

A General-Purpose Tagger with Convolutional Neural Networks

Xiang Yu and Agnieszka Falenska and Ngoc Thang Vu

Institut für Maschinelle Sprachverarbeitung

Universität Stuttgart

{xiangyu, falensaa, thangvu}@ims.uni-stuttgart.de

Abstract

We present a general-purpose tagger based on convolutional neural networks (CNN), used for both composing word vectors and encoding context information. The CNN tagger is robust across different tagging tasks: without task-specific tuning of hyper-parameters, it achieves state-of-the-art results in part-of-speech tagging, morphological tagging and supertagging. The CNN tagger is also robust against the out-of-vocabulary problem; it performs well on artificially unnormalized texts.

1 Introduction

Recently, character composition models have shown great success in many NLP tasks, mainly because of their robustness in dealing with out-of-vocabulary (OOV) words by capturing subword informations. Among the character composition models, bidirectional long short-term memory (LSTM) models and convolutional neural networks (CNN) are widely applied in many tasks, e.g. part-of-speech (POS) tagging (dos Santos and Zadrozny, 2014; Plank et al., 2016), named entity recognition (dos Santos and Guimarães, 2015), language modeling (Ling et al., 2015; Kim et al., 2016), machine translation (Costa-jussà and Fonollosa, 2016) and dependency parsing (Ballesteros et al., 2015; Yu and Vu, 2017).

In this paper, we present a state-of-the-art general-purpose tagger that uses CNNs both to compose word representations from characters and to encode context information for tagging.¹ We show that the CNN model is more capable than

the LSTM model for both functions, and more stable for unseen or unnormalized words, which is the main benefit of character composition models.

Yu and Vu (2017) compared the performance of CNN and LSTM as character composition model for dependency parsing, and concluded that CNN performs better than LSTM. In this paper, we show that this is also the case for POS tagging. Furthermore, we extend the scope to morphological tagging and supertagging, in which the tag set is much larger or long-distance dependencies between words are more important.

In these three tagging tasks, we compare our tagger with the bilstm-aux tagger (Plank et al., 2016) and the CRF-based morphological tagger MarMot (Müller et al., 2013) as baselines. The CNN tagger shows robust performance across the three tasks, and achieves the highest average accuracies in all tasks. It considerably outperforms the LSTM tagger in morphological tagging and both baselines in supertagging.

To test the robustness of the taggers against the OOV problem, we also conduct experiments on unnormalized text by artificially corrupting words in the normal dev sets. With the increasing degree of unnormalization, the performance of the CNN tagger degrades much slower than the other two, which suggests that the CNN tagger is more robust against unnormalized text.

Therefore we conclude that our CNN tagger is a robust state-of-the-art general-purpose tagger that can effectively compose word representation from characters and encode context information.

2 Model

Our proposed CNN tagger has two main components: the character composition model and the context encoding model. Both components are essentially very similar CNN models, capturing dif-

¹The tagger is available at <http://www.ims.uni-stuttgart.de/institut/mitarbeiter/xiangyu/index.en.html>

ferent levels of information: the first CNN captures morphological information from character n-grams, the second one captures contextual information from word n-grams. Figure 1 shows a diagram of both models of the tagger.




Figure 1: Diagram of the CNN tagger.

2.1 Character Composition Model

The character composition model is similar to Yu and Vu (2017), where several convolution filters are used to capture character n-grams of different sizes. The outputs of each convolution filter are fed through a max pooling layer, and the pooling outputs are concatenated to represent the word.

2.2 Context Encoding Model

The context encoding model captures the context information of the target word by scanning through the word representations of its context window. The word representation could be only word embeddings (\vec{w}), only composed vectors (\vec{c}), or the concatenation of both ($\vec{w} + \vec{c}$).

A context window consists of N words to both sides of the target word and the target word itself. To indicate the target word, we concatenate a binary feature to each of the word representations with 1 indicating the target and 0 otherwise, similar to Vu et al. (2016). Additional to the binary feature, we also concatenate a position embedding to encode the relative position of each context word, similar to Gehring et al. (2017).

2.3 Hyper-parameters

For the character composition model, we take a fixed input size of 32 characters for each word, with padding on both sides or cutting from the middle if needed. We apply four convolution filters with sizes of 3, 5, 7, and 9. Each filter has an output channel of 25 dimensions, thus the composed vector is 100-dimensional. We apply Gaussian noise with standard deviation of 0.1 on the composed vector during training.

For the context encoding model, we take a context window of 15 (7 words to both sides of the target word) as input and predict the tag of the target word. We also apply four convolution filters with sizes of 2, 3, 4 and 5, each filter is stacked by another filter with the same size, and the output has 128 dimensions, thus the context representation is 512-dimensional. We apply one 512-dimensional hidden layer with ReLU non-linearity before the prediction layer. We apply dropout with probability of 0.1 after the hidden layer during training.

The model is trained with averaged stochastic gradient descent with a learning rate of 0.1, momentum of 0.9 and mini-batch size of 100. We apply L2 regularization with a rate of 10^{-5} on all the parameters of the network except the embeddings.

3 Experiments

3.1 Data

We use treebanks from version 1.2 of Universal Dependencies² (UD), and in the case of several treebanks for one language, we only use the canonical one. There are in total 22 treebanks, as in Plank et al. (2016).³ Each treebank splits into train, dev, and test sets, we use the dev sets for early stop training.

In order to compare to more previous works on POS tagging, we additionally experiment POS tagging on the more established Penn Treebank Wall Street Journal (WSJ) data set (Marcus et al., 1993). We use the standard splitting, where sections 0-18 are used for training, 19-21 for tuning, and 22-24 for testing (Collins, 2002).

3.2 Tasks

We evaluate the taggers on three tagging tasks: POS tagging (**POS**), morphological tagging (**MORPH**) and supertagging (**SUPER**).

For POS tagging we use Universal POS tags, which are an extension of Petrov et al. (2012). The universal tag set tries to capture the “universal” properties of words and facilitate cross-lingual learning. Therefore the tag set is very coarse and leaves out most of the language-specific properties to morphological features.

Morphological tags encode the language-specific morphological features of the words, e.g., number, gender, case. They are represented in the

²<http://universaldependencies.org>

³We use all training data for Czech, while Plank et al. (2016) only use a subset.

UD treebanks as one string which contains several key-value pairs of morphological features.⁴

Supertags (Joshi and Bangalore, 1994) are tags that encode more syntactic information than standard POS tags, e.g. the head direction or the sub-categorization frame. We use dependency-based supertags (Foth et al., 2006) which are extracted from the dependency treebanks. Adding such tags into feature models of statistical dependency parsers significantly improves their performance (Ouchi et al., 2014; Faleńska et al., 2015). Supertags can be designed with different levels of granularity. We use the standard Model 1 from Ouchi et al. (2014), where each tag consists of head direction, dependency label and dependent directions. The SUPER task is more difficult than POS and MORPH because it generally requires taking long-distance dependencies between words into consideration.

These three tagging tasks differ strongly in tag set sizes. Generally, the POS set sizes for all the languages are no more than 17 and SUPER set sizes are around 200. When treating morphological features as a string (i.e. not splitting into key-value pairs), the sizes of the MORPH tag sets range from about 100 up to 2000.

3.3 Setups

As baselines to our models, we take the two state-of-the-art taggers MarMot⁵ (denoted as **CRF**) and bilstm-aux⁶ (denoted as **LSTM**). We train the taggers with the recommended hyper-parameters from the documentations.

To ensure a fair comparison (especially between LSTM and CNN), we generally treat the three tasks equally, and do not apply task-specific tuning on them, i.e., using the same features and same model hyper-parameters in each single task. Also, we do not use any pre-trained word embeddings.

For the LSTM tagger, we use the recommended hyper-parameters from the documentation⁷ including 64-dimensional word embeddings (\vec{w}) and 100-dimensional composed vectors (\vec{c}). We train the \vec{w} , \vec{c} and $\vec{w} + \vec{c}$ models as in Plank et al. (2016).

⁴German, French and Indonesian do not have MORPH tags in UD-1.2, thus not evaluated in this task.

⁵<http://cistern.cis.lmu.de/marmot/>

⁶<https://github.com/bplank/bilstm-aux>

⁷We use the most recent version of the tagger and stacking 3 layers of LSTM as recommended. The average accuracy for POS in our evaluation is slightly lower than reported in the paper, presumably due to different versions of the tagger, but it does not influence the conclusion.

We train the CNN taggers with the same dimensionalities for word representations.

For the CRF tagger, we predict POS and MORPH jointly as in the standard setting, which performs much better than with separate predictions, as shown in Müller et al. (2013). Also, the CRF tagger splits the morphological tags into key-value pairs, whereas the two neural-based taggers treat the whole string as a tag.⁸ We predict SUPER as a separate task.

3.4 Results

The test results for the three tasks are shown in Table 1 in three groups. The first group of seven columns are the results for POS, where both LSTM and CNN have three variations of input features: word only (\vec{w}), character only (\vec{c}) and both ($\vec{w} + \vec{c}$). For MORPH and SUPER, we only use the $\vec{w} + \vec{c}$ setting for both LSTM and CNN.

On macro-average, three taggers perform close in the POS task, with the CNN tagger being slightly better. In the MORPH task, CNN is again slightly ahead of CRF, while LSTM is about 2 points behind. In the SUPER task, CNN outperforms both taggers by a large margin: 2 points higher than LSTM and 8 points higher than CRF.

While considering the input features of the LSTM and CNN taggers, both taggers perform close with \vec{w} as input, which suggests that the two taggers are comparable in encoding context for POS. However, with only \vec{c} , CNN performs much better than LSTM (95.54 vs. 92.61), and close to $\vec{w} + \vec{c}$ (96.18). Also, \vec{c} consistently outperforms \vec{w} for all languages with CNN. This suggests that the CNN model alone is capable of learning most of the information that the word-level model can learn, while the LSTM model is not.

The more interesting cases are MORPH and SUPER, where CNN performs much higher than LSTM. One potential explanation for the considerably large difference is that the LSTM tagger may be more sensitive to hyper-parameters and requires task specific tuning. We use the same setting which is tuned for the POS task, thus it underperforms in the other tasks. Another factor could be the large tag sets in MORPH tagging task, which are larger than POS in orders of magnitudes, especially for *cs*, *eu*, *fi*, *hr*, *pl*, and *sl*, all of which have more than 500 distinct tags, and the LSTM

⁸Since we use the CRF tagger as a non-neural baseline, we prefer to use the settings which maximize its performances rather than rigorously equal but suboptimal settings.

	CRF	Pos						MORPH			SUPER		
		\vec{w}	LSTM \vec{c}	$\vec{w} + \vec{c}$	\vec{w}	CNN \vec{c}	$\vec{w} + \vec{c}$	CRF	LSTM $\vec{w} + \vec{c}$	CNN $\vec{w} + \vec{c}$	CRF	LSTM $\vec{w} + \vec{c}$	CNN $\vec{w} + \vec{c}$
ar	98.83	95.05	98.35	98.88	95.30	98.89	99.00	98.11	97.91	98.45	79.67	83.70	85.51
bg	98.11	94.96	96.94	98.07	95.25	97.79	98.20	95.12	92.28	94.85	78.91	85.91	87.64
cs	98.74	96.12	92.98	98.40	96.36	98.35	98.79	93.81	90.21	94.45	76.33	81.43	87.46
da	95.96	91.74	94.29	96.06	92.08	95.24	95.92	95.50	94.15	95.14	73.83	81.00	81.82
de	92.77	89.91	88.97	92.57	90.21	92.44	92.73	-	-	-	70.56	77.58	79.69
en	94.49	91.58	88.99	94.17	92.64	93.76	94.76	95.69	95.45	95.88	75.57	83.27	85.87
es	95.28	93.27	91.41	94.62	93.95	95.36	95.65	96.14	95.26	96.34	78.07	83.80	86.27
eu	94.79	88.70	89.80	93.99	89.69	94.31	94.94	89.60	84.32	89.06	70.44	77.88	80.43
fa	96.82	95.67	94.73	96.95	95.97	96.12	97.12	96.56	96.37	96.50	76.76	83.21	83.25
fi	95.79	87.78	84.41	94.16	88.24	94.33	95.31	94.33	87.33	93.82	70.69	76.65	82.63
fr	95.98	94.34	91.82	95.85	94.56	95.68	96.27	-	-	-	78.36	84.01	85.44
he	95.48	93.81	92.96	95.62	93.81	94.68	96.04	92.92	91.27	93.29	76.73	82.56	85.44
hi	96.36	95.66	91.12	96.23	96.04	95.77	96.69	90.93	90.78	92.11	85.54	89.62	90.08
hr	95.56	88.10	94.47	94.69	88.92	94.76	95.05	87.25	84.56	87.73	71.42	77.77	79.27
id	93.51	90.40	90.76	92.97	91.15	92.32	93.44	-	-	-	75.37	80.55	81.63
it	97.74	96.04	94.64	97.55	96.54	97.08	97.62	97.63	97.13	97.47	84.02	89.10	90.89
nl	91.03	85.09	86.52	92.23	83.74	92.05	93.11	92.32	91.26	93.12	67.04	77.71	79.68
no	97.61	94.39	93.32	97.49	94.60	97.05	97.65	96.03	94.85	95.74	79.99	86.45	89.41
pl	96.92	89.53	95.05	96.30	90.48	96.41	96.83	87.74	82.34	87.13	76.09	80.00	83.45
pt	97.78	94.20	94.95	97.53	94.41	97.22	97.46	94.99	94.75	95.76	78.68	86.02	87.42
sl	96.60	90.43	96.35	97.42	91.02	96.89	97.16	90.41	86.47	91.94	76.35	85.67	86.45
sv	96.23	93.04	94.48	96.20	93.27	95.38	96.28	95.65	94.08	95.30	73.81	81.04	83.34
avg	96.02	92.26	92.61	95.82	92.65	95.54	96.18	93.72	91.62	93.90	76.10	82.50	84.69

Table 1: Tagging accuracies of the three taggers in the three tasks on the test sets of UD 1.2, the highest accuracy for each task on each language is marked in boldface.

tagger performs poorly on these languages. In the SUPER task, where the information from long-distance context is more important, CNN performs much better than both CRF and LSTM. CRF simply has a much smaller context window, thus the poor performance. The LSTM model theoretically can model long-distance contexts, but the information may gradually fade away during the recurrence, whereas the CNN model treat all words equally as long as they are in the context window.

On the more established WSJ data set, Table 2 shows the tagging performances of the CNN model along with some previous works as reference. Generally, the differences among the taggers are very small, we could not conclude any one being considerably better on this data set. This result is expected since English is not a morphologically rich language and WSJ is large data set and has a relatively low OOV rate. Note that the Convnet tagger by dos Santos and Zadrozny (2014) used pre-trained word embeddings while our CNN tagger does not.

3.5 Unnormalized Text

It is a common scenario to use a model trained with news data to process text from social media, which could include intentional or unintentional

WSJ	Accuracy
CRF (Müller et al., 2013)	97.30
Convnet (dos Santos and Zadrozny, 2014)	97.32
bi-LSTM (Ling et al., 2015)	97.36
bi-LSTM (Plank et al., 2016)	97.22
CNN (this work)	97.30

Table 2: Tagging accuracy on the WSJ test set.

misspellings. Unfortunately, we do not have social media data for all the languages. However, we design an experiment to simulate unnormalized text, by systematically editing the words in the dev sets with one of the four operations: insertion, deletion, substitution, and swap. For example, if we modify a word *abcdef* at position 2 (0-based), the modified words would be *abxcdef*, *abdef*, *abxdef*, and *abdcef*, where *x* is a random character from the alphabet of the language.

For each operation, we create a group of modified dev sets, where all words longer than two characters are edited by the operation with a probability of 0.25, 0.5, 0.75, or 1. For each language, we use the models trained on the normal training sets and predict POS for the modified dev sets. The average accuracies are shown in Figure 2.

Generally, all models suffer from the increasing degrees of unnormalization, but CNN always de-

grades the least and slowest. In the extreme case where almost all words are unnormalized, CNN performs 4 to 8 points higher than LSTM and 4 to 12 points higher than CRF. This suggests that the CNN is more robust to misspelt words.

While looking into the specific cases of misspelling, CNN is less sensitive to substitution, while insertion and deletion have stronger effect, and swap degrades its performance the most. In the case of substitution, the distortion to the character n-gram patterns are smaller than varying the lengths, i.e. insertion and deletion, thus has smaller negative impact. However, in the case of swap, the effect is similar to substituting two characters instead of one, thus larger degradation. LSTM and CRF on the other hand, are affected the most by substitution.




Figure 2: POS tagging accuracies on the dev sets with the four modifications of different degrees.

4 Conclusion

In this paper, we propose a general-purpose tagger that uses two CNNs for both character composition and context encoding. On the universal dependency treebanks (v1.2), the tagger achieves state-of-the-art or comparable results for POS tagging and morphological tagging, and to the best of our knowledge, it performs by far the best for supertagging. The tagger works well across different tagging tasks without tuning the hyper-parameters, and it is also robust against unnormalized text.

Our tagger uses a greedy window-based approach, which mainly aims at showing the effectiveness of CNN in composing word representations and encoding contexts. However, a globally normalized decoding method, e.g. beam-search

or sentence-level inference as in Collobert et al. (2011), could potentially further improve the tagger’s performance, which is left for future work.

Acknowledgments

This work was supported by the German Research Foundation (DFG) in project D8 of SFB 732. We also thank our colleagues in the IMS and the anonymous reviewers for the suggestions.

References

- Miguel Ballesteros, Chris Dyer, and A. Noah 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>.
- Michael Collins. 2002. Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing-Volume 10*. Association for Computational Linguistics, pages 1–8.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12(Aug):2493–2537.
- R. Marta Costa-jussà and R. José A. Fonollosa. 2016. Character-based neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, pages 357–361. <https://doi.org/10.18653/v1/P16-2058>.
- Cicero dos Santos and Victor Guimarães. 2015. Boosting named entity recognition with neural character embeddings. In *Proceedings of the Fifth Named Entity Workshop*. Association for Computational Linguistics, pages 25–33. <https://doi.org/10.18653/v1/W15-3904>.
- Cicero dos Santos and Bianca Zadrozny. 2014. Learning character-level representations for part-of-speech tagging. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*. pages 1818–1826.
- Agnieszka Faleńska, Anders Björkelund, Özlem Çetinoğlu, and Wolfgang Seeker. 2015. Stacking or supertagging for dependency parsing – what’s the difference? In *Proceedings of the 14th International Conference on Parsing Technologies*. Association for Computational Linguistics, Bilbao, Spain, pages 118–129. <http://www.aclweb.org/anthology/W15-2215>.

- Kilian A. Foth, Tomas By, and Wolfgang Menzel. 2006. [Guiding a constraint dependency parser with supertags](#). In *ACL 2006, 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, Sydney, Australia, 17-21 July 2006*. <http://aclweb.org/anthology/P06-1037>.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. 2017. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*.
- Aravind K. Joshi and Srinivas Bangalore. 1994. [Disambiguation of Super Parts of Speech \(or Supertags\): Almost Parsing](#). In *Proceedings of the 15th Conference on Computational Linguistics - Volume 1*. Association for Computational Linguistics, Stroudsburg, PA, USA, COLING '94, pages 154–160. <https://doi.org/10.3115/991886.991912>.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. Character-aware neural language models. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*. AAAI Press, pages 2741–2749.
- Wang Ling, Chris Dyer, W. Alan Black, Isabel Trancoso, Ramon Fernandez, Silvio Amir, Luis Marujo, and Tiago Luis. 2015. [Finding function in form: Compositional character models for open vocabulary word representation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 1520–1530. <https://doi.org/10.18653/v1/D15-1176>.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. [Building a large annotated corpus of english: The penn tree-bank](#). *Computational Linguistics* pages 313–330. <http://dl.acm.org/citation.cfm?id=972470.972475>.
- Thomas Müller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient Higher-Order CRFs for Morphological Tagging. In *In Proceedings of EMNLP*.
- Hiroki Ouchi, Kevin Duh, and Yuji Matsumoto. 2014. [Improving Dependency Parsers with Supertags](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, Volume 2: Short Papers*. Association for Computational Linguistics, Gothenburg, Sweden, pages 154–158. <http://www.aclweb.org/anthology/E14-4030>.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. A universal part-of-speech tagset. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*. European Language Resources Association (ELRA), Istanbul, Turkey.
- Barbara Plank, Anders Søgaard, and Yoav Goldberg. 2016. [Multilingual part-of-speech tagging with bidirectional long short-term memory models and auxiliary loss](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, pages 412–418. <https://doi.org/10.18653/v1/P16-2067>.
- Ngoc Thang Vu, Heike Adel, Pankaj Gupta, and Hinrich Schütze. 2016. [Combining recurrent and convolutional neural networks for relation classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pages 534–539. <https://doi.org/10.18653/v1/N16-1065>.
- Xiang Yu and Ngoc Thang Vu. 2017. Character composition model with convolutional neural networks for dependency parsing on morphologically rich languages. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Vancouver, Canada.