

A Span-level Bidirectional Network for Aspect Sentiment Triplet Extraction

Yuqi Chen, Keming Chen*, Xian Sun and Zequn Zhang

Aerospace Information Research Institute

Key Laboratory of Network Information System Technology(NIST)

School of Electronic, Electrical and Communication Engineering

University of Chinese Academy of Sciences

chenyuqi19@mails.ucas.ac.cn, ckmdejob@hotmail.com, {sunxian,zqzhang1}@mail.ie.ac.cn

Abstract

Aspect Sentiment Triplet Extraction (ASTE) is a new fine-grained sentiment analysis task that aims to extract triplets of aspect terms, sentiments, and opinion terms from review sentences. Recently, span-level models achieve gratifying results on ASTE task by taking advantage of the predictions of all possible spans. Since all possible spans significantly increases the number of potential aspect and opinion candidates, it is crucial and challenging to efficiently extract the triplet elements among them. In this paper, we present a span-level bidirectional network which utilizes all possible spans as input and extracts triplets from spans bidirectionally. Specifically, we devise both the aspect decoder and opinion decoder to decode the span representations and extract triples from aspect-to-opinion and opinion-to-aspect directions. With these two decoders complementing with each other, the whole network can extract triplets from spans more comprehensively. Moreover, considering that mutual exclusion cannot be guaranteed between the spans, we design a similar span separation loss to facilitate the downstream task of distinguishing the correct span by expanding the KL divergence of similar spans during the training process; in the inference process, we adopt an inference strategy to remove conflicting triplets from the results base on their confidence scores. Experimental results show that our framework not only significantly outperforms state-of-the-art methods, but achieves better performance in predicting triplets with multi-token entities and extracting triplets in sentences contain multi-triplets¹.

1 Introduction

Aspect-based sentiment analysis (ABSA) is an important field in natural language processing (NLP).

*Corresponding author

¹We release our code at <https://github.com/chen1310054465/SBN>

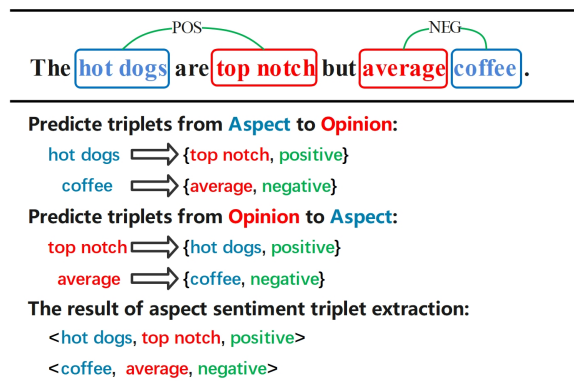


Figure 1: An example of ABSA subtasks. The spans highlighted in blue are aspect terms. The spans in red are opinion terms. Sentiments are marked with green.

The ABSA task contains various fundamental subtasks, such as aspect term extraction (ATE), opinion term extraction (OTE), and aspect-level sentiment classification (ASC). Recent studies focus on solving these tasks individually or doing a combination of two subtasks, such as aspect term polarity co-extraction (APCE), aspect opinion co-extraction (AOCE), and aspect-opinion pair extraction (AOPE). However, none of these subtasks aims to extract the aspect terms (AT) with their corresponding opinion terms (OT) and sentiment polarity (SP) simultaneously. To tackle this problem, (Peng et al., 2020) propose the aspect sentiment triplet extraction (ASTE) task which aims to extract (AT, OT, SP) triplets such as (hot dogs, top notch, positive) and (coffee, average, negative) in the example of Figure 1.

To solve the ASTE task, recent works (Peng et al., 2020; Wu et al., 2020; Mao et al., 2021) use sequential token-level methods and formulate this task as a sequence tagging problem. Although these works achieve competitive results, their token-level models suffer from cascading errors due to sequential decoding. Therefore, (Xu et al., 2021) propose a span-level model to capture

the span-to-span interactions among ATs and OTs by enumerating all possible spans as input. Despite the exciting results their work has yielded, several challenges remain with the existing span-level model. **First**, since both aspect terms and opinion terms can trigger triplets, it is a challenge to identify triplets bidirectionally. **Second**, unlike token-level methods, span-level input cannot guarantee mutual exclusivity among the spans, so the similar spans (spans that have shared tokens) such as *hot dogs*, *dogs*, and *the hot dogs*, may cause confusion in downstream tasks. Thus, it is challenging for span-level models to effectively distinguish these similar span. **Third**, the existence of similar spans enables span-level models to generate conflicting triples in the results, such as (*hot dogs*, *top notch*, *positive*), (*hot dogs,are top notch*, *positive*), and (*the hot dogs*, *top notch*, *positive*). How to properly extract non-conflicting triplets is also challenging.

To address these challenges, we propose a span-level bidirectional network for ASTE task. Unlike prior span-level works (Xu et al., 2021), our network decodes all possible span representations from both aspect-to-opinion and opinion-to-aspect directions through the cooperation of the aspect decoder and opinion decoder. In the aspect-to-opinion direction, the aspect decoder aims to extract ATs such as {*hot dogs*, *coffee*}, and the opinion decoder aims to extract OTs such as {*top notch*} for each specific AT like {*hot dogs*}. Analogously, in the opinion-to-aspect direction, the opinion decoder and aspect decoder are utilized to extract OTs and their corresponding ATs, respectively. Furthermore, we design the similar span separation loss to direct the model deliberately distinguishing similar span representations during the training process; and an inference strategy employed in the prediction process is also proposed for eliminating the conflicting triplets in the extraction results. To verify the effectiveness of our framework, we conduct a series of experiments based on four benchmark datasets. The experimental results show our framework substantially outperforms the existing methods. In summary, our contributions are as follows:

- We design a span-level bidirectional network to extract triplets in both aspect-to-opinion and opinion-to-aspect directions in a span-level model. By this design, our network can identify triplets more comprehensively.
- We propose the similar span separation loss to separate the representations of spans that

contain shared tokens. Based on these differentiated span representations, downstream models can discriminate the span representation more precisely.

- We design an inference strategy to eliminate the potential conflicting triplets due to the lack of mutual exclusivity among spans.

2 Related Work

Aspect based sentiment analysis (ABSA) is a fine-grained sentiment analysis task that consists of various subtasks, including aspect term extraction (ATE) (Wang et al., 2016; Li and Lam, 2017; Xu et al., 2018; Li et al., 2018; Ma et al., 2019), opinion term extraction (OTE) (Poria et al., 2016; Fan et al., 2019; Wu et al., 2020), aspect-level sentiment classification (ASC) (Dong et al., 2014; Tang et al., 2016; He et al., 2018; Li et al., 2019b). Since these subtasks are solved individually, recent studies attempted to couple two subtasks as a compound task, such as aspect term polarity co-extraction (APCE) (Li and Lu, 2017; He et al., 2019; Li et al., 2019a), aspect and opinion co-extraction (Qiu et al., 2011; Liu et al., 2013; Yu et al., 2019), aspect category and sentiment classification (Hu et al., 2019), and aspect-opinion pair extraction (AOPE) (Chen et al., 2020; Zhao et al., 2020; Gao et al., 2021), and aspect-opinion pair extraction (AOPE) (Gao et al., 2021; Zhao et al., 2020; Wu et al., 2021). Although many works have achieved great progress on these tasks, none of these tasks aims to identify the aspect terms as well as their corresponding opinion term and sentiment polarity.

To tackle this issue, (Peng et al., 2020) proposed the aspect sentiment triplet extraction (ASTE) task, which aimed to extract aspect terms, the sentiments of the aspect terms, and the opinion terms causing the sentiments. Some methods (Xu et al., 2020; Wu et al., 2020) designed a unified tagging scheme to solve this task. Some others (Chen et al., 2021; Mao et al., 2021) formulated this task as a multi-turn machine reading comprehension task and solve it with machine reading comprehension frameworks. Recently, (Xu et al., 2021) had propose a span-level model to extract ATs and OTs first and then predict the sentiment relation for each (AT, OT) pairs.

3 Methodology

As shown in Figure 2, our network consists of four parts: span generation, similar span separation loss,

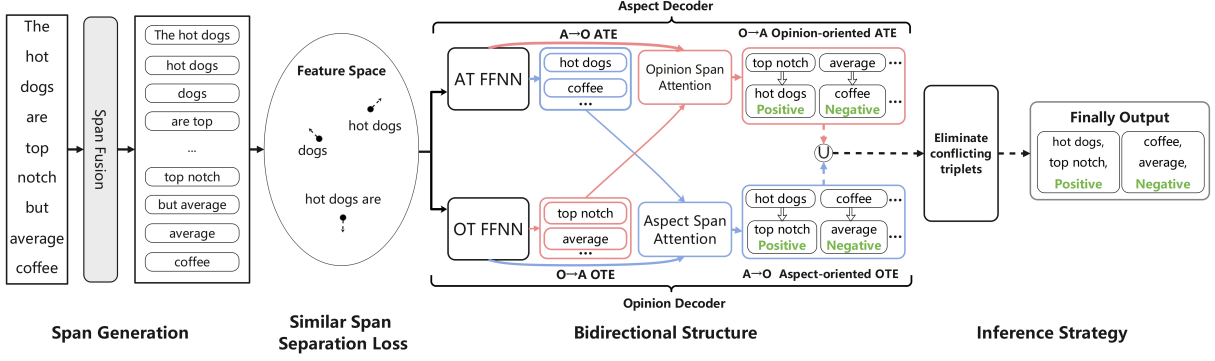


Figure 2: The overall architecture of our span-level bidirectional network. The blue arrows and modules as well as red arrows and modules indicate the extraction of aspect-to-opinion direction and the opinion-to-aspect direction, respectively. The process shown in the dotted line only proceeds in the inference.

bidirectional structure, and the inference strategy. In the following subsections, we first give the definition of ASTE tasks and then detail our network structure.

3.1 Task Definition

For a sentence $S = \{w_1, w_2, \dots, w_n\}$ consisting n words, the goal of the ASTE task is to extract a set of aspect sentiment triplets $\mathcal{T} = \{(a, o, c)_k\}_{k=1}^{|\mathcal{T}|}$ from the given sentence S , where (a, o, c) refers to (aspect term, opinion term, sentiment polarity) and $c \in \{Positive, Neutral, Negative\}$.

3.2 Span Generation

Given a sentence S with n tokens, there are m possible spans in total. Each span $s_i = \{w_{start(i)}, \dots, w_{end(i)}\}$ is defined by all the tokens from $start(i)$ to $end(i)$ inclusive, and the maximum length of span s_i is l_s :

$$1 \leq start(i) \leq end(i) \leq n \quad (1)$$

$$end(i) - start(i) \leq l_s \quad (2)$$

To obtain span representations, we need to get the token-level representations first. In this paper, we utilize BERT (Devlin et al., 2018) as a sentence encoder to obtain token-level contextualized representations $\{\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n\}$ of the given sentence S . Then, the token-level representations are combined by max pooling. Note that various methods can be applied to generate the representations for spans, the effectiveness of these span generation methods will be investigated in the ablation study in Appendix. We define the representation of span s_i as:

$$\mathbf{g}_i = \text{Max}(\mathbf{h}_{start(i)}, \mathbf{h}_{start+1(i)}, \dots, \mathbf{h}_{end(i)}) \quad (3)$$

where Max represents max pooling.

3.3 Similar Span Separation Loss

After generating the representation of span, most previous models directly use the span representations for downstream tasks. However, enumerating all possible spans in a sentence inevitably generates lots of spans that have same tokens with each other, and the model may suffer from the limitations in processing these similar spans due to their adjacent distribution. To separate the spans with similar distributions, we propose a similar span separation loss based on KL divergence to separate similar spans, as shown in Figure 2. The similar span separation loss is defined as:

$$KL(\mathbf{g}_i || G_i) = \sum_j^{G_i} \text{softmax}(\mathbf{g}_i) \log \frac{\text{softmax}(\mathbf{g}_i)}{\text{softmax}(\mathbf{g}_j)} \quad (4)$$

$$KL(G_i || \mathbf{g}_i) = \sum_j^{G_i} \text{softmax}(\mathbf{g}_j) \log \frac{\text{softmax}(\mathbf{g}_j)}{\text{softmax}(\mathbf{g}_i)} \quad (5)$$

$$\mathcal{J}_{KL} = \sum_i^m \log \left(1 + \frac{2}{KL(G_i || \mathbf{g}_i) + KL(\mathbf{g}_i || G_i)} \right) \quad (6)$$

where G_i indicates the set of the representations of spans which share at least one token with s_i . Note that we have not directly used the KL divergence as the separation loss but in combination with the $\log(1 + 1/x)$ function to achieve the effect that when KL divergence is small the separation loss is large and vice versa.

3.4 Bidirectional Structure

As the aspect sentiment triplet can be triggered by an aspect terms or an opinion terms, we propose a bidirectional structure to decode the span representations. As shown in Figure 2, the bidirectional structure consists of an aspect decoder and an opinion decoder. The details of each component in the bidirectional structure are given in the following subsections.

3.4.1 Aspect-to-opinion Direction

In aspect-to-opinion direction (Blue arrows and modules in Figure 2), the aspect decoder aims to extract all ATs along with their sentiment from the sentence. We can obtain the confidence score as well as the probability of the sentiment of AT as follows:

$$u_i^a = FFNN_a(\mathbf{g}_i, \theta_a) \quad (7)$$

$$q_i^{a \rightarrow o, a} = \mathbf{w}_{a \rightarrow o, a} u_i^a \quad (8)$$

$$\mathbf{p}_i^{a \rightarrow o, a} = \text{softmax}(q_i^{a \rightarrow o, a}) \quad (9)$$

where $FFNN_A$ represents the FFNN of aspect decoder, θ_a is the parameter for the FFNN, $\mathbf{w}_{a \rightarrow o, a} \in \mathbb{R}^{m \times c^0}$ is a trainable weight vector, and $c^0 \in \{Valid, Invalid\}$ is the number of categories.

Then, giving a set G_a of original span representations of all valid ATs $\mathbf{g}_j^a \in G_a$, we apply the opinion decoder to identify all OTs along with their sentiment for each particular valid AT by exploiting attention mechanism. Similarly, we obtain the probability distribution of the OT's sentiment along with its confidence score via:

$$u_i^o = FFNN_o(\mathbf{g}_i, \theta_o) \quad (10)$$

$$\alpha_{i,j}^{a \rightarrow o} = \frac{\exp(u_i^o)}{\exp(\mathbf{g}_j^a)} \quad (11)$$

$$q_{i,j}^{a \rightarrow o, o} = \mathbf{w}_{a \rightarrow o, o} (u_i^o + \alpha_{i,j}^{a \rightarrow o} \cdot \mathbf{g}_j^a) \quad (12)$$

$$\mathbf{p}_{i,j}^{a \rightarrow o, o} = \text{softmax}(q_{i,j}^{a \rightarrow o, o}) \quad (13)$$

where $FFNN_o$ represents the FFNN of opinion decoder, θ_o is the parameter for the FFNN, $\mathbf{w}_{a \rightarrow o, o} \in \mathbb{R}^{m \times c^*}$ is a trainable weight vector, and $c^* \in \{Positive, Neutral, Negative, Invalid\}$ is the number of sentiment polarity. Furthermore, we define the loss of aspect-to-opinion direction as:

$$\begin{aligned} \mathcal{J}_{a \rightarrow o} = & - \sum_i y_i^{a \rightarrow o, a} \log(q_i^{a \rightarrow o, a}) \\ & - \sum_i \sum_j^{G_a} y_{i,j}^{a \rightarrow o, o} \log(q_{i,j}^{a \rightarrow o, o}) \end{aligned} \quad (14)$$

where $y_i^{a \rightarrow o, a}$ and $y_{i,j}^{a \rightarrow o, o}$ are ground truth labels of the sentiments for AT and OT given a specific valid AT, respectively.

3.4.2 Opinion-to-aspect Direction

As for opinion-to-aspect direction (Red arrows and modules in Figure 2), the opinion decoder is deployed first to extract all the OTs along with their sentiment from the sentence. To minimize the number of model parameters, the opinion decoder in both aspect-to-opinion and opinion-to-aspect directions shares the FFNN features, as described in Equation (10). The probability distribution of the sentiments of OTs as well as the confidence scores can be obtained as:

$$q_i^{o \rightarrow a, o} = \mathbf{w}_{o \rightarrow a, o} u_i^o \quad (15)$$

$$\mathbf{p}_i^{o \rightarrow a, o} = \text{softmax}(q_i^{o \rightarrow a, o}) \quad (16)$$

where $\mathbf{w}_{o \rightarrow a, o} \in \mathbb{R}^{m \times c^0}$ is a trainable weight vector.

Given a set G_o of original span representations of all valid OTs $\mathbf{g}_j^o \in G_o$, the aspect decoder is deployed to identify the ATs and their sentiment for each particular valid OTs. Note that the aspect decoder in opinion-to-aspect direction also shares same FFNN features described in Equation (7) with the aspect decoder in aspect-to-opinion direction. The logits of ATs and their confidence scores in opinion-to-aspect direction can be obtained by:

$$\alpha_{i,j}^{o \rightarrow a} = \frac{\exp(u_i^a)}{\exp(\mathbf{g}_j^o)} \quad (17)$$

$$q_{i,j}^{o \rightarrow a, a} = \mathbf{w}_{o \rightarrow a, a} (u_i^a + \alpha_{i,j}^{o \rightarrow a} \cdot \mathbf{g}_j^o) \quad (18)$$

$$\mathbf{p}_{i,j}^{o \rightarrow a, a} = \text{softmax}(q_{i,j}^{o \rightarrow a, a}) \quad (19)$$

where $\mathbf{w}_{o \rightarrow a, a} \in \mathbb{R}^{m \times c^*}$ is a trainable weight vector.

Finally, the loss for opinion-to-aspect direction is defined as:

$$\begin{aligned} \mathcal{J}_{o \rightarrow a} = & - \sum_i y_i^{o \rightarrow a, o} \log(q_i^{o \rightarrow a, o}) \\ & - \sum_i \sum_j^{G_o} y_{i,j}^{o \rightarrow a, a} \log(q_{i,j}^{o \rightarrow a, a}) \end{aligned} \quad (20)$$

where $y_i^{o \rightarrow a, o}$ and $y_{i,j}^{o \rightarrow a, a}$ are the ground truth labels. Then, we combine the above loss functions to form the loss objective of the entire model:

$$\mathcal{J} = \mathcal{J}_{KL} + \mathcal{J}_{a \rightarrow o} + \mathcal{J}_{o \rightarrow a} \quad (21)$$

Algorithm 1 Inference Strategy

Input: $\mathcal{T}_{a \rightarrow o}, \mathcal{T}_{o \rightarrow a}$ $\mathcal{T}_{a \rightarrow o}$ denotes the triplet extraction results in aspect-to-opinion direction $\mathcal{T}_{o \rightarrow a}$ denotes the triplet extraction results in opinion-to-aspect direction

- 1: Get the overall triplets in both extract directions $\mathcal{T} = \mathcal{T}_{a \rightarrow o} \cup \mathcal{T}_{o \rightarrow a}$
 - 2: **for** $t_i \in \mathcal{T}$ **do**
 - 3: **for** $t_j \in (\mathcal{T} - \{t_i\})$ **do**
 - 4: $t_i = (a_i, o_i, c_i, s_i), t_j = (a_j, o_j, c_j, s_j)$,
 s_i and s_j are the confidence score of the corresponding triplets
 - 5: **if** $a_i \cap a_j \neq \emptyset$ **and** $o_i \cap o_j \neq \emptyset$ **then**
 - 6: **if** $s_i > s_j$ **then**
 - 7: $\mathcal{T} = \mathcal{T} - \{t_j\}$
 - 8: **else**
 - 9: $\mathcal{T} = \mathcal{T} - \{t_i\}$
 - 10: **end if**
 - 11: **end if**
 - 12: **end for**
 - 13: **end for**
 - 14: **return** \mathcal{T}
-

3.5 Inference

In contrast to the mutual exclusivity of the triplets in the token-level method, span-level model cannot guarantee that there are no conflicts between any two triples. Therefore, we propose an inference strategy to eliminate the potential conflicting triplets during the inference process. As illustrated in Algorithm 1, We first combine the extraction results in both directions by taking the union set \mathcal{T} (line 1). Afterwards, for each pair of triplets in the overall triplets set \mathcal{T} that have duplicates in both aspect a and opinion o (line 5), the conflicting results are eliminated by discarding the triplets with lower confidence scores s (line 6-9). Note that in the condition of determining whether two triplets conflict with each other (line 5), the determination of whether the union set is empty is performed on the position index, rather than on the tokens.

4 Experiments

4.1 Datasets

To verify the effectiveness of our network, we conduct experiments on four benchmark datasets²(Xu et al., 2020), which are constructed based on the original SemEval ABSA Challenges and the datasets of (Fan et al., 2019). Table 1 lists the

Datasets		#S	POS	NEU	NEG	#SW	#MW
14LAP	Train	1266	1692	166	480	1586	752
	Dev	310	404	54	119	388	189
	Test	492	773	66	155	657	337
14RES	Train	906	817	126	517	824	636
	Dev	219	169	36	141	190	156
	Test	328	364	63	116	291	252
15RES	Train	605	783	25	205	678	335
	Dev	148	185	11	53	165	84
	Test	322	317	25	143	297	188
16RES	Train	857	1015	50	329	918	476
	Dev	210	252	11	76	216	123
	Test	326	407	29	78	344	170

Table 1: Statistics of the datasets. ‘#S’ denotes the numbers of sentence, ‘POS’, ‘NEU’, and ‘NEG’ denote the numbers of positive, neutral, and negative triplets respectively. ‘#SW’ denotes the numbers of triplets where the ATs and OTs are single word spans. ‘#MW’ denotes the numbers of triplets that at least one of the ATs or the OTs are multi-word spans.

statistics of these datasets.

4.2 Experimental Setting

We adopt the cased base version of BERT (Devlin et al., 2018) in our experiments, which contains 110M parameters. During training, we use AdamW (Loshchilov and Hutter, 2017) to optimize the model parameters. The fine-tuning rate for BERT and the learning rate for other models are set to 1e-5 and 1e-4, respectively. Meanwhile, the mini-batch size is set to 16 and the dropout rate is set to 0.1. The maximum length of generated spans is set to 8. We train our framework in a total of 120 epochs on a NVIDIA Tesla V100 GPU.

4.3 Evaluation

To comprehensively evaluate the performance of different methods, we use *precision*, *recall*, *F1-score* as the evaluation metrics. The extracted ATs and OTs are considered correct if and only if predicted spans exactly match the ground truth spans. In the experiments, we select the testing results when the model achieves the best performance on the development set.

4.4 Baselines

To demonstrate the effectiveness of our network, we compare our method with the following baselines:

- **Peng-two-stage** (Peng et al., 2020) is a two-stage pipeline model. Peng-two-stage extracts

²<https://github.com/xuuuuuu/SemEval-Triplet-data/tree/master/ASTE-Data-V2-EMNLP2020>

	14LAP			14RES			15RES			16RES		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
PENG-two-stage	40.40	47.24	43.50	44.18	62.99	51.89	40.97	54.68	46.79	46.76	62.97	53.62
JETt	51.48	42.65	46.65	70.20	53.02	60.41	62.14	47.25	53.68	72.12	57.20	63.41
JETo	58.47	43.67	50.00	67.97	60.32	63.92	58.35	51.43	54.67	64.77	61.29	62.98
GTS-BERT	57.52	51.92	54.58	70.92	69.49	70.20	59.29	58.07	58.67	68.58	66.60	67.58
Dual-MRC	57.39	53.88	55.58	71.55	69.14	70.32	63.78	51.87	57.21	68.60	66.24	67.40
B-MRC	70.89*	50.20*	58.78*	75.41*	64.04*	69.26*	69.83*	56.04*	58.74*	69.03*	66.02*	67.49*
Span-ASTE	63.44	55.84	59.38	72.89	70.89	71.85	62.18	64.45	63.27	69.45	71.17	70.26
Ours	65.68	59.88	62.65	76.36	72.43	74.34	69.93	60.41	64.82	71.59	72.57	72.08

Table 2: Precision (%), Recall (%) and F1 score (%) on the test set of the ASTE tasks. State-of-the-art results are marked bold. * indicates that the result is reproduced by us.

both aspect-sentiment pairs and opinion terms in the first stage. In the second stage, Peng-two-stage pairs up the extraction results into triplets via an relation classifier.

- **JET** (Xu et al., 2020) is an end-to-end model which proposes a novel position-aware tagging scheme to jointly extracting the triplets. It also designs factorized feature representations so as to effectively capture the interaction among the triplet factors.
- **GTS** (Wu et al., 2020) is an end-to-end model which formulates ASTE as a unified grid tagging task. It first extracts the sentiment feature of each token, and then gets the initial prediction probabilities of token pairs based on these token-level features. It also designs an inference strategy to exploit the potential mutual indications between different opinion factors and performs the final prediction.
- **Dual-MRC** (Mao et al., 2021) is a joint training model which consists of two machine reading comprehensions. One of the MRC is for aspect term extraction, and another is for aspect-oriented opinion term extraction and sentiment classification.
- **B-MRC** (Chen et al., 2021) formalizes the ASTE task as a multi-turn machine reading comprehension task, and proposes three types of queries to extract targets, opinions and the sentiment polarities of aspect-opinion pairs, respectively.
- **Span-ASTE** (Xu et al., 2021) considers all possible spans in a sentence and build the interaction between the whole spans of aspect terms and opinion terms when predicting their sentiment relation. They also propose a dual-channel span pruning strategy to ease the high

	KL Loss	JS Loss	EM Loss	CS Loss	no Loss
Ours	62.65	62.34	61.17	61.51	60.82
w/o IS	61.48	61.05	60.40	60.82	60.08
Ours _{a→o}	62.14	61.86	60.08	61.70	60.66
w/o IS	61.09	61.18	60.08	60.67	58.78
Ours _{o→a}	62.38	62.03	60.88	61.70	60.68
w/o IS	61.75	60.15	59.90	60.41	59.38

Table 3: Experimental results of the ablation study in 14LAP dataset (*F1-score*, %). ‘w/o IS’ denotes not using the inference strategy described in Section 3.5. ‘KL Loss’, ‘JS Loss’, ‘EM Loss’, and ‘CS Loss’ denote that similar span separation loss is constructed base on KL divergence, JS divergence, Euclidean Metric, and Cosine Similarity, respectively. ‘no Loss’ means no application of similar span separation loss.

computational cost caused by span enumeration.

4.5 Main Results

Table 2 reports the results of our framework and baseline models. According to the results, our framework achieves state-of-the-art performance on all datasets. Specifically, our framework surpasses the best baselines by an average of 2.3 *F1-score* on ASTE. This result demonstrates that our framework can take advantage of bidirectional decoding and efficiently distinguish the span representation. Although some of the recall scores are slightly lower than Span-ASTE, the increase in precision significantly outperforms the previous baselines in most datasets, which shows the higher prediction accuracy of our network. It is worth noting that BMRC and Dual-MRC achieve better performance than JET and PENG-two-stage. This is probably because BMRC and Dual-MRC formalize the ASTE task as a multi-turn machine reading comprehension task and benefit from asking the model questions. Unlike those approaches, Span-ASTE and our method both utilize the span-level interactions to handle the ASTE task and avoid the cascading errors. Moreover, our model outperforms Span-ASTE because our method identify the

triplets from both aspect-to-opinion and opinion-to-aspect directions, rather than matching each aspect span and opinion span. Besides, our network also take advantage of similar span separation loss and inference strategy to overcome the drawback of mutual exclusivity absence among spans.

4.6 Ablation Study

To validate the origination of the significant improvement in our network, we conduct ablation experiments on 14LAP datasets. As shown in Table 3, our bidirectional model yields better results than unidirectional models, which clearly indicates the superiority of the collaboration in both two directions on decoding the span representations. And the inference results from opinion terms to aspect terms are better than the other direction, which may due to the simplicity of extracting opinion terms in the 14LAP dataset.

Moreover, the inference strategy has exhibited the enhancement on model performance. However, the improvement brought by the inference strategy is not significant, because conflicting triplets tend to exist among multi-token results, and only a small percentage of triplets containing multi-token terms in 14LAP dataset. We believe the effect of the inference strategy will be more obvious in datasets enriched with multi-token triplets.

In addition, to demonstrate the effectiveness of our proposed similar span separation loss based on KL divergence, we further design similar range separation losses based on JS divergence, Euclidean distance and cosine similarity. The experimental results show that all these loss functions have a boosting effect on our network, and the separation loss based on KL divergence performs the best. Note that numerous similarity measures can be used to separate similar spans, among which there may be some better measures that can bring more improvement to the model.

4.7 Effect of Entity Length

To investigate the performance of different methods on the ATE and OTE with different entity lengths, we report the F1 scores of our framework, Span-ASTE, GTS, and B-MRC on the extraction task with different lengths of entities. The results are illustrated in Figure 3. As the entity length increases, the performance gap between our framework and other models becomes more obvious. Since our method directly models span-level feature for each entity and alleviates the drawback of no mu-

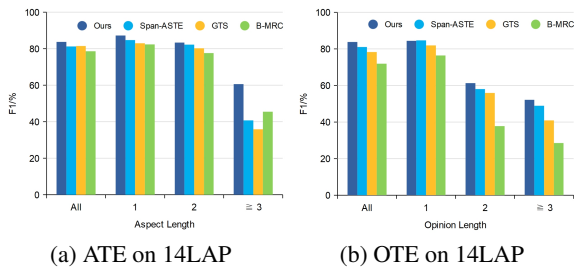


Figure 3: Effects of entity length for aspect term extraction and opinion term extraction (*F1-score*, %).

tual exclusivity among spans, our method will not be greatly affected with entity lengths increasing. In fact, most of the contribution to the improvement in our model comes from the performance in multi-token entities.

4.8 Effect of Multiple Triplets

	All	1	2	3	4	≥ 5
Ours	62.65	62.80	63.54	64.81	57.14	60.46
Span-ASTE	59.65	61.18	61.01	61.67	56.00	35.00
GTS	58.57	58.21	60.26	65.00	56.57	30.77
B-MRC	58.78	58.10	64.75	55.31	35.56	0.00

Table 4: Effects of multiple triplets in a sentence in 14LAP (*F1-score*, %).

To further verify the ability of our framework to handle multiple triplets, we compare the performance of our network and other baselines on ASTE task with different number of triplets in the sentences, and the results are shown in Table 4. We divide the sentences in 14LAP testset into 5 subclasses. Each subclass contains sentences with 1, 2, 3, 4, or ≥ 5 triplets, respectively. When extracting triplets from sentences that contain 1 or 2 triplets, the performance of our framework is competitive to other models. However, when the number of triplets increases, the performance of Span-ASTE, GTS, and B-MRC decrease significantly, while the performance of our network remains stable or even slightly increases. These experimental results demonstrate the efficiency and stability of our framework in handling multiple triplets in a sentence.

5 Conclusions

In this work, we propose a span-level bidirectional network for ASTE tasks. This span-level model can take advantages from both aspect-to-opinion and opinion-to-aspect directions. The bidirectional decoding can ensure that either an AT or an OT can

trigger an aspect sentiment triplet, which is more in line with human perception. For the shortcoming that mutual exclusivity cannot be guaranteed among spans, we deploy the similar span separation loss to guide the model in discriminating similar spans. We further design an inference strategy to eliminate conflicting triplet results that are specific to span-level models. The experimental results demonstrate that our network significantly outperforms the compared baselines and achieves state-of-the-art performance.

Limitations

Although in the previous section we showed the advanced performance of the network we designed, there are still some weaknesses in our model.

	MACs(G)	Params(M)
Ours	120.044	129.884
Span-ASTE	444.55	110
B-MRC	19.624	85.611
GTS	520.765	88.006

Table 5: Efficiency Comparison.

First, our model uses spans as input, and enumerating all possible spans inevitably increases the input size of the model, so span-level models tend to have larger computations than token-level models. As shown in Table 5, our network requires about 6 times more floating-point computations than the B-MRC model. While the Span-ASTE and GTS models require more computation, this is because Span-ASTE needs to match every aspect terms and opinion terms and GTS model extracts triplets by classifying the internal elements of a square matrix consisting of sentences in rows and columns.

Second, to reduce the input size of the model, we set the maximum span length of the spans to 8 to include as many potential aspect terms and opinion terms as possible. However, in some datasets with long extraction targets, the span-level model must increase the maximum span limit, thus affecting the performance of the model. Therefore, our model is suitable only for the tasks of extracting short targets.

Third, both the similar span separation loss and inference strategy proposed in this paper are used to alleviate the shortcoming of the missing mutual exclusivity in span-level models, while the inputs of token-level models are naturally mutually exclusive. So the similar range separation loss and in-

ference strategies are not applicable to token-level models.

Ethics Statement

This article does not contain any study with human participants or animals performed by any of the authors. And all authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- Shaowei Chen, Jie Liu, Yu Wang, Wenzheng Zhang, and Ziming Chi. 2020. [Synchronous double-channel recurrent network for aspect-opinion pair extraction](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 6515–6524. Association for Computational Linguistics.
- Shaowei Chen, Yu Wang, Jie Liu, and Yuelin Wang. 2021. Bidirectional machine reading comprehension for aspect sentiment triplet extraction. *arXiv preprint arXiv:2103.07665*.
- 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*.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. [Adaptive recursive neural network for target-dependent twitter sentiment classification](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 2: Short Papers*, pages 49–54. The Association for Computer Linguistics.
- 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.
- Lei Gao, Yulong Wang, Tongcun Liu, Jingyu Wang, Lei Zhang, and Jianxin Liao. 2021. [Question-driven span labeling model for aspect-opinion pair extraction](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 12875–12883. AAAI Press.

- Ruidan He, Wee Sun Lee, Hwee Tou Ng, and Daniel Dahlmeier. 2018. [Exploiting document knowledge for aspect-level sentiment classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 579–585. Association for Computational Linguistics.
- 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 *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 504–515. Association for Computational Linguistics.
- Mengting Hu, Shivan Zhao, Li Zhang, Keke Cai, Zhong Su, Renhong Cheng, and Xiaowei Shen. 2019. [CAN: constrained attention networks for multi-aspect sentiment analysis](#). 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 4600–4609. Association for Computational Linguistics.
- Hao Li and Wei Lu. 2017. [Learning latent sentiment scopes for entity-level sentiment analysis](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 3482–3489. AAAI Press.
- Xin Li, Lidong Bing, Piji Li, and Wai Lam. 2019a. [A unified model for opinion target extraction and target sentiment prediction](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6714–6721. AAAI Press.
- Xin Li, Lidong Bing, Piji Li, Wai Lam, and Zhimou Yang. 2018. [Aspect term extraction with history attention and selective transformation](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 4194–4200. ijcai.org.
- Xin Li and Wai Lam. 2017. [Deep multi-task learning for aspect term extraction with memory interaction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2886–2892. Association for Computational Linguistics.
- Zheng Li, Ying Wei, Yu Zhang, Xiang Zhang, and Xin Li. 2019b. [Exploiting coarse-to-fine task transfer for aspect-level sentiment classification](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 4253–4260. AAAI Press.
- Kang Liu, Heng Li Xu, Yang Liu, and Jun Zhao. 2013. [Opinion target extraction using partially-supervised word alignment model](#). In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 2134–2140. IJCAI/AAAI.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- Dehong Ma, Sujian Li, Fangzhao Wu, Xing Xie, and Houfeng Wang. 2019. [Exploring sequence-to-sequence learning in aspect term extraction](#). 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 3538–3547. Association for Computational Linguistics.
- Yue Mao, Yi Shen, Chao Yu, and Longjun Cai. 2021. [A joint training dual-mrc framework for aspect based sentiment analysis](#). In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 13543–13551. AAAI Press.
- 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.
- Soujanya Poria, Erik Cambria, and Alexander F. Gelbukh. 2016. [Aspect extraction for opinion mining with a deep convolutional neural network](#). *Knowl. Based Syst.*, 108:42–49.
- Guang Qiu, Bing Liu, Jiajun Bu, and Chun Chen. 2011. [Opinion word expansion and target extraction through double propagation](#). *Comput. Linguistics*, 37(1):9–27.
- Duyu Tang, Bing Qin, and Ting Liu. 2016. [Aspect level sentiment classification with deep memory network](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 214–224. The Association for Computational Linguistics.

- Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. 2016. [Recursive neural conditional random fields for aspect-based sentiment analysis](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 616–626. The Association for Computational Linguistics.
- Shengqiong Wu, Hao Fei, Yafeng Ren, Donghong Ji, and Jingye Li. 2021. [Learn from syntax: Improving pair-wise aspect and opinion terms extraction with rich syntactic knowledge](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, pages 3957–3963. ijcai.org.
- Zhen Wu, Chengcan Ying, Fei Zhao, Zhifang Fan, Xinyu Dai, and Rui Xia. 2020. [Grid tagging scheme for aspect-oriented fine-grained opinion extraction](#). *CoRR*, abs/2010.04640.
- Hu Xu, Bing Liu, Lei Shu, and Philip S. Yu. 2018. [Double embeddings and cnn-based sequence labeling for aspect extraction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 592–598. Association for Computational Linguistics.
- Lu Xu, Yew Ken Chia, and Lidong Bing. 2021. [Learning span-level interactions for aspect sentiment triplet extraction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 4755–4766. 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.
- Jianfei Yu, Jing Jiang, and Rui Xia. 2019. [Global inference for aspect and opinion terms co-extraction based on multi-task neural networks](#). *IEEE ACM Trans. Audio Speech Lang. Process.*, 27(1):168–177.
- 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.