Multi-document summarization using semantic discourse models

Resumen multidocumento utilizando teorías semántico-discursivas

Paula C. F. Cardoso
Universidade Federal de Lavras
Lavras, MG, Brazil
paula.cardoso@dcc.ufla.br

Thiago A. S. Pardo
Universidade de São Paulo
São Carlos, SP, Brazil
taspardo@icmc.usp.br

Resumen: El resumen automático tiene por objetivo reducir el tamaño de los textos, preservando el contenido más importante. En este trabajo, proponemos algunos métodos de resumen basados en dos teorías semántico-discursivas: Teoría de la Estructura Retórica (Rhetorical Structure Theory, RST) y Teoría de la Estructura Inter-Documento (Cross-document Structure Theory, CST). Han sido elegidas ambas teorías con el fin de abordar de un modo más relevante de un texto, los fenómenos relacionales de inter-documentos y la distribución de subtopicos en los textos. Los resultados muestran que el uso de informaciones semánticas y discursivas para la selección de contenidos mejora la capacidad informativa de los resúmenes automáticos.
Palabras clave: Resumen multidocumento, Cross-document Structure Theory, Rhetorical Structure Theory

Abstract: Automatic multi-document summarization aims at reducing the size of texts while preserving the important content. In this paper, we propose some methods for automatic summarization based on two semantic discourse models: Rhetorical Structure Theory (RST) and Cross-document Structure Theory (CST). These models are chosen in order to properly address the relevance of information, multi-document phenomena and subtopical distribution in the source texts. The results show that using semantic discourse knowledge for content selection improve the informativeness of automatic summaries.
Keywords: Multi-document Summarization, Cross-document Structure Theory, Rhetorical Structure Theory

1 Introduction

Due to the increasing amount of online available information, automatic Multi-Document Summarization (MDS) appears as a tool that may assist people in acquiring relevant information in a short time. MDS aims at producing automatic summaries from a collection of documents, possibly from different sources, on the same topic (Mani, 2001). Despite the importance of MDS, automatic summaries still have problems to be solved.

It is common to see approaches to MDS that make uniform use of the sentences in different texts. However, in a source text, some sentences are more important than others because of their position in the text or in a rhetorical structure, thus, this feature must be considered during the content selection phase. In the case of news texts, selecting sentences from the beginning of the text could form a good summary (Saggion and Poibeau, 2013). Sophisticated techniques use analysis of the discourse structure of texts for determining the most important sentences (Marcu, 1999; Da Cunha, Wanner, and Cabré, 2007; Uzêda, Pardo, and Nunes, 2010).

Another challenge is how to treat similarities and differences across texts that represent the multi-document phenomena. In order to deal with them, approaches that achieve good results use semantic relations (Radev, 2000; Zhang, Goldenshon, and Radev, 2002; Castro Jorge and Pardo, 2010; Kumar et al., 2014). However, those works have ignored the relevance of sentences in each text together with multi-document phenomena.

It is known that a set of related texts discussing a particular topic (a particular subject that we write about or discuss) usually contains information related to different subtopics (pieces of text that cover different aspects of the main topic) (Hearst, 1997; Salton et al., 1997; Henning, Umbrath, and Wetzker, 2008). For example, a set of news texts related to a natural disaster typically contains information about the type of disaster, damages, casualties and rescue efforts. Some MDS systems combine the subtopical structure and multi-document relationship (Salton et al., 1997; Harabagiu and Lacatusu, 2010; Wan, 2008) to find important information, but do not treat the salience of sentences in the corresponding texts.

We observe that there are not studies that jointly deal with (1) relevance of information, (2) multi-document phenomena and (3) subtopical distribution as humans do when writing summaries. As a result, the automatic summaries are not representative of the subtopics and less informative than they could be. In order to properly treat these criteria for MDS, we propose to model the MDS process using semantic discourse theories. To do that, we choose the theories RST (Rhetorical Structure Theory) (Mann and Thompson, 1987) and CST (Cross-document Structure Theory) (Radev, 2000) due to their importance for automatic summarization described in many works (Marcu, 1999; Da Cunha, Wanner, and Cabré, 2007; Uzêda, Pardo, and Nunes, 2010; Castro Jorge and Pardo, 2010; Zhang, Goldenshon, and Radev, 2002; Ribaldo, 2013; Kumar et al., 2014). The RST model details major aspects of the organization of a text and indicates relevant discourse units. The CST model, in turn, describes semantic connections among units of related texts. The theories' relations are domain-independent.

We present some methods for MDS, aiming at producing more informative and representative summaries from the source texts. The methods were developed over a multi-document corpus manually annotated with RST and CST. The results are satisfactory, improve the state of the art and indicate that the use of semantic discourse knowledge positively affects the production of informative extracts.

This paper is organized as follows: the next section (Section 2) reviews the two semantic discourse models and some related approaches for MDS; Section 3 describes the multi-document corpus; Section 4 defines new methods for MDS using RST and CST; Section 5 addresses evaluations and results; Section 6 concludes the paper.

2 Related work

2.1 Semantic discourse models

RST represents relations among propositions in a text (usually represented by clauses) and differentiates nuclear (i.e., important propositions) from satellite (i.e., additional information) propositions. Each sentence may be formed by one or more propositions. Relations composed of one nucleus and one satellite are named mononuclear relations. On the other hand, in multinuclear relations, two or more nuclei participate and have the same importance. The relationships are traditionally structured in a tree-like form. RST is probably the most used discourse model in computational linguistics and has influenced works in all language processing fields. Particularly for automatic summarization, it takes advantage of the fact that text segments are classified according to their importance: nuclei are more informative than satellites.

Inspired by RST, CST appears as a theory for relating text passages from different texts (multi-document organization) on the same topic. It is composed by a set of relations that detect similarities and differences among related texts. The relations are commonly identified between pairs of sentences, coming from different sources, which are related by a lexical similarity significantly higher than random. The result of annotating a group of texts is a graph, which is probably disconnected, since not all segments present relations with other segments. Researches that have used this theory in MDS take advantage of the CST relationships indicate relevant information in the sources and facilitate the processing of multi-document phenomena (Castro Jorge and Pardo, 2010; Kumar et al., 2014; Ribaldo, 2013; Zhang, Goldenshon, and Radev, 2002).

2.2 Document summarization

We briefly introduces some works that have used semantic knowledge to find relevant content in a collection of texts. Zhang, Goldenshon, and Radev (2002) replace low-salience

sentences with sentences that maximize the total number of CST relations in the summary. Afantenos et al., (2008) propose a summarization method based on pre-defined templates and ontologies. Kumar et al., (2014) take into account the generic components of a news story within a specific domain, such as *who*, *what* and *when*, to provide contextual information coverage, and use CST to identify the most important sentences.

For news texts in Brazilian Portuguese, the state of the art consists in three different summarization approaches (Castro Jorge and Pardo, 2010; Ribaldo, 2013; Castro Jorge, 2015). Castro Jorge and Pardo (2010) developed the CSTSumm system that take into account semantic relations (following CST) to produce preference-based summaries. Sentences are ranked according to the number of CST relationship they hold. Ribaldo (2013), in turn, developed the RSumm system, which segments texts into subtopics and group the subtopics using measures of similarity. After clustering, a relationship map is created where it is possible to visualize the structure of subtopics and to select the relevant content by the segmented bushy path (Salton et al., 1997). In the segmented bushy path, at least one sentence of each subtopic is selected to compose the summary. Following a statistical approach, Castro Jorge (2015) incorporated features given by RST to generative modelling approaches. The author considers that the number of times a sentence has been annotated as nucleus or satellite may indicate a pattern of summarization that humans follow. The model aims to capture these patterns, by computing the likelihood of sentences being selected to compose a summary. This method was named as MT-RST (which stands for Model of text-summary Transformation with RST).

As we can see, those works do not combine semantic discourse knowledge such as RST and CST for content selection. In this study, we argue that the combination of this two (RST and CST) semantic discourse knowledges improve the process of MDS.

3 The CSTNews corpus

The main resource used in this paper is the CSTNews corpus[1] (Cardoso et al., 2011; Car-

doso, Taboada, and Pardo, 2013), composed of 50 clusters of news articles written in Brazilian Portuguese, collected from several sections of mainstream news agencies: Politics, Sports, World, Daily News, Money, and Science. The corpus contains 140 texts altogether, amounting to 2,088 sentences and 47,240 words. On average, the corpus conveys in each cluster 2.8 texts, 41.76 sentences and 944.8 words. Besides the original texts, each cluster conveys single document manual summaries and multi-document manual and automatic summaries.

The size of the summaries corresponds to 30% of the number of words of the longest text of the cluster. All the texts in the corpus were manually annotated with subtopics, RST and CST structures in a systematic way. The corpus is used for evaluating the proposed methods for MDS, as we introduce in what follows.

4 A semantic discourse approach to MDS

In this section, we describe how RST, CST and subtopics may be arranged in new methods for content selection. The study was organized in three groups: (1) methods based solely on RST,(2) methods that combine RST and CST, and (3) methods that integrate RST, CST and subtopics. Subtopic segmentation, clustering and CST/RST annotation may be done manually or automatic; they may be independent steps from automatic summarization process. It is considered that the texts are previously segmented and clustered into similar subtopics, and annotated with CST and RST.

4.1 Methods based solely on RST

Prior work in single document summarization has developed content selection methods using properties of the RST tree, such as notions of salience and the level of units in the tree. The first group of methods we present is based on this literature, specifically on Marcu (1999), which associates a score for each node in the RST tree depending on its nuclearity and the depth of the tree where it occurs. The author put forward the idea of a promotion set, consisting of salient units of a text span. The salient units of the leaves are the leaves themselves. The salient units of each internal node is the union of the promotion sets of its nuclear children. Salient

---

[1]http://www2.icmc.usp.br/∼taspardo/sucinto/cst news.html

units that are in the promotion sets of the top nodes of a discourse tree are more important than salient units in the nodes found at the bottom. For scoring each textual unit, the method attributes to the root of the tree a score corresponding to the number of levels in the tree and, then, traverses the tree towards the unit under evaluation: each time the unit is not in the promotion set of a node during the traversing, it has the score decreased by one. Following the same idea, we proposed a method (which we refer to as RST-1) to compute a score for each sentence as the sum of its nodes' scores (propositions), given by Marcu's method. It does this for all texts of a collection and, then, a multi-document rank of sentences is organized. From the rank, the next step is to select only nuclear units of the best sentences.

As an example, consider that there are 3 sentences in part A of Figure 1: sentence 1 is formed by proposition 1; sentence 2, by 2; sentence 3, by 3 to 5. The symbols N and S indicate the nucleus and satellite of each rhetorical relation. Applying RST-1 method, the score (in bold) of sentences 1 and 2 is 4, and for sentence 3 is 6 (the sum of three propositions). As sentence 3 has the highest score, its nuclei are selected to compose the summary: just the text span in node 3. Since RST relations do not indicate if there is redundancy between nodes (sentences from different texts), we control it in the summary using the cosine measure (Salton, 1989) (i.e., we discard selected sentences that are too similar with previously selected sentences already in the summary).


Figure 1: Example of rhetorical structure

Because all these scores depend on the length of the text (Louis, Joshi, and Nenkova, 2010) and on the number of propositions in a sentence, a rank based on the sum of propositions' scores may insert discrepancies in the method and does not mirror the important sentences in a set of documents. In addition, as we work on news texts, it is expected that

first sentences are more relevant, differently from Figure 1 (part A), where the last sentence was more important than the former. As a solution, we proposed to compute the score for sentences, not for propositions, and to normalize each score by the height of the tree, resulting in a number ranged from 0 to 1. In Figure 1 (part B), each node represents a sentence; the bold numbers are sentences' scores before normalization. From this new sentence rank, we create two possibilities of content selection: only nuclear units (propositions) of sentences (we refer to as RST-2) or full sentences (RST-3).

4.2 Methods that combine RST and CST

We present two methods that combine RST and CST. We assume that the relevance of a sentence is influenced by its salience, given by RST, and its correlation with multi-document phenomena, indicated by CST. In this way, there are several different ways to combine the knowledge levels to content selection. As some authors write (Zhang, Goldenshon, and Radev, 2002; Castro Jorge and Pardo, 2010; Kumar et al., 2014), the more repeated and elaborated sentences between sources are, more relevant they are, and likely contain more CST relations. If we find the relevant sentences in a set of related documents, we may use RST to eliminate their satellites and make room for more information. In the following methods, redundancy is controlled by means of CST relationships. For example, if there is an IDENTITY relation (when the same content appears in more than one location) between two sentences, only one must be selected to the summary (usually, the shorter one).

Based on that, we propose an enhanced version of CSTSumm system (Castro Jorge and Pardo, 2010) with RST, which we refer to as RC-1. In RC-1 method, we rank the sentences according to the number of CST relationships one sentence has. The more relevant a sentence is, the higher in the rank it is. The best sentence is selected and, if it has satellites, they are removed. Two more variations for RC-1 that did not produce satisfactory results were tested, thus they are not described here (Cardoso, 2014).

The second method (we refer to as RC-4) is a combination of the number of CST relationships and RST-3 method (where the RST

score of a sentence is normalized by its tree's height), constituting a score that represents the salience of the sentence and its relevance for a collection. In other words, RST and CST scores are added to form the final score of a sentence. In contrast to RC-1, RC-4 selects sentences.

To illustrate RC-1 and RC-4 methods, consider Figure 2, where there are two discourse trees representing two texts (D1 and D2); D1 is upside down for better visualization; each node is a sentence with its RST score normalized in bold; dashed lines between texts are CST relationships. When we apply RST-3 method to the tree of document D1, which has height 3, we obtain the scores 3, 1, 2 and 2, for sentences 1, 2, 3 and 4, respectively. After normalizing by the depth of the tree, we obtain the scores 1, 0.3, 0.6 and 0.6.

By applying RC-1, the rank sentence is D1_1 > {D2_1, D2_3} > {D1_2, D1_3, D2_2} > D1_4, where DX_Y indicates the sentence Y in the document X. Sentences inside brackets have the same score/importance. Using RC-4 method, the rank is organized as follows: D1_1 > D2_1 > D2_3 > D1_3 > {D1_2, D2_2} > D1_4.


Figure 2: Example of RST and CST relationships for two texts

4.3 Methods that integrate RST, CST and subtopics

This group of methods combines RST, CST and subtopics and is based on lessons learned from the previous methods. Texts are segmented in subtopics (Cardoso, Taboada, and Pardo, 2013) and similar subtopics are clustered (Ribaldo, Cardoso, and Pardo, 2013). We assume that a subtopic discussed in several documents is more significant than one that was discussed in only one (Ercan and Cicekli, 2008; Chen et al., 2013), thus, sentences of repeated subtopics are relevant. With that in mind, to give preference to those

subtopics during content selection, the sentences receive an additional score.

We propose a method (we refer to as RCT-1) that considers that importance of a sentence as the sum of its number of CST relations, RST score (similar to RST-3 method without normalization) and the relevance of subtopic to which it belongs. From the sentence rank, important content is selected without satellite propositions. Also using the same rank, it was created the second variation, called RCT-2, which selects sentences.

Two other variations are the RCT-3 and RCT-4 methods. For these methods, the final score for each sentence is similar to the first two, with the difference that the RST score is normalized by the size (height) of its discourse tree, as in RST-3 and RC-4. RCT-1 and RCT-3 only select nuclear propositions of the best sentences, while RCT-2 and RCT-4 pick out sentences.

Figure 3 illustrates RCT-4. As we can see, there are three subtopics (separated by vertical lines) in the 2 source texts, which are identified by T1, T2 and T3. As only subtopic T1 is repeated in the sources, sentences belonging to it are preferred to compose the summary. By applying RCT-4, the sentence rank is: D1_1 > D2_1 > D1_2 > D2_4 > D2_2 > D2_3 > {D1_3, D1_4}.


Figure 3: Example of RST, CST and subtopics relationship between texts

5 Evaluation and discussion

We describe the results using ROUGE (Lin, 2004), a set of standard evaluation metrics used in text summarization, which produces scores that often correlate quite well with human judgments for ranking summarization systems. It automates the comparison between model and system summaries based on n-gram overlap. This benefit has made ROUGE immensely popular. The results are given in terms of Recall (R), Precision (P) and F-measure (F). Our methods are compared to CSTSumm (Castro Jorge and

Pardo, 2010), RSumm (Ribaldo, 2013) and MT-RST (Castro Jorge, 2015), which have used the same corpus as here and represent the state of the art in the area.

Among the RST group, the results in Table 1 (ordered by F-measure) show that sentence selection is better than only proposition selection: RST-3 has the best ROUGE evaluation (for unigrams comparison, since it is already enough for distinguishing systems). This is a particularly interesting result, because the decision to keep sentences was due to an attempt to soften the language quality problems observed empirically in the summaries of the RST-1 and RST-2. It is also possible to wonder that maybe RST is too refined for MDS needs, with a coarser discourse structure being more suitable for this task. We believe that RST may be used for improve abstractive summarization approaches.

	Methods	R	P	F
1	RC-4	0.4374	0.4511	0.4419
2	RC-1	0.4270	0.4557	0.4391
3	RCT-4	0.4279	0.4454	0.4346
4	RCT-3	0.4151	0.4446	0.4274
5	RCT-2	0.4199	0.4399	0.4269
6	RSumm	0.3517	0.5472	0.4190
7	RCT-1	0.3987	0.4313	0.4128
8	CSTSumm	0.3557	0.4472	0.3864
9	RST-3	0.3874	0.3728	0.3781
10	RST-2	0.3579	0.3809	0.3671
11	MT-RST	0.3453	0.3534	0.3482
12	RST-1	0.3198	0.3238	0.3206

Table 1: ROUGE evaluation

In the RC group, RC-4 is slightly better in F-measure compared to RC-1. As for RST-3, the result of RC-4 enhances that selecting sentences instead of propositions produces more informative summaries. RC-4 was also better than all other methods for recall and F-measure; it means that the relevance of sentences within their correspondent source texts leads to the production of summaries with content closer to human summary content.

In the evaluation of methods that combine three knowledge types (RST, CST and subtopics), RCT-4 had better performance. However, RC-4 is slightly better than RCT-4. Several factors may contribute to this: (1) the segmentation and clustering of subtopics may not be as good as expected; (2) the way to deal with relevant subtopics may not be

the most appropriate one (since there are several possible ways to merge the models); or (3) it may not be advantageous to invest in subtopics. Besides that, summaries produced using subtopics are similar to the ones based only on RST and CST.

One interesting point is that all methods of RC and RCT groups were better than those that used the models in isolation (RST group and CSTSumm) in terms of recall and F-measure. With the exception of RCT-1, those methods also outperform RSumm in terms of F-measure. This shows that the combination of semantic discourse knowledge positively affects the production of summaries. It is also interesting to see that most of the methods were better than the statistical approach of the MT-RST method.

Considering only F-measure, the three methods with better performance are: RC-4, RC-1 and RCT-4, in this order. However, summaries produced by the RC-1 method present eventual low linguistic quality due to cutting satellites, difficulting its comprehension.

We have run t-tests for the pair of methods for which we wanted to check the statistical difference. The F-measure difference is not significant when comparing RC-4 and RCT-4 with RSumm (with 95% confidence), but is for CSTSumm and MT-RST. When comparing RC-4 to RCT-4, there is not statistical difference.

As illustration, Figure 4 shows an automatic summary (translated from the original language - Portuguese) produced by RC-4 method. The source texts contain news about the facts related to the floods that hit North Korea. It may be noticed that RC-4 introduces sentences that are related to the central facts of the topic that is being narrated. This example reveals the power of RST to capture the main or most salient information from a topic.

6 Conclusions

We have introduced some new methods for MDS that combine different knowledge: RST, CST and subtopics. To the best of our knowledge, this is the first time that RST and CST are integrated for MDS. From their isolated study, we observe that those models may enhance the MDS process if they are used together.

The hypothesis that RST contributes to

> [S1]At least 549 people were killed and 295 are still missing as a result of floods that hit North Korea in July, according to a pro-Pyongyang Japanese newspaper.
>
> [S2]According to the newspaper Choson Sinbo, published by the Association of Korean Residents in Japan (which is close to the communist regime in North Korea), the heavy rains that flooded much of this country in the second half of July caused much damage.
>
> [S3]North Korea has refused offers from international agencies to launch campaigns to help the country, but a local officer said last week that Pyongyang would accept aid from South Korea if it was given without conditions.

Figure 4: A summary produced by RC-4

indicate relevant units for MDS is confirmed. The results are more informative summaries than previous approaches. Despite the intervention of the RST, with the CST, which is one of the most theories employed in MDS, it was possible to treat multi-document phenomena, identifying redundant, contradictory and complementary information. The information on subtopics and how to use it needs more investigation; summaries produced using subtopics are similar to the ones based only on RST and CST. We compared the performance of our methods with the state of the art for MDS, and the results indicate that the use of semantic discourse knowledge positively affects the production of informative summaries.

As a future work, we plan to evaluate the linguistic quality of the automatic summaries.

7   Acknowledgments

References

Afantenos, S.D., V. Karkaletsis, P. Stamatopoulos, and C. Halatsis. 2008. Using synchronic and diachronic relations for summarizing multiple documents describing evolving events. *Journal of Intelligent Information Systems*, 30(3):183–226.

Cardoso, P .C. F., M. Taboada, and T. A. S. Pardo. 2013. Subtopics annotation in a corpus of news texts: steps towards automatic subtopic segmentation. In *Proceedings of the Brazilian Symposium in Information and Human Language Technology.*

Cardoso, P.C.F. 2014. *Exploração de metódos de sumarização automática multidocumento com base em conhecimento semântico-discursivo.* Ph.D. thesis, Universidade de São Paulo.

Cardoso, P.C.F., E.G. Maziero, M.L.R. Castro Jorge, E.M.R. Seno, A. Di Felippo, L.H.M. Rino, M.G.V. Nunes, and T.A.S. Pardo. 2011. CSTNews - A discourse-annotated corpus for single and multi-document summarization of news texts in Brazilian Portuguese. In *Proceedings of the 3rd RST Brazilian Meeting*, pages 88–105.

Cardoso, P.C.F., M. Taboada, and T.A.S. Pardo. 2013. On the contribution of discourse structure to topic segmentation. *Proceedings of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 92–96.

Castro Jorge, M.L.R. 2015. *Modelagem gerativa para sumarização automática multidocumento.* Ph.D. thesis, Universidade de São Paulo.

Castro Jorge, M.L.R. and T.A.S. Pardo. 2010. Experiments with CST-based multidocument summarization. In *Proceedings of the 2010 Workshop on Graph-based Methods for Natural Language Processing*, pages 74–82.

Chen, F., Y. Liu, M. Zhang, S. Ma, and L. Chen. 2013. A subtopic taxonomy-aware framework for diversity evaluation. *Proceedings of EVIA*, pages 9–16.

Da Cunha, Iria, L. Wanner, and T. Cabré. 2007. Summarization of specialized discourse: The case of medical articles in Spanish. *Terminology*, 13(2):249–286.

Ercan, G. and I. Cicekli. 2008. Lexical cohesion based topic modeling for summarization. In *Computational Linguistics and Intelligent Text Processing*. Springer, pages 582–592.

Harabagiu, S. and F. Lacatusu. 2010. Using topic themes for multi-document summarization. *ACM Transactions on Information Systems (TOIS)*, 28(3):13.

Hearst, M.A. 1997. Texttiling: Segmenting text into multi-paragraph subtopic

passages. *Computational linguistics*, 23(1):33–64.

Henning, L., W. Umbrath, and R. Wetzker. 2008. An ontology-based approach to text summarization. In *Proceedings of the International Conference on Web Intelligence and Intelligent Agent Technology - IEEE/WIC/ACM*, pages 291–294.

Kumar, Y.J., N. Salim, A. Abuobieda, and A.T. Albaham. 2014. Multi document summarization based on news components using fuzzy cross-document relations. *Applied Soft Computing*, 21:265–279.

Lin, C-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81.

Louis, A., A. Joshi, and A. Nenkova. 2010. Discourse indicators for content selection in summarization. In *Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 147–156. Association for Computational Linguistics.

Mani, I. 2001. *Automatic summarization.* John Benjamins Publishing.

Mann, W.C. and S.A. Thompson. 1987. Rhetorical Structure Theory: A theory of text organization. In *University of Southern California, Information Sciences Institute*, number ISI/RS-87-190.

Marcu, D. 1999. Discourse trees are good indicators of importance in text. *Advances in automatic text summarization*, pages 123–136.

Radev, D.R. 2000. A common theory of information fusion from multiple text sources step one: cross-document structure. In *Proceedings of the 1st SIGdial workshop on Discourse and dialogue-Volume 10*, pages 74–83. Association for Computational Linguistics.

Ribaldo, R. 2013. *Investigação de mapas de relacionamento para Sumarização Multidocumento*. Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo.

Ribaldo, R., P.C.F. Cardoso, and T.A.S. Pardo. 2013. Investigação de métodos de segmentação e agrupamento de subtópicos para sumarização multidocumento. In *Anais do 3o Workshop de Iniciação Científica em Tecnologia da Informação e da Linguagem Humana*, pages 25–27.

Saggion, H. and T. Poibeau. 2013. *Automatic text summarization: Past, present and future.* Springer.

Salton, G. 1989. *Automatic Text Processing: The Transformation, Analysis, and Retrieval of Information by Computer.* Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.

Salton, G., A. Singhal, M. Mitra, and C. Buckley. 1997. Automatic Text Structuring and Summarization. *Information Processing & Management*, 33(2):193–207.

Uzêda, V.R., T.A.S. Pardo, and M.G.V. Nunes. 2010. A comprehensive comparative evaluation of RST-based summarization methods. *ACM Transactions on Speech and Language Processing (TSLP)*, 6(4):20.

Wan, X. 2008. An exploration of document impact on graph-based multi-document summarization. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 755–762. Association for Computational Linguistics.

Zhang, Z., S.B. Goldenshon, and D.R. Radev. 2002. Towards CST-Enhanced Summarization. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pages 439–446.