

FCKT: Fine-Grained Cross-Task Knowledge Transfer with Semantic Contrastive Learning for Targeted Sentiment Analysis

Wei Chen¹, Zhao Zhang², Meng Yuan¹, Kepeng Xu³ and Fuzhen Zhuang^{1,4†}

¹School of Artificial Intelligence, Beihang University, China

²School of Computer Science and Engineering, Beihang University, China

³Xidian University, China

⁴Zhongguancun Laboratory, China

{chenwei23, zhuangfuzhen}@buaa.edu.cn

Abstract

In this paper, we address the task of targeted sentiment analysis (TSA), which involves two sub-tasks, i.e., identifying specific aspects from reviews and determining their corresponding sentiments. Aspect extraction forms the foundation for sentiment prediction, highlighting the critical dependency between these two tasks for effective cross-task knowledge transfer. While most existing studies adopt a multi-task learning paradigm to align task-specific features in the latent space, they predominantly rely on coarse-grained knowledge transfer. Such approaches lack fine-grained control over aspect-sentiment relationships, often assuming uniform sentiment polarity within related aspects. This oversimplification neglects contextual cues that differentiate sentiments, leading to negative transfer. To overcome these limitations, we propose FCKT, a fine-grained cross-task knowledge transfer framework tailored for TSA. By explicitly incorporating aspect-level information into sentiment prediction, FCKT achieves fine-grained knowledge transfer, effectively mitigating negative transfer and enhancing task performance. Experiments on three datasets, including comparisons with various baselines and large language models (LLMs), demonstrate the effectiveness of FCKT. The source code is available on <https://github.com/cwei01/FCKT>.

1 Introduction

Unlike traditional sentiment analysis tasks [Liu, 2012], aspect-based sentiment analysis (ABSA) [Wang *et al.*, 2016; Tang *et al.*, 2020] requires a deeper contextual understanding and the extraction of more fine-grained information. ABSA has been widely applied in tasks such as customer feedback analysis and academic paper review evaluation [Nath and Dwivedi, 2024]. ABSA typically involves identifying (1) aspects (entities or attributes discussed), (2) opinions (sentiment expressions), and (3) sentiment polarity. However, research indicates that over 30% of sentiment expressions are

The lens and autofocus are impressive, except low-light capability.

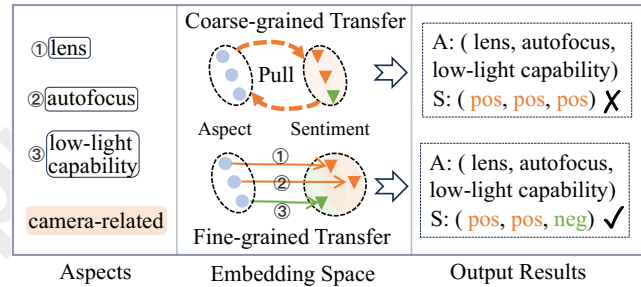


Figure 1: Illustration comparing coarse-grained and fine-grained transfer methods. Coarse-grained transfer assumes uniform sentiment polarity for related aspects, while fine-grained transfer captures nuanced aspect-sentiment relationships for precise alignment.

implicit [Cai *et al.*, 2021; Chen *et al.*, 2024a], meaning the ABSA must infer aspects and sentiments from the context without predefined opinions. In this paper, we address these challenges by developing targeted sentiment analysis (TSA). It involves two sub-tasks, i.e., identifying relevant aspects from reviews and explicitly associating them to their corresponding sentiments [Chen *et al.*, 2022b; Zhou *et al.*, 2024].

When TSA was first proposed, a two-stage pipeline method was developed [Hu *et al.*, 2019; Kalbhor and Goyal, 2023]. However, this approach often lead to error propagation between sub-tasks. Moreover, aspects are intrinsically linked with sentiment, and this staged pipeline approach disrupts this crucial interaction [Xu *et al.*, 2021]. To address these issues, many end-to-end approaches have been proposed [Chen *et al.*, 2022b; 2022a; Li *et al.*, 2023; Zhu *et al.*, 2024], aiming to jointly extract aspect terms and classify their sentiments by modeling the interactions between aspects and sentiments. A widely studied technique in this field is the task-specific feature alignment approach [Chen *et al.*, 2024a]. This method involves two key steps: first, encoding task-specific features for both aspects and sentiments; second, aligning these features in the latent space, even when some aspects differ in sentiments, to enhance cross-task knowledge interaction.

Despite recent advancements, these methods still rely on coarse-grained knowledge transfer, which fails to provide explicit and fine-grained control over aspect-sentiment relation-

[†] Corresponding author.

ships. This limitation often leads to negative transfer, as coarse-grained approaches assume uniform sentiment polarity across related aspects, ignoring key contextual differences. For instance, as illustrated in Figure 1, coarse-grained methods treat all camera-related aspects (“lens”, “autofocus”, and “low-light capability”) share the same sentiment representation. Consequently, these methods incorrectly classify “low-light capability” as positive, overlooking the negative sentiment implied by the contrastive cue “except”. This misalignment, stemming from the lack of explicit control over knowledge transfer, is particularly problematic when aspect-sentiment relationships are subtle or context-dependent.

Along this research line, explicit fine-grained knowledge transfer has emerged as a promising solution to enhance the interaction between aspects and sentiments. It disentangles aspect-sentiment relationships and enables dynamic alignment tailored to each aspect. By isolating sentiment representations for each aspect, it accurately associates “lens” and “autofocus” with positive sentiment and “low-light capability” with negative sentiment, leveraging contextual cues like the contrast implied by “except”. This ensures more precise and context-aware sentiment classification. However, despite its potential, explicit fine-grained cross-task knowledge transfer faces the more severe challenge of error accumulation that must be addressed: (1) In fine-grained knowledge transfer, the performance of tasks such as aspect extraction and sentiment classification is highly interdependent, as information flows directly between tasks. While leveraging aspect extraction to assist sentiment classification fosters cross-task collaboration, it also introduces the risk of error propagation. Errors or ambiguities in aspect extraction, such as incorrectly identified or missing aspects, are directly transferred to the sentiment classification phase, compounding inaccuracies in the final predictions. To address this, a more effective approach is required to constrain aspect extraction, ensuring more accurate identification and minimizing the influence of ambiguity. (2) Moreover, even with perfect aspect extraction, fine-grained knowledge transfer may still struggle due to insufficient supervisory signals for sentiment classification. As the sentiment classifier relies heavily on aspect extraction, its training can become overly dependent on these features, potentially overlooking other crucial contextual cues that are not explicitly tied to aspect information. This lack of comprehensive supervision can hinder the classifier’s ability to generalize effectively. This underscores the need for balanced, holistic supervision to optimize fine-grained knowledge transfer.

To address these challenges, we propose FCKT, a fine-grained cross-task knowledge transfer framework tailored for the TSA task. Specifically, we design a token-level semantic contrastive learning mechanism. In this approach, the start and end tokens of the same aspect are treated as positive pairs, while tokens from unrelated aspects serve as negative pairs. This refines the model’s understanding of aspects and enhances its ability to capture subtle contextual dependencies. To mitigate the lack of supervisory signals in the sentiment classifier, we introduce an alternating learning strategy. This approach trains a fixed proportion of samples using real labels, while the remaining samples are updated based on predictions from the previous model. This enables FCKT

to leverage both real labeled information and predicted transfer knowledge, synergistically improving the effectiveness of TSA. Furthermore, while Large Language Models (LLMs) like GPT-3.5 and GPT-4 excel in many general-purpose scenarios, they often struggle to adapt to task-specific challenges in fine-grained TSA tasks. Our proposed framework surpasses these models in both few-shot learning and chain-of-thought reasoning scenarios by leveraging task-specific knowledge transfer mechanisms and tailored training strategies. These results highlight the capability of FCKT to address the unique complexities of TSA tasks, providing valuable insights for advancing task-specific methods in the LLM era. Our key contributions are summarized as follows:

- We propose a token-level semantic contrastive learning mechanism that enhances the model’s understanding of aspects by treating start and end tokens as positive pairs and unrelated tokens as negative pairs, enabling it to effectively capture subtle contextual dependencies.
- To address insufficient supervisory signals in sentiment classification, we introduce an alternating learning strategy integrates real labeled data with predicted transfer knowledge, significantly improving task-specific performance.
- We conduct extensive experiments on three real-world datasets, demonstrating the superiority of FCKT over existing approaches. We also demonstrate the effectiveness of our approach across various LLM models, showcasing its superior performance in the era of LLMs.

2 Related Work

Targeted Sentiment Analysis. Targeted sentiment analysis (TSA) involves extracting aspect terms and their associated sentiment polarities, a key focus in recent research. Early work approached aspect extraction and sentiment analysis as separate tasks, employing methods such as Long Short-Term Memory (LSTM) networks [Luo *et al.*, 2019] and Bidirectional Encoder Representations from Transformers (BERT) [Yang *et al.*, 2020]. Recent studies have shifted toward end-to-end solutions using multi-task learning frameworks [Luo *et al.*, 2019; Lin and Yang, 2020; Chen *et al.*, 2022b], highlighting the interdependence between aspects and sentiments. Techniques such as shared-private feature interaction [Lin and Yang, 2020] and task-regularization interaction [Chen *et al.*, 2022b] have been introduced to improve performance by leveraging cross-task knowledge transfer. However, current approaches mainly rely on coarse-grained alignment approaches, which lacks fine-grained control over aspect-sentiment relationships. This often leads to negative transfer, as coarse-grained approaches assume uniform sentiment polarity within related aspects, overlooking contextual cues that differentiate sentiments. More recently, large language models (LLMs) like GPT have demonstrated strong generalization capabilities, outperforming fine-tuned BERT in some low-resource ABSA scenarios [Wang *et al.*, 2023; Zhou *et al.*, 2024]. Nevertheless, they continue to underperform compared to state-of-the-art task-specific models, underscoring the need for further adaptation to address the unique challenges of resource-constrained TSA tasks.

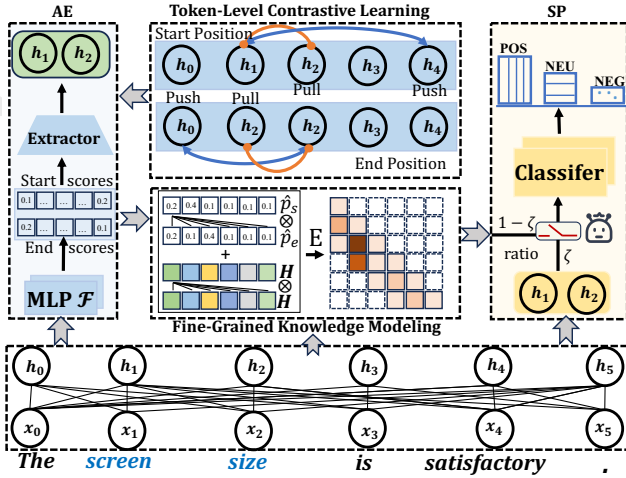


Figure 2: A depiction of the peopose FCKT framework.

Contrastive Learning. Contrastive learning has recently achieved significant success in various domains [Jaiswal *et al.*, 2020; Zhang *et al.*, 2022; Luo *et al.*, 2024; Zhong *et al.*, 2024b; 2024a; 2025; Yuan *et al.*, 2025a; 2025b; 2023]. In ABSA, several studies have explored its integration into model training to improve performance. For example, Liang [Liang *et al.*, 2021] leveraged contrastive learning to distinguish between aspect-invariant and aspect-dependent features, leading to better sentiment classification. Similarly, Xiong [Xiong *et al.*, 2022] proposed a triplet contrastive learning network that combines syntactic and semantic information via an aspect-oriented sub-tree and sentence-level contrastive learning. Chen [Chen *et al.*, 2024a] further utilized contrastive learning to model adaptive task-relatedness between aspect and sentiment co-extraction. Different from these works, we leverage contrastive learning specifically for aspect extraction to enhance knowledge transfer in TSA.

3 Methodology

In this section, we will introduce proposed FCKT framework in detail, which is depicted in Figure 2. The comprehensive task consists of two base components: first, extracting opinion aspects (AE), and second, predicting the sentiment polarity (i.e., positive, neutral, or negative) of each aspect (SP).

3.1 Problem Description

Formally, given a sentence $\mathbf{x} = \{x_0, x_1, \dots, x_{n-1}\}$, where x_i represents the i^{th} token in \mathbf{x} and n denotes the sentence length, the main objective of FCKT is to extract all possible aspects and predict their corresponding sentiments in \mathbf{x} .

To enhance the representation of contextual semantics, we leverage the widely recognized BERT model [Devlin *et al.*, 2019], an effective multi-layer bidirectional Transformer designed to construct rich contextual embeddings by simultaneously capturing both left and right word dependencies. BERT comprises a series of T Transformer layers. The working principle of BERT is summarized as follows:

$$\mathbf{H}^0 = \mathbf{W}_i \mathbf{E} + \mathbf{P}, \quad \mathbf{H}^t = \mathcal{T}(\mathbf{H}^{t-1}), t \in [1, T], \quad (1)$$

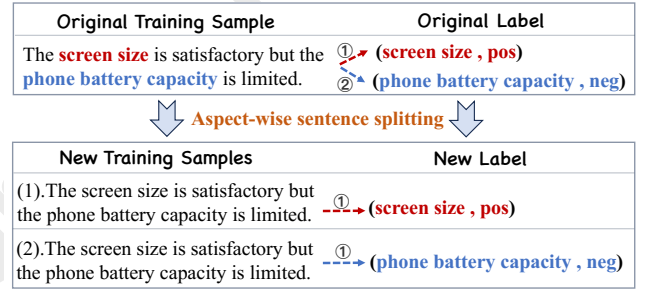


Figure 3: An example of the sentence splitting process.

where \mathbf{W}_i is the one-hot representation of the sub-word in the input sentence, \mathbf{E} denotes the sub-word embedding matrix, \mathbf{P} represents the positional embedding, \mathcal{T} is the Transformer block, and \mathbf{H}^t is the output of the t^{th} Transformer layer. For the given sentence $\mathbf{x} = \{x_0, x_1, \dots, x_{n-1}\}$, we assume the output of the final Transformer layer is $\mathbf{H} = \mathbf{H}^T = \{\mathbf{h}_0, \mathbf{h}_1, \dots, \mathbf{h}_{n-1}\} \in \mathbb{R}^{n \times d}$, where d is the embedding size.

Remark. In our modeling process, sentences containing multiple aspect terms are split into separate sentences during training, with each focusing on a single aspect, as illustrated in Figure 3. This preprocessing step allows the model to effectively learn aspect-specific representations by isolating each aspect, avoiding interference from others in the same sentence. It is important to note that this splitting strategy is only applied during training to facilitate the learning process. During testing, the original sentence structure is preserved, ensuring consistency in evaluation and practical application. Our approach, FCKT, leverages this strategy without altering the original optimization objectives or knowledge transfer mechanisms. **Further details are provided in the Appendix A**, where we demonstrate that this strategy maintains optimization consistency while enabling fine-grained modeling of aspect-specific interactions.

3.2 Semantic Contrastive Learning for AE

In this article, aspects are detected by predicting their start and end boundaries through a linear transformation operation [Lin and Yang, 2020]. Specifically, the start and end boundary distributions are predicted as follows:

$$\hat{\mathbf{p}}_s = \mathcal{F}_{\phi_1}(\mathbf{H}), \hat{\mathbf{p}}_e = \mathcal{F}_{\phi_2}(\mathbf{H}), \quad (2)$$

where \mathcal{F}_{ϕ_1} and \mathcal{F}_{ϕ_2} are MLP layers used for extracting aspects. The length of $\hat{\mathbf{p}}_s$ matches that of the sentence, and each element in $\hat{\mathbf{p}}_s$ represents the probability of a word being the start of an aspect. It is important to note that this step only predicts the boundary distributions of aspects. To determine the final aspects, a heuristic extraction algorithm [Hu *et al.*, 2019] is subsequently employed. During optimization, the learning objective for boundary distributions is formulated as:

$$\mathcal{L}_{ae} = - \sum_{i=1}^n \{ \mathbf{p}_{i,s}^T \log(\hat{\mathbf{p}}_{i,s}) + \mathbf{p}_{i,e}^T \log(\hat{\mathbf{p}}_{i,e}) \}, \quad (3)$$

where $\mathbf{p}_{i,s}^T \in \mathbb{R}^n$ and $\mathbf{p}_{i,e}^T \in \mathbb{R}^n$ are the boundary ground truths (i.e., 0-1 vectors), $\hat{\mathbf{p}}_{i,s}$ and $\hat{\mathbf{p}}_{i,e}$ are the predicted boundary distributions, n is the length of sentence.

Compared to conventional coarse-grained transfer methods [Chen *et al.*, 2024a; Sun *et al.*, 2024], fine-grained knowledge transfer imposes stricter requirements on the precision of aspect extraction. Any incorrectly identified or missing aspects are directly transferred to the sentiment classification phase, accumulating prediction inaccuracies. To mitigate this issue, we introduce a token-level semantic contrastive learning framework to improve aspect representation quality.

Intuitively, tokens within an aspect inherently exhibit strong semantic compatibility, as they frequently co-occur in similar contexts. For instance, as illustrated in Figure 2, the end boundary distribution suggests that both the 2nd and 3rd words have comparable probabilities of being the end of an aspect. However, from a semantic perspective, the word “is” is unrelated to the preceding words and is therefore less likely to be part of the aspect boundary. Semantic information can serve as a complementary signal to boundary distributions by verifying whether different tokens form a “semantically reasonable” aspect. To leverage this property, we treat the start and end tokens of the same aspect as positive pairs, encouraging their representations to be closely aligned. For negative pairs, we construct two types: (1) the start token of the aspect paired with end tokens from other, unrelated aspects, and (2) the end token of the aspect paired with start tokens from unrelated aspects. The InfoNCE loss [Chen *et al.*, 2024b; 2024c] is then employed to optimize these token-level embeddings. Formally, the contrastive loss is defined as:

$$\mathcal{L}_{cl} = - \sum_{(s,e) \in \mathcal{E}} \log \frac{\exp(\mathbf{s}(\mathbf{h}_s, \mathbf{h}_e)/\tau)}{\sum \exp(\mathbf{s}(\mathbf{h}_s, \mathbf{h}_i)/\tau) + \exp(\mathbf{s}(\mathbf{h}_e, \mathbf{h}_j)/\tau)}, \quad (4)$$

where $(\mathbf{h}_s, \mathbf{h}_e)$ are positive pair embeddings, while $(\mathbf{h}_s, \mathbf{h}_i)$ and $(\mathbf{h}_e, \mathbf{h}_j)$ are negative pairs. τ is the temperature parameter, $\mathbf{s}(\cdot, \cdot)$ denotes the cosine similarity function, and \mathcal{E} represents the set of all positive pairs in each sentence.

3.3 Fine-Grained Knowledge Transfer for SP

In sentiment prediction task, the sentiment polarity for a given aspect is determined based on the words within its boundaries. Let the start and end boundaries be denoted as s and e , respectively. The sentiment is then predicted as:

$$\hat{\mathbf{y}}(\psi) = \mathcal{C}_\theta(\mathbf{H}_{s \rightarrow e}), \quad (5)$$

where $\hat{\mathbf{y}}(\psi)$ represents the predicted class probability distribution, \mathcal{C}_θ is a nonlinear classifier that maps input features to the range $[0, 1]$, effectively modeling the likelihood of each sentiment class. $\mathbf{H}_{s \rightarrow e}$ denotes the aspect representation, aggregated from the start boundary s to the end boundary e .

As previously mentioned in Section 1, these two sub-tasks exhibit a logical sequential dependence, implying that aspects can provide valuable signals for sentiment prediction. However, sentiment prediction in previous coarse-grained transfer methods relies on ground truth boundaries, which essentially disconnects the sequential connections between these two sub-tasks. Therefore, we employ aspect features to supervise the sentiment prediction process. In this way, aspect extraction can provide sequentially useful supervised signals for sentiment prediction, while sentiment information can influence aspect detection through backpropagation.

To achieve this, we refrain from using the ground truth boundaries as input. Instead, we derive an expectation of word embedding based on the boundary distributions. Consequently, Eq. (5) can be redefined as:

$$\hat{\mathbf{y}}(\ell) = \mathcal{C}_\theta(\mathbb{E}_{(i \sim \hat{\mathbf{p}}_s, j \sim \hat{\mathbf{p}}_e)} \mathbf{H}_{i \rightarrow j}), \quad (6)$$

where \mathbb{E} denotes the expectation operator over the indices i and j , which are sampled from the predicted start boundary distribution $\hat{\mathbf{p}}_s$ and the end boundary distribution $\hat{\mathbf{p}}_e$, respectively. Furthermore, since $\hat{\mathbf{p}}_s$ and $\hat{\mathbf{p}}_e$ are both discrete distributions, we have the following derivation:

$$\begin{aligned} \hat{\mathbf{y}}(\ell) &= \mathcal{C}_\theta \left(\mathbb{E}_{i \sim \hat{\mathbf{p}}_s} \left[\mathbb{E}_{j \sim \hat{\mathbf{p}}_e} \left[\sum_{k=i}^j \mathbf{H}_k \right] \right] \right) \\ &= \mathcal{C}_\theta \left(\mathbb{E}_{i \sim \hat{\mathbf{p}}_s} \left[\sum_{j=1}^n \hat{p}_{e,j} \sum_{k=i}^j \mathbf{H}_k \right] \right) \\ &= \mathcal{C}_\theta \left(\sum_{i=1}^n \sum_{j=1}^n \hat{p}_{s,i} \cdot \hat{p}_{e,j} \sum_{k=i}^j \mathbf{H}_k \right), \end{aligned} \quad (7)$$

where we assume n is the sentence’s length. In this way, these two sub-tasks enable genuine end-to-end learning.

In practice, to compute Eq. (7), one needs to traverse the whole sentence, and the computational complexity is $\mathcal{O}(n^2)$, which is very difficult to calculate. Fortunately, this complexity can be reduced based on two facts: (i) The end point is not smaller than the start point. (ii) The length of an aspect term is usually not large. We then revise $\hat{\mathbf{y}}(\psi)$ to $\hat{\mathbf{y}}(\ell)$, as follows:

$$\hat{\mathbf{y}}(\ell) = \mathcal{C}_\theta \left(\sum_{i=1}^n \sum_{j=i}^{i+h} \hat{p}_{s,i} \cdot \hat{p}_{e,j} \sum_{k=i}^j \mathbf{H}_k \right), \quad (8)$$

where h is the maximum length of aspect. Thus, the computational complexity is reduced to $\mathcal{O}(nh)$ (i.e., $h \ll n$).

Actually, in the testing phase, the sentiment is inferred based on the predicted aspects, which is more aligned with our fine-grained transfer modeling. By inputting the distributional boundaries, the sentiment information can be back propagated to supervise the aspect detection process, and enhance its accuracy. However, if we only base on the distributional input, we may still face challenges due to insufficient supervisory signals for sentiment classifier. In such a scenario, providing some ground truth data can reduce the error propagation path, and thus alleviate this problem. Specifically, during training, we use an alternate training strategy: a fixed ratio (ξ) of samples are trained based on Eq. (5), while the remaining samples $(1 - \xi)$ are optimized using Eq. (8). The model parameters are optimized through the following cross-entropy loss formulation [Chen *et al.*, 2024a]:

$$\mathcal{L}_{sp} = - \sum_{i=1}^N \sum_{j=1}^K \mathbf{y}_{i,j} \log (\xi \cdot \hat{\mathbf{y}}(\psi)_{i,j} + (1 - \xi) \cdot \hat{\mathbf{y}}(\ell)_{i,j}), \quad (9)$$

where N represents the number of samples, K indicates the number of sentiment types, $\hat{\mathbf{y}}(\psi)$ and $\hat{\mathbf{y}}(\ell)$ are predicted sentiment distributions, and \mathbf{y} is the ground truth. This strategy effectively combines real labels with transferred knowledge, leveraging their synergistic integration to enhance FCKT.

Algorithm 1: Training Framework

Input: Training dataset \mathcal{D} , learning rate α , sampling ratio ξ , model parameters $\Theta = \{\phi_1, \phi_2, \theta\}$

while the convergence criterion is not met **do**

for each batch \mathcal{B} in $\text{Dataloader}(\mathcal{D})$ **do**

 Compute \hat{p}_s and \hat{p}_e using Eq. (2);

 Evaluate aspect extraction loss \mathcal{L}_{ae} via Eq. (3);

 Evaluate contrastive loss \mathcal{L}_{cl} via Eq. (4);

 Sample a random value $p \sim \text{Uniform}(0, 1)$;

if $p > \xi$ **then**

 Construct predicted $\mathbb{E}_{(i \sim \hat{p}_s, j \sim \hat{p}_e)} H_{i \rightarrow j}$;

 Compute the predicted $\hat{y}(\ell)$ via Eq. (8);

else

 Select real aspect embedding $H_{s \rightarrow e}$;

 Compute the real $\hat{y}(\psi)$ via Eq. (5);

 Evaluate sentiment loss \mathcal{L}_{sp} via Eq. (9);

 Aggregate total loss $\mathcal{L} = \mathcal{L}_{ae} + \mathcal{L}_{sp} + \lambda \mathcal{L}_{cl}$;

 Update parameter $\Theta = \Theta - \alpha \nabla_{\Theta} \mathcal{L}$.

3.4 Model Optimization

Overall Training. To enhance the performance of FCKT framework, we jointly optimize the aspect extraction loss \mathcal{L}_{ae} , sentiment prediction loss \mathcal{L}_{sp} , and contrastive loss \mathcal{L}_{cl} . The combined objective function is defined as:

$$\mathcal{L} = \mathcal{L}_{ae} + \mathcal{L}_{sp} + \lambda \mathcal{L}_{cl}, \quad (10)$$

where λ is a trade-off parameter for contrastive loss. The entire optimization process is illustrated in Algorithm 1, with all parameters trained in an end-to-end manner.

Complexity Analysis. We further analyze the time and space complexities of FCKT. The time complexity is composed of three main components: (1) the aspect extraction module, with a complexity of $\mathcal{O}(n + n^2)$ per training epoch; (2) the sentiment prediction module, with a complexity of $\mathcal{O}(h)$; and (3) the expectation input module, with a complexity of $\mathcal{O}(nh)$. Consequently, the overall time complexity of FCKT is $\mathcal{O}(n^2 + nh)$. In terms of space complexity, the model incurs costs primarily from word embeddings and the contrastive learning process, resulting in a total space complexity of $\mathcal{O}(nd + n^2)$, where d represents the embedding size.

4 Experiments

In this section, we conduct experiments to answer the following five research questions: **RQ1:** How does the proposed FCKT perform overall compared to baseline approaches? **RQ2:** How effectively does FCKT perform across different sub-tasks? **RQ3:** How do the various components of FCKT contribute to the final results? **RQ4:** How do different parameter configurations in the proposed method impact its performance? **RQ5:** Does fine-grained knowledge transfer help address the aforementioned challenges effectively?

4.1 Experimental Setup

Datasets. To fairly assess our proposed FCKT, we conducted our experiments using three public datasets, Laptop, Restaurant, and Tweets. The data statistics are presented in Table 1. More dataset details are shown in Appendix B.

Dataset	#Sentences	#Aspects	#+	#-	#0
Laptop	1869	2936	1326	900	620
Restaurant	3900	6603	4134	1538	931
Tweets	2350	3243	703	274	2266

Table 1: Statistics of three datasets. “+/-/0” denote the positive, negative, and neutral sentiment polarities, respectively.

Baselines. We compare three categories of models: **Pipeline-based models**, **End-to-End models**, **LLM-based models**. More baseline details are listed in Appendix C.

Evaluation. We adopt three widely used metrics, precision, recall, and F1 score, to evaluate the FCKT. Please refer to Appendix D for more metric details.

Implementation Details. They are provided in Appendix E, with the full Appendix available in our arXiv version.

4.2 Main Experimental Results (RQ1)

We conduct experiments on three public datasets and the comparison results between FCKT and the baselines are presented in Table 2. This table shows the following:

Among the baselines, we find that our proposed FCKT consistently achieves the best results in all cases on F1 score. More precisely, when compared to the state-of-the-art approach AIFI, FCKT demonstrates notable enhancements in F1 score performance across three distinct datasets, with an average improvement of approximately 1.38%¹. These findings strongly suggest that the meticulously crafted FCKT exhibits a potential for superior performance. This can be attributed to two primary factors: first, the inclusion of a fine-grained knowledge transfer framework that effectively captures mutual information between sub-tasks, and second, the introduction of a token-level contrastive learning mechanism that facilitates the extraction of more fitting aspects. However, the recall of FCKT is lower than several baseline models. One possible reason is that our heuristic extraction method filters out semantically irrational aspects based on the cumulative start/end scores, which may inadvertently exclude some false negatives, leading to a lower recall score.

FCKT demonstrates notable advantages in both performance and computational efficiency compared to advanced LLMs. As illustrated in Tables 2, even the state-of-the-art LLM, GPT-4, with its vast number of parameters, fails to deliver satisfactory results for TSA, despite leveraging Few-Shot learning and Chain-of-Thought (CoT) [Wei *et al.*, 2022] enhancement. Moreover, the use of LLMs incurs substantial computational overhead, further underscoring the efficiency of our approach. This suggests that improving performance in these low-resource domains with LLMs remains a challenging task. Despite advancements in LLMs, some traditional methods for TSA are still valuable, offering insights that complement the limitations of current techniques.

4.3 Analysis on Both Sub-Tasks (RQ2)

To verify model’s performance in individual tasks, we conduct a comparative analysis tailored for both sub-tasks in Ta-

¹Noted TSA remains a challenging task with minimal progress in recent research, making a 1.38% improvement significant.

Model		Laptop			Restaurant			Tweets		
		Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
Pipeline	CRF-Pipeline [†] [Mitchell <i>et al.</i> , 2013]	0.5969	0.4754	0.5293	0.5228	0.5101	0.5164	0.4297	0.2521	0.3173
	NN-CRF-Pipeline [†] [Zhang <i>et al.</i> , 2015]	0.5772	0.4932	0.5319	0.6009	0.6193	0.6100	0.4371	0.3712	0.4006
	TAG-Pipeline [†] [Hu <i>et al.</i> , 2019]	0.6584	0.6719	0.6651	0.7166	0.7645	0.7398	0.5424	0.5437	0.5426
	SPAN-Pipeline [†] [Hu <i>et al.</i> , 2019]	0.6946	0.6672	0.6806	0.7614	0.7334	0.7492	0.6072	0.5502	0.5769
End-to-End	SPJM [†] [Zhou <i>et al.</i> , 2019]	0.6140	0.5820	0.5976	0.7620	0.6820	0.7198	0.5484	0.4844	0.5144
	SPAN-Joint [†] [Hu <i>et al.</i> , 2019]	0.6741	0.6199	0.6459	0.7232	0.7261	0.7247	0.5703	0.5269	0.5455
	S-AESC [†] [Lv <i>et al.</i> , 2021]	0.6687	0.6492	0.6588	0.7826	0.7050	0.7418	0.5586	0.5374	0.5473
	HI-ASA [‡] [Chen <i>et al.</i> , 2022b]	0.6796	0.6625	0.6709	0.7915	0.7621	0.7765	0.5732	0.5622	0.5676
	DCS [‡] [Li <i>et al.</i> , 2023]	0.6812	0.6640	0.6725	0.7835	0.7751	0.7793	0.5862	0.5835	0.5848
	MiniConGTS [‡] [Sun <i>et al.</i> , 2024]	0.7206	0.6725	0.6957	0.7926	0.7952	0.7939	0.6164	0.5728	0.5938
	PDGN [‡] [Zhu <i>et al.</i> , 2024]	0.7025	0.6812	0.6921	0.8036	0.7985	0.8010	0.6235	0.5924	0.6076
	AIFI [‡] [Chen <i>et al.</i> , 2024a]	0.7105	0.6915	0.7009	0.7925	0.8034	0.7979	0.6342	0.5911	0.6119
LLM-based	GPT-3.5-turbo Zero-Shot [‡]	0.3462	0.4065	0.3739	0.6221	0.6605	0.6407	0.3750	0.2868	0.3250
	GPT-3.5-turbo Few-Shot [‡]	0.3389	0.4452	0.3855	0.5847	0.6605	0.6203	0.3812	0.2841	0.3256
	GPT-3.5-turbo CoT [‡]	0.3430	0.4581	0.3923	0.6624	0.6420	0.6520	0.4233	0.2752	0.3335
	GPT-3.5-turbo CoT + Few-Shot [‡]	0.3532	0.4581	0.3989	0.6215	0.6790	0.6490	0.3752	0.3012	0.3342
	GPT-4o Zero-Shot [‡]	0.3214	0.4065	0.3590	0.6242	0.5741	0.5981	0.2326	0.3704	0.2857
	GPT-4o Few-Shot [‡]	0.3299	0.4194	0.3693	0.6571	0.5679	0.6093	0.2532	0.3631	0.2984
	GPT-4o CoT [‡]	0.3371	0.3806	0.3576	0.6643	0.5741	0.6159	0.2823	0.3891	0.3272
	GPT-4o CoT + Few-Shot [‡]	0.3622	0.4323	0.3941	0.6842	0.6420	0.6624	0.4022	0.3842	0.3930
Ours	FCKT	0.7599	0.6740	0.7144*	0.8449	0.7877	0.8153*	0.6512	0.5962	0.6225*

Table 2: The overall performance comparison is conducted on three real-world datasets. The “†” denotes results directly taken from the original papers, while “‡” indicates results reproduced following the methods described in the original papers. Bold values highlight the best performance, and * indicates statistical significance with a p -value ≤ 0.05 compared to the best-performing baseline.

Task	Method	Laptop	Restaurant	Tweets
AE	HI-ASA(COLING22)	0.8424	0.8511	0.7546
	DCS(EMNLP23)	0.8455	0.8462	0.7543
	MiniConGTS(EMNLP24)	0.8374	0.8432	0.7514
	PDGN(ACL24)	0.8485	0.8546	0.7593
	GPT-4o CoT + Few-Shot	0.5120	0.6010	0.4724
	AIFI(AAAI24)	0.8511	0.8631	0.7634
	FCKT	0.8534	0.8685	0.7754
SP	HI-ASA(COLING22)	0.8531	0.9257	0.8451
	DCS(EMNLP23)	0.8467	0.9212	0.8429
	MiniConGTS(EMNLP24)	0.8541	0.9194	0.8349
	PDGN(ACL24)	0.8519	0.9224	0.8496
	GPT-4o CoT + Few-Shot	0.7045	0.7538	0.6323
	AIFI(AAAI24)	0.8594	0.9305	0.8496
	FCKT	0.8612	0.9310	0.8524

Table 3: The performance comparisons with different methods on aspect extraction (F1 score) and sentiment prediction (accuracy).

ble 3. In AE, our proposed FCKT is able to achieve the best overall performances in terms of F1 score. Compared with the other two datasets, the improvement of model on Tweets is more obvious (+2%). Since the context length of Tweets is usually shorter, our proposed fine-grained transfer methods are more adequate between two sub-tasks, leading to more significant improvements compared to other two datasets.

For SP, FCKT consistently outperforms other approaches across the three datasets. However, its improvement on the Restaurant dataset is less pronounced compared to the other datasets. We hypothesize that this discrepancy may stem from the relatively high sample diversity in the Restaurant dataset, where the limited interaction information might be insuffi-

Task	AKT	TCL	Laptop	Restaurant	Tweets	Avg. Drop
AE	✓	✓	0.8534	0.8685	0.7754	-
	✓	✗	0.8420	0.8612	0.7713	▼ 0.91%
	✗	✓	0.8451	0.8603	0.7714	▼ 0.82%
	✗	✗	0.8412	0.8568	0.7630	▼ 1.45%
SP	✓	✓	0.8612	0.9310	0.8524	-
	✓	✗	0.8502	0.9214	0.8413	▼ 1.20%
	✗	✓	0.8492	0.9213	0.8381	▼ 1.36%
	✗	✗	0.8334	0.9189	0.8293	▼ 2.38%
TSA	✓	✓	0.7144	0.8153	0.6225	-
	✓	✗	0.6942	0.8064	0.6153	▼ 1.67%
	✗	✓	0.6775	0.7842	0.6072	▼ 3.86%
	✗	✗	0.6724	0.7792	0.5942	▼ 4.93%

Table 4: The results of different modules on each task. Note that “✓/✗ AKT” indicates whether the aspect knowledge transfer strategy is used, while “✓/✗ TCL” denotes the inclusion or exclusion of the token-level contrastive learning mechanism.

cient to substantially enhance sentiment prediction.

4.4 Ablation Study (RQ3)

To examine the contributions of various components, we delve further into FCKT and carry out ablation studies. The results are shown in Table 4. It is evident that the removal of specific modules leads to a decrease in model performance, highlighting the indispensable nature of each module. This underscores the critical role that various model components play in achieving optimal performance. If we delve deeper into the comparisons in TSA task, we can find the “AKT” module holds greater significance, indicated by its lower F1




Figure 4: The results of FCKT w.r.t different parameter ξ on two datasets. “FCKT w/o RAI” refers to removing real aspect inputs.




Figure 5: The F1 score of FCKT w.r.t varying parameter h and λ .

score compared to the other module. This is not surprising, as cross-task knowledge transfer is the core focus of our work, emphasizing the importance of effectively leveraging mutual information between tasks to enhance performance in TSA.

4.5 Parameter Analysis (RQ4)

Effect of parameter ξ . In FCKT, a critical parameter is the ratio ξ , which controls the connection between different tasks. This parameter determines how much of the predicted aspect knowledge is transferred to the second task. We tune ξ within the range of 0 to 1, and the results, based on the TSA task, are presented in Figure 4. Across these datasets, we observe that the F1 score initially increases rapidly, then stabilizes, and finally declines as ξ increases from 0.8 or 0.9 to 1. The reason lies in that when ξ is small, the model lacks sufficient supervision from the ground truth, leading to suboptimal performance. And when ξ is very large, the model fails to leverage the valuable interactions between the two sub-tasks, also resulting in unsatisfactory performance. We observe that the best performances are usually achieved when ξ is moderate on two datasets, which implies that neither the ground truth boundary nor distributional boundary is dominantly superior, and a mixture of them can be more favorable.

Effect of parameter h . In this study, we investigate the influence of the parameter h as it regulates the length of the aspect in Eq. (8). Theoretically, when h takes on a large value, the model ends up with an enormous number of computing memory, making it challenging to achieve convergence. Conversely, when h is excessively small, accurately extracting aspects becomes non-trivial. Considering that aspect term lengths are typically not extensive, we conduct experiments by varying h within the range of $\{1, 2, 3, 4\}$. In our experimental results, presented in Figure 5, we showcase F1 scores for both the Restaurant and Laptop domains. As h increases, we observe a gradual rise in performance up to a peak, followed by a decline. This suggests that maintaining a moderate h value (such as 2 or 3) can not only reduce the parameter

Input Reviews	AIFI Model	FCKT (Ours)
(1) All the money went into the <u>interior decoration</u> , none of it went to the <u>chefs</u> .	[interior decoration, pos] [chefs, pos] X ✓	[interior decoration, pos] ✓ [chefs, neg] ✓
(2) You will obtain a gift if you buy the separate <u>ram memory</u> .	[separate ram memory, pos] X	[ram memory, pos] ✓
(3) I must say I am surprised by the bad reviews of the <u>restaurant</u> , though the <u>menu's font</u> is small.	[restaurant, neg] X [menu's font, neg] ✓	[restaurant, neg] X [menu's font, neg] ✓

Figure 6: Some examples of different models. The positive aspects are marked in green, and negative aspects are marked in orange. ✓ and X denote correct and incorrect predictions.

count but also ensure accurate aspect extraction.

Effect of parameter λ . We also analyze the effect of the contrastive weight λ on F1-score across two datasets: Restaurant and Laptop. The results indicate that the optimal performance is achieved when $\lambda = 0.1$. As λ increases beyond this point, the F1-score decreases steadily for both datasets, with a sharper decline in the Restaurant domain. This suggests that a smaller λ effectively balances the contrastive and task-specific objectives, while higher values overly emphasize the contrastive loss, negatively impacting overall performance.

4.6 Case Study and Error Analysis (RQ5)

Case Study. To offer a deeper understanding of FCKT’s exceptional performance, we present several case studies in a qualitative manner in Figure 6. In 1st case, the aspect detection and sentiment prediction in AIFI may be compromised due to weak sequential associations. For instance, it successfully identifies the aspect “chefs” but falters in accurately predicting the sentiment associated with “chefs” in the 1st case. Conversely, FCKT allows for the backpropagation of sentiment information into the boundary detection process, thereby bolstering correlations and leading to accurate predictions. By modeling semantic compatibility for aspect extraction via token-level contrastive learning, FCKT can extract more precise aspects. For example, in the 2nd case, word “separate” should not be included in the aspect of “ram memory”, a distinction that FCKT successfully identifies.

Error Analysis. While the FCKT model has shown favorable performance, it still encounters challenges when handling complex sentences. For example, in the 3th sentence, due to the ambiguously expressed sentiment associated with “restaurant”, the model struggles to make precise predictions.

5 Conclusion

In this paper, we addressed the problem of targeted sentiment analysis (TSA) by proposing a fine-grained cross-task knowledge transfer (FCKT) framework. By explicitly integrating aspect-level information into sentiment prediction, FCKT enables fine-grained knowledge transfer, effectively mitigating negative transfer and improving task performance. Extensive experiments conducted on three real-world benchmark datasets, including comparisons with diverse baselines and state-of-the-art LLMs validate the effectiveness of FCKT.

Acknowledgments

We sincerely thank all the anonymous reviewers for their valuable comments to improve this paper. This work was supported by the National Key Research and Development Program of China under Grant Nos. 2024YFF0729003, the National Natural Science Foundation of China under Grant Nos. 62176014, 62276015, 62206266, the Fundamental Research Funds for the Central Universities.

References

- [Cai *et al.*, 2021] Hongjie Cai, Rui Xia, and Jianfei Yu. Aspect-category-opinion-sentiment quadruple extraction with implicit aspects and opinions. In *ACL-IJCNLP*, 2021.
- [Chen *et al.*, 2022a] Hao Chen, Zepeng Zhai, Fangxiang Feng, Ruifan Li, and Xiaojie Wang. Enhanced multi-channel graph convolutional network for aspect sentiment triplet extraction. In *ACL*, 2022.
- [Chen *et al.*, 2022b] Wei Chen, Jinglong Du, Zhao Zhang, Fuzhen Zhuang, and Zhongshi He. A hierarchical interactive network for joint span-based aspect-sentiment analysis. In *COLING*, 2022.
- [Chen *et al.*, 2024a] Wei Chen, Yuxuan Liu, Zhao Zhang, Fuzhen Zhuang, and Jiang Zhong. Modeling adaptive inter-task feature interactions via sentiment-aware contrastive learning for joint aspect-sentiment prediction. In *AAAI*, 2024.
- [Chen *et al.*, 2024b] Wei Chen, Yiqing Wu, Zhao Zhang, Fuzhen Zhuang, Zhongshi He, Ruobing Xie, and Feng Xia. Fairgap: Fairness-aware recommendation via generating counterfactual graph. *TOIS*, 2024.
- [Chen *et al.*, 2024c] Wei Chen, Meng Yuan, Zhao Zhang, Ruobing Xie, Fuzhen Zhuang, Deqing Wang, and Rui Liu. Fairdgc: Fairness-aware recommendation with dynamic graph contrastive learning. *arXiv preprint arXiv:2410.17555*, 2024.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.
- [Hu *et al.*, 2019] Minghao Hu, Yuxing Peng, Zhen Huang, Dongsheng Li, and Yiwei Lv. Open-domain targeted sentiment analysis via span-based extraction and classification. In *ACL*, 2019.
- [Jaiswal *et al.*, 2020] Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 2020.
- [Kalbhor and Goyal, 2023] Shraddha Kalbhor and Dinesh Goyal. Survey on absa based on machine learning, deep learning and transfer learning approach. In *AIP Conference Proceedings*, 2023.
- [Li *et al.*, 2023] Pan Li, Ping Li, and Kai Zhang. Dual-channel span for aspect sentiment triplet extraction. In *EMNLP*, 2023.
- [Liang *et al.*, 2021] Bin Liang, Wangda Luo, Xiang Li, Lin Gui, Min Yang, Xiaoqi Yu, and Ruifeng Xu. Enhancing aspect-based sentiment analysis with supervised contrastive learning. In *CIKM*, 2021.
- [Lin and Yang, 2020] Peiqin Lin and Meng Yang. A shared-private representation model with coarse-to-fine extraction for target sentiment analysis. In *EMNLP: Findings*, 2020.
- [Liu, 2012] Bing Liu. Sentiment analysis and opinion mining. In *Synthesis Lectures on Human Language Technologies.*, 2012.
- [Luo *et al.*, 2019] Huaishao Luo, Tianrui Li, Bing Liu, and Junbo Zhang. Doer: Dual cross-shared rnn for aspect term-polarity co-extraction. In *ACL*, 2019.
- [Luo *et al.*, 2024] Haoran Luo, Haihong E, Yuhao Yang, Tianyu Yao, Yikai Guo, Zichen Tang, Wentai Zhang, Shiyao Peng, Kaiyang Wan, Meina Song, Wei Lin, Yifan Zhu, and Anh Tuan Luu. Text2nkg: Fine-grained n-ary relation extraction for n-ary relational knowledge graph construction. In *NeurIPS*, 2024.
- [Lv *et al.*, 2021] Yanxia Lv, Fangna Wei, Ying Zheng, Cong Wang, Cong Wan, and Cuirong Wang. A span-based model for aspect terms extraction and aspect sentiment classification. *Neural Computing and Applications*, 2021.
- [Mitchell *et al.*, 2013] Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme. Open domain targeted sentiment. In *EMNLP*, 2013.
- [Nath and Dwivedi, 2024] Deena Nath and Sanjay K Dwivedi. Aspect-based sentiment analysis: approaches, applications, challenges and trends. *KIS*, 2024.
- [Sun *et al.*, 2024] Qiao Sun, Liujia Yang, Minghao Ma, Nanyang Ye, and Qinying Gu. MiniConGTS: A near ultimate minimalist contrastive grid tagging scheme for aspect sentiment triplet extraction. In *EMNLP*, 2024.
- [Tang *et al.*, 2020] Hao Tang, Donghong Ji, Chenliang Li, and Qiji Zhou. Dependency graph enhanced dual-transformer structure for aspect-based sentiment classification. In *ACL*, 2020.
- [Wang *et al.*, 2016] Wenya Wang, Sinno Jialin Pan, Daniel Dahlmeier, and Xiaokui Xiao. Recursive neural conditional random fields for aspect-based sentiment analysis. *EMNLP*, 2016.
- [Wang *et al.*, 2023] Zengzhi Wang, Qiming Xie, Zixiang Ding, Yi Feng, and Rui Xia. Is chatgpt a good sentiment analyzer? a preliminary study. *arXiv preprint arXiv:2304.04339*, 2023.
- [Wei *et al.*, 2022] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*, 2022.
- [Xiong *et al.*, 2022] Haoliang Xiong, Zehao Yan, Hongya Zhao, Zhenhua Huang, and Yun Xue. Triplet contrastive learning for aspect level sentiment classification. *Mathematics*, 2022.

- [Xu *et al.*, 2021] Lu Xu, Yew Ken Chia, and Lidong Bing. Learning span-level interactions for aspect sentiment triplet extraction. In *ACL*, 2021.
- [Yang *et al.*, 2020] Yunyi Yang, Kun Li, Xiaojun Quan, Weizhou Shen, and Qinliang Su. Constituency lattice encoding for aspect term extraction. In *COLING*, 2020.
- [Yuan *et al.*, 2023] Meng Yuan, Fuzhen Zhuang, Zhao Zhang, Deqing Wang, and Jin Dong. Knowledge-based multiple adaptive spaces fusion for recommendation. In *RecSys*, 2023.
- [Yuan *et al.*, 2025a] Meng Yuan, Yutian Xiao, Wei Chen, Chu Zhao, Deqing Wang, and Fuzhen Zhuang. Hyperbolic diffusion recommender model. *WWW*, 2025.
- [Yuan *et al.*, 2025b] Meng Yuan, Zhao Zhang, Wei Chen, Chu Zhao, Tong Cai, Deqing Wang, Rui Liu, and Fuzhen Zhuang. Hek-cl: Hierarchical enhanced knowledge-aware contrastive learning for recommendation. *TOIS*, 2025.
- [Zhang *et al.*, 2015] Meishan Zhang, Yue Zhang, and Duy Tin Vo. Neural networks for open domain targeted sentiment. In *EMNLP*, 2015.
- [Zhang *et al.*, 2022] Rui Zhang, Yangfeng Ji, Yue Zhang, and Rebecca J Passonneau. Contrastive data and learning for natural language processing. In *NAACL: Human Language Technologies: Tutorial Abstracts*, 2022.
- [Zhong *et al.*, 2024a] Yan Zhong, Xingyu Wu, Li Zhang, Chenxi Yang, and Tingting Jiang. Causal-iqa: Towards the generalization of image quality assessment based on causal inference. In *ICML*, 2024.
- [Zhong *et al.*, 2024b] Yan Zhong, Ruoyu Zhao, Chao Wang, Qinghai Guo, Jianguo Zhang, Zhichao Lu, and Luziwei Leng. Spike-ssm: A sparse, precise, and efficient spiking state space model for long sequences learning. *arXiv preprint arXiv:2410.17268*, 2024.
- [Zhong *et al.*, 2025] Yan Zhong, Xinping Zhao, Guangzhi Zhao, Bohua Chen, Fei Hao, Ruoyu Zhao, Jiaqi He, Lei Shi, and Li Zhang. Ctd-inpainting: Towards the coherence of text-driven inpainting with blended diffusion. *Information Fusion*, 2025.
- [Zhou *et al.*, 2019] Yan Zhou, Longtao Huang, Tao Guo, Jizhong Han, and Songlin Hu. A span-based joint model for opinion target extraction and target sentiment classification. In *IJCAI*, 2019.
- [Zhou *et al.*, 2024] Changzhi Zhou, Dandan Song, Yuhang Tian, Zhijing Wu, Hao Wang, Xinyu Zhang, Jun Yang, Ziyi Yang, and Shuhao Zhang. A comprehensive evaluation of large language models on aspect-based sentiment analysis. *arXiv preprint arXiv:2412.02279*, 2024.
- [Zhu *et al.*, 2024] Linan Zhu, Xiangfan Chen, Xiaolei Guo, Chenwei Zhang, Zhechao Zhu, Zehai Zhou, and Xiangjie Kong. Pinpointing diffusion grid noise to enhance aspect sentiment quad prediction. In *ACL: Findings*, 2024.