

# Eliminating Sentiment Bias for Aspect-Level Sentiment Classification with Unsupervised Opinion Extraction

Bo Wang<sup>1</sup>, Tao Shen<sup>2</sup>, Guodong Long<sup>2</sup>, Tianyi Zhou<sup>3,4</sup>, Yi Chang<sup>1,5,✉</sup>

<sup>1</sup>School of Artificial Intelligence, Jilin University

<sup>2</sup>Australian AI Institute, School of CS, FEIT, University of Technology Sydney

<sup>3</sup>University of Washington, Seattle; <sup>4</sup>University of Maryland, College Park

<sup>5</sup>International Center of Future Science, Jilin University

bowang19@mails.jlu.edu.cn, {tao.shen, guodong.long}@uts.edu.au,  
tianyizh@uw.edu, yichang@jlu.edu.cn

## Abstract

Aspect-level sentiment classification (ALSC) aims at identifying the sentiment polarity of a specified aspect in a sentence. ALSC is a practical setting in aspect-based sentiment analysis due to no opinion term labeling needed, but it fails to interpret why a sentiment polarity is derived for the aspect. To address this problem, recent works fine-tune pre-trained Transformer encoders for ALSC to extract an aspect-centric dependency tree that can locate the opinion words. However, the induced opinion words only provide an intuitive cue far below human-level interpretability. Besides, the pre-trained encoder tends to internalize an aspect’s intrinsic sentiment, causing sentiment bias and thus affecting model performance. In this paper, we propose a span-based anti-bias aspect representation learning framework. It first eliminates the sentiment bias in the aspect embedding by adversarial learning against aspects’ prior sentiment. Then, it aligns the distilled opinion candidates with the aspect by span-based dependency modeling to highlight the interpretable opinion terms. Our method achieves new state-of-the-art performance on five benchmarks, with the capability of unsupervised opinion extraction.

## 1 Introduction

Aspect-based sentiment analysis (ABSA) (Jiang et al., 2011) aims to determine sentiment polarity w.r.t. a specified aspect term in a piece of text. For example, in “*The food is tasty, but the service is terrible*”, the sentiment towards aspect term (AT) “*food*” is positive according to the opinion term (OT) “*tasty*”, while the sentiment towards “*service*” is negative according to “*terrible*”. The most comprehensive setting of ABSA is aspect sentiment triplet extraction (ASTE) (Peng et al., 2020) consisting of a series of subtasks, i.e., aspect extraction, aspect-level sentiment classification and opinion extraction. Thereby, given a piece of text, ASTE can produce a set of triples, i.e., (*Aspect*

*Term, Sentiment, Opinion Term*), to describe sentiment with details of What, How and Why, so it enjoys full interpretability. Continuing the above example, ASTE can generate triples like (*food, Positive, tasty*). However, the human annotation on opinion terms is much more label-intensive than traditional sentiment analysis task.

Therefore, by following many recent works, we target the practical subtask of ASTE, called aspect-level sentiment classification (ALSC). It predicts a three-categorical sentiment (i.e., *positive, neutral* or *negative*) of a given aspect term in a sentence. Most recent works capture the modification relation between aspect and opinion terms in an implicit manner, which is usually achieved by integrating graph neural networks (GNNs) over dependency parsing tree into text representation learning (Zhang et al., 2019a; Sun et al., 2019b; Tang et al., 2020; Wang et al., 2020). Further, the performance can be significantly boosted when incorporating pre-trained Transformer encoder, e.g., BERT (Devlin et al., 2019) in a fine-tuning paradigm (Sun et al., 2019a; Tang et al., 2020; Wang et al., 2020; Chen et al., 2020). Although these works achieve excellent results even close to humans, they cannot derive the interpretability to explain *why* an aspect is associated with the polarity prediction.

Luckily, pre-trained Transformer encoder can also be used to explain linguistic knowledge underlying the given text via dependency probing (Clark et al., 2019; Wu et al., 2020). This has been exploited by Dai et al. (2021) to reveal that, after fine-tuning the pre-trained encoder on ALSC, an aspect-centric dependency tree can be induced to highlight the modifier of an aspect. Intuitively, the highlighted modifier is viewed as an opinion word of the corresponding aspect, which thus, to some extent, brings the interpretability back.

Nevertheless, compared with the span-formatted opinion terms in ASTE, the opinion word illustrated by the induced dependency tree can only pro-

vide an intuitive, noisy, word-level sentiment cue but is far from the human-level interpretability as in opinion extraction. What’s worse, as verified by Huang et al. (2020), the Transformer pre-trained on large-scale raw corpora tends to internalize terms’ intrinsic attributes, so it causes the problem of *sentiment bias* when generating text given a particular term prompt. In ALSC scenario, we found that *sentiment bias* also exists and affects a model to determine the sentiment of an aspect term regardless of its contextual information (e.g., opinion terms). The bias is especially obvious for the aspects that can imply strong sentiment themselves. For example, for “*There’s candlelight and music*”, a model based on pre-trained Transformer is likely to misclassify the sentiment towards “*music*” as positive, whereas the oracle label is neutral.

In this work, based on the pre-trained Transformer encoder to ensure state-of-the-art ALSC performance, we aim to eliminate sentiment bias in the ALSC scenario while equip the model with human-level aspect-opinion interpretability.

To this end, we propose a Span-based Anti-bias aspect Representation Learning (SARL) framework for ALSC with unsupervised opinion extraction. First, instead of widely feeding a concatenation of sentence and aspect into a pre-trained Transformer, we adopt a span-level paradigm (Hu et al., 2019; Zhao et al., 2020) which focuses on deriving span representation of an aspect term. Then, we propose an anti-bias aspect encoding module to eliminate the sentiment bias existing in aspect representations, which is achieved by an adversarial learning against the aspect’s prior sentiment indicated in SentiWordNet (Esuli and Sebastiani, 2006). Next, built upon sentiment-agnostic aspect representation from the above encoder, we propose an aspect-opinion dependency alignment module to capture explicit modifications from opinion term candidates to the targeted aspect, and integrate the modifications into aspect representation via gating. The integrated representation is lastly passed into a neural classifier for sentiment prediction.

For pairwise aspect-opinion alignment via a span-based model, our work share a high-level inspiration with SpanMlt (Zhao et al., 2020) but differs in that, SpanMlt targets fully-supervised aspect and opinion terms extraction whereas ours leverages the alignment to empower aspect-opinion interpretability even without opinion supervisions.

Our main contributions are as follows:

- We propose a span-based aspect encoding

module to alleviate the sentiment bias problem from which pre-trained Transformers suffer.

- By a dependency aligner, our model can derive human-level aspect-opinion interpretability w/o opinion extraction supervisions.
- We achieve new state-of-the-art results on 5 ALSC datasets with extensive analyses and present metrics to measure the interpretability.

## 2 Methodology

This section begins with a task definition of ALSC. Then, we present our Span-based Anti-bias aspect Representation Learning (SARL) framework (as in Figure 1) consisting of an adversarial anti-bias aspect encoder (§2.1) and a distilled aspect-opinion dependency aligner (§2.2). Lastly, we detail the training and inference of the proposed model (§2.3), including unsupervised opinion extraction.

**Task Definition.** Given a sentence with  $n$  words,  $\mathbf{x} = [x_1, \dots, x_n]$ , ALSC aims to predict a three-categorical sentiment from  $\{Positive, Neutral, Negative\}$  of an aspect term,  $a$ , where  $a$  is a span of  $\mathbf{x}$ , from  $s^a$  and to  $e^a$ , i.e.,  $\mathbf{a} = x_{s^a:e^a}$ .

### 2.1 Adversarial Anti-bias Aspect Encoder

As suggested by Devlin et al. (2019) and verified by many works, a common practice to tackle pair inputs is feeding the pre-trained Transformer with a concatenation of them. Hence, many ALSC works (Jiang et al., 2020; Hou et al., 2021) feed a concatenation of the sentence  $\mathbf{x}$  and an aspect  $\mathbf{a}$  into the encoder. Nonetheless, considering that every aspect is a word span of the sentence, it is natural to employ a span-level paradigm for aspect representation, which has been proven as effective as the concatenation models (Zhou et al., 2019; Zhao et al., 2020). Also, a span-level model is more efficient as it predicts the sentiment for all aspect terms in a single feed-forward process.

**Span-level Aspect Representation.** We thereby adopt the span-level paradigm to generate aspect representation. Formally, a Transformer is fed with a sentence without additional aspect terms,

$$\mathbf{H} = \text{Transformer-Enc}(\mathbf{x}; \theta^{(ptm)}) \quad (1)$$

where,  $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_n] \in \mathbb{R}^{d \times n}$  denotes contextualized representations corresponding to all the words. Then, given span  $[s^a, e^a]$  of each aspect term  $\mathbf{a}$ , we derive its representation by

$$\mathbf{c}^a = [\mathbf{h}_{s^a}^a; \mathbf{h}_{e^a}^a; \text{Attn-Pool}(\mathbf{H}_{s^a:e^a}; \theta^{(aa)})], \quad (2)$$

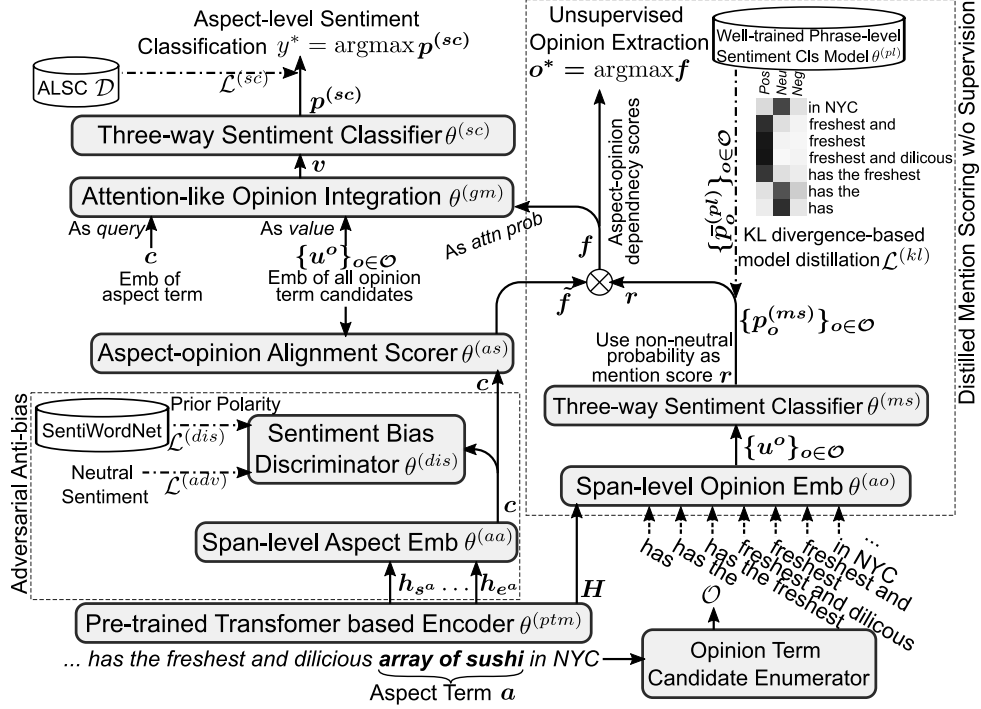


Figure 1: An overview of proposed Span-based Anti-bias aspect Representation Learning (SARL) framework.

where  $[\cdot]$  is vector concatenation,  $\text{Attn-Pool}(\cdot)$  is attention pooling to generate sequence-level embedding with multi-layer perceptron (MLP)-derived weights (Lin et al., 2017; Liu et al., 2016). And  $c^a \in \mathbb{R}^{3d}$  is the resulting span-level aspect representation. In the remainder, we omit the aspect indicator,  $a$ , if no confusion caused. But, as briefed in §1, sentiment bias will occur since the Transformer encoders pre-trained on large-scale text corpora incline to internalize the sentiment polarity of aspect terms. Hence, directly applying a classifier to the aspect representation  $c$  for its sentiment prediction will be affected.

**Adversarial Anti-Bias Module.** A promising way is leveraging adversarial learning to eliminate the sentiment bias of a particular aspect. Open questions still remain about how to define the *sentiment bias* as discriminator’s objective and how to fool the discriminator for our anti-bias purpose. As an answer to the first answer, we resort to external sentiment knowledge, SentiWordNet (Esuli and Sebastiani, 2006), for its prior three-categorical sentiment for common lexicons and phrases. Given an aspect  $a$  from a sentence, we can easily obtain its prior sentiment polarity  $y^{(pr)}$  by querying SentiWordNet. The prior sentiment polarity is the intrinsic attribute of an aspect, which thus can be viewed as its sentiment bias. Hence, we train a sen-

timent bias discriminator towards  $y^{(pr)}$  to exploit bias information underlying the aspect representation. Formally, we present an MLP-based neural classifier upon  $c$  as the discriminator, i.e.,

$$p^{(pr)} = \text{softmax}(\text{MLP}(c^a; \theta^{(dis)})) \in \mathbb{R}^3. \quad (3)$$

Next, training loss of this discriminator is

$$\mathcal{L}_{\theta^{(dis)}}^{(dis)} = - \sum_{\mathcal{D}} \sum_{\mathcal{A}} \log p_{[\hat{y}=y^{(pr)}]}^{(pr)}, \quad (4)$$

where  $\mathcal{D}$  denotes ALSC dataset with sentence-level samples, where  $\mathcal{A}$  denotes all aspects in a sentence, and  $p_{[\hat{y}=y^{(pr)}]}^{(pr)}$  denotes fetching the probability value corresponding to prior sentiment  $y^{(pr)}$ .

On the other side, to eliminate sentiment bias, the span-level aspect encoder aims to fool the discriminator towards generating neutral-sentiment representation. The adversarial loss is written as

$$\mathcal{L}_{\theta^{(ptm)}, \theta^{(aa)}}^{(adv)} = - \sum_{\mathcal{D}} \sum_{\mathcal{A}} \log p_{[\hat{y}=\text{Neutral}]}^{(pr)}. \quad (5)$$

With the adversarial learning between  $\mathcal{L}^{(dis)}$  and  $\mathcal{L}^{(adv)}$ , the aspect representation  $c$  can escape from sentiment bias and become sentiment-agnostic.

## 2.2 Aspect-Opinion Dependency Aligner

As opinion term is also a span of words, it is intuitive to use a span-based alignment model (Lee et al., 2017) for explicit modification from opinion

terms to the aspect. This model is first proposed for co-reference resolution to align an entity mention with its antecedent, but usually requires full supervisions. For example, it has been adapted to extract pair-wise aspect and opinion terms (Zhao et al., 2020) in ABSA, and the supervisions include both opinion terms and aspect-opinion associations but none of them is available in ALSC.

Fortunately, built upon the sentiment-agnostic aspect representation derived from the above encoder, the span-based alignment model is likely to automatically learn the dependency between opinion and aspect spans, which is guided by the aspect’s sentiment label. To this end, we present an aspect-opinion dependency aligner w/o supervisions.

**Opinion Term Candidates.** Similar with Lee et al. (2017) and Zhao et al. (2020), we first enumerate all the possible spans as candidates of opinion terms. Given the sentence  $\mathbf{x} = [x_1, \dots, x_n]$ , an opinion term candidate  $\mathbf{o}$  is a span of words from  $s^o$  and to  $e^o$ , i.e.,  $\mathbf{o} = x_{s^o:e^o}$ . A set of all the opinion term candidates can be written as

$$\mathcal{O} = \{x_{s^o:e^o} | 1 \leq s^o \leq e^o \leq n \wedge e^o - s^o \leq l \wedge (e^o < s^a \vee s^o > e^a)\}, \quad (6)$$

where the last condition means no overlap between an arbitrary oracle aspect term and its opinion candidates, and  $m := |\mathcal{O}|$ . Then, we can easily derive the representation of each candidate  $\mathbf{o} \in \mathcal{O}$  given the contextualized embedding  $\mathbf{H}$  by the encoder in Eq.(1). We adopt the same scheme to generate span-level opinion embedding as in Eq.(2), i.e.,

$$\mathbf{u}^o = [\mathbf{h}_{s^o}; \mathbf{h}_{e^o}; \text{Attn-Pool}(\mathbf{H}_{s^o:e^o}; \theta^{(ao)}),] \quad (7)$$

where the resulting vector  $\mathbf{u}^o \in \mathbb{R}^{3d}$  represents the candidate  $\mathbf{o}$ . Thus, we can get a set of candidate representations,  $\mathbf{U} = \{\mathbf{u}^o\}_{\mathbf{o} \in \mathcal{O}} \in \mathbb{R}^{3d \times m}$ .

Following previous models (Lee et al., 2017; Zhao et al., 2020), two kinds of scores are generated for the dependency between opinions and aspects: (1) a sentiment mention score determines if a span candidate is an opinion mention; and (2) an aspect-opinion alignment score estimates the modification relation in an aspect-opinion pair.

**Distilled Sentiment Mention Scoring.** Most previous span-based alignment models employ an one-dim-out neural module (e.g., MLP with Sigmoid) to determine the confidence of a mention but require full supervision for accurate predictions. Considering such supervision is unavailable

in ALSC, we consequently weaken the supervision from “whether a mention is a gold opinion term” to “whether a mention expresses a sentiment polarity”. The weak supervision can be readily obtained from a well-trained phrase-level sentiment classification model via distilling model (Hinton et al., 2015). Hence, we first employ an MLP-based classifier built upon an opinion term candidate  $\mathbf{u}^o$  to derive a three-categorical sentiment distribution,

$$\mathbf{p}_o^{(ms)} = \text{softmax}(\text{MLP}(\mathbf{u}^o; \theta^{(ms)})) \in \mathbb{R}^3; \quad (8)$$

and a phrase-level sentiment classification model is

$$\bar{\mathbf{p}}_o^{(pl)} = \text{Senti-Model}(\mathbf{o}; \theta^{(pl)}), \quad (9)$$

which is also based on a pre-trained Transformer and trained on a popular phrase-level sentiment analysis dataset, Stanford Sentiment Treebank (Socher et al., 2013). Then, we can define a soft loss of sentiment distillation based on Kullback–Leibler (KL) divergence, i.e.,

$$\mathcal{L}^{(kl)} = \text{KL}(\bar{\mathbf{p}}_o^{(pl)} || \mathbf{p}_o^{(ms)}). \quad (10)$$

Last, we use the non-neutral probability as the confidence of a mention expressing sentiment polarity, i.e., the sentiment mention score,

$$r^o := 1 - \mathbf{p}_o^{(ms)}[\hat{y} = \text{Neutral}], \quad \forall \mathbf{o} \in \mathcal{O}, \quad (11)$$

$$\mathbf{r} = \{r^o\}_{\mathbf{o} \in \mathcal{O}} \in \mathbb{R}^m. \quad (12)$$

**Aspect-Opinion Dependency Modeling.** Besides the mention score, we use an MLP to estimate the alignment between the aspect and opinion candidates. First, we obtain a relationship representation by an interactive concatenation (Reimers and Gurevych, 2019) of their span-level embeddings,

$$\mathbf{q}^o = [\mathbf{c}; \mathbf{u}^o; \mathbf{c} \odot \mathbf{u}^o; \mathbf{z}^o], \quad \forall \mathbf{o} \in \mathcal{O}, \quad (13)$$

where “ $\odot$ ” is Hadamard product and  $\mathbf{z}$  is a learnable relative-position embedding indicating their distance over the syntactic dependency parsing tree of the sentence. Again, please note we omitted the superscript  $a$  in the equation for simplification. Next,  $\mathbf{q}^o$  is passed into an MLP-based scorer to calculate pairwise alignment score, i.e.,

$$\tilde{f}^o = \text{MLP}(\mathbf{q}^o; \theta^{(as)}) \in \mathbb{R}, \quad \forall \mathbf{o} \in \mathcal{O}, \quad (14)$$

$$\tilde{\mathbf{f}} = \{\tilde{f}^o\}_{\mathbf{o} \in \mathcal{O}} \in \mathbb{R}^m. \quad (15)$$

Then, we apply softmax to  $\tilde{\mathbf{f}}$  for normalized alignment scores, which are subsequently weighted by the corresponding sentiment mention scores  $\mathbf{r}$  in Eq.(12) to derive the dependency scores between the aspect and all opinion candidates, i.e.,

$$\mathbf{f} = \text{softmax}(\tilde{\mathbf{f}}) \odot \mathbf{r}. \quad (16)$$

Dataset	Positive		Neutral		Negative	
	Train	Test	Train	Test	Train	Test
Laptop14	994	341	464	169	870	128
Rest14	2164	728	637	196	807	196
Rest15	912	326	36	34	256	182
Rest16	1240	469	69	30	439	117
Twitter	1561	173	3127	346	1560	173

Table 1: Summary statistics of five benchmark datasets.

**“Dummy” Opinion Term.** However, a special scenario has been often ignored when capturing the aspect-opinion dependency or alignment. That is, an aspect term may not correspond to any opinion (statistically,  $> 65\%$  of neutral-labeled aspects in benchmarks w/o opinion). To remedy this, we adopt a concept of “*dummy span*” that indicates no opinion term in the sentence towards an aspect. We can imply aspect-“dummy” dependency score by considering the neutral-sentiment probability of the aligned opinion term candidates, i.e.,

$$f^{(d)} = \delta \sum (\text{softmax}(\tilde{\mathbf{f}}) \odot (\mathbf{1} - \mathbf{r})), \quad (17)$$

where  $\delta$  denotes a re-scaling hyperparameter. So we rewrite the normalized dependency scores as

$$\mathbf{f} \leftarrow [\mathbf{f}; f^{(d)}] \in \mathbb{R}^{m+1}. \quad (18)$$

It is also essential to integrate the aligned opinion term candidates (including the dummy opinion) into the sentiment-agnostic aspect representation  $\mathbf{c}$ . Specifically, we first obtain opinion representation by an attention-like operation, i.e.,

$$\mathbf{u} = [\mathbf{U}, \mathbf{c}] \cdot \mathbf{f}, \quad (19)$$

where “ $\cdot$ ” denotes matrix multiplication and, as following Lee et al. (2017), the aspect representation itself corresponds to the dummy opinion, i.e.,  $[\mathbf{U}, \mathbf{c}] \in \mathbb{R}^{d \times (m+1)}$ . Lastly, we integrate the opinion representation to the aspect one by gating,

$$\mathbf{g} = \text{Sigmoid}(\text{MLP}([\mathbf{c}; \mathbf{u}]; \theta^{(gm)})) \in \mathbb{R}^{3d}, \quad (20)$$

$$\mathbf{v} = \mathbf{g} \cdot \mathbf{c} + (\mathbf{1} - \mathbf{g}) \cdot \mathbf{u}. \quad (21)$$

As a result,  $\mathbf{v}$  stands for opinion-enrich aspect representation and is ready for sentiment classification.

### 2.3 Model Training and Inference

**Aspect-level Sentiment Classification.** On top of  $\mathbf{v}$ , we define a neural classifier for the final three-categorical sentiment prediction as

$$\mathbf{p}^{(sc)} = \text{softmax}(\text{MLP}(\mathbf{v}; \theta^{(sc)})) \in \mathbb{R}^3. \quad (22)$$

Dataset	#AT	#OT	# $\theta$	Dataset	#AT	#OT	# $a^{neu/o}$
Laptop14	444	527	112	Rest14	832	974	137

Table 2: Summary statistics of unsupervised opinion extraction test set. # $a^{neu/o}$  is the number of neutral-sentiment aspects without corresponding opinion terms.

LR ( $\times 10^{-5}$ )	size	Laptop14	Rest14	Rest15	Rest16	Twitter
lr- $\theta^{(ptm)}$	base	1.3	2	2	2	2
	large	1	1	1	1	2
lr-others	base	20	1	20	20	20
	large	20	1	1	20	1

Table 3: Settings of learning rate. “others” includes  $\theta^{(aa)}$ ,  $\theta^{(dis)}$ ,  $\theta^{(ao)}$ ,  $\theta^{(ms)}$ ,  $\theta^{(as)}$ ,  $\theta^{(gm)}$  and  $\theta^{(sc)}$ .

And the training loss of ALSC task is written as

$$\mathcal{L}^{(sc)} = - \sum_{\mathcal{D}} \sum_{\mathcal{A}} \log \mathbf{p}_{[\hat{y}=y^{(sc)}]}^{(sc)}, \quad (23)$$

where  $y^{(sc)}$  denotes the oracle label of an aspect.

**Training and Inference.** Besides the discriminator loss  $\mathcal{L}^{(dis)}$  in Eq.(4), we train the learnable parameters in our proposed SARL model towards a linear combination of the other three losses, i.e.,

$$\mathcal{L}^{(alsc)} = \mathcal{L}^{(sc)} + \beta \mathcal{L}^{(adv)} + \gamma \mathcal{L}^{(kl)}. \quad (24)$$

We also set a hyper-parameter  $\alpha$  to control the proportion of discriminator learning (i.e.,  $\mathcal{L}^{(dis)}$ ) against the ALSC model learning (i.e.,  $\mathcal{L}^{(alsc)}$ ). The inference procedure can be simply written as

$$y^* = \arg \max \mathbf{p}^{(sc)}. \quad (25)$$

**Unsupervised Opinion Extraction.** A well-trained SARL model is equipped with the capability to extract opinion term(s) for an aspect, based on its intermediate variable, i.e.,

$$\mathbf{o}^* = \arg \max_{\{o, \text{dummy}\}} \mathbf{f}, \quad (26)$$

where  $\mathbf{f} \in \mathbb{R}^{m+1}$  from Eq.(16) denotes aspect-opinion dependency scores, including the last dim for dummy opinion, i.e., an aspect w/o any opinion in the sentence. Thereby,  $\mathbf{o}^*$  is the extracted opinion term for a specified aspect term in the sentence, and it is worth mentioning again that the opinion extraction is learned in an unsupervised manner.

## 3 Experiment

**Datasets.** For ALSC task, we evaluate our model on five datasets<sup>1</sup>, whose statistics are listed in Table 1, including (i) Laptop14 (SemEval-2014T4)

<sup>1</sup>The source code and datasets are available at [https://github.com/wangbo9719/SARL\\_ABSA](https://github.com/wangbo9719/SARL_ABSA)

Embedding	Method	Laptop14		Rest14		Rest15		Rest16		Twitter	
		Accu	Ma-F1	Accu	Ma-F1	Accu	Ma-F1	Accu	Ma-F1	Accu	Ma-F1
Static Embedding	ASGCN	75.55	71.01	80.86	72.19	79.89	61.89	88.99	67.48	72.15	70.40
	CDT	77.19	72.99	82.30	74.02	-	-	85.58	69.93	74.66	73.66
	BiGCN	74.59	71.84	81.97	73.48	81.16	64.79	88.96	70.84	74.16	73.35
BERT <sub>base</sub>	DGEDT	79.80	75.60	86.30	80.00	84.00	71.00	91.90	79.00	77.90	75.40
	RGAT	78.21	74.07	86.60	81.35	-	-	-	-	76.15	74.88
	kumaGCN	81.98	78.81	86.43	80.30	80.69	65.99	89.39	73.19	77.89	77.03
RoBERTa <sub>base</sub>	†ASGCN-FT-RoBERTa	83.33	80.32	86.87	80.59	-	-	-	-	76.10	75.07
	†PWCN-FT-RoBERTa	84.01	81.08	87.35	80.85	-	-	-	-	77.02	75.52
	†RGAT-FT-RoBERTa	83.33	79.95	87.52	81.29	-	-	-	-	75.81	74.91
	†MLP	83.78	80.73	87.37	80.96	-	-	-	-	77.17	76.20
	<b>SARL (ours)</b>	85.42	<b>82.97</b>	88.21	82.44	88.19	73.83	94.62	81.92	78.03	76.97
RoBERTa <sub>large</sub>	<b>SARL (ours)</b>	<b>85.74</b>	<b>82.97</b>	<b>90.45</b>	<b>85.34</b>	<b>91.88</b>	<b>78.88</b>	<b>95.76</b>	<b>84.29</b>	<b>78.32</b>	<b>77.32</b>

Table 4: ALSC results on the five datasets. †Numbers are from Dai et al. (2021), and others are from the original papers, i.e., ASGCN (Zhang et al., 2019a), CDT (Sun et al., 2019b), BiGCN (Zhang and Qian, 2020), DGEDT (Tang et al., 2020), RGAT (Wang et al., 2020), kumaGCN (Chen et al., 2020), PWCN-FT-RoBERTa (Zhang et al., 2019b).

(Pontiki et al., 2014) with laptop reviews, (ii) Rest14 (SemEval-2014T4), Rest15 (SemEval-2015T12) and Rest16 (SemEval-2016T5) (Pontiki et al., 2014, 2015, 2016) with restaurant reviews, and (iii) Twitter (Mitchell et al., 2013) with tweets. Following most competitors including the models compared in Table 4, we do not split training set.

To evaluate unsupervised opinion extraction, we employ the test set from Xu et al. (2020) where annotations of opinion terms<sup>2</sup> are from Fan et al. (2019). The statistics are listed in Table 2.

**Training Setups.** We use a mini-batch Stochastic Gradient Descent (SGD) to minimize the loss functions, with Adam optimizer, 10% warm-up, and a linear decay of the learning rate. To initialize the Transformer, we alternate between pre-trained RoBERTa<sub>base</sub> and <sub>large</sub>. We set batch size = 16 and max sequence length = 64 based on experience, and conduct grid searches for the other hyperparameters. Then, we set the max width of candidate span  $l = 15$ , the number of training epochs = 7 for Twitter and 10 for other datasets, the training proportion  $\alpha = 1/3$  for Laptop14 and Twitter, 0.2 for Rest14 and Rest15, 0.1 for Rest16, the loss weight in Eq.(24)  $\beta = 0.05$  for Rest15 and Twitter and 0.1 for the remains,  $\gamma = 1$ , and  $\delta = 1/m$  in Eq.(17). The learning rates are listed in Table 3.

### 3.1 Overall Performance

ALSC results of competitive approaches and our SARL on the five benchmarks are shown in Table 4. Following prior works, we adopt accuracy (Accu) and macro-F1 (Ma-F1) to evaluate the performance, and the results of SARL are the best values from

<sup>2</sup>The sentences in this dataset cannot completely match the standard ALSC dataset so we only use the overlap part.

ten runs. It is observed that our proposed SARL achieves state-of-the-art performance on all these datasets. Compared to static embedding-based methods, the methods with pre-trained Transformer gain better results. In the same embedding genre RoBERTa<sub>base</sub>, SARL outperforms others by an average of 1% on accuracy and 1.3% on macro-F1. Furthermore, the RoBERTa<sub>large</sub>-based SARL derives more significant progress.

### 3.2 Sentiment Bias Elimination

**Bias Statistics.** As in Figure 2 (left), the sentiment bias is common since aspects with bias polarity account for nearly 50% in most datasets.

**The Success of Adversarial Learning.** To explore the effectiveness of our adversarial learning, we first train a model without adversarial, “w/o adv” for short (please refer to below ablation study for more). Then, based on this trained model, we add a sentiment bias detector that has the same architecture as the discriminator to detect the bias existing in aspects representations. Here, we use the ratio of neutral predictions to measure the accuracy of detector since the smaller ratio means more bias de-




Figure 2: Sentiment bias analysis on five datasets. The “Baseline” in right denotes the span-based Transformer baseline. Suppose an aspects set  $\mathcal{S} = \{a | (y^{(pr)} \neq Neutral) \wedge (y^{(pr)} \neq y^{(sc)})\}$  denotes all potential aspects who will be misclassified due to the sentiment bias, the vertical axis in right represents the proportion of  $|\mathcal{Q}|/|\mathcal{S}|$ , where  $\mathcal{Q} = \{a | a \in \mathcal{S} \wedge y^* = y^{(pr)} \neq y^{(sc)}\}$ . And the number on the top of bin is  $|\mathcal{Q}|$ .

Supervision	Method	Laptop14								Rest14								
		@1				@3				@1				@3				
		EM	P	R	F	EM	P	R	F	EM	P	R	F	EM	P	R	F	
✓	SpanMlt	-	-	-	80.6	-	-	-	-	-	-	-	-	-	-	-	-	84.0
	Peng et al. (2020)	-	81.8	84.8	83.2	-	-	-	-	-	76.9	75.3	76.0	-	-	-	-	-
✗	Induced Tree	8.5	29.4	20.6	22.6	-	-	-	-	10.6	30.2	21.8	23.5	-	-	-	-	-
	SARL (ours)	<b>42.3</b>	<b>62.2</b>	<b>68.7</b>	<b>62.8</b>	<b>60.5</b>	<b>76.3</b>	<b>81.4</b>	<b>77.1</b>	<b>49.4</b>	<b>65.9</b>	<b>67.1</b>	<b>64.8</b>	<b>76.0</b>	<b>86.0</b>	<b>88.2</b>	<b>86.1</b>	-
	w/o distillation	16.3	37.5	32.1	31.8	27.5	52.8	50.7	47.8	19.0	43.9	45.3	41.1	34.8	65.3	68.1	62.8	-
	w/o sentiment score	15.0	32.7	29.7	28.6	27.7	54.2	51.9	49.4	30.9	52.5	50.5	48.9	52.1	73.7	74.1	71.2	-

Table 5: Unsupervised opinion extraction results (%) on Laptop14/Rest14. The resulting numbers of SpanMlt (Zhao et al., 2020) and Peng et al. (2020) are from their original papers which the metrics under @1 are the standard metrics. The results of induced tree (Dai et al., 2021) are calculated by regarding the aspects’ sub-nodes in the tree as its extracted opinions.




Figure 3: ALSC results of SARL, and “SARL with neutral reinforce”, i.e., w/  $\mathcal{L}_{\theta^{(dis)}, \theta^{(aa)}, \theta^{(ptm)}}^{(adv)}$  while w/o  $\mathcal{L}^{(dis)}$ .

tected. Compared to 98.75% by full SARL model, the detector’s ratio 69.64% is much smaller, which proves both necessity and success of adversarial learning. Further, as in Figure 3, we train a model w/o adv but with neutral reinforce and there are obvious decreases compared with the full SARL.

**Results of Elimination.** Comparing to a span-level baseline that simply feeds the aspect representation  $c$  in Eq.(2) into an MLP to get the sentiment predictions, ours alleviates the problem as shown in Figure 2 (right). A possible reason is that the aspects in Laptop14 (e.g., *screen* and *size*) merely have a little sentiment in themselves so that the bias labels from SentiWordNet are noisier.

### 3.3 Unsupervised Opinion Extraction

**Explicit Opinion Extraction.** To measure the model’s ability of opinions extraction w/o supervision, we present the following novel metrics including the top-N-based Exact Match (EM@N), Precision (P@N), Recall (R@N) and F1-Score (F@N). Specifically, EM@N denotes the gold opinion term appears in the top-N opinion term candidates. And the Precision@N, Recall@N and F1@N are also employed to describe the maximum char-level overlap. As shown in Table 5, SARL achieves promising performance on unsupervised opinion extrac-

Laptop14		Rest14	
Hits@1	Hits@3	Hits@1	Hits@3
30.36	56.25	45.99	61.31

Table 6: The Hits@N results about dummy opinions.

tion. Without any opinion annotation data, the top-3 based metrics achieve similar or even better performance compared with the models under opinion supervision. Further, comparing with Induced Tree, SARL achieves far better performance.

**Dummy Opinion Extraction.** To measure the dummy opinion ranking about neutral-sentiment aspects without opinion terms, we apply Hits@N that stands for the ratio of such aspects that the dummy opinion is ranked in top-N. As shown in Table 6, the performance is adequate to support the purpose of introducing dummy opinions.

### 3.4 Ablation Study

To explore each module’s contribution, we conduct an extensive ablation study. For ALSC, (1) The results of “model w/o adv” are slightly higher than the “model w/o all”, which indicates that the AT-OT aligner is helpful for classification due to our opinion terms integration; and (2) The “model w/o aligner” achieves sub-optimal results and completely loses the ability of opinion extraction. In addition, the component dropping also severely affects the performance of opinion extraction as in Table 5. The extraction performance decrease of SARL “w/o distillation” is larger than “w/o sentiment score”. A potential reason is that the sentiment scorer performs poorly without any supervision and thus generates incorrect sentiment scores.

Method	Laptop14		Rest14	
	Accu	Ma-F1	Accu	Ma-F1
Full model	85.74	82.97	90.45	85.34
w/o adv	83.70	80.55	89.46	82.16
w/o AT-OT aligner	84.64	81.75	89.28	83.51
w/o all	83.86	80.53	89.02	82.59

Table 7: Ablation results of ALSC. The full model denotes our proposed SARL. “w/o adv” removes anti-bias module in §2.1. “w/o AT-OT aligner” removes the proposed module in §2.2. “w/o all” degrades our SARL model to Transformer-based span-level ALSC (i.e., classification based on Eq.(2)).

Index	Example	“w/o all”	SARL	Top-3 candidates
1	After dinner I heard [music] <sub>neu</sub> playing and discovered that there is a lounge downstairs.	<i>pos</i>	<i>neu</i>	dummy; playing; playing and discovered
2	Desserts include [flan] <sub>neu</sub> and sopaipillas.	<i>pos</i>	<i>neu</i>	dummy; and; include
3	How is this [place] <sub>neg</sub> still open?	<i>pos</i>	<i>neg</i>	open; still open; dummy
4	This place has beautiful [sushi] <sub>pos</sub> , and it’s delicious CHEAP.	<i>pos</i>	<i>pos</i>	delicious; delicious CHEAP; beautiful
5	The wait [staff] <sub>pos</sub> was loud and inconsiderate.	<i>neg</i>	<i>neg</i>	loud; was loud; loud and inconsiderate
6	The sauce is excellent (very fresh) with [dabs of real mozzarella] <sub>neu</sub> .	<i>pos</i>	<i>pos</i>	excellent; excellent (very fresh); fresh
7	15% gratuity automatically added to the [bill] <sub>neg</sub> .	<i>neu</i>	<i>neu</i>	dummy; automatically; 15
8	The [quality of the meat] <sub>neg</sub> was on par with your local grocery store.	<i>pos</i>	<i>neu</i>	dummy; par; was on par

Table 8: Case study for sentiment bias (row 1-3), unsupervised opinion extraction (row 4-5) and error analysis (6-8).

### 3.5 Case Study

**Does SARL eliminate sentiment bias?** As the top-3 rows in Table 8, those aspects with intrinsic positive bias are mis-classified by our span-level baseline but correctly classified by SARL, which verifies the effectiveness of SARL.

**How does SARL obtain interpretability from unsupervised opinion extraction?** As listed in Table 8, SARL can exactly extract the opinion terms of the targeted aspect so provide promising interpretability for sentiment prediction under the unsupervised opinion setting. In addition, for the neutral aspect without explicit opinion terms in a sentence, the dummy opinion always ranks first like the 2<sup>nd</sup> row in Table 8, which explains the reason for neutral prediction.

**Error Analysis.** To analyze the limitation of ALSC models including ours, we investigate all the examples mis-classified by SARL on Rest14, and summarize two main problems: (1) The major (up to 66%) problem is neutral-related mis-classifying because a neutral aspect term is affected by the polarity words associated with other aspects (e.g., 6<sup>rd</sup> row). The other problem (23%) refers to that it is infeasible to determine aspects’ sentiment without commonsense knowledge (e.g., 7<sup>th</sup> row) or additional information (e.g., 8<sup>th</sup> row). In summary, how to accurately classify the aspect without explicit opinion is still an open problem.

## 4 Related Work

**Aspect-Level Sentiment Classification.** ALSC relies heavily on modification relations between aspect term and opinion term in a sentence, so recent progresses mainly fall into modeling the relations by applying graph neural networks to dependency

parsing tree (Sun et al., 2019b; Zhang et al., 2019a; Wang et al., 2020). Despite their effectiveness, they lack interpretability to the sentiment prediction. For further boosted, many methods (Sun et al., 2019a; Hu et al., 2019; Chen et al., 2020; Mao et al., 2021; Dai et al., 2021) introduce the pre-trained Transformer, which also brings a little interpretability due to the highlight modification between aspect and opinion terms derived from potential syntax knowledge (Wu et al., 2020; Dai et al., 2021). However, the derived interpretability is far from human-level. Furthermore, the Transformer tends to internalize terms’ intrinsic sentiment bias, which is harmful to ALSC.

**Aspect-Opinion Alignment.** Fan et al. (2019) first defined the aspect-oriented opinion extraction task in which the aspect terms are given in advance. Later, Zhao et al. (2020) proposed the aspect-opinion co-extraction task. More recently, Peng et al. (2020) proposed the aspect sentiment triplet extraction and further explorations have conducted by Xu et al. (2020) and Mao et al. (2021). These models extract the opinion terms for aspects and thus provide explicit interpretability. However, they require manually labeled opinions data for training, which is much more expensive than three-categorical labeling in ALSC.

## 5 Conclusion

In this work, we propose SARL framework for ALSC. Specifically, we first present an adversarial anti-bias aspect encoder to eliminate sentiment bias in aspects and then propose an aspect-opinion dependency aligner to unsupervisedly extract opinions. The experiments on 5 benchmarks can greatly support our motivations and empirical results show state-of-the-art performance with interpretability.



## Acknowledgement

This work is supported by the National Natural Science Foundation of China (No.61976102, No.U19A2065).

## References

- Chenhua Chen, Zhiyang Teng, and Yue Zhang. 2020. [Inducing target-specific latent structures for aspect sentiment classification](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 5596–5607. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of bert’s attention](#). *CoRR*, abs/1906.04341.
- Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu, and Xipeng Qiu. 2021. [Does syntax matter? A strong baseline for aspect-based sentiment analysis with roberta](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 1816–1829. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Andrea Esuli and Fabrizio Sebastiani. 2006. [SENTIWORDNET: A publicly available lexical resource for opinion mining](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006*, pages 417–422. European Language Resources Association (ELRA).
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 2509–2518. Association for Computational Linguistics.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- Xiaochen Hou, Peng Qi, Guangtao Wang, Rex Ying, Jing Huang, Xiaodong He, and Bowen Zhou. 2021. [Graph ensemble learning over multiple dependency trees for aspect-level sentiment classification](#). *CoRR*, abs/2103.11794.
- Minghao Hu, Yuxing Peng, Zhen Huang, Dongsheng Li, and Yiwei Lv. 2019. [Open-domain targeted sentiment analysis via span-based extraction and classification](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 537–546. Association for Computational Linguistics.
- Po-Sen Huang, Huan Zhang, Ray Jiang, Robert Stanforth, Johannes Welbl, Jack Rae, Vishal Maini, Dani Yogatama, and Pushmeet Kohli. 2020. [Reducing sentiment bias in language models via counterfactual evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 65–83. Association for Computational Linguistics.
- Bin Jiang, Jing Hou, Wanyue Zhou, Chao Yang, Shihan Wang, and Liang Pang. 2020. [Metnet: A mutual enhanced transformation network for aspect-based sentiment analysis](#). In *Proceedings of the 28th International Conference on Computational Linguistics, COLING 2020, Barcelona, Spain (Online), December 8-13, 2020*, pages 162–172. International Committee on Computational Linguistics.
- Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. 2011. [Target-dependent twitter sentiment classification](#). In *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*, pages 151–160. The Association for Computer Linguistics.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. [End-to-end neural coreference resolution](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 188–197. Association for Computational Linguistics.
- Zhouhan Lin, Minwei Feng, Cícero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. [A structured self-attentive sentence embedding](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Yang Liu, Chengjie Sun, Lei Lin, and Xiaolong Wang. 2016. [Learning natural language inference using bidirectional LSTM model and inner-attention](#). *CoRR*, abs/1605.09090.

- Yue Mao, Yi Shen, Chao Yu, and Longjun Cai. 2021. [A joint training dual-mrc framework for aspect based sentiment analysis](#). *CoRR*, abs/2101.00816.
- Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme. 2013. [Open domain targeted sentiment](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1643–1654. ACL.
- Haiyun Peng, Lu Xu, Lidong Bing, Fei Huang, Wei Lu, and Luo Si. 2020. [Knowing what, how and why: A near complete solution for aspect-based sentiment analysis](#). In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 8600–8607. AAAI Press.
- 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 V. Loukachevitch, Evgeniy V. Kotelnikov, Núria Bel, Salud María Jiménez Zafra, and Gülsen Eryigit. 2016. [Semeval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2016, San Diego, CA, USA, June 16-17, 2016*, pages 19–30. The Association for Computer Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. [Semeval-2015 task 12: Aspect based sentiment analysis](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation, SemEval@NAACL-HLT 2015, Denver, Colorado, USA, June 4-5, 2015*, pages 486–495. The Association for Computer Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [Semeval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation, SemEval@COLING 2014, Dublin, Ireland, August 23-24, 2014*, pages 27–35. The Association for Computer Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentencebert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3980–3990. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1631–1642. ACL.
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019a. [Utilizing BERT for aspect-based sentiment analysis via constructing auxiliary sentence](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 380–385. Association for Computational Linguistics.
- Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, and Xudong Liu. 2019b. [Aspect-level sentiment analysis via convolution over dependency tree](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5678–5687. Association for Computational Linguistics.
- Hao Tang, Donghong Ji, Chenliang Li, and Qiji Zhou. 2020. [Dependency graph enhanced dual-transformer structure for aspect-based sentiment classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6578–6588. Association for Computational Linguistics.
- Kai Wang, Weizhou Shen, Yunyi Yang, Xiaojun Quan, and Rui Wang. 2020. [Relational graph attention network for aspect-based sentiment analysis](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3229–3238. Association for Computational Linguistics.
- Zhiyong Wu, Yun Chen, Ben Kao, and Qun Liu. 2020. [Perturbed masking: Parameter-free probing for analyzing and interpreting BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4166–4176. Association for Computational Linguistics.
- Lu Xu, Hao Li, Wei Lu, and Lidong Bing. 2020. [Position-aware tagging for aspect sentiment triplet extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2339–2349. Association for Computational Linguistics.

- Chen Zhang, Qiuchi Li, and Dawei Song. 2019a. [Aspect-based sentiment classification with aspect-specific graph convolutional networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 4567–4577. Association for Computational Linguistics.
- Chen Zhang, Qiuchi Li, and Dawei Song. 2019b. [Syntax-aware aspect-level sentiment classification with proximity-weighted convolution network](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21-25, 2019*, pages 1145–1148. ACM.
- Mi Zhang and Tiejun Qian. 2020. [Convolution over hierarchical syntactic and lexical graphs for aspect level sentiment analysis](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 3540–3549. Association for Computational Linguistics.
- He Zhao, Longtao Huang, Rong Zhang, Quan Lu, and Hui Xue. 2020. [Spanmlt: A span-based multi-task learning framework for pair-wise aspect and opinion terms extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 3239–3248. Association for Computational Linguistics.
- Yan Zhou, Longtao Huang, Tao Guo, Jizhong Han, and Songlin Hu. 2019. [A span-based joint model for opinion target extraction and target sentiment classification](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5485–5491. ijcai.org.