

# Cross-Lingual Dependency Parsing for Closely Related Languages – Helsinki’s Submission to VarDial 2017

**Jörg Tiedemann**

Department of Modern Languages

University of Helsinki

first.lastname@helsinki.fi

## Abstract

This paper describes the submission from the University of Helsinki to the shared task on cross-lingual dependency parsing at VarDial 2017. We present work on annotation projection and treebank translation that gave good results for all three target languages in the test set. In particular, Slovak seems to work well with information coming from the Czech treebank, which is in line with related work. The attachment scores for cross-lingual models even surpass the fully supervised models trained on the target language treebank. Croatian is the most difficult language in the test set and the improvements over the baseline are rather modest. Norwegian works best with information coming from Swedish whereas Danish contributes surprisingly little.

## 1 Introduction

Cross-lingual parsing is interesting as a cheap method for bootstrapping tools in a new language from resources in another language. Various approaches have been proposed in the literature, which can mainly be divided into data transfer (i.e. annotation projection, e.g. (Hwa et al., 2005)) and model transfer approaches (e.g. delexicalized models such as (McDonald et al., 2013)). We will focus on data transfer in this paper using annotation projection and machine translation to transform source language treebanks to be used as training data for dependency parsers in the target language. Our previous work has shown that these techniques are quite robust and show better performance than simple transfer models based on delexicalized parsers (Tiedemann and Agić, 2016). This is especially true for real-world test

cases in which part-of-speech (PoS) labels are predicted instead of given as gold standard annotation while testing the parsing models (Tiedemann, 2015a).

Cross-lingual parsing assumes strong syntactic similarities between source and target language which can be seen at the degradation of model performance when using distant languages such as English and Finnish (Tiedemann, 2015b). The task at VarDial, therefore, focuses on closely related languages, which makes more sense also from a practical point of view. Many pools of closely related languages and language variants exist and, typically, the support in terms of resources and tools is very biased towards one of the languages in such a pool. Hence, one can say that the task at VarDial simulates real-world cases using existing resources from the universal dependencies project (Nivre et al., 2016) and promotes the ideas for practical application development. The results show that this test is, in fact, not only a simulation but actually improves the results for one of the languages in the test set: Slovak. Cross-lingual models outperform the supervised upper bound, which is a great result in favor of the transfer learning ideas.

More details about the shared task on cross-lingual parsing at VarDial 2017 can be found in (Zampieri et al., 2017). In the following, we will first describe our methodology and the data sets that we have used, before jumping to the results and some discussions in relation to our main findings.

## 2 Methodology and Data

Our submission is based on previous work and basically applies models and techniques that have been proposed by (Hwa et al., 2005; Tiedemann, 2014; Tiedemann et al., 2014). We made very lit-




Figure 1: Annotation projection heuristics with *dummy* nodes: One-to-many alignments create dummy nodes that govern the linked target language tokens. Many-to-one alignments are resolved by removing links from lower nodes in the source language tree. Non-aligned source language tokens are covered by additional dummy nodes that take the same incoming and outgoing relations. The final picture to the right illustrates that dummy leaf nodes can safely be deleted and internal dummy nodes with single daughters can be removed by collapsing relations.

tle changes to the basic algorithms but emphasized a systematic evaluation of different methods and parameters which we tested on the development data provided by VarDial 2017. All our results are, hence, scored on data sets with predicted PoS labels. In particular, we used three different cross-lingual model types:

**Projection:** Annotation projection across word-aligned parallel corpora using the data sets provided by the workshop organizers (the subtitle corpora). Source language parts are tagged and parsed automatically with supervised taggers and parsers.

**PBSMT:** Treebank translation using a phrase-based model of statistical machine translation (SMT). Annotation are then projected from the original source language treebank to the translations to create a synthetic target language treebank. Alignment is taken directly from the translation model. The translation and language models are trained on the provided parallel corpora. No extra resources are used.

**SyntaxSMT:** Treebank translation using a tree-to-string hierarchical SMT model. Dependencies are transformed into constituency representations using the spans defined by the yield of each word in the sentence with respect to the dependency relations. Constituency labels are taken from the dependency relations and PoS labels are used as well for the labels of leaf nodes. After translation, we project the annotation of the source

language treebank using the same procedures as for the other two approaches. Translation models are trained on the provided parallel corpora with automatically parsed source language sentences.

There are various improvements and heuristics for the projection of dependency trees. We applied two of them: (1) *collapseDummy*, which deletes leaf nodes that are labeled as “dummy” and also removes dummy nodes with just one daughter node by collapsing the parent and daughter relations. (2) *noDummy*, which discards all sentences that still include dummy nodes after applying *collapseDummy*. Dummy nodes appear with the projection heuristics introduced by (Hwa et al., 2005), which we also use for handling non-one-to-one word alignments. For example, unaligned source language tokens are projected on dummy target nodes to ensure the proper connection of the projected dependency tree. This can lead to dummy leaf nodes that can be ignored or dummy nodes with single daughters, which can be removed by collapsing the relations to head and dependent. Figure 1 illustrates the projection heuristics and the *collapseDummy* procedures. More details and examples are given in (Tiedemann and Agić, 2016). Overall, *noDummy* leads to a drop in performance and, therefore, we do not consider those results in this paper. The differences were small and most problems arose with smaller data sets where the reduction of training data has a negative impact on the performance.

For annotation projection, we used various sizes of parallel data to test the impact of data on parsing

performance. The translation models are trained on the entire data provided by VarDial. Language models are simply trained on the target side of the parallel corpus.

We also tested cross-lingual models that exploit language similarities on the lexical level without translating or projecting annotation. The idea is similar to delexicalized models that are trained on generic features on the source language treebank, which are then applied to the target languages without further adaptation. With closely related languages, we can assume substantial lexical overlaps, which can be seen at the relative success of the second baseline in the shared task (also shown in Table 1). In particular, we used substrings such as prefixes (simulating simple stemming) and suffixes (capturing inflectional similarities) to add lexical information to delexicalized models. However, those models did not perform very well and we omit the results in this paper.

For training the parsers, we used mate-tools (Bohnet, 2010), which gave us significantly better results than UDPipe (Straka et al., 2016) without proper parameter optimization except for some delexicalized models. Table 1 compares the baseline models with the two different toolkits. We still apply UDPipe for PoS and morphological tagging using the provided tagger models for the target languages and similar ones trained on the UD treebanks for the source languages except for Czech, which did not work with standard settings due to the complexity of the tagset and limitations of the implementation of UDPipe. Instead, we apply Marmot (Müller and Schütze, 2015) for Czech, which also provides efficient model training for PoS and morphology.

The only “innovation” compared to our previous work is the inclusion of target language tagging on top of annotation projection. Earlier, we only used projected annotation even for PoS information. In this paper, we also test the use of target language taggers (which are part of the provided setup) to (i) over-rule projected universal PoS tags and (ii) add morphological information to the data. Especially the latter makes a lot of sense especially for highly-inflecting languages like Slovak and Croatian. However, the risk of this procedure is that noisy projection of dependency label may not fit well together with the tags created by independent tools that are probably less noisy and make different kinds of mistakes. This may mis-

lead the training algorithm to learn the wrong connections and we can see that effect in our experiments especially in connection with the tagging of universal PoS labels. This actually degrades the parsing performance in most cases. More details will be presented in the following section in connection with the results of our experiments.

### 3 Results

We considered all language pairs from the VarDial campaign and here we present the relevant results from our experiments. First of all, we need to mention that we created new baselines using the mate-tools to have fair comparisons of the cross-lingual models with respect to baseline approaches. The new figures (on development data) are given in Table 1. The same table also summarizes our basic results for all language pairs using the three approaches for data transfer as introduced in the previous section. All projections are made in collapseDummy mode as explained above.

Target Source	Croatian Slovenian	Slovak Czech	Norwegian Danish	Swedish
<i>UDPipe</i>				
supervised	74.27	70.27	78.10	
delex	53.93	53.66	54.54	56.71
cross	56.85	54.61	54.11	55.85
<i>mate-tools</i>				
supervised	79.68	71.89	81.37	
delex	53.39	55.80	50.07	56.27
cross	60.29	62.21	56.85	59.63
<i>Projected</i>				
100,000	58.82	60.29	57.19	63.03
500,000	59.86	62.23	57.58	64.61
1,000,000	<b>62.92</b>	63.57	57.82	64.59
PBSMT	60.81	<b>65.97</b>	57.87	65.96
SyntaxSMT	58.57	63.13	58.36	<b>66.31</b>

Table 1: Basic results of cross-lingual parsing models in terms of labeled attachment scores (LAS) on development data: Annotation projection on automatically parsed bitexts of varying sizes (projected: number of sentence pairs); treebank translation models (PBSMT and SyntaxSMT); compared to three baselines: delexicalized models (delex), source language models without adaptation (cross) and fully-supervised target language models (supervised).

The first observation is that all cross-lingual models beat the delexicalized baseline by a large margin. This is, at least, self-assuring and motivates further developments in the direction of an-

notation projection and treebank translation. Another observation is that Croatian is surprisingly hard to improve in comparison to the cross-lingual model that applies a parser for Slovenian without any adaptation.

Another surprise is the quality of the Norwegian models coming from Danish. Both languages are very close to each other especially in writing (considering that we use bokmål in our data sets for Norwegian). Projection and translation should work well and should at least be on-par with using Swedish as the source language. However, the differences are quite significant between Danish and Swedish as the source language and this points to some substantial annotation differences between Danish and the other two languages that must be the reason behind this mystery. This conclusion is even more supported by the results of the cross-lingual baseline model without adaptation, which should perform better for Danish as the lexical overlap is large, greater than the overlap with Swedish. Yet another indication for the annotation differences is the result of the delexicalized parsers. There is also a big gap between Danish and Swedish as the source language. The result of these experiments demonstrate the remaining difficulties of cross-linguistically harmonized data sets, which is a useful outcome on its own.

We can also see, that treebank translation works rather well. For most language pairs, the performance is better than for annotation projection but the differences are rather small in many cases. An exception is Croatian for which annotation projection on parallel corpora works best, whereas translation is on par with Slovenian models applied to Croatian data.

In contrast to our previous findings, we can also see that the amount of data that is useful for annotation projection is bigger. Our prior work indicated that small corpora of around 40,000 sentence pairs are sufficient and that the learning curve levels out after that (Tiedemann and Agić, 2016). In this paper, we see increasing model performance until around one million sentence pairs before the scores converge (additional runs confirm this, even though they are not reported in the paper). A reason for this behaviour is that we now rely on movie subtitles instead of sentences from the European parliament proceedings. Subtitles are shorter in general and the domain may be even further away than parliament data, which explains the increased

amount of data to obtain reasonable lexical coverage.

Our next study looks at the impact of tagging the target language with supervised models. Our previous work on annotation projection and treebank translation relied entirely on annotation transfer from source to target when training target language parsing models. This means that we discarded any language-specific features and modeled parsing exclusively around universal PoS tags and lexical information. For highly-inflecting languages, this is not very satisfactory and the performance drops significantly compared to models that have access to morphological features. Therefore, we now test models that use projected data with additional annotation from automatic taggers. Table 2 summarizes the results of those experiments.

	projected PoS	target-tagged morph	PoS+morph
Projected			
sl-hr	<b>62.92</b>	62.10	56.42
cs-sk	63.57	–	70.68
da-no	57.82	58.08	61.40
sv-no	64.59	64.78	62.35
PBSMT			
sl-hr	60.81	61.60	61.10
cs-sk	67.81	–	<b>73.90</b>
da-no	57.87	58.46	63.67
sv-no	65.96	66.44	64.15
SyntaxSMT			
sl-hr	58.57	60.15	56.85
cs-sk	63.13	64.05	65.02
da-no	58.36	58.59	64.74
sv-no	66.31	66.64	65.43
da+sv-no	–		<b>67.80</b>

Table 2: Added PoS and morphological tagging to projected data sets: LAS scores on development data. Only morphological tagging added (morph) or tagging both, PoS and morphology (PoS+morph).

There are two models that we evaluate: (i) A model that adds morphological features to the projected annotation, and (ii) a model that even overwrites the universal PoS tags created through projection. The first variant adds information that may contradict the PoS labels transferred from the source. For example, it may assign nominal inflection categories to a word labeled as verb through projection. The latter model should be more consistent between PoS and morphology but has the problem that those categories may not fit the dependency relations attached to the corresponding words when projecting from the source. This can

also greatly confuse the learning procedures.

As it turns out, overwriting the projected PoS labels is more severe in most cases except Slovak and Norwegian (only when projected from Danish). There, it seems to be beneficial to run complete tagging after projection. In almost all other cases the performance drops, often quite dramatic. On the other hand, adding morphology always helps, except for Croatian annotation projection (which is a bit surprising again).

There is no clear winner between phrase-based and syntax-based SMT. For Slovak and Croatian, phrase-based systems seem to work best whereas Norwegian performs better with syntax-based models. A combination of Danish and Swedish data gives another significant boost (retagging projected Danish including PoS and adding morphology to projected Swedish).

We then used the best results on development data for each of the three target languages to run the cross-lingual models on the test set. No further adjustments were done after tuning the models on development data. The final results of the official test are shown in Table 3.

LAS	hr	no	sk
supervised	73.37	81.77	71.41
delex	50.05	58.13	53.87
cross	56.91	60.22	61.17
CUNI	60.70	70.21	78.12
our model	57.98	68.60	73.14

UAS	hr	no	sk
supervised	80.16	85.59	78.73
delex	63.29	67.86	64.55
cross	68.52	69.31	70.60
CUNI	69.73	77.13	84.92
our model	69.57	76.77	82.87

Table 3: Final results on the test set (*our model*) compared to baselines and fully supervised models. *CUNI* refers to a competing system – the winning team of VarDial. For the Norwegian baselines we report the results for Swedish as the source language, which is much better than using Danish.

The results on test data mainly confirm the findings from the development phase. Slovak performs clearly best in the cross-lingual scenario. This is the only language pair for which the cross-lingual model even outperforms the fully super-

vised “upper bound”. This is quite fascinating and rather unexpected. Certainly, the Czech treebank is by far the largest one in the collection and much bigger than the corresponding Slovak treebank. The languages are also very close to each other and their morphological complexity requires sufficient resources. This may explain why the large Czech training data can compensate for the shortcomings of the small Slovak training data. Other factors for the positive result may also include the similarity in domains covered by both treebanks and the closeness of annotation principles. The performance for the other target languages is less impressive. Norwegian performs similar to the scores that we have seen in related work on annotation projection and cross-lingual parsing. Croatian is rather disappointing even though it also beats the cross-lingual baselines.

The main scores in our evaluations is LAS but it is also interesting to look at unlabelled attachment scores (UAS). Table 3 lists those scores as well and we can see that labelling seems to be a major problems for our models. The difference to LAS scores is dramatic, much more than the absolute difference we see between UAS and LAS in the fully supervised models. Compared to the winning submission at VarDial (*CUNI*, see (Rosa et al., 2017)), we can also see that the main difference is in LAS whereas UAS are rather similar. This seems to be a shortcoming of our approach that we should investigate more carefully.

## 4 Conclusions

Our experiments demonstrate the use of annotation projection and treebank translation techniques. The models perform well, especially for Slovak, which even outperforms the fully supervised “upper bound” model. In this paper, we have discussed the use of target language tagging on top of annotation projection with the conclusion that adding morphological information is almost always useful. We observe a large gap between LAS and UAS, which would require some deeper investigations. A possible reason is the use of language-specific dependency labels that are not available from the projection. However, we actually doubt that explanation looking at the success of the winning team. In their results, LAS did not suffer that much. Some surprising results could be seen as well, for example, the fact that Danish does not work as well as a source for Norwegian as

Swedish does. This cannot be explained in terms of linguistic grounds but need to refer to unexpected annotation differences or possibly a larger domain mismatch. Croatian as a target language was also surprisingly difficult and the performance is the worst in the final among all test cases. This improvement over the non-adapted Slovenian parser is only very modest whereas large gains can be observed for the other language pairs.

## References

- Bernd Bohnet. 2010. Top Accuracy and Fast Dependency Parsing is not a Contradiction. In *Proceedings of COLING*, pages 89–97.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping Parsers via Syntactic Projection across Parallel Texts. *Natural Language Engineering*, 11(3):311–325.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, N ria Bertomeu Castell , and Jungmee Lee. 2013. Universal Dependency Annotation for Multilingual Parsing. In *Proceedings of ACL*, pages 92–97.
- Thomas M ller and Hinrich Sch tze. 2015. Robust morphological tagging with word representations. In *Proceedings of NAACL*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of Language Resources and Evaluation (LREC)*, pages 1659–1666, Portoroz, Slovenia.
- Rudolf Rosa, Daniel Zeman, David Mare ek, and Zden k  abokrtsk y. 2017. Slavic Forest, Norwegian Wood. In *Proceedings of the VarDial Workshop*.
- Milan Straka, Jan Haji , and Strakov . 2016. UDPipe: trainable pipeline for processing CoNLL-U files performing tokenization, morphological analysis, pos tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, Paris, France, May. European Language Resources Association (ELRA).
- J rg Tiedemann and  eljko Agi . 2016. Synthetic treebanking for cross-lingual dependency parsing. *Journal of Artificial Intelligence Research*, 55:209–248, January.
- J rg Tiedemann,  eljko Agi , and Joakim Nivre. 2014. Treebank translation for cross-lingual parser induction. In *Proceedings of the 18th Conference Natural Language Processing and Computational Natural Language Learning (CoNLL)*, Baltimore, Maryland, USA.
- J rg Tiedemann. 2014. Rediscovering annotation projection for cross-lingual parser induction. In *Proceedings of COLING 2014*, Dublin, Ireland, August.
- J rg Tiedemann. 2015a. Cross-lingual dependency parsing with universal dependencies and predicted PoS labels. In *Proceedings of the Third International Conference on Dependency Linguistics (Depling 2015)*, pages 340–349, Uppsala, Sweden, August. Uppsala University, Uppsala, Sweden.
- J rg Tiedemann. 2015b. Improving the cross-lingual projection of syntactic dependencies. In *Proceedings of the 20th Nordic Conference of Computational Linguistics (NODALIDA 2015)*, pages 191–199, Vilnius, Lithuania, May. Link ping University Electronic Press, Sweden.
- Marcos Zampieri, Shervin Malmasi, Nikola Ljube i , Preslav Nakov, Ahmed Ali, J rg Tiedemann, Yves Scherrer, and No mi 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, Spain.