

Exploiting BERT for End-to-End Aspect-based Sentiment Analysis*

Xin Li¹, Lidong Bing², Wenxuan Zhang¹ and Wai Lam¹

¹Department of Systems Engineering and Engineering Management
The Chinese University of Hong Kong, Hong Kong

²R&D Center Singapore, Machine Intelligence Technology, Alibaba DAMO Academy
{lixin, wxzhang, wlam}@se.cuhk.edu.hk
l.bing@alibaba-inc.com

Abstract

In this paper, we investigate the modeling power of contextualized embeddings from pre-trained language models, e.g. BERT, on the E2E-ABSA task. Specifically, we build a series of simple yet insightful neural baselines to deal with E2E-ABSA. The experimental results show that even with a simple linear classification layer, our BERT-based architecture can outperform state-of-the-art works. Besides, we also standardize the comparative study by consistently utilizing a hold-out development dataset for model selection, which is largely ignored by previous works. Therefore, our work can serve as a BERT-based benchmark for E2E-ABSA.¹

1 Introduction

Aspect-based sentiment analysis (ABSA) is to discover the users' sentiment or opinion towards an aspect, usually in the form of explicitly mentioned aspect terms (Mitchell et al., 2013; Zhang et al., 2015) or implicit aspect categories (Wang et al., 2016), from user-generated natural language texts (Liu, 2012). The most popular ABSA benchmark datasets are from SemEval ABSA challenges (Pontiki et al., 2014, 2015, 2016) where a few thousand review sentences with gold standard aspect sentiment annotations are provided.

Table 1 summarizes three existing research problems related to ABSA. The first one is the original ABSA, aiming at predicting the sentiment polarity of the sentence towards the given aspect. Compared to this classification problem, the second one and the third one, namely, Aspect-oriented Opinion Words Extraction (AOWE) (Fan

et al., 2019) and End-to-End Aspect-based Sentiment Analysis (E2E-ABSA) (Ma et al., 2018a; Schmitt et al., 2018; Li et al., 2019a; Li and Lu, 2017, 2019), are related to a sequence tagging problem. Precisely, the goal of AOWE is to extract the aspect-specific opinion words from the sentence given the aspect. The goal of E2E-ABSA is to jointly detect aspect terms/categories and the corresponding aspect sentiments.

Many neural models composed of a task-agnostic pre-trained word embedding layer and task-specific neural architecture have been proposed for the original ABSA task (i.e. the aspect-level sentiment classification) (Tang et al., 2016; Wang et al., 2016; Chen et al., 2017; Liu and Zhang, 2017; Ma et al., 2017, 2018b; Majumder et al., 2018; Li et al., 2018; He et al., 2018; Xue and Li, 2018; Wang et al., 2018; Fan et al., 2018; Huang and Carley, 2018; Lei et al., 2019; Li et al., 2019b)², but the improvement of these models measured by the accuracy or F1 score has reached a bottleneck. One reason is that the task-agnostic embedding layer, usually a linear layer initialized with Word2Vec (Mikolov et al., 2013) or GloVe (Pennington et al., 2014), only provides context-independent word-level features, which is insufficient for capturing the complex semantic dependencies in the sentence. Meanwhile, the size of existing datasets is too small to train sophisticated task-specific architectures. Thus, introducing a context-aware word embedding³ layer pre-trained on large-scale datasets with deep LSTM (McCann et al., 2017; Peters et al., 2018; Howard and Ruder, 2018) or Transformer (Radford et al., 2018, 2019; Devlin et al., 2019; Lample

*The work described in this paper is substantially supported by a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14204418).

¹Our code is open-source and available at: <https://github.com/lixin4ever/BERT-E2E-ABSA>

²Due to the limited space, we can not list all of the existing works here, please refer to the survey (Zhou et al., 2019) for more related papers.

³In this paper, we generalize the concept of "word embedding" as a mapping between the word and the low-dimensional word representations.

Sentence:	<Great> [food] _P but the [service] _N is <dreadful>.	
Settings	Input	Output
1. ABSA	sentence, aspect	aspect sentiment
2. AOWE	sentence, aspect	opinion words
3. E2E-ABSA	sentence	aspect, aspect sentiment

Table 1: Different problem settings in ABSA. Gold standard aspects and opinions are wrapped in [] and <> respectively. The subscripts N and P refer to aspect sentiment. Underline $\tilde{*}$ or $\tilde{*}$ indicates the association between the aspect and the opinion.

and Conneau, 2019; Yang et al., 2019; Dong et al., 2019) for fine-tuning a lightweight task-specific network using the labeled data has good potential for further enhancing the performance.

Xu et al. (2019); Sun et al. (2019); Song et al. (2019); Yu and Jiang (2019); Rietzler et al. (2019); Huang and Carley (2019) have conducted some initial attempts to couple the deep contextualized word embedding layer with downstream neural models for the original ABSA task and establish the new state-of-the-art results. It encourages us to explore the potential of using such contextualized embeddings to the more difficult but practical task, i.e. E2E-ABSA (the third setting in Table 1).⁴ Note that we are not aiming at developing a task-specific architecture, instead, our focus is to examine the potential of contextualized embedding for E2E-ABSA, coupled with various simple layers for prediction of E2E-ABSA labels.⁵

In this paper, we investigate the modeling power of BERT (Devlin et al., 2019), one of the most popular pre-trained language model armed with Transformer (Vaswani et al., 2017), on the task of E2E-ABSA. Concretely, inspired by the investigation of E2E-ABSA in Li et al. (2019a), which predicts aspect boundaries as well as aspect sentiments using a single sequence tagger, we build a series of simple yet insightful neural baselines for the sequence labeling problem and fine-tune the task-specific components with BERT or deem BERT as feature extractor. Besides, we standardize the comparative study by consistently utilizing the hold-out development dataset for model selection, which is ignored in most of the existing

⁴Both of ABSA and AOWE assume that the aspects in a sentence are given. Such setting makes them less practical in real-world scenarios since manual annotation of the fine-grained aspect mentions/categories is quite expensive.

⁵Hu et al. (2019) introduce BERT to handle the E2E-ABSA problem but their focus is to design a task-specific architecture rather than exploring the potential of BERT.

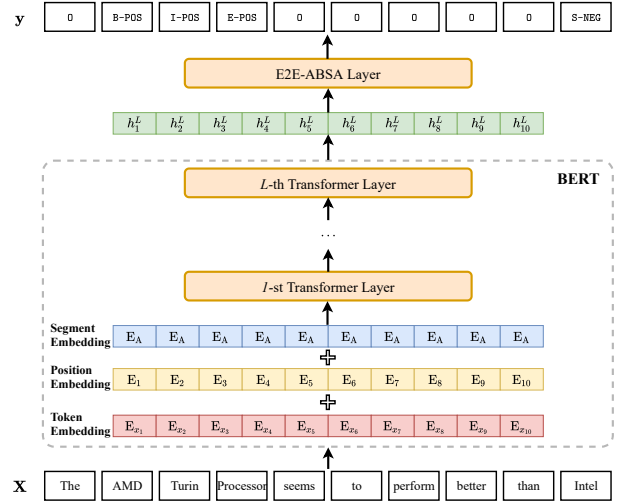


Figure 1: Overview of the designed model.

ABSA works (Tay et al., 2018).

2 Model

In this paper, we focus on the aspect term-level End-to-End Aspect-Based Sentiment Analysis (E2E-ABSA) problem setting. This task can be formulated as a sequence labeling problem. The overall architecture of our model is depicted in Figure 1. Given the input token sequence $\mathbf{x} = \{x_1, \dots, x_T\}$ of length T , we firstly employ BERT component with L transformer layers to calculate the corresponding contextualized representations $H^L = \{h_1^L, \dots, h_T^L\} \in \mathbb{R}^{T \times \dim_h}$ for the input tokens where \dim_h denotes the dimension of the representation vector. Then, the contextualized representations are fed to the task-specific layers to predict the tag sequence $\mathbf{y} = \{y_1, \dots, y_T\}$. The possible values of the tag y_t are B- $\{\text{POS}, \text{NEG}, \text{NEU}\}$, I- $\{\text{POS}, \text{NEG}, \text{NEU}\}$, E- $\{\text{POS}, \text{NEG}, \text{NEU}\}$, S- $\{\text{POS}, \text{NEG}, \text{NEU}\}$ or O, denoting the beginning of aspect, inside of aspect, end of aspect, single-word aspect, with positive, negative or neutral sentiment respectively, as well as outside of aspect.

2.1 BERT as Embedding Layer

Compared to the traditional Word2Vec- or GloVe-based embedding layer which only provides a single context-independent representation for each token, the BERT embedding layer takes the sentence as input and calculates the token-level representations using the information from the entire sentence. First of all, we pack the input features as $H^0 = \{e_1, \dots, e_T\}$, where e_t ($t \in [1, T]$) is

the combination of the token embedding, position embedding and segment embedding corresponding to the input token x_t . Then L transformer layers are introduced to refine the token-level features layer by layer. Specifically, the representations $H^l = \{h_1^l, \dots, h_T^l\}$ at the l -th ($l \in [1, L]$) layer are calculated below:

$$H^l = \text{Transformer}_l(H^{l-1}) \quad (1)$$

We regard H^L as the contextualized representations of the input tokens and use them to perform the predictions for the downstream task.

2.2 Design of Downstream Model

After obtaining the BERT representations, we design a neural layer, called E2E-ABSA layer in Figure 1, on top of BERT embedding layer for solving the task of E2E-ABSA. We investigate several different design for the E2E-ABSA layer, namely, linear layer, recurrent neural networks, self-attention networks, and conditional random fields layer.

Linear Layer The obtained token representations can be directly fed to linear layer with softmax activation function to calculate the token-level predictions:

$$P(y_t|x_t) = \text{softmax}(W_o h_t^L + b_o) \quad (2)$$

where $W_o \in \mathbb{R}^{\dim_h \times |\mathcal{Y}|}$ is the learnable parameters of the linear layer.

Recurrent Neural Networks Considering its sequence labeling formulation, Recurrent Neural Networks (RNN) (Elman, 1990) is a natural solution for the task of E2E-ABSA. In this paper, we adopt GRU (Cho et al., 2014), whose superiority compared to LSTM (Hochreiter and Schmidhuber, 1997) and basic RNN has been verified in Jozefowicz et al. (2015). The computational formula of the task-specific hidden representation $h_t^T \in \mathbb{R}^{\dim_h}$ at the t -th time step is shown below:

$$\begin{aligned} \begin{bmatrix} r_t \\ z_t \end{bmatrix} &= \sigma(\text{LN}(W_x h_t^L) + \text{LN}(W_h h_{t-1}^T)) \\ n_t &= \tanh(\text{LN}(W_{xn} h_t^L) + r_t * \text{LN}(W_{hn} h_{t-1}^T)) \\ h_t^T &= (1 - z_t) * n_t + z_t * h_{t-1}^T \end{aligned} \quad (3)$$

where σ is the sigmoid activation function and r_t, z_t, n_t respectively denote the reset gate, update gate and new gate. $W_x, W_h \in \mathbb{R}^{2\dim_h \times \dim_h}$, $W_{xn}, W_{hn} \in \mathbb{R}^{\dim_h \times \dim_h}$ are the parameters of

GRU. Since directly applying RNN on the output of transformer, namely, the BERT representation h_t^L , may lead to unstable training (Chen et al., 2018; Liu, 2019), we add additional layer-normalization (Ba et al., 2016), denoted as LN, when calculating the gates. Then, the predictions are obtained by introducing a softmax layer:

$$p(y_t|x_t) = \text{softmax}(W_o h_t^T + b_o) \quad (4)$$

Self-Attention Networks With the help of self attention (Cheng et al., 2016; Lin et al., 2017), Self-Attention Network (Vaswani et al., 2017; Shen et al., 2018) is another effective feature extractor apart from RNN and CNN. In this paper, we introduce two SAN variants to build the task-specific token representations $H^T = \{h_1^T, \dots, h_T^T\}$. One variant is composed of a simple self-attention layer and residual connection (He et al., 2016), dubbed as ‘‘SAN’’. The computational process of SAN is below:

$$\begin{aligned} H^T &= \text{LN}(H^L + \text{SLF-ATT}(Q, K, V)) \\ Q, K, V &= H^L W^Q, H^L W^K, H^L W^V \end{aligned} \quad (5)$$

where SLF-ATT is identical to the self-attentive scaled dot-product attention (Vaswani et al., 2017). Another variant is a transformer layer (dubbed as ‘‘TFM’’), which has the same architecture with the transformer encoder layer in the BERT. The computational process of TFM is as follows:

$$\begin{aligned} \hat{H}^L &= \text{LN}(H^L + \text{SLF-ATT}(Q, K, V)) \\ H^T &= \text{LN}(\hat{H}^L + \text{FFN}(\hat{H}^L)) \end{aligned} \quad (6)$$

where FFN refers to the point-wise feed-forward networks (Vaswani et al., 2017). Again, a linear layer with softmax activation is stacked on the designed SAN/TFM layer to output the predictions (same with that in Eq(4)).

Conditional Random Fields Conditional Random Fields (CRF) (Lafferty et al., 2001) is effective in sequence modeling and has been widely adopted for solving the sequence labeling tasks together with neural models (Huang et al., 2015; Lample et al., 2016; Ma and Hovy, 2016). In this paper, we introduce a linear-chain CRF layer on top of the BERT embedding layer. Different from the above mentioned neural models maximizing the token-level likelihood $p(y_t|x_t)$, the CRF-based model aims to find the globally most

Model		LAPTOP			REST		
		P	R	F1	P	R	F1
Existing Models	(Li et al., 2019a) [‡]	61.27	54.89	57.90	68.64	71.01	69.80
	(Luo et al., 2019) [‡]	-	-	60.35	-	-	72.78
	(He et al., 2019) [‡]	-	-	58.37	-	-	-
LSTM-CRF	(Lample et al., 2016) [‡]	58.61	50.47	54.24	66.10	66.30	66.20
	(Ma and Hovy, 2016) [‡]	58.66	51.26	54.71	61.56	67.26	64.29
	(Liu et al., 2018) [‡]	53.31	59.40	56.19	68.46	64.43	66.38
BERT Models	BERT+Linear	62.16	58.90	60.43	71.42	75.25	73.22
	BERT+GRU	61.88	60.47	61.12	70.61	76.20	73.24
	BERT+SAN	62.42	58.71	60.49	72.92	76.72	74.72
	BERT+TFM	63.23	58.64	60.80	72.39	76.64	74.41
	BERT+CRF	62.22	59.49	60.78	71.88	76.48	74.06

Table 2: Main results. The symbol [‡] denotes the numbers are officially reported ones. The results with [#] are retrieved from Li et al. (2019a).

Dataset		Train	Dev	Test	Total
LAPTOP	# sent	2741	304	800	4245
	# aspect	2041	256	634	2931
REST	# sent	3490	387	2158	6035
	# aspect	3893	413	2287	6593

Table 3: Statistics of datasets.

probable tag sequence. Specifically, the sequence-level scores $s(\mathbf{x}, \mathbf{y})$ and likelihood $p(\mathbf{y}|\mathbf{x})$ of $\mathbf{y} = \{y_1, \dots, y_T\}$ are calculated as follows:

$$s(\mathbf{x}, \mathbf{y}) = \sum_{t=0}^T M_{y_t, y_{t+1}}^A + \sum_{t=1}^T M_{t, y_t}^P \quad (7)$$

$$p(\mathbf{y}|\mathbf{x}) = \text{softmax}(s(\mathbf{x}, \mathbf{y}))$$

where $M^A \in \mathbb{R}^{|\mathcal{Y}| \times |\mathcal{Y}|}$ is the randomly initialized transition matrix for modeling the dependency between the adjacent predictions and $M^P \in \mathbb{R}^{T \times |\mathcal{Y}|}$ denote the emission matrix linearly transformed from the BERT representations H^L . The softmax here is conducted over all of the possible tag sequences. As for the decoding, we regard the tag sequence with the highest scores as output:

$$\mathbf{y}^* = \arg \max_{\mathbf{y}} s(\mathbf{x}, \mathbf{y}) \quad (8)$$

where the solution is obtained via Viterbi search.

3 Experiment

3.1 Dataset and Settings

We conduct experiments on two review datasets originating from SemEval (Pontiki et al., 2014, 2015, 2016) but re-prepared in Li et al. (2019a). The statistics are summarized in Table 3. We use the pre-trained ‘bert-base-uncased’ model⁶,

⁶<https://github.com/huggingface/transformers>




Figure 2: Performances on the Dev set of REST.

where the number of transformer layers $L = 12$ and the hidden size dim_h is 768. For the downstream E2E-ABSA component, we consistently use the single-layer architecture and set the dimension of task-specific representation as dim_h . The learning rate is $2e-5$. The batch size is set as 25 for LAPTOP and 16 for REST. We train the model up to 1500 steps. After training 1000 steps, we conduct model selection on the development set for very 100 steps according to the micro-averaged F1 score. Following these settings, we train 5 models with different random seeds and report the average results.

We compare with **Existing Models**, including tailor-made E2E-ABSA models (Li et al., 2019a; Luo et al., 2019; He et al., 2019), and competitive **LSTM-CRF** sequence labeling models (Lample et al., 2016; Ma and Hovy, 2016; Liu et al., 2018).

3.2 Main Results

From Table 2, we surprisingly find that only introducing a simple token-level classifier, namely, BERT-Linear, already outperforms the existing

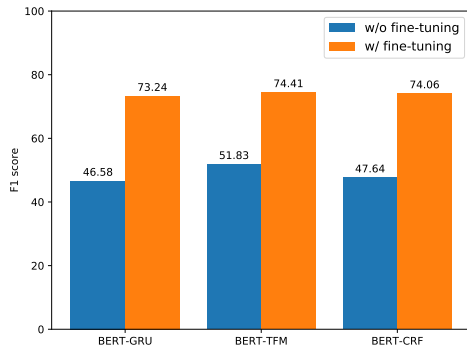


Figure 3: Effect of fine-tuning BERT.

works without using BERT, suggesting that BERT representations encoding the associations between arbitrary two tokens largely alleviate the issue of context independence in the linear E2E-ABSA layer. It is also observed that slightly more powerful E2E-ABSA layers lead to much better performance, verifying the postulation that incorporating context helps to sequence modeling.

3.3 Over-parameterization Issue

Although we employ the smallest pre-trained BERT model, it is still over-parameterized for the E2E-ABSA task (110M parameters), which naturally raises a question: does BERT-based model tend to overfit the small training set? Following this question, we train BERT-GRU, BERT-TFM and BERT-CRF up to 3000 steps on REST and observe the fluctuation of the F1 measures on the development set. As shown in Figure 2, F1 scores on the development set are quite stable and do not decrease much as the training proceeds, which shows that the BERT-based model is exceptionally robust to overfitting.

3.4 Finetuning BERT or Not

We also study the impact of fine-tuning on the final performances. Specifically, we employ BERT to calculate the contextualized token-level representations but kept the parameters of BERT component unchanged in the training phase. Figure 3 illustrate the comparative results between the BERT-based models and those keeping BERT component fixed. Obviously, the general purpose BERT representation is far from satisfactory for the downstream tasks and task-specific fine-tuning is essential for exploiting the strengths of BERT to improve the performance.

4 Conclusion

In this paper, we investigate the effectiveness of BERT embedding component on the task of End-to-End Aspect-Based Sentiment Analysis (E2E-ABSA). Specifically, we explore to couple the BERT embedding component with various neural models and conduct extensive experiments on two benchmark datasets. The experimental results demonstrate the superiority of BERT-based models on capturing aspect-based sentiment and their robustness to overfitting.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450*.
- Mia Xu Chen, Orhan Firat, Ankur Bapna, Melvin Johnson, Wolfgang Macherey, George Foster, Llion Jones, Mike Schuster, Noam Shazeer, Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Zhifeng Chen, Yonghui Wu, and Macduff Hughes. 2018. The best of both worlds: Combining recent advances in neural machine translation. In *ACL*, pages 76–86.
- Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. Recurrent attention network on memory for aspect sentiment analysis. In *EMNLP*, pages 452–461.
- Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memory-networks for machine reading. In *EMNLP*, pages 551–561.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *EMNLP*, pages 1724–1734.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *arXiv preprint arXiv:1905.03197*.
- Jeffrey L Elman. 1990. Finding structure in time. *Cognitive science*, 14(2):179–211.
- Feifan Fan, Yansong Feng, and Dongyan Zhao. 2018. Multi-grained attention network for aspect-level sentiment classification. In *EMNLP*, pages 3433–3442.

- Zhifang Fan, Zhen Wu, Xin-Yu Dai, Shujian Huang, and Jiajun Chen. 2019. Target-oriented opinion words extraction with target-fused neural sequence labeling. In *NAACL*, pages 2509–2518.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *CVPR*, pages 770–778.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. Exploiting document knowledge for aspect-level sentiment classification. In *ACL*, pages 579–585.
- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2019. An interactive multi-task learning network for end-to-end aspect-based sentiment analysis. In *ACL*, pages 504–515.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. In *ACL*, pages 328–339.
- Minghao Hu, Yuxing Peng, Zhen Huang, Dongsheng Li, and Yiwei Lv. 2019. Open-domain targeted sentiment analysis via span-based extraction and classification. In *ACL*, pages 537–546.
- Binxuan Huang and Kathleen Carley. 2018. Parameterized convolutional neural networks for aspect level sentiment classification. In *EMNLP*, pages 1091–1096.
- Binxuan Huang and Kathleen M Carley. 2019. Syntax-aware aspect level sentiment classification with graph attention networks. *arXiv preprint arXiv:1909.02606*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional lstm-crf models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Rafal Jozefowicz, Wojciech Zaremba, and Ilya Sutskever. 2015. An empirical exploration of recurrent network architectures. In *ICML*, pages 2342–2350.
- John D Lafferty, Andrew McCallum, and Fernando CN Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *ICML*, pages 282–289.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural architectures for named entity recognition. In *NAACL*, pages 260–270.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *arXiv preprint arXiv:1901.07291*.
- Zeyang Lei, Yujiu Yang, Min Yang, Wei Zhao, Jun Guo, and Yi Liu. 2019. A human-like semantic cognition network for aspect-level sentiment classification. In *AAAI*, pages 6650–6657.
- Hao Li and Wei Lu. 2017. Learning latent sentiment scopes for entity-level sentiment analysis. In *AAAI*, pages 3482–3489.
- Hao Li and Wei Lu. 2019. Learning explicit and implicit structures for targeted sentiment analysis. *arXiv preprint arXiv:1909.07593*.
- Xin Li, Lidong Bing, Wai Lam, and Bei Shi. 2018. Transformation networks for target-oriented sentiment classification. In *ACL*, pages 946–956.
- Xin Li, Lidong Bing, Piji Li, and Wai Lam. 2019a. A unified model for opinion target extraction and target sentiment prediction. In *AAAI*, pages 6714–6721.
- Zheng Li, Ying Wei, Yu Zhang, Xiang Zhang, and Xin Li. 2019b. Exploiting coarse-to-fine task transfer for aspect-level sentiment classification. In *AAAI*, pages 4253–4260.
- Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. In *ICLR*.
- Bing Liu. 2012. Sentiment analysis and opinion mining. *Synthesis lectures on human language technologies*, 5(1):1–167.
- Jiangming Liu and Yue Zhang. 2017. Attention modeling for targeted sentiment. In *EACL*, pages 572–577.
- Liyuan Liu, Jingbo Shang, Xiang Ren, Frank F Xu, Huan Gui, Jian Peng, and Jiawei Han. 2018. Empower sequence labeling with task-aware neural language model. In *AAAI*, pages 5253–5260.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Huaishao Luo, Tianrui Li, Bing Liu, and Junbo Zhang. 2019. DOER: Dual cross-shared RNN for aspect term-polarity co-extraction. In *ACL*, pages 591–601.
- Dehong Ma, Sujian Li, and Houfeng Wang. 2018a. Joint learning for targeted sentiment analysis. In *EMNLP*, pages 4737–4742.
- Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2017. Interactive attention networks for aspect-level sentiment classification. In *IJCAI*, pages 4068–4074.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional LSTM-CNNs-CRF. In *ACL*, pages 1064–1074.

- Yukun Ma, Haiyun Peng, and Erik Cambria. 2018b. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive lstm. In *AAAI*.
- Navonil Majumder, Soujanya Poria, Alexander Gelbukh, Md. Shad Akhtar, Erik Cambria, and Asif Ekbal. 2018. IARM: Inter-aspect relation modeling with memory networks in aspect-based sentiment analysis. In *EMNLP*, pages 3402–3411.
- Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. 2017. Learned in translation: Contextualized word vectors. In *NeurIPS*, pages 6294–6305.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, pages 3111–3119.
- Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme. 2013. Open domain targeted sentiment. In *EMNLP*, pages 1643–1654.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL*, pages 2227–2237.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. SemEval-2016 task 5: Aspect based sentiment analysis. In *SemEval*, pages 19–30.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. SemEval-2015 task 12: Aspect based sentiment analysis. In *SemEval*, pages 486–495.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. SemEval-2014 task 4: Aspect based sentiment analysis. In *SemEval*, pages 27–35.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8).
- Alexander Rietzler, Sebastian Stabinger, Paul Opitz, and Stefan Engl. 2019. Adapt or get left behind: Domain adaptation through bert language model finetuning for aspect-target sentiment classification. *arXiv preprint arXiv:1908.11860*.
- Martin Schmitt, Simon Steinheber, Konrad Schreiber, and Benjamin Roth. 2018. Joint aspect and polarity classification for aspect-based sentiment analysis with end-to-end neural networks. In *EMNLP*, pages 1109–1114.
- Tao Shen, Tianyi Zhou, Guodong Long, Jing Jiang, Shirui Pan, and Chengqi Zhang. 2018. Disan: Directional self-attention network for rnn/cnn-free language understanding. In *AAAI*.
- Youwei Song, Jiahai Wang, Tao Jiang, Zhiyue Liu, and Yanghui Rao. 2019. Attentional encoder network for targeted sentiment classification. *arXiv preprint arXiv:1902.09314*.
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence. In *NAACL*, pages 380–385.
- Duyu Tang, Bing Qin, and Ting Liu. 2016. Aspect level sentiment classification with deep memory network. In *EMNLP*, pages 214–224.
- Yi Tay, Luu Anh Tuan, and Siu Cheung Hui. 2018. Learning to attend via word-aspect associative fusion for aspect-based sentiment analysis. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*, pages 5998–6008.
- Shuai Wang, Sahisnu Mazumder, Bing Liu, Mianwei Zhou, and Yi Chang. 2018. Target-sensitive memory networks for aspect sentiment classification. In *ACL*, pages 957–967.
- Yequan Wang, Minlie Huang, Xiaoyan Zhu, and Li Zhao. 2016. Attention-based LSTM for aspect-level sentiment classification. In *EMNLP*, pages 606–615.
- Hu Xu, Bing Liu, Lei Shu, and Philip Yu. 2019. BERT post-training for review reading comprehension and aspect-based sentiment analysis. In *NAACL*, pages 2324–2335.
- Wei Xue and Tao Li. 2018. Aspect based sentiment analysis with gated convolutional networks. In *ACL*, pages 2514–2523.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.

- Jianfei Yu and Jing Jiang. 2019. Adapting bert for target-oriented multimodal sentiment classification. In *IJCAI*, pages 5408–5414.
- Meishan Zhang, Yue Zhang, and Duy-Tin Vo. 2015. Neural networks for open domain targeted sentiment. In *EMNLP*, pages 612–621.
- Jie Zhou, Jimmy Xiangji Huang, Qin Chen, Qinmin Vivian Hu, Tingting Wang, and Liang He. 2019. Deep learning for aspect-level sentiment classification: Survey, vision and challenges. *IEEE Access*.