

Exploiting the Contribution of Morphological Information to Parsing: the BASQUE_TEAM system in the SPRML'2013 Shared Task

Iakes Goenaga, Nerea Ezeiza

IXA NLP Group

Faculty of Computer Science

Univ. of the Basque Country UPV/EHU

iakesg@gmail.com, n.ezeiza@ehu.es

Koldo Gojenola

IXA NLP Group

Technical School of Engineering, Bilbao

Univ. of the Basque Country UPV/EHU

koldo.gojenola@ehu.es

Abstract

This paper presents a dependency parsing system, presented as BASQUE_TEAM at the SPMRL'2013 Shared Task, based on the analysis of each morphological feature of the languages. Once the specific relevance of each morphological feature is calculated, this system uses the most significant of them to create a series of analyzers using two freely available and state of the art dependency parsers, MaltParser and Mate. Finally, the system will combine previously achieved parses using a voting approach.

1 Introduction

Morphologically rich languages present new challenges, as the use of state of the art parsers for more configurational and non-inflected languages like English does not reach similar performance levels in languages like Basque, Greek or Turkish (Nivre et al., 2007). Using morphological information as features in parsing has been a commonly used method for parsing MRLs (Tsarfaty et al., 2010). In some cases the effect of this information is positive but in others it does not help or causes a negative effect.

In most of the work on dependency parsing, the specific relevance of each morphological feature in the final result is unknown. The authors include all the morphological features¹ in their systems with the aim of taking advantage of the diversity of the used information. This approach commonly produces very good results but they are not always the best ones (see table 2).

On the other hand, some authors have made experiments to specify which is the real impact of the morphological features. Ambati et al. (2010) explore ways of integrating local morphosyntactic features into Hindi dependency parsing. They experiment with different sets of features on a graph-based and a transition-based dependency parser. They show that using some morphological features (root, case, and suffix) outperforms a baseline using POS as the only feature, with both gold and predicted settings .

Bengoetxea and Gojenola (2010) make use of MaltParser's feature configuration file to take advantage of morphological features in parsing with gold data. Their experiments show that case and subordination type considerably increase parsing accuracy.

¹That is, they treat all the morphological features in the same way in the feature specification, and let the learning algorithms decide the weight assigned to each one.

Marton et al. (2013) also explore which morphological features could be useful in dependency parsing of Arabic. They observe the effect of features by adding them one at a time separately and comparing the outcomes. Experiments showed that when gold morphology is provided, case markers help the most, whereas when the morphology is automatically predicted the outcome is the opposite: using case harms the results the most. When features are combined in a greedy heuristic, using definiteness, person, number, and gender information improves accuracy.

Similarly, Seeker and Kuhn (2013) also determine that the use of case is specially relevant for parsing, demonstrating that morpho-syntactic constraints can delimit the search space of a statistical dependency parser to outperform state-of-the-art baselines for Czech, German and Hungarian.

Following this line of research, our first step will be to determine which is the concrete value of each feature on dependency parsing, adding one of the morphological features at a time starting with an empty *FEATS* column.

Çetinoğlu and Kuhn (2013) have shown that some parsers tend to improve the results when swapping or replacing *POS* by some of the morphological features. They have made use of the METU-Sabancı Turkish Treebank (Oflazer et al., 2003) for training and the ITU validation set (Eryigit, 2007) for testing. In their work, it is observed that moving *CASE* to the *POS* field helps with a 0.3% LAS absolute increase in the gold pipeline settings and using *CASE* instead of nominal *POS* improves the labelled accuracy by 0.3% absolute for the training set.

These experiments suggest that in some way the parser is not making an optimal use of all the available morpho-syntactic information, and that the parser algorithm (or the feature

specification for the learning phase) is geared towards *POS* and *CPOS*, giving a lower status to other types of information. Although this strategy is good in general, it seems that, at least for some languages, specific features (e.g. *CASE*) are crucial in obtaining a high parsing performance. Taking these ideas into consideration, we will work on three different approaches:

- We will experiment the effect of using only the best three morphological features in the *FEATS* column (see table 1), compared to working with the full set of morpho-syntactic features. This can have the effect of speeding the learning and parsing processes, as the number of features can be smaller. On the other hand, the elimination of non-relevant features can also help to improve the parser's results, because some features can even be detrimental for parsing.
- Following Çetinoğlu and Kuhn (2013), once our system resolves which feature is the most significant, it will be used to replace the *POS* and *CPOS* fields one by one and we will test the effect of these variants on the parsers. Finally, we will also try right-to-left versions of those 3 variants (baseline, and replacing *POS* and *CPOS*) completing a set of 6 different parsers.
- Finally, we will experiment the combination of the different or parsers with a voting approach (Hall et al., 2010) using the MaltBlender tool².

All of the experiments will be performed on automatically predicted *POS* and morphosyntactic data, taking the tags given in the Shared

²<http://w3.msi.vxu.se/users/jni/blend/>

Task data, that is, we will not made use of any specifically trained morphological tagger.

In the rest of this paper we will first present the resources we have used to carry out our experiments in section 2, followed by a study of the contribution of the morphological information to parsing in section 3 and the effect of this information on the individual parsers in subsection 4.1. The final results of the best parser combinations are showed in subsection 4.2 and the main conclusions of the work in section 5.

2 Resources

This section will describe the main resources that have been used in the experiments. Subsection 2.1 will describe the languages we have used in our experiments, subsection 2.2 will explain the parsers we use, while subsection 2.3 will present briefly the *MaltBlender* tool.

2.1 Selected Languages

Although the *SPMRL'2013 Shared Task* (Seddah et al., 2013) offers the opportunity to parse nine morphologically rich languages, to carry out our experiments we have selected five of them, due in part to time constraints, but also taking into account the relevance of the morpho-syntactic information (*FEATS* column, see table 1) . The selected five languages are: Basque (Aduriz et al., 2003), French (Abeillé et al., 2003), German (Seeker and Kuhn, 2012), Hungarian (Vincze et al., 2010) and Swedish (Nivre et al., 2006).

2.2 Parsers

We have made use of MaltParser (Nivre et al., 2007b) and Mate (Bohnet and Nivre, 2012), two state of the art dependency parsers³ repre-

³Due to time constraints, we did not have enough time to experiment with other options such as the MST parser or the EasyFirst parser.

senting the dominant approaches in data-driven dependency parsing, and that have been successfully applied to typologically different languages and treebanks.

MaltParser is a representative of local, greedy, transition-based dependency parsing models, where the parser obtains deterministically a dependency tree in a single pass over the input using two data structures: a stack of partially analyzed items and the remaining input sequence. To determine the best action at each step, the parser uses history-based feature models and discriminative machine learning. The specification of the learning configuration can include any kind of information (such as word-form, lemma, category, subcategory or morphological features). We will use one of its latest versions (MaltParser version 1.7).

To fine-tune Maltparser we have used MaltOptimizer (Ballesteros and Nivre, 2012a; Ballesteros and Nivre, 2012b). This tool is an interactive system that first performs an analysis of the training set in order to select a suitable starting point for optimization and then guides the user through the optimization of parsing algorithm, feature model, and learning algorithm. Empirical evaluation on data from the CoNLL 2006 and 2007 shared tasks on dependency parsing shows that MaltOptimizer consistently improves over the baseline of default settings and sometimes even surpasses the result of manual optimization.

The Mate parser (Bohnet and Nivre, 2012) is a development of the algorithms described in (Carreras, 2007; Johansson and Nugues, 2008). It basically adopts the second order maximum spanning tree dependency parsing algorithm. In particular, this parser exploits a hash kernel, a new parallel parsing and feature extraction algorithm that improves accuracy as well as parsing speed (Bohnet, 2010).

2.3 Parser Combinations

The MaltBlender tool makes a two-stage optimization of the result of several parser outcomes, based on the work of Sagae and Lavie (2006), and it was used for the first time for the ten languages in the multilingual track of the CoNLL 2007 shared task on dependency parsing (Hall et al., 2010). The first stage consists in tuning several single-parser systems. The second stage consists in building an ensemble system that will combine the different parsers. When this system was evaluated on the official test sets at the CoNLL 2007 shared task, the ensemble system significantly outperformed the single-parser system and achieved the highest average labelled attachment score of all participating systems.

3 Contribution of Morphological Information to Parsing

We examined the effect of each type of morphological information, contained in the *FEATS* column, to investigate their overall contribution to parsing. This will help us to determine which are the most relevant features for parsing. To carry out this task we have used the Mate parser, due to lack of time for testing, and also taking into consideration that it gives better results than MaltParser for all the languages's baselines. Firstly, we will obtain the baseline for each language parsing the files with an empty *FEATS* column. This baseline will help us to determine the contribution of each morphological feature to parsing. Next, we trained the parsers using one feature at a time obtaining as many results as features for each language. Table 1 shows the effect of each information on the Mate parser.

In this table we can observe that Basque is one of the most sensitive languages regarding

the influence of its features. Using case (KAS) as a unique feature improves the labelled attachment score over using an empty *FEATS* column by almost 5.7%. The next two better features are number (NUM) and type of subordinate sentence (ERL). They help with a 1.1% and 0.6% increase, respectively. The rest of the features do not contribute much in isolation, with a maximum of 0.2%. On the other hand, including all the features results in an improvement of 6.5%.

If we analyze the results for French we see that, in contrast to Basque, the influence of the features on the parser is minimum. The most significant feature is gender (g), which helps with a 0.1% increase. With respect to the improvement using the other features, although they do not provide big increases all of them contribute positively. In closing, including all the features we obtain a 84.6% labelled attachment score with a 0.4% improvement over not using any features.

As with French, the German morphological features provide small increases. The most two significant features are case and gender, which obtain increases of 0.2%, 0.13%, respectively. It is interesting to observe how including all the features we obtain worse results than using only the case, although the difference is not significant. That could occur due to the weak influence of its features in the final result and the negative influence of some of them.

Hungarian is the language which offers more features, 14 altogether. This language, in line with Basque, tends to vary significantly its labelled attachment score depending on the used morphological feature. If we focus on the three most significant features, the case (Cas) helps with a 5.6% increase, person of possessor (PerP) with a 1%, while number of possessor helps with a 0.9%. The grammatical subcate-

Basque	French	German	Hungarian	Swedish
all feats 83.0	all feats 84.6	all feats 91.0	all feats 82.8	all feats 76.7
no feats 76.5	no feats 84.2	no feats 90.9	no feats 75.3	no feats 76.9
KAS 82.2	g 84.3	case 91.0	Cas 80.9	verbform 77.0
NUM 77.7	n 84.3	gender 91.0	PerP 76.3	definiteness 76.8
ERL 77.1	p 84.3	number 90.9	NumP 76.3	degree 76.8
DADUDIO 76.8	c 84.2	person 90.9	SubPOS 75.9	case 76.8
NORK 76.7	m 84.2	tense 90.9	Def 75.7	number 76.3
MDN 76.6	s 84.2	degree 90.8	Num 75.7	perfectform 76.3
NOR 76.6	t 84.2	mood 90.8	PerP 75.7	abbrv 76.3
ASP 76.4	-	-	Mood 75.5	mood 76.2
NORI 76.2	-	-	NumPd 75.4	pronounform 76.1
ADM 76.5	-	-	Coord 75.3	gender 76.0
-	-	-	Form 75.3	-
-	-	-	Tense 75.3	-
-	-	-	Type 75.3	-
-	-	-	Deg 75.0	-

Table 1: The effect of each feature sorted by language (MATE parser)

gory within the main part of speech (SubPOS) improves the baseline in a 0.6% and the number and person in a 0.4%. The remaining features do not contribute very appreciatively even obtaining negative results. Including all the features we obtain a labelled attachment score of 82.83%. That means the real contribution of all the features is 7.5%, this improvement being the most important among all the used languages.

In common with French and German, the Swedish morphological features do not seem to help the parsers to achieve significant improvements in terms of *LAS*. However, we can observe some interesting phenomena. While in the other languages the case is one of the best features, in Swedish it does not help, achieving a negative result. In general, excluding the verb form (verbform), all the features obtain negative results with respect to not using any feature. In this scenario it is not surprising to verify that including all the features does not help

the Mate parser. Having said this, the best three features are the verb form (verbform), definiteness (definiteness) and degree (degree).

4 Testing the Effect of Different Morphosyntactic features on parsers

We examined the effect of the most significant morphological features, examined in the previous step, to investigate their overall contribution to parsing. For this task, we created three variants for each parser, apart from the baseline using all the morphosyntactic features. We obtain these variants by: i) using the most 3 relevant features in the *FEATS* column (see table 1 in previous section), ii) moving the most relevant feature for each language to the *POS* column and iii) moving the most relevant feature to the *CPOS* column. Next, we have tested parser combinations including all the baselines and their variants in subsection 4.2.

	Basque	French	German	Hungarian	Swedish
Baselines					
<i>All – Feats_{Malt}</i>	80.0	79.9	87.6	77.3	73.4
<i>All – Feats_{Mate}</i>	83.0	84.6	91.0	82.3	76.7
Left2right					
<i>3 – best_{Malt}</i>	79.9	79.9	87.6	75.9	73.4
<i>CPOS – best_{Malt}</i>	80.3	79.7	87.5	76.6	72.9
<i>POS – best_{Malt}</i>	78.7	78.7	86.6	77.2	72.8
<i>3 – best_{Mate}</i>	83.4	84.3	90.8	82.4	76.6
<i>CPOS – best_{Mate}</i>	82.7	84.3	91.0	82.7	76.8
<i>POS – best_{Mate}</i>	82.2	83.4	90.5	82.5	76.5
Right2left					
<i>3 – best_{Malt}</i>	80.1	78.9	86.9	75.3	69.3
<i>CPOS – best_{Malt}</i>	80.0	79.0	86.7	76.6	69.3
<i>POS – best_{Malt}</i>	81.0	77.8	85.4	74.9	70.2
<i>3 – best_{Mate}</i>	83.3	84.3	90.9	82.1	76.5
<i>CPOS – best_{Mate}</i>	83.1	84.6	91.0	82.6	77.0
<i>POS – best_{Mate}</i>	81.6	83.5	90.6	82.4	76.4

Table 2: Testing the effect of features on MaltParser and Mate

4.1 Individual Parsers

Table 2 shows the effect of each information on both parsers, Maltparser and Mate parser. If we analyze the results on Basque, the difference between the two parsers is noticeable, as Mate obtains on average a 3 point improvement with respect to MaltParser. A similar difference occurs on all the used languages. The best *LAS* in Basque is acquired using the 3 best features in the *FEATS* column with the Mate parser (83.4%). On a comparison with the *LAS* obtained by the Mate baseline (All-Feats), that means a 0.4 improvement. Regarding Maltparser’s results for Basque, we get the best *LAS* (81.0%) moving the best feature (case) to *POS* in its right-to-left version, increasing the *LAS* baseline (All-Feats) by 1.0. We notice that Maltparser and Mate tend to improve their baseline scores using some of the presented variants.

On the other hand, the best score for French

is obtained using the baseline (All-Feats and the Mate parser, 84,6%). Contrary to Basque, in French, although some of the used variants achieve similar scores with respect to their baselines (All-Feats), they do not give noticeable increases. The unique variant that equals its baseline (79,9%) is *3 – best_{Malt}* using the left-to-right version and the three best features (gender, number and person) in the *FEATS* column using Maltparser.

With respect to German, the only variant that equals the baseline is *CPOS – best_{Mate}* with 91.0% *LAS*. . If we focus on Maltparser’s (MaltOptimizer) scores, we get the best result among the variants with *3 – best_{Malt}* (87.6%) using the left-to-right version. The variants do not improve Maltparser’s baseline.

Although some of the Hungarian variant scores are very similar to their baselines, they give some improvements over the baseline. The best two results on the Mate parser are 82.7% and 82.6%. We obtain the first score mov-

ing the best feature (case) to *CPOS* in its left-to-right version, and the second one using the same configuration in its right-to-left version. The best two scores on Maltparser without taking the baseline into account are 77.2% and 76.6%, obtained when moving the best feature to *POS* and moving the best feature to *CPOS* in its right-to-left version, respectively.

The best two results for Swedish on the Mate parser are 77.0% and 76.8%. We get the first result moving the best feature (verbform) to *CPOS* in its right-to-left version and the second one in its standard version. These two results are the only variants that improve the baseline (76.7% *LAS*) with a 0.30 and 0.17 increase, respectively. On the other hand, if we focus on Maltparser, the variants do not improve the baseline (73.4% *LAS*) where the best two results are 73.4% and 72.9% *LAS*. For the best result we use the three best features (verbform, definiteness and degree) in the *FEATS* column, while for the second one the best feature (verbform) has been moved to *CPOS*.

Despite that only the Basque and Swedish variants have been able to significantly improve their baselines, in the next subsection we present a combination system expecting to take advantage on the variety of the parsed files (Surdeanu and Manning, 2010).

4.2 Parser Combinations

Although in several cases the use of specific morphosyntactic information does not give noticeable increases, we also tested the effect on parser combinations. Table 3 presents the result of combining the extended parsers with the baselines (using all the features) obtained in individual parsers. The table shows that the Basque language has achieved the biggest increase. Parser combination in Basque helps with an improvement of 3.2 with respect to the

Mate baseline. Contrary to Basque, French is the language that has obtained the smallest increases in parser combination if we compare it with the Mate (highest) parser baseline. The combined system improves the Mate parser baseline by 0.5. Parser combination in German gives a 0.8 increase with respect to the best single parser (Mate, 91.0). Our system achieves a 1.3 increase for Hungarian with respect to the Mate parser’s baseline. Finally, if we focus on Swedish, the parser combination helps with a 1.4 increase with respect to the Mate parser.

After examining the parsers involved in parser combinations we noticed that there are always several variants included in the best parser combinations, although the only variant that appears in all the best parser combinations is *CPOS – best_{Mate}* in its left-to-right version. Taking into account that the most relevant feature for Basque, German and Hungarian is the case, it would be interesting to use the *CPOS – case_{Mate}* variant for other languages. Finally, the presented results suggest that the introduced variants contribute positively on parsing and they help to improve the scores obtained by the base parsers.

5 Conclusion and Future Work

We have presented a combined system that was designed after analyzing the relevance of the morphological features in order to take advantage on the effect of those features on some parsers. In general the improvements have been noticeable, specially for Basque. We can point out some interesting avenues for research:

- Use of new parsing algorithms for testing the effect of different morphological features. The results of this work show that the used techniques are specially useful for languages where the *FEATS* column,

	Basque	French	German	Hungarian	Swedish
MaltParser baseline	80.0	79.9	87.6	77.3	73.4
Mate parser baseline	83.0	84.6	91.0	82.8	76.7
Parser combination	86.2	85.1	91.8	84.1	78.1

Table 3: Results of parser combinations

containing morpho-syntactic information, gives the biggest increments with respect to not using the features, like Basque and Hungarian. We expect that similar improvements could be obtained for languages like Turkish or Czech, which share many characteristics with Basque and Hungarian.

- Experimenting different models for parser combinations using new parsers. Several of the parser variants we have used give only slight modifications over the base algorithms, even though when combined they give significant increases. Widening the spectrum of parsers and adding new algorithms can imply an important boost in parser combination.
- Application to the rest of the languages of the *SPMRL 2013 Shared Task*: Korean, Hebrew, Arabic and Polish.

Acknowledgements

This research was supported by the Department of Industry of the Basque Government (IT344-10, S PE11UN114), the University of the Basque Country (GIU09/19) and the Spanish Ministry of Science and Innovation (MICINN, TIN2010-20218).

References

Anne Abeillé, Lionel Clément, and François Toussenet. 2003. Building a treebank for french.

In Anne Abeillé, editor, *Treebanks*. Kluwer, Dordrecht.

I. Aduriz, M. J. Aranzabe, J. M. Arriola, A. Atutxa, A. Díaz de Ilarraza, A. Garmendia, and M. Oronoz. 2003. Construction of a Basque dependency treebank. pages 201–204.

Bharat Ram Ambati, Samar Husain, Sambhav Jain, Dipti Misra Sharma, and Rajeev Sangal. 2010. Two methods to incorporate local morphosyntactic features in hindi dependency parsing. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 22–30.

Miguel Ballesteros and Joakim Nivre. 2012a. Maltoptimizer: A system for maltparser optimization. In *LREC*, pages 2757–2763.

Miguel Ballesteros and Joakim Nivre. 2012b. Maltoptimizer: an optimization tool for maltparser. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 58–62.

Kepa Bengoetxea and Koldo Gojenola. 2010. Application of different techniques to dependency parsing of basque. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*, pages 31–39.

Bernd Bohnet and Joakim Nivre. 2012. A transition-based system for joint part-of-speech tagging and labeled non-projective dependency parsing. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1455–1465.

Bernd Bohnet. 2010. Very high accuracy and fast dependency parsing is not a contradiction. In

- Proceedings of the 23rd International Conference on Computational Linguistics*, pages 89–97.
- Xavier Carreras. 2007. Experiments with a higher-order projective dependency parser. In *Proceedings of the CoNLL Shared Task Session of EMNLP-CoNLL 2007*, Prague, Czech Republic, June.
- Özlem Çetinoğlu and Jonas Kuhn. 2013. Towards joint morphological analysis and dependency parsing of turkish. In *Proceedings of the Second International Conference on Dependency Linguistics (DepLing 2013)*, pages 23–32, Prague, Czech Republic, August. Charles University in Prague, Matfyzpress, Prague, Czech Republic.
- Gülsen Eryigit. 2007. Itu validation set for metusabancı turkish treebank. URL: <http://www3.itu.edu.tr/gulsenc/papers/validationset.pdf>.
- Johan Hall, Jens Nilsson, and Joakim Nivre. 2010. Single malt or blended? a study in multilingual parser optimization. In *Trends in Parsing Technology*, pages 19–33. Springer.
- Richard Johansson and Pierre Nugues. 2008. Dependency-based syntactic-semantic analysis with propbank and nombank. In *Proceedings of the Twelfth Conference on Computational Natural Language Learning*, pages 183–187.
- Yuval Marton, Nizar Habash, and Owen Rambow. 2013. Dependency parsing of modern standard arabic with lexical and inflectional features. *Computational Linguistics*, 39(1):161–194.
- Joakim Nivre, Jens Nilsson, and Johan Hall. 2006. Talbanken05: A Swedish treebank with phrase structure and dependency annotation. In *Proceedings of LREC*, pages 1392–1395, Genoa, Italy.
- Joakim Nivre, Johan Hall, Sandra Kübler, Ryan McDonald, Jens Nilsson, Sebastian Riedel, and Deniz Yuret. 2007. The CoNLL 2007 shared task on dependency parsing. In *Proceedings of the CoNLL Shared Task Session of EMNLP-CoNLL 2007*, Prague, Czech Republic, June.
- Joakim Nivre, Johan Hall, Jens Nilsson, Atanas Chanev, Gülsen Eryigit, Sandra Kübler, Svetoslav Marinov, and Erwin Marsi. 2007b. Malt-parser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135.
- Kemal Oflazer, Bilge Say, Dilek Zeynep Hakkani-Tür, and Gökhan Tür. 2003. Building a turkish treebank. *Building and Exploiting Syntactically-annotated Corpora*.
- Kenji Sagae and Alon Lavie. 2006. Parser combination by reparsing. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the ACL*.
- Djamé Seddah, Reut Tsarfaty, Sandra Kübler, Marie Candito, Jinho Choi, Richárd Farkas, Jennifer Foster, Iakes Goenaga, Koldo Gojenola, Yoav Goldberg, Spence Green, Nizar Habash, Marco Kuhlmann, Wolfgang Maier, Joakim Nivre, Adam Przepiorkowski, Ryan Roth, Wolfgang Seeker, Yannick Versley, Veronika Vincze, Marcin Woliński, Alina Wróblewska, and Eric Villemonte de la Clérgerie. 2013. Overview of the spmrl 2013 shared task: A cross-framework evaluation of parsing morphologically rich languages. In *Proceedings of the 4th Workshop on Statistical Parsing of Morphologically Rich Languages: Shared Task*, Seattle, WA.
- Wolfgang Seeker and Jonas Kuhn. 2012. Making Ellipses Explicit in Dependency Conversion for a German Treebank. In *Proceedings of the 8th International Conference on Language Resources and Evaluation*, pages 3132–3139, Istanbul, Turkey. European Language Resources Association (ELRA).
- Wolfgang Seeker and Jonas Kuhn. 2013. Morphological and syntactic case in statistical dependency parsing. *Computational Linguistics*, 39(1):23–55.
- Mihai Surdeanu and Christopher D. Manning. 2010. Ensemble models for dependency parsing: Cheap and good? In *Proceedings of the North American Chapter of the Association for Computational Linguistics Conference (NAACL-2010)*, Los Angeles, CA, June.

Reut Tsarfaty, Djam Seddah, Yoav Goldberg, Sandra Kübler, Marie Candito, Jennifer Foster, Yannick Versley, Ines Rehbein, and Lamia Tounsi. 2010. Statistical parsing of morphologically rich languages (spmrl) what, how and whither. In *Proceedings of the NAACL HLT 2010 First Workshop on Statistical Parsing of Morphologically-Rich Languages*.

Veronika Vincze, Dóra Szauter, Attila Almási, György Móra, Zoltán Alexin, and János Csirik. 2010. Hungarian dependency treebank. In *LREC*.