Paraphrase Makes Perfect: Leveraging Expression Paraphrase to Improve Implicit Sentiment Learning

Xia Li[1,2], Junlang Wang[1]*, Yongqiang Zheng[1], Yuan Chen[1] and Yangjia Zheng[1]
[1]School of Information Science and Technology,
[2]Center for Linguistics and Applied Linguistics,
Guangdong University of Foreign Studies, Guangzhou, China
{xiali,junlangwang,yqzheng,yuanchen,yjzheng}@gdufs.edu.cn

Abstract

Existing implicit sentiment learning methods mainly focus on capturing implicit sentiment knowledge individually, without paying more attention to the potential connection between implicit and explicit sentiment. From a linguistic perspective, implicit and explicit sentiment expressions are essentially similar when conveying the same sentiment polarity for a specific aspect. In this paper, we present an expression paraphrase strategy and a novel sentiment-consistent contrastive learning mechanism to learn the intrinsic connections between implicit and explicit sentiment expressions and integrate them into the model to enhance implicit sentiment learning. We perform extensive experiments on public datasets, and the results show the significant efficacy of our method on implicit sentiment analysis.[1]

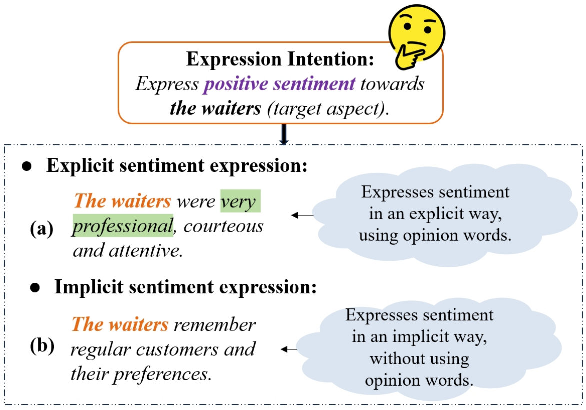

Figure 1: Illustration of explicit and implicit sentiment expressions and their potential connection.

1 Introduction

Aspect-based sentiment analysis (ABSA) is a fine-grained sentiment analysis task that focuses on inferring the sentiment polarities towards specific aspects in texts. It can be divided into explicit sentiment analysis (ESA) and implicit sentiment analysis (ISA) based on whether the text contains explicit opinion expressions or polarity markers for the target aspect (Liu, 2012; Pontiki et al., 2014; Russo et al., 2015). For example, as shown in Figure 1, sentences (a) and (b) are explicit and implicit sentiment expressions, respectively. In this work, we focus on ISA, which is much more challenging as the texts contain no explicit opinion expression.

Considering that the external knowledge (e.g., sentiment lexical knowledge and commonsense knowledge) is crucial for helping understand sentiment cues, how to leverage and interpret this external information is important for implicit sentiment analysis. Previous studies have explored the ability of models to learn and interpret implicit sentiment knowledge from different perspectives: 1) Leveraging lexical knowledge to assist in acquiring implicit sentiment information, where the core idea is using the lexicons to find opinion terms related to certain terms in implicit sentiment expressions (Deng and Wiebe, 2014; Choi and Wiebe, 2014). 2) Constructing specific patterns for learning implicit sentiment, where the main idea is to integrate linguistic information in human's expression (Wang et al., 2022, 2023a). 3) Distilling the knowledge from the external corpus or LLMs to improve the perception of implicit sentiment. For example, Li et al. (2021) design additional pre-training for the models to update them with external sentiment knowledge. Fei et al. (2023) propose a chain-of-thought (CoT) based method to encourage LLMs to learn implicit sentiment from intermediate reasoning responses.

Although the above works have achieved promising results, they focus more on modeling the knowledge of implicit sentiment individually, paying little attention to the potential connection between implicit and explicit sentiment. In fact, though people can express sentiment in direct and indirect ways, the intention of expressing sentiment is the

*Corresponding author.
[1]Code is released in https://github.com/gdufsnlp/SECP.

Proceedings of the 31st International Conference on Computational Linguistics, pages 3631–3647
January 19–24, 2025. ©2025 Association for Computational Linguistics

same. As shown in Figure 1, both sentences (a) and (b) can be used by humans to express a positive sentiment towards the target aspect "*the waiters*", where (a) is expressed in an explicit form and (b) is expressed in an implicit form. This shows that although the forms of expression are different, the sentiments are expressed consistently. It can be regarded as an intrinsic connection between explicit and implicit sentiment expressions, revealing that their semantics are similar and can transform into each other when expressing the same sentiment polarities towards the aspects.

So a natural question arises: can we leverage this sentiment consistency connection between explicit and implicit expressions to help identify implicit sentiment effectively? To address this issue, we propose a novel **S**entiment **E**xpression **C**onversion based **P**araphrase method (SECP) for implicit sentiment learning, which contains two components: expression paraphrase module and sentiment-consistent contrastive learning module. Expression paraphrase module aims to construct pairs of sentences containing both explicit and implicit sentiment expressions. Based on these pairs, sentiment-consistent contrastive learning module is designed to learn the intrinsic connections between implicit and explicit sentiment expressions and integrate them into the model to improve implicit sentiment learning.

In expression paraphrase module, we paraphrase sentences by transforming their implicit (explicit) sentiment expressions into explicit (implicit) ones while preserving the aspect's original sentiment polarity. By doing this, we can collect pairs of original and paraphrased sentences for each ABSA example, containing both explicit and implicit sentiment expressions. Motivated by the paraphrasing and instruction following capabilities of LLMs (Xue et al., 2023; Ouyang et al., 2022), we use the LLM to implement the paraphrase with our designed prompts, which contain paraphrased-based demonstrations and linguistic hints, guiding the LLM to generate the paraphrased sentences effectively.

In sentiment-consistent contrastive learning module, to capture the connection between explicit and implicit sentiment expressions with the above pairs of sentences, we design sentiment consistency supervision signals to determine the relation between different ABSA examples. It considers both the sentiment polarities towards the aspects and the forms of sentiment expressions (i.e. explicit or implicit sentiment expressions). We define four types of sentiment consistency hierarchically and further conduct contrastive learning (Khosla et al., 2020), pulling sentiment representations with close consistency together and pushing apart non-consistent sentiment expressions. In addition, we design the contextualized self-alignment module to promote the model to capture the connection between explicit and implicit sentiment expressions while learning contextual information of input sentences. To sum up, the contributions of this work are as follows:

- We propose an expression paraphrase strategy to construct pairs of sentences containing explicit and implicit sentiment expressions to explore the potential connection between them.
- We propose sentiment-consistent contrastive learning mechanism to better learn the connection between implicit and explicit sentiment expressions and improve implicit sentiment learning in ABSA.
- We conduct comprehensive experiments on public datasets, and the results demonstrate the efficacy and versatility of our method.

2 Our Approach

2.1 Overall Architecture

As shown in Figure 2, the overall architecture of our approach includes the expression paraphrase strategy module, encoding module, sentiment-consistent contrastive learning module, and contextualized self-alignment module. Our work is applicable to various PLMs, and we will take BERT (Devlin et al., 2019) as an example and describe the modules in the following subsections.

2.2 Task Definition

In ABSA, each sentence contains one or more aspects corresponding to multiple sentiments. In this paper, we focus on analyzing the sentiment polarity towards a target aspect at each step. Given a sentence $x_i = \{w_1, \dots, w_t, w_{a1}, \dots, w_{am}, w_{t+1}, \dots, w_n\}$, w_t is the t-th word and the target aspect a_i is denoted as $a_i = \{w_{a1}, \dots, w_{am}\}$. The goal of our model is to predict the sentiment polarity y_i towards the target aspect a_i.

2.3 Expression Paraphrase Strategy Module

To explore the connection between explicit and implicit sentiment expressions, we propose a novel expression paraphrase strategy using ChatGPT (Ope-

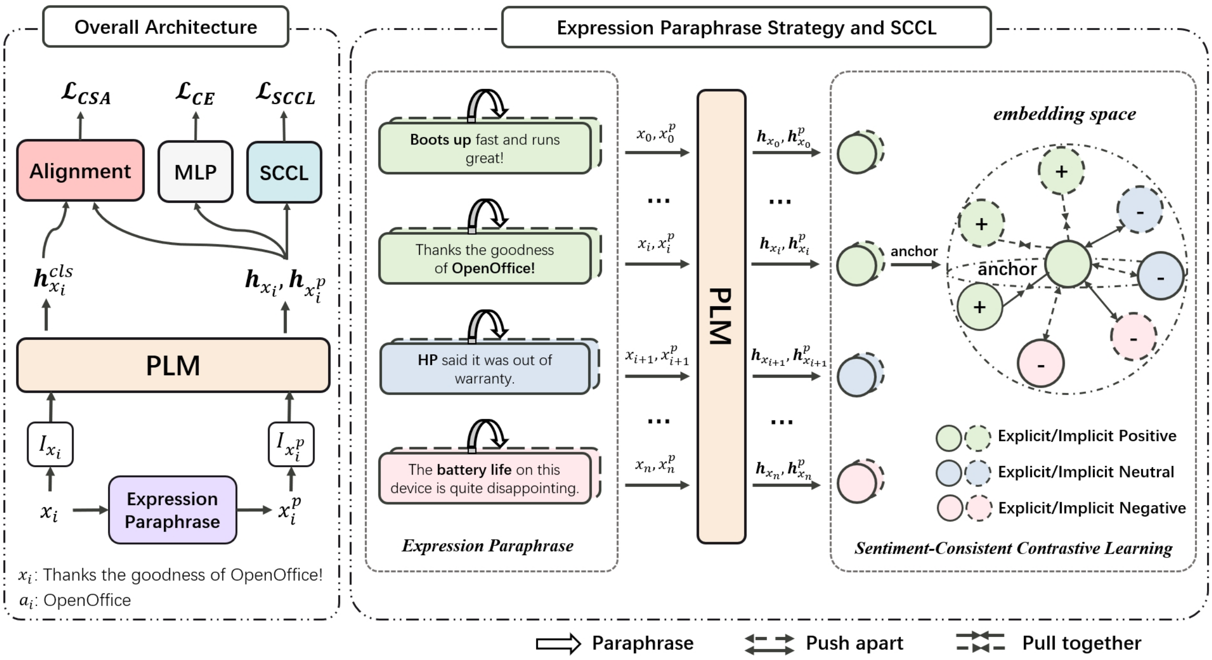

Figure 2: Overall architecture of the proposed approach. We illustrate the expression paraphrase strategy and sentiment-consistent contrastive learning (**SCCL**) as a whole and take the paraphrase from explicit to implicit sentiment expressions as examples (Actually, the paraphrase in the opposite direction will also be conducted). "Alignment" is the contextualized self-alignment. x_i and x_i^p are the original and paraphrased sentences of the i-th example in the mini-batch and a_i is the target aspect. I_{x_i} and $I_{x_i^p}$ are input sequences of x_i and x_i^p in the encoding module. $\mathbf{h}_{x_i}$ and $\mathbf{h}_{x_i^p}$ are the representations of "[MASK]" from I_{x_i} and $I_{x_i^p}$. $\mathbf{h}_{x_i}^{cls}$ is the representation of "[CLS]" from I_{x_i}. We illustrate them in Section 2.4. The process of the paraphrase is detailed in Figure 3. "+" and "-" are positive and negative instances for the anchor. The relation between the anchor and different instances is marked by different arrows, which is detailed in Section 2.5.

nAI, 2023). It transforms sentences that convey implicit (explicit) sentiment into paraphrased sentences that express explicit (implicit) sentiment, resulting in pairs of original and paraphrased sentences. We design effective prompts to instruct ChatGPT to conduct the paraphrase, which consist of two components: (*i*) paraphrase-based demonstrations and (*ii*) paraphrase-relevant linguistic hints. Details of the prompts are in Appendix D.

Paraphrase-based Demonstrations Generally, when prompting LLMs for sentiment analysis, recent efforts adopt k examples for each class accompanied by their gold labels (Zhang et al., 2023a; Wang et al., 2023c). Alternatively, we treat pairs of sentences as demonstrations, which guide the LLM in generating the paraphrased sentences effectively. For example, for a sentence expressing explicit positive sentiment towards "*OpenOffice*", we select another sentence that conveys implicit positive sentiment towards "*OpenOffice*" and pair them. Then, we design the instruction of the paraphrase from explicit positive to implicit positive sentiment for this pair and combine them as follows:

> *Explicit positive sentiment towards the aspect* "OpenOffice": "Thank goodness for OpenOffice!"=> *Implicit positive sentiment towards the aspect* "OpenOffice": "OpenOffice has been my lifesaver for a while!"

The process of the paraphrase is also shown in Figure 3. It is worth noting that few examples in the training data convey the same sentiment polarity towards a specific aspect but use different implicit and explicit expressions. Thus, we only consider 1-shot paraphrase-based demonstrations and manually proofread them, as shown in Appendix D.

Paraphrase-relevant Linguistic Hints Besides the demonstrations, we believe that LLMs can also benefit from linguistic knowledge in the paraphrase. Thus, we incorporate paraphrase-relevant linguistic hints into the designed prompts, which explain the difference between explicit and implicit sentiment expressions. Inspired by Yu et al. (2023), we employ ChatGPT to generate the attributes in distinguishing explicit and implicit sentiment expressions. It can be achieved by prompting ChatGPT

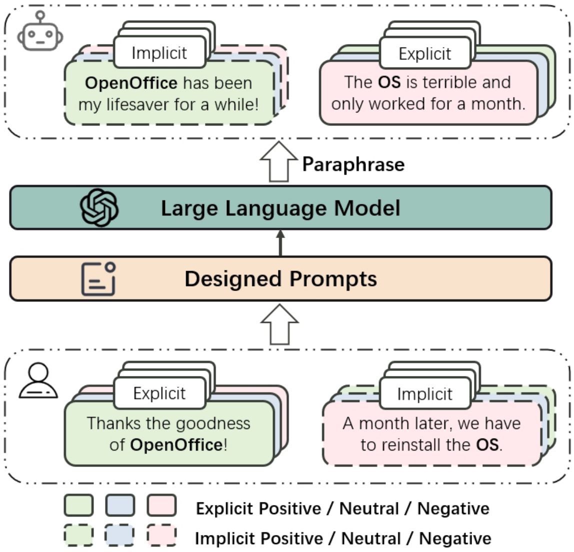

Figure 3: The process of the proposed paraphrase. We show examples from the domain of laptops.

such as "*What attributes do you think are crucial for distinguishing explicit sentiment from implicit sentiment?*". As a result, several attributes and their explanations are generated as the response. Among them, we select "Directness of expression", "Use of emotional cues" and "Level of recommendation or rejection" as the hints and inject them into the designed prompts. For example, the first hint is:

> *Directness of expression: Explicit sentiment directly expresses opinions and emotions, while implicit sentiments may use indirect or suggestive language.*

The prompts for this process are illustrated in Appendix D and the details of selected hints are shown in Table 11 of Appendix D.

2.4 Encoding Module

Considering that utilizing the LLM to paraphrase may lead to a drawback that the aspect terms in sentences might be replaced by their synonyms, we append a cloze prompt to the sentence and reconstruct the input sequence to help recognize aspects that are implied or worded differently (Seoh et al., 2021). For the original sentence x_i, we first use our proposed expression paraphrase strategy to get its paraphrased sentence x_i^p and treat it as a specific sentence, sharing a_i and y_i with x_i. We then reconstruct the sentences x_i and x_i^p as I_{x_i} and $I_{x_i^p}$ as equations (1) and (2). We use BERT to encode the two input sequences and get $\mathbf{h}_{x_i} \in \mathbb{R}^d$ and $\mathbf{h}_{x_i^p} \in \mathbb{R}^d$ as equations (3) and (4), where d is the hidden size.

$$I_{x_i} = \text{[CLS] } x_i \text{ [SEP] The } a_i \text{ is [MASK]} \quad (1)$$

$$I_{x_i^p} = \text{[CLS] } x_i^p \text{ [SEP] The } a_i \text{ is [MASK]} \quad (2)$$

$$\mathbf{h}_{x_i} = \text{BERT}(I_{x_i})_{\text{[MASK]}} \quad (3)$$

$$\mathbf{h}_{x_i^p} = \text{BERT}(I_{x_i^p})_{\text{[MASK]}} \quad (4)$$

2.5 Sentiment-consistent Contrastive Learning Module

To model implicit sentiment with the connection between explicit and implicit sentiment expressions, we propose sentiment-consistent contrastive learning (SCCL) and design a supervision named "sentiment consistency", which considers both the sentiment polarity and the sentiment expression category $e_i \in \{explicit, implicit\}$ of the sentence. Intuitively, sentiment polarities towards aspects from the sentences are the main factor distinguishing different examples in ABSA. The sentiment expression category also plays an important role, which can help to find examples that express sentiment in similar and different styles. Thus, for different examples, we define four types of sentiment consistency to determine their relation. Specifically, for each example in a mini-batch $\mathcal{B}$, we determine its relation with other examples by measuring their sentiment consistency considering both sentiment polarity and sentiment expression category. The functions for calculating them are denoted by $cp(\cdot)$ and $ce(\cdot)$, with their outputs limited to 1 or -1, indicating identical or different sentiment polarities and sentiment expression categories respectively. For the example x_i and x_j, the consistency of sentiment polarity and sentiment expression category are denoted by $\alpha_i^j = cp(y_i, y_j)$ and $\beta_i^j = ce(e_i, e_j)$, which are both involved in determining the relation γ_j^i between x_i and x_j. In this setting, γ_j^i is a hierarchical coefficient that determines the positive and negative instances of x_i, which is defined as:

$$\gamma_j^i = \begin{cases} 1.00 & \alpha_i^j = 1, \beta_i^j = 1 \\ \delta_1 & \alpha_i^j = 1, \beta_i^j = -1 \\ \delta_2 & \alpha_i^j = -1, \beta_i^j = 1 \\ 0.00 & \alpha_i^j = -1, \beta_i^j = -1 \end{cases} \quad (5)$$

where δ_1 and δ_2 are set as hyperparameters with values between 0 and 1, representing two types of weighted supervision. Considering the prioritization of sentiment consistency mentioned above, we set two types of positive and negative instances

Dataset	Train				Test				Total			
	Pos	Neu	Neg	IS(%)	Pos	Neu	Neg	IS(%)	Pos	Neu	Neg	IS(%)
Laptop	987	460	866	30.87	341	169	128	27.27	1328	629	994	30.09
Restaurant	2164	633	805	28.59	728	196	196	23.84	2892	829	1001	27.47

Table 1: Statistics on two benchmarks where "Pos", "Neu" and "Neg" are *positive*, *neutral* and *negative* respectively. "IS" represents the percentage of samples that convey implicit sentiment to the total.

corresponding to the first two terms and the last two terms in equation (5) respectively. For each example consisting of the original sentence, the proposed paraphrase strategy explicitly constructs a positive instance for it as shown in Figure 2 while extending its positives and negatives by increasing the in-batch instances (When performing SCCL, the size of the mini-batch $\mathcal{B}$ doubles from N to $2N$ since both original and paraphrased sentences are considered simultaneously. In other words, x_i and x_i^p are regarded as two independent examples in this module). Thus, the SCCL loss $\mathcal{L}_{SCCL}$ is designed as follows:

$$P(i,j) = \frac{e^{sim(\mathbf{h}_{x_i},\mathbf{h}_{x_j})/\tau}}{\sum_{k=1}^{2N} e^{sim(\mathbf{h}_{x_i},\mathbf{h}_{x_k})/\tau}} \quad (6)$$

$$\mathcal{L}_{SCCL} = -\sum_{i=1}^{2N} \frac{1}{\sum_{j=1}^{2N} \gamma_j^i} \sum_{j=1}^{2N} \gamma_j^i \log P(i,j) \quad (7)$$

Here, $P(i,j)$ is the normalized similarity between x_i and x_j where $sim(\mathbf{h}_{x_i}, \mathbf{h}_{x_j})$ is the cosine similarity and τ is the temperature. Moreover, γ_j^i scales the loss contribution of each pair of examples.

2.6 Contextualized Self-alignment Module

While boosting the representation modeling for explicit and implicit sentiment, the supervised contrastive loss brings the potential risk of representation collapse (Graf et al., 2021; Chen et al., 2022). In particular, the model might focus too much on other sentiment expressions under SCCL, neglecting the contextual information from the sentences. To alleviate this deviation, we propose contextualized self-alignment (CSA) based on the Kullback-Leibler (KL) divergence and integrate it into our approach. Specifically, within the mini-batch $\mathcal{B}$, this alignment measures the KL divergence between the representations of "[CLS]" and "[MASK]" tokens from the input sequences. Then the model would pay more attention to contextual information, treating $\mathcal{L}_{CSA}$ as one of the training losses:

$$\mathcal{L}_{CSA} = \frac{1}{N} \sum_{i=1}^{N} \mathcal{D}_{KL}(\mathbf{h}_{x_i} \,||\, \mathbf{h}_{x_i}^{cls}) \quad (8)$$

where $\mathbf{h}_{x_i}^{cls}$ is the representation of "[CLS]" and $\mathcal{D}_{KL}(\cdot)$ measures KL divergence between two representations. Notably, examples derived from the original sentences but not paraphrased sentences are involved, which is complementary to SCCL.

2.7 Joint Training

Besides the losses mentioned above, cross-entropy loss $\mathcal{L}_{CE}$ is employed for sentiment classification:

$$\mathcal{L}_{CE} = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{3} y_i^j \log \hat{y}_i^j + \lambda ||\theta||^2 \quad (9)$$

where y_n is the gold label distribution and $\hat{y}_n$ is the estimated distribution. λ and θ represent the coefficient of L_2 regularization and trainable parameters. It is worth noting that only training samples corresponding to original sentences are involved here. The overall training loss $\mathcal{L}$ for the model is:

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda_1 \mathcal{L}_{SCCL} + \lambda_2 \mathcal{L}_{CSA} \quad (10)$$

Here, λ_1 and λ_2 are the weights of SCCL and CSA.

In addition, the training and inference stages of our approach are different. During the training stage, the training sentences need to be paraphrased using our proposed expression paraphrase strategy to train the model. During the inference stage, a test sentence and the target aspect can be input into the trained model to get the predicted sentiment polarity directly.

3 Experiments

3.1 Datasets

The experiments are conducted on SemEval 2014 Laptop and Restaurant (Pontiki et al., 2014). Their Explicit Sentiment Expression (ESE) and Implicit

Sentiment Expression (ISE) slices are used to evaluate the explicit and implicit sentiment predictions, which are provided by Li et al. (2021). We also use the annotations of explicit and implicit sentiment expressions from them to construct paraphrased-based demonstrations in Section 2. The statistics of datasets are shown in Table 1.

3.2 Implementation Details

In the experiments, for the proposed paraphrase, we set the temperature coefficient of ChatGPT as 0.7. For the proposed approach, we fine-tune BERT-base-uncased and RoBERTa-base pre-trained by HuggingFace Transformers (Wolf et al., 2020) and implemented by PyTorch (Paszke et al., 2019). During the training, the learning rate is set as 5e-5 and the batch size is 16. We adopt AdamW (Loshchilov and Hutter, 2017) for the training and the coefficient of L_2 regularization is 0.015 (SECP-BERT) and 1e-5 (SECP-RoBERTa). In the proposed SCCL, the temperature coefficient τ is 0.1 and the values of δ_1 and δ_2 in the hierarchical supervised signal γ_j^i are set as 0.90 and 0.05. For joint learning, the coefficients λ_1 and λ_2 are 0.8 and 0.2. We report the average performance of models over five runs in the experiments.

3.3 Baselines

We evaluate our models (referred to as "SECP-BERT" and "SECP-RoBERTa") by comparing them to various models: **Large Language Models**[2], which consider zero-shot and few-shot (1-shot is set following Section 2.3) inference for LLMs such as ChatGPT (OpenAI, 2023), GPT-4 (Achiam et al., 2023) and Qwen2.5 (Bai et al., 2023). The input prompts to LLMs are detailed in Appendix D. **PLM-based Models**, which adapt PLMs to ABSA, including BERT+ISAIV (Wang et al., 2022), IPOS-BERT (Wang et al., 2023a), model from Wang et al. (2023b) (we select the distillation-based model and denote it by "BERT-ED"), LSA$_p$-BERT and LSA$_p$-RoBERTa (Yang and Li, 2024). Besides, some models are also pre-trained with additional large-scale annotated corpora, such as BERT-PT (Xu et al., 2019), BERT-ADA (Rietzler et al., 2020) and TransEncAsp+SCAPT (Li et al., 2021). The state-of-the-art (SoTA) ISA model Flan-T5+THOR (250M) (Fei et al., 2023) is also considered.

[2]We utilize the gpt-3.5-turbo-1106, gpt-4-0613, qwen-max-0428 version of ChatGPT, GPT-4 and Qwen2.5.

3.4 Main Results

Table 2 shows the performance of baselines and our models. We denote the accuracy and macro F1 score by "ACC" and "MF1". "ESE" and "ISE" indicate the accuracy on ESE and ISE slices. Following the results, some findings can be observed:

Our models show a notable advantage in implicit sentiment prediction compared to LLMs. While ChatGPT performs well in ESE, its ability in implicit sentiment prediction is less satisfactory, with ISE of 51.43% and 56.18% on two benchmarks on the few-shot setting. However, our model SECP-BERT shows significant boosts in implicit sentiment prediction compared with ChatGPT (few-shot), with increases of around 27% and 15% on the ISE slices of Laptop and Restaurant. Moreover, we also consider the comparison with other powerful LLMs such as GPT-4 and Qwen2.5. Our model SECP-RoBERTa achieves competitive and SoTA performance, as shown by its results of accuracy, MF1 and ISE on two benchmarks. However, it has significantly fewer parameters than these LLMs and is more feasible to deploy. The selected LLMs show strong explicit sentiment prediction abilities on Laptops. We think that they might benefit from the large-scale pre-training corpus. An unexpected observation is that three LLMs show minimal improvement on few-shot setting, which might be limited by the number of demonstrations.

Our models achieve strong performance in implicit sentiment prediction without additional pre-training. Despite only conducting fine-tuning, our proposed models outperform in implicit sentiment prediction. Our proposed SECP-BERT achieves 78.86%/71.00% while TransEncAsp+SCAPT achieves 72.82%/68.55% on ISE on Laptop and Restaurant. Moreover, SECP-BERT achieves higher MF1 and ISE compared to BERT-ADA, which is also based on BERT, with improvements of 4%/1.5% (MF1) and 8%/5% (ISE) in Laptop and Restaurant. The results indicate the feasibility of improving implicit sentiment learning without relying on additional pre-training.

Our models outperform other PLM-based models in most cases. Compared with the SoTA BERT-based models BERT+ISAIV and LSA$_p$-BERT, our SECP-BERT achieves the 1.4%/0.5% gains of accuracy and 0.8%/0.3% gains of MF1 on Laptop. SECP-BERT also outperforms BERT+ISAIV on ISE slices of Laptop and Restaurant by around 0.6% and 1.3%, indicating its advantage in im-

Model	Laptop				Restaurant			
	ACC	MF1	ESE	ISE	ACC	MF1	ESE	ISE
Large Language Models								
ChatGPT (zero-shot)	76.65	67.61	87.08	48.57	81.25	62.84	91.91	47.19
ChatGPT (few-shot)	76.80	68.43	86.42	51.43	83.39	68.67	91.91	56.18
GPT-4 (zero-shot)	83.86	79.65	**87.26**	74.86	87.14	79.59	92.36	70.63
GPT-4 (few-shot)	83.07	78.34	**87.26**	72.00	87.14	78.45	93.07	68.40
Qwen2.5 (zero-shot)	82.60	78.48	86.83	71.43	87.41	79.50	92.61	70.79
Qwen2.5 (few-shot)	80.56	75.15	86.39	65.14	87.50	78.80	92.73	70.79
PLM-based Models								
BERT-PT (Xu et al., 2019)†	78.07	75.08	81.47	71.27	84.95	76.96	92.15	64.79
BERT-ADA (Rietzler et al., 2020)†	78.96	74.18	82.76	70.11	87.14	80.05	91.14	65.92
TransEncAsp+SCAPT (Li et al., 2021)†	77.17	73.23	78.70	72.82	83.39	74.53	88.04	68.55
BERT+ISAIV (Wang et al., 2022)†	80.41	77.25	81.21	78.29	87.05	81.40	92.50	69.66
IPOS-BERT (Wang et al., 2023a)	80.56	76.99	81.21	76.00	85.83	79.41	91.44	66.66
BERT-ED (Wang et al., 2023b)‡	78.68	75.19	-	-	84.66	76.18	-	-
LSA$_p$-BERT (Yang and Li, 2024)‡	81.35	77.79	-	-	87.23	81.06	-	-
SECP-BERT (Ours)	81.82	78.11	82.94	78.86	86.88	81.52	91.89	71.00
Flan-T5+THOR (Fei et al., 2023)	-	79.75	-	67.63	-	**82.98**	-	71.70
LSA$_p$-RoBERTa (Yang and Li, 2024)‡	83.39	80.47	-	-	88.04	82.96	-	-
SECP-RoBERTa (Ours)	**84.48**	**81.81**	84.67	**84.00**	**88.21**	81.89	**93.20**	**72.28**

Table 2: Overall results (%) of the models on two benchmarks. †: Results retrieved from Wang et al. (2022). ‡: Results on ESE and ISE are missed since they do not focus on implicit sentiment learning. Our proposed models are based on the base versions of the corresponding PLMs. The best results among the models are highlighted in bold. ESE and ISE are used to evaluate explicit and implicit sentiment prediction and their values are for accuracy.

Variants	Restaurant			
	ACC	MF1	ESE	ISE
SECP-BERT	86.88	81.52	91.89	71.00
w/o $\mathcal{L}_{SCCL}$	85.00	77.30	92.26	61.80
w/o $\mathcal{L}_{CSA}$	85.45	78.03	92.15	64.04
BERT+PRs	85.71	79.69	91.79	65.92

Table 3: Ablation study on Restaurant.

plicit sentiment learning. Our SECP-RoBERTa achieves higher MF1 and ISE on Laptop compared to Flan-T5+THOR, with increases of around 2.1% and 16.3%. It reveals that learning intrinsic connections between explicit and implicit sentiment expressions contributes to the improvement in ABSA. **Our method can better leverage the knowledge from ChatGPT in ABSA**. This could be verified by the comparisons of the proposed SECP-BERT with BERT-ED, which treats the explanations generated from ChatGPT as additional training data. Though both use the same LLM, our SECP-BERT outperforms it by large margins, achieving around 3.1% and 2.2% gains in accuracy as well as 2.9% and 5.3% gains in MF1 on Laptop and Restaurant. The difference might be attributed to paraphrased sentences being more suitable for fine-tuning than explanations, as the latter are declarative long texts and less similar to ABSA training instances.

4 Discussion

4.1 Ablation Studies

We explore the efficacy of the modules from our approach. To mine the connection between the proposed paraphrase and SCCL, we also train a comparative model named "BERT+PRs", which directly treats the paraphrased sentences as training samples with cross-entropy loss, similar to traditional data augmentation strategies. As shown in Table 3, removing SCCL degrades the proposed SECP-BERT on Restaurant drastically, which is more significant than removing CSA. It reveals that SCCL enhances our model primarily and CSA plays the role of complementing it. Besides, SECP-BERT achieves 81.52% (MF1) and 71.00% (ISE), while BERT+PRs reaches 79.69% (MF1)

Model	Params	Laptop				Restaurant			
		ACC	MF1	ESE	ISE	ACC	MF1	ESE	ISE
ALBERT (Lan et al., 2019)	12M	78.68	74.88	81.21	72.00	81.79	73.94	88.04	61.80
SECP-ALBERT (Ours)	12M	80.09	76.75	81.90	75.43	85.80	79.37	92.26	65.17
BERT (Devlin et al., 2019)	110M	77.90	73.37	82.11	66.86	84.20	76.24	90.62	63.67
SECP-BERT (Ours)	110M	81.82	78.11	82.94	78.86	86.88	81.52	91.89	71.00
RoBERTa (Liu et al., 2019)	125M	81.97	78.38	83.18	78.86	87.23	81.00	93.08	68.54
SECP-RoBERTa (Ours)	125M	84.48	81.81	84.67	84.00	88.21	81.89	93.20	72.28

Table 4: Comparison of different PLMs and their variants equipped with our approach. For the selected PLMs, we use their base versions. "**Params**" represents the trainable parameters of the model.

Model	Laptop	Restaurant
	Ori→New	Ori→New
BERT-PT	78.53→53.29	86.70→59.29
BERT+ISAIV	80.41→59.45	87.05→58.77
SECP-BERT	81.82→69.85	**86.88→80.45**
SECP-RoBERTa	**84.48→76.29**	88.21→80.57

Table 5: The results of robustness evaluation.

and 65.92% (ISE) on Restaurant. It indicates that directly treating the paraphrased sentences as training samples is unsatisfactory and sheds light on the synergy between the paraphrase and SCCL. In Appendix A, we provide the ablation studies about effects of the prompt components on the paraphrase.

4.2 Generalization to Different PLMs

To verify the generalization of the proposed approach, we equip it in different PLMs and evaluate them on the benchmark datasets Laptop and Restaurant (Pontiki et al., 2014). Besides BERT, we select ALBERT (Lan et al., 2019) and RoBERTa (Liu et al., 2019) as the backbones of our approach. As shown in Table 4, for all selected PLMs, our method boosts their performance significantly including predicting both explicit and implicit sentiment, as shown by the improvements of all metrics. It demonstrates that our approach is adaptable to models with different parameter scales.

4.3 Analysis of Robustness

We analyze the robustness of our models on the aspect robustness test sets, which are based on the datasets mentioned above and proposed by Xing et al. (2020). Since they contain a variety of perturbed examples, the robustness of the model can be fully evaluated. We compare the performance of several models on original and aspect robustness test sets, which are evaluated by accuracy and shown in Table 5. It reveals that our models suffer the least performance degradation, showing 8.19% (SECP-RoBERTa) and 6.43% (SECP-BERT) decline on Laptop and Restaurant, which might be attributed to our models learning diverse expressions from the proposed paraphrase. Appendix B shows that our models are also domain robust.

4.4 Case Study

To investigate the advantage of our model compared with LLMs in ABSA and implicit sentiment prediction, we show several examples from Laptop and Restaurant in Table 6. For the first example, ChatGPT, GPT-4 and our SECP-RoBERTa correctly predict the sentiment polarities conveyed by explicit sentiment expressions. For the third example, ChatGPT misunderstands the sentence and might not capture the opinion term "gripe" which expresses negative sentiment towards the aspect "**RAM**". Consequently, it predicts the neutral sentiment. In example 6, both ChatGPT and GPT-4 fail to predict the sentiment polarity towards "**price**" but our proposed model correctly infers it as positive, which can perceive that the price drop of the laptop pleases customers. The comparisons indicate the ability of our model in implicit sentiment prediction. In example 2, GPT-4 unexpectedly predicts neutral sentiment towards "**rose roll**". We think that it might be affected by the problem of "over-alignment" (Zhang et al., 2023a).

5 Related Work

Implicit Sentiment Learning in ABSA Recent efforts on ABSA focus on opinion terms related to the aspects through syntactical analysis (Zhang et al., 2019; Wang et al., 2020) and utilize sen-

No.	Example	ChatGPT	GPT-4	SECP-RoBERTa
1	The **food**$_{[exp]}$ is good, the **teriyaki**$_{[exp]}$ I recommend.	(Pos$_\checkmark$, Pos$_\checkmark$)	(Pos$_\checkmark$, Pos$_\checkmark$)	(Pos$_\checkmark$, Pos$_\checkmark$)
2	Try the **rose roll**$_{[exp]}$ (not on **menu**$_{[imp]}$).	(Pos$_\checkmark$, Neg$_\times$)	(Neu$_\times$, Neu$_\checkmark$)	(Pos$_\checkmark$, Neu$_\checkmark$)
3	My only gripe would be the need to add more **RAM**$_{[exp]}$.	(Neu$_\times$)	(Neg$_\checkmark$)	(Neg$_\checkmark$)
4	This one still has the **CD slot**$_{[imp]}$.	(Pos$_\times$)	(Neu$_\checkmark$)	(Neu$_\checkmark$)
5	**Apple "Help"**$_{[imp]}$ is a mixed bag.	(Neu$_\times$)	(Neu$_\times$)	(Neg$_\checkmark$)
6	The **price**$_{[imp]}$ is 200 dollars down.	(Neg$_\times$)	(Neg$_\times$)	(Pos$_\checkmark$)

Table 6: Several examples from Laptop and Restaurant. The aspects of the sentences are marked in bold. "$[exp]$" and "$[imp]$" indicate the aspect convey explicit and implicit sentiment respectively. We use "Pos", "Neu" and "Neg" to represent *positive*, *neutral* and *negative*. ✓ and × symbolize correct and incorrect prediction.

timent lexicons (Baccianella et al., 2010; Cambria et al., 2020) with external affective knowledge (Liang et al., 2022; Zheng et al., 2023). Despite these successes, satisfactory performance has yet to be achieved considering the ubiquitous implicit sentiment expressions, making implicit sentiment learning a crucial research area. In this area, Liu (2012) highlights key characteristics and early methods for mining this pattern, while Li et al. (2021) introduced sentiment-aware pre-training to boost implicit sentiment learning. Fei et al. (2023) design THOR prompting and a supervised fine-tuning method, extending the CoT idea (Wei et al., 2022) to ABSA. However, these methods struggle with the computational requirements of fine-tuning LLMs such as Flan-T5-11B (Chung et al., 2024) and the reliance on large-scale annotated corpora. Several studies concentrate on enhancing pre-trained language models (PLMs) and constructing specific patterns for implicit sentiment learning (Cai et al., 2021; Wang et al., 2022, 2023a). Unlike these methods, we improve implicit sentiment learning by exploring the connection between explicit and implicit sentiment expressions.

LLMs as Teachers Since LLMs have attracted great attention for their instruction following and text generation abilities (Ouyang et al., 2022), recent works have endeavored to treat LLMs as versatile teachers and transfer their abilities to smaller models (Meng et al., 2022; Ho et al., 2023), which can be regarded as a variant of knowledge distillation (Hinton et al., 2015). A paradigm to treat LLMs as teachers is employing them for data augmentation in various tasks (Yoo et al., 2021; Ye et al., 2024; Zhang et al., 2024), such as text classification and named entity recognition. In ABSA, Wang et al. (2023b) utilize LLMs to generate the explanations and treat them as training samples.

Different from the above studies, we leverage LLMs for paraphrasing sentences to explore the connection between explicit and implicit sentiment expressions. By modeling different sentiment expressions in original and paraphrased sentences and learning from their connection, the knowledge from LLMs can be distilled to the model.

6 Conclusion

Based on the fact that people can express the same sentiment in both direct and indirect ways towards the specific aspect, in this paper, we explore the connection between explicit and implicit sentiment expression and propose a novel sentiment expression conversion- based paraphrase method (SECP) for implicit sentiment learning. We propose an expression paraphrase strategy and a sentiment-consistent contrastive learning mechanism to construct pairs of sentences containing both explicit and implicit sentiment expressions and learn the intrinsic connections between implicit and explicit sentiment expressions and integrate them into the model to improve implicit sentiment learning. We perform extensive experiments on public datasets, and the results show the significant efficacy of our method on implicit sentiment analysis.

7 Limitations

7.1 Limited Experimental Datasets

Although comprehensive experiments including ablation studies, robustness evaluation and case studies have been conducted, they are mostly based on SemEval 2014 Laptop and Restaurant benchmarks, as the annotations of explicit and implicit sentiment expressions are only available on them. Similarly, due to this constraint, most experiments from existing works focusing on implicit sentiment learning in ABSA are also limited to these two datasets, making it challenging to diversify results analysis and model evaluation. Therefore, we call for more accessible annotated datasets to develop implicit sentiment learning in ABSA, including manually annotated datasets and corpus that are automatically annotated using sentiment lexicons or predefined dictionaries of opinion terms.

7.2 Task Transferability

In this paper, we propose to leverage the paraphrase to exploit the connection between explicit and implicit sentiment to improve implicit sentiment learning in ABSA. Although the proposed expression paraphrase strategy focuses on ABSA, we believe that it can also be applied to different tasks under certain conditions. For example, when a predefined dictionary of swearing words is available or the annotations of swearing words are provided in the corpus, this strategy can be leveraged to hate speech detection, achieving the transformation of sentences with and without swearing words to enhance the model's perception of hate speeches. Future work would explore the task transferability of the paraphrase and extend its applicability to a wider range of tasks.

7.3 Reliability of Paraphrased Sentences

Given that the paraphrased sentences are generated by the LLM, their quality remains uncertain. Particularly, when the LLM is required to paraphrase complex sentences such as sentences that convey different sentiment polarities towards the same aspect, the paraphrased sentences may contain errors. Additionally, the LLM may be also susceptible to hallucinations (Zhang et al., 2023b) during the paraphrase. An important method to assess the reliability of paraphrased sentences is to evaluate them. However, manual evaluation of these sentences is time-consuming and suffers from a lack of standardization among evaluators, as individuals may vary in their criteria for assessing explicit and implicit sentiment expressions. Future work would focus on applying available metrics to the automated evaluation of the paraphrased sentences.

8 Ethics Statement

We conduct experiments on two publicly available datasets from SemEval 2014. These datasets do not include personal information and do not contain sensitive content. In the process of the paraphrase, we use the ChatGPT service from OpenAI and follow their policies. Owing to the lack of ethics and bias constraints in the paraphrase, the generated paraphrased sentences may contain sensitive content or biases. It is necessary to manually check the paraphrased sentences in real-world applications.

Acknowledgments

This work is supported by National Natural Science Foundation of China (No. 61976062).

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774v6*.

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta. European Language Resources Association (ELRA).

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609v1*.

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.

Erik Cambria, Yang Li, Frank Z. Xing, Soujanya Poria, and Kenneth Kwok. 2020. Senticnet 6: Ensemble application of symbolic and subsymbolic ai for sentiment analysis. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, CIKM '20, page 105–114, New York, NY, USA. Association for Computing Machinery.

Mayee Chen, Daniel Y Fu, Avanika Narayan, Michael Zhang, Zhao Song, Kayvon Fatahalian, and Christopher Re. 2022. Perfectly balanced: Improving transfer and robustness of supervised contrastive learning. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 3090–3122. PMLR.

Yoonjung Choi and Janyce Wiebe. 2014. +/-EffectWordNet: Sense-level lexicon acquisition for opinion inference. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1181–1191, Doha, Qatar. Association for Computational Linguistics.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.

Lingjia Deng and Janyce Wiebe. 2014. Sentiment propagation via implicature constraints. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 377–385, Gothenburg, Sweden. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Hao Fei, Bobo Li, Qian Liu, Lidong Bing, Fei Li, and Tat-Seng Chua. 2023. Reasoning implicit sentiment with chain-of-thought prompting. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1171–1182, Toronto, Canada. Association for Computational Linguistics.

Florian Graf, Christoph Hofer, Marc Niethammer, and Roland Kwitt. 2021. Dissecting supervised contrastive learning. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3821–3830. PMLR.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531v1*.

Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. Large language models are reasoning teachers. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14852–14882, Toronto, Canada. Association for Computational Linguistics.

Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 18661–18673. Curran Associates, Inc.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942v6*.

Zhengyan Li, Yicheng Zou, Chong Zhang, Qi Zhang, and Zhongyu Wei. 2021. Learning implicit sentiment in aspect-based sentiment analysis with supervised contrastive pre-training. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 246–256, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Bin Liang, Hang Su, Lin Gui, Erik Cambria, and Ruifeng Xu. 2022. Aspect-based sentiment analysis via affective knowledge enhanced graph convolutional networks. *Knowledge-Based Systems*, 235:107643.

Bing Liu. 2012. *Sentiment Analysis and Opinion Mining*, volume 5 of *Synthesis Lectures on Human Language Technologies*. Morgan & Claypool Publishers.

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 arXiv:1907.11692v1*.

Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101v3*.

Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.

Yu Meng, Jiaxin Huang, Yu Zhang, and Jiawei Han. 2022. Generating training data with language models: Towards zero-shot language understanding. In *Advances in Neural Information Processing Systems*, volume 35, pages 462–477. Curran Associates, Inc.

OpenAI. 2023. Chatgpt: Optimizing language models for dialogue. Accessed: 2024-02-10.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward

Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Ethan Perez, Douwe Kiela, and Kyunghyun Cho. 2021. True few-shot learning with language models. In *Advances in Neural Information Processing Systems*, volume 34, pages 11054–11070. Curran Associates, Inc.

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.

Alexander Rietzler, Sebastian Stabinger, Paul Opitz, and Stefan Engl. 2020. Adapt or get left behind: Domain adaptation through BERT language model finetuning for aspect-target sentiment classification. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4933–4941, Marseille, France. European Language Resources Association.

Irene Russo, Tommaso Caselli, and Carlo Strapparava. 2015. SemEval-2015 task 9: CLIPEval implicit polarity of events. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 443–450, Denver, Colorado. Association for Computational Linguistics.

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 *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6311–6322, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Junlang Wang, Xia Li, Junyi He, Yongqiang Zheng, and Junteng Ma. 2023a. Enhancing implicit sentiment learning via the incorporation of part-of-speech for aspect-based sentiment analysis. In *Chinese Computational Linguistics: 22nd China National Conference, CCL 2023, Harbin, China, August 3–5, 2023, Proceedings*, page 382–399, Berlin, Heidelberg. Springer-Verlag.

Kai Wang, Weizhou Shen, Yunyi Yang, Xiaojun Quan, and Rui Wang. 2020. Relational graph attention network for aspect-based sentiment analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3229–3238, Online. Association for Computational Linguistics.

Qianlong Wang, Keyang Ding, Bin Liang, Min Yang, and Ruifeng Xu. 2023b. Reducing spurious correlations in aspect-based sentiment analysis with explanation from large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2930–2941, Singapore. Association for Computational Linguistics.

Siyin Wang, Jie Zhou, Changzhi Sun, Junjie Ye, Tao Gui, Qi Zhang, and Xuanjing Huang. 2022. Causal intervention improves implicit sentiment analysis. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6966–6977, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Zengzhi Wang, Qiming Xie, Zixiang Ding, Yi Feng, and Rui Xia. 2023c. Is chatgpt a good sentiment analyzer? a preliminary study. *arXiv preprint arXiv:2304.04339v2*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Xiaoyu Xing, Zhijing Jin, Di Jin, Bingning Wang, Qi Zhang, and Xuanjing Huang. 2020. Tasty burgers, soggy fries: Probing aspect robustness in aspect-based sentiment analysis. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3594–3605, Online. Association for Computational Linguistics.

Hu Xu, Bing Liu, Lei Shu, and Philip Yu. 2019. BERT post-training for review reading comprehension and aspect-based sentiment analysis. 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 2324–2335, Minneapolis, Minnesota. Association for Computational Linguistics.

Mingfeng Xue, Dayiheng Liu, Wenqiang Lei, Jie Fu, Jian Lan, Mei Li, Baosong Yang, Jun Xie, Yidan Zhang, Dezhong Peng, and Jiancheng Lv. 2023. Unifying discrete and continuous representations for unsupervised paraphrase generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13805–13822, Singapore. Association for Computational Linguistics.

Heng Yang and Ke Li. 2024. Modeling aspect sentiment coherency via local sentiment aggregation. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 182–195, St. Julian's, Malta. Association for Computational Linguistics.

Junjie Ye, Nuo Xu, Yikun Wang, Jie Zhou, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. Llm-da: Data augmentation via large language models for few-shot named entity recognition. *arXiv preprint arXiv:2402.14568v1*.

Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. 2021. GPT3Mix: Leveraging large-scale language models for text augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2225–2239, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander J Ratner, Ranjay Krishna, Jiaming Shen, and Chao Zhang. 2023. Large language model as attributed training data generator: A tale of diversity and bias. In *Advances in Neural Information Processing Systems*, volume 36, pages 55734–55784. Curran Associates, Inc.

Chen Zhang, Qiuchi Li, and Dawei Song. 2019. Aspect-based sentiment classification with aspect-specific graph convolutional networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4568–4578, Hong Kong, China. Association for Computational Linguistics.

Wenxuan Zhang, Yue Deng, Bing-Quan Liu, Sinno Jialin Pan, and Lidong Bing. 2023a. Sentiment analysis in the era of large language models: A reality check. *arXiv preprint arXiv:2305.15005v1*.

Yanyue Zhang, Pengfei Li, Yilong Lai, and Deyu Zhou. 2024. Large, small or both: A novel data augmentation framework based on language models for debiasing opinion summarization. *arXiv preprint arXiv:2403.07693v2*.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023b. Siren's song in the ai ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219v2*.

Yongqiang Zheng, Xia Li, and Jian-Yun Nie. 2023. Store, share and transfer: Learning and updating sentiment knowledge for aspect-based sentiment analysis. *Information Sciences*, 635:151–168.

A Analysis of prompts on the paraphrase

Recent efforts show that different responses would be generated when LLMs receive different prompts, even they are similar (Perez et al., 2021; Lu et al., 2022). In other words, different paraphrased sentences can be derived and affect our models. Thus, we use different prompts for the paraphrase considering the ablation of their components and obtain different paraphrased sentence sets for fine-tuning the proposed SECP-BERT. As shown in Figure 4, both components are crucial to the performance of our model implied by the accuracy and ISE declines. Moreover, ISE is more sensitive to the changes in prompts, showing that appropriate prompts are essential to elicit LLMs' ability to understand implicit sentiment and paraphrase.

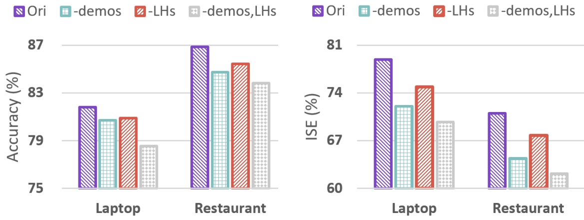

Figure 4: Effects of different prompt components on the paraphrase, which are evaluated by the performance of our SECP-BERT. "Ori" represents the original performance. "-demons", "-LHs" and "-demos, LHs" are removing demonstrations, linguistic hints, both demonstrations and linguistic hints from the prompts.

Model	Lap→Rest		Rest→Lap	
	ACC	MF1	ACC	MF1
BERT-ED	79.11	64.93	75.55	70.99
BERT+PRs	80.54	70.01	72.73	68.88
SECP-BERT	78.93	71.84	77.74	74.42
SECP-RoBERTa	**86.17**	**77.91**	**80.56**	**77.10**

Table 7: Experimental results in the cross-domain setting. "Lap→Rest" ("Rest→Lap") is the settings that specify Laptop (Restaurant) as the training set and Restaurant (Laptop) as the test set. The highest scores are marked in bold.

B Cross-domain Evaluation

Considering that ABSA models would handle sentences from various domains in realistic scenarios, we explore our models in cross-domain evaluation, which requires the models to be trained and evaluated on data from different domains, showing the robustness of the models in domain adaptation. Experimental results are shown in Table 7. It can be observed that our models (SECP-BERT and SECP-RoBERTa) outperform the compared models in most cases. Their performance shows the feasibility of exploiting the connection between explicit

No.	sentence	Ground Truth
1	**Ori**: It is in the best condition and has a really high **quality** (*EXP*). **Par**: Can you find any other laptop that matches the **quality** of this one?	**Aspect**: quality **Polarity**: positive
2	**Ori**: The **case** is carved out of a single block of aluminum. (*IMP*) **Par**: The **case** is decently made and has a simplistic design.	**Aspect**: case **Polarity**: neutral
3	**Ori**: The machine is slow to **boot up** and occasionally crashes completely. (*EXP*) **Par**: The machine takes forever to start up and occasionally crashes completely.	**Aspect**: boot up **Polarity**: negative
4	**Ori**: I will never go back to **Windows**! (*IMP*) **Par**: I have tried **Windows** in the past and I absolutely despise it!	**Aspect**: Windows **Polarity**: negative

Table 8: Examples of the proposed paraphrase from Laptop. "**Ori**" and "**Par**" represent original and paraphrased sentences respectively. We list the sentiment polarities towards the target aspects in the "**Ground Truth**" column. "*EXP*" and "*IMP*" in the brackets denote the given sentiment expression categories of the original sentences.

Category	Prompt Format
ABSA inference	Given the sentence, please infer the sentiment towards the aspect "{aspect}". Please select a sentiment label from ['negative', 'neutral', 'positive']. Return the predicted label only without any other text. *Sentence: {neg demo} (sentiment towards "{neg demo asp}")* *Label: negative* *Sentence: {neu demo} (sentiment towards "{neu demo asp}")* *Label: neutral* *Sentence: {pos demo} (sentiment towards "{pos demo asp}")* *Label: positive* Sentence: {original sentence} Label:

Table 9: The format of prompts to LLMs for ABSA inference, which are designed by Zhang et al. (2023a). Text in italics indicates the given demonstrations in few-shot inference and would be removed in zero-shot inference. Those placeholders would be replaced with the given demonstrations, the aspect and the review of the input example.

and implicit sentiment expressions to improve implicit sentiment learning. We believe that some common explicit and implicit sentiment expressions are shared across domains, which accounts for the benefit of our method in cross-domain setting. We leave the further analysis to future work.

C Analysis of the Paraphrase

We analyze the proposed paraphrase and show several examples as shown in Table 8. They are training examples from the Laptop benchmark and expressed by both explicit and implicit sentiment expressions. By comparing the sentences with their paraphrased variants, ChatGPT shows its abilities in paraphrasing and sentiment expression conversion. For example, though the opinion term towards the aspect "*quality*" is absent, the paraphrased sentence can also convey positive sentiment towards "*quality*" in the first example. In example 4, ChatGPT appends appropriate opinion "*absolutely despise*" to the paraphrased sentences, which explicitly expresses negative sentiment towards the aspect "*Windows*". However, it is worth noting that ChatGPT would substitute the aspects for their synonyms while converting sentiment expressions, such as substituting "*boot up*" for "*start up*" in the paraphrased sentence of example 3. We think that it might be due to the language modeling mechanism. To alleviate the impact of this problem, we use the prompt-based input sequences in Section 2.4.

D The Prompt Format

Prompts to LLMs for ABSA Inference Considering that different prompts will induce different responses from LLMs, we employ the prompts formulated by Zhang et al. (2023a) to facilitate zero-shot and few-shot inference for ABSA using LLMs in Section 3, which are designed to be relatively consistent across different datasets and LLMs. Their format is detailed in Table 9.

Prompts for Generating Linguistic Hints Table 10 presents the complete prompts for ChatGPT to generate paraphrase-relevant linguistic hints, mentioned in "**Paraphrase-relevant Linguistic Hints**" of Section 2.3.

Prompts for the Paraphrase Besides the prompt template and the proposed components in Section 2.3, our designed prompts for the expression paraphrase strategy are also composed of the domain, original sentence, target aspect, sentiment polarity towards the aspect. We illustrate the format of these prompts in Table 11.

Category	Prompt
Generate linguistic hints: explicit→implicit	Given following three pairs of reviews, which attribute dimensions do you consider vital in distinguishing implicit sentiment towards a specific aspect from explicit sentiment towards the same aspect? Explicit positive sentiment towards the aspect "4 GB stick of RAM": I highly recommend purchasing this model with a cost-effective 4 GB stick of RAM, which is a good deal to save 10. Implicit positive sentiment towards the aspect "4 GB stick of RAM": One more tip, please purchase this model and get a 4 GB stick of RAM to save you 10. Explicit neutral towards the aspect "battery life": Like most other laptops, it lasts for 5-6 hours with an ordinary battery life. Implicit neutral towards the aspect "battery life": Has a 5-6 hour battery life. Explicit negative towards the aspect "tech support": Waste a long time calling the unresponsive tech support and get nothing. Implicit negative towards the aspect "tech support": Just forget the tech support! You will feel better and fix the problem by yourself.
Generate linguistic hints: implicit→explicit	Given following three pairs of reviews, which attribute dimensions do you consider vital in distinguishing explicit sentiment towards a specific aspect from implicit sentiment towards the same aspect? Explicit positive sentiment towards the aspect "4 GB stick of RAM": I highly recommend purchasing this model with a cost-effective 4 GB stick of RAM, which is a good deal to save 10. Implicit positive sentiment towards the aspect "4 GB stick of RAM": One more tip, please purchase this model and get a 4 GB stick of RAM to save you 10. Explicit neutral towards the aspect "battery life": Like most other laptops, it lasts for 5-6 hours with an ordinary battery life. Implicit neutral towards the aspect "battery life": Has a 5-6 hour battery life. Explicit negative towards the aspect "tech support": Waste a long time calling the unresponsive tech support and get nothing. Implicit negative towards the aspect "tech support": Just forget the tech support! You will feel better and fix the problem by yourself.

Table 10: Prompts for generating paraphrase-relevant linguistic hints in Section 2.3. The "**Category**" column represents two different sentiment conversions in different directions.

Category	Prompt Format
Paraphrase: implicit$\rightarrow$explicit	Suppose you are a {domain} consumer. You should paraphrase the review you had written that expresses explicit sentiment towards a particular aspect of the {domain} and express implicit sentiment towards this aspect in your new review. The sentiment polarity will be only one of these 3 classes: ["negative", "neutral", "positive"]. **Here are some examples**: Explicit positive sentiment towards the aspect "{pos demo asp}": "{pos explicit demo}" => Implicit positive sentiment towards the aspect "{pos demo asp}": "{pos implicit demo}" Explicit neutral sentiment towards the aspect "{neu demo asp}": "{neu explicit demo}" => Implicit neutral sentiment towards the aspect "{neu demo asp}": "neu implicit demo" Explicit negative sentiment towards the aspect "{neg demo asp}": "{neg explicit demo}" => Implicit negative sentiment towards the aspect "{neg demo asp}": "{neg implicit demo}" **And here are some linguistic hints for paraphrasing**: 1. Directness of expression: For implicit sentiment, the mention of the aspect might be more subtle or incidental compared to explicit sentiment, which directly describes the aspect. Implicit sentiment often uses indirect or suggestive language. 2. Use of emotional cues: Implicit sentiment might contains fewer emotional markers and linguistic cues such as adverbs, intensifiers and emotional words, making it challenging to detect the sentiment. 3. Level of recommendation or rejection: For implicit sentiment, suggestions and recommendations without clearly stating the sentiment are frequently involved rather than direct recommendations or rejection towards a particular aspect. Now given a {gold label} review and a particular aspect "{aspect}", please paraphrase it and let it express the same sentiment polarity implicitly. Explicit {gold label} sentiment towards the aspect "{aspect}": {original review} => Implicit {gold label} sentiment towards the aspect "{aspect}":
Paraphrase: implicit$\rightarrow$explicit	Suppose you are a {domain} consumer. You should paraphrase the review you had written that expresses explicit sentiment towards a particular aspect of the {domain} and express implicit sentiment towards this aspect in your new review. The sentiment polarity will be only one of these 3 classes: ["negative", "neutral", "positive"]. **Here are some examples**: Implicit positive sentiment towards the aspect "{pos demo asp}": "{pos implicit demo}" => Explicit positive sentiment towards the aspect "{pos demo asp}": "{pos explicit demo}" Implicit neutral sentiment towards the aspect "{neu demo asp}": "{neu implicit demo}" => Explicit neutral sentiment towards the aspect "{neu demo asp}": "{neu explicit demo}" Implicit negative sentiment towards the aspect "{neg demo asp}": "{neg implicit demo}" => Explicit negative sentiment towards the aspect "{neg demo asp}": "{neg explicit demo}" **And here are some linguistic hints for paraphrasing**: 1. Directness of expression: Explicit sentiment directly expresses opinions and emotions, while implicit sentiments may use indirect or suggestive language. 2. Use of emotional cues: Explicit sentiment often includes explicit emotional expressions like adverbs, superlatives, and emotional words. According to the opinion words and emotional expressions, explicit sentiment can be easily and obviously detected and classified. 3. Level of recommendation or rejection: Explicit sentiment tends to directly recommend or reject a particular aspect and usually present a clear opinion about the aspect. Now given a {gold label} review and a particular aspect "{aspect}", please paraphrase it and let it express the same sentiment polarity explicitly. Implicit {gold label} sentiment towards the aspect "{aspect}": {original review} => Implicit {gold label} sentiment towards the aspect "{aspect}":

Table 11: The format of prompts for prompting LLMs to paraphrase. The bold parts list the proposed paraphrase-based demonstrations and paraphrase-relevant linguistic hints in Section 2.3. Those placeholders would be replaced with the domain, aspects and reviews of demonstrations, gold label (the sentiment polarity towards target aspect), target aspect and original review.