

A Contrastive Cross-Channel Data Augmentation Framework for Aspect-Based Sentiment Analysis

Bing Wang^{1*}, Liang Ding^{2†}, Qihuang Zhong⁴, Ximing Li¹, Dacheng Tao^{2,3}

¹College of Computer Science and Technology, Jilin University

²JD Explore Academy ³The University of Sydney ⁴Wuhan University

dingliang1@jd.com, liximing86@gmail.com

Abstract

Aspect-based sentiment analysis (ABSA) is a fine-grained sentiment analysis task, which focuses on detecting the sentiment polarity towards the aspect in a sentence. However, it is always sensitive to the *multi-aspect challenge*, where features of multiple aspects in a sentence will affect each other. To mitigate this issue, we design a novel training framework, called Contrastive Cross-Channel Data Augmentation (C³DA), which leverages an in-domain generator to construct more multi-aspect samples and then boosts the robustness of ABSA models via contrastive learning on these generated data. In practice, given a generative pretrained language model and some limited ABSA labeled data, we first employ some parameter-efficient approaches to perform the in-domain fine-tuning. Then, the obtained in-domain generator is used to generate the synthetic sentences from two channels, i.e., Aspect Augmentation Channel and Polarity Augmentation Channel, which generate the sentence condition on a given aspect and polarity respectively. Specifically, our C³DA performs the sentence generation in a cross-channel manner to obtain more sentences, and proposes an Entropy-Minimization Filter to filter low-quality generated samples. Extensive experiments show that our C³DA can outperform those baselines without any augmentations by about 1% on accuracy and Macro-F1. Code and data are released in <https://github.com/wangbing1416/C3DA>.

1 Introduction

Sentiment Analysis (SA) is a typical Natural Language Understanding (NLU) task to predict the sentence-level sentiment polarities (Pang and Lee, 2007; Liu, 2012). However, since a single review sentence always exists multiple polarities, Aspect-based Sentiment Analysis (ABSA), a fine-grained

sentiment analysis task (Ma et al., 2017; Fan et al., 2018; Sun et al., 2019; Wang et al., 2020; Li et al., 2021b; Zhang et al., 2022a), is further introduced to detect the sentiment polarities towards given aspects (entities) in a review sentence. Gathering a review "while the ambiance and atmosphere were great, the food and service could have been a lot better.", the sentiment polarities of "ambiance" and "atmosphere" are *Positive* and we get *Negative* polarity for aspects "food" and "service".

It is crucial to capture the aspect-specific contextual features for an ABSA model. Unfortunately, as shown in the aforementioned example, the multi-aspect challenge that there are multiple aspects in a sentence always deteriorates the model's generalization, especially when the expressed polarities of multiple aspect words in a sentence are opposite. Meanwhile, since ABSA is a low-resource task, it is hard to train a robust model under fewer samples.

Some existing works focused on the multi-aspect challenge, such as Lu et al. (2011); Hu et al. (2019) in which novel model structures were designed to capture aspect-specific sentiment information. Additionally, Jiang et al. (2019) manually collected a large-scale high-quality multi-aspect multi-sentiment (MAMS) dataset. Undoubtedly, the data-centric MAMS is an effective approach to tackle the multi-aspect problem, but the human-annotated and non-expandable limitations still hinder the robust training for ABSA models. Therefore, given limited labeled data, it is critical to investigate how to collect more in-domain multi-aspect samples automatically. On the other hand, capturing the aspect-specific information from multi-aspect sentences is also necessary for robust ABSA models.

In response to the aforementioned problems, we propose a novel training framework, namely Contrastive Cross-Channel Data Augmentation (C³DA), to generate more in-domain multi-aspect samples and train robust ABSA models based on these generated data. Firstly, inspired by successful

*Work was done when Bing and Qihuang were interning at JD Explore Academy.

†Liang Ding is the corresponding author.

generative pretraining models (Raffel et al., 2020; Lewis et al., 2020), we employ an representative pretraining models T5 (Raffel et al., 2020) as the multi-aspect data generator. In practice, due to the limited ABSA labeled data, it is hard to fine-tuning the entire T5 model effectively. Thus, we further introduce some parameter-efficient methods, e.g. prompt (Lester et al., 2021a), prefix (Li and Liang, 2021) and LoRA (Hu et al., 2022), to tune the T5 generator. In this way, a domain -specific generator is obtained and we can apply it to collect more in-domain multi-aspect samples.

To be more specific, there are two channels to generate expected multi-aspect samples in our framework, i.e. **Aspect Augmentation Channel (AAC)** and **Polarity Augmentation Channel (PAC)**. In the AAC, the generator is encouraged to generate the synthetic sentence towards the given sentence and aspect, while in the PAC, the synthetic sentence towards the given sentence and polarity is obtained. To further boost the generated samples, we attempt to collect the multi-aspect samples in a cross-channel manner. Specifically, in the first generation, the source sentence is fed to the double channels and obtain an aspect-specific sentence and a polarity-inverted sentence respectively. In the second generation, both generated sentences are injected to the another channel to get the ultimate sentences. Finally, we propose an Entropy-Minimization Filter (EMF) to filter some low-quality generated samples. And a contrastive training objective that can draw away the different aspect’s embeddings in a sentence is also leveraged to alleviate the multi-aspect problem and help train the final robust ABSA model.

We conduct sufficient experiments on three popular ABSA datasets, i.e. *Restaurant*, *Laptop* (Pontiki et al., 2014) and *Twitter* (Dong et al., 2014), to prove that our C³DA framework can boost the model’s robustness and predictive performance, and outperform other NLP data augmentation strategies. Specifically, with the help of our C³DA, RoBERTa-based models can achieve averaged performance improvement 1.28% and 1.62% in terms of accuracy and macro-F1 respectively, while the improvements of BERT-based models are also over 0.87% and 1.10%. Some in-depth discussions and case studies are also executed. Contributions of this paper are threefold:

- We recast the vanilla ABSA training scheme with a data augmentation-based training

framework and a contrastive training objective to tackle the multi-aspect challenge.

- We design a novel cross-channel data augmentation method based on generative large-scale pretrained language models to generate high-quality in-domain multi-aspect samples, which has great potential to benefit other fine-grained NLU tasks.
- Extensive experiments on three widely-used datasets show the effectiveness of our C³DA.

2 Related Work

2.1 Language Models and Parameter Efficient

Large Pre-trained Language Models (PLMs) are still research cores in various natural language processing tasks. The PLMs tend to be variants of Transformer (Vaswani et al., 2017), and are trained on massive unlabeled raw sentences under some linguistic unsupervised objectives. According to the model structure, these Transformer-based PLMs are roughly divided into (1) Encoder-only LMs (Devlin et al., 2019; Liu et al., 2019) can be utilized to capture token-level and sentence-level features, (2) Decoder-only LMs (Radford et al., 2018, 2019) tend to design an auto-regressive objective to cater to text generation tasks and (3) Encoder-Decoder LMs (Lewis et al., 2020; Raffel et al., 2020) need two sentences (source and target) to perform conditional generation.

The above PLMs always follow the *pre-train + fine-tune (FT)* paradigm to adapt the downstream tasks, however, tuning a whole model is a tricky challenge and is not applicable to low-resource NLP tasks. Therefore, a novel *pre-train + prompt-tune (PT)* (Liu et al., 2021c) concept is proposed to break this challenge, PT-based models can also be divided into (1) manual discrete prompts design human-made templates to adapt different pre-trained objective (Liu et al., 2021a), and (2) soft prompts inject some learnable pseudo tokens or matrices to frozen PLMs (Liu et al., 2021d; Lester et al., 2021b; Li and Liang, 2021; Zhong et al., 2022b).

2.2 Aspect-based Sentiment Analysis

Aspect-based Sentiment Analysis (ABSA) is a kind of Text Classification task, that benefits from a better aspect-aware text representation. Therefore, an amount of neural network-based models were

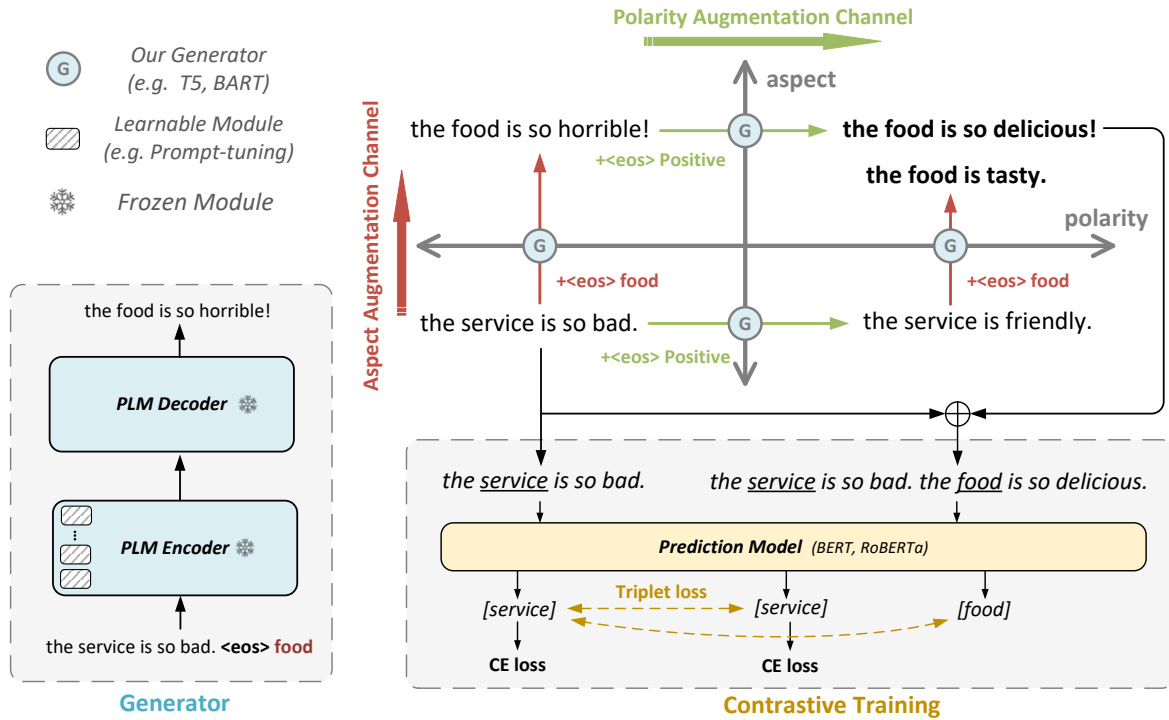


Figure 1: The overall C³DA framework. In our generator, learnable modules can be optimized in the fine-tuning stage and will also be frozen in the generation stage, and the PLM encoder-decoder structure is always fixed.

proposed to break the ABSA challenge, for example, graph-based models (Sun et al., 2019; Wang et al., 2020; Li et al., 2021b; Zhong et al., 2022a) conduct graph convolutional operations on dependency trees to encode the semantic information, and attention-based models (Ma et al., 2017; Chen et al., 2017; Fan et al., 2018; Song et al., 2019; Liu et al., 2021b) focus on interaction between aspect terms and context tokens, *etc.*

In recent years, pre-trained language models changed the situation of ABSA because they are better text encoders and can capture context features. Therefore, researchers explored more views to improve ABSA performance, except that PLM is regarded as an encoder to replace the vanilla *Glove* + *BiLSTM* paradigm. For example, Dai et al. (2021) leverages RoBERTa to re-construct dependency trees, Yan et al. (2021); Li et al. (2021a); Zhang et al. (2021) try to integrate various ABSA subtasks into a unified generative framework, and Zhou et al. (2021); Seoh et al. (2021) combine ABSA with other NLP tasks, *etc.*

2.3 Data Augmentation

To improve the scale of training samples, some Data Augmentation (DA) methods in the NLP community are proposed (Zhang et al., 2018; Wei

and Zou, 2019; Kobayashi, 2018; Wu et al., 2019; Anaby-Tavor et al., 2020; Ding et al., 2021; Cao et al., 2021; Wang et al., 2022; Zhang et al., 2022b). EDA (Wei and Zou, 2019) transforms source sentences with Synonym Replacement approaches, *etc.* CBERT (Wu et al., 2019) introduces a new conditional masked language model task. Meanwhile, some PLM-based generative augmentation frameworks (Anaby-Tavor et al., 2020; Wang et al., 2022; Zhang et al., 2022b) also have been explored. Also, recent work reveals the complementarity between PLMs and classical approaches, e.g., back translation (Liu et al., 2021e) and random initialization (Zan et al., 2022), to augment both the accuracy and generalization.

3 Our C³DA Framework

Task Description of ABSA Given a group of triplets $\Omega^{src} = \{s_i, \mathbf{a}_i, \mathbf{p}_i\}_{i=1}^N$, a sentence $s_i = \{s_{i1}, \dots, s_{iL}\}$ and an aspect-span indicator $\mathbf{a}_i \in \{0, 1\}^L$ (1 denotes this token is an aspect, vice versa) will be fed to a trainable model to obtain a hidden embedding $\mathbf{h}_i = \mathcal{F}(s_i, \mathbf{a}_i, \Theta) \in \mathbb{R}^{L \times D}$, where N , L and D denote the instance scale, sentence length and hidden embedding size, and $\mathcal{F}(\cdot, \Theta)$ is a prediction model (*e.g.*, BERT (Devlin et al., 2019)). Then, a ground-truth polarity label

$\mathbf{p}_i \in \{0, 1, 2\}$ is utilized to conduct supervised learning to estimate task-specific parameters Θ .

3.1 Overall Learning Objective

In this section, we introduce the overall process of our proposed method, and more details will be further described in Section 3.2 and 3.3, and our C³DA framework is shown in Figure 1.

Our C³DA framework consists of two steps: (1) Data Augmentation and (2) Contrastive Training. We utilize the representative pre-trained T5 model (Raffel et al., 2020) as our generator in the data augmentation stage¹. Given a sentence \mathbf{s}_i , we expect to generate an another sentence $\hat{\mathbf{s}}_i = \mathcal{G}(\mathbf{s}_i; \Pi)$ that expresses the opposite polarity and has a different aspect from the source sentence, where $\mathcal{G}(\cdot; \Pi)$ is the T5 generator.

Basically, we conduct the Supervised Classification Training (SCT) on source sentences and augmented sentences, the objective is as follows:

$$\mathcal{L}_{SCT} = \frac{1}{N} \sum_{i=1}^N \ell_{CE}(\mathbf{h}_i \mathbf{W}_s + \mathbf{b}_s, \mathbf{p}_i) + \alpha \ell_{CE}(\mathbf{h}_i^p \mathbf{W}_s + \mathbf{b}_s, \mathbf{p}_i), \quad (1)$$

$$\mathbf{h}_i^p = \mathcal{F}([\mathbf{s}_i, \hat{\mathbf{s}}_i], \mathbf{a}_i, \Theta), \quad (2)$$

where $[\cdot, \cdot]$ and $\ell_{CE}(\cdot)$ are the concatenate operation and a Cross-Entropy loss function, $\mathbf{W}_s \in \mathbb{R}^{D \times M}$ and α denote the classification head and a hyper-parameter, respectively.

To obtain more robust performance, we leverage the Contrastive Training (CT) formulated as the triplet loss (Schultz and Joachims, 2003) as follows:

$$\mathcal{R}_{CT} = \frac{1}{N} \sum_{i=1}^N \max\{d(\mathbf{h}_i, \mathbf{h}_i^p) - d(\mathbf{h}_i, \mathbf{h}_i^n) + \xi, 0\}, \quad (3)$$

where \mathbf{h}_i^n is the average-pooled sentence representation of $\hat{\mathbf{s}}_i$ in the sentence $[\mathbf{s}_i, \hat{\mathbf{s}}_i]$, ξ is a hyper-parameter called margin that can control whether an instance should be trained, and $d(\cdot, \cdot)$ denotes a distance measure function. Specially, we apply a negative cosine similarity to measure the embedding distance $d(\cdot, \cdot)$.

$$d(\mathbf{h}_i, \mathbf{h}_i^p) = -\frac{\mathbf{h}_i \mathbf{h}_i^p}{\|\mathbf{h}_i\| \times \|\mathbf{h}_i^p\|}. \quad (4)$$

¹Note that we can employ various generative pretraining models as the generator. T5 is used in this work, since it is widely-used and powerful.

Finally, the overall learning objective can be formulated as:

$$\mathcal{L} = \mathcal{L}_{SCT} + \beta \mathcal{R}_{CT}, \quad (5)$$

where β is a controllable hyper-parameter.

3.2 Cross-Channel Data Augmentation

Our motivation is to generate high-quality in-domain multi-aspect samples. In practice, the cross-channel data augmentation consists of two steps: 1) tune the pretrained backbone model to adapt our cross-channel sentence generation; 2) generate sentences with opposite polarity and different aspects.

Pre-Trained Generator We employ pre-trained encoder-decoder model T5 as our backbone. ABSA is a low-resource task, therefore, tuning a whole large language model (*e.g.*, T5) with scarce samples will cause an over-fitting problem. Inspired by Lester et al. (2021a); Wang et al. (2022), we utilize the parameter-efficient strategies to tune the generation model. In particular, we consider several efficient tuning methods, including prefix-tuning (Li and Liang, 2021), prompt-tuning (Lester et al., 2021a), and LoRA (Hu et al., 2022). We grid-search the optimal efficient tuning method for each dataset, detailed comparisons can be found at Appendix A.

Model Fine-Tuning To make the T5 model capable of generating new sentences conditioned on the original sentence and given aspect or polarity, we first construct the fine-tuning data based on the ABSA training set. Concretely, we randomly sample two instances $\{\mathbf{s}_i\}$ and $\{\mathbf{s}_j, \mathbf{a}_j, \mathbf{p}_j\}$ from the source triplets Ω^{src} . Then, we can tune the T5 model to generate \mathbf{s}_j conditioned on $[\mathbf{s}_i, \mathbf{a}_j]$ or $[\mathbf{s}_i, \mathbf{p}_j]$. As aforementioned, parameter-efficient strategies are used to avoid over-fitting. Taking the cutting-edge prompt-tuning (Lester et al., 2021a) as an example, the massive parameters in the generator should be fixed in this stage, and we only tune the parameters of prompts.

Sentence Generation The sentence generation for data augmentation stage consists of Aspect Augmentation Channel (AAC) and Polarity Augmentation Channel (PAC).

AAC generates an aspect-specific sentence $\mathbf{s}_i^a = \mathcal{G}([\mathbf{s}_i, \hat{\mathbf{A}}], \Pi)$ ², where Π is the trainable parameters

²AAC can generate one or more sentences, we take generating a single sentence as an example.

Table 1: An example of the concatenation operation.

\mathbf{s}	<i>I've been here 3 times for lunch and it is one of my favorites in the city.</i>
$[\mathbf{s}, \text{service}]$	<i>I've been here 3 times for lunch and it is one of my favorites in the city. < eos > service</i>
$[\mathbf{s}, \text{Negative}]$	<i>I've been here 3 times for lunch and it is one of my favorites in the city. < eos > so bad</i>

of the T5 generator $\mathcal{G}(\cdot)$. PAC generates a polarity-inverted sentence $\mathbf{s}_i^p = \mathcal{G}([\mathbf{s}_i, \hat{\mathbf{P}}], \Pi)$, where $\hat{\mathbf{A}}$ is a random-sampled in-domain aspect and $\hat{\mathbf{P}}$ is a group of seed spans that have the inverted polarity from \mathbf{s}_i , such as *so good*. Then, these sentences will be fed into the another channel to obtain $\mathbf{s}_i^{pa} = \mathcal{G}([\mathbf{s}_i^p, \hat{\mathbf{A}}], \Pi)$ and $\mathbf{s}_i^{ap} = \mathcal{G}([\mathbf{s}_i^a, \hat{\mathbf{P}}], \Pi)$ that can satisfy our motivation. We give an example of the concatenation operation $[\mathbf{s}, \cdot]$ in Table 1.

Note that for sentence generation, both the generator backbone and prompts will be frozen. The detailed Cross-Channel Data Augmentation method is shown in Algorithm 1.

Algorithm 1 Cross-Channel Data Augmentation.

Require: Source triplets $\Omega^{src} = \{\mathbf{s}_i, \mathbf{a}_i, \mathbf{p}_i\}_{i=1}^N$; Initial pre-trained T5 generator with soft prompt $\mathcal{G}(\cdot, \Pi^0)$; Pre-defined aspect set \mathcal{A} ; Fine-tune iterations τ .

Ensure: Augmented sentences $\Omega^{aug} = \{\hat{\mathbf{s}}_i\}_{i=1}^N$.

- 1: **for** $iter$ in τ **do** ▷ Stage I: Model Fine-tuning
- 2: $\{\mathbf{s}_i\}, \{\mathbf{s}_j, \mathbf{a}_j, \mathbf{p}_j\} \leftarrow \text{Sample}(\Omega^{src})$
- 3: $\nabla^a \leftarrow \text{loss}(\mathbf{s}_j, \mathcal{G}([\mathbf{s}_i, \mathbf{a}_j], \Pi^{iter}))$
- 4: $\nabla^p \leftarrow \text{loss}(\mathbf{s}_j, \mathcal{G}([\mathbf{s}_i, \mathbf{p}_j], \Pi^{iter}))$
- 5: $\Pi^{iter+1} \leftarrow \text{Train}(\Pi^{iter}, \nabla^a, \nabla^p)$
- 6: **end for**
- 7: $\Pi \leftarrow \Pi^\tau$
- 8: **for** $\{\mathbf{s}_i, \mathbf{a}_i, \mathbf{p}_i\}$ in Ω^{src} **do**
- 9: ▷ Stage II: Sentence Generation
- 10: $\hat{\mathbf{A}} \leftarrow \text{Sample}(\mathcal{A})$
- 11: $\hat{\mathbf{P}} \leftarrow -\mathbf{p}_i$ ▷ $-\mathbf{p}$ denotes opposite polarity from \mathbf{p}
- 12: $\mathbf{s}_i^a \leftarrow \mathcal{G}([\mathbf{s}_i, \hat{\mathbf{A}}], \Pi), \mathbf{s}_i^p \leftarrow \mathcal{G}([\mathbf{s}_i, \hat{\mathbf{P}}], \Pi)$
- 13: $\mathbf{s}_i^{pa} \leftarrow \mathcal{G}([\mathbf{s}_i^p, \hat{\mathbf{A}}], \Pi), \mathbf{s}_i^{ap} \leftarrow \mathcal{G}([\mathbf{s}_i^a, \hat{\mathbf{P}}], \Pi)$
- 14: $\Omega^{aug} \leftarrow \Omega^{aug} + \text{Filter}(\mathbf{s}_i^a, \mathbf{s}_i^p, \mathbf{s}_i^{pa}, \mathbf{s}_i^{ap})$
- 15: **end for**
- 16: **return** Ω^{aug} ;

3.3 Entropy-Minimization Filtering

Cross-Channel Data Augmentation generates four sentences for each source sentence, and we can leverage these sentences to train a prediction model $\mathcal{F}(\cdot, \Theta)$. However, it is necessary to filter out some low-quality samples that carry noisy (or difficult) information, which can be estimated by language modeling (Moore and Lewis, 2010) and norm of word embedding (Liu et al., 2020).

To ensure the certainty of synthetic sentences, we design an Entropy-Minimization Filter (EMF) to filter those noisy sentences. We establish a hypothesis that *noisy sentences always express an ambiguous polarity* such that our prediction model gives a more smooth polarity probability distribution. To be specific, the information entropy of prediction distribution as Equation 6 is an appropriate measure. A sentence with larger prediction entropy should be filtered, and a hyper-parameter k can control how many augmented sentences are selected to participate in training for each sample.

$$\mathcal{H}_i = -\mathbb{E}_{\mathbf{y}_i} [\log_2 p(\mathbf{y}_i)] = -p(\mathbf{y}_i) \log_2 p(\mathbf{y}_i), \tag{6}$$

$$p(\mathbf{y}_i) = \text{softmax}(\mathbf{h}_i \mathbf{W}_s + \mathbf{b}_s). \tag{7}$$

4 Experiment

4.1 Experimental Settings and Baselines

We conduct our experiments on three public datasets, i.e. *Restaurant*, *Laptop* from SemEval 2014 ABSA task (Pontiki et al., 2014) and *Twitter* (Dong et al., 2014). In practice, we run the experiments over 5 random seeds and report the average values of accuracy (acc) and Macro-F1 (F1) to avoid stochasticity.

Two widely-used pretraining models, i.e. BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), are regarded as our prediction backbone, and we compare our framework C³DA with the following four data augmentation methods:

- **EDA** (Wei and Zou, 2019) is a token-level transformation method. It follows some patterns such as Synonym Replacement, Random Insertion, *etc.* For the ABSA task, we deliberately avoid corrupting the aspect words.
- **Back Translation (BT)** (Sennrich et al., 2016) is a sentence-level augmentation method, which translates a sentence to another language and translates it back to the original language. We follow Fan et al. (2021) that translates English to French and translates it back to English.

Table 2: Empirical results of C³DA and other cutting-edge data augmentation methods. The best scores of each metric are indicated in bold, and \uparrow denotes the improvement over BERT-*base* and RoBERTa-*base* baselines.

Model	<i>Restaurant</i>		<i>Laptop</i>		<i>Twitter</i>	
	acc	F1	acc	F1	acc	F1
BERT- <i>base</i>	86.31	80.22	79.66	76.11	76.50	75.23
+ EDA (Wei and Zou, 2019)	86.42	79.63	79.59	75.79	76.26	75.16
+ BT (Fan et al., 2021)	86.47	80.29	79.67	76.35	76.63	75.47
+ CBERT (Wu et al., 2019)	86.27	80.00	79.83	76.12	76.44	75.36
+ SCon (Liang et al., 2021)	86.51	80.55	80.23	76.48	-	-
+ C ³ DA (Ours)	86.93 \uparrow _{0.62}	81.23 \uparrow _{1.01}	80.61 \uparrow _{0.95}	77.11 \uparrow _{1.00}	77.55 \uparrow _{1.05}	76.53 \uparrow _{1.30}
RoBERTa- <i>base</i>	86.38	80.29	80.10	76.24	76.63	75.37
+ EDA (Wei and Zou, 2019)	86.43	80.21	80.38	76.59	76.47	75.36
+ BT (Fan et al., 2021)	86.50	80.59	80.22	76.73	76.59	75.47
+ CBERT (Wu et al., 2019)	86.77	80.51	80.54	76.57	76.73	75.37
+ C ³ DA (Ours)	87.11 \uparrow _{0.73}	81.63 \uparrow _{1.34}	81.83 \uparrow _{1.73}	78.46 \uparrow _{2.22}	78.31 \uparrow _{1.38}	76.67 \uparrow _{1.30}

- **CBERT** (Wu et al., 2019) fully excavates the power of mask language model (MLM) objective to replace some tokens, so it is also a token-level replacement approach.
- **SCon** (Liang et al., 2021) design aspect-invariant/-dependent data augmentation for ABSA and deploy a supervised contrastive learning objective. We reproduce it according to their released code and default best settings.

4.2 More Empirical Details

We implement our experiment by following the released ABSA Pytorch³ repository and our pre-trained models, such as *bert-base-uncased*, *roberta-base* and *t5-base*, are from *HuggingFace*⁴. Meanwhile, we utilize a flexible toolkit *OpenDelta*⁵ to adapt various parameter-efficient methods to our T5 generator.

In the data augmentation stage, we leverage an *Adafactor* optimizer. A T5 generator is fine-tuned for 100 epochs, and batch size is fixed to 16. As training a prediction model BERT or RoBERTa, we adapt an *Adam* optimizer with a learning rate of 2×10^{-5} , and the dropout rate is set to 0.3. Additionally, we train the final prediction model for 15 epochs with the batch size as 16.

A single NVIDIA A100 is used to conduct our all experiments. The model fine-tuning and sentence generation stages will generally spend 2 - 3

hours, and it spends about 1 hour to train a prediction model.

4.3 Main Result

Table 2 shows our main experimental results. To be fair, for each data augmentation method, we generate 1 synthetic sample per sentence. Overall, our C³DA framework consistently and significantly improve the performance of both BERT-based and RoBERTa-based models, e.g. the result of *Restaurant* even outperforms the RoBERTa-*base* baseline 2.22% in term of F1. Moreover, we observe that the performance order of various data augmentation methods is C³DA > CBERT > BT > EDA, which indicates: (1) EDA semantic-regardlessly removes or inserts tokens, which will inject more noise into final prediction models, thus leading to the poor performance. (2) Comparing BT and CBERT, BT always achieves higher F1, and CBERT will get more advantageous accuracy. Since BT is a sentence-level method, it focuses on how to transform a source sentence into a form-different and semantic-reserved sentence, but neglects the useful aspect information. On the contrary, CBERT will only change some tokens, which is a semantic-reversed and aspect-reserved approach. Table 3 summarizes the characteristics of some augmentation methods.

Turn to compare two metrics, another finding is that C³DA has more significant performance in term of Macro-F1. Proverbially, since Macro-F1 is a metric to measure the generalization of a model, the higher improvements on Macro-F1 prove that our framework can effectively improve the model’s robustness in this view. Concurrently, the effect

³<https://github.com/songyouwei/ABSA-PyTorch>

⁴<https://huggingface.co/models>

⁵<https://github.com/thunlp/OpenDelta>




Figure 2: Sensitivity Analysis of the contrastive training margin ξ . The ribbon upper-side and below-side denote the maximum and minimum performance in 5 seeds respectively, and the centerline is the average value.




Figure 3: Sensitivity Analysis of the number of selected sentences k in EMF.

Table 3: A comparison of various data augmentation methods. Semantic-R. and Aspect-R. denote semantic reserved methods and aspect reserved methods, and T.- / S.- discriminates token-level and sentence-level methods respectively.

Methods	Semantic-R.	Aspect-R.	T.- / S.-
EDA	×	✓	T.
BT	✓	×	S.
CBERT	✓	✓	T.
C ³ DA	✓	✓	S.

of C³DA is better on *Restaurant* and *Twitter*. The difference between these datasets is that *Laptop* is a small dataset, indicating that our C³DA is sensitive to the data scale and C³DA may have better improvement if there is a larger-scale dataset.

4.4 Ablation Study

The core modules in our framework C³DA are Data Augmentation (DA), Contrastive Learning (CL), and Entropy-Minimization Filter (EMF). In this section, we perform ablation study to investigate the effect of these modules, and the results are listed in Table 5. BERT-base is regarded as the prediction model in this study, and *w/o* denotes *without* some modules.

Firstly, *w/o DA & CL* removes our data augmentation framework, it is equivalent to a vanilla BERT-base-only prediction model. As shown in Table 5, *w/o DA & CL* always presents poor perfor-

mance, even decrease by about 1% on *Laptop*'s acc and F1 and *Restaurant*'s F1. It proves our C³DA framework imports more high-quality data items and is effective to boost the performance. Then, *w/o CL* remains augmented sentences and revokes the contrastive training objective \mathcal{R}_{CL} . Overall, this objective improves the results by about 0.3% on each dataset, conforming that contrastive learning can powerfully optimize the aspect-specific features. Finally, *w/o EMF* conducts both supervised training and contrastive training on all generated sentences. Intuitively, more samples can consolidate the predictive power of a classification model, but some emotionless and illogical sentences also carry plenty of noise. As we speculate, *w/o EMF* degrades the model's performance by 0.5% - 1.0%, the outcome on *Restaurant*'s accuracy is even worse than the result without DA & CL. More sensitivity studies about EMF will be shown in Section 4.5.

Additionally, we also investigate the performance of generating sentences with **only AAC** and **only PAC** settings and report their results in Table 5. We show that both them are beneficial, where only AAC slightly outperforms that of the only PAC, demonstrating that the multi-aspect data are more crucial than the multi-polarity data.

4.5 Sensitivity Analysis

As mentioned above, we employ 4 hyper-parameters in our C³DA framework, and ξ and k are more crucial. To investigate the influence of

Table 4: Case study. The aspect words and corresponding polarities are noted as red color and blue color, respectively.

case						
	<i>the falafal was rather over cooked and dried but the chicken was fine.</i>					
Aspect & Polarity	<i>falafal</i>			<i>chicken</i>		
	<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>	<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>
w/o C ³ DA	0.098	0.109	0.793	0.643	0.047	0.310
with C ³ DA	0.071	0.092	0.837	0.698	0.076	0.226
case						
	<i>good food but the service was dreadful!</i>					
Aspect & Polarity	<i>food</i>			<i>service</i>		
	<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>	<i>Positive</i>	<i>Neutral</i>	<i>Negative</i>
w/o C ³ DA	0.502	0.173	0.325	0.191	0.236	0.573
with C ³ DA	0.589	0.211	0.200	0.177	0.186	0.637

Table 5: Ablation study of C³DA. The prediction model is BERT-base in this study.

Model	<i>Restaurant</i>	<i>Laptop</i>	<i>Twitter</i>
	acc		
C³DA	86.93	80.61	77.08
w/o DA & CL	86.31	79.66	76.50
w/o CL	86.69	80.35	76.87
w/o EMF	86.45	79.84	76.50
only AAC	86.63	80.10	77.16
only PAC	86.44	79.94	77.19
Model	F1		
C³DA	81.23	77.11	75.76
w/o DA & CL	80.22	76.11	75.23
w/o CL	81.00	76.88	75.63
w/o EMF	80.21	76.45	75.17
only AAC	80.61	76.69	75.93
only PAC	80.24	76.17	75.91

both hyper-parameters, we further conduct sensitivity analysis in this section. Additionally, we fix α and β to 0.5 and 2.0, respectively. Limited by space, their sensitivity analysis will be described and analyzed in Appendix B.

Contrastive Training Margin ξ Margin ξ in Equation 3 is an important hyper-parameter to control the model’s distinctive ability to discriminate negative samples and positive samples. Under a little margin, the prediction model will indolently balance $d(\mathbf{h}_i, \mathbf{h}_i^p)$ and $d(\mathbf{h}_i, \mathbf{h}_i^n)$ to cause under-optimization, while a large margin makes a model converge hard.

To investigate the impact of the margin ξ , we implement experiments on *Restaurant* and *Laptop*, and range ξ from 0.1 to 0.6. As shown in Figure 2, our C³DA always gets the best performance when ξ is 0.3, and as it expands or shrinks, the results

consistently show a downward trend.

Number of Selected Sentences k in EMF As the description in Section 4.4, k in EMF is a key hyper-parameter to decide to introduce more useful sentences or more noisy information, we select k sentences for each source sentence with the entropy-minimization filter method.

We evaluate C³DA on $k \in \{1, 2, 3, 4\}$ and show the results in Figure 3. The best performance invariably appears in $k = 2$ or $k = 1$. Meanwhile, the performance decreases significantly when the value of k is 3, 4.

4.6 Case Study

C³DA is a DA-based framework, and our purpose is to break the multi-aspect challenge. Therefore, we implement a case study to evaluate whether C³DA can solve this problem, and more cases with different DA methods and cases with varying fine-tuning epochs will be shown in Appendix C.

We select two samples from the *Restaurant* test set, and these two sentences all have two aspects with different sentiment polarities. We will compare the performance of our C³DA on their polarity prediction distribution. To observe the Table 4, we summarize that (1) C³DA can consistently and significantly enhance the prediction model’s robustness, such as the *Positive* probability of the aspect *food* improve about 0.077. Therefore, our framework can correct the polarity prediction when a sentence expresses an ambiguous polarity. (2) When a sentence is short, the interaction between two aspects with contradictory polarities will be more intimate. In this scenario, sentences will express more ambiguous sentiment and the effectiveness of our framework will be more significant.

4.7 Computing Cost Analysis

Our method will prolong each sentence, longer sentences will undoubtedly increase computing costs. Therefore, we attempt to compare the computing time of our framework and other augmentation methods. In the data preparing stage, C³DA will fine-tune a generator and inference on the training set, it will spend 2 - 3 hours, and it is a once-for-all process. In the prediction training stage, we find that EDA, Back Translation and CBERT do not change the sentence length, their prediction time cost $0.98\times \sim 1.05\times$ (without any augmentation is $1\times$). By comparison, C³DA spends $1.14\times$ time cost, which is completely acceptable.

5 Conclusion

In our work, we focus on solving the multi-aspect problem in the ABSA task. To address this challenge, we propose a data-centric training framework and design a novel C³DA method to implement data augmentation. The proposed framework conduct both supervised training and contrastive training on augmented samples. The key idea of C³DA is to generate some sentences that express the opposite sentiment polarity from the source sentences. Extensive experiments on widely-used benchmarks demonstrate that our framework can effectively improve the prediction model's robustness and predictive performance. In the future, it will be interesting to apply our C³DA framework to improve other fine-grained natural language understanding tasks.

Acknowledgements

We are grateful to the anonymous reviewers and the area chair for their insightful comments and suggestions. Bing Wang and Ximing Li were supported by the project from the National Natural Science Foundation of China (NSFC) (No.62006094).

References

- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. [Do not have enough data? deep learning to the rescue!](#) In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 7383–7390.
- Yu Cao, Liang Ding, Zhiliang Tian, and Meng Fang. 2021. [Towards efficiently diversifying dialogue generation via embedding augmentation](#). *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7443–7447.
- Peng Chen, Zhongqian Sun, Lidong Bing, and Wei Yang. 2017. [Recurrent attention network on memory for aspect sentiment analysis](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 452–461.
- Junqi Dai, Hang Yan, Tianxiang Sun, Pengfei Liu, and Xipeng Qiu. 2021. [Does syntax matter? A strong baseline for aspect-based sentiment analysis with roberta](#). In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1816–1829.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: pre-training of deep bidirectional transformers for language understanding](#). In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186.
- Liang Ding, Di Wu, and Dacheng Tao. 2021. [Improving neural machine translation by bidirectional training](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3278–3284, Online and Punta Cana, Dominican Republic.
- Li Dong, Furu Wei, Chuanqi Tan, Duyu Tang, Ming Zhou, and Ke Xu. 2014. [Adaptive recursive neural network for target-dependent twitter sentiment classification](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 49–54.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22:107:1–107:48.
- Feifan Fan, Yansong Feng, and Dongyan Zhao. 2018. [Multi-grained attention network for aspect-level sentiment classification](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3433–3442.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The International Conference on Learning Representations*.
- Mengting Hu, Shiwan Zhao, Li Zhang, Keke Cai, Zhong Su, Renhong Cheng, and Xiaowei Shen. 2019. [CAN: constrained attention networks for multi-aspect sentiment analysis](#). In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 4600–4609.

- Himanshu Jain, Yashoteja Prabhu, and Manik Varma. 2016. [Extreme multi-label loss functions for recommendation, tagging, ranking & other missing label applications](#). In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 935–944.
- Qingnan Jiang, Lei Chen, Ruifeng Xu, Xiang Ao, and Min Yang. 2019. [A challenge dataset and effective models for aspect-based sentiment analysis](#). In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 6279–6284. Association for Computational Linguistics.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 452–457, New Orleans, Louisiana.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021a. [The power of scale for parameter-efficient prompt tuning](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021b. [The power of scale for parameter-efficient prompt tuning](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Chengxi Li, Feiyu Gao, Jiajun Bu, Lu Xu, Xiang Chen, Yu Gu, Zirui Shao, Qi Zheng, Ningyu Zhang, Yongpan Wang, and Zhi Yu. 2021a. [Sentiprompt: Sentiment knowledge enhanced prompt-tuning for aspect-based sentiment analysis](#). *arXiv preprint*, abs/2109.08306.
- Ruifan Li, Hao Chen, Fangxiang Feng, Zhanyu Ma, Xiaojie Wang, and Eduard H. Hovy. 2021b. [Dual graph convolutional networks for aspect-based sentiment analysis](#). In *Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 6319–6329.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). In *Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 4582–4597.
- Bin Liang, Wangda Luo, Xiang Li, Lin Gui, Min Yang, Xiaoqi Yu, and Ruifeng Xu. 2021. [Enhancing aspect-based sentiment analysis with supervised contrastive learning](#). In *ACM International Conference on Information and Knowledge Management*, pages 3242–3247.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. [Focal loss for dense object detection](#). In *IEEE International Conference on Computer Vision*, pages 2980–2988.
- Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- Jian Liu, Zhiyang Teng, Leyang Cui, Hanmeng Liu, and Yue Zhang. 2021a. [Solving aspect category sentiment analysis as a text generation task](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 4406–4416.
- Juhua Liu, Qihuang Zhong, Liang Ding, Hua Jin, Bo Du, and Dacheng Tao. 2021b. [Unified instance and knowledge alignment pretraining for aspect-based sentiment analysis](#). *arXiv preprint*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021c. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *arXiv preprint*, abs/2107.13586.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021d. [GPT understands, too](#). *arXiv preprint*, abs/2103.10385.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020. [Norm-based curriculum learning for neural machine translation](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 427–436, Online.
- Xuebo Liu, Longyue Wang, Derek F. Wong, Liang Ding, Lidia S. Chao, Shuming Shi, and Zhaopeng Tu. 2021e. [On the complementarity between pre-training and back-translation for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP*, pages 2900–2907.
- 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 pretraining approach](#). *arXiv preprint*, abs/1907.11692.
- Bin Lu, Myle Ott, Claire Cardie, and Benjamin K. Tsou. 2011. [Multi-aspect sentiment analysis with topic models](#). In *IEEE International Conference on Data Mining workshops*, pages 81–88.
- Dehong Ma, Sujian Li, Xiaodong Zhang, and Houfeng Wang. 2017. [Interactive attention networks for aspect-level sentiment classification](#). In *International Joint Conference on Artificial Intelligence*, pages 4068–4074.

- Robert C. Moore and William Lewis. 2010. [Intelligent selection of language model training data](#). In *Proceedings of the ACL Conference Short Papers*, pages 220–224, Uppsala, Sweden.
- Bo Pang and Lillian Lee. 2007. [Opinion mining and sentiment analysis](#). *Found. Trends Inf. Retr.*, 2(1-2):1–135.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Harris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [Semeval-2014 task 4: Aspect based sentiment analysis](#). In *International Workshop on Semantic Evaluation*, pages 27–35.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#). *OpenAI blog*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- 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](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Matthew Schultz and Thorsten Joachims. 2003. [Learning a distance metric from relative comparisons](#). *Advances in neural information processing systems*, 16.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 86–96, Berlin, Germany.
- Ronald Seoh, Ian Birle, Mrinal Tak, Haw-Shiuan Chang, Brian Pinette, and Alfred Hough. 2021. [Open aspect target sentiment classification with natural language prompts](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 6311–6322.
- Youwei Song, Jiahai Wang, Tao Jiang, Zhiyue Liu, and Yanghui Rao. 2019. [Attentional encoder network for targeted sentiment classification](#). *arXiv preprint*, abs/1902.09314.
- Kai Sun, Richong Zhang, Samuel Mensah, Yongyi Mao, and Xudong Liu. 2019. [Aspect-level sentiment analysis via convolution over dependency tree](#). In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 5678–5687.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Annual Conference on Neural Information Processing Systems*, pages 5998–6008.
- Kai Wang, Weizhou Shen, Yunyi Yang, Xiaojun Quan, and Rui Wang. 2020. [Relational graph attention network for aspect-based sentiment analysis](#). In *Annual Meeting of the Association for Computational Linguistics*, pages 3229–3238.
- Yufei Wang, Can Xu, Qingfeng Sun, Huang Hu, Chongyang Tao, Xiubo Geng, and Daxin Jiang. 2022. [Promda: Prompt-based data augmentation for low-resource NLU tasks](#). *arXiv preprint*, abs/2202.12499.
- Jason W. Wei and Kai Zou. 2019. [EDA: easy data augmentation techniques for boosting performance on text classification tasks](#). In *Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, pages 6381–6387.
- Xing Wu, Shangwen Lv, Liangjun Zang, Jizhong Han, and Songlin Hu. 2019. [Conditional BERT contextual augmentation](#). In *International Conference on Computational Science*, volume 11539 of *Lecture Notes in Computer Science*, pages 84–95.
- Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, and Zheng Zhang. 2021. [A unified generative framework for aspect-based sentiment analysis](#). In *Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 2416–2429.
- Changdong Zan, Liang Ding, Li Shen, Yu Cao, Weifeng Liu, and Dacheng Tao. 2022. [On the complementarity between pre-training and random-initialization for resource-rich machine translation](#). In *International Conference on Computational Linguistics (COLING)*.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. 2018. [mixup: Beyond empirical risk minimization](#). In *International Conference on Learning Representations*.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2021. [Towards generative aspect-based sentiment analysis](#). In *Annual Meeting of the Association for Computational Linguistics and the International Joint Conference on Natural Language Processing*, pages 504–510.
- Wenxuan Zhang, Xin Li, Yang Deng, Lidong Bing, and Wai Lam. 2022a. [A survey on aspect-based sentiment analysis: Tasks, methods, and challenges](#). *arXiv preprint*, abs/2203.01054.
- Zheng Zhang, Liang Ding, Dazhao Cheng, Xuebo Liu, Min Zhang, and Dacheng Tao. 2022b. [Bliss: Robust sequence-to-sequence learning via self-supervised input representation](#). *arXiv preprint*.
- Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, Hua Jin, and Dacheng Tao. 2022a. [Knowledge graph augmented network towards multiview representation learning for aspect-based sentiment analysis](#). *arXiv preprint*.

Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. 2022b. [PANDA: prompt transfer meets knowledge distillation for efficient model adaptation](#). *CoRR*, abs/2208.10160.

Deyu Zhou, Jianan Wang, Linhai Zhang, and Yulan He. 2021. [Implicit sentiment analysis with event-centered text representation](#). In *Conference on Empirical Methods in Natural Language Processing*, pages 6884–6893.

A Various Parameter-Efficient Methods Result

We adapt three parameter efficient methods for our T5 generator:

- **Prefix-tuning** (Li and Liang, 2021) fixes PLMs parameters, and inject an additional trainable prefix tokens for each transformer hidden layer. We set the number of prefix tokens to default 6.
- **Prompt-tuning** (Lester et al., 2021b) only prepends a prompt family to the input embedding layer. The length of the soft tokens is 100 in our experiments.
- **Low-Rank Adaptation (LoRA)** (Hu et al., 2022) trains a low-rank matrix to lightweight the PLMs. To be specific, we fix the hyper-parameter rank r to 8 and cancel the dropout operation ($dropout = 0$).
- **Full-tuning** fine-tunes all parameters in the T5 backbone.
- **No-tuning** directly generates sentences with a pre-trained T5 model without fine-tuning.




Figure 4: Convergence Analysis.

A.1 Case Study

More instances from different fine-tuning methods are listed in Table 6, the domains and polarities are attached to the source sentences. Under our observation, we found some phenomena: (1) the sentences from Full-tuning are always longer, and Full-tuning habitually outputs the same sentences, which proves that Full-tuning will cause an over-fit problem by conjecture. (2) three parameter efficient methods consistently generate multi-aspect samples, however, it is doubtful whether more aspects in a sentence will promote the model’s robustness more significantly. (3) most of the generated sentences can achieve our motivation (express an opposite polarity).

A.2 Convergence Analysis

We analyze the fine-tuning loss for those aforementioned fine-tuning methods (without No-tuning) with `T5ForConditionGeneration.loss` interface in *Hugging Face* on *Restaurant*. Figure 4 shows the loss curve, we adapt a percentile filter with 20 pts to conduct the curve smooth, and the losses of ultimate convergence are listed in Table 7. Obviously, Full-tuning will achieve a lower loss, because all parameters are tuned, and the scale of parameters is far greater than the number of data items, it is a precursor of an over-fitting problem. By comparing three fine-tuning methods, we prefer a method with lower convergence loss, Prefix-tuning and Prompt-tuning are indistinguishable, however, LoRA obtains an obvious superiority.

B More Sensitivity Analysis

Objective coefficients α and β are also controllable hyper-parameters, so we conduct sensitivity analysis on them. Our experimental results are shown in Figure 5 and Figure 6, the prediction model is consistently BERT-base and $\alpha \in \{0.1, 0.3, 0.5, 1, 2\}$, $\beta \in \{0.5, 1, 2, 5, 10\}$. According to the figures, the best performance always when $\alpha = 0.3$ or 0.5 and $\beta = 1$ or 2 .

C More Case Studies

Different DA methods We investigate different data augmentation frameworks from the view of the model’s prediction results in Section 4.3. We will also list an example to compare these DA baselines and throw more generated sentences from our C³DA in Table 8. The example in the upper table is from *Restaurant*, and we adapt LoRA to the T5

Table 6: Case study. The upper table compares results of different data augmentation methods, the bold spans is changed part; The below table shows some generated sentences with C³DA.

Method	Examples [dataset, polarity]
source	<i>not only was the food outstanding, but the little 'perks' were great.</i> [Restaurant, Positive]
No-tuning	<i>the kitchen is very well equipped with all the equipment you need to cook up a storm.</i>
Full-Tuning	<i>The food is okay and the prices here are mediocre.</i>
Prefix-tuning	<i>The kitchen is a little small, but the food is good.</i>
Prompt-tuning	<i>The kitchen is a little dated, but the food is good.</i>
LoRA	<i>The seats are a bit cramped, but the food was delicious.</i>
source	<i>it is easy to start up and does not overheat as much as other laptops.</i> [Laptop, Positive]
No-tuning	<i>a laptop that is easy to start up and does not overheat as much as other laptops.</i>
Full-Tuning	<i>I also purchased iWork to go with it which has programs for word processing, spreadsheets, and presentations (similar to Microsoft Office).</i>
Prefix-tuning	<i>The battery isn't very strong, but it is very light.</i>
Prompt-tuning	<i>but the screen is so dark, and the service center is a little numb.</i>
LoRA	<i>The software is very easy to use, the start up is very fast, the graphics are fantastic.</i>




Figure 5: Sensitivity Analysis of the objective coefficient α in supervised classification training.

Table 7: Convergence loss. "-T." is short for "-Tuning".

Method	Full-T.	Prefix-T.	Prompt-T.	LoRA
Loss	0.05102	0.70009	0.69182	0.46218

generator, and more examples of C³DA below are from *Restaurant* and *Laptop*.

We summarize that **EDA** deletes "my", transforms "co-workers" to its synonym "workers" and randomly inserts a comma to an unreasonable position. This method adds some corruptions without any semantic relations, and it is possible to generate an illogical sample; Compared to EDA, **CBERT** is also a token-level method, but a semantic-reserved sentence is consistently created; **Back Translation** is a sentence-level method, and a correct sentence can be generated. However, it erases the aspect words for ABSA task. By comparison, our C³DA not only reserves the origin semantic information but also introduces some high-quality in-domain corpus-aware information and adversarial signals. We also observe that most generated sentences from C³DA exist two or more aspects, and this phenomenon is widespread in several datasets.

Cases with Varying Fine-tuning Epochs τ We also explore the generation results with varying fine-tuning epochs τ . As shown in Table 9, given a source sentence "but the staff was so horrible to us.", the T5 generator tends to generate an irrelative sentence when we have not yet fine-tuned it. In the middle of the fine-tuning stage, the generator is inclined to retell the source sentences or generate some irrational sentences. Finally, our generation model will eventually converge when $\tau = 80$.

D Instance-level Loss Re-weight for Long-tail Aspects

There is an another problem when we implement sentence generation, our generated sentences in one dataset often fit to few aspects, such as *food* in *Restaurant*. To investigate this issue, we show the aspect information for three public ABSA datasets in Table 10, and we found that *food* in *Restaurant* actually maintain maximum frequency. Similarly, *service*, *battery* and *screen* etc. report this long-tail distribution problem. And most of aspects in *Twitter* is [UNK], it means that extensive aspect information will be unusable, this causes some models don't work, especially the attention-based models.

Similar to focal loss (Lin et al., 2017), we adapt




Figure 6: Sensitivity Analysis of the contrastive training objective coefficient β in our overall objective.

Table 8: Case study. The upper table compares results of different data augmentation methods, the blue spans is changed part; The below table shows some generated sentences with C^3DA .

DA methods	examples
source	<i>my co-workers had to wait almost an hour for delivery.</i>
EDA	<i>workers had to, wait an hour for delivery.</i>
B.T.	<i>my colleague waited nearly an hour to deliver it.</i>
CBERT	<i>my mother had to wait half an hour for delivery.</i>
C^3DA	<i>my co-workers had to wait almost an hour for delivery. but the perks are great, and you get to enjoy all the perks.</i>

source sentences	C^3DA generated sentences
<i>if you love seafood, you would love this place!</i>	<i>the sauce was good, but the meats were a little fatty.</i>
<i>we had a girls, night dinner here for restaurant week.</i>	<i>try the samosas at a reasonable price.</i>
<i>it is easy to start up and does not overheat as much as other laptops.</i>	<i>but the screen is so dark, and the service center is a little numb.</i>

Table 9: Case study. Given a source sentence “*but the staff was so horrible to us.*”, we can observe the trend of its augmented sentence with fine-tune epochs τ .

epochs τ	sentence
source	<i>but the staff was so horrible to us.</i>
0	<i>i cried out for the suffragellah woman.</i>
20	<i>the staff is so horrible to us.</i>
40	<i>the as-is, the service isn't great.</i>
60	<i>the food is very good.</i>
80	<i>the food was great, the service was excellent.</i>

Table 10: The aspect words and its frequency statistic of three public datasets. We only show the top-10 aspect words (with the T5 tokenizer) and hide some special tokens such as ‘a’, ‘i’, etc.

dataset	aspect words: frequency
Restaurant	<i>food</i> : 419, [UNK]: 404, <i>service</i> : 204, <i>wait</i> : 94, <i>menu</i> : 80, <i>dinner</i> : 75, <i>wine</i> : 72, <i>staff</i> : 68, <i>pizza</i> : 65, <i>place</i> : 60
Laptop	[UNK]: 267, <i>battery</i> : 97, <i>screen</i> : 81, <i>use</i> : 58, <i>life</i> : 56, <i>windows</i> : 55, <i>price</i> : 55, <i>software</i> : 54, <i>keyboard</i> : 54, <i>drive</i> : 53
Twitter	[UNK]: 5328, <i>spear</i> : 894, <i>bri</i> : 887, <i>ney</i> : 887, <i>lo</i> : 394, <i>han</i> : 393, <i>lind</i> : 388, <i>say</i> : 388, <i>in</i> : 370, <i>ry</i> : 360

an instance-level loss re-weight method to our fine-tuning stage. Inspired by Jain et al. (2016), we multiple a multiplier to each instance’s objective, the multiplier is shown in Equation (8).

$$\Delta_j = \frac{1}{1 + Ce^{-A \log(M_j + B)}}, \quad (8)$$

$$C = (\log M - 1)(B + 1)^A, \quad (9)$$

where M_j , M are the number of data points annotated with aspect j , and aspect items, respectively.