

# STAGE: Span Tagging and Greedy Inference Scheme for Aspect Sentiment Triplet Extraction

Shuo Liang<sup>1,2</sup>, Wei Wei<sup>†,1,2</sup>, Xian-Ling Mao<sup>3</sup>, Yuanyuan Fu<sup>2,4</sup>, Rui Fang<sup>2,4</sup>, Dangyang Chen<sup>2,4</sup>

<sup>1</sup> Cognitive Computing and Intelligent Information Processing (CCIIP) Laboratory, School of Computer Science and Technology, Huazhong University of Science and Technology

<sup>2</sup> Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL)

<sup>3</sup> Department of Computer Science and Technology, Beijing Institute of Technology

<sup>4</sup> Ping An Property & Casualty Insurance company of China, Ltd

shuoliang@hust.edu.cn, weiw@hust.edu.cn, maoxl@bit.edu.cn,

fuyuan83@gmail.com, fangrui051@pingan.com.cn, chendangyang273@pingan.com.cn

## Abstract

Aspect Sentiment Triplet Extraction (ASTE) has become an emerging task in sentiment analysis research, aiming to extract triplets of the aspect term, its corresponding opinion term, and its associated sentiment polarity from a given sentence. Recently, many neural networks based models with different tagging schemes have been proposed, but almost all of them have their limitations: heavily relying on 1) prior assumption that each word is only associated with a single role (*e.g.*, aspect term, *or* opinion term, *etc.*) and 2) word-level interactions and treating each opinion/aspect as a set of independent words. Hence, they perform poorly on the complex ASTE task, such as a word associated with multiple roles or an aspect/opinion term with multiple words. Hence, we propose a novel approach, **Span Tagging and Greedy infErence** (STAGE), to extract sentiment triplets in span-level, where each span may consist of multiple words and play different roles simultaneously. To this end, this paper formulates the ASTE task as a multi-class **span** classification problem. Specifically, STAGE generates more accurate aspect sentiment triplet extractions via exploring span-level information and constraints, which consists of two components, namely, *span tagging* scheme and *greedy inference* strategy. The former tag all possible candidate spans based on a newly-defined tagging set. The latter retrieves the aspect/opinion term with the maximum length from the candidate sentiment snippet to output sentiment triplets. Furthermore, we propose a simple but effective model based on the STAGE, which outperforms the state-of-the-arts by a large margin on four widely-used datasets. Moreover, our STAGE can be easily generalized to other pair/triplet extraction tasks, which also demonstrates the superiority of the proposed scheme STAGE.

## Introduction

Aspect Sentiment Triplet Extraction (ASTE) aims to identify all sentiment triplets from an input sentence, namely, the aspect term (*a*) with its corresponding opinion term (*o*) and sentiment polarity (*s*) (as shown in Figure 1). It can be applied in many downstream tasks, such as dialogue generation (Liu et al. 2022; Wei et al. 2021, 2019) and recommendation system (Wang et al. 2022; Zou et al. 2022; Zhao et al.

Sentences		Triplets		
		Aspect Term	Opinion Term	Sentiment
(a)	Love the <b>stability</b> of the Mac software and operating system .	stability	Love	Positive
		Mac software	Love	Positive
		operating system	Love	Positive
			stability	Positive
(b)	Keyboard <b>good sized</b> and <b>easy to use</b> .	Keyboard	good sized	Positive
			easy to use	Positive
		use	easy	Positive
(c)	After <b>replacing the hard drive</b> the <b>battery</b> stopped working ( 3 months of use ) which was frustrating.	hard drive	replacing	Neutral
		battery	stopped working	Negative

Figure 1: Examples of ASTE task from the real-world dataset. In each sentence, aspect terms are in blue and opinion terms are in rounded rectangles.

2022). Indeed, ASTE can be decomposed into several sub-tasks, such as Aspect Term Extraction (ATE) (Li and Lam 2017; Xu et al. 2018; Dai and Song 2019), Opinion Term Extraction (OTE) (Fan et al. 2019; Dai and Song 2019), Aspect-Based Sentiment Classification (ABSC) (Zhang, Li, and Song 2019; Tang et al. 2020; Liang et al. 2022), and Aspect-Opinion Pair Extraction (AOPE) (Wang et al. 2017; Zhao et al. 2020). And a straightforward method is to independently extract elements of sentiment triplet through those subtasks in a pipeline fashion (Peng et al. 2020; Chen et al. 2021). However, the primary obstacle for those pipeline methods is the well-known error propagation problem.

Later, many efforts resort to extracting sentiment triplets in an end-to-end framework (Xu et al. 2020; Wu et al. 2020; Zhang et al. 2020; Yu Bai Jian et al. 2021; Xu, Chia, and Bing 2021; Mukherjee et al. 2021; Zhang et al. 2021; Chen et al. 2022; Liu, Li, and Li 2022). Some works (Mukherjee et al. 2021; Zhang et al. 2021) formulate it as a generative problem and achieve good performance under the benefit of existing generative pre-trained models, such as BART (Lewis et al. 2020) and T5 (Raffel et al. 2020). In addition, most of end-to-end approaches focus on design-

ing a new tagging scheme, converting ASTE into a classification problem towards each *word* (Xu et al. 2020) or each *word pair* (Wu et al. 2020; Chen et al. 2022). Despite their success, those *word-level* tagging schemes faces several problems: heavily relying on (1) **Prior assumption that restricts the diversity of word roles**. Those tagging schemes only work when each word is only associated with a single role (the aspect, *or* the opinion, *etc.* ) *or* when each word pair corresponds to one specific relation. As a result, they cannot perform well when it comes to a word playing multiple roles or a word pair having multiple relations. (2) **Word-level interactions, thus falling short of utilizing span-level information**. Those methods only focus on the interactions between words and treat the multi-word term as a set of independent words, thus separating the semantic information of the whole span when facing multi-word aspect/opinion terms. It also causes difficulty in guaranteeing the sentiment consistency when extracting triplets. Without loss of generality, we illustrate those problems with examples in Figure 1: (1) **Multiple roles of a word**. “Stability” from sentence (a) has two roles simultaneously, that is, being (part of) the aspect term and the opinion term, which confuses the *word-level* tagging schemes when figuring out the proper tag; (2) **Multiple relations of a word-pair**. In sentence (b), not only “easy” and “use” are words of the same opinion term (*i.e.*, easy to use), but they form a valid aspect-opinion pair as well, resulting in multiple relations of the word pair “easy-use”, which also challenges the existing tagging scheme; (3) **Multi-word term**. It is common that an aspect/opinion term contains multiple words, and a significant semantic difference may exist between the whole term and the words in it, such as “hard drive” from sentence (c). Treating those multi-word terms as a set of independent words, the *word-level* tagging schemes fail to utilize their complete semantics and consequently deteriorate the performance of the extraction task. Moreover, in *word-level* tagging methods, words that make up a term are classified independently and thus may contain different sentiment polarities. Therefore, extra effort is required to deal with this sentiment inconsistency problem when outputting triplets.

The aforementioned problems motivate us to explore a *span-level* tagging method and formulate it as a multi-class classification problem towards each *span*. Generally, it is observed that a span in the given sentences may have three roles: aspect term, opinion term, and a snippet that includes an aspect term and the corresponding opinion term, forming its boundaries (*i.e.*, sentiment snippet), such as “love the stability” from Figure 1 (a). Note that a span can play different roles simultaneously.

Based on the above viewpoint, we pre-define three role dimensions (*i.e.*, aspect term, opinion term, and sentiment snippet) and propose a novel *span-level* tagging-based approach, **Span TA**gging and **Greedy infE**rence (STAGE), to solve ASTE in an end-to-end fashion. By exploring span-level information and constraints, STAGE can efficiently generate more accurate sentiment triplets. Specifically, it consists of two components: (1) *Span tagging* scheme: which tags all possible spans based on the defined role dimensions, guaranteeing the diversity of span roles. Thus, it

can depict a complete picture of each span, like how many and what roles a span plays. (2) *Greedy inference* strategy: which decodes all triplets simultaneously in a more accurate and efficient way by (i) considering the mutual span constraints between the sentiment snippets and aspect/opinion terms and (ii) only retrieving the aspect/opinion term with the maximum length from a sentiment snippet, compared to previous decoding algorithms that look into every possible pair combination of candidate aspects and opinions to see whether they are valid. Therefore, naturally performing at the span-level, STAGE can overcome the problems of existing *word-level* tagging schemes. Based on STAGE, we propose a simple model to demonstrate its effectiveness.

Moreover, with a minor tagging set change, STAGE is very easily generalized to other pair/triplet extraction tasks, such as entity and relation extraction (Zhong and Chen 2021; Yan et al. 2021), a fundamental task aiming to extract (head entity, tail entity, relation) triplets from the given sentence<sup>1</sup>.

In summary, the key contributions are as follows:

(1) To the best of our knowledge, we make the first effort to explore the *span-level* tagging method and formulate the ASTE task as a multi-class **span** classification problem.

(2) We propose a novel end-to-end tagging-based approach, **Span TA**gging and **Greedy infE**rence (STAGE). Specifically, it consists of the *span tagging* scheme that considers the diversity of span roles, overcoming the limitations of existing tagging schemes, and the *greedy inference* strategy that considers the span-level constraints, generating more accurate triplets efficiently.

(3) We propose a simple but effective model based on STAGE and extensive experiments on four widely-used English datasets verify its effectiveness. Furthermore, it can be easily generalized to other pair/triplet extraction tasks, which also demonstrates the superiority of STAGE.

## Related Work

We present some commonly used tagging schemes in the ASTE task, which can be categorized into two main kinds.

### Sequence Tagging

The base tagging scheme is BIESO tagging<sup>2</sup>. Most pipeline models firstly adopt it to extract aspects and opinions separately and then determines whether any two of them can form a valid aspect-opinion pair with the corresponding sentiment. Generally, they view ASTE as two sequence labeling problems (*i.e.*, classification problems towards each *word*) plus one classification problem towards each candidate aspect-opinion pair. Peng et al. (2020) propose a two-stage pipeline approach and utilize a unified tagging scheme to tag the boundary of aspect term with its sentiment and a common BIESO tagging scheme to extract opinion term.

However, the pipeline approaches ignore the interactions between three triplet elements and commonly suffer from error propagation. Besides, those tagging schemes encode no positional information and thus fail to model rich interactions among aspect and opinion terms. So Xu et al. (2020)

<sup>1</sup>More details are shown in Section Analysis.

<sup>2</sup>BIESO means “begin, inside, end, single, other”, respectively.

Sub Tags	Meaning
A	an aspect term
O	an opinion term
NEG	a sentiment snippet containing <i>negative</i> aspect-opinion pair.
NEU	a sentiment snippet containing <i>neutral</i> aspect-opinion pair.
POS	a sentiment snippet containing <i>positive</i> aspect-opinion pair.
N	nothing in the specific dimension

Table 1: Descriptions of sub tags.

propose JET, an extended BIESO tagging scheme with position index information. However, it still fails to simultaneously tackle the situations that one aspect/opinion term relates to multiple opinion/aspect terms.

### Table Tagging

The basic idea is to tag a  $|n| \times |n|$  table  $T$ , where  $n$  is the length of the input text and  $T[i][j]$  corresponds to the relation between  $i$ -th word and  $j$ -th word. Wu et al. (2020) firstly propose a grid tagging scheme, utilizing a tag set  $\{A, O, POS, NEU, NEG, N\}$  to donate five possible relations of a word-pair. In order to utilize more relations between words, Chen et al. (2022) define ten relations and consider diverse linguistic features. The above approaches both perform in an end-to-end way, avoiding error prorogation and can simultaneously handle the case of one aspect term corresponding to multiple opinion terms and vice versa. However, as we mentioned before, they still face several limitations.

## Methodology

In this section, we first introduce the ASTE task and then explain our proposed span tagging and greedy inference (STAGE) scheme in detail. Finally, we present our implemented model.

### Problem Statement

Given the input sentence  $X = \{w_1, w_2, \dots, w_n\}$ , where  $w_i$  denotes a word and  $n$  is the sentence length. The goal of ASTE task is to extract a set of sentiment triplets  $\mathcal{T} = \{(a, o, s)_m\}_{m=1}^{|\mathcal{T}|}$  from the sentence  $X$ , where  $a$ ,  $o$  and  $s$  represent aspect term, opinion term and sentiment polarity separately and  $s \in \{POS, NEG, NEU\}$ . We use  $SP = \{SP_{1,1}, SP_{1,2}, \dots, SP_{i,j}, \dots, SP_{n,n}\}$  to illustrate all possible enumerated spans in  $X$ , where  $i$  and  $j$  represent the start and end positions in  $X$ , respectively. So the formal definition for the sentiment snippet is as follows,  $SP_{i,j}$  is a sentiment snippet means  $(SP_{i,k}, SP_{l,j})$  is a valid (aspect, opinion) or (opinion, aspect) pair, where  $i \leq k, l \leq j$ .

### STAGE: Span Tagging and Greedy inference

**Span Tagging** Considering the diversity of span roles, our span tagging scheme depict each span (*i.e.*,  $SP_{i,j}$ ) in three role dimensions: whether  $SP_{i,j}$  is a valid aspect term, a valid opinion term and a valid sentiment snippet. We use  $role_{i,j}^a$ ,

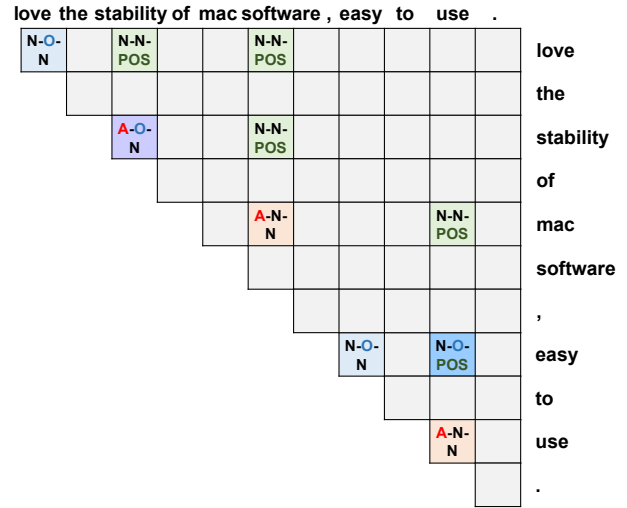


Figure 2: Example of 3D-version span tagging. Cells with gray background represent “N-N-N”.

$role_{i,j}^o, role_{i,j}^s$  to represent them separately, where  $role_{i,j}^a \in \{N, A\}$ ,  $role_{i,j}^o \in \{N, O\}$ ,  $role_{i,j}^s \in \{N, NEG, NEU, POS\}$ . The meanings of those sub tags are shown in Table 1.

By considering those three dimensions independently, we propose a **3D-version** span tagging method with the tag set  $\{N, A\}-\{N, O\}-\{N, NEG, NEU, POS\}$ . The tagging results are like an upper triangular table  $T$ , where  $T[i][j]$  represents the tag for span  $SP_{i,j}$ , as illustrated by Figure 2. Note that the tag for span “easy to use” is “N-O-POS” because the span is not only an opinion term but also contains “(use, easy)”, a valid aspect-opinion pair with positive sentiment.

Apart from considering the three dimensions independently, we also investigate two variants:

1) **2D-version**: With tag set  $\{N, O, A\}-\{N, NEG, NEU, POS\}$ , it can be viewed as projecting the aspect and opinion role dimensions into one. The tag “N” in the first dimension denotes the span is neither an aspect nor an opinion. This variant cannot handle the situations that a span being aspect and opinion simultaneously.

2) **1D-version**: With tag set  $\{N, NEG, NEU, POS, O, A\}$ , it can be viewed as projecting the three role dimensions into one. The tag “N” means the span plays no roles. This variant fails to deal with the cases that a span has multiple roles.

**Greedy Inference** Given the tagging results in our span tagging scheme, in order to generate more accurate triplets efficiently, we propose a greedy inference strategy, where the “**greedy**” means retrieving the aspect/opinion term with the maximum length separately in a valid sentiment snippet.

The details are shown in Algorithm 1. Firstly, we obtain all possible aspects, opinions and sentiment snippets from the tagging results (line 1 to line 4). Then we look into every sentiment snippet (line 5 to line 24) and consider two cases:

(1) The aspect term appears before the opinion term (line 7 to line 14), which means the former shares the start index and the latter shares the end index with the current sen-

**Input:**

The tagging results  $P$  of the sentence  $X$  with length  $n$ .  $P_{i,j}$  represents the tag label of the span  $SP_{i,j}$  and can be decomposed into three dimensions  $role_{i,j}^a$ ,  $role_{i,j}^o$  and  $role_{i,j}^s$ .

**Output:**

Triples  $\mathcal{T}$  of the given sentence;

- 1: Initialize  $\mathcal{A} = \{\}, \mathcal{O} = \{\}, \mathcal{D} = \{\}, \mathcal{T} = \{\}$ ;
- 2:  $\mathcal{A} = \{(i, j) | role_{i,j}^a = A, 0 \leq i \leq j \leq n\}$
- 3:  $\mathcal{O} = \{(i, j) | role_{i,j}^o = O, 0 \leq i \leq j \leq n\}$
- 4:  $\mathcal{D} = \{(i, j, role_{i,j}^s) | role_{i,j}^s \neq N, 0 \leq i \leq j \leq n\}$
- 5: **for** each  $(i, j, s)$  in  $\mathcal{D}$  **do**
- 6:   // CASE-1: aspect term before opinion term;
- 7:    $CA = \{k | i \leq k \leq j, (i, k) \in \mathcal{A}\}$
- 8:    $CO = \{l | i \leq l \leq j, (l, j) \in \mathcal{O}\}$
- 9:   **if**  $CA \neq \emptyset$  and  $CO \neq \emptyset$  **then**
- 10:     remove  $j$  (if exists) from  $CA$  when  $|CA| > 1$
- 11:     remove  $i$  (if exists) from  $CO$  when  $|CO| > 1$
- 12:      $k = \max(CA), l = \min(CO)$
- 13:      $\mathcal{T} \leftarrow \mathcal{T} \cup (SP_{i,k}, SP_{l,j}, s)$
- 14:   **end if**
- 15:   // CASE-2: aspect term after opinion term;
- 16:    $CO = \{k | i \leq k \leq j, (i, k) \in \mathcal{O}\}$
- 17:    $CA = \{l | i \leq l \leq j, (l, j) \in \mathcal{A}\}$
- 18:   **if**  $CO \neq \emptyset$  and  $CA \neq \emptyset$  **then**
- 19:     remove  $j$  (if exists) from  $CO$  when  $|CO| > 1$
- 20:     remove  $i$  (if exists) from  $CA$  when  $|CA| > 1$
- 21:      $k = \max(CO), l = \min(CA)$
- 22:      $\mathcal{T} \leftarrow \mathcal{T} \cup (SP_{l,j}, SP_{i,k}, s)$
- 23:   **end if**
- 24: **end for**
- 25: **return**  $\mathcal{T}$ ;

timent snippet (line 7 to line 8). If both candidate aspects and opinions exist, we then separately retrieve the one with the maximum length, forming a valid triplet (line 9 to line 14). Considering that the sentiment snippet itself can be the aspect/opinion, we only view it as an aspect/opinion when there is no other counterpart candidates (line 10/11);

(2) The aspect term appears after the opinion term (line 16 to line 23), which is very similar to the previous case, except the candidate aspects share the end index and the candidate opinions share the start index with current sentiment snippet.

Finally, we can obtain all valid sentiment triplets after examining all sentiment snippets. Compared to previous decoding algorithms, our strategy has advantages in:

(1) **Accuracy:** The mutual span-level constraints guarantee the accuracy of the extracted sentiment triplets. For the sentiment snippet, which also indicates the pairing of the aspect and opinion, not only itself should have correct sentiment polarity, but also it should contain valid aspects and opinions that simultaneously sharing one of its boundaries. For the aspect/opinion term, not only themselves should be classified correctly, but also the boundaries of them should form a valid sentiment snippet.

(2) **Efficiency:** Contrary to the common practice that

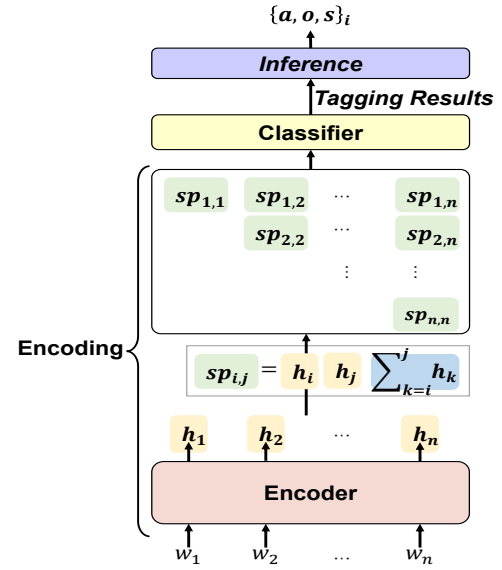


Figure 3: Model Architecture.

looks into every possible pair combination of candidate aspects and opinions to see whether they are valid, which contributes to heavy computational overhead, our method extracts the triplets based on the predicted sentiment snippets, the number of which is much less. More specifically, the time complexity is  $O(n^2)$  with greedy inference and  $O(n^4)$  with common practice, where  $n$  is length of input text.

### STAGE Model

Based on STAGE, we implement a simple BERT-based model. Shown in Figure 3, the overall model architecture contains an encoder that outputs contextual word representations for generating span representations and a classifier that predicts the tag of each span. The inference module is used to generate triplets from the tagging results.

**Encoding** The key point of this process is to obtain span representations. Firstly, we adopt BERT (Devlin et al. 2019) plus an additional fully connected layer as our encoder to generate contextual word representations  $\mathbf{h}_i$  by,

$$\hat{\mathbf{h}} = \text{BERT}([\text{CLS}] + \{w_i\} + [\text{SEP}]), \quad (1)$$

$$\mathbf{h}_i = \mathbf{W}_{\text{fc}} \left( \frac{\hat{\mathbf{h}}}{|\text{BertT}(w_i)|} \sum_{k \in \text{BertT}(w_i)} \hat{\mathbf{h}}_k \right) + \mathbf{b}_{\text{fc}}, \quad (2)$$

where  $\text{BertT}(w_i)$  returns the index set of  $w_i$ 's sub-words in BERT sequence, and  $|\cdot|$  returns its length.  $\mathbf{W}_{\text{fc}}$  and  $\mathbf{b}_{\text{fc}}$  are the parameters of the fully connected layer, which is used to project the word representations from BERT to a lower dimension space. Then we form the span representations by,

$$\mathbf{sp}_{i,j} = \mathbf{h}_i \oplus \mathbf{h}_j \oplus \sum_{k=i}^j \mathbf{h}_k, \quad (3)$$

where  $\oplus$  means vector concatenation. In this way, the model can utilize both the boundary information and the whole semantics of the span.

Dataset		#S	#T	#Pos.	#Neu.	#Neg.	Role distributions of spans (aspect-opinion-sentiment snippet)						
							A-N-N	N-O-N	N-N-S	A-O-N	A-N-S	N-O-S	A-O-S
14lap	Train	906	1460	817	126	517	1279	1263	1454	1	1	4	0
	Dev	219	345	169	36	140	296	304	344	0	0	0	0
	Test	328	541	364	63	114	463	473	540	0	0	1	0
14res	Train	1266	2337	1691	166	480	2043	2072	2313	0	7	13	1
	Dev	310	577	404	54	119	497	498	569	0	3	5	0
	Test	492	994	773	66	155	840	850	984	1	7	3	0
15res	Train	605	1013	783	25	205	861	941	1012	0	1	0	0
	Dev	148	249	185	11	53	213	236	249	0	0	0	0
	Test	322	485	317	25	143	432	461	485	0	0	0	0
16res	Train	857	1394	1015	50	329	1197	1307	1393	0	1	0	0
	Dev	210	339	252	11	76	296	319	339	0	0	0	0
	Test	326	514	407	29	78	451	473	513	1	0	1	0

Table 2: Statistics of datasets. #S, #T mean the total number of sentences and triplets. #Pos., #Neu. and #Neg. denote the number of positive, neutral, and negative sentiment triplets respectively. Note that we remove the duplicate ones. Role distributions of spans show the number of the spans playing different roles from three dimensions (aspect, opinion and sentiment snippet), where N means nothing in that dimension, A, O and S denote aspect, opinion and sentiment snippet, respectively.

**Training** The span representations are fed to a fully connected layer (*i.e.*, classifier) with a softmax activation function, generating the probabilities over role tags:

$$p(\mathbf{sp}_{i,j}) = \text{softmax}(\mathbf{W}_r \mathbf{sp}_{i,j} + \mathbf{b}_r), \quad (4)$$

where  $\mathbf{W}_p$ ,  $\mathbf{b}_p$  are parameters of the classifier.

The training loss is defined as the cross-entropy loss between golden labels and predicted label distributions of all spans:

$$L(\theta) = - \sum_{i=1}^n \sum_{j=i}^n \text{loss}(p(\mathbf{sp}_{i,j}), y_{i,j}), \quad (5)$$

where  $y_{i,j}$  represent the gold label for  $SP_{i,j}$ ,  $\text{loss}$  is the standard cross-entropy loss and  $\theta$  represents model parameters.

## Experiment

### Datasets and Metrics

To evaluate our proposed STAGE model, we conduct experiments on four commonly-used datasets released by Xu et al. (2020)<sup>3</sup> that includes one dataset in laptop domain and three datasets in restaurant domain. It is a refined version of the dataset released by Peng et al. (2020) as the latter contains unannotated triplets. Those datasets are derived from the SemEval Challenge (Pontiki et al. 2014, 2015, 2016) and the annotations of the opinion terms are retrieved from Fan et al. (2019). Some detailed statistics are shown in Table 2. The numbers in the A-O-N, A-N-S, N-O-S and A-O-S columns indicate that it is common that a span plays different roles simultaneously in the sentence, implying that a word may have multiple roles and a word-pair may have multiple relations, which contradicts the assumption that previous tagging schemes (Wu et al. 2020; Chen et al. 2022) are based. Following previous works, “Precision( $P$ )”, “Recall( $R$ )” and “ $F_1$  scores ( $F_1$ )” are the evaluation metrics.

<sup>3</sup>ASTE-Data-V2 from <https://github.com/xuuluuuu/SemEval-Triplet-data>

### Implementation Details

The proposed STAGE<sup>4</sup> model contains a BERT-base-uncased model<sup>5</sup> and a fully connected layer as encoder, with hidden state dimension of 768 and 200 respectively. The dropout rate is 0.5. The model is trained in 100 epochs with a batch size of 16. AdamW optimizer (Loshchilov and Hutter 2017) is adopted with a learning rate of  $2 \times 10^{-5}$  for BERT fine-tuning and  $10^{-3}$  for other trainable parameters. For each dataset, we select the model with the best  $F_1$  scores on the development set and report the average results of five runs with different random seeds.

### Baselines

We compare our model with the following baselines:

- Pipeline: **CMLA+** (Wang et al. 2017) considers the aspect-opinion interactions by an attention mechanism, **RI-NANTE+** (Dai and Song 2019) uses a BiLSTM-CRF and models the dependency relations, and **Li-unified-R** (Li et al. 2019) adopts a unified tagging method to jointly extract the aspect term and its sentiment, all modified by Peng et al. (2020) to perform ASTE task. **Peng-two-stage** (Peng et al. 2020) models the interaction between aspects and opinions.
- End-to-End: **OTE-MTL** (Zhang et al. 2020) adopts a multi-task learning framework with a biaffine scorer. **JET-BERT** (Xu et al. 2020) introduces position information into the sequence tagging scheme. **GTS-BERT** (Wu et al. 2020) performs ASTE task by predicting the sentiment relations of word pairs. **Span-ASTE** (Xu, Chia, and Bing 2021) learns span interactions between aspects and opinions. **BMRC** (Chen et al. 2021) views ASTE as an multi-turn machine reading comprehension problem. **ASTE-RL** (Yu Bai Jian et al. 2021) adopts a hierarchical reinforcement learning framework by considering the aspect and opinion terms as arguments of the expressed sentiment. **EMC-GCN** (Chen et al. 2022) proposes to fully utilize the relations between words and diverse linguistic features.

<sup>4</sup>Code is available at <https://github.com/CCIPLab/STAGE>

<sup>5</sup><https://github.com/huggingface/transformers>

Model	14lap			14res			15res			16res		
	<i>P</i>	<i>R</i>	<i>F</i> <sub>1</sub>	<i>P</i>	<i>R</i>	<i>F</i> <sub>1</sub>	<i>P</i>	<i>R</i>	<i>F</i> <sub>1</sub>	<i>P</i>	<i>R</i>	<i>F</i> <sub>1</sub>
CMLA+ <sup>#</sup>	30.09	36.92	33.16	39.18	47.16	42.79	34.56	39.84	37.01	41.34	42.10	41.72
RINANTE+ <sup>#</sup>	21.71	18.66	20.07	31.42	39.38	34.95	29.88	30.06	29.97	25.68	22.30	23.87
Li-unified-R <sup>#</sup>	40.56	44.28	42.34	41.04	67.35	51.00	44.72	51.39	47.82	37.33	54.51	44.31
Peng-two-stage <sup>#</sup>	37.38	50.38	42.87	43.24	63.66	51.46	48.07	57.51	52.32	46.96	64.24	54.21
OTE-MTL <sup>b</sup>	49.53	39.22	43.42	62.00	55.97	58.71	56.37	40.94	47.13	62.88	52.10	56.96
JET-BERT <sup>#</sup>	55.39	47.33	51.04	70.56	55.94	62.40	64.45	51.96	57.53	70.42	58.37	63.83
GTS-BERT <sup>†</sup>	58.62	52.35	55.29	68.14	68.77	68.45	62.37	59.71	60.98	66.16	68.81	67.44
BMRC <sup>b</sup>	70.55	48.98	57.82	75.61	61.77	67.99	68.51	53.40	60.02	71.20	61.08	65.75
ASTE-RL	64.80	54.99	59.50	70.60	68.65	69.71	65.45	60.29	62.72	67.21	69.69	68.41
Span-ASTE <sup>†</sup>	62.74	56.75	59.58	74.17	68.27	71.00	64.15	62.10	63.05	69.31	71.32	70.29
EMC-GCN <sup>†</sup>	59.61	56.30	57.90	70.37	72.84	71.58	60.45	62.72	61.55	63.43	72.63	67.69
STAGE-1D	71.48	53.97	61.49	79.54	68.47	73.58	72.05	58.23	64.37	78.38	69.10	<b>73.45</b>
STAGE-2D	70.56	55.16	<b>61.88</b>	78.51	69.30	<u>73.61</u>	72.33	58.93	<b>64.94</b>	77.67	68.44	72.75
STAGE-3D	71.98	53.86	<u>61.58</u>	78.58	69.58	<b>73.76</b>	73.63	57.90	<u>64.79</u>	76.67	70.12	<u>73.24</u>

Table 3: Experimental results on the test set. The best  $F_1$  results are in bold. The second best  $F_1$  results are underlined. The “<sup>#</sup>” and “<sup>b</sup>” mean that the results are retrieved from Xu, Chia, and Bing (2021) and Chen et al. (2022), respectively. The “<sup>†</sup>” means we reproduce the results using the official released implementation code and configuration.

## Overall Performance

The overall performances are shown in Table 3. The observations are that: (1) Under the  $F_1$  scores, our STAGE models (including two variants) all outperform state-of-the-art baselines by large margins in four datasets. And compared to previous best results, STAGE-3D significantly exceeds their  $F_1$  scores by an average of 2.22%. And two variants (STAGE-2D and STAGE-1D) also exceed by an average of 2.17% and 2.10%, respectively. Considering our model architecture is quite simple, those significant improvements can well verify the effectiveness of our STAGE scheme. (2) Compared with STAGE-3D, the variant STAGE-2D achieves competitive results, even though it cannot deal with the cases that a span is the aspect and opinion simultaneously. It may result from the very limited training samples in those cases (as shown in Table 2), which causes STAGE-3D cannot learn such knowledge well. (3) STAGE-1D exceeds GTS-BERT (Wu et al. 2020) by an average 5.18% under the  $F_1$  score. Considering that the two models have the same tag set (*i.e.*, {A, O, POS, NEU, NEG, N}) but with different tagging schemes, it verifies that *span-level* tagging is superior to *word-level* tagging. Moreover, all our models outperform two state-of-the-art tagging-based baselines GTS-BERT and EMC-GCN significantly and consistently. (4) Although exceeding most baselines, our recall results are slightly lower than the best. But our precision results are better. One reason is that considering the mutual span-level constraints makes the process of retrieving and pairing much stricter during the inference.

## Experiments on Subtasks

To further investigate the effectiveness of STAGE, we conduct experiments on three subtasks of ASTE, that is, aspect term extraction (ATE), opinion term extraction (OTE) and aspect-opinion pair extraction (AOPE). Note that Our method can directly address those subtasks without additional modifications. Specifically, we compare our

Task	Model	14lap	14res	15res	16res
ATE	GTS-BERT	80.44	82.76	78.27	81.13
	EMC-GCN	81.67	84.68	77.62	80.57
	STAGE-1D	<b>82.96</b>	<b>85.24</b>	80.91	83.60
	STAGE-2D	<u>82.64</u>	<u>85.23</u>	<b>81.39</b>	<b>83.95</b>
OTE	STAGE-3D	82.62	84.99	<u>81.21</u>	<u>83.86</u>
	GTS-BERT	77.82	84.85	77.53	84.36
	EMC-GCN	78.85	85.62	78.97	85.33
	STAGE-1D	80.68	<b>86.36</b>	79.47	<u>85.97</u>
AOPE	STAGE-2D	<b>80.82</b>	<u>86.02</u>	<u>79.57</u>	<b>86.51</b>
	STAGE-3D	80.39	85.70	<b>80.03</b>	85.72
	GTS-BERT	66.46	74.63	67.52	74.20
	EMC-GCN	67.94	76.33	67.26	74.15
AOPE	STAGE-1D	<u>69.80</u>	<u>77.47</u>	70.80	<u>79.49</u>
	STAGE-2D	<b>69.98</b>	77.42	<b>71.35</b>	79.41
	STAGE-3D	69.70	<b>77.87</b>	70.60	<b>79.98</b>

Table 4: Test  $F_1$  scores on ATE, OTE and AOPE tasks, with the best in bold and second best underlined.

method with two state-of-the-art tagging-based methods GTS-BERT (Wu et al. 2020) and EMC-GCN (Chen et al. 2022) by using their official released codes and configurations, for a fair comparison. The results are shown in Table 4 and the observations are that: (1) On ATE and OTE task, our models achieve the best results consistently on four datasets, which indicates that considering the complete semantics of each span via span tagging scheme, our method benefits the extraction of aspect and opinion terms. (2) On AOPE task, our models also exceed the compared models by a large margin. It proves that not only our span tagging scheme can better capture aspect/opinion terms, but the greedy inference that considers the span-level constraints can improve the pairing performance as well.

Setting	Model	14lap	14res	15res	16res
<i>Single.</i>	GTS-BERT	63.33	75.67	67.51	72.21
	EMC-GCN	64.88	77.88	68.59	71.64
	Span-ASTE	67.85	77.25	69.61	74.11
	STAGE-1D	70.25	79.67	70.53	<b>78.11</b>
	STAGE-2D	<b>70.71</b>	80.17	<b>71.21</b>	76.98
	STAGE-3D	70.35	<b>80.40</b>	70.69	77.30
<i>Multi.</i>	GTS-BERT	45.16	52.61	50.58	57.27
	EMC-GCN	48.23	57.87	49.37	59.22
	Span-ASTE	49.10	57.02	53.08	62.50
	STAGE-1D	49.73	<b>60.24</b>	54.53	63.40
	STAGE-2D	<b>50.00</b>	59.08	54.70	63.55
	STAGE-3D	49.50	59.20	<b>55.25</b>	<b>64.88</b>
<i>Multi. A.</i>	GTS-BERT	48.22	52.79	54.35	63.29
	EMC-GCN	51.57	<b>59.82</b>	51.98	64.42
	Span-ASTE	52.24	59.22	<b>58.41</b>	66.62
	STAGE-1D	52.87	59.51	56.87	70.26
	STAGE-2D	<b>53.05</b>	58.86	57.98	68.98
	STAGE-3D	52.75	59.20	57.32	<b>71.48</b>
<i>Multi. O.</i>	GTS-BERT	28.44	46.89	40.27	32.98
	EMC-GCN	<b>31.86</b>	47.07	38.52	35.66
	Span-ASTE	29.38	45.02	38.37	<b>41.17</b>
	STAGE-1D	29.30	<b>57.45</b>	45.47	33.26
	STAGE-2D	31.18	53.90	43.55	38.43
	STAGE-3D	28.52	54.34	<b>46.60</b>	37.44

Table 5: Test  $F_1$  scores of ASTE under four settings, with the best in bold. *Single.* denotes that triplets with single-word aspect and opinion. *Multi.* denotes that triplets with multi-word aspects or multi-word opinions. *Multi. A./Multi. O.* denotes triplets with multiple-word aspects/opinions.

## Additional Experiments

We also compare the performance of STAGE with previous models GTS-BERT (Wu et al. 2020), EMC-GCN (Chen et al. 2022) and Span-ASTE (Xu, Chia, and Bing 2021) for the following four settings : 1) Single-Word: Both aspect and opinion in a triplet are single-word terms. 2) Multi-Word: At least one of the aspect or opinion in a triplet is a multi-word term. 3) Multi-word Aspect: Aspect in a triplet is a multi-word term. 4) Multi-word Opinion: Opinion in a triplet is a multi-word term. The results are shown in Table 5. The main observations are that: (1) Generally, compared to the three baselines, our models show superiority under all settings. (2) Compared with Single-Word setting, Multi-Word setting poses challenges in the ASTE task, as the performances of all models drop dramatically. However, considering span-level information, Span-ASTE and our models perform better than GTS-BERT and EMC-GCN, which heavily rely on word-level interactions. It indicates the significance of modeling span-level information in the ASTE task.

## Analysis

### Scheme Generalization

In this part, we extend our STAGE in entity and relation extraction task. Similar to the ASTE task, the span in the given sentence may have two roles: an entity term or a relation

Model	SciERC		
	<i>Ent</i>	<i>Rel</i>	<i>Rel+</i>
UniRE (Wang et al. 2021)	68.4	-	36.9
PURE-S (Zhong and Chen 2021)	66.6	48.2	35.6
PURE-C (Zhong and Chen 2021) <sup>†</sup>	68.9	<b>50.1</b>	36.8
PFN (Yan et al. 2021)	66.8	-	38.4
STAGE	<b>69.1</b>	48.8	<b>38.9</b>

Table 6: Test  $F_1$  scores on SciERC dataset with the best in bold. <sup>†</sup> means using cross-sentence context. *Ent* denotes the entity boundaries evaluation. *Rel* denotes the entity boundaries and relation being correct and *Rel+* denotes the entity boundaries, entity types and the relation being correct.

snippet that includes a pair of entities with a specific relation, forming its boundaries. So, we can define two role dimensions (entity and relation snippet) for entity and relation extraction task. But unlike the ASTE task, the “relation” here has “direction” (*i.e.*, from head entity to tail entity). Thus, we adopt the span tagging scheme with the following tag set:  $\{N, E_1, \dots, E_{|E|}\} - \{N, R_1^{(H-T)}, R_1^{(T-H)}, \dots, R_{|R|}^{(T-H)}\}$ , where  $E = \{E_1, E_2, \dots, E_{|E|}\}$  denotes entity categories,  $R = \{R_1, R_2, \dots, R_{|R|}\}$  denotes relation types,  $N$  means nothing in that role dimension,  $(H - T)$  and  $(T - H)$  discriminate the relation directions as the former indicates head entity appears before the tail entity and the latter shows the opposite, and  $|\cdot|$  represents the set size.

We conduct experiments on SciERC dataset (Luan et al. 2018), a widely-used dataset including 500 scientific articles with annotations on the scientific terms and their relations. We follow the experiment settings of previous work (Zhong and Chen 2021) and utilize micro F1 score ( $F_1$ ) as the evaluation metric. The experimental results in Table 6 show our model also achieves good performance, which verifies the good generalization ability of STAGE.

## Conclusion

In this paper, we propose a novel approach, **Span TAGging and Greedy infERENCE** (STAGE), to solve the ASTE task in an end-to-end manner. To the best of our knowledge, it is the first effort to explore a span-level tagging method and formulate the ASTE task as a multi-class **span** classification problem. Naturally performing on *span-level*, STAGE can overcome the limitations of previous tagging methods and effectively generate more accurate sentiment triplets by exploring span-level information and constraints. Specifically, STAGE consists of two components, namely *span tagging* scheme and *greedy inference* strategy. The former considers the diversity of span roles and tags the span based on three pre-defined role dimensions. The latter considers the mutual span-level constraints and retrieves the aspect/opinion term with the maximum length from the candidate sentiment snippet. Based on STAGE, we propose a simple but effective model, which outperforms the state-of-the-arts by a large margin on four widely-used datasets. Moreover, it can be easily generalized to other pair/triplet extraction tasks, indicating the superiority of the scheme STAGE.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62276110, Grant No.61772076, in part by CCF-AFSG Research Fund under Grant No.RF20210005, and in part by the fund of Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL). The authors would also like to thank the anonymous reviewers for their comments on improving the quality of this paper.

## References

- Chen, H.; Zhai, Z.; Feng, F.; Li, R.; and Wang, X. 2022. Enhanced Multi-Channel Graph Convolutional Network for Aspect Sentiment Triplet Extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2974–2985. Dublin, Ireland: Association for Computational Linguistics.
- Chen, S.; Wang, Y.; Liu, J.; and Wang, Y. 2021. Bidirectional machine reading comprehension for aspect sentiment triplet extraction. In *Proceedings Of The AAAI Conference On Artificial Intelligence*, volume 35, 12666–12674.
- Dai, H.; and Song, Y. 2019. Neural Aspect and Opinion Term Extraction with Mined Rules as Weak Supervision. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5268–5277. Florence, Italy: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 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, Volume 1 (Long and Short Papers)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Fan, Z.; Wu, Z.; Dai, X.-Y.; Huang, S.; and Chen, J. 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, Volume 1*, 2509–2518. Minneapolis, Minnesota: Association for Computational Linguistics.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7871–7880. Online: Association for Computational Linguistics.
- Li, X.; Bing, L.; Li, P.; and Lam, W. 2019. A unified model for opinion target extraction and target sentiment prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 6714–6721.
- Li, X.; and Lam, W. 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*, 2886–2892. Copenhagen, Denmark: Association for Computational Linguistics.
- Liang, S.; Wei, W.; Mao, X.-L.; Wang, F.; and He, Z. 2022. BiSyn-GAT+: Bi-Syntax Aware Graph Attention Network for Aspect-based Sentiment Analysis. In *Findings of the Association for Computational Linguistics: ACL 2022*, 1835–1848. Dublin, Ireland: Association for Computational Linguistics.
- Liu, S.; Li, K.; and Li, Z. 2022. A Robustly Optimized BMRC for Aspect Sentiment Triplet Extraction. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 272–278. Seattle, United States: Association for Computational Linguistics.
- Liu, Y.; Wei, W.; Liu, J.; Mao, X.; Fang, R.; and Chen, D. 2022. Improving Personality Consistency in Conversation by Persona Extending. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management, CIKM '22*, 1350–1359. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392365.
- Loshchilov, I.; and Hutter, F. 2017. Fixing Weight Decay Regularization in Adam. *ArXiv*, abs/1711.05101.
- Luan, Y.; He, L.; Ostendorf, M.; and Hajishirzi, H. 2018. Multi-Task Identification of Entities, Relations, and Coreference for Scientific Knowledge Graph Construction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3219–3232. Brussels, Belgium: Association for Computational Linguistics.
- Mukherjee, R.; Nayak, T.; Butala, Y.; Bhattacharya, S.; and Goyal, P. 2021. PASTE: A Tagging-Free Decoding Framework Using Pointer Networks for Aspect Sentiment Triplet Extraction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9279–9291. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Peng, H.; Xu, L.; Bing, L.; Huang, F.; Lu, W.; and Si, L. 2020. Knowing what, how and why: A near complete solution for aspect-based sentiment analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 8600–8607.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Manandhar, S.; and Androutsopoulos, I. 2015. SemEval-2015 Task 12: Aspect Based Sentiment Analysis. In *\*SEMEVAL*.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Manandhar, S.; and Androutsopoulos, I. 2016. SemEval-2016 Task 5: Aspect Based Sentiment Analysis. In *\*SEMEVAL*.
- Pontiki, M.; Galanis, D.; Pavlopoulos, J.; Papageorgiou, H.; Androutsopoulos, I.; and Manandhar, S. 2014. SemEval-2014 Task 4: Aspect Based Sentiment Analysis. In *COLING 2014*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21(140): 1–67.
- Tang, H.; Ji, D.; Li, C.; and Zhou, Q. 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*, 6578–6588. Online: Association for Computational Linguistics.
- Wang, W.; Pan, S. J.; Dahlmeier, D.; and Xiao, X. 2017. Coupled multi-layer attentions for co-extraction of aspect and opinion terms. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Wang, Y.; Sun, C.; Wu, Y.; Zhou, H.; Li, L.; and Yan, J. 2021. UniRE: A Unified Label Space for Entity Relation Extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 220–231. Online: Association for Computational Linguistics.
- Wang, Z.; Liu, H.; Wei, W.; Hu, Y.; Mao, X.-L.; He, S.; Fang, R.; and Chen, D. 2022. Multi-Level Contrastive Learning Framework for Sequential Recommendation. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management*, CIKM '22, 2098–2107. New York, NY, USA: Association for Computing Machinery. ISBN 9781450392365.
- Wei, W.; Liu, J.; Mao, X.; Guo, G.; Zhu, F.; Zhou, P.; and Hu, Y. 2019. Emotion-Aware Chat Machine: Automatic Emotional Response Generation for Human-like Emotional Interaction. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, CIKM '19, 1401–1410. New York, NY, USA: Association for Computing Machinery. ISBN 9781450369763.
- Wei, W.; Liu, J.; Mao, X.; Guo, G.; Zhu, F.; Zhou, P.; Hu, Y.; and Feng, S. 2021. Target-Guided Emotion-Aware Chat Machine. *ACM Trans. Inf. Syst.*, 39(4).
- Wu, Z.; Ying, C.; Zhao, F.; Fan, Z.; Dai, X.; and Xia, R. 2020. Grid Tagging Scheme for Aspect-oriented Fine-grained Opinion Extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2576–2585. Online: Association for Computational Linguistics.
- Xu, H.; Liu, B.; Shu, L.; and Yu, P. S. 2018. Double Embeddings and CNN-based Sequence Labeling for Aspect Extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 592–598. Melbourne, Australia: Association for Computational Linguistics.
- Xu, L.; Chia, Y. K.; and Bing, L. 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 (Volume 1: Long Papers)*, 4755–4766. Online: Association for Computational Linguistics.
- Xu, L.; Li, H.; Lu, W.; and Bing, L. 2020. Position-Aware Tagging for Aspect Sentiment Triplet Extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2339–2349. Online: Association for Computational Linguistics.
- Yan, Z.; Zhang, C.; Fu, J.; Zhang, Q.; and Wei, Z. 2021. A Partition Filter Network for Joint Entity and Relation Extraction. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 185–197. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Yu Bai Jian, S.; Nayak, T.; Majumder, N.; and Poria, S. 2021. Aspect sentiment triplet extraction using reinforcement learning. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 3603–3607.
- Zhang, C.; Li, Q.; and Song, D. 2019. 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)*, 4568–4578. Hong Kong, China: Association for Computational Linguistics.
- Zhang, C.; Li, Q.; Song, D.; and Wang, B. 2020. A Multi-task Learning Framework for Opinion Triplet Extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, 819–828. Online: Association for Computational Linguistics.
- Zhang, W.; Li, X.; Deng, Y.; Bing, L.; and Lam, W. 2021. Towards Generative 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 Processing (Volume 2: Short Papers)*, 504–510. Online: Association for Computational Linguistics.
- Zhao, H.; Huang, L.; Zhang, R.; Lu, Q.; and Xue, H. 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*, 3239–3248. Online: Association for Computational Linguistics.
- Zhao, S.; Wei, W.; Zou, D.; and Mao, X. 2022. Multi-view intent disentangle graph networks for bundle recommendation. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*.
- Zhong, Z.; and Chen, D. 2021. A Frustratingly Easy Approach for Entity and Relation Extraction. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 50–61. Online: Association for Computational Linguistics.
- Zou, D.; Wei, W.; Mao, X.-L.; Wang, Z.; Qiu, M.; Zhu, F.; and Cao, X. 2022. Multi-Level Cross-View Contrastive Learning for Knowledge-Aware Recommender System. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '22, 1358–1368. New York, NY, USA: Association for Computing Machinery. ISBN 9781450387323.