

Performance Impact Caused by Hidden Bias of Training Data for Recognizing Textual Entailment

Masatoshi Tsuchiya

Toyohashi University of Technology
1-1 Hibarigaoka, Tempaku-cho, Toyohashi, Aichi, Japan
tsuchiya@imc.tut.ac.jp

Abstract

The quality of training data is one of the crucial problems when a learning-centered approach is employed. This paper proposes a new method to investigate the quality of a large corpus designed for the recognizing textual entailment (RTE) task. The proposed method, which is inspired by a statistical hypothesis test, consists of two phases: the first phase is to introduce the predictability of textual entailment labels as a null hypothesis which is extremely unacceptable if a target corpus has no hidden bias, and the second phase is to test the null hypothesis using a Naive Bayes model. The experimental result of the Stanford Natural Language Inference (SNLI) corpus does not reject the null hypothesis. Therefore, it indicates that the SNLI corpus has a hidden bias which allows prediction of textual entailment labels from hypothesis sentences even if no context information is given by a premise sentence. This paper also presents the performance impact of NN models for RTE caused by this hidden bias.

Keywords: Recognizing Textual Entailment; Evaluation Methodology; Naive Bayes

1. Introduction

The quality of a training data is one of the crucial problems when a learning-centered approach including neural network (NN) is employed. (Reidsma and Carletta, 2008) demonstrated that annotation errors of the dialog act corpus, which follows a certain systematic pattern, mislead the learning result of Bayesian network. (Zhang et al., 2017) described that the capacity of a NN model is large enough for brute-force memorizing the entire data set, even if its labels are random. Thus, an influence for a NN model caused by a certain systematic pattern may become more serious than the influence for other learning-centered methods.

Both a method to improve the quality of a training data and a metric to evaluate its reliability are important. As the former method, majority vote of human annotators was widely used (Sabou et al., 2014; Bowman et al., 2015; Al Khatib et al., 2016) As the latter metric, many kind of inter-annotator agreement metrics based on multiple human annotation results were employed, because direct assessment of data quality is difficult (Craggs and Wood, 2005; Artstein and Poesio, 2008; Mathet et al., 2012).

This paper proposes a new empirical method to investigate a quality of a large corpus designed for the recognizing textual entailment (RTE) task (Condoravdi et al., 2003; Bos and Markert, 2005; MacCartney and Manning, 2009; Marelli et al., 2014a). The proposed method, which is inspired by a statistical hypothesis test, assesses a quality of a target corpus directly and does not depend on multiple human annotation results unlike the existing metrics. The proposed method consists of two phases: the first phase is to introduce the predictability of textual entailment (TE) labels as a null hypothesis which is extremely unacceptable when a target corpus has no hidden bias, and the second phase is to test the null hypothesis for the target corpus using a Naive Bayes (NB) model.

In order to demonstrate the efficiency of the proposed method, we investigate two RTE corpora: the Stanford Natural Language Inference (SNLI)

s_1	Two boys are swimming in the pool.
s_2	Two girls are playing basketball.
s_3	Two women are swimming in the pool.
s_h	Two children are swimming in the pool.

Figure 1: Example sentences of RTE. The textual entailment label of s_h is determinable if and only if context information is given by a premise sentence.

corpus (Bowman et al., 2015) and the Sentences Involving Compositional Knowledge (SICK) corpus (Marelli et al., 2014b). Although the experimental result of the SICK corpus rejects the null hypothesis, the result of the SNLI corpus does not reject it. Thus, the SNLI corpus has a hidden bias to allow prediction of TE labels from hypothesis sentences even if no context information is given by a premise sentence. The other experiment shows that this hidden bias causes the risk that a NN model for RTE works as a entirely different model than its constructor expects.

The major contributions of this paper are following three points:

- This paper proposes a new empirical method to reveal a hidden bias of a large RTE corpus (Section 2.).
- This paper applies the proposed method on the SNLI corpus and the SICK corpus, and reveals the hidden bias of the SNLI corpus (Section 3.).
- This paper also presents that this hidden bias causes the risk that a NN model proposed for RTE works as a entirely different model than its constructor expects (Section 4.).

2. Proposed Method

This section describes the detail of the proposed method, which consists of two phases.

Predicted labels	Corpus labels		
	Entailment	Neutral	Contradiction
Entailment	2275	644	706
Neutral	508	1976	563
Contradiction	585	599	1968

(a) The confusion matrix obtained by the TE label prediction model trained and tested on the SNLI corpus. It tries to predict an appropriate TE label for each individual hypothesis sentence.

Predicted labels	Corpus labels		
	Entailment	Neutral	Contradiction
Entailment	3	3	2
Neutral	1411	2790	718
Contradiction	0	0	0

(b) The confusion matrix obtained by the TE label prediction model trained and tested on the SICK corpus. It simply tries the major TE label ‘neutral’ for almost all hypothesis sentences without prediction.

Figure 2: Confusion matrices of TE Label Prediction Models

2.1. Predictability of TE Labels without Premise Sentences

The first phase of the proposed method is to derive a null hypothesis which is extremely unacceptable when a target corpus has no hidden bias. We focus the task definition of the target corpus for this phase.

(Marelli et al., 2014a) defined RTE as a task to partition relationships between a premise sentence and a hypothesis sentence into three categories: entailment, neutral and contradiction. Consider the example sentences shown in Figure 1. When s_1 is given as a premise sentence and s_h is given as a hypothesis sentence, the relationship between s_1 and s_h is labeled entailment. The relationship between s_2 and s_h is labeled neutral, and the relationship between s_3 and s_h is labeled contradiction. These examples indicate that the TE label is determinable if and only if context information is given by a premise sentence. Based on this observation, the null hypothesis of the proposed method is defined as follows:

Definition of the null hypothesis

The TE label of the hypothesis sentence is determinable without the premise sentence.

Because the null hypothesis looks extremely unacceptable, it is appropriate to reveal a hidden bias of the target RTE corpus.

2.2. TE Label Prediction Model

The second phase of the proposed method is to test statistical significance of the null hypothesis for a target corpus. This phase requires two models: the first model is the statistical model of the null hypothesis, henceforth referred to as the *TE label prediction model*, and the second model is the statistical model of the alternative hypothesis, henceforth referred to as the *baseline model*.

The TE label prediction model is a model which predicts TE labels for hypothesis sentences without context information of premise sentences. This paper employs a multinomial NB model (Wang and Manning, 2012) as the TE label prediction model, which is defined by the following equation:

$$\hat{y} = \operatorname{argmax}_y P(y) \prod_{i=1}^n P(x_i|y), \quad (1)$$

where y is a TE label and x_i is a feature. This paper simply employs all word unigrams of hypothesis sentences as features.

Corpus	TE label prediction model	Baseline model
SNLI	63.3%	34.3%
SICK	56.7%	56.7%

Table 1: Performance of the TE label prediction model trained and tested on two RTE corpora. The TE label prediction model trained and tested on the SNLI corpus achieves statistically significant accuracy than the baseline model.

The baseline model assigns TE labels for hypothesis sentences when no information is given by either premise sentences or hypothesis sentences but only the TE label distribution $P(y)$ of the target corpus is available. In such case, it is reasonable to assign the TE label which occurs most frequently in the target corpus for all hypothesis sentences. This baseline assignment is defined as follows:

$$\check{y} = \operatorname{argmax}_y P(y) \quad (2)$$

If there is a statistically significant difference between the TE label prediction model and the baseline model, the null hypothesis is not rejected for the target corpus, and it indicates that the target corpus contains a hidden bias. Otherwise, the null hypothesis is rejected for the target corpus,

3. Experiment

This section presents the detailed experimental conditions and the experimental results. The highlight of these results is that the TE label prediction model achieves 63.3% accuracy for the SNLI corpus.

Table 1 shows the performances of the TE label prediction models trained and tested on two RTE corpora¹. The TE label prediction model, which is trained on the SNLI training hypothesis sentences and their TE labels, achieves 63.3% accuracy on the SNLI test hypothesis sentences without premise sentences. The baseline model based on the SNLI TE label distribution achieves 34.3% accuracy on the same hypothesis sentences. The sign test indicates that there is a statistically significant difference between these models ($p = 5.7e^{-202}$). On the other hand, the performance of the TE label prediction model trained and tested on the SICK corpus is close to the performance of the baseline model (56.7%). The sign test indicates that there is no statistically significant difference between these models ($p = 0.65$).

¹(Pedregosa et al., 2011) is employed to implement the TE label prediction model.

	E_e	H_e
Entailment	2,275 (36.6%)	1,093 (30.3%)
Neutral	1,976 (31.8%)	1,243 (34.5%)
Contradiction	1,968 (31.6%)	1,269 (35.2%)
Total	6,219 (63.3%)	3,605 (36.7%)

Table 2: Empirical classification of the SNLI corpus using the TE label prediction model

Figure 2 clearly shows the difference between the behavior of the TE label prediction model trained on the SNLI corpus and the model trained on the SICK corpus. The left matrix is obtained by the model trained and tested on the SNLI corpus, and the right matrix is obtained by the model trained and tested on the SICK corpus. The model of the SICK corpus simply tries the major TE label ‘neutral’ for almost all hypothesis sentences without prediction, although the model of the SNLI corpus tries to predict an appropriate TE label for each individual hypothesis sentence.

These results indicate that the null hypothesis is rejected for the SICK corpus, but it is not rejected for the SNLI corpus. Therefore, hypothesis sentences of the SNLI corpus have a hidden bias to allow prediction of their TE labels without premise sentences.

4. Discussion

As described in Section 3., hypothesis sentences of the SNLI corpus have a hidden bias to allow prediction of their TE labels without premise sentences. The magnitude of the performance impact caused by the hidden bias is important, because the SNLI corpus is widely used as training data by many NN models for RTE (Bowman et al., 2015; Rocktäschel et al., 2015; Yin et al., 2016; Mou et al., 2016a; Wang and Jiang, 2016; Liu et al., 2016a; Liu et al., 2016b; Cheng et al., 2016; Parikh et al., 2016; Sha et al., 2016). This section discusses the performance impact of the NN models caused by the hidden bias.

4.1. Empirical Classification of SNLI Corpus

The test pairs of the SNLI corpus are classified into two subsets using the TE label prediction model trained on the SNLI corpus. The first subset is the empirical easy test set E_e , which consists of all test pairs whose TE labels are predicted correctly by the TE label prediction model. The second subset is the empirical hard test set H_e , which consists of the rest pairs. Table 2 shows the classification result. 63.3% test pairs of the SNLI corpus were classified as E_e , and the rest pairs were classified as H_e .

4.2. Definitions of NN Models for RTE

Two NN models are prepared to evaluate performance impacts caused by the hidden bias. The first model (henceforth denoted as the parallel LSTM model) was proposed by (Bowman et al., 2015) for RTE, and was evaluated on the performance difference between RTE corpora by (Mou et al., 2016b). This model is defined by the following equations.

$$\mathbf{h}_{p,i} = \text{LSTM}_p(W_e x_{p,i} + W_{hp} \mathbf{h}_{p,i-1})$$

$$\begin{aligned} \mathbf{h}_{h,i} &= \text{LSTM}_h(W_e x_{h,i} + W_{hh} \mathbf{h}_{h,i-1}) \\ \mathbf{l}_1 &= \tanh(W_1[\mathbf{h}_{p,|\mathbf{x}_p|}, \mathbf{h}_{h,|\mathbf{x}_h|}] + B_1) \\ \mathbf{l}_2 &= \tanh(W_2 \mathbf{l}_1 + B_2) \\ \mathbf{l}_3 &= \tanh(W_3 \mathbf{l}_3 + B_3) \\ \mathbf{y} &= \text{softmax}(\mathbf{l}_3) \end{aligned}$$

The first step is to convert a premise sentence \mathbf{x}_p and a hypothesis sentence \mathbf{x}_h into embedding vectors using the word embedding matrix W_e , which is initialized with the 300-dimension reference GloVe vectors (Pennington et al., 2014). The second step is to convert embedding vectors into two 100-dimension sentence vectors with LSTMs, and they are concatenated into a 200-dimension vector. The remaining steps are to predict a TE label with three tanh fully connected layers and then to apply the softmax function.

The second model (henceforth denoted as the sequential LSTM model), which was proposed by (Rocktäschel et al., 2015) for RTE, is defined as follows.

$$\begin{aligned} \mathbf{h}_{p,i} &= \text{LSTM}_p(W_e x_{p,i} + W_{hp} \mathbf{h}_{p,i-1}) \\ \mathbf{h}_{h,0} &= \text{LSTM}_h(W_{hh} \mathbf{h}_{p,|\mathbf{x}_p|}) \\ \mathbf{h}_{h,i} &= \text{LSTM}_h(W_e x_{h,i} + W_{hh} \mathbf{h}_{h,i-1}) \\ \mathbf{l} &= \tanh(W_l \mathbf{h}_{h,|\mathbf{x}_h|} + B_l) \\ \mathbf{y} &= \text{softmax}(\mathbf{l}) \end{aligned}$$

In the second model, two LSTMs are sequentially connected. Thus, it is possible to consider that the memory cells of these LSTMs are directly modeling a recognition process unlike the parallel LSTM model. All vectors of the sequential LSTM model are 100-dimension. Although (Rocktäschel et al., 2015) proposed the variants with attentions between a premise sentence and a hypothesis sentence, the attention-less model is employed in this experiment, because of its simplicity.

4.3. Performance Impact of NN Models for RTE

This subsection presents the big performance drop of the NN models caused by the hidden bias.

Table 3 shows the experimental results of these NN models trained and tested on the SNLI corpus. Although both NN models achieve high accuracy for the whole test set and for the empirical easy test set E_e , they achieve drastic low accuracy for the empirical hard test set H_e . These performance drops mean that a large portion of the high accuracy achieved by both NN models benefits from the empirical easy test set E_e .

Table 4 shows the performance which is achieved by the same NN models when all words of premise sentences are replaced by unknown word symbols. Because this replacement removes all context information from premise sentences, thus the performance of the NN models must drop close to the chance ratio, if the NN models decide TE labels based on context information of premise sentences. Despite this expectation, both NN models achieve obviously higher performance than the chance ratio for E_e (36.8%, shown in Table 2). This result indicates that both NN models do not

Model	Related models	$E_e \cup H_e$	E_e	H_e
Parallel LSTM model	76.3% (Mou et al., 2016b)	76.8%	87.8%	57.8%
Sequential LSTM model	80.9% (Rocktäschel et al., 2015)	81.4%	90.1%	65.6%

Table 3: Performance of the NN models for RTE. Performance drops from the empirical easy test set E_e to the empirical hard test set H_e are observed for both NN models.

Model	$E_e \cup H_e$	E_e	H_e
Parallel LSTM model	54.1%	66.0%	33.7%
Sequential LSTM model	48.6%	56.7%	34.7%

Table 4: Performance of the NN models for RTE, when all words of premise sentences are replaced by unknown word symbols. Although all context information is removed by this replacement, both NN models achieve obviously higher performance than the chance ratio for the empirical easy test set E_e . This result indicates that both NN models do not work as RTE models for E_e , but work as TE label prediction models for E_e .

	SNLI			SICK		
	Training	Development	Test	Training	Development	Test
Entailment	183,416 (33.4%)	3,329 (33.8%)	3,368 (34.3%)	1,299 (28.9%)	144 (28.8%)	1,414 (28.7%)
Neutral	182,764 (33.3%)	3,235 (32.9%)	3,219 (32.8%)	2,536 (56.4%)	282 (56.4%)	2,793 (56.7%)
Contradiction	183,187 (33.4%)	3,278 (33.3%)	3,237 (33.0%)	665 (14.8%)	74 (14.8%)	720 (14.6%)
Total	549,367	9,842	9,824	4,500	500	4,927

Table 5: Statistics of TE labels of two RTE corpora. Although TE labels of the SNLI corpus are balanced, TE labels of the SICK corpus are not balanced.

	SNLI	SICK
# of training pairs	55k	4,500
# of development pairs	10k	500
# of test pairs	10k	4,927
Premise mean token count	14.1	9.8
Hypothesis mean token count	8.3	9.5
Vocabulary size of training pairs	36,427	2,178
Vocabulary size of test pairs	6,548	2,188
OOV ratio of test pairs	0.25%	0.32%

Table 6: Comparison of the SNLI and SICK corpora. Both corpora are extremely similar from the view point of descriptive statistics.

work as RTE models for E_e , but work as TE label prediction models for E_e . This behavior of NN models for E_e must be entirely different than their constructor expected.

4.4. Comparison of SNLI and SICK corpora

The SNLI and SICK corpora, are entirely similar in their sentence domains, English scene descriptions. Both of them use the Flickr30k corpus (Young et al., 2014) as origins of their sentences. It is also exhibited by the small differences of sentence token mean counts as shown in Table 6. The second is about their vocabulary. The out-of-vocabulary (OOV) ratio of SICK test pairs is 0.15%, when words of SNLI training pairs are regarded as known. This small OOV ratio indicates that SICK test pairs and SNLI training pairs are quite close from the view point of their vocabulary.

The SNLI and SICK corpora are different in the method of composing sentences. Hypothesis sentences of the SNLI corpus are composed by human workers, but all sentences of the SICK corpus are derived from original sentences using hand-crafted rules. We think that this difference may be

a cause of the hidden bias revealed by this paper.

5. Conclusion

This paper proposes a new empirical method to investigate the quality of large RTE corpus. The proposed method consists of two phases: the first phase is to introduce the predictability of TE labels as a null hypothesis, and the second phase is to test the null hypothesis using a NB model. The proposed method reveals a hidden bias of the SNLI corpus, which allows prediction of TE labels from hypothesis sentences without context information given by premise sentences.

This paper also presents that this hidden bias makes a large performance impact on the NN models for RTE. The experimental result shows that a large portion of the high accuracy achieved by the NN models benefits from the hidden bias. The other experimental result shows that a NN model trained on the SNLI corpus does not work as an RTE model, but works as a TE label prediction model, when biased test pairs are given. These results arise a risk that a complex NN model works as an entirely different model than its constructor expects.

6. Acknowledgments

A part of this research was supported by JSPS KAKENHI Grant No. 15K12097. I would like to express my sincere appreciation to Dr. Mitsuo Yoshida and Dr. Adam Meyers for their valuable comments.

7. Bibliographical References

Al Khatib, K., Wachsmuth, H., Kiesel, J., Hagen, M., and Stein, B. (2016). A news editorial corpus for mining argumentation strategies. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3433–3443,

- Osaka, Japan, December. The COLING 2016 Organizing Committee.
- Artstein, R. and Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Computational Linguistics*, 34(4):555–596.
- Bos, J. and Markert, K. (2005). Recognising textual entailment with logical inference. In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 628–635, Vancouver, British Columbia, Canada, October. Association for Computational Linguistics.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015). A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal, September. Association for Computational Linguistics.
- Cheng, J., Dong, L., and Lapata, M. (2016). Long short-term memory-networks for machine reading. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 551–561, Austin, Texas, November. Association for Computational Linguistics.
- Condoravdi, C., Crouch, D., de Paiva, V., Stolle, R., and Bobrow, D. G. (2003). Entailment, intensionality and text understanding. In Graeme Hirst et al., editors, *Proceedings of the HLT-NAACL 2003 Workshop on Text Meaning*, pages 38–45.
- Craggs, R. and Wood, M. M. (2005). Evaluating discourse and dialogue coding schemes. *Computational Linguistics*, 31(3):289–296, September.
- Liu, P., Qiu, X., Chen, J., and Huang, X. (2016a). Deep fusion lstms for text semantic matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1034–1043, Berlin, Germany, August. Association for Computational Linguistics.
- Liu, P., Qiu, X., Zhou, Y., Chen, J., and Huang, X. (2016b). Modelling interaction of sentence pair with coupled-lstms. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1703–1712, Austin, Texas, November. Association for Computational Linguistics.
- MacCartney, B. and Manning, C. D. (2009). An extended model of natural logic. In *Proceedings of the Eight International Conference on Computational Semantics*, pages 140–156, Tilburg, The Netherlands, January. Association for Computational Linguistics.
- Marelli, M., Bentivogli, L., Baroni, M., Bernardi, R., Menini, S., and Zamparelli, R. (2014a). Semeval-2014 task 1: Evaluation of compositional distributional semantic models on full sentences through semantic relatedness and textual entailment. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 1–8, Dublin, Ireland, August. Association for Computational Linguistics and Dublin City University.
- Marelli, M., Menini, S., Baroni, M., Bentivogli, L., Bernardi, R., and Zamparelli, R. (2014b). A SICK cure for the evaluation of compositional distributional semantic models. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 216–223, Reykjavik, Iceland, may. European Language Resources Association (ELRA).
- Mathet, Y., Widłocher, A., Fort, K., François, C., Galibert, O., Grouin, C., Kahn, J., Rosset, S., and Zweigenbaum, P. (2012). Manual corpus annotation: Giving meaning to the evaluation metrics. In *Proceedings of COLING 2012: Posters*, pages 809–818, Mumbai, India, December. The COLING 2012 Organizing Committee.
- Mou, L., Men, R., Li, G., Xu, Y., Zhang, L., Yan, R., and Jin, Z. (2016a). Natural language inference by tree-based convolution and heuristic matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 130–136, Berlin, Germany, August. Association for Computational Linguistics.
- Mou, L., Meng, Z., Yan, R., Li, G., Xu, Y., Zhang, L., and Jin, Z. (2016b). How transferable are neural networks in nlp applications? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 479–489, Austin, Texas, November. Association for Computational Linguistics.
- Parikh, A., Täckström, O., Das, D., and Uszkoreit, J. (2016). A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, Austin, Texas, November. Association for Computational Linguistics.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Reidsma, D. and Carletta, J. (2008). Reliability measurement without limits. *Computational Linguistics*, 34(3):319–326, September.
- Rocktäschel, T., Grefenstette, E., Hermann, K. M., Kočiský, T., and Blunsom, P. (2015). Reasoning about entailment with neural attention. In *Proceedings of The International Conference on Learning Representations (ICLR2016)*.
- Sabou, M., Bontcheva, K., Derczynski, L., and Scharl, A. (2014). Corpus annotation through crowdsourcing: Towards best practice guidelines. pages 859–866.
- Sha, L., Chang, B., Sui, Z., and Li, S. (2016). Reading and thinking: Re-read lstm unit for textual entailment recognition. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2870–2879, Osaka, Japan, De-

- ember. The COLING 2016 Organizing Committee.
- Wang, S. and Jiang, J. (2016). Learning natural language inference with lstm. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1442–1451, San Diego, California, June. Association for Computational Linguistics.
- Wang, S. and Manning, C. (2012). Baselines and bigrams: Simple, good sentiment and topic classification. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 90–94, Jeju Island, Korea, July. Association for Computational Linguistics.
- Yin, W., Sch \tilde{A} $\frac{1}{4}$ tze, H., Xiang, B., and Zhou, B. (2016). Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics*, 4:259–272.
- Young, P., Lai, A., Hodosh, M., and Hockenmaier, J. (2014). From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. In *Proceedings of The International Conference on Learning Representations (ICLR2016)*.