

# A Hierarchical Interactive Network for Joint Span-based Aspect-Sentiment Analysis

Wei Chen<sup>1,2</sup>, Jinglong Du<sup>3</sup>, Zhao Zhang<sup>4</sup>, Fuzhen Zhuang<sup>2,5\*</sup>, Zhongshi He<sup>1\*</sup>

<sup>1</sup>College of Computer Science, Chongqing University, Chongqing, China

<sup>2</sup>Institute of Artificial Intelligence, Beihang University, Beijing 100191, China

<sup>3</sup>College of Medical Informatics, Chongqing Medical University, Chongqing, China

<sup>4</sup>Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

<sup>5</sup>SKLSDE, School of Computer Science, Beihang University, Beijing 100191, China  
{mlg\_cwei,zshe}@cqu.edu.cn, zhuangfuzhen@buaa.edu.cn

## Abstract

Recently, some span-based methods have achieved encouraging performances for joint aspect-sentiment analysis, which first extract aspects (aspect extraction) by detecting aspect boundaries and then classify the span-level sentiments (sentiment classification). However, most existing approaches either sequentially extract task-specific features, leading to insufficient feature interactions, or they encode aspect features and sentiment features in a parallel manner, implying that feature representation in each task is largely independent of each other except for input sharing. Both of them ignore the internal correlations between the aspect extraction and sentiment classification. To solve this problem, we novelly propose a hierarchical interactive network (HI-ASA) to model two-way interactions between two tasks appropriately, where the hierarchical interactions involve two steps: shallow-level interaction and deep-level interaction. First, we utilize cross-stitch mechanism to combine the different task-specific features selectively as the input to ensure proper two-way interactions. Second, the mutual information technique is applied to mutually constrain learning between two tasks in the output layer, thus the aspect input and the sentiment input are capable of encoding features of the other task via back-propagation. Extensive experiments on three real-world datasets demonstrate HI-ASA's superiority over baselines<sup>1</sup>.

## 1 Introduction

Aspect-sentiment analysis (ASA) (Yan et al., 2021; Birjali et al., 2021) aims at extracting all the aspects and their corresponding sentiments within the text simultaneously. And it can be divided into two tasks, i.e., aspect extraction (AE) and sentiment classification (SC). AE is to extract the aspects (Jakob and Gurevych, 2010; Poria et al., 2016;

Karimi et al., 2021) in the sentence, and SC aims to predict the sentiments (Jiang et al., 2011; Lin et al., 2019; Karimi et al., 2021) for the extracted aspects.

In recent years, the span-based models (Hu et al., 2019; Zhou et al., 2019; Lin and Yang, 2020; Lv et al., 2021) are increasingly becoming an alternative for ASA because of their inherent advantages (e.g., they can avoid sentiment inconsistency and huge search space problems in tagging-based models (Luo et al., 2019; Wang et al., 2021)), where the aspects are extracted by directly predicting the boundary distributions, and the sentiment polarities are classified based on the aspect-level words. For example, in the sentence “The screen size is satisfactory but the phone battery capacity is limited.”, the aspect spans are “screen size” and “the phone battery”, and span-level sentiments are positive and negative, respectively. Formally, most of the existing methods can be divided into two types: sequential encoding (Jebbara and Cimiano, 2016; Zhou et al., 2019) and parallel encoding (Hu et al., 2019; Lv et al., 2021) due to different ways in encoding task-specific features. In sequential encoding, the task-specific features are extracted sequentially, i.e., features extracted later have no direct associations with previous ones, which is a unidirectional interaction strategy. In parallel encoding, task-specific features are extracted independently except for using shared input, i.e., the interaction is only present in input sharing. Hence, both encoding methods above fail to model two-way interactions between AE and SC appropriately.

In practice, the learning of the AE and SC may mutually influence each other. On the one hand, sentiment words can be understood better if given the desired aspects. For example, in the two sentences: “The battery has a **large** capacity.” and “The computer case is too **large** to carry.”, we can find that the sentiment word “large” expresses positive sentiment when describing the aspect “hard

\*Corresponding authors.

<sup>1</sup>The source codes are released in <https://github.com/cwei01/HI-ASA>.

drive” while negative sentiment towards “computer”. This actually shows that incorporating the aspect features is conducive to SC. On the other hand, since sentimental expressions are usually close to aspects, the potential sentiment features may provide useful signals for AE. For example, if we observe the word “**spicy**” appears in a restaurant review, it is likely to exist an aspect related to food. Thus, incorporating the sentiment feature also facilitates aspect detection. These two examples illustrate that it is essential to reasonably establish the two-way associations between two tasks.

Inspired by the above analysis, in this work, we propose a novel Hierarchical Interactive model for joint span-based Aspect-Sentiment Analysis (HI-ASA). Specifically, the hierarchical interactions are achieved in two steps: shallow-level interaction and deep-level interaction. The former learns semantic-level interactive features to facilitate information sharing between two tasks in encoding layer. And the latter takes advantage of the technique of mutual information maximization to ensure two-way interaction between the two tasks. Extensive experimental comparisons against state-of-the-art solutions demonstrate the effectiveness of our proposed methods.

## 2 Model

In this section, we formally define the problem of joint span-based aspect-sentiment analysis. Giving a training set  $\{(\mathbf{x}_i, \{(s_i^j, e_i^j, p_i^j)\}_{j=1}^{l_i})\}_{i=1}^N$ , in each sample  $(\mathbf{x}_i, \{(s_i^j, e_i^j, p_i^j)\}_{j=1}^{l_i})$ ,  $\mathbf{x}_i = \{x_{i,1}, x_{i,2}, \dots, x_{i,n_i}\}$  is a  $n_i$ -length sentence, and  $(s_i^j, e_i^j, p_i^j)$  identifies an aspect with  $s_i^j$  and  $e_i^j$  as the start and end boundaries, and  $p_i^j$  as sentiment polarity.  $l_i$  is the number of aspects in the sentence.

The overall architecture is illustrated in Figure 1, which mainly consists of two parts: AE and SC. We model the hierarchical interactions between two tasks appropriately to enhance the correlations. For the input section, firstly, we leverage BERT (Devlin et al., 2018) to extract the semantic information for input sentence. The output through the transformer layers is  $\mathbf{B} = \{b_1, b_2, \dots, b_n\} \in \mathbb{R}^{n \times d}$ , where  $d$  is the BERT’s embedding size,  $n$  is the length of sentence. Then, we stack GRUs upon  $\mathbf{B}$  for different tasks, here, we define the outputs of AE and SC as  $\mathbf{H}_a^0 \in \mathbb{R}^{n \times \hat{d}}$  and  $\mathbf{H}_s^0 \in \mathbb{R}^{n \times \hat{d}}$ , respectively, where  $\hat{d}$  is the dimension size of GRU.

### 2.1 Shallow-Level Interaction

In previous models (Hu et al., 2019; Lin and Yang, 2020; Lv et al., 2021), the feature representations are extracted independently except for using shared input. In other words, two tasks have no associations with each other, which is not in line with human intuition. Intuitively, the features of SC are not only derived from its own features of previous layers, but also come from the features of AE, and vice versa. Therefore, we design a shallow-level interaction strategy inspired by the idea of cross-stitch mechanism (Misra et al., 2016). Its core is to selectively combine the different features to reasonably model two-way interactions. Specifically, the calculation is as follows:

$$\begin{bmatrix} \mathbf{H}_a^1 \\ \mathbf{H}_s^1 \end{bmatrix} = \begin{bmatrix} \gamma_{aa} & \alpha_{sa} \\ \alpha_{as} & \gamma_{ss} \end{bmatrix} \begin{bmatrix} \mathbf{H}_a^0 \\ \mathbf{H}_s^0 \end{bmatrix} \quad (1)$$

where  $\gamma_{aa}$ ,  $\gamma_{ss}$  are the task-specific parameters, and  $\alpha_{sa}$ ,  $\alpha_{as}$  are task-sharing parameters.  $\mathbf{H}_a^1$  and  $\mathbf{H}_s^1$  are the output features of AE and SC in the last encoding layer, respectively. We can observe that the cross-stitch unit is beneficial for the interaction between AE and SC features. To speed up the training process, we define the following constraints, (1):  $\gamma_{aa} = \gamma_{ss}$ ,  $\alpha_{sa} = \alpha_{as}$  (2):  $\gamma_{aa} + \alpha_{sa} = 1$ . Thus, the Eq.1 is simplified to:

$$\begin{aligned} \mathbf{H}_a^1 &= \alpha \cdot \mathbf{H}_s^0 + (1 - \alpha) \cdot \mathbf{H}_a^0 \\ \mathbf{H}_s^1 &= \alpha \cdot \mathbf{H}_a^0 + (1 - \alpha) \cdot \mathbf{H}_s^0 \end{aligned} \quad (2)$$

where  $\alpha$  controls information transferred from the other task.

### 2.2 Deep-Level Interaction

In the last section, we model the two-way interactions from the encoding layer. Although some semantic-level associations can be captured, we believe that task-level associations cannot be adequately modeled. To better understand our proposed method, we give an intuitive explanation in Figure 2. We observe the span-based aspect extraction models put higher scores on the entity words (e.g., “Windows”, “7”). But with aspect-level sentiment learning, higher attention weights are not only put on some polarity words (e.g., “love”), but also on the entity words. Therefore, it would be interesting if we perform aspect-sentiment mutual learning on both tasks. Thus, we propose the deep-level interactive strategy based on mutual information maximization (Kong et al., 2019), where the task-level information is shared by the two tasks in the output layer. We will explain




Figure 1: The architecture of our HI-ASA.

	I love Windows 7 .					I love Windows 7 .				
Aspect Scores:	0.2	0.3	0.7	0.8	0.1	0.2	0.5	0.7	0.9	0.1
Attention Scores:	0.1	0.9	0.4	0.5	0.1	0.1	1.0	0.5	0.7	0.1

Figure 2: The example of our mutual information maximization strategy in sentence “I love Windows 7.”

in detail how to obtain the aspect scores (i.e., the more a word acts like an aspect word, the higher its score) and attention scores for these two tasks.

**Aspect Extraction ( $\mathcal{E}$ ).** We extract the aspects by predicting the boundaries, i.e., start position and end position. Specifically, we employ a linear classifier to predict the start scores  $g_s = \text{sigmoid}(\mathbf{V}_s \mathbf{H}_a^1)$ , i.e.,  $\mathbf{V}_s$  is a trainable vector, and end scores  $g_e$  can be derived in a similar manner. During training phase, we define the learning object:  $\mathcal{J}_{ae} = -\sum_{i=1}^n \mathbf{p}_s^i \log(\mathbf{g}_s^i) - \sum_{i=1}^n \mathbf{p}_e^i \log(\mathbf{g}_e^i)$ , where  $\mathbf{p}_s^i$  and  $\mathbf{p}_e^i$  are the ground truths of the boundaries.  $\mathbf{g}_s^i$  and  $\mathbf{g}_e^i$  are the predicted distributions.

For the aspect score  $\mathcal{E}$ , we cannot use the boundary distributions of aspect (i.e.,  $g_s$  and  $g_e$ ) directly due to their different meanings. In fact, the span of aspect is relatively short, we can use the average boundary distributions to approximate the aspect scores, so we have:  $\mathcal{E} = \frac{\text{pooling}(g_s) + \text{pooling}(g_e)}{2}$ , where  $\text{pooling}(\cdot)$  is the mean-pooling function, in our work, our set window size as  $1 \times 3$ . Therefore, suppose the start score  $g_s$  is  $[0.2, 0.3, 0.7, 0.8, 0.1]$ , the first item of  $\text{pooling}(g_s)$  is  $\frac{0+0.2+0.3}{3} = 0.17$ , and the  $\text{pooling}(g_s)$  is  $[0.17, 0.4, 0.63, 0.57, 0.33]$ , such strategy can guarantee that the aspect distribution approximates the boundary distributions as much as possible.

**Sentiment Classification ( $\mathcal{A}$ ).** In aspect-level

sentiment classification, the attention weights of words produced during the training towards the specific aspects could reveal word-polarity information. Thus, we can leverage the attention distribution as the sentiment features to help detect the boundaries of aspects. To this end, we choose the over-and-over attention model (Huang et al., 2018) to implement the sentiment classification. Here, we directly define the attention distribution as  $\mathcal{A}$ , (i.e.,  $\mathcal{A}$  is attention scores in Figure 2). During training phase, the optimization goal for SC is:  $\mathcal{J}_{sc} = -\sum_{i=1}^N \mathbf{y}_i^T \log \hat{\mathbf{y}}_i$ , where  $\hat{\mathbf{y}}_i$  is the predicted sentiment distribution, and  $\mathbf{y}_i^T$  is the corresponding ground truth,  $N$  is the number of samples.

**Mutual Information Maximization.** After getting the aspect distribution  $\mathcal{E}$  and sentiment attention distribution  $\mathcal{A}$ . As discussed above, it should be interesting if we perform aspect-sentiment mutual learning on both tasks. In Figure 2, we can find that the aspect scores and attention weights are both improved after applying the mutual learning technique. We need to maximize the similarity between two distributions during the training. An intuitive idea is to use Kullback Leibler (KL) divergence to measure the distance between distributions, considering the correlations are bidirectional, we define the following similarity measurement metrics using Jensen Shannon (JS) divergence:

$$JS(\mathcal{E} \parallel \mathcal{A}) = \frac{1}{2} KL \left( \mathcal{E} \parallel \frac{\mathcal{E} + \mathcal{A}}{2} \right) + \frac{1}{2} KL \left( \mathcal{A} \parallel \frac{\mathcal{E} + \mathcal{A}}{2} \right) \quad (3)$$

where the JS divergence is symmetric, thus we can leverage it to constrain two tasks to learn from each other, resulting in balanced interaction between AE

Dataset	#Sentences	#Aspects	#+	#-	#0
<i>Restaurant</i>	3900	6603	4134	1538	931
<i>Laptop</i>	1869	2936	1326	900	620
<i>Tweets</i>	2350	3243	703	274	2266

Table 1: Statistics of the datasets. “+/-/0” refer to the positive, negative, and neutral sentiment classes.

Span-based Models	<i>Laptop</i>	<i>Res</i>	<i>Tweets</i>
Zhou (Zhou et al., 2019)	59.76	71.98	51.44
Hu-pipeline (Hu et al., 2019)	68.06	74.92	57.69
Hu-joint (Hu et al., 2019)	64.59	72.47	54.55
Hu-collapsed(Hu et al., 2019)	48.66	57.85	48.11
SPRM (Lin and Yang, 2020)	68.72	79.17	59.45
S-AESC (Lv et al., 2021)	65.88	74.18	54.73
HI-ASA	<b>70.39</b>	<b>79.90</b>	<b>60.36</b>
- w/o SI	68.05	79.10	60.04
- w/o DI	67.45	79.50	60.13

Table 2: The performance (F1-score) comparisons with different methods. Baseline results are retrieved from published papers.

and SC. Whereas in sequential and parallel encoding, sentiment features have no direct impact on the information of aspect features. Finally, considering the mutual information maximization between two tasks, HI-ASA is optimized by combining three sections:

$$\arg \min_{\theta} (\mathcal{J}_{ae} + \mathcal{J}_{sc} + \beta \cdot JS(\mathcal{E}||\mathcal{A})) \quad (4)$$

where  $\theta$  is the set of model’s parameters,  $\mathcal{J}_{ae}$  and  $\mathcal{J}_{sc}$  are the optimization goals of AE and SC respectively.  $\beta$  is a balanced parameter. During the test phase, HI-ASA outputs two parts, one is the scores of boundaries, based on which we leverage the heuristic algorithm (Lin and Yang, 2020) to extract all the aspects; the other one is corresponding sentiment polarities of extracted aspects.

## 3 Experiments

### 3.1 Datasets and Settings

**Datasets.** In our experiments, we use three public datasets, including *Laptop* (Pontiki et al., 2014), *Restaurant* (Pontiki et al., 2014, 2015, 2016) and *Tweets* (Mitchell et al., 2013), which have been widely used in previous works (Lv et al., 2021). The statistic details of experimental dataset are refer to Table 1. For each sentence in datasets, the gold span boundaries and sentiment polarity labels are available. Specifically, *Restaurant* is the union set of the restaurant domain from SemEval2014, SemEval2015 and SemEval2016. *Laptop* contains costumer reviews in the electronic product domain, which is collected from SemEval

Challenge. *Tweets* is composed of twitter posts from different users.

**Settings.** In the experiments, the commonly used metric F1-score (F1) is selected to evaluate for aspect-sentiment analysis, accuracy is applied to sentiment classification. A correct predicted aspect only when it matches the gold aspect and the corresponding polarity. We utilize the BERT-Large model as the backbone network, where the number of transformers is 24 and the hidden size is 784. In addition, we use Adam optimizer with a learning rate of 3e-5, the batch size is 32 and the dropout probability of 0.1 is used.

### 3.2 Main Results

The comparisons between HI-ASA and the baselines are presented in Table 2. Specifically, “- w/o SI” means we remove the shallow-interaction layer, “- w/o DI” denotes removing the deep-interaction layer. We can observe: **(1)** Overall, our proposed HI-ASA consistently achieves the best F1 scores across all the baselines. More specifically, compared to the state-of-the-art approach SPRM, HI-ASA improves the performance by about 1.67%, 0.63%, and 0.91% on three datasets, respectively. These observations indicate the carefully designed HI-ASA is capable of achieving better performances. The reason can be concluded two folds: one is that we selectively combine the task-specific features of each task to reasonably model two-way interactions in the encoding layer; the second benefit is that we model a balanced task-level interactions. Under this framework, aspect and sentiment associations are captured appropriately in a hierarchical manner. **(2)** Besides, we investigate the effectiveness of each single module, i.e., “- w/o SI” and “- w/o DI”. We can see that when a certain module is removed, the performance of our model decreases, which indicates the indispensability of each module.

### 3.3 Ablation Study

In this section, we go deeper into HI-ASA and analyse the results on both tasks. The results are shown in Table 3. Generally, we observe that HI-ASA outperforms baseline competitors on both tasks, which indicates the effectiveness of the two-way interaction on both aspect extraction and sentiment classification. For the AE task, HI-ASA can enhance the performance of most of the baselines. On the SC task, HI-ASA outperforms SPRM by 3.98%, 2.83%, and 5.16% on three datasets, respec-



Aspect Extraction	<i>Laptop</i>	<i>Res</i>	<i>Tweets</i>
Hu (Hu et al., 2019)	83.35	82.38	75.28
SPRM (Lin and Yang, 2020)	84.72	86.71	69.85
S-AESC (Lv et al., 2021)	85.19	84.20	76.04
HI-ASA	<b>86.30</b>	86.93	<b>76.36</b>
- w/o SI	84.50	<b>87.33</b>	75.81
- w/o DI	85.10	87.09	75.99
Sentiment Classification	<i>Laptop</i>	<i>Res</i>	<i>Tweets</i>
Hu (Hu et al., 2019)	81.39	89.95	75.16
SPRM (Lin and Yang, 2020)	81.50	90.35	78.34
HI-ASA	<b>85.02</b>	<b>93.18</b>	<b>83.50</b>
- w/o SI	84.07	92.53	82.90
- w/o DI	83.75	92.62	83.49

Table 3: The top part is the performance (F1-score) comparisons with different methods on aspect extraction. And the bottom part is the performance (Accuracy) comparisons with different methods on sentiment classification.

Dataset/ $\alpha$	0	0.1	0.2	0.3	0.4	0.5
<i>Laptop</i>	68.05	<b>70.39</b>	67.62	68.63	68.58	70.36
<i>Restaurant</i>	79.14	78.93	79.73	79.45	<b>79.90</b>	78.48
<i>Tweets</i>	60.04	<b>60.36</b>	59.70	59.89	58.79	58.19

Table 4: The performance (F1-score) of different sharing ratios ( $\alpha$ ) on three datasets.

tively, which implies that it is more able to boost the performance of SC compared to AE.

### 3.4 Parameter Analysis $\alpha$

In HI-ASA, the parameter  $\alpha$  controls information transfer between two tasks. For a special case, if  $\alpha = 0$ , our approach will degenerate to the parallel encoding. We tune  $\alpha$  in the range of [0.0,0.1,0.2,0.3,0.4,0.5] and the results are presented in Table 4. It is worth noting that we set the value of  $\alpha$  below 0.5. This setting is inspired by the intuition that the features for each task should come more from the task itself, rather than the other task. We can notice the best results are achieved when  $\alpha$  is 0.1, 0.4, and 0.1 on three datasets, respectively, rather than 0.0. This actually shows our interaction is successful.

Furthermore, we find an interesting phenomenon is that  $\alpha$  is small (i.e., 0.1) on *Laptop* and *Tweets*, and large on *Restaurant* when performing the best performance. We conjecture the reason lies in that *Restaurant* has more samples than the other two datasets, requiring more knowledge interactions for a better learning process.

## 4 Related Work

Aspect-sentiment Analysis (ASA) (Yan et al., 2021; Birjali et al., 2021) is an essential task in sentiment analysis and can be separated into two tasks,

i.e., aspect extraction (AE) and sentiment classification (SC). Over the past years, some span-based methods have achieved promising results for ASA, which first extract aspects by detecting aspect boundaries (AE) and then predict the span-level sentiments (SC). AE has been widely studied by traditional machine (Jakob and Gurevych, 2010) and deep learning algorithms (Karimi et al., 2021). However, the absence of sentiment information may result in redundant and noisy detection. SC is to predict the sentiment expressed on some specific aspects in a sentence, which has been studied extensively in NLP community (Tang et al., 2015; Chen et al., 2019, 2020; Karimi et al., 2021). However, these aspects must be annotated before the AE task. (Hu et al., 2019) proposed an extract-then-classify framework, which extracted aspects with a heuristic decoding algorithm and then correspondingly classified the span-level sentiments. (Lin and Yang, 2020) designed share-private representation for each task to capture the correlations between two tasks in encoding layer. Besides, (Zhou et al., 2019) introduced a joint model based on span-aware attention mechanism to predict the sentiment polarity. Although achieved improved performances, these works fail to model two-way interactions between AE and SC appropriately.

## 5 Conclusion

In this paper, we proposed a novel aspect-sentiment analysis model named HI-ASA. The proposed model is equipped with a hierarchical interactive network to facilitate information sharing between the aspect extraction task and the sentiment classification task. The experimental results on three benchmark datasets demonstrated HI-ASA’s effectiveness and generality.

## Acknowledgments

We are grateful to all the anonymous reviewers for their helpful advice on various aspects of this work. The research work is supported by the National Key Research and Development Program of China under Grant No. 2021ZD0113604, the National Natural Science Foundation of China under Grant No. 62176014, Graduate Research and Innovation Foundation of Chongqing, China under Grant No.CYC21072.

## References

- Marouane Birjali, Mohammed Kasri, and Abderrahim Beni-Hssane. 2021. A comprehensive survey on sentiment analysis: Approaches, challenges and trends. *Knowledge-Based Systems*, 226:107134.
- Wei Chen, Wenxin Yu, Gang He, and Ning Jiang. 2020. Coarse-to-fine attention network via opinion approximate representation for aspect-level sentiment classification. In *International Conference on Neural Information Processing*, pages 704–715. Springer.
- Wei Chen, Wenxin Yu, Zhiqiang Zhang, Yunye Zhang, Kepeng Xu, Fengwei Zhang, Yibo Fan, Gang He, and Zhuo Yang. 2019. Target-based attention model for aspect-level sentiment analysis. In *International Conference on Neural Information Processing*, pages 259–269. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- 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 Annual Meeting of the Association for Computational Linguistics*, pages 537–546.
- Binxuan Huang, Yanglan Ou, and Kathleen M Carley. 2018. Aspect level sentiment classification with attention-over-attention neural networks. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pages 197–206.
- Niklas Jakob and Iryna Gurevych. 2010. Extracting opinion targets in a single and cross-domain setting with conditional random fields. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, pages 1035–1045.
- Soufian Jebbara and Philipp Cimiano. 2016. Aspect-based sentiment analysis using a two-step neural network architecture. In *Semantic Web Evaluation Challenge*, pages 153–167. Springer.
- Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. 2011. Target-dependent twitter sentiment classification. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 151–160.
- Akbar Karimi, Leonardo Rossi, and Andrea Prati. 2021. Adversarial training for aspect-based sentiment analysis with bert. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 8797–8803. IEEE.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2019. A mutual information maximization perspective of language representation learning. In *International Conference on Learning Representations*.
- Peiqin Lin and Meng Yang. 2020. A shared-private representation model with coarse-to-fine extraction for target sentiment analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 4280–4289.
- Peiqin Lin, Meng Yang, and Jianhuang Lai. 2019. Deep mask memory network with semantic dependency and context moment for aspect level sentiment classification. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 5088–5094.
- Huaishao Luo, Tianrui Li, Bing Liu, and Junbo Zhang. 2019. Doer: Dual cross-shared rnn for aspect term-polarity co-extraction. *arXiv preprint arXiv:1906.01794*.
- Yanxia Lv, Fangna Wei, Ying Zheng, Cong Wang, Cong Wan, and Cuirong Wang. 2021. A span-based model for aspect terms extraction and aspect sentiment classification. *Neural Computing and Applications*, 33(8):3769–3779.
- Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. 2016. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3994–4003.
- 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*, pages 1643–1654.
- Maria Pontiki, Dimitrios 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 2014)*, pages 27–35.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, AL-Smadi Mohammad, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, et al. 2016. Semeval-2016 task 5: Aspect based sentiment analysis. In *Proceedings of the 10th international workshop on semantic evaluation (SemEval-2016)*, pages 19–30.
- 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 2015)*, pages 486–495.
- Soujanya Poria, Erik Cambria, and Alexander Gelbukh. 2016. Aspect extraction for opinion mining with a deep convolutional neural network. *Knowledge-Based Systems*, 108:42–49.

- Duyu Tang, Bing Qin, Xiaocheng Feng, and Ting Liu. 2015. Effective lstms for target-dependent sentiment classification. *arXiv preprint arXiv:1512.01100*.
- Xinyi Wang, Guangluan Xu, Zequn Zhang, Li Jin, and Xian Sun. 2021. End-to-end aspect-based sentiment analysis with hierarchical multi-task learning. *Neurocomputing*, 455:178–188.
- Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, and Zheng Zhang. 2021. A unified generative framework for aspect-based sentiment analysis. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language*, pages 2416–2429.
- 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 28th International Joint Conference on Artificial Intelligence*, pages 5485–5491.