

Fake News Detection as Natural Language Inference

Kai-Chou Yang
National Cheng Kung University
Tainan, Taiwan
zake7749@gmail.com

Timothy Niven
National Cheng Kung University
Tainan, Taiwan
tim.niven.public@gmail.com

Hung-Yu Kao
National Cheng Kung University
Tainan, Taiwan
hykao@mail.ncku.edu.tw

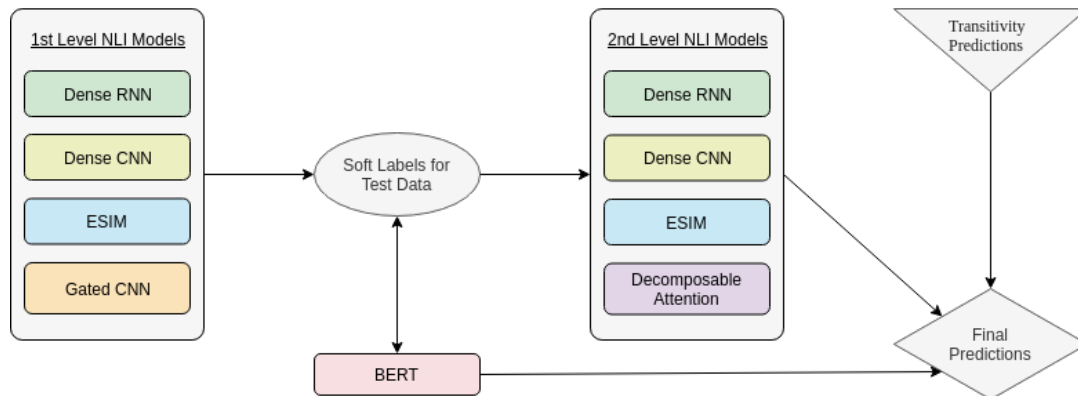


Figure 1: Overview of our method. High performing NLI models are independently trained and ensembled with a fine-tuned BERT model to determine soft labels, which are then used to fine-tune the original NLI models, BERT, and the Decomposable Attention model. These are then ensembled and combined with predictions made via observing transitivity relations.

ABSTRACT

This report describes the entry by the Intelligent Knowledge Management (IKM) Lab in the WSDM 2019 Fake News Classification challenge. We treat the task as natural language inference (NLI). We individually train a number of the strongest NLI models as well as BERT. We ensemble these results and retrain with noisy labels in two stages. We analyze transitivity relations in the train and test sets and determine a set of test cases that can be reliably classified on this basis. The remainder of test cases are classified by our ensemble. Our entry achieves test set accuracy of 88.063% for 3rd place in the competition.

KEYWORDS

fake news, natural language inference, natural language processing

ACM Reference Format:

Kai-Chou Yang, Timothy Niven, and Hung-Yu Kao. 2019. Fake News Detection as Natural Language Inference. In *WSDM '19, February 11–15, 2019, Melbourne, Australia*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/1122445.1122456>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '19, February 11–15, 2019, Melbourne, Australia

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9999-9/18/06...\$15.00
<https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

The WSDM 2019 Fake News Classification challenge presents pairs of sentences requiring three-class prediction. The first sentence is the title of an article already known to be fake news. The second sentence is the title of another article, and the task is to decide whether it agrees with the original fake news, disagrees with it, or is unrelated. Sentences are in Mandarin and are drawn from news sources in China. English translations are provided, however they are noisy machine translations and we did not obtain good results using them. For this reason, we ignored the English sentences and just used the Mandarin.¹

This task can be viewed as a natural language inference task [2], where the first sentence corresponds to the “premise” (P) and the second to the “hypothesis” (H). An example of a contradiction (disagreed) case taken from the training data:

- P Over 1,000 foreigners dissect children and steal their organs
H The rumour of 1,000 foreigners harvesting children’s organs strikes again

Based on this observation we focused on the best performing neural models for SNLI² as well as BERT [5]. We used these models in an ensemble and fine-tuned with soft labels in a process summarized in figure 1, that will be outlined in the following section.

We additionally explored transitive relations in the train and test datasets and found that classification on this basis can be used to reliably increase our accuracy as compared to using our final classifier’s predictions.

¹Examples given in this report are in English and come from our own translations.

²<https://nlp.stanford.edu/projects/snli/>

2 METHOD

Our overall method is summarized in figure 1 and is explained in detail in this section, which aims to provide step-by-step instructions to accompany our published code³ for reproducing of our results.

In general we optimize with Adam [7], use early stopping conditioned on validation set accuracy, and dropout [12] for regularization. As we have a large number of models and hyperparameter settings, we refer the reader to our published code for further details.

2.1 Embeddings

We prepare multiple kinds of embedding at both the word and character level. We train models on all types independently and ensemble the results. This approach is motivated by the fact that Chinese word segmentation is a hard problem often resulting in noise and many out-of-vocabulary tokens. Ensembling the different information captured by these embedding techniques is also expected to reduce bias and variance.

At the word level, we use Tencent [11] word embedding and the SGNS version of Chinese word embeddings [8]. These word embeddings are trained in an unsupervised manner on a large corpus drawn from different sources and have acceptable word coverage for this task. At the character level, we train a skip-gram model [9], CBOW model, and FastText model [1] with window sizes 5, 7, and 3, respectively. The training data for character embedding training is composed of the set of sentences in train and test datasets. These three embeddings are then concatenated to form the final character-level embedding.

2.2 First Level NLI Models

We independently train a number of relatively high performing NLI models:

- Dense RNN
- Dense CNN
- ESIM [3]
- Gated CNN [4]
- Decomposable Attention [10]

The dense NLI architecture is an abstract version of [6], which densely concatenates the features from different levels and repeats the comparison loop several times. The dense CNN shares the same general architecture as the dense RNN (figure 2) but uses a CNN encoder instead. For ESIM, Gated CNN and decomposable attention, we slightly modified the architectures to make them fit well on this task. Due to space requirements, we refer the reader to corresponding researches for details about their architectures.

2.3 First Level Ensemble

To ensemble the first level weak models, we use LightGBM⁴ and a densely connected feed-forward network. For each of our 10 models, we concatenate the 3 class label probabilities, yielding a 30-dimensional input vector for each training and testing data point. For preventing data leakage, we choose the same validation set for

³<https://github.com/zake7749/WSDM-Cup-2019>

⁴<https://github.com/Microsoft/LightGBM>

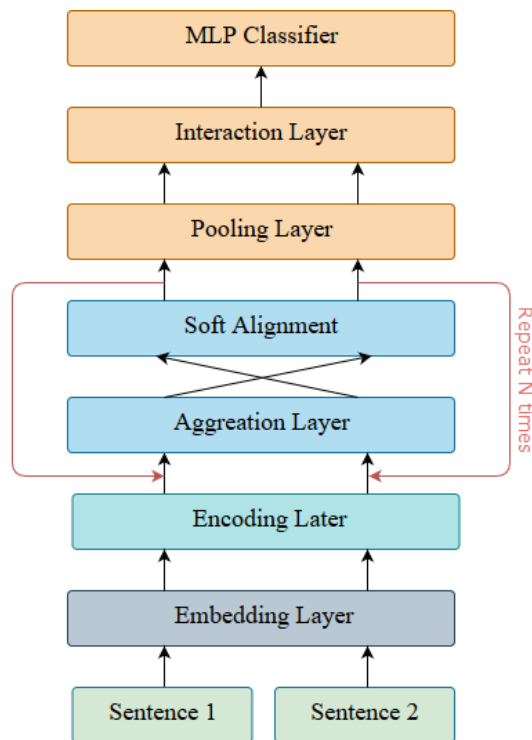


Figure 2: The general architecture of the Dense RNN and Dense CNN models.

early-stopping. The ensemble result achieves a test set accuracy of 86.741%.

2.4 BERT

We train the BERT base Chinese model⁵ for three epochs on the full training set. The learning rate was $5e-5$, maximum sequence length was 128, and batch size was 32. This achieved a test set accuracy of 86.689%. We did not perform an exhaustive search of hyperparameter space due to time constraints.

2.5 Blending

For the final output in the first level, we blend the predictions of (2.3) and (2.4). The blended predictions are the weighted sum of the BERT and ensemble predictions, with weights 0.42 and 0.58, respectively. We chose these weights by threshold search to make sure the blended result is an optimal mixture of both predictions. The test set accuracy is 86.963%.

2.6 Second Level: Fine Tune NLI Models

We took the predictions generated by step (2.5) as soft pseudo-labels and used them to fine-tune the pretrained NLI models in step (2.1). The training set was concatenated with the pseudo-labeled test set and the same validation set was used for early stopping. The fine-tuned results are reported in table 1. Note that the Decomposable attention model was added at this stage, initialized with random

⁵<https://github.com/huggingface/pytorch-pretrained-BERT>

Table 1: Test set accuracy of for NLI models

Model	First Level			Fine-tuned		
	Tencent	SGNS	Character	Tencent	SGNS	Character
Decomposable Attention				0.86730	0.86721	
Dense RNN	0.85529	0.85620	0.85170	0.87704	0.87248	0.86809
Dense CNN	0.85082	0.84803	0.85287	0.87479	0.87114	0.86302
ESIM	0.85738	0.85622	0.86053	0.87788	0.87334	0.87436
Gated CNN	0.84711					

Table 2: Test set accuracy of for ensembles, BERT, blended models, and transitivity rules

Model	Accuracy
Ensemble (1)	0.86741
BERT (1)	0.86689
Blended (1)	0.86963
Ensemble (2)	0.87990
BERT (2)	0.87484
Blended (2)	0.88019
Blended (2) + Transitivity	0.88063

weights. Due to time constraints we did not train every combination of models and embeddings.

2.7 Second Level Ensemble

Once again, we used LightGBM and a multi-layer perceptron to perform ensembling with the output of the second level NLI models in the same manner as described in section (2.3). This achieved a test set accuracy of 87.990%.

2.8 BERT Pseudo-Label Training

We used the psueod-labels to continue fine tuning the BERT model from step (2.4). Whereas we tuned the NLI models on the pseudo-labeled test set only, BERT was trained on the whole concatenation of the entire training set and the pseudo-labeled test set, without any validation. We trained BERT for three epochs with the same hyperparameter settings as in step (2.4). This model achieved a test set score of 87.484%.

2.9 Final Blending

After fine-tuning BERT and the based NLI models, we once again blended their predictions to obtain our final predictions. The blending weights were 0.79 and 0.21, respectively. Just as in step (2.5) these weights were determined by threshold search. The blended result was 88.019%.

2.10 Post-Processing: Transitive Relations

We investigated transitivity relations in the data and found they were reliable enough to use as test set predictions. Figure 3 demonstrates the two types of relation we considered. For positive relations, we observed that if sentence *A* agrees with sentence *B*, and *B* agrees with sentence *C*, then *A* should also agree with *C*. In the

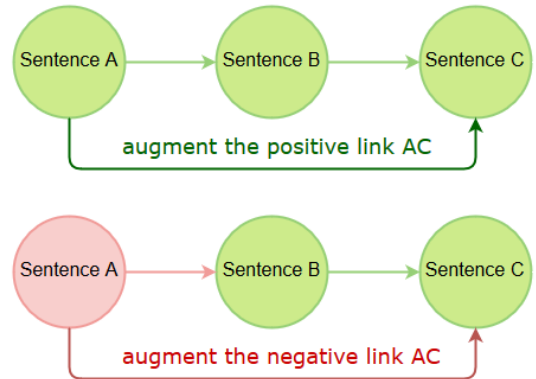


Figure 3: Positive and negative transitivity relations in the labeled data.

negative case, if *A* disagrees with *B*, and *B* agrees with *C*, the *A* should also disagree with *C*.

We found that the positive case held for 99.9% of the training data. The negative case held 99.7% of the time. As there are sentence overlaps between the train and test sets, we are able to apply these rules recursively to generate predictions for 6,888 data points in the test set. We expected these should be more reliable than our trained classifiers. We made a late submission with these predictions, labeling all other test set samples with a “fake” prediction label so only the samples labeled by transitivity rules were considered by the scorer. This submission achieved 93% accuracy, a level of performance dramatically higher than any trained classifier in the competition, validating our expectation. Overall, we observed a 0.04% increase in our test set accuracy using this method as compared to using our best classifier for all test set data points. The final accuracy was 88.063%.

2.11 Results

The results from all levels are summarized in table 2. We see a consistent improvement from successive ensembling and pseudo-label fine tuning. We note again that, due to time constraints we were unable to fully exploit all models we identified, nor conduct a suitably thorough hyperparameter search. We therefore expect our method could easily improve upon the results presented here.

3 CONCLUSION

NLI models appear to be generally effective for this task. Ensemble via gradient boosting and fine-tuning with noisy labels proved to be very beneficial. However, due to time constraints we were unable to cover all the combinations we targeted and expect we could improve further. As the train and test sets have overlapping sentences, we were able to exploit transitive relations between them to reliably improve our performance.

For future work we intend to further investigate the transitivity method for data augmentation. Initial investigation revealed we can create an additional 700,000 sentence “agreed” pairs and 19,000 “disagreed”. That is 7 and 2 times the number of original training data points, respectively.

REFERENCES

- [1] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* 5 (2017), 135–146.
- [2] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. *CoRR* abs/1508.05326 (2015). arXiv:1508.05326 <http://arxiv.org/abs/1508.05326>
- [3] Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, and Hui Jiang. 2016. Enhancing and Combining Sequential and Tree LSTM for Natural Language Inference. *CoRR* abs/1609.06038 (2016). arXiv:1609.06038 <http://arxiv.org/abs/1609.06038>
- [4] Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. 2016. Language Modeling with Gated Convolutional Networks. *CoRR* abs/1612.08083 (2016). arXiv:1612.08083 <http://arxiv.org/abs/1612.08083>
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018). arXiv:1810.04805 <http://arxiv.org/abs/1810.04805>
- [6] Seonhoon Kim, Jin-Hyuk Hong, Inho Kang, and Nojun Kwak. 2018. Semantic Sentence Matching with Densely-connected Recurrent and Co-attentive Information. *CoRR* abs/1805.11360 (2018). arXiv:1805.11360 <http://arxiv.org/abs/1805.11360>
- [7] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *CoRR* abs/1412.6980 (2014). arXiv:1412.6980 <http://arxiv.org/abs/1412.6980>
- [8] Shen Li, Zhe Zhao, Renfen Hu, Wensi Li, Tao Liu, and Xiaoyong Du. 2018. Analogical Reasoning on Chinese Morphological and Semantic Relations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, 138–143. <http://aclweb.org/anthology/P18-2023>
- [9] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR* abs/1301.3781 (2013). arXiv:1301.3781 <http://arxiv.org/abs/1301.3781>
- [10] Ankur P. Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A Decomposable Attention Model for Natural Language Inference. *CoRR* abs/1606.01933 (2016). arXiv:1606.01933 <http://arxiv.org/abs/1606.01933>
- [11] Yan Song, Shuming Shi, Jing Li, and Haisong Zhang. 2018. Directional Skip-Gram: Explicitly Distinguishing Left and Right Context for Word Embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. Association for Computational Linguistics, 175–180. <https://doi.org/10.18653/v1/N18-2028>
- [12] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research* 15 (2014), 1929–1958. <http://jmlr.org/papers/v15/srivastava14a.html>