

LetSum, an automatic Legal Text Summarizing system

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. This paper presents our work on the development of a new methodology for automatic summarization of justice decision. We describe LetSum (Legal text Summarizer), a prototype system, which determines the thematic structure of a judgment in four themes INTRODUCTION, CONTEXT, JURIDICAL ANALYSIS and CONCLUSION. Then it identifies the relevant sentences for each theme. We discuss the evaluation of produced summaries with statistical method and also human evaluation based on jurist judgment. The results so far indicate good performance of the system when compared with other summarization technologies.

1 Introduction

Legal experts perform difficult and responsible legal clerical work which requires accuracy and speed. This task includes understanding, interpreting, explaining and researching in a wide variety of legal documents. A summary of a judgement, as a compressed but accurate restatement of its content, helps in organizing a large volume of documents and in finding the relevant judgments for their case. For this reason, the decisions are frequently manually summarized by legal experts. But the human time and expertise required to provide manual summaries for legal resources, makes human summaries very expensive. This problem shows the interest in automated text summarization to process the ever increasing amount of documents. In this paper, we describe our method of text summarization which produces short summaries for legal documents. We have attempted to reproduce the results of human expert reasoning by organizing and manipulating factual and heuristic knowledge.

We present our approach for summarizing the legal record of the proceedings of federal courts in Canada and presenting it as a table-style summary for the needs of lawyers and experts in the legal domain. The FLEXICON [10], SALOMON [8] and SUM [4] projects and the works of Borges et al. [1] attest the importance of the exploration of legal knowledge for sentence categorisation and summarisation. Our method investigates the extraction of the most important units based on the identification of the thematic structure in the document and the determination of argumentative themes of the textual units in the judgment (see Farzindar et al. [3]).

In Canada, the Canadian Legal Information Institute (CANLII) gathers legislative and judicial texts in order to make a virtual library of Canadian law accessible for free on the Internet (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.

The lawyers need to process previous legal decisions to find a solution to a legal problem not directly indicated in the law, they look for *precedents* of similar cases. Each decision contains the reasons which justify the solution for a legal problem. They constitute a *law*

jurisprudence precedent from which it is possible to extract a legal rule that can be applied to similar cases.

One reason for the difficulty of the work in legal field is the complexity of the domain: specific terminology of the legal domain and legal interpretations of expressions produce many ambiguities. For example, the word *disposition* means nature, effort, mental attitude or property in general English but in legal terms it means the final part of a judgement indicating the nature of a decision: acceptance of an inquiry or dismissal. That is why we have constructed our conceptual dictionary containing 200 concepts of legal domain. In this project we collaborate with legal experts of CanLII.

Thematic segments give information which can be used to answer specific questions about the theme of the segment such as: what are the themes of a document? How is a theme used? How are the problem and the facts presented? How does a judge reason to reach a conclusion?

In this paper, we will describe how we deal with the problem of the exploration of structure of document and content selection, according to the themes of a judgement.

1.1 Legal text summarization

Our approach to produce the summary is based on the identification of the thematic structure to find the argumentative themes of the judgement. This approach is a result of our corpus analysis in which we compared model summaries written by humans with the texts of the original judgments. The textual units considered as important by the professional abstractors were aligned manually with one or more elements of the source text. We look for a match between the information considered important in the professional abstract and the information in the source documents.

We extract the relevant sentences for each theme and present them as a table-style summary. Showing the information considered important could help the user read and navigate easily between the summary and the source judgment. 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 identification of these structures separates the key ideas from the details of a judgement and improves readability and coherency in the summary. Therefore, in the presentation of a final summary, we propose to preserve this organization of the structure of the text in order to build a table-style summary.

Our corpus contains 3500 judgments of the Federal Court of Canada, which are available in HTML on www.canlii.org/ca/cas/fct. For some of these decisions, their summaries written by professional legal abstractors are available. We analyzed manually 50 judgments in English as well as their human written summaries. These judgments were suggested by the lawyers of CanLII project, as representing the *standard* judgements with *ideal* summaries. The rest of the corpus is used for statistical computations.

2 Components of LetSum

To process an input decision to LetSum, the system first does some pre-processing. 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 (see Farzindar and Lapalme [2]).

Pre-Processing splits the input judgment into main units. First the body of the text of the decision are identified. Some keywords like *Reasons for order*, *Reasons for judgment and order* separate the basic data (date, name of court, etc.), placed in the head of document, from the beginning of the judgment. The features used for the end of the decision are the

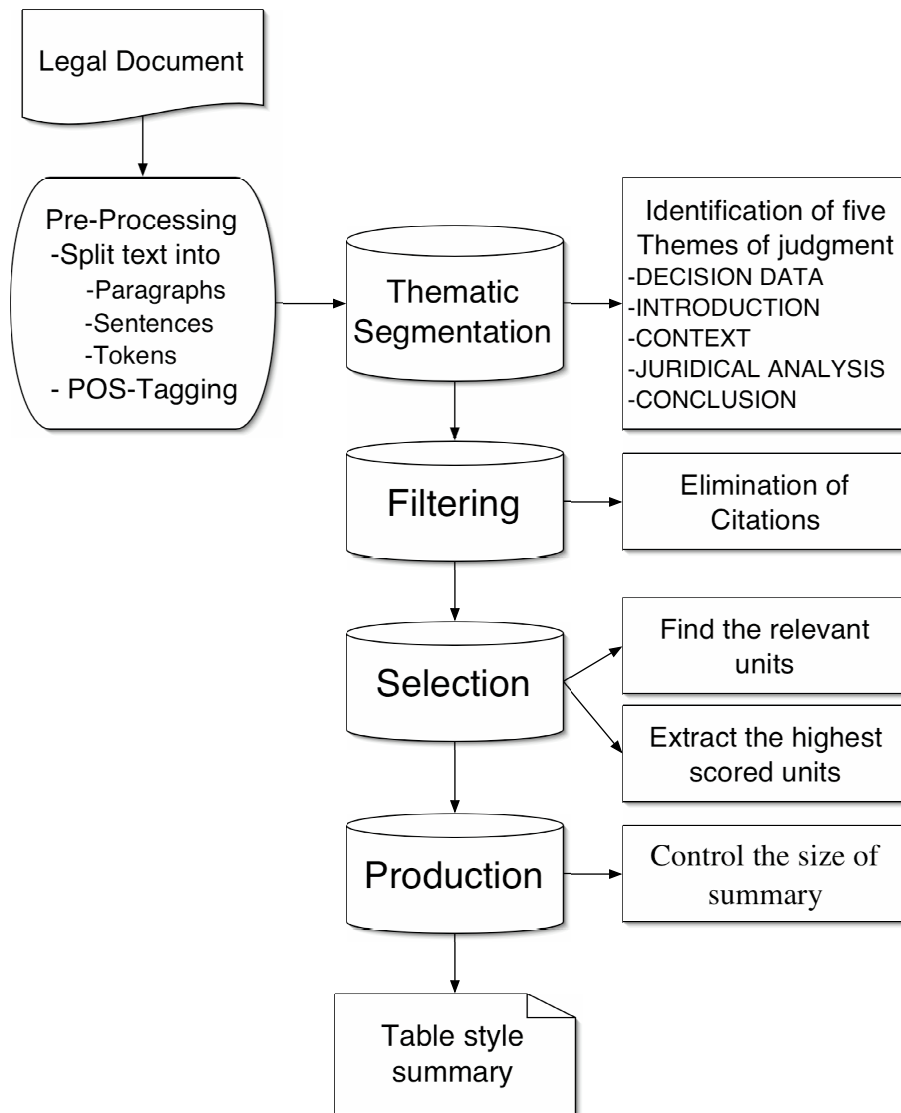


Figure 1: The procedural steps for generating of table-style summary

date and place of hearing, name and signature of the judge. Then the document is divided into: section titles, paragraphs, sentences and tokens. To determine the Part-of-Speech tags, the tagger described by Hepple [5] is used.

Thematic segmentation is based on the specific knowledge of the legal field. According to our analysis, the texts of jurisprudence have a thematic structure, independently of the category of judgment. Textual units dealing with the same subject form a thematic segment set. In this context, we distinguish four themes which divide the legal decisions into thematic segments, based on the experimental work of judge Mailhot [7]:

INTRODUCTION describes the situation before the court and answers these questions: who? did what? 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 thematic segmentation the following information are used: the presence of significant section titles, the positions of a segment, the identification of direct or narrative style (as the transition of CONTEXT and JURIDICAL ANALYSIS segments), certain linguistic markers. We present the heuristics and some examples of linguistic markers for each theme:

INTRODUCTION presents a short description of a case with cue phrases such as: *application for judicial review, application to review a decision, motion filed by, Statement of Claim.*

This segment is at the beginning of the judgement after a title like: *Reasons for order, Reasons for judgment and order.*

Significant section titles are *Introduction* and *Summary*.

CONTEXT introduces the parties (eg. *The Applicant is a 52 year old citizen of Sri Lanka*).

It describes the application request using markers such as: *advise, indicate, concern, request, etc.* (eg. *Motion concerns Air Canada's failure to provide ground services in the French language at the Halifax airport*).

It explains the situation in the past tense and narration form.

Section titles are *Facts, Background, Factual background and Agreed statement of the facts.*

JURIDICAL ANALYSIS, in which the judge gives his explanation on the subject thus the style of expression is direct using *I*.

Cue phrases are *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.*

Section titles are *Analysis* and *Decision of the court.*

CONCLUSION contains the final result of the court decision using phrases such as: *The motion is dismissed, the application must be granted.*

This segment is at the end of the judgement before the signature of the judge (judge's name, date, place of hearing, etc.).

Section titles are *Conclusion, Costs* and *Disposition.*

Filtering identifies parts of the text which can be eliminated, without losing relevant information for the summary. In a judgment, citation units (sentence or paragraph) occupy a large volume in the text, up to 30%, of the judgment, whereas their content is less important for the summary (according to our manual alignments between summaries and sources). 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 articles), 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,

Themes	Important sentences
INTRODUCTION	[1] This is an application by Her Majesty the Queen (Crown) for an order striking out the Statement of Claim or, in the alternative, an extension of time to allow the Crown to file a Statement of Defence in the present action.[7] I believe, that before I recite the facts of the present case, it is important to note that on a motion to strike a Statement of Claim due to the fact that the Statement of Claim discloses no reasonable cause of action, it must be plain and obvious that the claim will not succeed notwithstanding the fact that the allegations in the Statement of Claim must be deemed to be true.
Context	[11] The plaintiff (Riabko) was a member of the Royal Canadian Mounted Police (RCMP) from November 6, 1978 to September 14, 1994, almost 16 years. On May 6, 1994 an Adjudication Board created under sections 43 and 44 of the According to the Crown "These actions arose from certain incidents in which the plaintiff was involved in and occurred in 1992".[13] As a result of the Board's decision of May 6, 1994, Riabko was sanctioned by requesting or ordering his resignation from the RCMP Force within 14 days.[16] On April 30, 1996, Riabko filed a Statement of Claim in this action in the Federal Court of Canada.
Juridical Analysis	[35] Because of the alleged breach of the RCMP Code of Conduct, a formal disciplinary hearing took place pursuant to section 43 of the RCMP Act , that is, an Adjudication Board was appointed to conduct a hearing into the alleged complaint.[42] It is obvious that the plaintiff Riabko did not follow the procedure set out in the RCMP Act and he is now alleging that he is claiming against Her Majesty because the process wherein he was asked to resign was an abuse of power by the Board, that is, from the very start, the process of the Board was flawed and he would thus have the right to proceed in Court.[45] I am satisfied that by having resigned, she could not avail herself of the internal process as stated in the RCMP Act and could sue for damages for sexual harassment. It must be noted that before she commenced her action before the Federal Court she did not avail herself or never took part in the process set out in the "She never did anything wrong" while in the case at bar the plaintiff was found to have contravened the RCMP Code of Conduct.[47] I am satisfied that where it cannot be shown that the power with regard to the grievance process as set out in the RCMP Act has been exceeded or abused, then there would be no cause of action.[49] I am satisfied there would be no purpose for Parliament to set out a grievance procedure by statute if a party could, after taking part in the procedure, decide to circumvent the statutory procedure.
Conclusion	[50] As well, after a plain reading of the Statement of Claim, and particularly paragraphs 5 and 6, I am satisfied that there is no allegation that the Adjudication Board of the RCMP abused or exceeded its jurisdiction.[51] Plaintiff's claim is struck with costs.

Table 1: A table-style summary produced by LetSum, the original judgment has 3500 words and the summary is 15% of the source

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 prepositions, relative clauses and typographic marks (colon, quotation marks). The indirect citations are the neighboring units of a quoted phrase. For example, citation segment in the phrase *paragraph 78(1), which reads as follows*: is identified using direct markers but it points to the textual units with no direct marker which are also quotations. We thus identify the enumerated 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 an 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.

System ID	ROUGE-1	ROUGE-2	ROUGE-3	ROUGE-4	ROUGE-L
LetSum	0.57500	0.31381	0.20708	0.15036	0.45185
Baseline	0.47244	0.27569	0.19391	0.14472	0.34683
Mead	0.45581	0.22314	0.14241	0.10064	0.32089
Word	0.44473	0.21295	0.13747	0.09727	0.29652
Pertinence Mining	0.32833	0.15127	0.09798	0.07151	0.22375

Table 2: Result of statistical evaluation with ROUGE, LetSum is the first with the best evaluation scores

Production of the final summary controls the size of the summary and displays the selected sentences in tabular format. The final summary is about 10% of the source document. The elimination of the unimportant sentences takes into account length statistics based on our observation from human abstracts. In the INTRODUCTION segment, units with the highest score are kept within 10% of the size of the summary. In the CONTEXT segment, the selected units occupy 25% of the summary length. The contribution of the JURIDICAL ANALYSIS segment is 60% and the units with the role CONCLUSION occupy 5% of the summary.

Table 1 shows an example of a table-style summary generated by LetSum. The summary is 15% of the source judgment.

3 Evaluating LetSum

We evaluate LetSum in two steps; first the evaluation of the modules of the system and second, a global evaluation of produced summaries. The evaluations of components of LetSum are very promising; we obtained 90% correct segmentation for thematic segmentation module and 97% correct detection for filtering stage (correct detection of 57 quoted segments over 60).

The final global evaluation includes the intrinsic and extrinsic tests [11]. Intrinsic evaluations test the system in of itself and extrinsic evaluation test the system in relation to some other task.

For intrinsic evaluations, we conducted a pilot study on 10 long judgements (average 8 pages). We compare the summaries produced by machine with the reference summaries written by professional abstractors. This evaluation is Recall-based, which measures how many of the reference summary sentences the machine summary contains. For measuring recall, we used the ROUGE software [6], which determines the quality of a summary by comparing it to ideal summaries created by humans. The score of ROUGE-N is based on the number of n-grams occurring at the reference summary side. For example, ROUGE-2 computes the number of two successive words occurring between the machine summary and ideal summary. For measuring ROUGE-L, we view a summary sentence as a sequence of words. This evaluation computes the longest common subsequence of words to estimate the similarity between two summaries.

We compared LetSum with some other systems. A baseline system is a simple system which other system can be compared. Traditionally, for a newspaper article a baseline is the first paragraphs of the text. In our case, we defined a baseline with compression rate of 15% of the source document with the following algorithm:

- Choose 8% words of the beginning of the judgment. According our thematic segmentation, it takes the sentences from the themes INTRODUCTION and CONTEXT. If the last sentence is cut with this limit, complete it.
- Choose the last 4% words of the judgment with themes JURIDICAL ANALYSIS and CONCLUSION. If the first sentence is cut with this limit, complete it.

For automatic evaluation, we have compared the produced sentences by LetSum with: the baseline, the commercial automatic summaries produced by Microsoft Word and Pertinence Mining (www.pertinence.net) and a state-of-the-art summarization system Mead [9], with the human reference summaries. Table 2 shows the result of this evaluation. The higher score means better score and more performance system. LetSum is ranked first with the best evaluation scores.

The evaluation results show the interest of developing a summarization system for a specific domain because it is more and more difficult to generate a general summary without consideration of the user profile and the domain.

Our extrinsic evaluations will be based on legal expert judgement. We have defined a series of specific questions for the judgment (with the help a lawyer of CanLII), which cover the main topics of the document. If a user is able to answer the questions correctly by only reading the summary, it means the summary contains all the necessary information of the source judgment. We are currently performing our human-based evaluations.

4 Conclusion

LetSum is a one the few systems developed specifically for summarization of legal documents. This system is implemented in an environment such as CanLII which has to deal with thousand of texts and produce summaries for each. We have presented our approach based on the extraction of relevant units in the source judgment by identifying the document's structure and determining the themes of the segments in the decision. The generation of the summary is done in four steps: thematic segmentation to detect the legal document structure in four themes INTRODUCTION, CONTEXT, JURIDICAL ANALYSIS and CONCLUSION, filtering to

eliminate unimportant quotations and noises, selection of the candidate units and production of table-style summary. The presentation of the summary is in a tabular form along the themes of the judgment.

The evaluation of the system includes the intrinsic and extrinsic tests. The result of intrinsic evaluation of LetSum is very promising. We are completing the extrinsic evaluation based on legal expert judgments.

Acknowledgements

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

- [1] Filipe Borges, Raoul Borges, and Danièle Bourcier. Artificial neural networks and legal categorization. In *The 16th Annual Conference on Legal Knowledge and Information Systems (JURIX'03)*, page 187, The Netherlands, 11 and 12 December 2003.
- [2] Atefeh Farzindar and Guy Lapalme. Legal text summarization by exploration of the thematic structures and argumentative roles. In *Text Summarization Branches Out Workshop held in conjunction with ACL'2004*, pages 27–34, Barcelona, Spain, 25–26 July 2004.
- [3] Atefeh Farzindar, Guy Lapalme, and Jean-Pierre Desclés. Résumé de textes juridiques par identification de leur structure thématique. *Traitement Automatique des Langues (TAL), Numéro spécial sur: Le résumé automatique de texte : solutions et perspectives*, 45(1):26 pages, 2004.
- [4] Claire Grover, Ben Hachey, and Chris Korycinski. *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 2003.
- [5] Mark Hepple. 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 2000.
- [6] Chin Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out Workshop held in conjunction with ACL'2004*, pages 74–81, Barcelona, Spain, 25–26 July 2004.
- [7] Louise Mailhot. *Decisions, Decisions: a handbook for judicial writing*. Editions Yvon Blais, Québec, Canada, 1998.
- [8] Marie-Francine Moens, C. Uyttendaele, and J. Dumortier. *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, 1999.
- [9] Dragomir Radev, Jahna Otterbacher, Hong Qi, and Daniel Tam. *Mead reduces: Michigan at duc 2003*. In DUC03, Edmonton, Alberta, Canada, May 31 – June 1 2003. Association for Computational Linguistics.
- [10] J. C. Smith and Cal Deedman. The application of expert systems technology to case-based law. *ICAIL*, pages 84–93, 1987.
- [11] Karen Spark-Jones and Julia R. Galliers. *Evaluating Natural Language Processing Systems: An Analysis and Review*. Number 1083 in Lecture Notes in Artificial Intelligence. Springer, 1995.