

Extending Automatic Discourse Segmentation for Texts in Spanish to Catalan

Iria da Cunha¹, Eric SanJuan², Juan-Manuel Torres-Moreno^{2,3}, Irene Castellón⁴, and Marina Lloberes⁴

¹ Universidad Nacional de Educación a Distancia (UNED), Madrid, Spain
`iriad@flog.uned.es`

² LIA, Université d'Avignon et des Pays de Vaucluse, France
`{juan-manuel.torres, eric.sanjuan}@univ-avignon.fr`

³ Ecole Polytechnique de Montréal - Département de Génie informatique, Montréal, Canada

⁴ Universitat de Barcelona - Departament de Lingüística General, Barcelona, Spain
`icastellon@ub.edu; mllobesa8@alumnes.ub.edu`

Abstract. At present, automatic discourse analysis is a relevant research topic in the field of NLP. However, discourse is one of the phenomena most difficult to process. Although discourse parsers have been already developed for several languages, this tool does not exist for Catalan. In order to implement this kind of parser, the first step is to develop a discourse segmenter. In this article we present the first discourse segmenter for texts in Catalan. This segmenter is based on Rhetorical Structure Theory (RST) for Spanish, and uses lexical and syntactic information to translate rules valid for Spanish into rules for Catalan. We have evaluated the system by using a gold standard corpus including manually segmented texts and results are promising.

Keywords: Discourse Parsing, Discourse Segmentation, Rhetorical Structure Theory, Shallow Parsing, Catalan

1 Introduction

Nowadays discourse parsing is a very prominent research area used in Natural Language Processing (NLP). Recently other NLP applications and approaches that underlie discourse parsing have arose, such as Machine Translation [1], Textual Similarity [2], and Sentiment Analysis and Opinion Mining [3, 4] for example.

In order to develop these applications, discourse segmenters and parsers are needed, as well as discourse level annotated corpora. Several resources have been developed for different languages. Regarding Iberian Peninsula's Romance languages, there are resources for Portuguese [5] and Spanish [2]. However, they have not been developed for Catalan yet. Catalan is a romance language that comes from Latin. It is spoken in several parts of Spain (mainly in Catalonia), Andorra, France (Roussillon) and Italy (Alghero), among others. Despite this

number of speakers, this is an under-resourced language and there isn't any discourse annotated corpus available.

Most of the related work mentioned in this paper relies on *Rhetorical Structure Theory* (RST) by [6]. According to this theory, a text can be considered as a hierarchical tree made of elementary discourse units (EDUs) that can work as nucleus or as satellite. While nuclei provide relevant information about the author's point of view, satellites give additional information associated to nuclei. RST discourse parsing is formed by three steps: 1) text segmentation in EDUs, 2) discourse relations analysis and 3) discourse tree building.

There are three strategic approaches to impulse the development of resources for under-resourced languages: a) using the crowd and collaborative platforms; b) using technologies of interoperability with well-developed languages; and c) using Semantic Web technologies and, more specifically, Linked Data. In this work, we focus on the second approach. Therefore, the main goal of this work is to develop and evaluate the first discourse segmenter for Catalan, by adapting the existing discourse segmenter for Spanish (a technologically advanced language).

We use the FreeLing Shallow parser for Catalan [7] as a linguistic resource. The system is formed by linguistic rules based on lexical units (conjunctions and adverbs), discourse markers, syntactic structures and punctuation marks. Furthermore, this work aims to present the annotated corpus developed for evaluating the segmenter. This corpus has been developed as a gold standard open to scientific community.

In section 2, we present the state of the art on discourse segmentation. In section 3, we explain the methodology used to develop our segmenter for Catalan and we go deeper into the system implementation. The details about the experiments and the results are described in section 4. Finally, the conclusions of the current research and the future work are presented in section 5.

2 State of the Art

As stated in [8], discourse is one of the most difficult language levels to process automatically due to its complexity. Actually, this difficulty makes automatic discourse analysis a challenging task because it can be applied to develop several NLP tools. In this sense, [9] and [10] have made surveys upon the current research on discourse parsing and its applications. Some examples are text generation [11–13] and automatic summarization [14–16]. Specifically, research on discourse segmentation has been proved to be useful for different NLP tasks. For example, in [17–19] authors study the relationship between discourse segmentation and compression for sentences in Spanish. They present a method for sentence compression that uses statistical information in order to delete intra-sentence discourse segments, obtaining very good results. This kind of systems has been successfully employed for cinema or TV subtitling, and for elaboration of short messages by mobile companies, among other applications. Also, discourse segmentation has been used for machine translation. For example, [20] use discourse segments to align passages of texts in different languages. Usu-

ally, in this field, alignment is done between sentences, but alignment between discourse segments (that is, parts of those sentences) can offer additional information that can be useful for machine translation and other fields.

Existing discourse segmenters employ several strategies; the most productive strategy has been the use of linguistic information (mainly, lexical and syntactic information). The segmenter for English by [21] is based on a statistical model that uses lexical and syntactic features to assign a probability to the insertion of a segment boundary after every word of a sentence. The segmenter for English by [22] is based on linguistic rules (lexical and syntactic) and uses a constituency-based parser. The segmenter for Spanish by [23] is based on linguistic rules adapted to this language and on the grammar of the shallow parser of Freeling.

Another possible strategy to deal with discourse segmentation is machine learning. It would be the case, for example, of the segmenter for French by [24]. However, these types of approaches need a high amount of annotated texts in order to learn and carry out an adequate segmentation. Therefore, nowadays their results are not better than the results of systems based on linguistic information.

Recently, the possibility of developing language-independent segmenters or systems using very few linguistic resources is being explored. It is the case of [25], whose system uses general statistical techniques based on morphological tags, and general linguistic rules. The only language-dependent linguistic resource is a list of discourse markers in the language of the text. The advantages of this type of strategies are that they are easy to implement and require very few resources. Therefore, they are especially adequate for languages without NLP tools. However, currently, the results of these segmenters are not better than the results of the systems designed specifically for particular languages.

In some cases, different strategies have been applied over the same monolingual corpus and have been compared among them, with the aim of determining the most productive strategy for a specific language. It is the case of [26], where both statistical and linguist strategies are applied to segment texts in Basque. Also, constituent and dependency parsing are used. In this work, the strategy based on dependency parsing is the most productive. This is due to the high complexity and particular syntactic characteristics of Basque, since it is an agglutinative language.

3 Methodology

We follow the linguistic strategy to develop a discourse segmenter for Catalan. There are two main reasons for applying this strategy. First, Catalan and Spanish are very similar languages, and the linguistic strategy has been applied for Spanish, obtaining good results. Second, nowadays, in general, this strategy is the most productive for all languages in the state of the art, especially if the used corpus is limited.

3.1 Shallow parsing

The used grammar is an extension of the grammar of the shallow parser for Catalan included in Freeling [7, 27]. The aim of this extension is to detect and re-categorize those words or groups of words that can indicate a boundary between discourse segments in sentences. These rules indicate units that can work as discourse markers. For that, two lexicons of Catalan discourse markers have been processed [28, 29] and 252 markers have been obtained; they were divided in two groups: ambiguous markers (118) and non-ambiguous markers (134). Finally, the 252 rules developed were added to the grammar.

Non-ambiguous markers have been introduced in a new category ‘disc-mk’ (discourse marker) and therefore have been re-categorized in the grammar of the shallow parser for Catalan. For example: adverbs and adverbial groups (*aleshores*, ‘so’; *així doncs*, ‘therefore’), prepositional groups (*per causa de*, ‘because of’) or sequences of lexical units (*tot seguit*, ‘next’; *tot i que*, ‘although’). On the other hand, the category ‘disc-mk-amb’ includes composed elements that have been re-categorized from different tags of the shallow parser. For example: *com a mostra* (‘as it is shown by’) or *després* (‘after’).

In the case of ambiguous markers, it is necessary to take into account the context where they appear and require advance parsing. For example, the marker *després* (‘after’) can be just an adverb (see example 1) or a discourse maker (see example 2):

1. *Els resultats mostren que **després** del test augmentaren els valors.* (‘The results show that after the test the values increased.’)
2. *Els jugadors de futbol de categoria juvenil van tenir fatiga del sistema nerviós **després** de realitzar un test de capacitat d’esprints repetits (CER).* (‘The football players of the youth category had fatigue of the nervous system after carrying out a repeated sprint test (RST).’)

In the grammar, rules related to discourse markers (both ambiguous and non-ambiguous) have priority over the other rules.

3.2 Implementation

The resulting segmenter for Catalan called DiSegCAT is then generated automatically from the segmenter for Spanish. These tools are implemented in Perl, and are based on regular expressions and the Twig XML library. Firstly, DiSeg calls the FreeLing library in order to apply the discourse markers grammar. This first module transforms the syntactic tree generated by the FreeLing Shallow Parser into XML format. The second module applies the rules that detect EDUs boundaries. This module reads all the leaves of the XML syntactic tree. When it recognizes a boundary, the syntactic tree is modified by adding a new node where the boundary is located. This task is iterated twice to find sub-EDUs boundaries inside the EDUs already detected. The third DiSeg module re-reads the new XML syntactic tree to split sentences into coherent EDUs that should contain one or more verbs.

These modules use regular expressions to recognize discourse markers, lemmas and grammatical categories. The idea was the development of a Perl script to translate the code for Spanish discourse segmentation to Catalan. This strategy presupposes that EDUs have the same grammatical structure and segmentation markers are the only elements that are altered. This assumption is not entirely true. There are discourse markers in Spanish that correspond to one or more discourse markers in Catalan. For example, the Spanish marker *para* ('to') is *per* ('to', 'by' or 'through') or *en* ('in') in Catalan. Furthermore, there are grammatical categories tags, such as 'vaux' for auxiliary verbs, that do not have a correspondence in the grammar version of Catalan. Therefore, a lexicon that transforms the Spanish DiSeg to the Catalan DiSeg cannot handle these cases but contextual rules can contribute to solve these asymmetries between both languages.

The final implementation is the `es2cat.pl` perl script that translates the original DiSeg into DiSegCAT. As a consequence, any expansion of the original DiSeg is applied automatically to the Catalan version. This strategy relies on the idea that translating a NLP software is much more easier than translating texts written in natural language. Therefore, the results provided by DiSegCAT are more reliable than the original DiSeg combined with a machine translation system Catalan \leftrightarrow Spanish.

4 Experiments and Results

We evaluate system performance over a corpus of manually segmented texts. This type of evaluation is used in previous work in this area [22, 24, 23].

4.1 Corpus

The corpus includes 20 abstracts of research articles in Catalan from the medical domain, extracted from the specialized Journal of Medicine and Physical Activity and Sport *Apunts: Medicina de l'esport*⁵. This journal publishes each article in two languages, Spanish and Catalan, and provides the abstract in these two languages and also in English. This would allow us to perform experiments with parallel corpora in the future. Specifically, for this work, texts published between 2010 and 2013 were selected, in order to have recent documents. Also, texts related to different subjects were selected, such as scoliosis, attention deficit, cardiology, nutrition, etc., in order to guarantee thematic diversity. The textual genre 'abstract' and the medical domain were selected to be able to compare adequately the results of our discourse segmenter for Catalan with the results obtained by the discourse segmenter for Spanish; in the evaluation of this Spanish segmenter, 20 texts with similar characteristics were used.

Once the segmentation criteria were defined for Catalan and the corpus was compiled, manual discourse segmentation was carried out. For that, two annotators were asked to segment each text of the corpus, following the mentioned

⁵ <http://www.raco.cat/index.php/Apunts/issue/archive>

criteria, individually and without questions between them, to avoid biases in the results. Both annotators are linguists and have a wide experience in corpus annotation. After the manual annotation of the 20 texts, both segmentations were compared, in order to determine the inter-annotator agreement. Annotators agreed on 264 discourse segment boundaries and they disagreed on 23 boundaries. Therefore, there was a boundary agreement of 92%.

Following [30] and [31], we have also calculated inter-annotator agreement by using Kappa Cohen in two ways: taking into account words as boundaries and taking into account clauses as boundaries. For the first one, the Kappa value is 0.9556 and, for the second one (that is more conservative), the Kappa value is 0.8674. We consider that these values show that agreement between annotators is high for the task of discourse segmentation.

After this quantitative analysis, a qualitative analysis of the disagreements was done, where we observed that nearly all the disagreements were due to human mistakes. Finally, in the line of work on this topic [23, 26] and [8] indicates, a debate was carried out between annotators to solve disagreements. Thus, agreement was obtained for every case. This final corpus segmented was used as gold standard and will be available online for the scientific community. Table 1 shows the gold standard statistics. As it can be observed in this table, the corpus includes 183 sentences; by contrast, it contains a higher number of discourse segments (280), which means that intra-sentence discourse segmentation is productive.

	Total	Longest text	Shortest text	Average
Num. of words	4 676	317	91	233.80
Num. of sentences	183	17	4	9.15
Num. of segments	280	24	8	14.00

Table 1. Gold Standard statistics.

4.2 Evaluation

We consider two baseline segmenters to compare our results:

- Baseline₁: it inserts boundaries before coordinating conjunctions.
- Baseline₂: it considers all complete sentences like discourse segments. This baseline will have a precision of 100%, because all detected segments will be correct, since sentences are considered discourse segments.

The results obtained are shown in Table 2. Differences between DiSegCAT and Baseline₂ scores are all significant based on 12-fold t-test with p -value < 0.05 .

Since there is not another discourse segmenter for Catalan we cannot compare our results with another system. This is a current situation when working on

System	F-Score	Precision	Recall
DiSegCAT	75%	68%	85%
Baseline₁	52%	44%	65%
Baseline₂	18%	100%	10%

Table 2. Results of our experiment.

local languages. We did try to combine DiSeg for Spanish with mainstream available translators from Spanish to Catalan. But translation is an even more complex problem than discourse analysis and in the case of our corpus in Spanish, translators were not efficient with high rate of errors including over short multi-word expressions.

However, we find that our results are similar to those obtained for other languages by using similar strategies of segmentation, such as for English (F-Score = 83%) and Spanish (F-Score = 80%).

As it can be observed in Table 2, our system (DiSegCAT) obtains the best F-Score (75%), in comparison with Baseline₁ (52%) and Baseline₂ (18%). As expected, Baseline₂ obtains 100% of precision, since it considers sentences as segments; however, it obtains only 10% of recall, since it does not detect intra-sentence segments. Baseline₁ has 44% of precision and 65% of recall. Although this baseline detects correctly some segments (because coordinated clauses can be also discourse segments), both results (precision and recall) are worst than the results obtained by DiSegCAT. These results mean that our algorithm outperforms baselines including linguistic information.

Regarding the results obtained by DiSegCAT, its performance is better for recall than for precision (85% and 68%, respectively). After obtaining these quantitative results, we have carried out a qualitative analysis in order to find different types of errors of the system.

With respect to precision, the main error is related to coordination. See for example the following segments, obtained automatically by DiSegCAT⁶:

[El nostre objectiu fou establir quins paràmetres antropomètrics]
 [i de maduració es correlacionen amb el rendiment
 en rem-ergòmetre en una mostra de 114 adolescents d'ambdós sexes,
 sense experiència prèvia en rem.]

[We aimed to establish which anthropometric]
 [and maturity offset parameters correlate with
 rowing ergometer performance in a sample of 114 adolescent,
 rowing-inexperienced boys and girls.]

Here, following our segmentation criteria and rules, the correct segmentation would be:

⁶ English translation of examples has been extracted from the papers published by the authors in the journal *Apunts: Medicina de l'esport*.

[El nostre objectiu fou establir quins paràmetres antropomètrics i de maduració es correlacionen amb el rendiment en rem-ergòmetre en una mostra de 114 adolescents d’ambdós sexes, sense experiència prèvia en rem.]

The passage *quins [...] rem* is the direct object of the main verb (*fou establir*, “was to establish”) of the sentence. Therefore, it should not be segmented. Nevertheless, this direct object includes a coordination that contains the conjunction *i* (“and”) and the finite verb *es* (“is”). Thus, the system segments after the conjunction, since one of the rules indicates that a passage written after this conjunction should be a discourse segment if it includes a finite verb. As it can be observed, in this case, the performance of this rule is not adequate. We should find a solution for this problem in the future. Coordination is also one of the main difficulties found in the performance of the discourse segmenter for Spanish [32].

With regard to recall, the main problem is related to segments that are not explicitly marked in the text. See for example the following segment, obtained automatically by the system:

Té un cost baix, és massiva i de fàcil aplicació.

It is low cost, it is massive and easy to use.

The adequate segmentation of this passage should be:

[Té un cost baix]
[és massiva i de fàcil aplicació.]

The second segment includes a finite verb, but there is no specific mark indicating that it is a discourse segment (the comma is not included in our rules as a boundary mark, since it would over-generate discourse segments). In the future, we plan to study strategies to solve this limitation, although it is a difficult issue.

5 Conclusions and Future Work

This paper presents the first discourse segmentation system for Catalan based on RST. The segmenter uses simple linguistic rules and promising results have been obtained in our experiments. This system could be used in different tasks in the context of NLP. For example, segmentation of (too) long sentences which are difficult to parse; elimination of text segments in Sentence Compression and Automatic Text Summarization; Alignment of text in different languages for Machine Translation, etc. Also, the segmenter can be the basis for further development of an automatic discourse parsing system for Catalan, since this tool does not exist so far.

The methodology used to develop this segmenter for Catalan is very similar to that used for Spanish. In fact, the good results obtained show us that this

methodology is probably valid in general for Romance languages. The future experiments will be conducted on this line. We also plan to expand the size of the corpus, adding texts of other genres (such as news reports) and other domains (such as Linguistics or Economy).

Acknowledgements

This work has been partially supported by a Ramón y Cajal research contract (RYC-2014-16935) and the research project APLE 2 (FFI2009-12188-C05-01) of the Institute for Applied Linguistics (IULA).

References

1. Guzmán, F., Joty, S., Márquez, L., Nakov, P.: Using discourse structure improves machine translation evaluation. In: ACL 2014. Volume 1., Association for Computational Linguistics (ACL) (2014) 687–698
2. da Cunha, I., Vivaldi, J., Torres-Moreno, J.M., Sierra, G.: SIMTEX: An Approach for Detecting and Measuring Textual Similarity based on Discourse and Semantics. *Computación y Sistemas* **18**(3) (2014) 505–516
3. Trnavac, R., Taboada, M.: Discourse structure and attitudinal valence of opinion words in sentiment extraction. In: 88th Meeting of the Linguistic Society of America, Poster, Minneapolis, MN (2014)
4. Chenlo, J., Hogenboom, A., Losada, D.: Sentiment-based Ranking of Blog Posts using Rhetorical Structure Theory. In: NLDB’13, Salford, UK (2013) 13–24
5. Pardo, T., Nunes, M.: On the Development and Evaluation of a Brazilian Portuguese Discourse Parser. *Journal of Theoretical and Applied Computing* **15**(2) (2008) 43–64
6. Mann, W., Thompson, S.: Rhetorical structure theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse* **8**(3) (1988) 243–281
7. Atserias, J., Casas, B., Comelles, E., González, M., Padró, L., Padró, M.: FreeLing 1.3. Syntactic and semantic services in an open-source NLP library. In: LREC’06, Genoa, Italy, ELRA (2006) 48–55
8. Hovy, E.: Annotation. In: Tutorial Abstracts of ACL 2010, Uppsala, Sweden, ACL (July 2010) 4
9. Taboada, M., Mann, W.: Applications of Rhetorical Structure Theory. *Discourse Studies* **8**(4) (2006) 567–588
10. da Cunha, I., Torres-Moreno, J.M., Sierra, G.: Aplicaciones lingüísticas del análisis discursivo automático. In Ruiz, L.; Álvarez, M.R., ed.: *Comunicación Social en el Siglo XXI Vol. II*, Centro de Lingüística Aplicada (2011) 919–923
11. Hovy, E.: Automated discourse generation using discourse structure relations. *Artificial Intelligence* **63** (1993) 341–385
12. Dale, R., Hovy, E., Rösner, D., O., S.: *Aspects of Automated Natural Language Generation*. Springer (1992)
13. O’Donnell, M., Mellish, C., Oberlander, J., Knott, A.: ILEX: An architecture for a dynamic Hypertext generation system. *Natural Language Engineering* **7** (2001) 225–250

14. Marcu, D.: The Theory and Practice of Discourse Parsing Summarization. Institute of Technology, Massachusetts (2000)
15. Radev, D.: A common theory of information fusion from multiple text sources. Step one: Cross document structure. In: 1st SIGdial Workshop on Discourse and Dialogue, Hong-Kong (2000) 74–83
16. Pardo, T., Rino, M.: DMSumm: Review and assessment. In: PorTAL’02, Faro, Portugal:Springer (2002) 263–274
17. Molina, A., Torres-Moreno, SanJuan, E., da Cunha, I., Sierra, G.J.M., Velázquez-Morales, P.: Discourse Segmentation for Sentence Compression. In: LNAI 7094. Volume abs/1212.3493., Berlin: Springer (2011) 316–327
18. Molina, A., Torres-Moreno, J.M., da Cunha, I., SanJuan, E., Sierra, G.: Sentence Compression in Spanish driven by Discourse Segmentation and Language Models. CoRR abs/1212.3493 (2012)
19. Molina, A., Torres-Moreno, J.M., SanJuan, E., da Cunha, I., Sierra, G.: Discursive Sentence Compression. In: LNCS. Volume 7817., Springer (2013) 394–407
20. Ghorbel, H., Ballim, A., Coray, G.: ROSETTA: Rhetorical and Semantic Environment for Text Alignment. In Rason, P., e.a., ed.: Proceedings of Corpus Linguistics 2001. (2001) 224–233
21. Soricut, R., Marcu, D.: Sentence Level Discourse Parsing Using Syntactic and Lexical Information. In: Conference of the North American Chapter of the ACL on Human Language Technology, Edmonton, Canada (2003) 149–156
22. Tofiloski, M., Brooke, J., Taboada, M.: A Syntactic and Lexical-based Discourse Segmenter. In: 47th Meeting of the ACL, Singapur (2009) 77–80
23. da Cunha, I., SanJuan, E., Torres-Moreno, J.M., Lloberes, M., Castellón, I.: DiSeg: Un segmentador discursivo automático para el español. *Procesamiento del Lenguaje Natural* **45** (2010) 145–152
24. Afantenos, S., Denis, P., Muller, P., Danlos, L.: Learning recursive segments for discourse parsing. In: LREC’10, Malta, ELRA (2010) 3578–3584
25. Saksik, R., Molina, A., Carneiro Linhares, A., Torres-Moreno, J.M.: Segmentacao discursiva automática: uma avaliação preliminar em francês. In: 4th Workshop RST and Discourse Studies, STIL’13, Fortaleza, Brasil (2013) 40–49
26. Iruskietia, M., Aranzabe, M., Díaz de Ilarraza, A., Gonzalez, I., Lersundi, M., Lopez de Lacalle, O.: The RST Basque TreeBank: an online search interface to check rhetorical relations. In: 4th Workshop RST and Discourse Studies, STIL’13, Fortaleza, Brasil (2013) 40–49
27. Padró, L., Collado, M., Reese, S., Lloberes, M., Castelló, I.: Freeling 2.1: Five years of open-source language processing tools. In: LREC’10, Malta (2010) 931–936
28. Alonso, L.: Representing Discourse for Automatic Text Summarization via Shallow. PhD thesis, Univ. de Barcelona (2005)
29. Generalitat-Catalunya: Redacció de documents: mots connectors. Barcelona: Departament de Justícia. (2014)
30. Iruskietia, M.: The relational discourse structure in pragmatics: description and evaluation in Computational Linguistics. PhD thesis, University of the Basque Country (2014)
31. Iruskietia, M., Díaz de Ilarraza, A., Lersundi, M.: Establishing criteria for RST-based discourse segmentation and annotation for texts in Basque. *Corpus Linguistics and Linguistic Theory (CLLT)* **9**(1) (2013) 1–32
32. da Cunha, I., SanJuan, E., Torres-Moreno, J.M., Cabré, M., Sierra, G.: A Symbolic Approach for Automatic Detection of Nuclearity and Rhetorical Relations among Intra-sentence Discourse Segments in Spanish. In: LNCS. Volume 7181., Springer (2012) 462–474