

Exploiting Document Knowledge for Aspect-level Sentiment Classification

Ruidan He^{†‡}, Wee Sun Lee[†], Hwee Tou Ng[†], and Daniel Dahlmeier[‡]

[†]Department of Computer Science, National University of Singapore

[‡]SAP Innovation Center Singapore

[†]{ruidanhe, leews, nght}@comp.nus.edu.sg

[‡]d.dahlmeier@sap.com

Abstract

Attention-based long short-term memory (LSTM) networks have proven to be useful in aspect-level sentiment classification. However, due to the difficulties in annotating aspect-level data, existing public datasets for this task are all relatively small, which largely limits the effectiveness of those neural models. In this paper, we explore two approaches that transfer knowledge from document-level data, which is much less expensive to obtain, to improve the performance of aspect-level sentiment classification. We demonstrate the effectiveness of our approaches on 4 public datasets from SemEval 2014, 2015, and 2016, and we show that attention-based LSTM benefits from document-level knowledge in multiple ways.

1 Introduction

Given a sentence and an opinion target (also called an aspect term) occurring in the sentence, aspect-level sentiment classification aims to determine the sentiment polarity in the sentence towards the opinion target. An opinion target or target for short refers to a word or a phrase describing an aspect of an entity. For example, in the sentence “*This little place has a cute interior decor but the prices are quite expensive*”, the targets are *interior decor* and *prices*, and they are associated with positive and negative sentiment respectively.

A sentence may contain multiple sentiment-target pairs, thus one challenge is to separate different opinion contexts for different targets. For this purpose, state-of-the-art neural methods (Wang et al., 2016; Liu and Zhang, 2017; Chen et al., 2017) adopt attention-based LSTM networks, where the LSTM aims to capture sequential patterns and the attention mechanism aims

to emphasize target-specific contexts for encoding sentence representations. Typically, LSTMs only show their potential when trained on large datasets. However, aspect-level training data requires the annotation of all opinion targets in a sentence, which is costly to obtain in practice. As such, existing public aspect-level datasets are all relatively small. Insufficient training data limits the effectiveness of neural models.

Despite the lack of aspect-level labeled data, enormous document-level labeled data are easily accessible online such as Amazon reviews. These reviews contain substantial linguistic patterns and come with sentiment labels naturally. In this paper, we hypothesize that aspect-level sentiment classification can be improved by employing knowledge gained from document-level sentiment classification. Specifically, we explore two transfer methods to incorporate this sort of knowledge – pretraining and multi-task learning. In our experiments, we find that both methods are helpful and combining them achieves significant improvements over attention-based LSTM models trained only on aspect-level data. We also illustrate by examples that additional knowledge from document-level data is beneficial in multiple ways. Our source code can be obtained from <https://github.com/ruidan/Aspect-level-sentiment>.

2 Related Work

Various neural models (Dong et al., 2014; Nguyen and Shirai, 2015; Vo and Zhang, 2015; Tang et al., 2016a,b; Wang et al., 2016; Zhang et al., 2016; Liu and Zhang, 2017; Chen et al., 2017) have been proposed for aspect-level sentiment classification. The main idea behind these works is to develop neural architectures that are able to learn continuous features and capture the intricate relation between a target and context words. However, to sufficiently train these models, substantial aspect-

level annotated data is required, which is expensive to obtain in practice.

We explore both pretraining and multi-task learning for transferring knowledge from document level to aspect level. Both methods are widely studied in the literature. Pretraining is a common technique used in computer vision where low-level neural layers can be usefully transferred to different tasks (Krizhevsky and Sutskever, 2012; Zeiler and Fergus, 2014). In natural language processing (NLP), some efforts have been initiated on pretraining LSTMs (Dai and Le, 2015; Zoph et al., 2016; Ramachandran et al., 2017) for sequence-to-sequence models in both supervised and unsupervised settings, where promising results have been obtained. On the other hand, multi-task learning simultaneously trains on samples in multiple tasks with a combined objective (Collobert and Weston, 2008; Luong et al., 2015a; Liu et al., 2016), which has improved model generalization ability in certain cases. In the work of Mou et al. (2016), the authors investigated the transferability of neural models in NLP applications with extensive experiments, showing that transferability largely depends on the semantic relatedness of the source and target tasks. For our problem, we hypothesize that aspect-level sentiment classification can be improved by employing knowledge gained from document-level sentiment classification, as these two tasks are highly related semantically.

3 Models

3.1 Attention-based LSTM

We first describe a conventional implementation of an attention-based LSTM model for this task. We use it as a baseline model and extend it with pretraining and multi-task learning approaches for incorporating document-level knowledge.

The inputs are a sentence $s = (w_1, w_2, \dots, w_n)$ consisting of n words, and an opinion target $x = (x_1, x_2, \dots, x_m)$ occurring in the sentence consisting of a subsequence of m words from s . Each word is associated with a continuous word embedding \mathbf{e}_w (Mikolov et al., 2013) from an embedding matrix $\mathbf{E} \in \mathbb{R}^{V \times d}$, where V is the vocabulary size and d is the embedding dimension.

LSTM is used to capture sequential information, and outputs a sequence of hidden vectors:

$$[\mathbf{h}_1, \dots, \mathbf{h}_n] = \text{LSTM}([\mathbf{e}_{w_1}, \dots, \mathbf{e}_{w_n}], \theta_{lstm}) \quad (1)$$

An attention layer assigns a weight α_i to each word in the sentence. The final target-specific representation of the sentence s is then given by:

$$\mathbf{z} = \sum_{i=1}^n \alpha_i \mathbf{h}_i \quad (2)$$

And α_i is computed as follows:

$$\alpha_i = \frac{\exp(\beta_i)}{\sum_{j=1}^n \exp(\beta_j)} \quad (3)$$

$$\beta_i = f_{score}(\mathbf{h}_i, \mathbf{t}) = \tanh(\mathbf{h}_i^T \mathbf{W}_a \mathbf{t}) \quad (4)$$

$$\mathbf{t} = \frac{1}{m} \sum_{i=1}^m \mathbf{e}_{x_i} \quad (5)$$

where \mathbf{t} is the target representation computed as the averaged word embedding of the target. f_{score} is a content-based function that captures the semantic association between a word and the target, for which we adopt the formulation used in (Luong et al., 2015b; He et al., 2017) with parameter matrix $\mathbf{W}_a \in \mathbb{R}^{d \times d}$.

The sentence representation \mathbf{z} is fed into an output layer to predict the probability distribution \mathbf{p} over sentiment labels on the target:

$$\mathbf{p} = \text{softmax}(\mathbf{W}_o \mathbf{z} + \mathbf{b}_o) \quad (6)$$

We refer to this baseline model as LSTM+ATT. It is trained via cross entropy minimization:

$$J = - \sum_{i \in D} \log \mathbf{p}_i(c_i) \quad (7)$$

where D denotes the overall training corpus, c_i denotes the true label for sample i , and $\mathbf{p}_i(c_i)$ denotes the probability of the true label.

3.2 Transfer Approaches

LSTM+ATT is used as our aspect-level model with parameter set $\theta_{aspect} = \{\mathbf{E}, \theta_{lstm}, \mathbf{W}_a, \mathbf{W}_o, \mathbf{b}_o\}$. We also build a standard LSTM-based classifier based on document-level training examples. This network is the same as the LSTM+ATT apart from the lack of the attention layer. The training objective is also cross entropy minimization as shown in equation (7) and the parameter set is $\theta_{doc} = \{\mathbf{E}', \theta'_{lstm}, \mathbf{W}'_o, \mathbf{b}'_o\}$.

Pretraining (PRET): In this setting, we first train on document-level examples. The last hidden vector from the LSTM outputs is used as the document representation. We initialize the relevant

	Dataset	Pos	Neg	Neu
D1	Restaurant14-Train	2164	807	637
	Restaurant14-Test	728	196	196
D2	Laptop14-Train	994	870	464
	Laptop14-Test	341	128	169
D3	Restaurant15-Train	1178	382	50
	Restaurant15-Test	439	328	35
D4	Restaurant16-Train	1620	709	88
	Restaurant16-Test	597	190	38

Table 1: Dataset description.

parameters \mathbf{E} , θ_{lstm} , \mathbf{W}_o , \mathbf{b}_o of LSTM+ATT with the pretrained weights, and train it on aspect-level examples to fine tune those weights and learn \mathbf{W}_a which is randomly initialized.

Multi-task Learning (MULT): This approach simultaneously trains two tasks – document-level and aspect-level classification. In this setting, the embedding layer (\mathbf{E}) and the LSTM layer (θ_{lstm}) are shared by both tasks, and a document is represented as the mean vector over LSTM outputs. The other parameters are task-specific. The overall loss function is then given by:

$$L = J + \lambda U \quad (8)$$

where U is the loss function of document-level classification. $\lambda \in (0, 1)$ is a hyperparameter that controls the weight of U .

Combined (PRET+MULT): In this setting, we first perform PRET on document-level examples. We use the pretrained weights for parameter initialization for both aspect-level model and document-level model, and then perform MULT as discussed above.

4 Experiments

4.1 Datasets and Experimental Settings

We run experiments on four benchmark aspect-level datasets, taken from SemEval 2014 (Pontiki et al., 2014), SemEval 2015 (Pontiki et al., 2015), and SemEval 2016 (Pontiki et al., 2016). Following previous work (Tang et al., 2016b; Wang et al., 2016), we remove samples with *conflicting* polarities in all datasets¹. Statistics of the resulting datasets are presented in Table 1.

We derived two document-level datasets from Yelp2014 (Tang et al., 2015) and the Amazon Electronics dataset (McAuley et al., 2015) respectively. The original reviews were rated on a 5-point scale. We consider 3-class classification and

¹We remove samples in the 2015/6 datasets if an opinion target is associated with different sentiment polarities.

thus label reviews with rating < 3 , > 3 , and $= 3$ as negative, positive, and neutral respectively. Each sampled dataset contains 30k instances with exactly balanced class labels. We pair up an aspect-level dataset and a document-level dataset when they are from a similar domain – the Yelp dataset is used by D1, D3, and D4 for PRET and MULT, and the Electronics dataset is only used by D2.

In all experiments, 300-dimension GloVe vectors (Pennington et al., 2014) are used to initialize \mathbf{E} and \mathbf{E}' when pretraining is not conducted for weight initialization. These vectors are also used for initializing \mathbf{E}' in the pretraining phase. Values for hyperparameters are obtained from experiments on development sets. We randomly sample 20% of the original training data from the aspect-level dataset as the development set and only use the remaining 80% for training. For all experiments, the dimension of LSTM hidden vectors is set to 300, λ is set to 0.1, and we use dropout with probability 0.5 on sentence/document representations before the output layer. We use RMSProp as the optimizer with the decay rate set to 0.9 and the base learning rate set to 0.001. The mini-batch size is set to 32.

4.2 Model Comparison

Table 2 shows the results of LSTM, LSTM+ATT, PRET, MULT, PRET+MULT, and four representative prior works (Tang et al., 2016a,b; Wang et al., 2016; Chen et al., 2017). Significance tests are conducted for testing the robustness of methods under random parameter initialization. Both accuracy and macro-F1 are used for evaluation as label distribution is unbalanced. The reported numbers are obtained as the average value over 5 runs with random initialization for each method.

We observe that PRET is very helpful, and consistently gives a 1–3% increase in accuracy over LSTM+ATT across all datasets. The improvements in macro-F1 scores are even more, especially on D3 and D4 where the labels are extremely unbalanced. MULT gives similar performance as LSTM+ATT on D1 and D2, but improvements can be clearly observed for D3 and D4. The combination (PRET+MULT) overall yields better results.

There are two main reasons why the improvements of macro-F1 scores are more significant on D3 and D4 than on D1: (1) D1 has much more neutral examples in the training set. A classifier

Methods	D1		D2		D3		D4	
	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1
Tang et al. (2016a)	75.37	64.51	68.25	65.96	76.39	58.70	82.16	54.21
Wang et al. (2016)	78.60	67.02	68.88	63.93	78.48	62.84	83.77	61.71
Tang et al. (2016b)	76.87	66.40	68.91	62.79	77.89	59.52	83.04	57.91
Chen et al. (2017)	78.48	68.54	72.08	68.43	79.98	60.57	83.88	62.14
LSTM	75.23	64.21	66.79	64.02	75.28	54.10	81.95	58.11
LSTM+ATT	76.83	66.48	68.07	64.82	77.38	60.52	82.73	59.12
Ours: PRET	78.28	68.55	71.32	68.53	80.67	68.31	84.87	70.73
Ours: MULT	77.41	66.68	68.65	64.57	81.05	65.69	83.27	64.56
Ours: PRET+MULT	79.11	69.73*	71.15	67.46	81.30*	68.74*	85.58*	69.76*

Table 2: Average accuracies and Macro-F1 scores over 5 runs with random initialization. The best results are in bold. * indicates that PRET+MULT is significantly better than Tang et al. (2016a), Wang et al. (2016), Tang et al. (2016b), Chen et al. (2017), LSTM, and LSTM+ATT with $p < 0.05$ according to one-tailed unpaired t-test.

Settings	D1		D2		D3		D4	
	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1	Acc.	Macro-F1
LSTM only	78.09	67.85	71.04	66.80	78.95	65.30	83.85	67.11
Embeddings only	77.12	67.19	69.12	65.06	80.13	67.04	84.12	70.11
Output layer only	76.88	66.81	69.63	66.07	78.30	64.49	82.55	62.83
Without LSTM	77.45	67.25	69.82	66.63	80.27	68.02	84.80	70.27
Without embeddings	77.97	67.96	70.59	67.16	79.08	65.56	83.94	68.79
Without output layer	78.36	68.06	71.10	67.87	80.82	67.68	84.71	70.48

Table 3: PRET with different transferred layers. Averaged results over 5 runs are reported.

without any external knowledge might still be able to learn some neutral-related features on D1 but it is very hard to learn from D3 and D4. (2) The numbers of neutral examples in the test sets of D3 and D4 are very small. Thus, the precision and recall on neutral class will be largely affected by even a small prediction difference (e.g., with 5 more neutral examples correctly identified, recall is increased by more than 10% on both datasets). As a result, the macro-F1 scores on D3 and D4 are affected more.

4.3 Ablation Tests

Table 2 indicates that a large percentage of the performance gain comes from PRET. To better understand the transfer effects of different layers – embedding layer (\mathbf{E}), LSTM layer (θ_{lstm}), and output layer ($\mathbf{W}_o, \mathbf{b}_o$) – we conduct ablation tests on PRET with different layers transferred from the document-level model to the aspect-level model. Results are presented in Table 3. “LSTM only” denotes the setting where only the LSTM layer is transferred, and “Without LSTM” denotes the setting where only the embedding and output layers are transferred (excluding the LSTM layer). The key observations are: (1) Transfer is helpful in all settings. Improvements over LSTM+ATT are observed even when only one layer is transferred. (2) Overall, transfers of the LSTM and embedding

layer are more useful than the output layer. This is what we expect, since the output layer is normally more task-specific. (3) Transfer of the embedding layer is more helpful on D3 and D4. One possible explanation is that the label distribution is extremely unbalanced on these two datasets. Sentiment information is not adequately captured by GloVe word embeddings. Therefore, with a small number of training examples in the negative and neutral classes, the embeddings trained by aspect-level classification still do not effectively capture the true semantics of the relevant opinion words. Transfer of the embedding layer can greatly help in this case.

4.4 Analysis

To show that aspect-level classification indeed benefits from document-level knowledge, we conduct experiments to vary the percentage of document-level training examples from 0.0 to 1.0 for PRET+MULT. The changes of accuracies and macro-F1 scores on the four datasets are shown in Figure 1. The improvements on accuracies with increasing number of document examples are stable across all datasets. For macro-F1 scores, the improvements on D1 and D2 are stable. We observe sharp increases in the macro-F1 scores of D3 and D4 when changing the percentage from 0 to 0.4. This may be related to their extremely

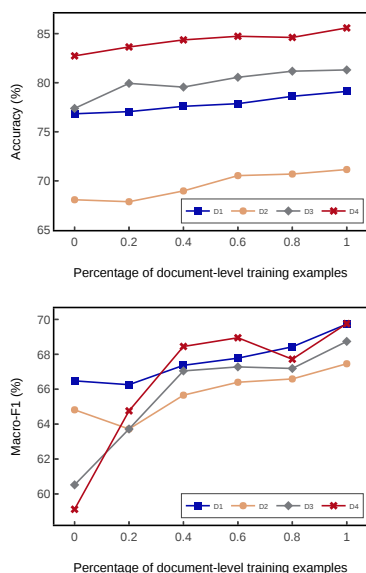


Figure 1: Results of PRET+MULT vs. percentage of document-level training data.

unbalanced label distribution. In such cases, with the knowledge gained from a small number of balanced document-level examples, aspect-level predictions on neutral examples can be significantly improved.

To better understand in which conditions the proposed method is helpful, we analyze a subset of test examples that are correctly classified by PRET+MULT but are misclassified by LSTM+ATT. We find that the benefits brought by document-level knowledge are typically shown in four ways.

First of all, to our surprise, LSTM+ATT made obvious mistakes on some instances with common opinion words. Below are two examples where the target is enclosed in [] with its true sentiment indicated in the subscript:

1. “I was highly disappointed in the [food]_{neg}.”
2. “This particular location certainly uses sub-standard [meats]_{neg}.”

In the above examples, LSTM+ATT does attend to the right opinion words, but makes the wrong predictions. One possible reason is that the word embeddings from GloVe without PRET do not effectively capture sentiment information, while the aspect-level training samples are not sufficient to capture that for certain words. PRET+MULT eliminates this kind of errors.

Another finding is that our method helps to better capture domain-specific opinion words due to additional knowledge from documents that are

from a similar domain:

1. “The smaller [size]_{pos} was a bonus because of space restrictions.”

2. “The [price]_{pos} is 200 dollars down.”

LSTM+ATT attends on *smaller* correctly for the first example but makes the wrong prediction as *smaller* can be negative in many cases. It does not even capture *down* in the second example.

Thirdly, we find that LSTM+ATT made a number of errors on sentences with negation words:

1. I have experienced no problems, [works]_{pos} as anticipated.

2. [Service]_{neg} not the friendliest to our party!

LSTMs typically only show their potential on large datasets. Without sufficient training examples, it may not be able to effectively capture various sequential patterns. Pretraining the network on larger document-level corpus addresses this problem.

Lastly, PRET+MULT makes fewer errors on recognizing neutral instances. This can also be observed from the macro-F1 scores in Table 2. The lack of training examples makes the prediction of neutral instances very difficult for all previous methods. Knowledge from document-level examples with balanced labels compensates for this disadvantage.

5 Conclusion

The effectiveness of existing aspect-level neural models is limited due to the difficulties in obtaining training data in practice. Our work is the first attempt to incorporate knowledge from document-level corpus for training aspect-level sentiment classifiers. We have demonstrated the effectiveness of our proposed approaches and analyzed the major benefits brought by the knowledge transfer. The proposed approaches can be potentially integrated with other aspect-level neural models to further boost their performance.

References

- Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. Recurrent attention network on memory for aspect sentiment analysis. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*.
- Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning (ICML 2008)*.

- Andrew M. Dai and Quoc V. Le. 2015. Semi-supervised sequence learning. In *Neural Information Processing Systems (NIPS 2015)*.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. Adaptive recursive neural network for target-dependent Twitter sentiment classification. In *Annual Meeting of the Association for Computational Linguistics (ACL 2014)*.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2017. An unsupervised neural attention model for aspect extraction. In *Annual Meeting of the Association for Computational Linguistics (ACL 2017)*.
- Alex Krizhevsky and Ilya Sutskever. 2012. Imagenet classification with deep convolutional neural networks. In *Neural Information Processing Systems (NIPS 2012)*.
- Jiangming Liu and Yue Zhang. 2017. Attention modeling for target sentiment. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL 2017)*.
- Yang Liu, Sujian Li, Xiaodong Zhang, and Zhifang Sui. 2016. Implicit discourse relation classification via multi-task neural networks. In *AAAI Conference on Artificial Intelligence (AAAI 2016)*.
- Minh-Tang Luong, Hieu Pham, and Christopher D. Manning. 2015b. Effective approaches to attention-based neural machine translation. In *Annual Meeting of the Association for Computational Linguistics (ACL 2015)*.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2015a. Multi-task sequence to sequence learning. In *International Conference on Learning Representation (ICLR 2015)*.
- Julian J. McAuley, Christopher Targett, Qinfeng Shi, and Anton van den Hengel. 2015. Image-based recommendations on styles and substitutes. In *The 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Neural Information Processing Systems (NIPS 2013)*.
- Lili Mou, Zhao Meng, Rui Yan, Ge Li, Yan Xu, Lu Zhang, and Zhi Jin. 2016. How transferable are neural networks in NLP applications? In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*.
- Thien Hai Nguyen and Kiyooki Shirai. 2015. PhraseRNN: Phrase recursive neural network for aspect-based sentiment analysis. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global vectors for word representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*.
- Maria Pontiki, Dimitrios Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. SemEval-2015 task 12: Aspect based sentiment analysis. In *International Workshop on Semantic Evaluation (SemEval 2015)*.
- Maria Pontiki, Dimitrios Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. SemEval-2014 task 4: Aspect based sentiment analysis. In *International Workshop on Semantic Evaluation (SemEval 2014)*.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammed AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Veronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Núria Bel, Salud Maria Jiménez-Zafra, and Gülşen Eryiğit. 2016. SemEval-2016 task 5: Aspect based sentiment analysis. In *International Workshop on Semantic Evaluation (SemEval 2016)*.
- Prajit Ramachandran, Peter J. Liu, and Quoc V. Le. 2017. Unsupervised pretraining for sequence to sequence learning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2017)*.
- Duyu Tang, Bing Qin, Xiaocheng Feng, and Ting Liu. 2016a. Effective LSTMs for target-dependent sentiment classification. In *International Conference on Computational Linguistics (COLING 2016)*.
- Duyu Tang, Bing Qin, and Ting Liu. 2015. Learning semantic representation of users and products for document level sentiment classification. In *Annual Meeting of the Association for Computational Linguistics (ACL 2015)*.
- Duyu Tang, Bing Qin, and Ting Liu. 2016b. Aspect level sentiment classification with deep memory network. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*.
- Duy-Tin Vo and Yue Zhang. 2015. Target-dependent Twitter sentiment classification with rich automatic features. In *International Joint Conference on Artificial Intelligence (IJCAI 2015)*.
- Yequan Wang, Minlie Huang, Li Zhao, and Xiaoyan Zhu. 2016. Attention-based LSTM for aspect-level sentiment classification. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*.
- Matthew D. Zeiler and Rob Fergus. 2014. Visualizing and understanding convolutional networks. In *European Conference on Computer Vision (ECCV 2014)*.

Meishan Zhang, Yue Zhang, and Duy-Tin Vo. 2016. Gated neural networks for targeted sentiment analysis. In *AAAI Conference on Artificial Intelligence (AAAI 2016)*.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2016)*.