

State-of-the-art Chinese Word Segmentation with Bi-LSTMs

Ji Ma Kuzman Ganchev David Weiss

Google AI Language

{maji, kuzman, djweiss}@google.com

Abstract

A wide variety of neural-network architectures have been proposed for the task of Chinese word segmentation. Surprisingly, we find that a bidirectional LSTM model, when combined with standard deep learning techniques and best practices, can achieve better accuracy on many of the popular datasets as compared to models based on more complex neural-network architectures. Furthermore, our error analysis shows that out-of-vocabulary words remain challenging for neural-network models, and many of the remaining errors are unlikely to be fixed through architecture changes. Instead, more effort should be made on exploring resources for further improvement.

1 Introduction

Neural networks have become ubiquitous in natural language processing. For the word segmentation task, there has been a growing body of work exploring novel neural network architectures for learning useful representation and thus better segmentation prediction (Pei et al., 2014; Ma and Hinrichs, 2015; Zhang et al., 2016a; Liu et al., 2016; Cai et al., 2017; Wang and Xu, 2017).

We show that properly training and tuning a relatively simple architecture with a minimal feature set and greedy search achieves state-of-the-art accuracies and beats more complex neural-network architectures. Specifically, the model itself is a straightforward stacked bidirectional LSTM (Figure 1) with just two input features at each position (character and bigram). We use three widely recognized techniques to get the most performance out of the model: pre-trained embeddings (Yang et al., 2017; Zhou et al., 2017), dropout (Srivastava et al., 2014), and hyperparameter tuning (Weiss et al., 2015; Melis et al., 2018). These results have important ramifications for further model development. Unless best practices are followed, it is

difficult to compare the impact of modeling decisions, as differences between models are masked by choice of hyperparameters or initialization.

In addition to the simpler model we present, we also aim to provide useful guidance for future research by examining the errors that the model makes. About a third of the errors are due to annotation inconsistency, and these can only be eliminated with manual annotation. The other two thirds are those due to out-of-vocabulary words and those requiring semantic clues not present in the training data. Some of these errors will be almost impossible to solve with different model architectures. For example, while 抽象概念 (abstract concept) appears as one word at test time, any model trained only on the MSR dataset will segment it as two words: 抽象 (abstract) and 概念 (concept), which are seen in the training set 28 and 90 times, respectively, and never together. Thus, we expect that iterating on model architectures will give diminishing returns, while leveraging external resources such as unlabeled data or lexicons is a more promising direction.

In sum, this work contributes two significant pieces of evidence to guide further development in Chinese word segmentation. First, comparing different model architectures requires careful tuning and application of best practices in order to obtain rigorous comparisons. Second, iterating on neural architectures may be insufficient to solve the remaining classes of segmentation errors without further efforts in data collection.

2 Model

Our model is relatively simple. Our approach uses long short-term memory neural networks architectures (LSTM) since previous work has found success with these models (Chen et al., 2015; Zhou et al., 2017, inter alia). We use two features: uni-

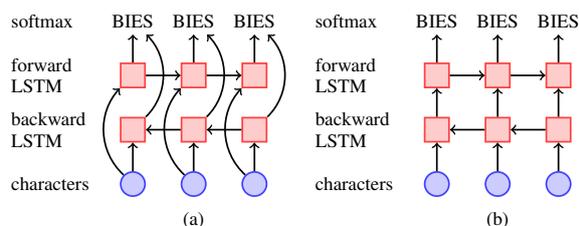


Figure 1: Bi-LSTM models: (a) non-stacking, (b) stacking. Blue circles are input (char and char bigram) embeddings. Red squares are LSTM cells. BIES is a 4-way softmax.

grams and bi-grams of characters at each position. These features are embedded, concatenated, and fed into a stacked bidirectional LSTM (see Figure 1) with two total layers of 256 hidden units each. The softmax layer of the bi-LSTM predicts Begin/Inside/End/Single tags encoding the relationship from characters to segmented words.

In the next sections we describe the best practices we used to achieve state-of-the-art performance from this architecture. Note that all of these practices and techniques are derived from related work, which we describe.

Recurrent Dropout. Contrary to the recommendation of Zaremba et al. (2014), we apply dropout to the recurrent connections of our LSTMs, and we see similar improvements when following the recipe of Gal and Ghahramani (2016) or simply sample a new dropout mask at every recurrent connection.

Hyperparameters. We use the momentum-based averaged SGD procedure from (Weiss et al., 2015) to train the model, with few additions. We normalized each gradient to be at most unit norm, and used asynchronous SGD updates to speed up training time. For each configuration we evaluated, we trained different settings of a manually tuned hyperparameter grid, varying the initial learning rate, learning rate schedule, and input and recurrent dropout rates. We fixed the momentum parameter $\mu = 0.95$. The full list of hyperparameters is given in Table 2. We show the impact of this tuning procedure in Table 7, which we found was crucial to measure the best performance of the simple architecture.

Pretrained Embeddings. Pre-training embedding matrices from automatically gathered data is a powerful technique that has been applied to many NLP problems for several years (e.g. Collobert et al. (2011); Mikolov et al. (2013)). We

	Train	Development	Test
AS	4,903,564	546,017	122,610
CTIYU	1,309,208	146,422	40,936
MSR	2,132,480	235,911	106,873
CTB6	641,368	59,954	81,578
CTB7	950,138	59,954	81,578
PKU	994,822	115,125	104,372
UD	98,608	12,663	12,012

Table 1: Statistics of training, development and test set.

pretrain the character embeddings and character-bigram embeddings using wang2vec¹ (Ling et al., 2015), which modifies word2vec by incorporating character/bigram order information during training. Note that this idea has been used in segmentation previously by Zhou et al. (2017), but they also augment the contexts by adding the predictions of a baseline segmenter as an additional context. We experimented with both treating the pre-trained embeddings as constants or fine-tuning on the particular datasets.

Other Related Work. Recently, a number of different neural network based models have been proposed for word segmentation task. One common approach is to learn word representation through the characters of that word. For example, Liu et al. (2016) runs bi-directional LSTM over characters of the word candidate and then concatenate bi-directional LSTM outputs at both end points. Cai et al. (2017) adopts a gating mechanism to control relative importance of each character in the word candidate.

Besides modeling word representation directly, sequential labeling is another popular approach. For instance, Zheng et al. (2013) and Pei et al. (2014) predict the label of a character based context of a fixed sized local window. Chen et al. (2015) extends the approach by using LSTMs to capture potential long distance information. Both Chen et al. (2015) and Pei et al. (2014) use a transition matrix to model interaction between adjacent tags. Zhou et al. (2017) conduct rigorous comparison and show that such transition matrix rarely improves accuracy. Our model is similar to Zhou et al. (2017), except that we stack the backward LSTM on top of the forward one, which improves accuracy as shown in later section.

Our model is also trained via a simple maximum likelihood objective. In contrast, other state-

¹<https://github.com/wlin12/wang2vec>

	AS	CITYU	CTB6	CTB7	MSR	PKU	UD
Liu et al. (2016)	—	—	95.9	—	97.3	96.8	—
Yang et al. (2017)	95.7	96.9	96.2	—	97.5	96.3	—
Zhou et al. (2017)	—	—	96.2	—	97.8	96.0	—
Cai et al. (2017)	—	95.6	—	—	97.1	95.8	—
Kurita et al. (2017)	—	—	—	96.2	—	—	—
Chen et al. (2017)	94.6	95.6	96.2	—	96.0	94.3	—
Qian and Liu (2017)	—	—	—	—	—	—	94.6
Wang and Xu (2017)	—	—	—	—	98.0	96.5	—
Ours (fix embedding)	96.2	97.2	96.7	96.6	97.4	96.1	96.9
Ours (update embedding)	96.0	96.8	96.3	96.0	98.1	96.1	96.0

Table 2: The state of the art performance on different datasets. For Kurita et al. (2017) and Chen et al. (2017) we report their best systems (segpos+dep and Model-I-ADV respectively). †Not directly comparable to the rest of the table due to the usage of an external dictionary. Our bolded results are significantly better ($p < 0.05$ bootstrap resampling) except on MSR.

	AS	CITYU	CTB6	CTB7	MSR	PKU	UD
Liu et al. (2016)	—	—	94.6	—	94.8	94.9	—
Zhou et al. (2017)	—	—	94.9	—	97.2	95.0	—
Cai et al. (2017)	95.2	95.4	—	—	97.0	95.4	—
Wang and Xu (2017)	—	—	—	—	96.7	94.7	—
Ours	95.5	95.7	95.5	95.6	97.5	95.4	94.6

Table 3: Performance of recent neural network based models without using pretrained embeddings. Our model’s wins are statistically significantly better than prior work ($p < 0.05$ bootstrap resampling), except on PKU.

of-the-art models use a non-greedy approach to training and inference, e.g. Yang et al. (2017) and Zhang et al. (2016b).

3 Experiments

Data. We conduct experiments on the following datasets: Chinese Penn Treebank 6.0 (CTB6) with data split according the official document; Chinese Penn Treebank 7.0 (CTB7) with recommended data split (Wang et al., 2011); Chinese Universal Treebank (UD) from the Conll2017 shared task (Zeman et al., 2017) with the official data split; Dataset from SIGHAN 2005 bake-off task (Emerson, 2005). Table 1 shows statistics of each data set. For each of the SIGHAN 2005 dataset, we randomly select 10% training data as development set. We convert all digits, punctuation and Latin letters to half-width, to handle full/half-width mismatch between training and test set. We train and evaluate a model for each of the dataset, rather than train one model on the union of all dataset. Following Yang et al. (2017), we convert AS and CITYU to simplified Chinese.

3.1 Main Results

Table 2 contains the state-of-the-art results from recent neural network based models, together with

the performance of our model. Table 3 contains results achieved without using any pretrained embeddings.

Our model achieves the best results among NN models on 6/7 datasets. In addition, while the majority of datasets work the best if the pre-trained embedding matrix is treated as constant, the MSR dataset is an outlier: fine-tuning embeddings yields a very large improvement. We observe that the likely cause is a low OOV rate in the MSR evaluation set compared to other datasets.

3.2 Ablation Experiments

To see which decisions had the greatest impact on the result, we performed ablation experiments on the holdout sets of the different corpora. Starting with our proposed system², we remove one decision, perform hyperparameter tuning, and see the change in performance. The results are summarized in Table 6. Negative numbers in Table 6 correspond to decreases in performance for the ablated system. Note that although each of the components help performance on average, there are cases where we observe no impact. For example using recurrent dropout on AS and MSR rarely af-

²Based on development set accuracy, we keep the pre-trained embedding fixed for all datasets except MSR and AS.

Parameter	Values
Char embedding size	[64]
Bigram embedding size	[16, 32, 64]
Learning rate	[0.04, 0.035, 0.03]
Decay steps	[32K, 48K, 64K]
Input dropout rate	[0.15, 0.2, 0.25, 0.3, 0.4, 0.5, 0.6]
LSTM dropout rate	[0.1, 0.2, 0.3, 0.4]

Table 4: Hyperparameter settings.

	AS	CITYU	CTB6	CTB7	MSR	PKU	UD
OOV %	4.2	7.5	5.6	5.0	2.7	3.6	12.4
Recall % (random embedding)	65.7	75.1	73.4	74.1	71.0	66.0	81.1
Recall % (pretrain embedding)	70.7	87.5	85.4	85.6	80.0	78.8	89.7

Table 5: Test set OOV rate, together with OOV recall achieved with randomly initialized and pretrained embeddings, respectively.

fects accuracy.

We next investigate how important the hyperparameter tuning is to this ablation. In the main result, we tuned each model separately for each dataset. What if instead, each model used a single hyperparameter configuration for all datasets? In Table 7, we compare fully tuned models with those that share hyperparameter configurations across dataset for three settings of the model. We can see that hyperparameter tuning consistently improves model accuracy across all settings.

3.3 Error Analysis

In order to guide future research on Chinese word segmentation, it is important to understand the types of errors that the system is making. To get a sense of this, we randomly selected 54 and 50 errors from the CTB-6 and MSR test set, respectively. We then manually analyzed them.

The model learns to remember words it has seen, especially for high frequency words. It also learns the notion of prefixes/suffixes, which aids predicting OOV words, a major source of segmentation errors (Huang and Zhao, 2007). Using pretrained embeddings enables the model to expand the set of prefixes/suffixes through their nearest neighbors in the embedding spaces, and therefore further improve OOV recall (on average, using pretrained embeddings contributes to 10% OOV recall improvement, also see Table 5 for more details).

Nevertheless, OOV remains challenging especially for those that can be divided into words fre-

quently seen in the training data, and most (37 out of 43) of the oversegmentation errors are due to this. For instance, the model incorrectly segmented the OOV word 抽象概念 (abstract concept) as 抽象 (abstract) and 概念 (concept). 抽象 and 概念 are seen in the training set for 28 times and 90 times, respectively. Unless high coverage dictionaries are used, it is difficult for any supervised model to learn not to follow this trend in the training data.

In addition, the model sometimes struggles when a prefix/suffix can also be a word by itself. For instance, 权 (right/power) frequently serves as a suffix, such as 管理权 (right of management), 立法权 (right of legislation) and 终审权 (right of final judgment). When the model encounters 下放 (delegate/transfer) 权 (power), it incorrectly merges them together.

Similarly, the model segments 居 (in/at) + 中 (middle) as 居中 (in the middle), since the training data contains words such as 居首 (in the first place) and 居次 (in the second place). This example also hints at the ambiguity of word delineation in Chinese, and explains the difficulty in keeping annotations consistent.

As another example, 县 is often attached to another proper noun to become a new word, e.g., 高雄 (Kaohsiung) + 县 becomes 高雄县 (county of Kaohsiung), 新竹 (Hsinchu) + 县 becomes 新竹县 (county of Hsinchu). When seeing 银行县支行 (bank’s county branch), which should be 银行 (bank) + 县支行 (county branch), the model outputs 银行县 + 支行 (i.e. a county named bank).

System	AS	CITYU	CTB6	CTB7	MSR	PKU	UD	Average
This work	98.03	98.22	97.06	97.07	98.48	97.95	97.00	97.69
-LSTM dropout	+0.03	-0.33	-0.31	-0.24	+0.04	-0.29	-0.76	-0.35
-stacked bi-LSTM	-0.13	-0.20	-0.15	-0.14	-0.17	-0.17	-0.39	-0.27
-pretrain	-0.13	-0.23	-0.94	-0.74	-0.45	-0.27	-2.73	-0.78

Table 6: Ablation results on development data. Top row: absolute performance of our system. Other rows: difference relative to the top row.

System	Fully tuned	Avg
This work	97.69	97.49
-Stacked	97.41	97.16
-Pretraining	96.90	96.81

Table 7: Hyperparameter ablation experiments. “Fully tuned” indicates per-system tuning for each dataset. “Avg” is the best setting when averaging across datasets.

Fixing the above errors requires semantic level knowledge such as ‘Bank’ (銀行) is unlikely to be the name of a county (县), and likewise, transfer power (下放权) is not a type of right (权).

Previous work (Huang and Zhao, 2007) also pointed out that OOV is a major obstacle to achieving high segmentation accuracy. They also mentioned that machine learning approaches together with character-based features are more promising in solving OOV problem than rule based methods. Our analysis indicate that learning from the training corpus alone can hardly solve the above mentioned errors. Exploring other sources of knowledge is essential for further improvement. One potential way to acquire such knowledge is to use a language model that is trained on a large scale corpus (Peters et al., 2018). We leave this to future investigation.

Unfortunately, a third (34 out of 104) of the errors we have looked at were due to annotation inconsistency. For example, 建筑系 (Department of Architecture) is once annotated as 建筑 (Architecture) + 系 (Department) and once as 建筑系 under exactly the same context 建筑系教授喻肇青 (Zhaoqing Yu, professor of Architecture). 高新技术 (advanced technology) is annotated as 高 (advanced) + 新 (new) + 技术 (technology) for 37 times, and is annotated as 高新 (advanced and new) + 技术 (technology) for 19 times.

In order to augment the manual verification we performed above, we also wrote a script to automatically find inconsistent annotations in the data. Since this is an automatic script, it cannot distin-

corpus	tokens	inconsistency %
AS	4,903,564	1.31
CITYU	1,309,208	0.62
CTB6	641,368	1.27
CTB7	950,138	1.64
MSR	2,132,480	0.28
PKU	994,822	0.53
UD	98,608	0.46

Table 8: Automatically computed inconsistency in the corpus training data. See text for methodology.

guish between genuine ambiguity and inconsistent annotations. The heuristic we use is the following: for all word bigrams in the training data, we see if they also occur as single words or word trigrams. We ignore the dominant analysis and count the number of occurrences of the less frequent analyses and report this number as a fraction of the number of tokens in the corpus. Table 8 shows the results of running the script. We see that the AS corpus is the least consistent (according to this heuristic) while MSR is the most consistent. This might explain why both our system and prior work have relatively low performance on AS even though this has the largest training set. By contrast results are much stronger on MSR, and this might be in part because it is more consistently annotated. The ordering of corpora by inconsistency roughly mirrors their ordering by accuracy.

4 Conclusion

In this work, we showed that further research in Chinese segmentation must overcome two key challenges: (1) rigorous tuning and testing of deep learning architectures and (2) more effort should be made on exploring resources for further performance gain.

References

- Deng Cai, Hai Zhao, Zhisong Zhang, Yuan Xin, Yongjian Wu, and Feiyue Huang. 2017. Fast and accurate neural word segmentation for chinese. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 608–615. Association for Computational Linguistics.
- Xinchi Chen, Xipeng Qiu, Chenxi Zhu, Pengfei Liu, and Xuanjing Huang. 2015. Long short-term memory neural networks for chinese word segmentation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1197–1206, Lisbon, Portugal. Association for Computational Linguistics.
- Xinchi Chen, Zhan Shi, Xipeng Qiu, and Xuanjing Huang. 2017. Adversarial multi-criteria learning for chinese word segmentation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1193–1203, Vancouver, Canada. Association for Computational Linguistics.
- 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.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1019–1027. Curran Associates, Inc.
- Chang-ning Huang and Hai Zhao. 2007. Chinese word segmentation: A decade review. 21(3):8.
- Shuhei Kurita, Daisuke Kawahara, and Sadao Kurohashi. 2017. Neural joint model for transition-based chinese syntactic analysis. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1204–1214. Association for Computational Linguistics.
- Wang Ling, Chris Dyer, Alan W Black, and Isabel Trancoso. 2015. Two/too simple adaptations of word2vec for syntax problems. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1299–1304, Denver, Colorado. Association for Computational Linguistics.
- Yijia Liu, Wanxiang Che, Jiang Guo, Bing Qin, and Ting Liu. 2016. Exploring segment representations for neural segmentation models. *CoRR*, abs/1604.05499.
- Jianqiang Ma and Erhard Hinrichs. 2015. Accurate linear-time chinese word segmentation via embedding matching. 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 1733–1743, Beijing, China. Association for Computational Linguistics.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. On the state of the art of evaluation in neural language models. In *International Conference on Learning Representations*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Wenzhe Pei, Tao Ge, and Baobao Chang. 2014. Max-margin tensor neural network for chinese word segmentation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 293–303. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *CoRR*, abs/1802.05365.
- Xian Qian and Yang Liu. 2017. A non-dnn feature engineering approach to dependency parsing – fbaml at conll 2017 shared task. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 143–151. Association for Computational Linguistics.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.
- Chunqi Wang and Bo Xu. 2017. Convolutional neural network with word embeddings for chinese word segmentation. *CoRR*, abs/1711.04411.
- Yiyou Wang, Jun’ichi Kazama, Yoshimasa Tsuruoka, Wenliang Chen, Yujie Zhang, and Kentaro Torisawa. 2011. Improving chinese word segmentation and pos tagging with semi-supervised methods using large auto-analyzed data. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 309–317, Chiang Mai, Thailand. Asian Federation of Natural Language Processing.
- David Weiss, Chris Alberti, Michael Collins, and Slav Petrov. 2015. Structured training for neural network transition-based parsing. 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)*, volume 1, pages 323–333.

- Jie Yang, Yue Zhang, and Fei Dong. 2017. Neural word segmentation with rich pretraining. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 839–849. Association for Computational Linguistics.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329*.
- Daniel Zeman, Martin Popel, Milan Straka, Jan Hajic, Joakim Nivre, Filip Ginter, Juhani Luotolahti, Sampo Pyysalo, Slav Petrov, Martin Potthast, Francis Tyers, Elena Badmaeva, Memduh Gokirmak, Anna Nedoluzhko, Silvie Cinkova, Jan Hajic jr., Jaroslava Hlavacova, Václava Kettnerová, Zdenka Uresova, Jenna Kanerva, Stina Ojala, Anna Missilä, Christopher D. Manning, Sebastian Schuster, Siva Reddy, Dima Taji, Nizar Habash, Herman Leung, Marie-Catherine de Marneffe, Manuela Sanguinetti, Maria Simi, Hiroshi Kanayama, Valeria de Paiva, Kira Droганova, Héctor Martínez Alonso, Çağrı Çöltekin, Umut Sulubacak, Hans Uszkoreit, Vivien Macketanz, Aljoscha Burchardt, Kim Harris, Katrin Marheinecke, Georg Rehm, Tolga Kayadelen, Mohammed Attia, Ali Elkahky, Zhuoran Yu, Emily Pitler, Saran Lertpradit, Michael Mandl, Jesse Kirchner, Hector Fernandez Alcalde, Jana Strnadová, Esha Banerjee, Ruli Manurung, Antonio Stella, Atsuko Shimada, Sookyoung Kwak, Gustavo Mendonca, Tatiana Lando, Rattima Nitisaroj, and Josie Li. 2017. Conll 2017 shared task: Multilingual parsing from raw text to universal dependencies. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–19, Vancouver, Canada. Association for Computational Linguistics.
- Meishan Zhang, Yue Zhang, and Guohong Fu. 2016a. Transition-based neural word segmentation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Meishan Zhang, Yue Zhang, and Guohong Fu. 2016b. Transition-based neural word segmentation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 421–431.
- Xiaoqing Zheng, Hanyang Chen, and Tianyu Xu. 2013. Deep learning for chinese word segmentation and pos tagging. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 647–657. Association for Computational Linguistics.
- Hao Zhou, Zhenting Yu, Yue Zhang, Shujian Huang, XIN-YU DAI, and Jiajun Chen. 2017. Word-context character embeddings for chinese word segmentation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*,