

# OATS: Opinion Aspect Target Sentiment Quadruple Extraction Dataset for Aspect-Based Sentiment Analysis

**Siva Uday Sampreeth Chebolu, Franck Dernoncourt, Nedim Lipka, Tamar Solorio**

University of Houston, Adobe Research, Adobe Research, University of Houston & MBZUAI  
4800 Calhoun Rd, Houston, TX 77004, USA, 325 Park Ave San Jose, CA 95110, USA,  
325 Park Ave San Jose, CA 95110, USA, 4800 Calhoun Rd, Houston, TX 77004, USA, Masdar City, UAE  
sivauday.sampreeth8@gmail.com, franck.dernoncourt@gmail.com,  
lipka@adobe.com, tsolorio@uh.edu

## Abstract

Aspect-based sentiment analysis (ABSA) delves into understanding sentiments specific to distinct elements within a user-generated review. It aims to analyze user-generated reviews to determine a) the target entity being reviewed, b) the high-level aspect to which it belongs, c) the sentiment words used to express the opinion, and d) the sentiment expressed toward the targets and the aspects. While various benchmark datasets have fostered advancements in ABSA, they often come with domain limitations and data granularity challenges. Addressing these, we introduce the OATS dataset, which encompasses three fresh domains and consists of 27,470 sentence-level quadruples and 17,092 review-level tuples. Our initiative seeks to bridge specific observed gaps in existing datasets: the recurrent focus on familiar domains like restaurants and laptops, limited data for intricate quadruple extraction tasks, and an occasional oversight of the synergy between sentence and review-level sentiments. Moreover, to elucidate OATS's potential and shed light on various ABSA subtasks that OATS can solve, we conducted experiments, establishing initial baselines. We hope the OATS dataset augments current resources, paving the way for an encompassing exploration of ABSA (<https://github.com/RiTUAL-UH/OATS-ABSA>).

**Keywords:** Aspect-based Sentiment Analysis, Aspect based Sentiment Analysis Dataset, ABSA, ABSA Dataset, ASQP, ACOS, ASTE, Quadruple Extraction Dataset

## 1. Introduction

The trend in analyzing online reviews has shifted from extracting a broad understanding of consumer opinions on overall product performance to a more granular examination of individual product aspects. This shift demands a different approach to analyzing reviews. Aspect-based sentiment Analysis (ABSA) emerged as the answer to this nuanced requirement, focusing on sentiment pertaining to specific aspects of an entity (Hu and Liu, 2004a). However, current datasets often fall short of capturing the complete spectrum of ABSA. The main goal of ABSA is to identify the target of the opinion, the aspect category it belongs to, the opinion phrase, and the sentiment polarity associated with the opinion. One of the significant limitations is the inability to perform joint detection of all ABSA elements due to the absence of at least one critical component in the human-annotated reviews. This limitation stunts the potential of ABSA tasks. Despite the SemEval datasets' popularity, many of their sentences often group multiple aspects under a single sentiment polarity. As noted by Jiang et al. (2019), this approach simplifies the ABSA task, reducing it to mere sentence-level sentiment analysis. As a solution, they proposed a new large-scale Multi-Aspect Multi-Sentiment dataset, where each phrase has at least two independent aspects

with different sentiment polarity. However, their introduction of a "miscellaneous" category or neutral sentiment when a sentence contains only one opinion tuple is impractical. This makes the task more challenging, but it is not realistic.

Recently released ABSA datasets such as ACOS (Cai et al., 2021) and ASQP (Zhang et al., 2021a) provide comprehensive annotations for all four elements. However, they are largely limited to the long-standing traditional domains of restaurants and laptops, a trend that has been observed since 2014. In contrast, newer datasets (i.e., Xu et al. (2023); Chia et al. (2023)) address domain diversity by including home appliances, fashion, groceries, and more, moving beyond the typical restaurant and laptop reviews. Yet, these innovative datasets lack the critical aspect category annotations required for a holistic joint detection of all elements.

Several crucial insights emerge in delving deep into the landscape of ABSA datasets (Chebolu et al., 2023). Firstly, there needs to be more consistency in how ABSA components are defined and structured across various sources, accentuating an urgent need for a universal standardized format. Secondly, the challenge of accurately detecting opinion words, which are central to discerning sentiment polarity and their corresponding aspect terms, still needs to be solved. For in-

Dataset	Tgt. Span	Opi. Ph. Span	Asp-Sen	Quadruple
Amazon_FF	72.57	69.72	85.43	65.78
Coursera	78.26	71.26	79.63	68.56
Hotels	74.78	72.05	87.32	73.84

Dataset	#sents	#pos	#neg	#neu	#Total
Rest-15	1,562	1,710	701	85	2,496
Rest-16	2,024	2,293	877	125	3,295
<b>Total</b>	<b>3,586</b>	<b>4,003</b>	<b>1,578</b>	<b>210</b>	<b>5,791</b>

Table 1: **(a) Left:** Inter-Annotator agreement F1-scores for the OATS Datasets. Tgt. Span: Target span extraction. Opi. Ph. Span: Opinion Phrase span extraction. Asp-Sen: Aspect Category and Sentiment combination categorization. **(b) Right:** Current ASQP Dataset Statistics

stance, in the hotel review *'the room was spotless and large, bathroom and kitchen were fine for our needs and the king size bed was great'*, the word 'spotless' indicates a positive sentiment polarity for the aspect 'rooms cleanliness', while 'large' pertains to 'rooms design\_features'. Correctly associating these opinion words with their respective aspects and targets is needed. Thirdly, the current datasets, while extensive, do not always reflect real-world complexities. Different domains, such as healthcare, education, or entertainment, have their unique terminologies, sentiment expressions, and contextual nuances. Moreover, the prevailing focus on English in ABSA datasets sidelines many low-resource languages, neglecting challenges tied to idiomatic expressions, linguistic structures, and cultural contexts inherent to these languages. This need is further emphasized by recent advancements in unified models using generative frameworks (Chebolu et al., 2021; Zhang et al., 2021b), highlighting the immense potential of models that jointly solve ABSA tasks. These models necessitate datasets with richer annotations, similar to (Zhang et al., 2021a). Lastly, current ABSA research largely adopts a sentence-centric approach, risking misinterpretation by neglecting the inter-sentence context in reviews. While one might suggest merging sentence-level predictions, this doesn't remedy the foundational issue: sentiments in one sentence can be influenced by others. Without the full review context, even merged predictions can be off-mark. Although (Xu et al., 2023) contains complete reviews of up to 250 words, they address only the ASTE task. Thus, there's a clear need for additional review-level ABSA datasets.

Given these challenges, the OATS dataset has been developed with a vision to rectify most of the existing gaps. The contributions of OATS are:

- Introduces both sentence-level quadruples and review-level tuples, capturing sentiments in all their granularity.
- The dataset spans multiple domains, ensuring its applicability across a diverse range of ABSA tasks.
- OATS will be publicly available in two formats: XML (as introduced by Pontiki et al. (2014) catering to detailed character-level annotations), and Text (following the format set

by Zhang et al. (2021a)).

- We provide baselines for our dataset, focusing on four primary tasks (ASD, ASTE, TASD, and ASQP) from the 14 ABSA subtasks.

## 2. Related Work

In this section, we describe the tasks of ABSA and their related datasets, followed by triplet and quadruple extraction task-specific methods.

### 2.1. Current Datasets

Traditionally, when ABSA was first introduced by Hu and Liu (2004b) as a task in NLP, the main aim then was to extract different aspects given a sentence followed by assigning polarities given those aspects. As a decade passed, Pontiki et al. (2014) added another important element of ABSA: aspect categories. In the SemEval-2014 shared task on ABSA, two sub-tasks were drafted in addition to the aspect term extraction and its sentiment polarity. They are the aspect category detection and their sentiment polarity. Two datasets from the restaurant and laptop domains were released as part of this shared task.

In the consecutive years, Pontiki et al. (2015, 2016) proposed two more ABSA shared tasks at the SemEval, redefining the subtasks of ABSA and giving them a more concrete structure. Aspect categories were a combination of entity and attribute pairs, aspect terms were called targets, and the sentiment polarity was assigned jointly for the targets and the aspect categories. Several datasets with the three elements from multiple languages and domains were published, including English, Dutch, Spanish, French, Russian, and Turkish restaurants, English laptops, Arabic hotels, and many others.

Later, Wan et al. (2020) utilized the data from the SemEval-2015 and SemEval-2016 to define the target-aspect-sentiment joint detection task (TASD). After that, researchers felt the importance of detecting opinion phrases to identify the sentiment polarity and understand its relationship with targets and aspect categories. This resulted in three new sub-tasks with new datasets that stemmed from the SemEval challenges, namely aspect-opinion-pair extraction

(AOPE) (Chen et al., 2020), target-opinion-word extraction (TOWE) (Fan et al., 2019), and aspect-sentiment-triplet-extraction (ASTE) (Peng et al., 2020). More recently, the ABSA research has shifted towards the joint detection of all four elements of ABSA, which proved to better identify the inter-relationships among the elements, thereby enhancing the performance of other sub-tasks. This task is called aspect sentiment quadruple prediction (ASQP) (Zhang et al., 2021a) or aspect-category-opinion-sentiment joint detection (ACOS) (Cai et al., 2021). The ASQP task is evaluated on the dataset with quadruples created from the SemEval challenges, ASTE triplets, and TOWE and AOPE tuples. The ACOS task introduced two datasets from the restaurant and the laptop domain with implicit and explicit aspect terms and opinion phrases.

The current ABSA dataset landscape is notably skewed towards particular domains such as restaurants and electronics. This stems from the easy availability and volume of data on review websites and other online platforms. This narrow focus restricts the wider application of ABSA, limiting its ability to deliver insights across different sectors. Expanding the diversity of ABSA datasets is essential to push research and applications beyond the boundaries of mere reviews. To address this limitation, we propose three new datasets for the quadruple extraction task from three new domains, along with implicit and explicit targets and opinion phrase annotations. We also include the review-level tuples for identifying the overall sentiment polarity for different aspect categories in a review. This dataset could be used to solve all the above-mentioned subtasks of ABSA.

## 2.2. Related Methods

We focus on a few joint detection tasks: the TASD, ASTE, and ASQP for our research. Chebolu et al. (2023) have provided a detailed survey of the other datasets, tasks, and their challenges.

**Triplet Extraction Tasks** In ASTE research, three primary methodologies have emerged: MRC-based techniques (Liu et al., 2022), methods anchored on BERT and table-filling (Zhang et al., 2022), and generative approaches (Peper and Wang, 2022; Zhang et al., 2021c,a). The MRC-based methods involve crafting a specific query for each component in the triplet, subsequently extracting them based on the response to this query. Generative methods, in contrast, frame the ASTE challenge as a sequence generation task and employ sequence-to-sequence (seq2seq) models. The triplets are then decoded using a specially tailored algorithm. In this study, we employ representative techniques from these three categories and investigate OATS using

these methods.

For the TASD task, there are two paths: BERT-based methods (Wan et al., 2020), and generative-based approaches (Chebolu et al., 2021; Zhang et al., 2021a). The BERT-based methods extract the aspect sentiment tuple from a sentence using the sentence-pair classification as a backbone (Sun et al., 2019) followed by extracting the targets for each pair using the token classification with the BIO or a softmax classifier. The generative-based approaches convert the task similar to an abstractive summarization and use the sequence-to-sequence models like the T5, BART, and others (Raffel et al., 2020; Lewis et al., 2019) to predict the triplets.

**Quadruple Extraction Task** Researchers have pointed out two promising directions. Cai et al. (2021) propose a two-stage method by extracting the aspect and opinion terms first. Then, these items are utilized to classify aspect categories and sentiment polarity. Another method is based on the generation model (Zhang et al., 2021a). By paraphrasing the input sentence, the quadruplet can be extracted end-to-end. In this work, we follow the generative direction and consider the order-free property of the quadruplet. A few other studies also proposed generative-based models with this paraphrasing as a backbone, including Hu et al. (2022); Peper and Wang (2022).

## 3. OATS Dataset

In this section, we start by discussing the sources for the OATS dataset. The following subsections discuss the annotation process, including inter-annotator agreement, relevant data statistics, and OATS relevance to ABSA and the wider NLP community.

We derived three distinct English review datasets from multiple sources. Our corpus is primarily built from text reviews available on Kaggle (consistent with the CCO and ODbL licenses) and resources provided by Yin et al. (2017).

**Amazon Fine Foods Dataset** This dataset, referred as Amazon\_FF in this work, is extracted from a Kaggle competition<sup>1</sup> containing around 500k reviews of fine foods from Amazon. These reviews span topics such as the quality of food or products, promptness of delivery, packaging standards, and product availability, among others. We curated 1,794 complete reviews from this dataset, consisting of over 8,900 sentences and approximately 8,200 opinion quadruples.

---

<sup>1</sup><https://www.kaggle.com/snap/amazon-fine-food-reviews>

Domain	#Revs.	#Sent.	#Rev.Op.	#Sent.Op.	#Total Op
Amazon_FF	1,794	8,913	4,326	8,260	12,586
Coursera	1,702	8,278	5,350	7,875	13,225
Hotels	1,497	7,963	7,416	11,335	18,751
<b>Total</b>	<b>4,993</b>	<b>25,154</b>	<b>17,092</b>	<b>27,470</b>	<b>44,562</b>

Stats/Domain	Amazon_FF	Coursera	Hotels
Avg. Sentences/Review	4.96	4.81	5.31
Avg. Length of Sentence	71.34	79.02	75.95
Avg. Length of Review	359.9	391.55	405.29
Avg. Opinions/Sentence	0.92 (1.25)	0.95 (1.27)	1.42 (1.71)
Avg. Opinions/Review	2.4 (2.48)	3.16 (3.18)	4.96 (5.37)

Table 2: **(a) Left:** OATS Dataset Overall Statistics **(b) Right:** OATS Average Statistics. The values inside the () are the average opinions per sentence/review, excluding those with zero opinions.

**Coursera Dataset** Originating from a Kaggle competition, this dataset contains reviews scraped from the Coursera website, totaling close to 100k reviews. These reviews reflect diverse perspectives on the course’s quality, content, comprehensiveness, and the alignment of faculty lessons with course content. For our work, we selected 1,702 comprehensive reviews, which include roughly 8,200 sentences and about 7,800 opinion quadruples.

**TripAdvisor Dataset** This dataset is based on data from Yin et al. (2017), featuring over 100k hotel reviews. The reviews comment on various facets of the hotel experience, such as pricing, design attributes, and more. We collated 1,497 complete reviews from this extensive collection, resulting in approximately 8,000 sentences and nearly 11,300 opinion quadruples.

### 3.1. Annotation Procedure

We relied on Upwork to identify three freelancers for the data annotation and employed the BRAT tool for annotation. We provided a comprehensive guideline document introducing them to the topic, dataset, the BRAT tool, and specific annotation requirements. Initially, annotator (A) took the lead in annotating a subset of the data (50 complete reviews), which was subsequently reviewed by another annotator (B) for corrections. This process was iterated twice, with pairs of annotators (A, C) and then (B, C) for checks and balances. We used 150 reviews (50 for each annotator) from each domain for this phase of the agreement task. Any emerging disagreements were addressed and resolved through consultations with one of our NLP experts, which are added as additional instructions to the annotators for future reviews.

This process is repeated multiple times until a decent inter-annotator agreement is reached among the three annotators. Then, the remaining dataset reviews were equally divided among the three annotators for the actual annotation process. When all three annotators held differing views, a consensus was achieved in collaboration with an expert annotator. Gold annotations are generated via majority voting. To ensure the high-quality of the annotations, for every 150 reviews, each annotator’s work is reviewed by another annotator. If

the agreement score is less than the initial score, that part of the data is re-annotated again until the score crosses the baseline.

According to Hripcsak and Rothschild (2005), the infamous Kappa metric (Tomanek and Hahn, 2009; Becker et al., 2005) may not be the best fit for span-extraction annotation in textual data to measure the inter-annotator agreement. The limitation arises from the requirement of Kappa to compute the number of negative cases, which is unidentifiable for spans as they constitute sequences of words without a predetermined quantity of items for annotation in a text. The F-measure is often more suitable for gauging inter-annotator agreement in span extraction annotation tasks such as target and opinion phrase extraction (Deléger et al., 2012). We can calculate the F1-score by considering one annotator’s annotations as the reference (or gold annotations or groundtruth) and the other’s as a system’s responses (or predictions).

We computed the inter-annotator agreement scores, following this F1 metric, for each domain using several combinations of the ABSA elements as shown in Table 1(a). The average quadruple extraction agreement F1 score for the three domains is 69.39%. When we observe the simple task of aspect category and sentiment polarity identification, the average agreement F1-score is 84.12%. If we go to challenging ones that include span extraction like the target and opinion phrase extraction, This score is significant in light of the inherent subjectivity and complexity involved in jointly identifying all four ABSA elements. This highlights the considerable level of consistency among our annotators and, on the other hand, indicates the task’s complexity.

### 3.2. The OATS dataset in Numbers

In Table 2 (a), we present some basic statistics of each domain in the datasets. It includes information about the total number of reviews, sentences, review-level opinion tuples, and sentence-level opinion quadruples for every domain. Table 2 (b) has the averaged statistics, such as the sentences per review, length of sentence (and review), and the number of opinions per sentence (and review) with the domains as columns. The numbers provided inside “()” for the number of opin-



Domain	Sentence-Level					Review-Level				
	0-Op	1-Op	2-Op	3-Op	>3-Op	0-Op	1-Op	2-Op	3-Op	>3-Op
Amazon_FF	1,957	4,529	890	170	42	47	175	626	490	183
Coursera	1,395	3,685	690	134	36	8	107	298	349	449
Hotels	1,454	2,910	1,167	499	337	116	11	58	138	883
<b>Total</b>	<b>4,806</b>	<b>11,124</b>	<b>2,747</b>	<b>803</b>	<b>415</b>	<b>171</b>	<b>293</b>	<b>982</b>	<b>977</b>	<b>1,515</b>

Table 3: OATS Dataset Statistics for the total number of sentences/reviews with the respective number of opinions

Domain	Sentence-Level			Review-Level			
	#pos	#neg	#neu	#pos	#neg	#neu	#conf
Amazon_FF	5,577	1,187	234	2,900	606	74	82
Coursera	4,403	1,008	213	2,910	721	129	71
Hotels	6,952	1,207	169	4,557	817	110	53
<b>Total</b>	<b>16,932</b>	<b>3,402</b>	<b>616</b>	<b>10,367</b>	<b>2,144</b>	<b>313</b>	<b>206</b>

Table 4: OATS Dataset Statistics for the total number of positive, negative, neutral, and conflict sentiment polarities observed for the aspect categories at Review-Level and the aspect terms, aspect category, and opinion phrase triplets at the Sentence-Level in each domain.

ions per sentence and review are computed after excluding the zero opinion sentences and reviews. We also provide fine-grained statistics in Table 3, such as the number of sentences and reviews with zero, one, two, three, and more than three opinions. The stats corresponding to the number of opinions with different sentiment polarities, including positive, negative, neutral, and conflict at both the sentence and review levels, are presented in Table 4.

The hotels domain has a higher number of opinions per sentence as well as per review when compared to the other two domains. Despite having the most reviews (from Table 2 (a)), the fine foods domain has the least number of opinions per sentence and review among the three. One of the main reasons is that the number of zero opinion sentences in the fine foods domain is higher than in the hotels and Coursera (from Table 3). Furthermore, the ratio of two-opinion sentences to one-opinion sentences for the Amazon\_FF and the Coursera domains ( $\approx 1 : 5$ ) pushed the average numbers closer to one opinion per sentence. We initially filtered the reviews for this dataset with a requirement of covering a minimum of 2 distinct aspect categories at the review level is clearly evident from the average number of opinions per review. However, it was lacking at the sentence level.

Another important factor for ABSA datasets is the distribution of aspect categories. Figure 1 shows that distribution for amazon\_FF domain. The *General* attribute for different entities in all the datasets has the highest number of opinions, such as the *FOOD#GENERAL* in

Distribution of Aspect Categories for Amazon\_FF in %




Figure 1: Distribution of Aspect Categories for Sentence-Level and Review-Level Annotations for Amazon\_FF Domain

the Amazon\_FF, *COURSE#GENERAL* and *FACULTY#GENERAL* for the Coursera domain, and *HOTEL#GENERAL*, *LOCATION#GENERAL*, and *SERVICE#GENERAL* in the hotel domain. The review-level and the sentence-level opinions have similar distributions for the aspect categories. We observed a similar distribution among the aspect categories in the other domains, so we chose to present only one domain plot for representation.

## 4. Experiments

This section will first list the tasks we conduct experiments on, followed by the baseline methods for each task. We divide the methods into task-specific and unified approaches. Finally, we discuss the evaluation metrics for the tasks.

### 4.1. Tasks

We selected three major joint detection tasks for our sentence-level ABSA experiments: target-aspect-sentiment detection (TASD) (Wan et al., 2020), aspect sentiment triplet extraction (ASTE) (Peng et al., 2020), and aspect sentiment quadruple prediction (ASQP) (Zhang et al., 2021a). The rationale for choosing these tasks is that they encompass all four elements of ABSA in different combinations. Our emphasis on joint extraction tasks over single-element extraction is motivated not just by the findings of Chebolu et al.

Method	Amazon_FF	Coursera	Hotels
GAS-T5	19.62	<b>22.23</b>	26.33
Paraphrase	20.84	19.78	<b>34.51</b>
Template-ILO	<b>21.01</b>	21.24	24.62
Template-DLO	20.39	21.96	26.59
GEN-SCL-NAT	20.36	20.20	28.58

Method	Amazon_FF	Coursera	Hotels
BDTF (BERT)	27.81	43.98	42.02
BMRC	<b>49.48</b>	<b>50.85</b>	<b>50.71</b>
GEN-SCL-NAT	20.36	20.20	28.58
GAS	46.46	40.15	45.96
Paraphrase	46.64	39.30	46.37

Table 5: **(a) Left:** F1 scores of ASQP task on OATS **(b) Right:** F1 scores of ASTE task on OATS. Best results are highlighted in bold font.

(2021), but also by broader research trends in the NLP community. For instance, the DREAM paper on entity-relation extraction successfully illustrates the advantages of jointly modeling evidence ranking, leading to enhanced performance (Wang et al., 2023). Such joint models, tailored to manage multiple intertwined elements simultaneously, have consistently been shown to surpass models fine-tuned for single-element extraction in various domains.

For the review-level (or text-level) ABSA, where we have to extract aspect category and sentiment tuples, we chose to use the aspect sentiment joint detection task (ASD) (Sun et al., 2019) for our experiments. The two main challenges with the review-level task are that we have to use the entire review context rather than simply using each sentence to detect the aspect categories and the sentiment polarity of the review. Secondly, there can be multiple opinions on the same aspect category with multiple polarities, from which we must assign the dominating sentiment to each category. In the review, “The room was mostly spotless and the bathroom pristine, but there was dust on the bedside table,” two positive sentiments outweigh a milder negative one for the room’s cleanliness, leading to an overall positive polarity.

## 4.2. Baseline Methods

We implement several representative models from various frameworks, including MRC-based approaches similar to Xu et al. (2021); Mao et al. (2021); Chen et al. (2021), generation-based, and BERT-based frameworks (Devlin et al., 2019), for the task evaluations.

**Task-specific methods** For *ASTE*, we use the following methods:

**BMRC** (Liu et al., 2022): a MRC-based method. It extracts aspect-oriented triplets and opinion-oriented triplets. Then, it obtains the final results by merging the two directions.

**BDTF (BERT)** (Zhang et al., 2022): a BERT-based method that uses table-filling to solve the problem. It transforms the *ASTE* task into detecting and classifying relation regions in the 2D table representing each triplet in addition to an effective

relation representation learning approach to understand word and relation interactions.

These are the methods for **ASQP** task: **Template-ILO** (Hu et al., 2022): a generative-based method. Similar to Zhang et al. (2021a) with an additional step to identify the best permutation of the four ABSA elements at the *instance level*, and further combine multiple proper templates as data augmentation, instead of having a fixed order template, that is passed as input to the generative model.

**Template-DLO** (Hu et al., 2022): a generative-based method. Similar to Zhang et al. (2021a), but has an additional step to identify the best permutation of the four ABSA elements at the *dataset level*, and further combine multiple proper templates as data augmentation, instead of having a fixed order template, that is passed as input to the generative model.

**TASD Task-specific methods** For the *TASD* task, we chose **TAS-BERT** (Wan et al., 2020), which is a BERT-based method. It fine-tunes the pre-trained BERT model to solve the aspect-sentiment detection task using the classification token and then detects the targets corresponding to those tuples using the token classification with CRF/softmax decoding.

**ASD Task-specific methods** The below methods are used for both review-level (R-ASD) and sentence-level (S-ASD) tasks:

**BERT-pair-NLI-B** (Sun et al., 2019): a BERT-based model that takes the review context as segment A and predicts the presence of the combination of aspect category and sentiment polarity, which is taken as an entailment in the segment B of BERT’s input, using the output from [CLS] token.

**BERT-pair-QA-B** (Sun et al., 2019): a BERT-based model that uses question-answering format instead of entailment in BERT-pair-NLI-B.

**QACG-BERT** (Wu and Ong, 2020): developed a CG-BERT, which utilizes context-guided (CG) softmax-attention by initially modifying a context-aware Transformer. Subsequently, they introduced an improved Quasi-Attention CG-BERT model that learns a compound attention mechanism, facilitating subtractive attention.

Method	Amazon_FF	Coursera	Hotels
TAS-BERT-BIO	45.12	44.41	45.92
TAS-BERT-TO	47.51	42.77	45.76
T5-ABSA	<b>51.61</b>	<b>44.57</b>	49.78
GAS	43.04	41.53	<b>50.69</b>
Paraphrase	44.89	40.24	49.81

Method	Amazon_FF		Coursera		Hotels	
	R-ASD	S-ASD	R-ASD	S-ASD	R-ASD	S-ASD
BERT-pair-NLI-B	88.22	93.25	91.31	95.92	91.84	96.82
BERT-pair-QA-B	<b>88.42</b>	<b>93.61</b>	<b>91.51</b>	<b>96.69</b>	<b>91.93</b>	<b>97.18</b>
QACG-BERT-NLI-M	87.15	92.13	90.41	95.81	90.05	95.68
T5-ABSA	58.26	56.85	40.76	48.81	54.09	59.85
GAS-T5	65.23	57.57	46.61	56.79	58.61	66.9
Paraphrase	66.46	57.25	47.88	55.27	56.11	66.38

Table 6: **(a) Left:** F1 scores of in-domain TASD on OATS. **(b) Right:** Results for Aspect-Sentiment Joint Detection task at Review-level (R-ASD) and Sentence-level (S-ASD) on OATS datasets. Best results are highlighted in bold font.

**T5-ABSA** (Chebolu et al., 2021): a generative approach that takes the review context along with a task prefix as input to generate the respective task outputs in an auxiliary sentence-based or phrase-based format defined in their work.

**Unified methods** The following generative frameworks can be applied to any sentence-level and review-level ABSA tasks:

**GEN-NAT-SCL** (Peper and Wang, 2022): a generative-based method. It combines a new generative format with a supervised contrastive learning objective to predict the ASTE’s triplets and ASQP’s quadruples.

**GAS** (Zhang et al., 2021c): a generation-based method. It transforms the ASTE, ASQP, TASD, and ASD tasks into a text generation problem that inputs the review and generates all the respective combinations of ABSA opinion elements as output.

**Paraphrase** (Zhang et al., 2021a): a generative-based method. It is similar to GAS but transforms output opinion elements into paraphrases that read as natural sentences. We substituted the corresponding element for implicit targets and opinions with the word "it."

### 4.3. Evaluation

Following Zhang et al. (2021a); Xu et al. (2023); Wan et al. (2020), we use the F1 score to measure the performance of different approaches on all the tasks from Section 4.1. All experimental results are reported using the average of 5 different runs using distinct random seeds. We divided each domain dataset into train, validation, and test sets with 80%, 10%, and 10% splits, respectively. A tuple, triplet, and quadruple is considered correct only if all the corresponding prediction elements match the gold standard labels. We consider any partial matches as wrong predictions following Zhang et al. (2021a). We adopt the base versions of all the transformer models for our experiments, including the BERT-base (Devlin et al., 2019), T5-base (Raffel et al., 2020), RoBERTa-base (Liu et al., 2019), and others.

## 5. Results and Discussion

In this section, we present the results of different baseline systems on the selected tasks, highlighting the benefits and challenges of current methods on the OATS dataset and the specific task.

**ASQP** Table 5 (a) reports the performance on the ASQP task for the quadruple extraction using five different baseline systems on the three datasets. On a high level, the overall performance of the hotels domain is better compared to the other two domains in all the baseline systems.

For the Amazon\_FF domain, even though Template-ILO gave the best performance, all the other methods are not far from it. Similarly, for the Coursera domain, GAS-T5 performed better than the Paraphrase and GEN-SCL-NAT, but the Template method is close to its performance. Paraphrase outperformed all the other methods in the hotels domain with a 34.51% F1 score.

Prior works have shown that the Template method significantly outperformed Paraphrase and GAS in the restaurant domain (Hu et al., 2022). That trend only translated to the Amazon\_FF domain, leaving out Coursera and Hotels. Similarly, GEN-SCL-NAT, which outperformed the Paraphrase method on the restaurants and laptops domain, failed to carry it to any of the three domains.

Domain/Task	TOWE	TSD	ASD	TASD	ASTE	ASQP
Amazon_FF	31.29	65.27	57.25	44.89	46.64	20.84
Coursera	38.64	67.75	55.27	40.24	39.30	19.78
Hotels	35.75	62.34	66.38	49.81	46.37	34.51

Table 7: Performance Analysis on six main ABSA joint tasks using the Paraphrase method (Zhang et al., 2021a) on OATS

**ASTE** In Table 5 (b), we present the performance of five baseline systems on the ASTE task, which focuses on extracting the triplet of aspect term, opinion term, and sentiment polarity across three datasets. Notably, the BMRC method emerges as the standout performer, surpassing other baselines across all domains for this task, while the GEN-SCL-NAT method consistently lags behind. The success of the BMRC method can be attributed to the approach taken by Liu et al. (2022). Instead of solely considering the relation-

Domain	ET / IT	EO / IO	ET-EO	ET-IO	IT-EO	IT-IO
Amazon_FF	2,999 / 5,261	6,780 / 1,480	2,491	508	4,298	972
Coursera	5,163 / 2,712	6,185 / 1,690	4,222	941	1,963	749
Hotels	8,654 / 2,820	10,193 / 1,281	7,927	727	2,266	554

Table 8: Distribution of Explicit and Implicit Targets (ET and IT), and Explicit and Implicit Opinion Phrases (EO and IO) in the OATS datasets with their different combinations.

ship from target to opinion phrases, they also incorporate the potential backward relation from opinion phrases to target expressions. This deviates from other methods like Mao et al. (2022) and Bao et al. (2022), which deploy two distinct paths in a tree structure for targets and opinion phrases to identify the quadruples.

**R-ASD** Table 6 (b) shows the aspect sentiment joint detection task results at the review-level (R-ASD) and sentence-level (S-ASD). BERT-based methods significantly outperform the generative approaches. One primary factor driving this difference is the augmentation strategy BERT-pair methods use, which expands the dataset based on the number of categories multiplied by the polarities in that dataset. For example, with 27 aspect categories and 4 different polarities in the Coursera domain for the ASQP task, this results in a total of 894024 instances ( $8278 * 27 * 4$ ). Additionally, these methods transform the multi-label, multi-class classification of the ASD task into a simpler binary classification, making the task inherently easier than the challenge generative models face when extracting categories and polarities directly from a review sentence.

**TASD** We present the results of the TASD task experiments on all the domains in Table 6 (a). The generative-based models performed better than the BERT-based approaches on all the domain datasets. The T5-ABSA method performed best for Amazon\_FF and Coursera domains, while the GAS method outperformed the hotels domain. All the generative-based approaches are not far from each other in the hotels domain, while the BERT-based methods have a significant difference in the Amazon\_FF and hotels domain when compared to the generative approaches.

#### Explicit and Implicit Targets and Opinions

We observe distinct patterns of explicit and implicit targets and opinions across the OATS domains from the statistics presented in Table 8. Amazon\_FF, for instance, showcases a predominance of implicit targets, suggesting users often leave their main subject of discussion implied. Conversely, Coursera leans more towards explicit target mentions, possibly indicative of the direct nature of feedback in educational settings. Notably, regardless of the domain, explicit opinions outnumber their implicit counterparts. This trend under-

lines that while users might leave their subjects (targets) implied, they tend to be direct about their sentiments or opinions on them. The Hotels domain stands out with the highest counts of explicit mentions for both targets and opinions, hinting at the straightforward nature of feedback in this domain. It is important to recognize these domain-specific trends as they aid in tailoring models appropriately, potentially improving their accuracy and adaptability.

**Challenging combination of elements in OATS** We tested the Paraphrase-T5 method (Zhang et al., 2021a) on all the OATS datasets for six different ABSA tasks to see how hard it is to detect various ABSA elements. The results are in Table 7. The TSD task did the best on all the domains, but TOWE was the hardest. This shows that connecting an opinion phrase to its target is especially tough in OATS. Interestingly, even though TOWE is kind of a part of ASTE, it does not do as well. This might be because having sentiment information in ASTE helps pick out opinion phrases better, giving it an edge.

Furthermore, integrating aspect categories into TSD creating the TASD task, saw a noticeable performance drop. This highlights the challenge of pinpointing aspect categories for specific targets. This challenge could also explain the lower scores of the ASD and ASQP tasks in the Amazon\_FF and Coursera domains compared to TSD. Interestingly, in the hotels domain, ASD emerged as the top-performing task, which contributed to a superior TASD and ASQP performance relative to other domains. This success in the hotels domain might also explain the lower scores of the ASD and ASQP tasks in the Amazon\_FF and Coursera domains compared to TSD. The complex nature of ASQP, which demands establishing relationships between categories, opinion phrases, and targets, might account for its sub-optimal performance stemming from the above two reasons.

## 6. Significance of OATS

The main differences between the current ASQP datasets and OATS are: **(a)** we annotated the data from scratch with all the elements together instead of aligning from several sources, **(b)** we include all the implicit aspect terms and opinion phrases, un-



like ASQP, which excluded implicit opinion terms, (c) we provide both the review-level and sentence-level annotations facilitating the analysis and experiments for the review/text-level aspect-based sentiment analysis that differentiates Cai et al. (2023) from our dataset, (d) there are  $\approx 5.8K$  opinion quadruples in total for the ASQP corpus, which includes two restaurant datasets that are just 80% of opinion quadruples of the Coursera domain itself, which is the smallest in the OATS corpus. (e) OATS has a total of  $\approx 27.5K$  opinion quadruples, almost five times larger than the ASQP corpus. Also, if we include the review-level annotations, we have nearly  $44.5K$  opinions in OATS.

The OATS corpus will help analyze and understand the inherent relationships among all the elements of ABSA and exploit them to solve the ABSA task holistically. The comprehensive nature of the OATS corpus not only allows the evaluation of the ASQP task but also unlocks the potential analysis and exploration of all the subtasks of ABSA, including TASD, ASTE, TOWE, TAD, and many more. As pointed out by (Zhang et al., 2021a), the characteristic of tackling various ABSA tasks in a unified framework enables the knowledge to be easily transferred across related subtasks, which is especially beneficial under low-resource settings. It also allows cross-task transfer for subtasks that underperform when trained using only a task-specific dataset. The provided review-level annotations will help many researchers who like to explore ABSA for whole reviews rather than individual sentences or single-sentence reviews.

## 7. Conclusion

ABSA still presents a lot of difficulties in the rapidly developing field of generative artificial intelligence. We present a thoroughly curated ABSA dataset including quadruples, both explicit and implicit attributes, and views, spanning three unique domains. In addition, the dataset incorporates review-level sentiment polarity for each aspect category, providing a comprehensive perspective of the sentiments expressed in the reviews. Our annotations outnumber those found in previous datasets. We do an in-depth study, give comprehensive annotation guidelines, and provide dataset statistics. We also enable evaluations of generative and non-generative benchmarks on a range of common ABSA tasks.

## 8. Limitations and Ethical Considerations

We identified the following limitations of this work:

1. While the OATS dataset spans multiple do-

main, including the relatively unexplored domains of Amazon fine foods and Coursera course reviews, it exclusively contains English-language reviews. The dataset currently lacks reviews in low-resource languages. We plan to address this in future works by releasing multi-lingual datasets, inclusive of the low-resource setting.

2. The experimental results showcased in this paper are purely in-domain; models were fine-tuned and tested within the same domain. In subsequent works, we aim to design experiments to tackle cross-domain and out-of-domain ABSA tasks.
3. Our dataset incorporates review-level ABSA tuples focusing only on aspect categories and sentiment polarities. It currently omits review-level target expressions and opinion phrase annotations, which would render the dataset more comprehensive for ABSA. We earmark this as an avenue for future exploration.

We point out the following ethical considerations while building and using the OATS dataset:

1. The data collected from platforms such as Amazon Finefoods, Coursera, and TripAdvisor Hotels come from public reviews provided by users. We ensured that all identifiable information, including usernames, avatars, and any other potentially identifying details, were removed to preserve the anonymity of the reviewers.
2. Our data processing methodology focused on extracting and analyzing the content of the reviews without altering the original sentiment or meaning. We handled this data with the utmost care to ensure that the sentiments and opinions of the original reviewers were not misrepresented.
3. We recognize that reviews from online platforms might not represent the complete spectrum of users' opinions, as they may be influenced by various factors like platform algorithms, user demographics, and more. We urge users of this dataset to be aware of potential biases and always consider the data in the proper context.
4. While the reviews were publicly available, the original authors might not have expected their content to be used in research. We've taken measures to respect their privacy, but future users of this dataset should also be aware of this consideration.

5. While the reviews are publicly accessible, we acknowledge the platforms (Amazon Finefoods, Coursera, TripAdvisor Hotels) as the source of this data. We have only used this data for research and academic purposes, ensuring that we respect the terms of use for each platform.

## Acknowledgements

This research was supported by Adobe Gift funding to the University of Houston collaboration and by the National Science Foundation (NSF) under grant #1910192. The authors acknowledge the use of the Sabine Cluster and the advanced support from the Research Community Data Core at the University of Houston to carry out the research presented here. We also want to acknowledge Viet Lai (an NLP researcher from the University of Oregon) for his constant support of the BRAT Annotation tool maintenance and help with all related issues.

## 9. Bibliographical References

- Xiaoyi Bao, Wang Zhongqing, Xiaotong Jiang, Rong Xiao, and Shoushan Li. 2022. [Aspect-based sentiment analysis with opinion tree generation](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 4044–4050. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Markus Becker, Ben Hachey, Beatrice Alex, and Claire Grover. 2005. Optimising selective sampling for bootstrapping named entity recognition. In *ICML-2005 Workshop on Learning with Multiple Views*, pages 5–11. Citeseer.
- Hongjie Cai, Nan Song, Zengzhi Wang, Qiming Xie, Qiankun Zhao, Ke Li, Siwei Wu, Shijie Liu, Jianfei Yu, and Rui Xia. 2023. [Memd-absa: A multi-element multi-domain dataset for aspect-based sentiment analysis](#). *ArXiv*, abs/2306.16956.
- Hongjie Cai, Rui Xia, and Jianfei Yu. 2021. [Aspect-category-opinion-sentiment quadruple extraction with implicit aspects and opinions](#). 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)*, pages 340–350, Online. Association for Computational Linguistics.
- Siva Uday Sampreeth Chebolu, Franck Dernoncourt, Nedim Lipka, and Tamar Solorio. 2021. [Exploring conditional text generation for aspect-based sentiment analysis](#). In *Proceedings of the 35th Pacific Asia Conference on Language, Information and Computation*, pages 119–129, Shanghai, China. Association for Computational Linguistics.
- Siva Uday Sampreeth Chebolu, Franck Dernoncourt, Nedim Lipka, and Tamar Solorio. 2023. [Survey of aspect-based sentiment analysis datasets](#).
- 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*, pages 6515–6524, Online. Association for Computational Linguistics.
- Shaowei Chen, Yu Wang, Jie Liu, and Yuelin Wang. 2021. [Bidirectional machine reading comprehension for aspect sentiment triplet extraction](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12666–12674.
- Yew Ken Chia, Hui Chen, Wei Han, Guizhen Chen, Sharifah Mahani Aljunied, Soujanya Poria, and Lidong Bing. 2023. [Domain-expanded aste: Rethinking generalization in aspect sentiment triplet extraction](#).
- Louise Deléger, Qi Li, Todd Lingren, Megan Kaiser, Katalin Molnár, Laura Stoutenborough, Michal Kouril, Keith A. Marsolo, and Imre Solti. 2012. Building gold standard corpora for medical natural language processing tasks. *AMIA ... Annual Symposium proceedings. AMIA Symposium*, 2012:144–53.
- J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.
- 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, Volume 1 (Long and Short Papers)*, pages 2509–2518, Minneapolis, Minnesota. Association for Computational Linguistics.
- George Hripcsak and Adam S Rothschild. 2005. Agreement, the f-measure, and reliability in information retrieval. *Journal of the American medical informatics association*, 12(3):296–298.

- Mengting Hu, Yike Wu, Hang Gao, Yinhao Bai, and Shiwan Zhao. 2022. [Improving aspect sentiment quad prediction via template-order data augmentation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 7889–7900, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Minqing Hu and Bing Liu. 2004a. [Mining and summarizing customer reviews](#). In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, page 168–177, New York, NY, USA. Association for Computing Machinery.
- Minqing Hu and Bing Liu. 2004b. [Mining and summarizing customer reviews](#). In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, page 168–177, New York, NY, USA. Association for Computing Machinery.
- Qingnan Jiang, Lei Chen, Ruifeng Xu, Xiang Ao, and Min Yang. 2019. [A challenge dataset and effective models for aspect-based sentiment analysis](#). pages 6281–6286.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *CoRR*, abs/1910.13461.
- Shu Liu, Kaiwen Li, and Zuhe Li. 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*, pages 272–278, Seattle, United States. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pre-training approach](#).
- Yue Mao, Yi Shen, Jingchao Yang, Xiaoying Zhu, and Longjun Cai. 2022. [Seq2Path: Generating sentiment tuples as paths of a tree](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2215–2225, Dublin, Ireland. 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](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13543–13551.
- 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](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8600–8607.
- Joseph J. Peper and Lu Wang. 2022. Generative aspect-based sentiment analysis with contrastive learning and expressive structure. In *Conference on Empirical Methods in Natural Language Processing*.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. [SemEval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. [SemEval-2015 task 12: Aspect based sentiment analysis](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 486–495, Denver, Colorado. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [SemEval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Chi Sun, Luyao Huang, and Xipeng Qiu. 2019. Utilizing bert for aspect-based sentiment analysis via constructing auxiliary sentence. *arXiv preprint arXiv:1903.09588*.
- Katrin Tomanek and Udo Hahn. 2009. Timed annotations—enhancing muc7 metadata by the

- time it takes to annotate named entities. In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 112–115.
- Hai Wan, Y. Yang, Jianfeng Du, Y. Liu, Kunxun Qi, and J. Pan. 2020. Target-aspect-sentiment joint detection for aspect-based sentiment analysis. In *AAAI*.
- Youwei Wang, Ying Wang, Zhongchuan Sun, Yinghao Li, Shizhe Hu, and Yangdong Ye. 2023. Deep purified feature mining model for joint named entity recognition and relation extraction. *Information Processing and Management*, 60(6):103511.
- Zhengxuan Wu and Desmond C. Ong. 2020. Context-guided bert for targeted aspect-based sentiment analysis.
- 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 (Volume 1: Long Papers)*, pages 4755–4766, Online. Association for Computational Linguistics.
- Ting Xu, Huiyun Yang, Zhen Wu, Jiaze Chen, Fei Zhao, and Xinyu Dai. 2023. Measuring your ASTE models in the wild: A diversified multi-domain dataset for aspect sentiment triplet extraction. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2837–2853, Toronto, Canada. Association for Computational Linguistics.
- Yichun Yin, Yangqiu Song, and Ming Zhang. 2017. Document-level multi-aspect sentiment classification as machine comprehension. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2044–2054, Copenhagen, Denmark. Association for Computational Linguistics.
- Wenxuan Zhang, Yang Deng, Xin Li, Yifei Yuan, Lidong Bing, and Wai Lam. 2021a. Aspect sentiment quad prediction as paraphrase generation.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2021b. 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)*, pages 504–510, Online. Association for Computational Linguistics.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2021c. 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)*, pages 504–510, Online. Association for Computational Linguistics.
- Yice Zhang, Yifan Yang, Yihui Li, Bin Liang, Shiwei Chen, Yixue Dang, Min Yang, and Ruifeng Xu. 2022. Boundary-driven table-filling for aspect sentiment triplet extraction. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6485–6498, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.