

A Multilingual Study of Compressive Cross-Language Text Summarization

Elvys Linhares Pontes^{1,2,3}, Stéphane Huet¹, and Juan-Manuel Torres-Moreno^{1,2,3}

¹ LIA, Université d'Avignon et des Pays de Vaucluse, Avignon, 84000 France
`elvys.linhares-pontes@alumni.univ-avignon.fr`

² Département de GIGL, École Polytechnique de Montréal,
C.P. 6079, succ. Centre-ville, Montréal (Québec) H3C 3A7 Canada

³ Laboratoire GDAC, Université du Québec à Montréal,
C.P. 8888, succ. Centre-ville Montréal (Québec) H3C 3P8 Canada

Abstract. Cross-Language Text Summarization generates summaries in a language different from the language of the source documents. Recent methods use information from both languages to generate summaries with the most informative sentences. However, these methods have performance that can vary according to languages, which can reduce the quality of summaries. In this paper, we propose a compressive framework to generate cross-language summaries. In order to analyze performance and especially stability, we tested our system and extractive baselines on a dataset available in four languages (English, French, Portuguese, and Spanish) to generate English and French summaries. An automatic evaluation showed that our method outperformed extractive state-of-art CLTS methods with better and more stable ROUGE scores for all languages.

1 Introduction

Cross-Language Text Summarization (CLTS) aims to generate a summary of a document, where the summary language differs from the document language. Many of the state-of-the-art methods for CLTS are of the extractive class. They mainly differ on how they compute sentence similarities and alleviate the risk that translation errors are introduced in the produced summary.

Previous works analyze the CLTS only between two languages for a given dataset, which does not demonstrate the stability of methods for different texts and languages. Recent works carried out compressive approaches based on neural networks, phrase segmentation, and graph theory [1–3]. Among these models, Linhares Pontes *et al.* introduced the use of chunks and two compression methods at the sentence and multi-sentence levels to improve the informativeness of cross-language summaries [3].

In this paper, we adapt the method presented in [3] to perform CLTS for several languages. More precisely, we modified the creation of chunks and we

simplified their multi-sentence compression method to be able to analyze several languages and compress small clusters of similar sentences. To demonstrate the stability of our system, we extend the MultiLing Pilot dataset [4] with two Romance languages (Portuguese and Spanish) to test our system to generate {French, Portuguese, Spanish}-to-English and {English, Portuguese, Spanish}-to-French cross-language summaries. Finally, we carried out an automatic evaluation to make a systematic performance analysis of systems, which details the characteristics of each language and their impacts on the cross-language summaries.

The remainder of this paper is organized as follows. Section 2 details our contributions to the compressive CLTS approach. In Section 3, we describe the most recent works about CLTS. Section 4 reports the results achieved on the extended version of the MultiLing 2011 dataset and the analysis of cross-language summaries. Finally, conclusions and future work are set out in Section 5.

2 Compressive Cross-Language Text Summarization

Following the approach proposed by Linhares Pontes *et al.* [3], we combined the analysis of documents in the source and the target languages, with Multi-Sentence Compression (MSC) to generate more informative summaries. We expanded this approach in three ways.

In order to simplify and to extend the analysis to several languages, we only use Multi-Word Expressions for the English target language and we replace the analysis of parallel phrases with syntactic patterns to create chunks. Then, we also optimized the MSC method for small clusters by removing the analysis of 3-grams. Unfortunately, we have not found any available dataset for sentence compression in other languages; therefore, we restrict the use of compressive methods to MSC. Finally, differently from [3], compressed versions of sentences were considered in the CoRank method instead of the only original versions, in order to estimate the relevance of sentences for summaries.

The following subsections highlight our contributions to the architecture of the method presented in [3].

2.1 Preprocessing

Initially, source texts are translated into English and French with the Google Translate system⁴, which was used in the majority of the state-of-the-art CLTS methods [5, 2, 3].

Then, a chunk-level tokenization is performed on the target language side [6, 7]. We applied two simple syntactic patterns to identify useful structures: $\langle (ADJ)^*(NP|NC)^+ \rangle$ for English and $\langle (ADJ)^*(NP|NC)^+(ADJ)^* \rangle$ for French, where ADJ stands for adjective, NP for proper noun and NC for common noun. We also use the Stanford CoreNLP tool [8] for the English translations.

⁴ <https://translate.google.com>

This tool detects phrasal verbs, proper names, idioms and so on. Unfortunately, we did not find a similar tool for French; consequently, the French chunk-level tokenization is limited to the syntactic pattern.

2.2 Multi-Sentence Compression

We aim to generate a single, short, and informative compression from clusters of similar sentences. Therefore, we use the Linhares Pontes *et al.*'s method [3] to create clusters of similar sentences based on their similarity in the source and the target languages.

As the majority of clusters are composed of few similar sentences (normally two or three sentences), 3-grams are not frequent and the associated score is of little interest for the compress process. Therefore, we simplify the Linhares Pontes *et al.*'s method [3, 9] to process MSC guided only by the cohesion of words and keywords.

Our MSC method looks for a sentence that has a good cohesion and the maximum of keywords, inside a word graph built for each cluster of similar sentences according to the method devised by Filippova [10]. In this graph, arcs between two vertices representing words (or chunks) are weighted by a cohesion score that is defined by the frequency of these words inside the cluster. Vertices are labeled depending on whether they are or not a keyword identified by the Latent Dirichlet Allocation (LDA) method inside the cluster (see [3] for more details). From these scores and labels, the MSC problem is expressed as the following objective:

$$\text{minimize} \left(\sum_{(i,j) \in A} w(i,j) \cdot x_{i,j} - k \cdot \sum_{l \in L} b_l \right) \quad (1)$$

where x_{ij} indicates the existence of the arc (i, j) in the solution, $w(i, j)$ is the cohesion of the words i and j , L is the set of labels (each representing a keyword), b_l indicates the existence of a chunk with a keyword l in the solution, k is the keyword bonus of the graph⁵. Finally, we generate the 50 best solutions according to the objective (1) and we select the compression with the lowest normalized score (Equation 2) as the best compression:

$$\text{score}(c) = \frac{e^{\text{opt}(c)}}{\|c\|}, \quad (2)$$

where $\text{opt}(c)$ is the score of the compression c from Equation 1. We restrict the MSC method to the sentences in the target language in order to avoid errors generated by machine translation, which would be applied in a post-processing step on compressed sentences.

⁵ The keyword bonus is defined by the geometric average of all weight arcs in the graph and aims at favoring compressions with several keywords.

2.3 CoRank Method

Sentences are scored based on their information in both languages using the CoRank method [5] which analyzes sentences in each language separately, but also between languages (Equations 3-7).

$$\mathbf{u} = \alpha \cdot (\tilde{\mathbf{M}}^{\text{sc}})^T \mathbf{u} + \beta \cdot (\tilde{\mathbf{M}}^{\text{tg,sc}})^T \mathbf{v} \quad (3)$$

$$\mathbf{v} = \alpha \cdot (\tilde{\mathbf{M}}^{\text{tg}})^T \mathbf{v} + \beta \cdot (\tilde{\mathbf{M}}^{\text{tg,sc}})^T \mathbf{u} \quad (4)$$

$$M_{ij}^{\text{tg}} = \begin{cases} \text{cosine}(s_i^{\text{tg}}, s_j^{\text{tg}}), & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

$$M_{ij}^{\text{sc}} = \begin{cases} \text{cosine}(s_i^{\text{sc}}, s_j^{\text{sc}}), & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

$$M_{ij}^{\text{tg,sc}} = \sqrt{\text{cosine}(s_i^{\text{sc}}, s_j^{\text{sc}}) \times \text{cosine}(s_i^{\text{tg}}, s_j^{\text{tg}})} \quad (7)$$

where M^{tg} and M^{sc} are normalized to \tilde{M}^{tg} and \tilde{M}^{sc} , respectively, to make the sum of each row equal to 1. \mathbf{u} and \mathbf{v} denote the relevance of the source and target language sentences, respectively. α and β specify the relative contributions to the final scores from the information in the source and the target languages, with $\alpha + \beta = 1$.

Finally, summaries are generated with the most relevant sentences in the target language. We add a sentence/compression to the summary only if it is sufficiently different from the sentences/compressions already in the summary.

3 Related Work

Cross-Language Text Summarization schemes can be divided in early and late translations, and joint analysis. The early translation first translates documents to the target language, then it summarizes these translated documents using only information of these translations. The late translation scheme does the reverse. The joint analysis combines the information from both languages to extract the most relevant information.

Regarding the analysis of machine translation quality, Wan *et al.* [11] and Boudin *et al.* [12] used sentence features (sentence length, number of punctuation marks, number of phrases in the parse tree) to estimate the translation quality of a sentence. Wan *et al.* used an annotated dataset made of pairs of English-Chinese sentences with translation quality scores to train their Support Vector Machine (SVM) regression method. Finally, sentences that have a high translation quality and a high informativeness were selected for the summaries. Similarly, Boudin *et al.* trained an *e*-SVR using a dataset composed of English and automatic French translation sentences to calculate the translation quality based on the NIST metrics. Then, they used the PageRank algorithm to estimate the relevance of sentences based on their similarities and translation

quality. Yao *et al.* devised a phrase-based model to jointly carry out sentence scoring and sentence compression [1]. They developed a scoring scheme for the CLTS task based on a submodular term of compressed sentences and a bounded distortion penalty term.

Wan [5] leverages the information in the source and in the target language for cross-language summarization. He proposed two graph-based summarization methods (SimFusion and CoRank) for the English-to-Chinese CLTS task. The first method linearly fuses the English-side and Chinese-side similarities for measuring Chinese sentence similarity. In a nutshell, this method adapts the PageRank algorithm to calculate the relevance of sentences, where the weight arcs are obtained by the linear combination of the cosine similarity of pairs of sentences for each language. The CoRank method was described in Section 2.3.

Recently, Wan *et al.* [2] carried out the cross-language document summarization task by extraction and ranking of multiple summaries in the target language. They analyzed many summaries in order to produce high-quality summaries for every kind of documents. Their method uses a top-K ensemble ranking for candidate summary, based on features that characterize the quality of a candidate summary. They used multiple text summarization and machine translation methods to generate the summaries.

In order to generate abstractive cross-lingual summaries, Zhang *et al.* [13] extended the work of Bing *et al.* [14] that constructs new sentences by exploring noun/verb phrases. Their method first constructs a pool of concepts and facts represented by phrases in English and Chinese translation sentences. Then new sentences are generated by selecting and merging informative phrases in both languages to maximize the salience of phrases and meanwhile satisfy the sentence construction constraints. They employ integer linear optimization for conducting phrase selection and merging simultaneously in order to generate informative cross-lingual summaries with a good translation quality. This method generates abstractive summaries; however, the framework to identify concepts and facts only works for English, which prevents this method from being extended for other languages.

4 Experimental Evaluation

We estimate the performance of our approach in relation to the early and the late translations, SimFusion and CoRank methods⁶⁷. All systems generate summaries containing a maximum of 250 words with the best scored sentences without redundant sentences. We regard a similarity score (cosine similarity) with a threshold θ of 0.6 to create clusters of similar sentences for the MSC⁸ and a threshold θ of 0.5 to remove redundant sentences in the summary generation.

⁶ We used the same configuration for SimFusion and CoRank as described in [3].

⁷ Unfortunately, the majority of state-of-the-art systems in CLTS are not available. Therefore, we only considered extractive systems in our analysis.

⁸ We use the same threshold θ of 0.5 described in [3] for French-to-English cross-language summaries.

4.1 Datasets

We used the English and French language versions of the MultiLing Pilot 2011 dataset [4]. This dataset contains 10 topics which have 10 source texts and 3 reference summaries per topic. These summaries are composed of 250 words. In order to extend the analysis to other languages, English source texts were translated into the Portuguese and Spanish languages by native speakers⁹. Specifically, we use English, French, Portuguese, and Spanish texts to test our system.

4.2 Evaluation

An automatic evaluation with ROUGE [15] was carried out to compare the differences between the distribution of n -grams of the candidate summary and a set of reference summaries. More specifically, we used unigram (ROUGE-1 or R-1), bigram (ROUGE or R-2), and skip-gram (ROUGE-SU4 or R-SU4) analyses.

Table 1 describes the ROUGE f-scores obtained by each system to generate French summaries from English, Portuguese, and Spanish source texts. Despite using the information from both languages, the SimFusion method achieved comparable results with respect to the early and late approaches. On the contrary, CoRank and our approach consistently obtained better results than other baselines, with at least an absolute difference of 0.035 in ROUGE-1 for all languages. The MSC method improved the CoRank method by generating more informative compressions for all languages. The last two lines show that chunks helped our MSC method to generate slightly more informative summaries (better ROUGE-1 scores).

Table 1. ROUGE f-scores for cross-language summaries from English, Portuguese, and Spanish languages to French language.

Methods	English			Portuguese			Spanish		
	R-1	R-2	R-SU4	R-1	R-2	R-SU4	R-1	R-2	R-SU4
Late	0.4190	0.0965	0.1588	0.4403	0.1128	0.1746	0.4371	0.1133	0.1738
Early	0.4223	0.1007	0.1631	0.4386	0.1110	0.1743	0.4363	0.1143	0.1729
SimFusion	0.4240	0.1004	0.1637	0.4368	0.1105	0.1735	0.4350	0.1125	0.1723
CoRank	0.4733	0.1379	0.1963	0.4723	0.1460	0.2006	0.4713	0.1387	0.1942
Our approach	0.4831	0.1460	0.2030	0.4784	0.1511	0.2045	0.4825	0.1481	0.2050
Our approach w/o chunks	0.4817	0.1463	0.2021	0.4784	0.1518	0.2044	0.4805	0.1486	0.2056

The Multiling dataset is composed of 10 topics in several languages; however, these topics are expressed in different ways for each language. These dissimilarities implies a variety of vocabulary sizes and sentence lengths, and, consequently,

⁹ The extension of the MultiLing Pilot 2011 dataset is available at: <http://dev.termwatch.es/~fresa/CORPUS/TS/>

of outputs of the MT system from each source language (Table 2). The biggest difference in the statistics is between English source texts and its French translation vocabulary. French translations significantly increased the vocabulary from English source texts and the number of words. These translations also are longer than source texts, except for the Spanish that has similar characteristics. Our simple syntactic pattern created similar numbers of chunks for all languages with the same average length. The addition of these simple chunks did not significantly improve the informativeness of our compressions.

Table 2. Statistics of datasets and their translation to French.

	English		Portuguese		Spanish	
	Source	Fr-Translation	Source	Fr-Translation	Source	Fr-Translation
#words	36,109	39,960	37,339	39,302	40,440	40,269
#vocabulary	8,077	8,770	8,694	8,572	8,808	8,744
#sentences	1,816	1,816	2,002	2,002	1,787	1,787
sentence length	19.9	22.0	18.6	19.6	22.6	22.5
#chunks	–	1,615	–	1,579	–	1,606
Average length of chunks	–	2.1	–	2.1	–	2.1

These differences also act on the clustering process and the MSC method. Table 3 details the number and the average size of clusters with at least two French sentences translated from each source language. French translations from Portuguese produced the shortest compressions (18.6 words) while compressions from Spanish had the highest compression ratio. With respect to other languages, the similarity of the sentences translated from English is lower, which leads to fewer clusters. Summaries from Spanish have a larger proportion of compressions in the summaries than other languages.

Table 3. Statistics about clusters and compressions for texts translated into French.

	English	Portuguese	Spanish
#clusters	50	70	75
Average size of clusters	2.2	2.7	2.8
Average length of clusters	29.1	25.6	35.4
Average length of compressions	21.7	18.6	23.5
Average number of compressions in summaries	0.7	0.9	1.3
Average compression rate of compressions	74.6%	72.6%	66.4%

We apply a similar analysis for the generation of English summaries from French, Portuguese, and Spanish source texts. As observed before for French summaries, the joint analysis still outperformed other baselines (Table 4). While CoRank obtained a large range of ROUGE scores among different languages (ROUGE-1 between 0.4602 and 0.4715), our approach obtained the best ROUGE scores for all languages with a small difference of ROUGE scores (ROUGE-1 between 0.4725 and 0.4743), which proves that our method generates more stable cross-language summaries for several languages. Chunks spot by the syntactic pattern and the Stanford CoreNLP helped our approach to produce more informative compressions, which results in better ROUGE scores.

Table 4. ROUGE f-scores for cross-language summaries from French, Portuguese, and Spanish languages to English language.

Methods	French			Portuguese			Spanish		
	R-1	R-2	R-SU4	R-1	R-2	R-SU4	R-1	R-2	R-SU4
Late	0.4149	0.1030	0.1594	0.4161	0.1010	0.1576	0.4107	0.1083	0.1603
Early	0.4163	0.1021	0.1602	0.4135	0.1003	0.1580	0.4148	0.1132	0.1644
SimFusion	0.4179	0.1042	0.1607	0.4157	0.0999	0.1582	0.4099	0.1103	0.1616
CoRank	0.4645	0.1326	0.1939	0.4715	0.1415	0.2015	0.4602	0.1414	0.1966
Our approach	0.4727	0.1375	0.1969	0.4743	0.1466	0.2047	0.4725	0.1458	0.2027
Our approach w/o chunks	0.4704	0.1391	0.1963	0.4731	0.1444	0.2037	0.4648	0.1393	0.1975

English translations have fewer words and a smaller vocabulary (difference bigger than 1,000 words) than source texts (Table 5). These translations also have shorter sentences and a more similar vocabulary size than French translations and source texts. The combination of syntactic patterns and the Stanford CoreNLP led to the same characteristics of chunks in terms of numbers and sizes.

Table 5. Statistics of datasets and their translation to English.

	French		Portuguese		Spanish	
	Source	En-Translation	Source	En-Translation	Source	En-Translation
#words	41,071	35,929	37,339	35,244	40,440	37,066
#vocabulary	8,837	7,718	8,694	7,615	8,808	7,703
#sentences	2,000	2,000	2,002	2,002	1,787	1,787
sentence length	20.5	18.0	18.6	17.6	22.6	20.7
#chunks	–	4,302	–	4,327	–	4,324
Average length of chunks	–	2.3	–	2.3	–	2.3

Table 6 details the clustering and the compression processes for the English translations. These translations from French source texts have more clusters because we used a smaller similarity threshold to consider two sentences as similar. English summaries from French have more compressions because of the large number of clusters.

Table 6. Statistics about clusters and compressions for English translated texts.

	French	Portuguese	Spanish
#clusters	128	69	84
Average size of clusters	2.7	2.7	2.8
Average length of clusters	22.1	19.2	27.0
Average length of compressions	16.4	16.3	21.1
Average number of compressions in summaries	2.5	0.9	1.5
Average compression rate of compressions	74.2%	84.9%	78.1%

French and Portuguese source texts have almost the same number of sentences, while English and Spanish source texts have fewer sentences. Comparing the results of English and French translations, English compressions are shorter than French compressions. The use of chunks in MSC improved the results of our cross-language summaries, especially for English translations that have chunks that are more numerous and complex than French translations.

To sum up, our approach has shown to be more stable than extractive methods, thus generating more informative cross-language summaries with consistent ROUGE scores measured in several languages.

5 Conclusion

Cross-Language Text Summarization (CLTS) produces a summary in a target language from documents written in a source language. It implies a combination of the processes of automatic summarization and machine translation. Unfortunately, this combination produces errors, thereby reducing the quality of summaries. A joint analysis allows CLTS systems to extract relevant information from source and target languages, which improves the generation of extractive cross-language summaries. Recent methods have proposed compressive and abstractive approaches for CLTS; however, these methods use frameworks or tools that are available in few languages, limiting the portability of these methods to other languages. Our Multi-Sentence Compression (MSC) approach generates informative compressions from several perspectives (translations from different languages) and achieves stable ROUGE results for all languages. In addition, our method can be easily adapted to other languages.

As future work, we plan to reduce the number of errors generated from the pipeline made of the compression and machine translation processes by developing a Neural Network method to jointly translate and compress sentences. It would also be interesting to include a neural language model to correct possible errors produced during the sentence compression process.

Acknowledgement

This work was granted by the European Project CHISTERA-AMIS ANR-15-CHR2-0001. We also like to acknowledge the support given by the Laboratoire VERIFORM from the École Polytechnique de Montréal and her coordinator Hanifa Boucheneb.

References

1. Yao, J., Wan, X., Xiao, J.: Phrase-based compressive cross-language summarization. In: EMNLP. (2015) 118–127
2. Wan, X., Luo, F., Sun, X., Huang, S., Yao, J.g.: Cross-language document summarization via extraction and ranking of multiple summaries. *Knowledge and Information Systems* (2018)
3. Linhares Pontes, E., Huet, S., Torres-Moreno, J.M., Linhares, A.C.: Cross-language text summarization using sentence and multi-sentence compression. In: 23rd International Conference on Natural Language & Information Systems (NLDB). (2018)
4. Giannakopoulos, G., El-Haj, M., Favre, B., Litvak, M., Steinberger, J., Varma, V.: TAC2011 multiling pilot overview. In: 4th Text Analysis Conference TAC. (2011)
5. Wan, X.: Using bilingual information for cross-language document summarization. In: ACL. (2011) 1546–1555
6. Moirón, B.V., Tiedemann, J.: Identifying idiomatic expressions using automatic word-alignment. In: EACL 2006 Workshop on Multiword Expressions in a Multilingual Context. (2006)
7. de Caseli, H.M., Ramisch, C., das Graças Volpe Nunes, M., Villavicencio, A.: Alignment-based extraction of multiword expressions. *Language Resources and Evaluation* **44** (2010) 59–77
8. Manning, C., Surdeanu, M., Bauer, J., Finkel, J., Bethard, S., McClosky, D.: The Stanford CoreNLP natural language processing toolkit. In: 52nd Annual Meeting of the Association for Computational Linguistics (ACL): System Demonstrations. (2014) 55–60
9. Linhares Pontes, E., Huet, S., Gouveia da Silva, T., Linhares, A.C., Torres-Moreno, J.M.: Multi-sentence compression with word vertex-labeled graphs and integer linear programming. In: TextGraphs-12: the Workshop on Graph-based Methods for Natural Language Processing, Association for Computational Linguistics (2018)
10. Filippova, K.: Multi-sentence compression: Finding shortest paths in word graphs. In: COLING. (2010) 322–330
11. Wan, X., Li, H., Xiao, J.: Cross-language document summarization based on machine translation quality prediction. In: ACL. (2010) 917–926
12. Boudin, F., Huet, S., Torres-Moreno, J.M.: A graph-based approach to cross-language multi-document summarization. *Polibits* **43** (2011) 113–118

13. Zhang, J., Zhou, Y., Zong, C.: Abstractive cross-language summarization via translation model enhanced predicate argument structure fusing. *IEEE/ACM Trans. Audio, Speech & Language Processing* **24** (2016) 1842–1853
14. Bing, L., Li, P., Liao, Y., Lam, W., Guo, W., Passonneau, R.J.: Abstractive multi-document summarization via phrase selection and merging. In: *ACL, The Association for Computer Linguistics* (2015) 1587–1597
15. Lin, C.Y.: ROUGE: A Package for Automatic Evaluation of Summaries. In: *Workshop Text Summarization Branches Out (ACL'04)*. (2004) 74–81