: **Thematic segmentation, Filtering and Selection.** Source judgment is divided into thematic blocks associated with rhetorical roles, citation block will be removed in the filtering phase and textual units (shown in bold) have been selected as relevant.

Law Research Center at the University of Montreal for their valuable suggestions. This project supported by Public Law Research Center and Natural Sciences and Engineering Research Council of Canada.

References

- L. Karl Branting, Charles B. Callaway, Bradford W. Mott, and James C. Lester. 1997. A framework for self-explaining legal documents. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-97)*, pages 72–81, University of Melbourne, Melbourne, Australia, June 30–July 3.
- Freddy Choi. 2000. Advances in domain independent linear text segmentation. In *Proceeding of the 1st North American Chapter of the Association for Computational Linguistics*, pages 26–33, Seattle, Washington.
- H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan. 2002. Gate: A framework and graphical development environment for robust nlp tools

- and applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02)*, Philadelphia, July.
- Atefeh Farzindar. 2004. D'éveloppement d'un système de résumé automatique de textes juridiques. In *TALN-RECITAL'2004*, pages 39–44, Fès, Maroc, 19-22 April.
- Claire Grover, Ben Hachey, and Chris Korycinski. 2003. Summarising legal texts: Sentential tense and argumentative roles. In Dragomir Radev and Simone Teufel, editors, *HLT-NAACL 2003 Workshop: Text Summarization (DUC03)*, pages 33–40, Edmonton, Alberta, Canada, May 31 - June 1.
- Marti A. Hearst. 1994. Multi-paragraph segmentation of expository text. In *the 32nd Meeting of the Association for Computational Linguistics*, Los Cruces, NM, June.
- Mark Hepple. 2000. Independence and commitment: Assumptions for rapid training and execution of rule-based part-of-speech taggers. In *the 38th Annual Meeting of the Association for Computational Linguistics (ACL-2000)*, pages 278–285, October.
- Louise Mailhot. 1998. *Decisions, Decisions: a handbook for judicial writing*. Editions Yvon Blais, Québec, Canada.
- Inderjeet Mani. 2001. *Automatic Text Summarization*. John Benjamins Publishing Company.
- Marie-Francine Moens, C. Uyttendaele, and J. Dumortier. 1999. Abstracting of legal cases: the potential of clustering based on the selection of representative objects. *Journal of the American Society for Information Science*, 50(2):151–161.
- Chris D. Paice. 1981. The automatic generation of literary abstracts: An approach based on identification of self-indicating phrases. In O. R. Norman, S. E. Robertson, C. J. van Rijsbergen, and P. W. Williams, editors, *Information Retrieval Research*, London: Butterworth.
- Horacio Saggion and Guy Lapalme. 2002. Generating indicative-informative summaries with sumum. *Computational Linguistics*, 28(4).
- J. C. Smith and Cal Deedman. 1987. The application of expert systems technology to case-based law. *ICAIL*, pages 84–93.
- Simone Teufel and Marc Moens. 2002. Summarising scientific articles - experiments with relevance and rhetorical status. *Computational Linguistics*, 28(4):409–445.