

# The DCU-EPFL Enhanced Dependency Parser at the IWPT 2021 Shared Task

James Barry<sup>1</sup>, Alireza Mohammadshahi<sup>2,3</sup>, Joachim Wagner<sup>1</sup>, Jennifer Foster<sup>1</sup>, James Henderson<sup>2</sup>

<sup>1</sup> ADAPT Centre, School of Computing, Dublin City University

<sup>2</sup> Idiap Research Institute

<sup>3</sup> École Polytechnique Fédérale de Lausanne—EPFL

<sup>1</sup> {james.barry, joachim.wagner, jennifer.foster}@adaptcentre.ie

<sup>2</sup> {alireza.mohammadshahi, james.henderson}@idiap.ch

## Abstract

We describe the *DCU-EPFL* submission to the IWPT 2021 Parsing Shared Task: *From Raw Text to Enhanced Universal Dependencies*. The task involves parsing Enhanced UD graphs, which are an extension of the basic dependency trees designed to be more facilitative towards representing semantic structure. Evaluation is carried out on 29 treebanks in 17 languages and participants are required to parse the data from each language starting from raw strings. Our approach uses the Stanza pipeline to preprocess the text files, XLM-RoBERTa to obtain contextualized token representations, and an edge-scoring and labeling model to predict the enhanced graph. Finally, we run a post-processing script to ensure all of our outputs are valid Enhanced UD graphs. Our system places 6th out of 9 participants with a coarse Enhanced Labeled Attachment Score (ELAS) of 83.57. We carry out additional post-deadline experiments which include using Trankit for pre-processing, XLM-RoBERTa<sub>LARGE</sub>, treebank concatenation, and multitask learning between a basic and an enhanced dependency parser. All of these modifications improve our initial score and our final system has a coarse ELAS of 88.04.

## 1 Introduction

The IWPT 2021 Parsing Shared Task: *From Raw Text to Enhanced Universal Dependencies* (Bouma et al., 2021) is the second task involving the prediction of Enhanced Universal Dependencies (EUD) graphs<sup>1</sup> following the 2020 task (Bouma et al., 2020). EUD graphs are an extension of basic UD trees, designed to be more useful in shallow natural language understanding tasks (Schuster and Manning, 2016) and lend themselves more easily to the

<sup>1</sup><https://universaldependencies.org/overview/enhanced-syntax.html>

representation of semantic structure than strict surface structure dependency trees. In the shared task, the enhanced graphs must be predicted from raw text, i.e. participants must segment the input into sentences and tokens. Participants are encouraged to predict lemmas, Part-of-Speech (POS) tags, morphological features and basic dependency trees as well.

Our system, *DCU-EPFL*, uses a single multilingual Transformer (Vaswani et al., 2017) encoder, namely XLM-RoBERTa (XLM-R) (Conneau et al., 2020), which is a multilingual RoBERTa model (Liu et al., 2019), to obtain contextualized token encodings. These are then passed to the enhanced dependency parsing model. The system is straightforward to apply to new languages with enhanced UD annotations. In the official submission, we use the same hyper-parameters for all languages. Our parsing component can produce arbitrary graphs, including graph structures where words may have multiple heads and cyclic graphs. Our system uses the following three components:

1. Stanza (Qi et al., 2020) for sentence segmentation, tokenization and the prediction of all UD features apart from the enhanced graph.
2. A Transformer-based dependency parsing model to predict Enhanced UD graphs.
3. A post-processor ensuring that every graph is a rooted graph where all nodes are reachable from the notional root token.

Our official system placed 6th out of 9 teams with a coarse Enhanced Labeled Attachment Score (ELAS) of 83.57. In a number of unofficial post-evaluation experiments, we make four incremental changes to our pipeline approach:

1. We replace the Stanza pre-processing pipeline with Trankit (Nguyen et al., 2021).

2. We use XLM-R<sub>Large</sub> instead of XLM-R<sub>Base</sub>.
3. We concatenate treebanks from the same language which have more than one training treebank and concatenating English treebanks to the Tamil training data.
4. We introduce a novel multitask model which parses the basic UD tree and enhanced graph in tandem.

All of these additional steps improved our evaluation scores, and for our final system, which incorporates all additional modifications, our evaluation score increases from 83.57 to 88.04. Our code is publicly available.<sup>2</sup>

## 2 Related Work

In this section, we discuss the relevant literature related to Enhanced Universal Dependencies.

### 2.1 Enhanced Universal Dependencies

Despite the recent wave of Deep Learning models and accompanying analyses that show that such models learn information about syntax, there is still interest and merit in utilizing hierarchically structured representations such as trees and semantic representations to provide greater supervision about what is taking place in a sentence (Oepen et al., 2019). While dependency trees are often used in downstream applications, their structural restrictions may hinder the representation of content words (Schuster and Manning, 2016). The Enhanced UD representation tries to fill this gap by enabling more expressive graphs in the UD format, which capture phenomena such as added subject relations in control and raising, shared heads and dependents in coordination, the insertion of null nodes for elided predicates, co-reference in relative clause constructions and augmenting modifier relations with prepositional or case-marking information.

Schuster and Manning (2016) build on the Stanford Dependencies (SD) initiative (de Marneffe et al., 2006) and extend certain flavors of the SD dependency graph representations to UD in the form of enhanced UD relations for English. They use a rule-based system that converts basic UD trees to enhanced UD graphs based on dependency structures identified to require enhancement. Nivre

<sup>2</sup><https://github.com/jbrry/IWPT-2021-shared-task>

et al. (2018) use rule-based and data-driven approaches in a cross-lingual setting for bootstrapping enhanced UD representations in Swedish and Italian and show that both techniques are capable of annotating enhanced dependencies in different languages.

### 2.2 The IWPT 2020 Shared Task on Parsing Enhanced Universal Dependencies

The first shared task on parsing Enhanced Universal Dependencies (Bouma et al., 2020) brought renewed attention to the problem of predicting enhanced UD graphs. Ten teams submitted to the task. The winning system (Kanerva et al., 2020) utilized the UDify model (Kondratyuk and Straka, 2019), which uses a BERT model (Devlin et al., 2019) as the encoder with multitask classifiers for POS-tagging, morphological prediction and dependency parsing built on top. They developed a system for encoding the enhanced representation into the basic dependencies so it can be predicted in the same way as a basic dependency tree but with enriched dependency types that can then be converted into the enhanced structure. In an unofficial submission shortly after the task deadline, Wang et al. (2020) outperform the winning system using second-order inference methods with Mean-Field Variational Inference.

Most systems used pretrained Transformers to obtain token representations, either by using the Transformer directly (Kanerva et al., 2020; Grünewald and Friedrich, 2020; He and Choi, 2020) or passing the encoded representation to BiLSTM layers where they are combined with other features such as context-free FastText word embeddings (Wang et al., 2020), character features and features obtained from predicted POS tags, morphological features and basic UD trees (Barry et al., 2020), or are used as frozen embeddings (Hershcovich et al., 2020). The only transition-based system among the participating teams (Hershcovich et al., 2020) used a stack-LSTM architecture (Dyer et al., 2015). Ek and Bernardy (2020) and Dehouck et al. (2020) combine basic dependency parsers and a rule-based system to generate EUD graphs from the predicted trees.

## 3 Official System Overview

This section describes our official system, which is the system we submitted prior to the competition deadline. The architecture of our system is

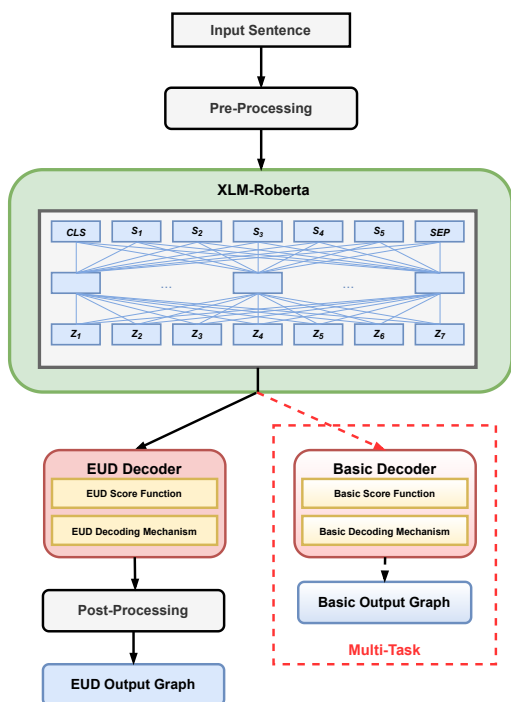


Figure 1: DCU-EPFL Architecture.

shown in Figure 1.<sup>3</sup> The raw text test files for each language contain a mixture of test data covering multiple treebanks, so participants do not know their exact domain. For our official system, we choose the model trained on the treebank with the most amount of training data in terms of sentences for each language to process the test files. This heuristic corresponds to using Czech-PDT for Czech, Dutch-Alpino for Dutch, English-EWT for English, Estonian-EDT for Estonian and Polish-PDB for Polish.

### 3.1 Pre-processing

For sentence segmentation, tokenization and the prediction of the base UD features (all UD features apart from the enhanced dependency graphs and miscellaneous items in CoNLL-U files), we use the Stanza library (Qi et al., 2020) trained on version 2.7 of the UD treebanks for each treebank released as part of the training data for the shared task.<sup>4</sup> Note that our parser does not pre-suppose any input features other than the input text but we predict the base features using our pre-processing pipeline

<sup>3</sup>For the official system, we did not include the basic dependency parser in a multitask setup.

<sup>4</sup>For Arabic, our Stanza Multi-word Token (MWT) expander predicted MWTs with a span of length 1 for two sentences. In the UD guidelines, MWT span lengths must be larger than one. To pass validation, we trained a UDPipe tokenizer (Straka and Straková, 2017) with Word2Vec embeddings for Arabic instead.

for completeness and to enable possible additional post-processing which involves altering enhanced dependency labels with lemma information.

### 3.2 Enhanced UD Parsing

For the enhanced UD parser, we use a Transformer encoder in the form of XLM-R (Conneau et al., 2020) with a first-order arc-factored model which utilizes the edge and label scoring method of (Kiperwasser and Goldberg, 2016). In initial experiments, we found this model to perform better than biaffine attention (Dozat and Manning, 2016) for the task of EUD parsing. This finding was also made by (Lindemann et al., 2019) and (Straka and Straková, 2019) for the task of semantic parsing across numerous Graphbanks (Oepen et al., 2019). Straka and Straková (2019) suggest that biaffine attention may be less suitable for predicting whether an edge exists between any pair of nodes using a predefined threshold and is perhaps more suited for dependency parsing, where words are competing with one another to be classified as the head in a softmax layer. The consistency of these findings across EUD and semantic parsing Graphbanks may provide evidence that enhanced UD is closer to semantic dependency parsing than basic UD parsing.

**Parser Implementation** Given a sentence  $x$  of length  $n$ , our model computes vector representations  $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n)$  for the predicted tokens  $(x_1, x_2, \dots, x_n)$ . Since the WordPiece tokenization (Wu et al., 2016) of XLM-R differs from the tokenization used in UD, we track the mapping  $I$  from XLM-R’s  $k$ -th sub-word unit of the  $j$ -th input token produced by Stanza to the sub-word unit’s position  $I_{j,k}$  in context of the sentence and we consider the output vector  $\mathbf{e}_{I_{j,1}}$  of the first sub-word unit of each word  $x_j$  as its vector representation ( $\mathbf{r}_j$ ):

$$\begin{aligned} \mathbf{E} &= \text{XLMR}(x_1, x_2, \dots, x_n) \\ \mathbf{R} &= \text{Filter}(\mathbf{E}, \mathbf{I}) \end{aligned} \quad (1)$$

where  $\mathbf{E} = (\mathbf{e}_1, \dots, \mathbf{e}_N)$  are the output vectors of all sub-word units,  $N$  being the total number of sub-word units in the sentence, and  $\text{Filter}()$  chooses the first embedding for each token. We add a dummy representation of the same dimensionality for the ROOT token to the sequence of vectors  $\mathbf{R}$  but mask out predictions from this token. Following Kiperwasser and Goldberg (2016), these representations  $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_n)$  are then passed to the dependency parsing component, where the feature

function  $\phi$  is the concatenation of the representations of a potential head word  $x_h$  and dependent word  $x_d$ , where  $\circ$  denotes concatenation:

$$\phi(h, d) = \mathbf{r}_h \circ \mathbf{r}_d \quad (2)$$

**Edge Prediction** We compute scores for all  $n(n - 1)$  potential edges  $(h, d)$ ,  $h \neq d$ , with an MLP:

$$s_{hd}^{(\text{arc})} = \text{MLP}^{(\text{arc})}(\phi(h, d)) \quad (3)$$

The edge classifier computes scores for all possible head-dependent pairs, and we compute a sigmoid on the resulting matrix of scores to obtain probabilities. We use an edge prediction threshold of 0.5, i.e. we include all edges with a score above 0.5 in the preliminary EUD graph. This enables words to have multiple heads but it can also lead to words receiving no head, where we manually select the edge that has the highest probability, and to fragmented graphs, see post-processing in Section 3.3.

**Label Prediction** To label the graph, we then choose a label for each edge using a separate classifier:

$$s_{hd}^{(\text{label})} = \text{MLP}^{(\text{label})}(\phi(h, d)) \quad (4)$$

The scores for all possible labels are passed to a softmax layer, which outputs the probability of each label for edge  $(h, d)$  and we select the label with the highest probability for each edge.

**Loss Function** For edge prediction, sigmoid cross-entropy loss is used, and for label prediction, as we want to select the label for each chosen edge, softmax cross-entropy loss is used (Dozat and Manning, 2018). We interpolate between the loss given by the edge classifier and the loss given by the label classifier (Dozat and Manning, 2018; Wang et al., 2020) with a constant  $\lambda$ :

$$\mathcal{L} = \lambda \mathcal{L}^{(\text{label})} + (1 - \lambda) \mathcal{L}^{(\text{edge})} \quad (5)$$

**Training details** For the empty nodes which are prevalent in enhanced UD graphs, we added them into the graph, and offset the head indices to account for the new token(s) added to the graph. At test time, we did not predict whether an elided token should be added to the graph. Due to time constraints, we trained using the full lexicalized enhanced dependency labels but intend to devise a delocalization and relexicalization procedure in future work.

Hyperparameter	Size
XLM-R <sub>Base</sub> Hidden Size	768
XLM-R <sub>Large</sub> Hidden Size	1024
Edge Feedforward	300
Label Feedforward	300
Input Dropout	0.35
Dropout	0.35
Edge Prediction Threshold	0.5
Loss interpolation $\lambda$	0.10

Table 1: Hyperparameters of our EUD parsing model.

### 3.3 Post-processing

In the Enhanced UD guidelines, the predicted structure must be a connected graph where all nodes are reachable from the notional root<sup>5</sup>. After predicting the test files, we use the graph connection tool in (Barry et al., 2020) to make sure that each sentence is a connected graph. Specifically, we repeatedly check for unreachable nodes and the number of unreachable nodes that can be reached from them. We choose the candidate which maximises this number (in the case there are ties, we choose the first node in surface order) and makes it a child of the notional ROOT, i.e. this node becomes an additional root node. System outputs are then validated at level 2 by the UD validator<sup>6</sup> to catch bugs prior to submission.

## 4 Experiments

In this section, we discuss our official results and then describe post-deadline experiments that improved our submission’s score. Model hyperparameters are listed in Table 1. The choice of XLM-R encoder (Base or Large) determines the hyperparameters of the encoder part of our model. In our official submission, we use XLM-R<sub>Base</sub>. A dropout value of 0.35 is used for the input embeddings as well as for the encoder and MLP networks. A loss interpolation constant  $\lambda$  of 0.1 is used as in (Wang et al., 2020).

### 4.1 Official Submission

For the official submission, we use the Stanza pre-processing pipeline and our dependency parsing model with XLM-R<sub>Base</sub>. The results are listed in

<sup>5</sup>In UD, the notional ROOT is the token with ID 0, whereas a root node is any node that has 0 as its head.

<sup>6</sup><https://github.com/UniversalDependencies/tools/blob/master/validate.py>



Language	[1]	[2]	[3]	[4]	[5]
	<b>Official</b>	<b>[1]+Trankit</b>	<b>[2]+XLM-R<sub>Large</sub></b>	<b>[3]+Concat</b>	<b>[4]+MTL</b>
Arabic	71.01	78.05(+24.2%)	79.51(+6.6%)	-	81.72(+10.8%)
Bulgarian	92.44	92.47(+0.4%)	93.26(+10.5%)	-	93.59(+4.9%)
Czech	89.93	90.28(+3.5%)	91.06(+8.1%)	91.43(+4.1%)	91.30(-1.5%)
Dutch	81.89	86.51(+25.5%)	87.67(+8.6%)	88.60(+7.5%)	89.51(+7.9%)
English	85.70	85.97(+1.8%)	86.94(+6.9%)	87.46(+3.9%)	87.28(-1.4%)
Estonian	84.35	84.54(+1.2%)	85.92(+8.9%)	86.68(+5.4%)	86.76(+0.6%)
Finnish	89.02	89.34(+2.9%)	90.79(+13.6%)	-	91.16(+4.1%)
French	86.68	86.80(+0.9%)	89.12(+17.6%)	-	90.38(+11.6%)
Italian	92.41	92.44(+0.4%)	93.35(+12.1%)	-	93.47(+1.8%)
Latvian	86.96	86.85(-0.8%)	88.81(+14.9%)	-	89.18(+3.3%)
Lithuanian	78.04	78.44(+1.8%)	82.09(+16.9%)	-	83.47(+7.7%)
Polish	89.17	89.30(+1.2%)	90.20(+8.4%)	91.15(+9.7%)	90.46(-7.8%)
Russian	92.83	93.06(+3.2%)	93.95(+12.8%)	-	94.09(+2.3%)
Slovak	89.59	90.81(+11.7%)	92.33(+16.5%)	-	92.73(+5.2%)
Swedish	85.20	85.98(+5.3%)	88.10(+15.1%)	-	88.64(+4.5%)
Tamil	39.32	40.64(+2.2%)	48.85(+13.8%)	61.14(+24.0%)	62.06(+2.4%)
Ukrainian	86.09	86.30(+1.5%)	89.44(+22.91%)	-	90.91(+13.9%)
<b>Average</b>	<b>83.57</b>	<b>84.58(+6.2%)</b>	<b>86.55(+12.7%)</b>	<b>87.48(+6.9%)</b>	<b>88.04(+4.5%)</b>

Table 2: Evaluation scores on the official test data on the language-specific test files. All runs after **Official** subsume **Trankit** pre-processing and all runs after **Trankit** subsume the **XLM-R<sub>Large</sub>** model. All numbers inside the parentheses are calculated as the relative error reduction of each column and its corresponding previous column.

column [1] of Table 2. Our official submission placed 6th of 9 participants. The overall scores submitted by each team are listed in Table 3. The scores of two teams are close to our overall score: Combo and Unipi placed 4th and 5th with scores of 83.79 and 83.64 compared to our score of 83.57. This grouping is outperformed by the top three submissions TGIF, ShanghaiTech and RobertNLP by a margin from 3.2 ELAS points (RobertNLP vs. Combo) to 5.7 ELAS points (TGIF vs. DCU-EPFL).

## 4.2 Trankit Pre-processing

In a post-deadline experiment, we replace the Stanza pre-processing pipeline (which uses Word2Vec and FastText embeddings as external input features and a BiLSTM encoder) with Trankit (Nguyen et al., 2021), which uses the Transformer XLM-R as the encoder. The results from adopting Trankit for sentence segmentation and tokenization are listed in column [2] of Table 2. We notice slight improvements for all languages, with notable exceptions being Arabic, Dutch and Slovak, where the better pre-processing accounts for a 24.2%, 25.5% and 11.7% relative error reduction.

## 4.3 XLM-R<sub>Large</sub>

Our next modification is to leverage the XLM-R<sub>Large</sub> model. This model has roughly twice as many parameters as the XLM-R<sub>Base</sub> model used in our official submission. The results for combining Trankit pre-processing and using XLM-R<sub>Large</sub> are listed in column [3] of Table 2. The larger capacity of the model translates to large relative error reductions particularly for Finnish, French, Latvian, Lithuanian, Swedish, Tamil and Ukrainian. Given the improvements seen by adopting both Trankit for pre-processing and the larger XLM-R<sub>Large</sub> model, we now incorporate these modifications into all further experiments.

## 4.4 Treebank Concatenation

In our official system, we used just one treebank per language. Our next experiment is to investigate the effect of concatenating all treebanks with enhanced UD annotations for a language. We hypothesize that there could be a positive transfer from learning similar (within-language) treebanks and that it would make our parser more robust to the multiple domains in the test data. This means that for Czech we concatenate the PDT, CAC and FicTree treebanks, for Dutch, Alpino and LassySmall, for English EWT and GUM, and for Estonian EDT

Language	combo	dcu-epfl	fastparse	grew	nuig	robertnlp	shanghaitech	tgif	unipi	off. reference	our best run
Arabic	76.39	71.01	53.74	71.13	0.0	81.58	<b>82.26</b>	81.23	77.17	67.35	<b>81.72</b>
Bulgarian	86.67	92.44	78.73	88.83	78.45	93.16	92.52	<b>93.63</b>	90.84	85.81	<b>93.59</b>
Czech	89.08	89.93	72.85	87.66	0.0	90.21	<b>91.78</b>	<b>92.24</b>	88.73	78.44	91.30
Dutch	87.07	81.89	68.89	84.09	0.0	88.37	88.64	<b>91.78</b>	84.14	82.48	<b>89.51</b>
English	84.09	85.70	73.00	85.49	65.40	87.88	87.27	<b>88.19</b>	87.11	83.68	<b>87.28</b>
Estonian	84.02	84.35	60.05	78.19	54.03	86.55	86.66	<b>88.38</b>	81.27	76.86	<b>86.76</b>
Finnish	87.28	89.02	57.71	85.20	0.0	91.01	90.81	<b>91.75</b>	89.62	78.26	<b>91.16</b>
French	87.32	86.68	73.18	83.33	0.0	88.51	88.40	<b>91.63</b>	87.43	98.80	<b>90.38</b>
Italian	90.40	92.41	78.32	90.98	0.0	93.28	92.88	<b>93.31</b>	91.81	80.20	<b>93.47</b>
Latvian	84.57	86.96	66.43	77.45	56.67	88.82	89.17	<b>90.23</b>	83.01	79.32	<b>89.18</b>
Lithuanian	79.75	78.04	48.27	74.62	59.13	80.76	80.87	<b>86.06</b>	71.31	75.26	<b>83.47</b>
Polish	87.65	89.17	71.52	78.20	0.0	89.78	<b>90.66</b>	<b>91.46</b>	88.31	81.59	<b>90.46</b>
Russian	90.73	92.83	78.56	90.56	66.33	92.64	93.59	<b>94.01</b>	90.90	79.63	<b>94.09</b>
Slovak	87.04	89.59	64.28	86.92	67.45	89.66	90.25	<b>94.96</b>	86.05	76.42	<b>92.73</b>
Swedish	83.20	85.20	67.26	81.54	63.12	88.03	86.62	<b>89.90</b>	84.91	80.98	<b>88.64</b>
Tamil	52.27	39.32	42.53	58.69	0.0	59.33	58.94	<b>65.58</b>	51.73	75.44	<b>62.06</b>
Ukrainian	86.92	86.09	63.42	83.90	0.0	88.86	88.94	<b>92.78</b>	87.51	77.24	<b>90.91</b>
<b>Average</b>	<b>83.79</b>	<b>83.57</b>	<b>65.81</b>	<b>81.58</b>	<b>30.03</b>	<b>86.97</b>	<b>87.07</b>	<b>89.24</b>	<b>83.64</b>	<b>79.87</b>	<b>88.04</b>

Table 3: Evaluation scores on the official test data on the language-specific test files submitted by each team. We also include the official reference system (**off. reference**) which copies the gold tree to the enhanced graph as well as (**our best run**) which is our best post-deadline run, which corresponds to the **+Concat+MTL** run in Table 2. The first and second top scoring models in each language are specified with black and blue color, respectively.

and EWT. For Tamil, we concatenate English EWT and GUM training data to Tamil to address the very poor evaluation score of our official submission, taking inspiration from Wang et al. (2020) who observe substantial positive effects when they add Czech and English data to the Tamil treebank.<sup>7</sup> The results are listed in column [4] of Table 2. Treebank concatenation helps for all languages but most notable is the improvement of over 12 points ELAS or a relative error reduction of 24% for Tamil, the language with the least amount of training data in the task.

#### 4.5 Joint Learning of Basic and Enhanced Dependency Parsing

The official reference system submitted by the shared task organizers which copies the gold trees to the enhanced representation performs very well with 79.87 ELAS (see Table 3). Thus, there is evidence that the basic tree and enhanced graph contain a lot of mutual information. Previous methods which have leveraged the basic representation for producing EUD graphs (see Sec. 2) have focused on using heuristic rules to convert the basic tree to EUD (Schuster and Manning, 2016; Ek and Bernardy, 2020; Dehouck et al., 2020), using the basic tree as input features to the enhanced parsing model (Barry et al., 2020) or converting the enhanced graph to a richer basic representation (Kanerva et al., 2020).

<sup>7</sup>We did not include Czech to reduce training time.

In our final experiment, we try to leverage the information from the basic tree by jointly learning to predict the enhanced graph and the basic tree, testing whether performing basic dependency parsing and EUD parsing in a multitask setup is beneficial for EUD parsing. Given the positive effects seen through concatenation, for those languages where we performed concatenation, we also train multitask models on the concatenated versions of treebanks. We use our EUD parsing model as in Section 3 and integrate with additional basic dependency parsing component (as shown in the right part of Figure 1) which is the biaffine parsing model of Dozat and Manning (2016) and train both parsers jointly. The losses of the two components are combined with equal weight. The results are listed in column [5] of Table 2.

**Single Treebanks** First, we compare the multitask model to the  $\text{XLM-R}_{\text{Large}}$  run for languages where we did not perform concatenation. Predicting the basic tree and the enhanced graph in a multitask setting yields improvements for all languages, particularly for Arabic, French and Ukrainian.

**Multitask Model and Treebank Concatenation** When used alongside treebank concatenation, multitask learning can help for Dutch, Estonian and Tamil where it provides additional performance gains. It is interesting to note that concatenation alone is more helpful for Czech and English where we see slight performance drops and multitask

learning is not helpful when trained on concatenated Polish treebanks.

The positive contribution of multitask learning for all languages when not performing treebank concatenation, could mean that it would be useful in settings where only one treebank with the enhanced representation is available for a language and the basic tree could be used as auxiliary information to predict the enhanced representation.

**Comparison to Official Systems** Our best unofficial run +Concat+MTL is added to Table 3. Compared to the other official runs, the ELAS scores of this run ranks in second place for 13/17 languages and places first for Italian and Russian.

## 5 Conclusion

We have described the *DCU-EPFL* submission to the IWPT 2021 Shared Task on Parsing into Enhanced Universal Dependencies. Our approach uses a single multilingual Transformer encoder as well as an enhanced dependency parsing component. Our official system placed 6th out of 9 teams. In post-deadline experiments, we show how our submission can be improved by leveraging better upstream pre-processing, a larger encoder, concatenating treebanks as well as introducing a multitask parser that can parse the basic tree and enhanced graphs jointly.

## Acknowledgments

This research is supported by Science Foundation Ireland through the ADAPT Centre for Digital Content Technology, which is funded under the SFI Research Centres Programme (Grant 13/RC/2106) and is co-funded under the European Regional Development Fund. This research is also supported by the Swiss National Science Foundation, grant CRSII5-180320. We would like to thank the shared task organizers as well as the anonymous reviewers for their helpful feedback.

## References

James Barry, Joachim Wagner, and Jennifer Foster. 2020. [The ADAPT enhanced dependency parser at the IWPT 2020 shared task](#). In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 227–235, Online. Association for Computational Linguistics.

Gosse Bouma, Djamé Seddah, and Daniel Zeman. 2020. [Overview of the IWPT 2020 shared task on](#)

[parsing into enhanced Universal Dependencies](#). In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 151–161, Online. Association for Computational Linguistics.

Gosse Bouma, Djamé Seddah, and Daniel Zeman. 2021. [From raw text to enhanced universal dependencies: The parsing shared task at iwpt 2021](#). In *Proceedings of the 17th International Conference on Parsing Technologies (IWPT 2021)*, pages 146–157, Bangkok, Thailand (online). Association for Computational Linguistics.

Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

Mathieu Dehouck, Mark Anderson, and Carlos Gómez-Rodríguez. 2020. [Efficient EUD parsing](#). In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 192–205, Online. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Timothy Dozat and Christopher D. Manning. 2016. [Deep biaffine attention for neural dependency parsing](#). *CoRR*, abs/1611.01734.

Timothy Dozat and Christopher D. Manning. 2018. [Simpler but more accurate semantic dependency parsing](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 484–490, Melbourne, Australia. Association for Computational Linguistics.

Chris Dyer, Miguel Ballesteros, Wang Ling, Austin Matthews, and Noah A. Smith. 2015. [Transition-based dependency parsing with stack long short-term memory](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 334–343, Beijing, China. Association for Computational Linguistics.

- Adam Ek and Jean-Philippe Bernardy. 2020. [How much of enhanced UD is contained in UD?](#) In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 221–226, Online. Association for Computational Linguistics.
- Stefan Grünewald and Annemarie Friedrich. 2020. [RobertNLP at the IWPT 2020 shared task: Surprisingly simple enhanced UD parsing for English.](#) In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 245–252, Online. Association for Computational Linguistics.
- Han He and Jinho D. Choi. 2020. [Adaptation of multilingual transformer encoder for robust enhanced Universal Dependency parsing.](#) In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 181–191, Online. Association for Computational Linguistics.
- Daniel Hershcovich, Miryam de Lhoneux, Artur Kuznetsov, Elham Pejman, and Joakim Nivre. 2020. [Køpsala: Transition-based graph parsing via efficient training and effective encoding.](#) In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 236–244, Online. Association for Computational Linguistics.
- Jenna Kanerva, Filip Ginter, and Sampo Pyysalo. 2020. [Turku enhanced parser pipeline: From raw text to enhanced graphs in the IWPT 2020 shared task.](#) In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 162–173, Online. Association for Computational Linguistics.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. [Simple and accurate dependency parsing using bidirectional LSTM feature representations.](#) *Transactions of the Association for Computational Linguistics*, 4:313–327.
- Dan Kondratyuk and Milan Straka. 2019. [75 languages, 1 model: Parsing Universal Dependencies universally.](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2779–2795, Hong Kong, China. Association for Computational Linguistics.
- Matthias Lindemann, Jonas Groschwitz, and Alexander Koller. 2019. [Compositional semantic parsing across graphbanks.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4576–4585, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach.](#)
- Marie-Catherine de Marneffe, Bill MacCartney, and Christopher D. Manning. 2006. [Generating typed dependency parses from phrase structure trees.](#) In *LREC*.
- Minh Van Nguyen, Viet Dac Lai, Amir Pouran Ben Veyseh, and Thien Huu Nguyen. 2021. [Trankit: A light-weight transformer-based toolkit for multilingual natural language processing.](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 80–90, Online. Association for Computational Linguistics.
- Joakim Nivre, Paola Marongiu, Filip Ginter, Jenna Kanerva, Simonetta Montemagni, Sebastian Schuster, and Maria Simi. 2018. [Enhancing Universal Dependency treebanks: A case study.](#) In *Proceedings of the Second Workshop on Universal Dependencies (UDW 2018)*, pages 102–107, Brussels, Belgium. Association for Computational Linguistics.
- Stephan Oepen, Omri Abend, Jan Hajic, Daniel Hershcovich, Marco Kuhlmann, Tim O’Gorman, Nianwen Xue, Jayeol Chun, Milan Straka, and Zdenka Uresova. 2019. [MRP 2019: Cross-framework meaning representation parsing.](#) In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*, pages 1–27, Hong Kong. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: A python natural language processing toolkit for many human languages.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 101–108, Online. Association for Computational Linguistics.
- Sebastian Schuster and Christopher D. Manning. 2016. [Enhanced English Universal Dependencies: An improved representation for natural language understanding tasks.](#) In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 2371–2378, Portorož, Slovenia. European Language Resources Association (ELRA).
- Milan Straka and Jana Straková. 2017. [Tokenizing, POS tagging, lemmatizing and parsing UD 2.0 with UDPipe.](#) In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99, Vancouver, Canada. Association for Computational Linguistics.
- Milan Straka and Jana Straková. 2019. [ÚFAL MRPipe at MRP 2019: UDPipe goes semantic in the meaning](#)



representation parsing shared task. In *Proceedings of the Shared Task on Cross-Framework Meaning Representation Parsing at the 2019 Conference on Natural Language Learning*, pages 127–137, Hong Kong. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Xinyu Wang, Yong Jiang, and Kewei Tu. 2020. [Enhanced Universal Dependency parsing with second-order inference and mixture of training data](#). In *Proceedings of the 16th International Conference on Parsing Technologies and the IWPT 2020 Shared Task on Parsing into Enhanced Universal Dependencies*, pages 215–220, Online. Association for Computational Linguistics.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#).