

A Morphology-based Representation Model for LSTM-based Dependency Parsing of Agglutinative Languages

Şaziye Betül Özateş*, Arzucan Özgür*, Tunga Güngör*, Balkız Öztürk†

*Department of Computer Engineering

†Department of Linguistics

Boğaziçi University

Bebek, 34342 İstanbul, Turkey

saziye.bilgin, arzucan.ozgur, gungort, balkiz.ozturk@boun.edu.tr

Abstract

We propose two word representation models for agglutinative languages that better capture the similarities between words which have similar tasks in sentences. Our models highlight the morphological features in words and embed morphological information into their dense representations. We have tested our models on an LSTM-based dependency parser with character-based word embeddings proposed by Ballesteros et al. (2015). We participated in the CoNLL 2018 Shared Task on multilingual parsing from raw text to universal dependencies as the BOUN team. We show that our morphology-based embedding models improve the parsing performance for most of the agglutinative languages.

1 Introduction

This paper describes our submission to the CoNLL 2018 Shared Task (Zeman et al., 2018b) on parsing of Universal Dependencies (UD) (Nivre et al., 2016). We propose morphologically enhanced character-based word embeddings to improve the parsing performance especially for agglutinative languages. We apply our approach to a transition-based dependency parser by Ballesteros et al. (2015) that uses stack Long Short Term Memory structures (LSTMs) to predict the parser state. This parser uses character-level word representation, which has been shown to perform better for languages with rich morphology (Ballesteros et al., 2015; Dozat et al., 2017). From our experiment results performed on UD version 2.2 data sets (Nivre et al., 2018; Zeman et al., 2018a) we observe that including morphological information to a character-based word embedding model yields

a better learning of relationships between words and increases the parsing performance for most of the agglutinative languages with rich morphology.

The rest of the paper is organized as follows: Section 2 provides a brief description of the LSTM-based dependency parser used in this study and introduces our embedding models. Section 3 gives the implementation details of our system and describes the training strategies we apply to different languages. Section 4 discusses our results on the shared task as well as the post-evaluation experiments and Section 5 concludes the paper.

2 Parsing Model

We use the LSTM-based parser by Ballesteros et al. (2015). It is an improved version of a state-of-the-art transition-based dependency parser proposed by Dyer et al. (2015) and uses stack LSTM structures with push and pop operations to learn representations of the parser state. Instead of lookup-based word representations, bidirectional LSTM modules are used to create character-based encodings of words. With this character-based modelling, the authors obtain improvements on the dependency parsing of many morphologically rich languages.

2.1 Character Embeddings of Words

The character-based word embedding model using bi-LSTMs in (Ballesteros et al., 2015) is depicted in Figure 1. The authors compute character-based vector representations of words using bi-LSTMs. Their embedding system reads each word character by character from the beginning to the end and computes an embedding vector of the character sequence, which is denoted as \vec{w} in Figure 1. The system also reads the word character by character from the end to the beginning and the produced embedding is denoted as \overleftarrow{w} . These two embedding vectors and the learned representation of the

POS-tag of the word t are concatenated to produce the vector representation of the word. A linear mapping of POS-tag words to integers is used to create a representation of the POS tags as in (Ballesteros et al., 2015).




Figure 1: Vector representation of the word *travel* with the character-based embedding model in (Ballesteros et al., 2015).

2.2 Morphology-based Character Embeddings

To improve the parsing performance of the LSTM parser with character-based word embeddings mentioned in Section 2.1, we include the morphological information of words to the embedding model. In agglutinative languages like Turkish, a stem usually takes different suffixes and by this way, different meanings are created using a single root-word. Words that share the same suffixes tend to have similar roles in a sentence. For instance, gerunds in Turkish are a kind of derivational suffixes. Verbs that take the same gerund as a suffix have usually the same role in sentences. Table 1 shows some statistics of verbs with gerunds in the development data set of Turkish-IMST treebank for demonstration purposes. The first column shows some example suffixes that attach to verbs and turn them to adverbs. The second column shows the number of verbs with the corresponding suffix in the development set. The third column shows the statistics of the dependency labels of these verbs. As it can be seen from the table, these suffixes help determining the role of the word they attach to. Therefore, representing each word using its corresponding lemma and suffixes separately and utilizing the morphological infor-

mation of words can improve the parsing performance in agglutinative languages.




Figure 2: Vector representation of the word *travel* with the character-based embedding model in (Ballesteros et al., 2015).

Lemma-Suffix Model

For agglutinative languages where the stem of a word does not change in different word forms, we created a model that uses lemma and suffix information of words in character-based embeddings. In this model, each word is separated to its lemma and suffixes. Then, the embedding system first reads the lemma of the word character by character from the beginning to the end and computes an embedding vector of the character sequence of the lemma which is denoted as \vec{r} . Secondly, the system reads the lemma character by character from the end to the beginning and the produced embedding is denoted as \overleftarrow{r} . A similar process is performed for the suffixes of the word and the produced vectors are denoted as \vec{s} and \overleftarrow{s} . These four embedding vectors and the vector representation of the POS-tag of the word t are then concatenated to produce the vector representation of the word. POS-tag representations are created by linearly mapping the POS-tag words to integers as in (Ballesteros et al., 2015). Vector representation of an example word using this model is depicted in Figure 3.

Morphological Features Model

The lemma-suffix model is suitable only for agglutinative languages which make use of suffixes to create different word forms. For languages that do not have this type of grammar, we created another model where the specific morphological fea-

Suffix	Number of Occurrences	Dependency Label				
-Ip	41	23 nmod	8 compound	5 conj	4 obj	1 root
-ArAk	32	26 nmod	3 conj	2 compound	1 root	
-ken	20	18 nmod	1 conj	1 acl		
-IncA	8	7 nmod	1 compound			
-mAdAn	7	4 nmod	2 compound	1 obj		
-DIkçA	3	3 nmod				

Table 1: Number of occurrences of some example suffixes and the corresponding dependency labels of verbs with these suffixes in the development data of *Turkish-IMST* treebank.

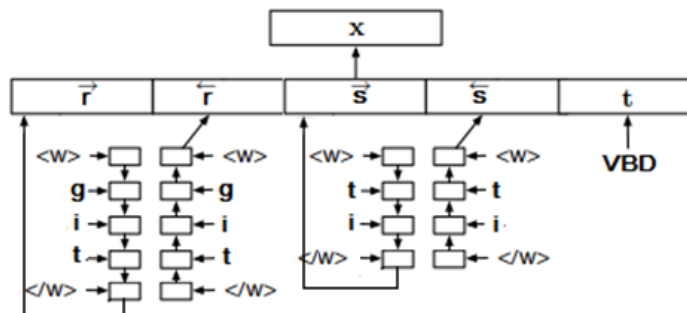


Figure 3: Character-based word embedding of a Turkish word *gitti* ("it went" in English) using lemma-suffix embedding model.

tures of each word are embedded to the dense representations of the words. The reason behind this choice is that some morphological features have a direct impact in identifying the dependency labels of words. For instance, if a word has a *case* feature and its value is *accusative*, then it is usually an object of the sentence. By extracting and utilizing such morphological features, we can improve the parsing accuracy for languages that suit this model.

In this model, the embedding of a word is created character by character as in Section 2.1. Then, the embedding vector of each of its selected morphological features are created by reading the feature value character by character from the beginning to the end. Finally, these embedding vectors and the vector representation of the POS-tag of the word are concatenated to produce the vector representation of the word.

The vector representation of an example word using its morphological features is shown in Figure 4.

3 Implementation

The systems participating in the CoNLL 2018 Shared Task on UD Parsing are expected to parse raw text without any gold-standard pre-processing operations such as tokenization, lemmatization,

and morphological analysis. However, the baseline pre-processed versions of the raw text by the UDPipe system (Straka et al., 2016) are available for the participants who want to focus only on the dependency parsing task. We used the automatically annotated version of the corpora provided by UDPipe, since our primary aim is to observe the effect of our embedding models on the dependency parsing of agglutinative languages.

Word	No morp.parser		With morp.parser	
	Lemma	Suffix	Lemma	Suffix
Her	her	-	her	-
şeyden	şey	den	şey	DAn
önce	önce	-	önce	-
sanatçıydı	sanat	çıydı	sanat	CHY DH

Table 2: Lemma and suffix separation example without using morphological analyzer and disambiguator and with using morphological analyzer and disambiguator on the Turkish sentence "Her şeyden önce sanatçıydı." (English meaning: "She was an artist before anything else.")

In the implementation of the lemma-suffix embedding model, we did not utilize any morphological analyzer and disambiguator tools to find the lemmas and the suffixes of the words. Instead, for each word in the treebank we extracted its corresponding *lemma* information from the conll-u ver-

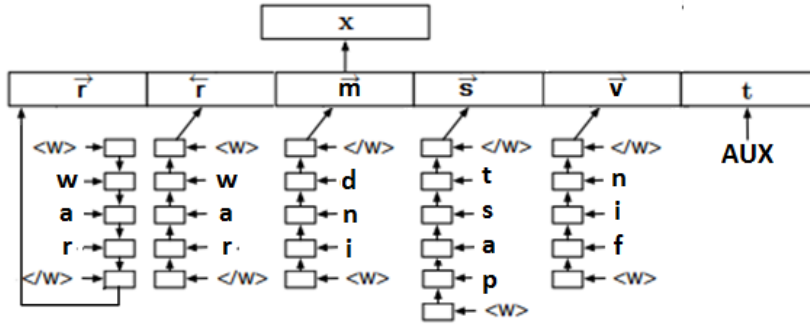


Figure 4: Character-based word embedding of a German word *war* ("was" in English) with its morphological features being $Mood = Ind|Number = Sing|Person = 3|Tense = Past|VerbForm = Fin$ using morphological features embedding model. The selected features for German are *Case*, *Mood*, *Tense*, and *VerbForm*. Since there is no *Case* feature in the morphological features of *war*, the *Case* feature is represented with an empty string in the word vector of *war*.

sion of the treebank data and subtracted the lemma from the word to find the suffix information. We compared these two approaches on the Turkish-IMST treebank. For this purpose, we utilized the Turkish morphological parser and disambiguator by (Sak et al., 2008). A comparison between the two approaches is shown on an example sentence in Table 2. We observed that finding the suffixes by subtracting the lemmas from the words gives the same parsing performance as using a morphological analyzer tool to find the lemma and suffixes of a word. So, we opted not to use a morphological analyzer and disambiguator for the languages with the lemma-suffix embedding model due to the additional costs of these tools.

3.1 Embedding Model Selection for Different Languages

We applied the lemma-suffix model in 2.2 to Buryat, Hungarian, Kazakh, Turkish, and Uyghur languages because these languages have agglutinative morphology, take suffixes, and the stem of a word usually does not change in different word forms. We also applied this model to Danish to observe the effect in parsing performance of a language with little inflectional morphology.

For the languages that do not follow this scheme, we applied the morphological features embedding model in 2.2. Table 3 shows the morphological features selected for these languages in the shared task. We selected four morphological features from the input conll-u files for most of the languages. For French, Indonesian, and Old French, we used less than four features because

there are less than four common morphological features in the conll-u files of these languages.

For Persian, Japanese, Korean, Vietnamese, and Chinese, we used the baseline embedding model due to the lack of representative morphological features in their corresponding conll-u files.

Languages without Training Data

We trained a mixed language parser model with morphological features embedding model for the languages with no training data. For training this parser model, we used the mixed language training data supplied by the organizers of the shared task. This data is created by including the first 200 sentences of each treebank.

In the shared task, this model is applied to the *Buryat-KEB*, *Czech-PUD*, *English-PUD*, *Faroese-OFT*, *Japanese-Modern*, *Naija-NSC*, *Swedish-PUD*, and *Thai-PUD* treebanks.

We trained parser models for the *Upper Sorbian-UFAL* and *Galician-TreeGal* treebanks using the morphological features embedding model and for the *Buryat-BDT* treebank using the lemma-suffix embedding model. However, we used the mixed language parser model for these treebanks in the shared task due to some software issues.

3.2 Training Specifications

Our model mostly uses the same hyper-parameter configuration with the original settings of the parser in (Ballesteros et al., 2015) with a few exceptions. We used stochastic gradient descent trainer with a learning rate of 0.13. We replaced

Language	Morphological Features			
Afrikaans	Aspect	Case	Tense	VerbForm
Ancient Greek	Aspect	Case	Tense	VerbForm
Arabic	Aspect	Case	Mood	VerbForm
Armenian	Aspect	Case	Tense	VerbForm
Basque	Aspect	Case	Tense	VerbForm
Bulgarian	Aspect	Case	Tense	VerbForm
Catalan	AdpType	Mood	Tense	VerbForm
Croatian	Case	Mood	Tense	VerbForm
Czech	Aspect	Case	Tense	VerbForm
Dutch	Degree	Case	Tense	VerbForm
English	Case	Mood	Tense	VerbForm
Estonian	Case	Mood	Tense	VerbForm
French	Mood	Tense	VerbForm	
Finnish	Case	Mood	Tense	VerbForm
Galician	Case	Mood	Tense	VerbForm
German	Case	Mood	Tense	VerbForm
Gothic	Case	Mood	Tense	VerbForm
Greek	Aspect	Case	Tense	VerbForm
Hebrew	HebBinyan	HebSource	Tense	VerbForm
Hindi	Aspect	Case	Tense	VerbForm
Indonesian	PronType	Degree		
Irish	Case	Mood	Tense	VerbForm
Italian	PronType	Mood	Tense	VerbForm
Kurmanji	Case	Mood	Tense	VerbForm
Latin	Case	Mood	Tense	VerbForm
Latvian	Aspect	Case	Tense	VerbForm
North Sami	Case	Mood	Tense	VerbForm
Norwegian	Case	Mood	Tense	VerbForm
Old Church Slavonic	Case	Mood	Tense	VerbForm
Old French	Tense	VerbForm		
Polish	Aspect	Case	Tense	VerbForm
Portuguese	PronType	Mood	Tense	VerbForm
Romanian	Case	Mood	Tense	VerbForm
Russian	Aspect	Case	Tense	VerbForm
Serbian	PronType	Mood	Tense	VerbForm
Slovak	Aspect	Case	Tense	VerbForm
Slovenian	Aspect	Case	Tense	VerbForm
Spanish	Case	Mood	Tense	VerbForm
Swedish	Case	Mood	Tense	VerbForm
Ukrainian	Aspect	Case	Tense	VerbForm
Upper Sorbian	Case	Mood	Tense	VerbForm
Mixed Language	Case	Mood	Tense	VerbForm

Table 3: List of morphological features used for the languages with the morphological features embedding model.

the original character-based embedding model with our embedding models. In the lemma-suffix model, the forward word vector and the backward word vector of the lemma of a word both have 50 dimensions. The forward and backward word vectors of the suffix of a word also have 50 dimensions each. In the morphological features model, each of the forward and backward word vectors of a word have 50 dimensions. Each of the four morphological feature vectors have 25 dimensions. If a morphological feature is absent in a word, an embedding vector of an empty string is created for that feature. So, we increased the dimension of the character-based representations to 200 in total.

The original parser is not compatible with UD parsing. We adapted it to be able to take input and produce output in conll-u format. The source code of our modified version of the LSTM-based parser by Ballesteros et al. (2015) can be found at <https://github.com/CoNLL-UD-2018/BOUN>.

A full run over the 82 test sets takes about 3 hours when no pre-trained embeddings are used, and 20 hours when the CoNLL-17 pre-trained word embeddings from (Ginter et al., 2017) are used on the TIRA virtual machine (Potthast et al., 2014). The largest of the test sets needs 4 GB memory without pre-trained word vectors. When the CoNLL-17 pre-trained vectors are used, memory usage can reach to 32 GB depending on the pre-trained vector sizes.

4 Results

This section presents the parsing performance of our parser models on the CoNLL-18 Shared Task as well as the post-evaluation scores of our models.

4.1 Shared Task

Table 4 shows our official LAS, MLAS, and BLEX results in the CoNLL-18 Shared Task. The models that use the CoNLL-17 pre-trained word embeddings from (Ginter et al., 2017) are indicated in *pre-trained vectors* column. We also trained parser models using pre-trained word embeddings for *Czech-PDT*, *German-GSD*, *English-EWT*, *English-GUM*, *English-LinES*, *Spanish-AnCora*, *Indonesian-GSD*, *Latvian-LVTB*, *Swedish-LinES*, *Swedish-Talbanken*, and *Turkish-IMST*. However, we could not run these models with their corresponding embedding files

inside the TIRA virtual machine due to some unknown memory and disk issues.

Although the parser we used does not obtain competitive performance when compared with the best performing systems in the shared task, it achieves better performance on the treebanks with no training data when compared to its performance on treebanks with training data. We exclude the parallel UD treebanks from this judgment because one can get better performance on parallel UD treebanks by training the parser using the training data of the treebanks that have the same language with the parallel UD treebanks (e.g., the training data of *English-EWT* for *English-PUD*, *Czech-PDT* for *Czech-PUD* etc.). Due to time-constraints, we did not focus on the parallel UD treebanks and treated them as unknown languages.

4.2 Post-Evaluation

We performed another set of experiments using our models on the test data of UD version 2.2 data sets. The purpose of these experiments is to investigate the effect of our embedding models on parsing performance. Here we used the gold-standard conll-u files instead of the automatically annotated corpora by UDPipe, since our aim in these experiments is to observe the performance difference between our embedding models and the baseline embedding model.

In Table 5, we compare our models with the baseline model proposed in (Ballesteros et al., 2015). Due to time constraints, we trained all models without pre-trained word embeddings.

From the comparative results shown in Table 5, we observe that on the languages that have rich inflectional and derivational processes mostly by adding suffixes to words, our morphological features model outperforms the baseline model in terms of parsing scores. This is the case for the Bulgarian, Croatian, Czech, Basque, Gothic, Latin, Polish, Russian, Slovak, Slovene, North Sami, and Ukrainian languages.

The morphological features model is not suitable for the grammatical structure of Arabic, which has derivational morphology and it also fails to outperform the baseline in Romanic languages like French, Spanish, Catalan, Galician, and Portuguese. The possible reason behind this failure might be the analytic structure of the grammar of these languages where every morpheme is

Treebank	LAS Rank	LAS	MLAS	BLEX	Pre-trained vectors	Treebank	LAS Rank	LAS	MLAS	BLEX	Pre-trained vectors
<i>af-afribooms</i>	23	72.09	58.08	59.42	-	<i>hy-armtdp</i>	18	21.22	5.66	11.03	-
<i>ar-padt</i>	19	66.84	55.68	58.22	ar.vec	<i>id-gsd</i>	21	74.01	63.22	62.50	-
<i>bg-btb</i>	21	82.74	72.76	71.14	-	<i>it-isdt</i>	23	84.96	75.13	75.21	-
<i>br-keb</i>	10	10.59	0.43	2.21	-	<i>it-postwita</i>	17	67.94	54.15	54.81	it.vec
<i>bxr-bdt</i>	19	9.12	1.01	2.46	-	<i>ja-gsd</i>	21	72.21	58.04	59.78	-
<i>ca-ancora</i>	21	85.03	75.73	76.15	ca.vec	<i>ja-modern</i>	19	18.82	5.02	6.27	-
<i>cs-cac</i>	21	83.22	70.50	76.98	cs.vec	<i>kk-ktb</i>	21	12.88	2.22	3.95	kk.vec
<i>cs-fictree</i>	20	82.00	68.75	74.49	cs.vec	<i>kmr-mg</i>	21	14.12	2.59	6.65	-
<i>cs-pdt</i>	20	83.24	74.04	78.45	-	<i>ko-gsd</i>	17	74.99	69.00	62.99	ko.vec
<i>cs-pud</i>	21	69.60	56.86	63.13	-	<i>ko-kaist</i>	17	81.45	74.22	68.61	ko.vec
<i>cu-proiel</i>	22	60.39	47.99	52.96	cu.vec	<i>la-ittb</i>	18	76.32	67.78	72.02	la.vec
<i>da-ddt</i>	21	73.03	63.21	63.69	da.vec	<i>la-perseus</i>	21	41.94	26.20	28.93	la.vec
<i>de-gsd</i>	24	56.85	27.51	45.83	-	<i>la-proiel</i>	21	58.20	45.91	52.11	la.vec
<i>el-gdt</i>	20	82.11	65.23	68.59	el.vec	<i>lv-lvtb</i>	20	68.47	54.24	57.43	-
<i>en-ewt</i>	22	73.61	64.00	66.11	-	<i>nl-alpino</i>	21	75.94	60.99	63.59	nl.vec
<i>en-gum</i>	22	72.07	60.46	59.93	-	<i>nl-lassysmall</i>	22	74.48	61.40	62.77	nl.vec
<i>en-lines</i>	22	68.92	59.03	59.74	-	<i>no-bokmaal</i>	21	81.47	72.64	73.73	-
<i>en-pud</i>	23	69.28	56.99	60.04	-	<i>no-nynorsk</i>	22	78.48	68.49	69.75	-
<i>es-ancora</i>	21	83.02	73.95	74.40	-	<i>no-nynorskklia</i>	21	46.98	35.58	38.54	-
<i>et-edt</i>	17	75.47	67.74	64.42	et.vec	<i>pcm-nsc</i>	19	11.60	3.84	9.60	-
<i>eu-bdt</i>	19	70.41	57.83	63.36	eu.vec	<i>pl-lfg</i>	22	85.11	71.71	75.43	-
<i>fa-seraji</i>	19	79.62	73.09	69.84	fa.vec	<i>pl-sz</i>	22	79.71	61.89	69.76	-
<i>fi-ftb</i>	20	75.34	65.24	61.69	fi.vec	<i>pt-bosque</i>	19	82.62	67.97	72.83	pt.vec
<i>fi-pud</i>	20	60.07	53.14	48.14	-	<i>ro-rrt</i>	20	80.18	71.00	71.44	ro.vec
<i>fi-tdt</i>	20	75.68	67.56	61.21	fi.vec	<i>ru-syntagrus</i>	17	84.69	76.43	77.43	ru.vec
<i>fo-oft</i>	19	23.73	0.34	5.81	-	<i>ru-taiga</i>	24	45.93	29.00	31.09	-
<i>fr-gsd</i>	21	80.07	71.01	72.70	fr.vec	<i>sk-snk</i>	21	74.37	53.95	59.97	sk.vec
<i>fr-sequoia</i>	21	80.03	70.04	73.03	fr.vec	<i>sl-ssj</i>	20	76.78	62.89	68.32	sl.vec
<i>fr-spoken</i>	25	58.88	45.56	46.02	-	<i>sl-sst</i>	20	44.43	32.07	36.20	sl.vec
<i>fro-srcmf</i>	22	76.56	68.00	71.17	-	<i>sme-giella</i>	20	52.97	42.13	39.10	-
<i>ga-idt</i>	21	55.57	28.92	34.56	-	<i>sr-set</i>	23	75.79	62.84	66.68	-
<i>gl-ctg</i>	18	76.36	62.94	66.00	-	<i>sv-lines</i>	22	72.04	57.75	64.68	-
<i>gl-treegal</i>	21	63.43	46.11	48.83	-	<i>sv-pud</i>	22	64.55	37.53	48.00	-
<i>got-proiel</i>	21	58.18	44.69	50.63	-	<i>sv-talbanken</i>	21	76.93	67.83	68.50	-
<i>grc-perseus</i>	21	54.57	28.46	35.31	grc.vec	<i>th-pud</i>	7	0.70	0.04	0.52	-
<i>grc-proiel</i>	20	64.77	46.68	52.86	grc.vec	<i>tr-imst</i>	22	50.33	40.54	42.00	-
<i>he-htb</i>	21	57.28	43.42	45.94	he.vec	<i>ug-udt</i>	19	55.61	35.98	43.63	ug.vec
<i>hi-hdtb</i>	21	85.88	67.93	78.21	-	<i>uk-iu</i>	20	74.34	56.38	63.46	uk.vec
<i>hr-set</i>	22	75.91	56.67	67.43	hr.vec	<i>ur-udtb</i>	21	77.04	50.47	63.40	ur.vec
<i>hsb-ufal</i>	10	29.04	7.18	15.67	-	<i>vi-vtb</i>	20	39.06	25.90	27.61	vi.vec
<i>hu-szeged</i>	22	63.47	50.91	54.57	hu.vec	<i>zh-gsd</i>	22	56.43	46.55	51.20	zh.vec

Table 4: Our official results in the CoNLL-18 Shared Task.

Treebank	Embedding model	LAS		MLAS		BLEX	
		Baseline	Our model	Baseline	Our model	Baseline	Our model
<i>af-afribooms</i>	MF	82.16	82.80	73.17	74.35	75.48	76.37
<i>ar-padt</i>	MF	78.80	78.60	73.59	73.33	74.66	74.39
<i>bg-btb</i>	MF	86.52	87.41	80.88	82.00	81.33	82.38
<i>ca-ancora</i>	MF	87.21	87.03	80.56	80.23	81.07	80.79
<i>cs-cac</i>	MF	87.37	87.85	84.12	84.94	84.78	85.49
<i>cs-fictree</i>	MF	83.49	86.03	77.80	81.64	78.58	82.34
<i>cs-pdt</i>	MF	86.66	88.47	83.39	85.89	83.91	86.38
<i>cu-proiel</i>	MF	75.73	75.59	69.85	69.62	72.09	71.97
<i>da-ddt</i>	LS	77.34	78.04	71.45	71.48	72.97	73.33
<i>de-gsd</i>	MF	77.45	77.79	69.79	70.26	72.40	73.24
<i>el-gdt</i>	MF	83.22	83.98	74.83	76.69	75.57	77.53
<i>en-ewt</i>	MF	83.88	83.58	79.44	78.95	79.96	79.51
<i>en-gum</i>	MF	80.34	81.46	73.30	74.34	73.87	75.03
<i>en-lines</i>	MF	75.89	73.83	71.06	67.75	72.53	69.28
<i>es-ancora</i>	MF	86.71	85.55	80.65	78.97	81.17	79.61
<i>et-edt</i>	MF	80.25	81.49	76.59	78.28	77.40	78.98
<i>eu-bdt</i>	MF	74.07	74.65	69.97	71.21	71.69	72.74
<i>fi-ftb</i>	MF	82.76	83.88	77.86	79.02	78.54	79.89
<i>fi-idt</i>	MF	80.39	80.46	76.33	76.60	76.96	77.31
<i>fr-gsd</i>	MF	84.69	83.77	78.56	77.59	79.13	78.39
<i>fr-sequoia</i>	MF	83.16	82.49	76.75	75.96	77.38	76.40
<i>fr-spoken</i>	MF	67.99	68.70	57.82	58.19	58.68	59.04
<i>fro-srcmf</i>	MF	83.01	82.54	76.90	76.44	77.78	77.43
<i>ga-idt</i>	MF	61.02	63.23	45.21	47.98	48.68	51.90
<i>gl-ctg</i>	MF	81.56	80.70	70.76	69.41	75.38	74.25
<i>got-proiel</i>	MF	71.88	74.95	64.13	67.99	66.78	70.78
<i>grc-perseus</i>	MF	61.22	60.10	50.75	49.95	53.92	52.79
<i>grc-proiel</i>	MF	79.26	79.28	64.54	64.12	67.09	67.16
<i>he-htb</i>	MF	80.06	79.80	71.56	70.97	72.02	71.52
<i>hi-hdtb</i>	MF	92.11	91.51	87.64	86.72	88.38	87.46
<i>hr-set</i>	MF	80.12	81.36	74.75	76.45	76.22	77.90
<i>hu-szeged</i>	LS	64.00	68.33	56.44	62.55	59.75	66.11
<i>hy-armtdp</i>	MF	29.60	28.56	21.65	25.14	24.04	28.56
<i>it-isdt</i>	MF	88.91	89.23	82.77	83.34	83.20	83.79
<i>it-postwita</i>	MF	79.19	79.10	72.30	72.26	72.84	72.91
<i>kk-ktb</i>	LS	35.92	35.34	25.89	25.19	30.09	30.18
<i>la-ittb</i>	MF	83.86	85.37	79.06	80.93	80.02	82.10
<i>la-perseus</i>	MF	47.48	51.82	41.00	46.56	44.56	51.79
<i>la-proiel</i>	MF	68.81	70.95	62.15	64.66	64.95	67.11
<i>lv-lvtb</i>	MF	73.48	75.43	66.19	68.67	67.37	69.66
<i>nl-alpino</i>	MF	80.60	79.11	72.53	70.60	73.25	71.44
<i>nl-lassysmall</i>	MF	81.15	79.19	74.84	72.37	75.52	73.26
<i>no-bokmaal</i>	MF	88.53	88.22	84.48	83.87	84.96	84.46
<i>no-nynorsk</i>	MF	86.64	85.42	82.00	80.39	82.94	81.21
<i>no-nynorskklia</i>	MF	66.27	64.76	58.40	56.92	60.12	58.55
<i>pl-lfg</i>	MF	92.02	92.68	88.93	89.80	89.16	90.01
<i>pl-sz</i>	MF	85.86	89.56	81.34	86.50	82.02	87.17
<i>pt-bosque</i>	MF	83.28	83.20	75.64	74.85	76.95	76.28
<i>ro-rrt</i>	MF	81.22	80.84	74.26	74.03	75.55	75.36
<i>ru-syntagrus</i>	MF	88.01	88.14	84.35	84.86	84.72	85.24
<i>ru-taiga</i>	MF	48.95	56.57	40.82	50.74	42.74	52.33
<i>sk-snk</i>	MF	78.49	82.66	73.94	79.61	74.58	80.46
<i>sl-ssj</i>	MF	86.23	88.82	81.84	85.33	82.23	85.80
<i>sl-sst</i>	MF	64.47	65.41	57.67	59.38	59.31	61.22
<i>sme-giella</i>	MF	66.21	71.55	58.73	66.87	61.22	69.03
<i>sr-set</i>	MF	80.71	80.36	75.36	74.84	76.74	76.38
<i>sv-lines</i>	MF	76.86	77.43	72.98	73.75	74.13	74.72
<i>sv-talbanken</i>	MF	83.03	82.39	78.36	77.68	79.27	78.58
<i>tr-imst</i>	LS	55.45	56.74	49.29	50.42	50.45	51.97
<i>ug-udt</i>	LS	58.02	56.97	47.47	45.52	50.18	47.86
<i>uk-iu</i>	MF	76.10	78.87	70.57	74.66	70.77	74.95
<i>ur-udtb</i>	MF	86.07	86.04	79.43	79.77	80.75	81.02

Table 5: Comparison of our embedding models with the baseline char-based word embedding model explained in Section 2.1. MF stands for the morphological features embedding model and LS stands for the lemma-suffix embedding model.

	Singular	Plural	Singular	Plural
Nominative	<i>ember</i>	<i>ember-ek</i>	<i>adam</i>	<i>adam-lar</i>
Accusative	<i>ember-et</i>	<i>ember-ek-et</i>	<i>adam-ı</i>	<i>adam-lar-ı</i>
Dative	<i>ember-nek</i>	<i>ember-ek-nek</i>	<i>adam-a</i>	<i>adam-lar-a</i>
Locative	<i>ember-ben</i>	<i>ember-ek-ben</i>	<i>adam-da</i>	<i>adam-lar-da</i>

Table 6: Word-morpheme structure on the Hungarian word *ember* and the Turkish word *adam* (English meaning: *man*).

Treebank	Number of words	Embedding dimension	LAS	MLAS	BLEX
<i>tr-imst</i> without pre-trained embeddings	-	-	56.74	50.42	51.97
<i>tr-imst</i> with CoNLL-17 ud-word-embeddings	3,633,786	100	59.11	53.02	54.51
<i>tr-imst</i> with Facebook word-embeddings	416,051	300	59.69	53.56	54.98

Table 7: The effect of using pre-trained word embeddings on parsing performance on Turkish-IMST test data set.

an independent word. English, Hebrew, Hindi and Urdu languages are also categorized as mostly analytic languages which do not use inflections and have a low morpheme-per-word ratio (Moravcsik, 2013). Dutch, Norwegian, and Swedish languages have a very simplified inflectional grammar. So, these languages are not represented well using our morphology-based embedding models. Besides, our model is not the best choice for the languages that have high ratio of morphophonological modifications to the root word like Old Church Slavonic.

The lemma-suffix embedding model is applied to the Danish, Hungarian, Kazakh, Turkish, and Uyghur languages. The best performance is reached in the Hungarian language with more than 4% increase in LAS score. Our model outperforms the baseline in Turkish too. These languages are highly agglutinative languages where words may consist of several morphemes and the boundaries between morphemes are clearcut. In this type of languages, there is a one-to-one form-meaning correspondence and shape of a morpheme is invariant (Moravcsik, 2013). An example word-morpheme relationship in Hungarian and Turkish languages is shown in Table 6. As it can be seen from the table, this structure is very suitable to the lemma-suffix embedding model.

However, the lemma-suffix model fails to reach better performance than the baseline system on the Kazakh and Uyghur treebanks. A possible reason might be that our embedding model increases the complexity of the system unnecessarily for these languages with very little training data. Although Danish can be considered as an analytic language with a simplified inflectional grammar, the lemma-suffix model outperforms the baseline for this lan-

guage.

Table 7 shows the parsing scores of the parser with lemma-suffix embedding model on the test data of *Turkish-IMST* treebank version 2.2. We compared the parsing performances when the parser does not use pre-trained word embeddings, when it uses pre-trained embeddings from CoNLL-17 UD word embeddings, and when it uses pre-trained embeddings from word vectors trained on Wikipedia by Facebook (Bojanowski et al., 2017). From the results, we observe that the usage of pre-trained word vectors increases the parsing performance by great extent for Turkish. We also observe that Facebook word vectors outperform the CoNLL-17 UD word vectors, although the number of words in the Facebook vectors data set is much smaller than the number of words in the CoNLL-17 UD word vectors data set.

5 Conclusion

We introduced two morphology-based adaptations of the character-based word embedding model in (Ballesteros et al., 2015) and experimented with these models on the UD version 2.2 data set. The experiment results suggest that our models utilizing morphological information of words increases the parsing performance in agglutinative languages.

Acknowledgments

This work was supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) under grant number 117E971 and as a graduate scholarship.

References

- Miguel Ballesteros, Chris Dyer, and Noah A. Smith. 2015. Improved transition-based parsing by modeling characters instead of words with lstms. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 349–359. <https://doi.org/10.18653/v1/D15-1041>.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics* 5:135–146. <http://aclweb.org/anthology/Q17-1010>.
- Timothy Dozat, Peng Qi, and Christopher D. Manning. 2017. Stanford’s graph-based neural dependency parser at the conll 2017 shared task. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, pages 20–30. <https://doi.org/10.18653/v1/K17-3002>.
- 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)*. Association for Computational Linguistics, pages 334–343. <https://doi.org/10.3115/v1/P15-1033>.
- Filip Ginter, Jan Hajič, Juhani Luotolahti, Milan Straka, and Daniel Zeman. 2017. CoNLL 2017 shared task - automatically annotated raw texts and word embeddings. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University. <http://hdl.handle.net/11234/1-1989>.
- E.A. Moravcsik. 2013. *Introducing Language Typology*. Cambridge introductions to language and linguistics. Cambridge University Press. <https://books.google.com.tr/books?id=9sdanQAACAAJ>.
- 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, Portoro, Slovenia, pages 1659–1666.
- Joakim Nivre et al. 2018. Universal Dependencies 2.2. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University, Prague, <http://hdl.handle.net/11234/1-1983xxx>. <http://hdl.handle.net/11234/1-1983xxx>.
- 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. https://doi.org/10.1007/978-3-319-11382-1_22.
- Haşim Sak, Tunga Güngör, and Murat Saraçlar. 2008. Turkish language resources: Morphological parser, morphological disambiguator and web corpus. In *Advances in natural language processing*, Springer, pages 417–427.
- 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, Portoro, Slovenia.
- Dan Zeman et al. 2018a. Universal Dependencies 2.2 CoNLL 2018 shared task development and test data. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University, Prague, <http://hdl.handle.net/11234/1-2184>. <http://hdl.handle.net/11234/1-2184>.
- Daniel Zeman, Jan Hajič, Martin Popel, Martin Potthast, Milan Straka, Filip Ginter, Joakim Nivre, and Slav Petrov. 2018b. CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies. In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*. Association for Computational Linguistics, Brussels, Belgium, pages 1–20.