

# A rule-based system for cross-lingual parsing of Romance languages with Universal Dependencies

**Marcos Garcia**

LyS Group  
Departamento de Letras  
Universidade da Coruña

marcos.garcia.gonzalez@udc.gal

**Pablo Gamallo**

Centro Singular de Investigación en  
Tecnoloxías da Información (CiTIUS)  
Universidade de Santiago de Compostela

pablo.gamallo@usc.es

## Abstract

This article describes *MetaRomance*, a rule-based cross-lingual parser for Romance languages submitted to *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. The system is an almost delexicalized parser which does not need training data to analyze Romance languages. It contains linguistically motivated rules based on PoS-tag patterns. The rules included in *MetaRomance* were developed in about 12 hours by one expert with no prior knowledge in Universal Dependencies, and can be easily extended using a transparent formalism. In this paper we compare the performance of *MetaRomance* with other supervised systems participating in the competition, paying special attention to the parsing of different treebanks of the same language. We also compare our system with a delexicalized parser for Romance languages, and take advantage of the harmonized annotation of Universal Dependencies to propose a language ranking based on the syntactic distance each variety has from Romance languages.

## 1 Introduction

This article describes the *MetaRomance* parser, which participated at *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies* (Zeman et al., 2017). *MetaRomance* is a rule-based parser for Romance languages adapted to Universal Dependencies (UD). The system relies on a basic grammar consisting on simple cross-lingual and (almost) delexicalized rules likely to be shared by most Romance lan-

guages. Rules are almost delexicalized because they are mainly applied on Universal PoS-tags, only containing few grammar words (some prepositions and conjunctions) together with a small list of verbs. The grammar was developed by one expert with no prior knowledge in UD in about 12 hours.<sup>1</sup>

As the Universal Dependencies initiative (Nivre et al., 2016) offers linguistic criteria providing a consistent representation across languages, it fits perfectly with our objective of defining cross-lingual rules. In fact, the availability of harmonized treebanks provides an interesting test bench for cross-lingual dependency parsing research (McDonald et al., 2011; Mcdonald et al., 2013; Vilares et al., 2016).

Our participation at this CoNLL 2017 shared task has several experimental objectives. First, we will compare our rule-based approach with the rest of participants, which are likely to be supervised systems, with regard to Romance languages. Namely, we will analyze the performance of several systems on different treebanks of the same language. Then, we will also evaluate the cross-lingual property of our system by comparing it with a supervised delexicalized parser. Last but not least, the analysis of the results in the shared task will allow us to check whether our method might be useful to measure the syntactic distance between Romance and non-Romance languages.

The results of different experiments show that, in spite of its simplicity, *MetaRomance* achieve reasonable results on Romance languages with no training data, and that its performance is relatively uniform across different treebanks of the same language. The delexicalized rules of this system al-

<sup>1</sup>The whole system, DepPattern and *MetaRomance*, is freely available at:  
<https://github.com/CoNLL-UD-2017/MetaRomance>.

lowed us to present a classification of all the languages present in the shared task ranked by their syntactic distance from Romance languages.

The remaining of the paper is organized as follows. Section 2 presents some related work on dependency and cross-lingual parsing. Then, we present the architecture of *MetaRomance* in Section 3, and several experiments on Section 4. Finally, we briefly discuss the results and present the conclusions of our work in Sections 5 and 6, respectively.

## 2 Related Work

In the last 15 years, most work on dependency parsing has been developed within two supervised (data-driven) approaches: deterministic parsing, which is also known as *transition-based parsing* (Yamada and Matsumoto, 2003; Nivre, 2004; Gómez-Rodríguez and Fernández-González, 2012), and non-deterministic strategies, known as *graph-based dependency parsing* (McDonald and Pereira, 2006; Carreras, 2007; Martins et al., 2010).

In opposition to data-driven approaches, many grammar-driven (or rule-based) parsers use finite-state technology, which has attractive properties for syntactic parsing, such as conceptual simplicity, flexibility, and efficiency in terms of space and time. It permits to build robust and deterministic parsers. Most finite-state based parsing strategies use cascades of transducers (Ait-Mokhtar et al., 2002; Oflazer, 2003).

Concerning cross-lingual parsing, there are two main approaches for parsing one language (the target) with resources from one or more source languages: (a) data transfer, and (b) model transfer methods. On the one hand, data transfer approaches obtain annotated treebanks of a target language by projecting the syntactic information from the source data. Some methods use parallel corpora (Hwa et al., 2005; Ganchev et al., 2009; Agić et al., 2016) while others create artificial data taking advantage of machine translation (Tiedemann and Agić, 2016).

On the other hand, the strategies based on model transfer train systems on the source data that can be used to parse a target language (Zeman and Resnik, 2008). The emergence of different initiatives promoting harmonized annotations allowed researchers to explore this approach, using delexicalized models and multi-source strategies (Mc-

Donald et al., 2011; Täckström et al., 2012).

More recently, some works addressed multilingual parsing using a single model (trained in a combination of various treebanks) to analyze different languages (Vilares et al., 2016; Ammar et al., 2016).

The growth in cross-lingual parsing research has given rise to a recent shared task at VarDial 2017 (Zampieri et al., 2017), Cross-lingual Dependency Parsing (CLP) (Tiedemann, 2017). CLP is a shared task whose aim is to develop models for parsing selected target languages without annotated training data, but having annotated data in one or two closely related languages (Rosa et al., 2017).

With the emergence of UD as the practical standard for multilingual PoS and syntactic dependency annotation, it is possible to develop universal rule-based strategies requiring no training data, and relying on basic rules exploiting the UD criteria. The Universal Dependency Parser, described in (Martínez Alonso et al., 2017), is a good example of this unsupervised strategy. Our work goes in that direction, but with two differences: the grammar is focused on Romance languages and the parser relies on basic rules implemented as cascades of finite-state transducers.

## 3 MetaRomance

### 3.1 The architecture

The core of *MetaRomance*, depicted in Figure 1, consists of the following modules:

- An *adapter* converting CoNLL-U into the format required by the rule-based parser.
- A *MetaRomance grammar* with 150 cross-lingual rules configured to work with tags, labels and linguistic constraints of UD.
- A *grammar compiler* that takes the grammar as input and generates a dependency parser, which is based on finite state transitions.

In order to allow *MetaRomance* to work on raw text, some scripts are provided in addition to the core architecture for converting the tags given by different PoS-taggers (namely, FreeLing (Padró and Stanilovsky, 2012; Garcia and Gamallo, 2010), TreeTagger (Schmid, 1994), and LinguaKit (Garcia and Gamallo, 2015)) into the CoNLL-U format. Thus, *MetaRomance* is able

to parse raw text which has been tokenized, lemmatized and PoS-tagged with several systems that provide high-quality analyses for different languages.

### 3.2 The MetaRomance grammar

The cost of writing the grammar is not high since its size is small and the rules are not language-specific. The strategy we followed to write the *MetaRomance* grammar is based on two methodological principles:

- Start with high-coverage rules.
- Otherwise, develop rules shared by as many Romance languages as possible.

The objective is to find a trade-off between high performance and low effort, i.e. we look for efficiency. Most rules satisfy these two principles, giving rise to a broad-coverage parser. We have not defined non-projective rules since, in general, they have low coverage and are language dependent. Some rules contain information on specific lexical units, but only to identify grammatical words: some prepositions, conjunctions, determiners, and pronouns (and a small and automatically extracted list of verbs). Most phenomena not covered by the grammar are related with some long distance dependencies, including subordinate clauses in non-canonical positions, or complex issues derived from coordination.

Cross-lingual rules were written with DepPattern (Gamallo and González, 2011), a high-level syntactic formalism aimed to write dependency-based grammars. This dependency formalism has been adapted so as to let it interpret Universal Dependencies, more specifically UDv2. All rules were written in about 12 hours by an expert linguist who has skills in the DepPattern formalism, but with no prior knowledge in UD. He took into account the syntactic structure of all Romance languages of the UDv2 treebanks except Romanian. In the following you can see an example of a DepPattern rule:

```
det: DET [ADJ]? NOUN
Agreement: gender, number
%
```

The first line contains, divided by the colon, the name of the dependency relation (*det*) together with the PoS context. Here, a determiner will be

linked as dependent of a noun (the head), with an optional adjective between them. The second line states that this rule will only be applied if both the dependent and the head agree in gender and number.

As the grammar is not complete, giving rise to partial parses, we implemented a post-editor script linking all tokens without head information to the corresponding sentence root. Moreover, in order to assign a label to each unknown dependency, the script associates dependency names to PoS-tags: e.g., PUNCT is associated with the dependency name “punct”, VERB with “xcomp”, and so on.

It is worth noting that the rules implemented in *MetaRomance* only make use 25 out of the 37 universal relations defined in the UDv2 guidelines.

### 3.3 A finite-state transition parser

The parser, automatically generated from the formal grammar, is based on a finite-state transition approach making use of a similar strategy to the shift-reduce algorithm. More precisely, it consists of a set of transducers/rules that *compress* the input sequence of tokens by progressively removing the dependent tokens as soon as dependencies are recognized (Gamallo, 2015). So, at each application of a rule, the system reduces the input and make it easier to find new dependencies in further rule applications. In particular, short dependencies are recognized first and, as a consequence, the input is simplified so as to make lighter the recognition of long distance dependencies. This is inspired by the easy-first strategy.

## 4 Experiments

This section presents several evaluations of *MetaRomance* using the data provided by the CoNLL 2017 shared task on UD parsing (Nivre et al., 2017). We will show the results of the following experiments:

- Comparison of *MetaRomance* with other supervised approaches on all the testing treebanks of Romance languages.
- Analysis of the performance of several parsers on different treebanks of the same language.
- Comparison of *MetaRomance* with a neural network delexicalized parser for Romance languages.

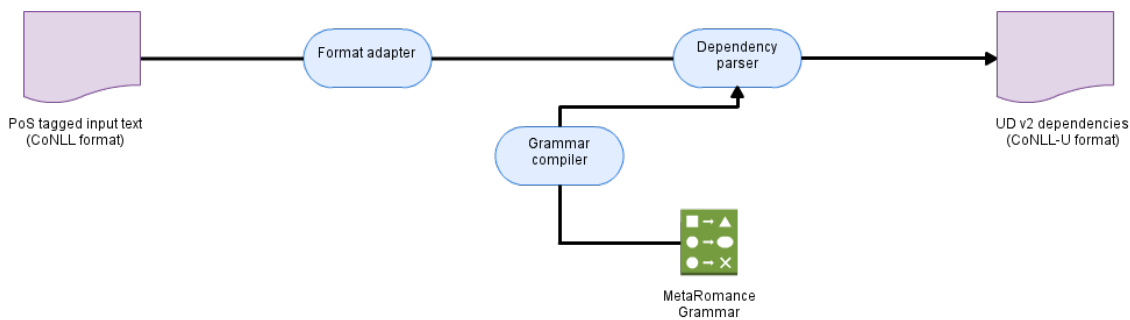


Figure 1: Architecture of *MetaRomance*

- Syntactic distance between Romance and non-Romance languages.

As we had several alignment issues concerning the evaluation of data pre-processed by *LinguaKit* and *FreeLing*, all the experiments presented in this paper (as well as the official *MetaRomance* results) used as input the tokenized, lemmatized and PoS-tagged data provided by the *UDPipe* baseline models.

#### 4.1 Results at CoNLL-2017 shared task

In general, our system obtained low LAS and UAS results in the whole dataset of the shared task (34.05% LAS, 42.55% UAS).<sup>2</sup> The results were mostly expected due to the characteristics of *MetaRomance*: an almost delexicalized parser which does not require training data, with simple rules only based on the syntactic structure of Romance languages.

*MetaRomance* needed 29 minutes and 155MB of memory to parse all the testing sets on the TIRA virtual machine provided by the shared task (Potthast et al., 2014).

Table 1 shows the official *MetaRomance* results on every treebank of a Romance language evaluated in the shared task. On average, our system achieved F1 results of 58.9 (LAS) and 66.1 (UAS). The worst results were obtained in Romanian; this fact was expected because (a) Romanian is linguistically more distant than the other Romance languages (Gamallo et al., 2017), and (b) we did not implement any dependency rule with this language in mind.

<sup>2</sup>After correcting a small bug in a script—which produced invalid treebanks for three languages—we obtained 34.98% LAS and 43.81% UAS. The new results, not present in the official ones, were (LAS/UAS): *bxr*: 19.51/30.22, *cs*: 41.63/47.92, and *tr*: 13.70/23.85.

Even if the values in Table 1 are not comparable with most supervised systems in the competition, our simple parser obtained competitive results in some languages, such as *es*, *it*, and *pt*. Interestingly, *MetaRomance* performed better in the *ptd* datasets than in the others treebanks of the same languages (with only one exception: UAS results in *pt* and *pt\_pud*), while most systems in the shared task decreased their performance in the *ptd* datasets in several points. In this respect, *MetaRomance* led some supervised approaches in treebanks such as *pt\_pud* or *gl\_treegal* (this last one with small training data).

Some of the results on different treebanks of the same language have noticeable LAS differences: more than 5 points between *es* and *es\_pud*, and about 10% between *pt\_br* and the two other treebanks of Portuguese.<sup>3</sup>

In this regard, our next experiment compares the cross-treebank performance of supervised models (i.e., parsing different treebanks of the same language with the same model). To carry out this experiment we trained a *UDPipe* model (Straka et al., 2016) in each training dataset of Spanish, Galician, and Portuguese. These models were trained using the default parameters of *UDPipe* 1.1, but removing the lemmas and the morphological features of the treebanks, with a view to building parsers with more robust performance among the different test sets.<sup>4</sup>

<sup>3</sup>Concerning Portuguese, it is important to note that (a) Brazilian Portuguese has some syntactic (as well as morphological, orthographic, etc.) differences regarding European Portuguese, and that (b) the *pt\_br* treebank does not contain lemmas neither morphological features (and also it has some tokenization issues: e.g., most contractions with the preposition *em* are tokenized as *en*).

<sup>4</sup>Preliminary tests using the baseline models provided by the shared task organization showed that some models trained in one treebank may obtain LAS results with drops of more than 26% when parsing a different treebank of the same lan-



Treebank	LAS	UAS
<i>ca</i>	57.71	65.57
<i>es</i>	59.80	67.20
<i>es_ancora</i>	60.99	69.63
<i>es_pud</i>	65.49	71.68
<i>fr</i>	54.10	62.20
<i>fr_partut</i>	56.17	63.10
<i>fr_sequoia</i>	55.16	60.76
<i>fr_pud</i>	58.67	65.94
<i>gl</i>	54.87	62.59
<i>gl_treegal</i>	57.20	63.87
<i>it</i>	62.96	70.35
<i>it_pud</i>	65.49	71.82
<i>pt</i>	65.50	71.77
<i>pt_br</i>	56.19	65.81
<i>pt_pud</i>	66.35	71.43
<i>ro</i>	45.04	53.90
average	58.86	66.10

Table 1: *MetaRomance* results on the Romance languages test sets (predicted tokens, lemmas, features, and PoS-tags).

Table 2 includes the LAS and UAS values of each model (in the columns) on the target treebanks (on each row). These numbers clearly show that the results of supervised models are very different when parsing a different treebank to the one used for training, even if both corpus belong to the same language. These differences are much higher than those reported for *MetaRomance*, exceeding 22% in *gl* parsing *gl\_treegal*, more than 15% in the analysis of *es* by *es\_ancora*, or more than 14 in *pt\_br* parsing *pt*. Note, however, that most supervised parsers (except *gl* analyzing *gl\_treegal*) achieved better results than those obtained by *MetaRomance*.

These results (both the UDPipe and the *MetaRomance* ones) suggest that careful analyses of the different treebanks are required, aimed at knowing whether these large variations are due to different domains, annotation issues, or linguistic differences.

#### 4.2 Comparison with a cross-lingual delexicalized parser

In the next experiment we compare the performance of *MetaRomance* with a delexicalized parser trained with a combined corpus which includes sentences from every Romance treebank.

guage.

Target	Source			
Spanish	<i>es</i>		<i>es_ancora</i>	
	LAS	UAS	LAS	UAS
	<i>es</i>	76.85 81.19	64.25	71.95
Galician	<i>gl</i>		<i>gl_treegal</i>	
	LAS	UAS	LAS	UAS
	<i>gl</i>	73.71 77.17	58.03	68.47
Portuguese	<i>pt</i>		<i>pt_br</i>	
	LAS	UAS	LAS	UAS
	<i>pt</i>	78.74 82.43	68.00	77.92
Spanish	<i>es</i>		<i>es_ancora</i>	
	LAS	UAS	LAS	UAS
	<i>es_ancora</i>	67.25 76.43	79.36	83.42
Galician	<i>gl</i>		<i>gl_treegal</i>	
	LAS	UAS	LAS	UAS
	<i>gl_treegal</i>	50.98 63.37	65.24	70.81
Portuguese	<i>pt</i>		<i>pt_br</i>	
	LAS	UAS	LAS	UAS
	<i>pt_br</i>	66.85 76.19	82.10	84.83
Portuguese	<i>pt</i>		<i>pt_br</i>	
	LAS	UAS	LAS	UAS
	<i>pt_pud</i>	71.59 77.58	67.75	77.87

Table 2: Results of UDPipe models trained in the source treebanks (columns) on the target test sets (rows).

This is a competitive supervised baseline in cross-lingual transfer parsing work, which gives us an indication of how our system compares to standard cross-lingual parsers.

We trained 50 UDPipe models by randomly selecting from 1 to 50 sentences of each Romance treebank in the training data. Then, we obtained the average results on all the Romance test treebanks, and plotted them together with the *MetaRomance* performance in Figure 2.

This figure shows that *MetaRomance* obtains similar results ( $\approx 59\%$  LAS) to those achieved with about 2,000 tokens of all the Romance treebanks. The learning curve also suggest that it is difficult for cross-lingual models with no lexical features (as *MetaRomance*, which is also delexicalized) to keep increasing their cross-lingual performance on Romance languages. Thus, UDPipe achieves 64% with about 5,000 tokens, but it cannot surpass 65% even with a training corpus of 20,000 tokens.

#### 4.3 Syntactic distance from Romance languages

The last experiment is an attempt to rank all the languages in the shared task with respect to the Romance family, aimed at knowing whether it is possible to use these results as a syntactic distance between Romance and non-Romance languages.

Table 3 includes the UAS values obtained by *MetaRomance* in every language of the dataset.

<i>Language</i>	<i>UAS</i>	<i>Family</i>	<i>Branch</i>
Italian	71	Ind	Romance
Portuguese	69	Ind	Romance
Spanish	69	Ind	Romance
Catalan	65	Ind	Romance
Galician	63	Ind	Romance
French	63	Ind	Romance
Bulgarian	54	Ind	Slavic
Romanian	53	Ind	Romance
Greek	52	Ind	European
Russian	52	Ind	Slavic
Indonesian	51	Non-Ind	Austronesian
Polish	51	Ind	Germanic
Old Bulgarian	50	Ind	Slavic
Gothic	50	Ind	Germanic
Ukrainian	48	Ind	Slavic
Dutch	48	Ind	Germanic
Croatian	48	Ind	Slavic
Norwegian	48	Ind	Germanic
Danish	46	Ind	Germanic
Hungarian	46	Non-Ind	Uralic
Czech	46	Ind	Slavic
Sorbian	45	Ind	Slavic
English	45	Ind	Germanic
Swedish	44	Ind	Germanic
Slovak	44	Ind	Slavic
German	43	Ind	Germanic
Old Greek	40	Ind	European
Irish	40	Ind	Celtic
Slovene	38	Ind	Slavic
Hebrew	38	Non-Ind	Semitic
Latin	37	Ind	European
Estonian	37	Non-Ind	Uralic
Arabic	37	Non-Ind	Semitic
Latvian	36	Ind	Germanic
Farsi	35	Ind	Indian
Norther Sami	35	Non-Ind	Uralic
Kurdish	35	Ind	Indian
Finnish	34	Non-Ind	Uralic
Russian Buryat	30	Non-Ind	Mongolic
Korean	27	Non-Ind	Koreanic
Turkish	21	Non-Ind	Turkic
Uyghur	18	Non-Ind	Turkic
Vietnamese	18	Non-Ind	Austroasiatic
Kazakh	18	Non-Ind	Turkic
Basque	18	Non-Ind	<i>isolated</i>
Hindi	15	Ind	Indian
Chinese	14	Non-Ind	Sino-tibetan
Urdu	13	Ind	Indian
Japanese	8	Non-Ind	Japonic

Table 3: *MetaRomance* results (UAS) on each language of the CoNLL 2017 dataset. Results in languages with more than one treebank are the average values.

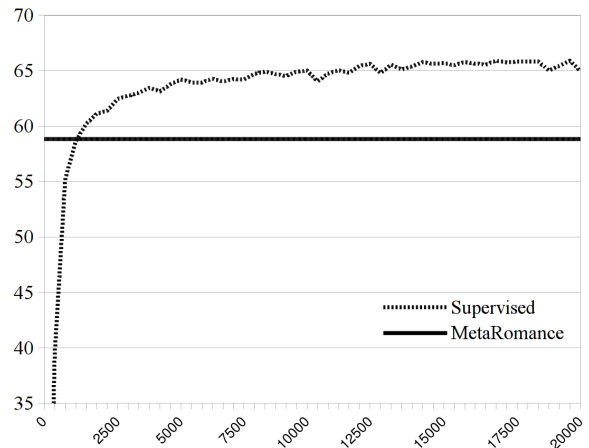


Figure 2: LAS values of *MetaRomance* versus the learning curve (0–20,000 tokens) of a delexicalized UDPipe model trained with random sentences from all the Romance treebanks. Results are average F1 values of all the testing Romance treebanks.

For those languages with more than one treebank we show the average results.<sup>5</sup>

As expected, at the top of the ranking we find Romance languages, on which *MetaRomance* achieves the best results (except on Romanian, slightly surpassed by Bulgarian). With few exceptions, such as the Indian varieties which obtained low values, Indo-european languages have the best results. In general, our system does not reach 40% UAS in Non-Indo-european languages, except in Hungarian and in Indonesian. In this regard, it is worth mentioning that Indonesian (with 51% UAS) has a Subject-Verb-Object word order similar to most European languages (Sneddon, 1996).

## 5 Discussion

The experiments performed in this paper provided some interesting results that claim for further research in cross-lingual parsing.

On the one hand, there are noticeable differences when parsing different treebanks of the same language, both using a rule-based system and harmonized supervised models. In this respect, it could be interesting to analyze the source of these variations, and *MetaRomance* could be

<sup>5</sup>Table 3 follows the language distinction provided by the shared task, even if we are aware that some linguistic varieties may be considered dialects of the same language (e.g., Galician as a variety of Portuguese, Old Bulgarian as a variety of Bulgarian, etc.).

useful for this purpose because it uses linguistically transparent rules based on PoS-tags.

On the other hand, the learning curve of a cross-lingual delexicalized model reinforces the idea that lexical features are required to obtain high-quality parsing results. In this respect, further experiments could compare this learning curve to lexicalized cross-lingual models, which seem to obtain good results in languages from the same linguistic family. Concerning *MetaRomance*, the addition of new rules (both lexicalized and without lexical information) could allow the parser to better analyze different languages.

Finally, and even if this is not a fair comparison, it is worth noting that *MetaRomance* obtained higher results in Romance languages than those achieved by UDP (Martínez Alonso et al., 2017). UDP is a training-free parser based on PageRank and a small set of head attachment rules, being more generic than *MetaRomance* (it can be applied to any language with more homogeneous results than our system). The differences on Romance languages vary between few decimals to more than 6% UAS, but the experiments were performed using different versions of the UD treebanks.<sup>6</sup>

## 6 Conclusions

This paper presented our submission to the *CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. The system, *MetaRomance*, is a fast rule-based parser suited to analyze Romance languages with no training data. It can be used on the top of several PoS-taggers such as *LinguaKit*, *FreeLing*, *TreeTagger*, or in a *CoNLL-U* file processed by tools such as *UDPipe*.

This cross-lingual parser contains 150 rules based on PoS-tags patterns, implemented by a linguist in about 12 hours. The *MetaRomance* grammar was written in *DepPattern*, a formalism that allows experts to easily modify and increase the rules to cover more syntactic phenomena.

Several experiments showed that a simple system such as the proposed in this paper can analyze in a uniform way different treebanks of Romance languages (and also from other linguistic families). Furthermore, a preliminary experiment on cross-lingual delexicalized parsing of Romance languages suggested that lexical features

are needed to increase the parsing performance. Lexical information can be added both to supervised systems and to our rule-based approach.

The grammar provided by *MetaRomance* was also used to present a classification of all the languages of the shared task datasets ranked by their syntactic distance with respect to Romance languages.

## Acknowledgments

This work has received financial support from a 2016 BBVA Foundation Grant for Researchers and Cultural Creators, the *TelePares* project (MINECO, ref:FFI2014-51978-C2-1-R), the Consellería de Cultura, Educación e Ordenación Universitaria (accreditation 2016-2019, ED431G/08), the European Regional Development Fund (ERDF), and from a *Juan de la Cierva-formación* grant (FJCI-2014-22853).

## References

- Željko Agić, Anders Johannsen, Barbara Plank, Héctor Martínez Alonso, Natalie Schluter, and Anders Søgaard. 2016. Multilingual projection for parsing truly low-resource languages. *Transactions of the Association for Computational Linguistics* 4:301–312.
- S. Ait-Mokhtar, J-P Chanod, and C. Roux. 2002. Robustness beyond Shallowness: Incremental Deep Parsing. *Natural Language Engineering* 8(2/3):121–144.
- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Cris Dyer, and Noah A. Smith. 2016. Many Languages, One Parser. *Transactions of the Association for Computational Linguistics* 4:431–444.
- Xavier Carreras. 2007. Experiments with a higher-order projective dependency parser. In *Proceedings of the CoNLL Shared Task Session of EMNLP-CoNLL 2007*. Association for Computational Linguistics, Prague, pages 957–961.
- Pablo Gamallo. 2015. Dependency parsing with compression rules. In *Proceedings of the 14th International Workshop on Parsing Technology (IWPT 2015)*. Association for Computational Linguistics, Bilbao, Spain, pages 107–117.
- Pablo Gamallo and Isaac González. 2011. A grammatical formalism based on patterns of part-of-speech tags. *International Journal of Corpus Linguistics* 16(1):45–71.
- Pablo Gamallo, José Ramon Pichel, and Iñaki Alegria. 2017. From language identification to language distance. *Physica A* 484:162–172.

<sup>6</sup>At this moment it is not possible to perform a better comparison of both systems, because UDP works with UDv1.2 and *MetaRomance* with UDv2.0.

- Kuzman Ganchev, Jennifer Gillenwater, and Ben Taskar. 2009. Dependency grammar induction via bitext projection constraints. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*. Association for Computational Linguistics, volume 1, pages 369–377.
- Marcos Garcia and Pablo Gamallo. 2010. Análise morfo-sintáctica para português europeu e galego: Problemas, soluções e avaliação. *Linguamática* 2(2):59–67.
- Marcos Garcia and Pablo Gamallo. 2015. Yet another suite of multilingual NLP tools. In *Languages, Applications and Technologies*. Springer, Switzerland, volume 563 of *Communications in Computer and Information Science*, pages 65–75. Revised Selected Papers of the Symposium on Languages, Applications and Technologies (SLATE 2015).
- Carlos Gómez-Rodríguez and Daniel Fernández-González. 2012. Dependency parsing with undirected graphs. In *13th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*. Avignon, France, pages 66–76.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering* 11(03):311–325.
- Héctor Martínez Alonso, Željko Agić, Barbara Plank, and Anders Søgaard. 2017. Parsing universal dependencies without training. In *15th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*. Valencia, Spain, pages 229–239.
- André F. T. Martins, Noah A. Smith, Eric P. Xing, Pedro M. Q. Aguiar, and Mário A. T. Figueiredo. 2010. Turboparsers: Dependency parsing by approximate variational inference. In *Empirical Methods in Natural Language Processing (EMNLP'10)*. Boston, USA.
- Ryan McDonald, Joakim Nivre, Yvonne Quirmbach-brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu, and Castelló Jungmee Lee. 2013. Universal dependency annotation for multilingual parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL 2013)*. Association for Computational Linguistics, Sofia, pages 92–97.
- Ryan McDonald and Fernando Pereira. 2006. Online Learning of Approximate Dependency Parsing Algorithms. In Association for Computational Linguistics, editor, *Proceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics (EACL 2006)*. Trento, pages 81–88.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source Transfer of Delexicalized Dependency Parsers. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2011)*. Association for Computational Linguistics, Edinburgh, United Kingdom, pages 62–72.
- Joakim Nivre. 2004. Incrementality in deterministic dependency parsing. In *ACL Workshop on Incremental Parsing: Bringing Engineering and Cognition Together*. Association for Computational Linguistics, Barcelona, pages 50–57.
- Joakim Nivre, Željko Agić, Lars Ahrenberg, et al. 2017. Universal dependencies 2.0 – CoNLL 2017 shared task development and test data. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University. <http://hdl.handle.net/11234/1-2184>.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, Reut Tsarfaty, and Daniel Zeman. 2016. Universal Dependencies v1: A multilingual treebank collection. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*. European Language Resources Association, Portorož, Slovenia, pages 1659–1666.
- Kemal Oflazer. 2003. Dependency parsing with an extended finite-state approach. *Computational Linguistics* 29(4):515–544.
- Lluís Padró and Evgeny Stanilovsky. 2012. FreeLing 3.0: Towards Wider Multilinguality. In *Proceedings of the Language Resources and Evaluation Conference (LREC 2012)*. European Language and Resources Association, Istanbul, Turkey, pages 2473–2479.
- Martin Potthast, Tim Gollub, Francisco Rangel, Paolo Rosso, Efstathios Stamatatos, and Benno Stein. 2014. Improving the reproducibility of PAN’s shared tasks: Plagiarism detection, author identification, and author profiling. In Evangelos Kanoulas, Mihai Lupu, Paul Clough, Mark Sanderson, Mark Hall, Allan Hanbury, and Elaine Toms, editors, *Information Access Evaluation meets Multilinguality, Multimodality, and Visualization. 5th International Conference of the CLEF Initiative (CLEF 14)*. Springer, Berlin Heidelberg New York, pages 268–299.
- Rudolf Rosa, Daniel Zeman, David Mareček, and Zdeněk Žabokrtský. 2017. Slavic forest, norwegian wood. In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial4)*. Association for Computational Linguistics, Stroudsburg, PA, USA, pages 210–219.
- Helmut Schmid. 1994. Probabilistic Part-of-Speech Tagging Using Decision Trees. In *International Conference on New Methods in Language Processing*. Manchester, pages 154–163.



- James Neil Sneddon. 1996. *Indonesian Reference Grammar*. Allen and Unwin, St. Leonards, Australia.
- Milan Straka, Jan Hajič, and Jana Straková. 2016. UD-Pipe: trainable pipeline for processing CoNLL-U files performing tokenization, morphological analysis, POS tagging and parsing. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC 2016)*. European Language Resources Association, Portorož, Slovenia.
- Oscar Täckström, Ryan McDonald, and Jakob Uszkoreit. 2012. Cross-lingual word clusters for direct transfer of linguistic structure. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human language technologies (NAACL - HLT 2012)*. Association for Computational Linguistics, pages 477–487.
- Jörg Tiedemann. 2017. Cross-lingual dependency parsing for closely related languages. In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2017)*. Association for Computational Linguistics, Valencia, pages 131–136.
- Jörg Tiedemann and Željko Agić. 2016. Synthetic Treebanking for Cross-Lingual Dependency Parsing. *Journal of Artificial Intelligence Research (JAIR)* 55:209–248.
- David Vilares, Miguel A. Alonso, and Carlos Gómez-Rodríguez. 2016. One model, two languages: training bilingual parsers with harmonized treebanks. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL 2016)*. Berlin, Germany, volume 2, pages 425–431.
- Hiroyasu Yamada and Yuji Matsumoto. 2003. Statistically Dependency Analysis with Support Vector Machines. In *Proceedings of 8th International Workshop on Parsing Technologies (IWPT)*. pages 195–206.
- Marcos Zampieri, Shervin Malmasi, Nikola Ljubešić, Preslav Nakov, Ahmed Ali, Jörg Tiedemann, Yves Scherrer, and Noëmi Aepli. 2017. Findings of the VarDial Evaluation Campaign 2017. In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*. Valencia.
- Daniel Zeman, Martin Popel, Milan Straka, Jan Hajič, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, Francis Tyers, Elena Badmaeva, Memduh Gökürmak, Anna Nedoluzhko, Silvie Cinková, Jan Hajič jr., Jaroslava Hlaváčová, Václava Kettnerová, Zdeňka Urešová, Jenna Kanerva, Stina Ojala, Anna Mäsilä, Christopher Manning, Sebastian Schuster, Siva Reddy, Dima Taji, Nizar Habash, Herman Leung, Marie-Catherine de Marneffe, Manuela Sanguinetti, Maria Simi, Hiroshi Kanayama, Valeria de Paiva, Kira Drohanova, Héctor Martínez Alonso, Hans Uszkoreit, Vivien Macketanz, Aljoscha Burchardt, Kim Harris, Katrin Marheinecke, Georg Rehm, Tolga Kayadelen, Mohammed Attia, Ali Elkahky, Zhuoran Yu, Emily Pitler, Saran Lertpradit, Michael Mandl, Jesse Kirchner, Hector Fernandez Alcalde, Jana Strnadova, Esha Banerjee, Ruli Manurung, Antonio Stella, Atsuko Shimada, Sookyoung Kwak, Gustavo Mendonça, Tatiana Lando, Rattima Nitisaroj, and Josie Li. 2017. CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics.
- Daniel Zeman and Philip Resnik. 2008. Cross-Language Parser Adaptation between Related Languages. In *Proceedings of the Workshop on NLP for Less Privileged Language at the 3rd International Joint Conference on Natural Language Processing (IJCNLP 2008)*. Asian Federation of Natural Language Processing, Hyderabad, pages 35–42.