

Improving Implicit Discourse Relation Recognition with Natural Language Explanations from LLMs

Heng Wang¹, Changxing Wu^{1*}

¹ School of Information and Software Engineering, East China Jiaotong University, Nanchang, China
2023218083500007@ecjtu.edu.cn, wuchangxing@ecjtu.edu.cn

Abstract

Implicit Discourse Relation Recognition (IDRR) remains a challenging task due to the requirement for deep semantic understanding in the absence of explicit discourse markers. A further limitation is that existing methods only predict relations without providing any supporting explanations. Recent advances in large language models (LLMs) have shown strong reasoning capabilities in both deep language understanding and natural language explanation generation. In this work, we propose a simple yet effective approach to distill the reasoning capabilities of LLMs into lightweight IDRR models to improve both performance and interpretability. Specifically, we first prompt an LLM to generate explanations for each training instance conditioned on its gold label. Then, we introduce a novel classification-generation framework that jointly performs relation prediction and explanation generation, and train it with the additional supervision of LLM-generated explanations. Our framework is plug-and-play, enabling easy integration with most existing IDRR models. Experimental results on PDTB demonstrate that our approach significantly improves IDRR performance, while human evaluation further confirms that the generated explanations enhance model interpretability. Furthermore, we validate the generality of our approach on sentiment classification and natural language inference.

Code, Appendix — <https://github.com/nlper-hub/EIDRR>

Introduction

To fully understand natural language text, it is essential not only to comprehend the meaning of individual sentences but also to grasp the semantic relationships that link them, known as discourse relations (e.g., *Comparison*). Discourse relation recognition has attracted considerable attention in natural language processing (NLP) due to its potential to enhance both language understanding and generation (Li, Wu, and Li 2020; Tang et al. 2021; Hu and Wan 2023; Li, Yin, and Carenini 2024). Among these, implicit discourse relations—which lack explicit connectives such as *but*, *and*, or *because*—pose a greater challenge for both humans and machines to identify. Despite recent progress in NLP and the emergence of LLMs, implicit discourse relation recognition

*Corresponding author.

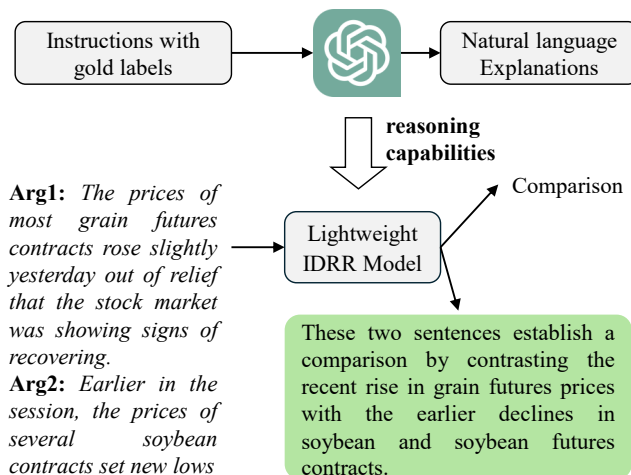


Figure 1: Our lightweight IDRR model distills LLM reasoning to predict discourse relations and generate interpretable explanations. **Arg1** and **Arg2** are two arguments.

(IDRR) remains a challenging task (Chan et al. 2024; Yung et al. 2024).

Implicit discourse relation recognition is typically formulated as a classification task, where the goal is to predict the discourse relation connecting two given arguments (e.g., sentences or clauses). With the advancement of deep learning, IDRR models that once relied on human-designed features (Park and Cardie 2012; Rutherford and Xue 2014) have transitioned to neural network-based models. Early research efforts primarily focused on designing task-specific neural network structures to better capture the semantics of arguments and model their interactions (Zhang et al. 2015; Bai and Zhao 2018; Liu et al. 2020). Recently, researchers have enhanced IDRR models by leveraging relation hierarchies, adopting either generation models (Wu et al. 2022) or contrastive learning approaches (Long and Webber 2022). More recently, prompt-based IDRR models have demonstrated superior performance by utilizing pre-trained language models to predict discourse connectives and then mapping them to corresponding relations (Xiang et al. 2022; Zhao et al. 2023; Zeng et al. 2024). Despite their notable performance gains, current IDRR models output only predicted relations

without providing any supporting explanations, which limits their interpretability and reduces user trust.

Recent advances in large language models (LLMs) have shown strong reasoning capabilities, including both deep language understanding and the ability to generate natural language explanations (Madsen, Chandar, and Reddy 2024; Bilal, Ebert, and Lin 2025). However, empirical results indicate that LLMs, when used with prompting or in-context learning, significantly underperform compared to lightweight IDRR models trained on human-annotated data (Chan et al. 2024), and also exhibit limited faithfulness (Miao et al. 2024). While fine-tuning LLMs can improve performance, their heavy computational demands and high deployment cost limit their applicability in low-resource environments. These limitations raise a natural question: *can we transfer the reasoning capabilities of LLMs into lightweight models to enable more accurate and interpretable IDRR?*

To address this question, we propose a simple yet effective distillation approach using natural language explanations as a bridge. As shown in Figure 1, our approach consists of two main stages. **First**, we prompt an LLM to generate natural language explanations for each training instance, conditioned on its gold label. To obtain high-quality outputs, we draw inspiration from the Chain-of-Thought concept (Wei et al. 2022). Specifically, we structure each explanation into two parts: the restatement of two arguments (serving as a thought step, not shown in Figure 1 for clarity) and a rationale of their discourse relationship. To some extent, these generated explanations not only offer a deep understanding of the input instances but also reflect the reasoning patterns inherent in LLMs. **In the second stage**, we introduce a classification-generation framework equipped with a transformer module. This framework learns to jointly predict discourse relations and generate explanations under the combined supervision of gold discourse labels and LLM-generated explanations. To facilitate effective interaction between the two tasks, we adopt a multi-task learning strategy with a shared encoder. The transformer module is introduced to alleviate the potential mismatch when using a shared encoder for both classification and generation objectives. Our framework is plug-and-play and can be readily combined with most existing IDRR models by adding a simple transformer module and a generation component, with minimal changes required.

The primary contributions of our work are as follows:

- We make an initial attempt toward interpretable IDRR by jointly predicting discourse relations and generating natural language explanations.
- We propose a plug-and-play classification-generation framework that facilitates the distillation of LLM reasoning capabilities into lightweight IDRR models.
- Experimental results and human evaluations confirm that our method significantly improves IDRR performance and delivers high-quality, interpretable explanations.
- We further validate the generalizability of our approach on two additional tasks: sentiment classification and natural language inference.

Related Work

Neural IDRR Methods

Neural network-based methods have significantly advanced this area. Previous researchers focused on designing diverse neural network architectures to effectively capture the interactions between arguments (Zhang et al. 2018; Bai and Zhao 2018; Zhang et al. 2021). During this phase, the utilization of pre-trained language models (PLMs) like BERT and RoBERTa significantly improved the performance (Liu et al. 2020). Meanwhile, some studies began exploring generation-based methods to better capture relational semantics. For example, Jiang et al. (2021) proposed CG-T5, which treats IDRR as a generation task that jointly predicts discourse relations and generates sentences conveying their meanings. Recently, researchers have enhanced IDRR models by integrating label hierarchy information. Wu et al. (2022) viewed multi-level IDRR as a label sequence generation task, enabling the effective utilization of inter-label dependencies. Long and Webber (2022) and Jiang, Zhang, and Wang (2023) introduced additional contrastive learning losses based on the hierarchical structure of labels. More recently, prompt-based learning has emerged as a powerful paradigm, achieving the SOTA performance in IDRR. Xi-ang et al. (2022) and Zhou et al. (2022b) created different prompt templates to encourage PLMs to predict connectives that link two arguments, which are then used to determine discourse relations. Zhao et al. (2023), Chan et al. (2023) and Zeng et al. (2024) infused hierarchical label information into prompt tuning. The scarcity of annotated data in IDRR has driven the exploration of data-augmented approaches. Researchers employed conditional variational autoencoders (Dou et al. 2021) or large language models (Omura, Cheng, and Kurohashi 2024) to create diverse, high-quality synthetic data to enhance the annotated dataset. Others sought to leverage abundant explicit discourse data, naturally annotated with connectives, to pre-train prompt-based models (Wang, Jian, and Huang 2023; Liu and Strube 2023). Unlike existing methods that focus primarily on classification without providing explanations for their predictions, our approach enhances both predictive performance and model interpretability.

Natural Language Explanations

There is an increasing interest in providing natural language explanations for the decisions made by neural models (Lyu, Apidianaki, and Callison-Burch 2024). Based on manually-annotated explanations, Liu, Yin, and Wang (2019) proposed a generative explanation framework for classification tasks, Kumar and Talukdar (2020) automatically generated explanations for each possible label of an instance and used them to make the decision with an explanation. Recently, with explanations generated by LLMs, Wang et al. (2023) mitigated spurious correlations in aspect-based sentiment classification, Ludan et al. (2023) proposed an explanation-based fine-tuning method to enhance robustness, Kroeger et al. (2024) showed that their in-context explainer framework enables LLMs to generate explanations on par with the best explainers. Liu et al. (2025) and Fan, Li, and Zhu (2025)

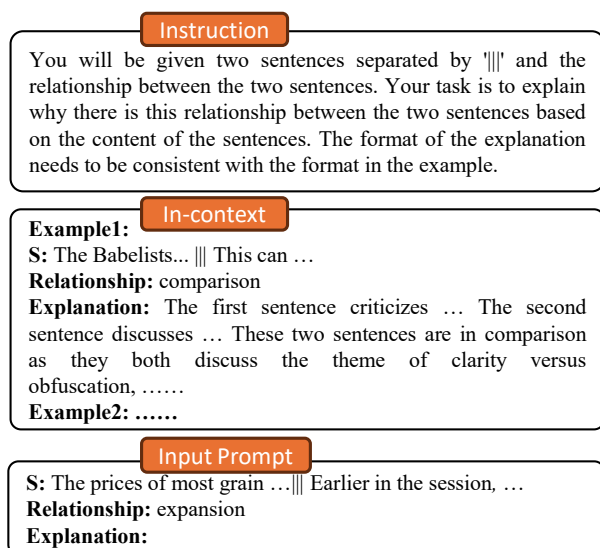


Figure 2: The prompt template for explanation generation. An explanation consists of two parts: the restatement of two arguments, and a rationale of their discourse relationship. S means two arguments.

leveraged explanatory texts generated by LLMs to enhance the ability to handle complex semantic phenomena in multi-party dialogues, which improved performance in discourse analysis. Inspired by these studies, we present the first exploration of interpretable IDRR by leveraging explanations generated by LLMs.

Method

Explanation-enhanced PDTB Construction

PDTB (Prasad et al. 2008) is widely recognized as the largest and most commonly used corpus for IDRR. An implicit discourse instance in PDTB can be depicted as (arg_1, arg_2, y) , where arg_1 and arg_2 are two arguments and y denotes the manually-annotated discourse label. Recent studies have shown that LLMs like GPT3 (Brown et al. 2020) can generate high-quality explanations for classification tasks, such as sentiment classification (Wang et al. 2023; Kroeger et al. 2024). Motivated by this, we use LLMs to generate a natural language explanation e for each PDTB instance, conditioned on its gold label.

In the explanation generation stage, we design a prompt template based on In-Context Learning (Brown et al. 2020) to activate the understanding and inference capabilities of LLMs. As illustrated in Figure 2, the template comprises three components: instructions, in-context examples, and the input prompt. The instructions clearly explain the generation task and provide guidance on the output format. In-context examples offer further support for LLMs to comprehend the task. The input prompt embeds the arguments and the gold label of an implicit discourse instance for processing. Following Chain-of-Thought, we format an explanation by first restating the two arguments (i.e., *the first sentence criticizes...*, *the second sentence discusses*) and then explaining

why they express a specific discourse relation (i.e., *discuss the theme of clarity versus obfuscation...*). Incorporating the restating of arguments serves two purposes: 1) it encourages the LLM to perform deeper semantic interpretation of the input arguments, and 2) it functions as a thought step in the Chain-of-Thought process, leading to improved explanation generation.

We compared the explanation quality of ChatGPT-4o, Qwen-long, and Ernie Bot-3.5-128K, and finally chose Qwen-long. We randomly selected 1,000 instances and found that most explanations provide reasonable rationales for the given gold labels, with an average human evaluation score of 4.53 out of 5 (See evaluation criteria on P6). Therefore, we only manually corrected explanations of instances that were inconsistent with their labels. We identified two main challenges in explanation generation. 1) Ethical concerns: LLMs reject instructions when encountering certain keywords. In such cases, we resort to alternative LLMs or manual processing. 2) Insufficient context: Short arguments often lead to incorrect or contradictory explanations. To mitigate this, we supplement the input with additional context from the source document. In the end, we obtained explanations for 15,004 PDTB instances with minimal human effort, requiring manual post-processing for fewer than 100 cases.

Our Classification-Generation Framework

We expect our framework to possess the following capabilities: 1) Strong performance on both the IDRR task and the associated explanation generation task, with the two tasks mutually enhancing each other; 2) Explanations that are well-aligned with the model’s classification decisions; 3) Plug-and-play compatibility, enabling seamless integration with most existing IDRR models. A straightforward strategy is to reformulate the classification task as a generation task, enabling a generation-based model (e.g., T5) to produce a label and its explanation, in flexible order. However, we identify two main drawbacks of this strategy: 1) it often fails to achieve performance comparable to classification-based or prompt-based IDRR models of similar scale; and 2) it lacks generality due to challenges in integrating with existing classification models.

We propose a simple yet effective framework that learns to jointly predict discourse relations and generate explanations, guided by both gold labels and LLM-generated explanations. As shown in Figure 3, our framework consists of a shared encoder, a module for classification, and a decoder paired with a transformer module for explanation generation. The two tasks can interact and benefit from each other through the shared encoder. Following prior work in IDRR, we adopt RoBERTa-like PLMs (Liu et al. 2019) as the shared encoder. To handle the limited training data, we employ the decoder component of T5-like PLMs (Rafael et al. 2020) for generation. However, directly combining the two types of PLMs causes a mismatch, as they are pre-trained for different purposes: one for natural language understanding and the other for generation. We introduce a transformer module—consisting of a randomly initialized transformer layer—to mitigate the mismatch, and find it effective in practice. To ensure better alignment between clas-

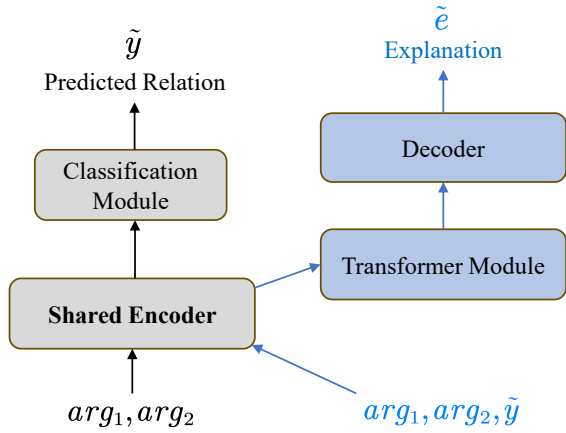


Figure 3: Our proposed classification-generation framework. For ease of understanding, we use the IDRR task as an example to illustrate our framework. It first predicts the discourse relation \tilde{y} , which is then used as extra input to generate the explanation \tilde{e} .

sification outputs and generated explanations, we first predict the discourse label and subsequently condition the explanation generation on this label.

Our plug-and-play framework is compatible with both classification-based and prompt-based IDRR models. We exemplify its integration using prompt-based models due to their superior performance. Given an instance (arg_1, arg_2) , a prompt-based IDRR model usually calculates the predicted relation \tilde{y} as follows:

$$\begin{aligned} \tilde{y} &= \text{Verbalizer}(h_{mask}), \\ h_{mask} &= \text{RoBERTa}(T_a(arg_1, arg_2)), \end{aligned} \quad (1)$$

where T_a is the prompt template for IDRR, h_{mask} is the representation of the $\langle mask \rangle$ token. RoBERTa is a pre-trained masked language model. The Verbalizer predicts the most likely discourse connective (e.g., *but*) at the masked position, which is subsequently mapped to a discourse relation (e.g., *Comparison*) based on handcrafted rules detailed in Appendix A. In our experiments, we use the template¹ defined in (Zhou et al. 2022b), and the Verbalizer in (Xiang et al. 2022). We omit the details of Verbalizer for brevity. During training, we adopt the standard cross-entropy loss for classification (as L_c in Equation 3).

For explanation generation, we only need to stack a minimal transformer module and a generation decoder on top of the shared RoBERTa encoder. Formally, the explanation \tilde{e} is generated as follows:

$$\begin{aligned} \tilde{e} &= \text{T5-Decoder}(\hat{H}), \\ \hat{H} &= \text{Transformers}(H_{last}), \\ H_{last} &= \text{RoBERTa}(T_b(arg_1, arg_2, \tilde{y})), \end{aligned} \quad (2)$$

where T_b is the prompt template² for explanation generation,

¹Arg1:arg1.Arg2:arg2.</s></s>The conjunction between Arg1 and Arg2 is <mask>.

²Arg1:arg1.Arg2:arg2.</s></s>The conjunction between Arg1 and Arg2 is \tilde{y} , the main reason is that.

H_{last} is the final layer outputs of RoBERTa, Transformers consists of several stacked Transformer layers and \hat{H} is the adaptive representation for generation. We adopt the standard auto-regressive cross-entropy loss for text generation as L_g in Equation 3.

During training, we adopt a multi-task learning approach. Given the prediction \tilde{y} and the generation explanation \tilde{e} , the total loss L_{total} is defined as:

$$L_{total} = \alpha L_c(y, \tilde{y}) + \beta L_g(e, \tilde{e}), \quad (3)$$

where L_c and L_g are the respective losses for classification and explanation generation, with y and e as the true label and explanation, α and β are the coefficients. During training, our framework usually exhibits good performance on the classification task after just a few epochs, but struggles with the generation task. To tackle this issue, we first emphasize training on the generation task ($\alpha < \beta$) and then shift focus to the classification task ($\alpha > \beta$).

The key difference from prior frameworks (Camburu et al. 2018; Liu, Yin, and Wang 2019) is that we introduce a Transformer module to bridge the gap between classification and generation modules — a crucial addition when combining separately pre-trained understanding and generation LMs.

Experiments

Dataset and Settings

To evaluate the effectiveness of our method, we carry out experiments on the PDTB corpus. Following previous work (Wu et al. 2022), we divide the corpus into three parts: Sections 2-20 as the training set (12,775 instances), Sections 0-1 as the validation set (1,183 instances), and Sections 21-22 as the test set (1,046 instances). Four primary top-level discourse relations are considered: *Temporal*, *Comparison*, *Contingency*, and *Expansion*. More details of these datasets can be found in Appendix B.

Our model is trained using a two-stage approach (Zhou et al. 2022a). In the first stage, we assign different learning rates: $5e-6$ for RoBERTa and $5e-5$ for the other modules. To emphasize the generation task, we set the loss weights to $\alpha = 0.4$ and $\beta = 0.6$, and trained for 30 epochs. In the second stage, we adjust the learning rate for all modules except RoBERTa to $3e-5$. The coefficient values $\alpha = 0.8$ and $\beta = 0.2$ are employed to shift the training emphasis towards the IDRR task. Our model reaches the optimal validation performance on IDRR within just 6 epochs. In both stages, we use the AdamW optimizer with a 0.3 dropout rate. The Transformer module has one randomly initialized layer, as adding more layers yields no gains. Following prior work, we evaluate using macro-averaged F_1 score and accuracy (Acc), averaging results over three random seeds on an NVIDIA 3090 GPU.

Performance of IDRR

We compare our method (named as EIDRR) with recent baselines, which primarily include the following:

LLM-based methods: ChatGPT is evaluated under the in-context learning setting (Chan et al. 2024). We formulate

Method	Acc (%)	F_1 (%)
ChatGPT	50.24	44.09
LLaMA-3B + IDRR _{gt}	65.11	57.11
LLaMA-8B + IDRR _{gt}	70.36	61.01
T5 + IDRR _{gt} w/o Exp	64.05	55.92
T5 + IDRR _{gt}	64.63	56.65
T5 + IDRR _{ct} w/o Exp	63.67	52.39
T5 + IDRR _{ct}	64.91	55.71
BMGF	69.06	63.39
CVAE	70.17	65.06
LDSGM	71.18	63.73
GOLF	72.52	65.76
SCIDER	72.11	67.00
PCP	70.84	64.95
PEMI	71.13	64.05
Discoprompt	71.70	65.79
EIDRR w/o Exp	71.80	66.87
EIDRR	73.14	68.01

Table 1: Comparison with baselines. IDRR_{gt} and IDRR_{ct} means that IDRR is formulated as a generation task and a classification task, respectively. Here w/ Exp and w/o Exp indicate whether the explanation generation task is included or omitted.

IDRR as a generation task and fine-tune³ LLaMA-3B and LLaMA-8B with LoRA.

T5-based methods: We reformulate IDRR as a generation task, utilizing T5 to produce label text and corresponding explanations sequentially (T5+IDRR_{gt}). We stack a classification layer on the encoder of T5 for IDRR and leverage the T5 decoder to generate explanations (T5+IDRR_{ct}) with the predicted labels as additional inputs.

Non-prompt-based methods: BMGF (Liu et al. 2020), CVAE (Dou et al. 2021), LDSGM (Wu et al. 2022), GOLF (Jiang, Zhang, and Wang 2023), SCIDER (Cai, Yang, and Jian 2024).

Prompt-based methods: PCP (Zhou et al. 2022b), PEMI (Zhao et al. 2023), and Discoprompt (Chan et al. 2023).

Based on the results in Table 1, we can draw the following conclusions. **Firstly**, a powerful model like ChatGPT struggles to achieve promising performance on IDRR without fine-tuning. While fine-tuned LLMs like LLaMA show some improvements (Part 1), they still lag behind prompt-based IDRR models (Part 4). **Secondly**, directly using an LM pre-trained for generation (T5-based methods, Part 2) still fails to obtain comparable performance to current models such as SCIDER and Discoprompt. The main reason is that LMs pre-trained for generation do not perform as well in classification tasks like IDRR as those pre-trained for language understanding. **Thirdly**, including the explanation generation task boosts the IDRR performance across all three settings (w/o Exp). These results strongly suggest that explanations

³<https://github.com/hiyouga/LLaMA-Factory>

Method	Acc (%)	F_1 (%)	Human-Avg
e-INFERSENT	73.04	65.79	1.71
EIDRR	73.14	68.01	4.20

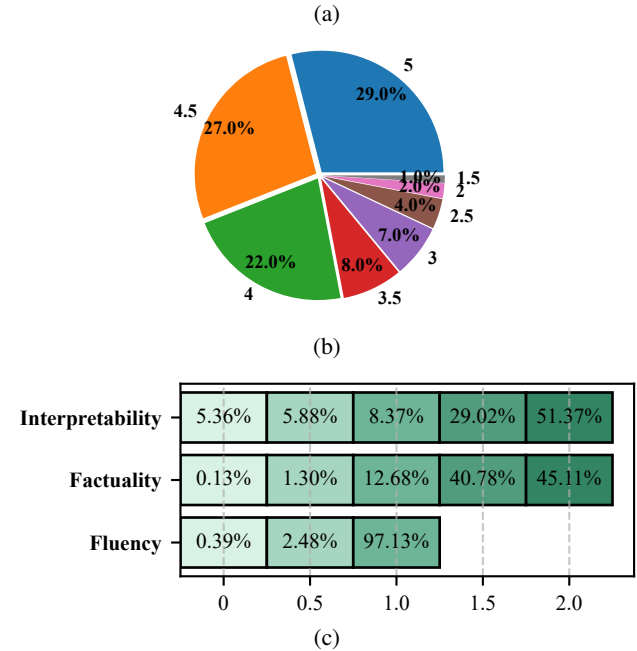


Figure 4: Quality of generated explanations. (a) Compared with e-INFERSENT, our method achieves a higher average score of 4.20 out of 5. (b) The distribution of scores assigned through human evaluation. (c) The distribution of scores for each aspect.

generated via LLMs are useful. **Lastly**, our EIDRR significantly outperforms the prompt-based Discoprompt, which leverages label hierarchy information. The improvement is primarily attributed to incorporating the explanation generation task in a multi-task learning way, with 1.34% and 1.14% gains in Acc and F_1 , respectively (EIDRR vs. w/o explanation). A closer look at the results reveals that EIDRR obtains higher F_1 scores for each relation. In addition, results of the ablation studies are provided in Appendix C.

Overall, by leveraging LLM-generated explanations as additional supervision, our classification-generation framework effectively improves IDRR performance.

Quality Evaluation on Explanation

Following Camburu et al. (2018) and Majumder et al. (2022), we restrict the evaluation to correctly predicted test instances, since explanations are conditioned on the predicted labels and are unlikely to be correct with wrong labels. Five human annotators were recruited for the evaluation. Each instance was independently scored by two annotators, and a third was involved if their scores differed by more than two points. The final score for each instance was computed as the average of the two closest scores. We asked annotators to evaluate the quality of explanations across three

aspects (more details are provided in *Appendix D*):

- **Interpretability:** *Does the rationale part justify the predicted label?* We assign scores of 2 (Yes), 1 (Weak-Yes), or 0 (No).
- **Factuality:** *Does the restatement part align with the facts stated in the arguments?* We assign scores of 2 (Yes), 1 (Weak-Yes), or 0 (No).
- **Fluency:** *Is the explanation grammatically correct and natural?* We assign scores of 1 (Yes) or 0 (No).

We compare our framework with e-INFERSENT (Camburu et al. 2018), which directly couples the classification encoder with the generation decoder. It also follows the PredictANDExplain strategy by prepending the predicted label to the input for explanation generation. For a fair comparison, we substitute the corresponding components in e-INFERSENT with RoBERTa and the T5 decoder, respectively. From the results in Figure 4 (a), our EIDRR outperforms e-INFERSENT in both IDRR and explanation generation. The poor explanation quality of e-INFERSENT (average score of 1.71) mainly stems from the incompatibility between the pre-trained T5 decoder and RoBERTa without the transformer module. A deeper analysis reveals that the explanations generated by e-INFERSENT often fail to align with the input arguments. In addition, we experimented with the ExplainThenPredict strategy, which first generates explanations and then uses them for classification. However, we found that generating explanations without access to labels often leads to inconsistencies, and using such explanations as additional input significantly degrades classification performance. This experimental observation is consistent with that reported in (Camburu et al. 2018).

As shown in Figure 4 (b), 93% of explanations scored 3 or higher, while only 3% scored 2 or below. More importantly, the average score of 4.20 out of 5 demonstrates that our framework trained with LLM-generated explanations produces fluent and interpretable outputs. Figure 4 (c) also presents the score distribution across evaluation aspects. In terms of interpretability, 88.76% of explanations scored 1 or above, indicating that most provide meaningful rationales. Regarding factuality, 98.57% scored 1 or higher, showing that our model effectively captures the key semantics of both arguments. The generated explanations are highly fluent, with 97.13% achieving a perfect score, benefiting from the strong generation capability of the T5 decoder.

Based on the results in Table 1 and Figure 4, we conclude that our approach effectively transfers the reasoning capabilities of LLMs to lightweight models, thereby improving IDRR performance while enhancing model interpretability.

Faithful Evaluation on Explanation

It remains unclear whether these explanations genuinely capture the IDRR model’s reasoning process. While evaluating the faithfulness of natural language explanations is inherently challenging, we address this by adopting *Feature Importance Agreement* and *Robust Equivalence* (Wiegrefe, Marasović, and Smith 2021; Majumder et al. 2022).

Feature Importance Agreement Features important for classification should also influence explanation generation,

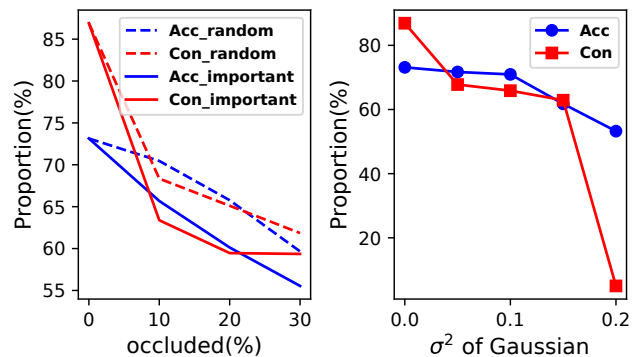


Figure 5: Feature Importance Agreement (Left) and Robustness Equivalence (Right). *Con* denotes the proportion of label–explanation alignment. *Random* and *Important* indicate that randomly selected or top-ranked important features are occluded, respectively.

and vice versa. To validate this, we employ a gradient-based attribution method to identify salient input features (tokens) ranked within the top- $\{10, 20, 30\}$ % attribution scores for the given task. Subsequently, we remove these features and quantify their impact on the model’s predictive performance. As illustrated in Figure 5 (Left), the removal of explanation-relevant features (blue solid line) causes a more pronounced reduction in classification accuracy than randomly removing features (blue dashed line). Likewise, eliminating classification-relevant features (red solid line) lowers label–explanation consistency to a greater extent than random removal (red dashed line). This indicates mutual importance of features across two tasks and supports the faithfulness of generated explanations.

Robustness Equivalence We inject zero-mean Gaussian noise $\mathcal{N}(0, \sigma^2)$ into internal features and observe changes in classification accuracy and label-explanation consistency. The goal is to assess whether both label prediction and explanation generation remain stable (or unstable) under similar noise conditions and exhibit similar trends in performance degradation. Figure 5 (Right) shows that, as the noise level increases, both metrics degrade gradually and in a similar trend, indicating stable model behavior. Under high noise (e.g., $\sigma^2 = 0.2$), both tasks fail: classification accuracy drops to 53.25%, comparable to that of a randomly initialized model, and the generated explanations become irrelevant to the input. This shared sensitivity to noise further supports the faithfulness of the generated explanations.

Case Study

We provide several examples in Table 2 to demonstrate the faithfulness of generated explanations. For the first example, the generated explanation captures the key *Temporal* cues, such as *reaching its highest point happened after the rally slowing and them moving forward*. However, the comparison model may have been influenced by the word *instead* in *Arg1* and incorrectly predicted a comparison relationship. The second example is very confusing even for humans. Is

Example and Predictions	Explanation
Arg1: Instead, the rally only paused for about 25 minutes and then steamed forward as institutions resumed buying. Arg2: The market closed minutes after reaching its high for the day. Relationship: Temporal EIDRR: Temporal✓ w/o explanation: Comparison ✗	The first sentence describes the slow start of a rally, which then moved forward with institutions resuming buying. The second sentence mentions the closing price of the market, which was just after reaching its highest point for the day. The temporal relationship is evident because the second event (the market closing after reaching its highest point) happened after the first event (the rally slowing and then moving forward). This sequence of events establishes the temporal connection between the two sentences.
Arg1: We view the market here as going through a relatively normal cycle Arg2: We continue to feel that the stock market is still the place to be for long-term appreciation. Relationship: Expansion EIDRR: Expansion✓ w/o explanation: Contingency ✗	The first sentence states that the market is seen as undergoing a relatively normal cycle. The second sentence provides additional information about the market’s continued positive outlook for long-term appreciation. The second sentence expands on the idea introduced in the first sentence by giving more detail about the market’s outlook and the reasons behind it. Thus, the relationship between the two sentences is one of expansion.
Arg1: It would have been too late to think about on Friday. Arg2: We had to think about it ahead of time. Relationship: Contingency EIDRR: Contingency✓ w/o explanation: Comparison ✗	The first sentence indicates a situation where thinking about something would have been too late on Friday. The second sentence states that there was a need to think about it before it could be decided. The contingency relationship is demonstrated here because the need for thinking ahead of time (second sentence) is contingent upon the situation described in the first sentence. In other words, the need for thinking ahead is based on the situation being discussed.
Arg1: Tokyo stocks closed off a significant but less-than-alarming 1.8 % on thin volume Arg2: Hong Kong stocks declined 6.5 % in orderly trading Relationship: Comparison EIDRR: Comparison✓ w/o explanation: Expansion ✗	The first sentence states that Tokyo stocks closed lower by 1.8% on a narrow volume, which is less than amalgamating, but less than amalgamating one. The second sentence mentions that Hong Kong stocks declined by 6.5% in stable trading. These two sentences form a comparison by contrasting the closing performance of Tokyo stocks with the performance of HK stocks in different markets. The comparison highlights the difference in performance between the two stocks in terms of their closing value and their performance in different markets.

Table 2: Examples of IDRR explanations. More examples are in *Appendix E* and on *Github* (all test instances).

Task		Acc (%)	F ₁ (%)	Human
ATSC	w/o explanation	85.59	70.18	-
	w/ explanation	87.44	71.30	4.01
NLI	w/o explanation	74.83	74.79	-
	w/ explanation	75.83	75.73	4.26

Table 3: The performance of ATSC and NLI.

it *Expansion* or *Contingency*? EIDRR generated a reasonable explanation as *giving more detail about the market’s outlook and the reasons behind it*. We can see that the explanation conveys both the *Expansion* relationship through words *more detail* and the *Contingency* relationship through *reasons*. EIDRR prioritizes predicting as the *Expansion* relationship. For the third example, the comparison model misclassified it as *Comparison*, possibly due to the presence of words *think about* in both arguments. EIDRR makes the correct prediction and provides a consistent explanation, stating that *thinking ahead is based on the situation being discussed*. Similarly, our EIDRR generated a faithful explanation for the *Comparison* relationship in the last instance. Overall, these examples clearly highlight our method’s ability to generate explanations.

Performance of ATSC and NLI

To demonstrate the generality of our method, we evaluate it on aspect term sentiment classification (ATSC) and natural language inference (NLI). Specifically, we use the Restaurant15 dataset (Pontiki et al. 2015) in our experiment, with

the training, validation, and test sets containing 963, 361, and 576 instances, respectively. For the NLI task, we randomly sampled 1,499 training, 400 validation, and 600 test instances from the corpus used in (Chen et al. 2017). The prompt templates used for ATSC and NLI are shown in *Appendix F*. From the results in Table 3, we can draw the following conclusions. **Firstly**, our method shows improvements of approximately 1% on two tasks, in terms of both *Acc* and *F₁*. Specifically, we achieve 1.85% gain in *Acc* and 1.12% in *F₁* score on ATSC. It is noteworthy that these gains are achieved with only limited explanations. **Secondly**, we manually evaluate the generated explanations and obtain average scores exceeding 4.0 on both tasks. Examples of ATSC and NLI explanations are shown in *Appendix G*. Overall, our method performs well on both the ATSC and NLI tasks, delivering higher classification performance and high-quality explanations, confirming its effectiveness and generality.

Conclusion

In this study, we successfully transferred the reasoning capabilities of LLMs to enhance both the performance and interpretability of lightweight IDRR models. Our classification-generation framework can be readily combined with existing classifiers and demonstrates effectiveness not only on the IDRR task but also shows potential for broader applications. Future work includes developing more effective methods to evaluate the faithfulness of generated explanations and further exploring the interpretability of model internals in discourse relation recognition (Miao and Kan 2025; Mondorf and Wold 2025).

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (Nos. 62266017 and 62166018), the Natural Science Foundation of Jiangxi Province of China (Nos. 20232BAB202050 and 20242BAB25117), and Jiangxi Province Key Laboratory of Advanced Network Computing under Grant No. 2024SSY03071.

References

- Bai, H.; and Zhao, H. 2018. Deep Enhanced Representation for Implicit Discourse Relation Recognition. In *Proceedings of COLING*.
- Bilal, A.; Ebert, D.; and Lin, B. 2025. LLMs for Explainable AI: A Comprehensive Survey. ArXiv:2504.00125 [cs].
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In *Proceeding of NIPS*.
- Cai, M.; Yang, Z.; and Jian, P. 2024. Improving Implicit Discourse Relation Recognition with Semantics Confrontation. In *Processings of LREC-COLING*.
- Camburu, O.-M.; Rocktäschel, T.; Lukasiewicz, T.; and Blunsom, P. 2018. e-SNLI: Natural Language Inference with Natural Language Explanations. In *Proceeding of NIPS*.
- Chan, C.; Jiayang, C.; Wang, W.; Jiang, Y.; Fang, T.; Liu, X.; and Song, Y. 2024. Exploring the Potential of ChatGPT on Sentence Level Relations: A Focus on Temporal, Causal, and Discourse Relations. In *Proceedings of EACL Findings*.
- Chan, C.; Liu, X.; Cheng, J.; Li, Z.; Song, Y.; Wong, G.; and See, S. 2023. DiscoPrompt: Path Prediction Prompt Tuning for Implicit Discourse Relation Recognition. In *Proceedings of ACL Findings*.
- Chen, Q.; Zhu, X.; Ling, Z.-H.; Wei, S.; Jiang, H.; and Inkpen, D. 2017. Recurrent Neural Network-Based Sentence Encoder with Gated Attention for Natural Language Inference. In *Proceedings of the 2nd Workshop on Evaluating Vector Space Representations for NLP*.
- Dou, Z.; Hong, Y.; Sun, Y.; and Zhou, G. 2021. CVAE-based Re-anchoring for Implicit Discourse Relation Classification. In *Proceeding of EMNLP Findings*.
- Fan, Y.; Li, P.; and Zhu, Q. 2025. Improving Dialogue Discourse Parsing through Discourse-aware Utterance Clarification. In *Proceedings of ACL*.
- Hu, X.; and Wan, X. 2023. Exploring Discourse Structure in Document-level Machine Translation. In *Proceedings of EMNLP*.
- Jiang, F.; Fan, Y.; Chu, X.; Li, P.; and Zhu, Q. 2021. Not Just Classification: Recognizing Implicit Discourse Relation on Joint Modeling of Classification and Generation. In *Proceedings of EMNLP*.
- Jiang, Y.; Zhang, L.; and Wang, W. 2023. Global and Local Hierarchy-aware Contrastive Framework for Implicit Discourse Relation Recognition. In *Findings of ACL*.
- Kroeger, N.; Ley, D.; Krishna, S.; Agarwal, C.; and Lakkaraju, H. 2024. In-Context Explainers: Harnessing LLMs for Explaining Black Box Models. *arXiv:2310.05797*.
- Kumar, S.; and Talukdar, P. 2020. NILE : Natural Language Inference with Faithful Natural Language Explanations. In *Proceedings of ACL*.
- Li, C.; Yin, Y.; and Carenini, G. 2024. Dialogue Discourse Parsing as Generation: A Sequence-to-Sequence LLM-based Approach. In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, 1–14.
- Li, Z.; Wu, W.; and Li, S. 2020. Composing Elementary Discourse Units in Abstractive Summarization. In *Proceedings of ACL*.
- Liu, H.; Yin, Q.; and Wang, W. Y. 2019. Towards Explainable NLP: A Generative Explanation Framework for Text Classification. In *Proceedings of ACL*.
- Liu, S.; Li, P.; Fan, Y.; and Zhu, Q. 2025. Enhancing Multi-party Dialogue Discourse Parsing with Explanation Generation. In *Proceedings of COLING*, 1531–1544.
- Liu, W.; and Strube, M. 2023. Annotation-Inspired Implicit Discourse Relation Classification with Auxiliary Discourse Connective Generation. In *Proceedings of ACL*.
- Liu, X.; Ou, J.; Song, Y.; and Jiang, X. 2020. On the Importance of Word and Sentence Representation Learning in Implicit Discourse Relation Classification. In *Proceeding of IJCAI*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv:1907.11692*.
- Long, W.; and Webber, B. 2022. Facilitating Contrastive Learning of Discourse Relational Senses by Exploiting the Hierarchy of Sense Relations. In *Proceedings of EMNLP*.
- Ludan, J. M.; Meng, Y.; Nguyen, T.; Shah, S.; Lyu, Q.; Apidianaki, M.; and Callison-Burch, C. 2023. Explanation-based Finetuning Makes Models More Robust to Spurious Cues. In *Proceedings of ACL*.
- Lyu, Q.; Apidianaki, M.; and Callison-Burch, C. 2024. Towards Faithful Model Explanation in NLP: A Survey. *Computational Linguistics*.
- Madsen, A.; Chandar, S.; and Reddy, S. 2024. Are self-explanations from Large Language Models faithful? In *Proceedings of ACL Findings*.
- Majumder, B. P.; Camburu, O.; Lukasiewicz, T.; and Mcauley, J. 2022. Knowledge-Grounded Self-Rationalization via Extractive and Natural Language Explanations. In *International Conference on Machine Learning*, 14786–14801. PMLR.
- Miao, Y.; and Kan, M.-Y. 2025. Discursive Circuits: How Do Language Models Understand Discourse Relations? In *Proceedings of EMNLP*.

- Miao, Y.; Liu, H.; Lei, W.; Chen, N.; and Kan, M.-Y. 2024. Discursive socratic questioning: Evaluating the faithfulness of language models' understanding of discourse relations. In *Proceedings of ACL*, 6277–6295.
- Mondorf, P.; and Wold, S. 2025. Circuit compositions: Exploring modular structures in transformer-based language models. In *Proceedings of ACL*.
- Omura, K.; Cheng, F.; and Kurohashi, S. 2024. An Empirical Study of Synthetic Data Generation for Implicit Discourse Relation Recognition. In *Proceedings of COLING*.
- Park, J.; and Cardie, C. 2012. Improving Implicit Discourse Relation Recognition Through Feature Set Optimization. In *Proceedings of SIGDIAL*.
- Pontiki, M.; Galanis, D.; Papageorgiou, H.; Manandhar, S.; and Androutsopoulos, I. 2015. SemEval-2015 Task 12: Aspect Based Sentiment Analysis. In *Proceedings of SemEval*.
- Prasad, R.; Dinesh, N.; Lee, A.; Miltsakaki, E.; Robaldo, L.; Joshi, A.; and Webber, B. 2008. The Penn Discourse TreeBank 2.0. In *Proceedings of LREC*.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*.
- Rutherford, A.; and Xue, N. 2014. Discovering Implicit Discourse Relations Through Brown Cluster Pair Representation and Coreference Patterns. In *Proceedings of EACL*.
- Tang, J.; Lin, H.; Liao, M.; Lu, Y.; Han, X.; Sun, L.; Xie, W.; and Xu, J. 2021. From Discourse to Narrative: Knowledge Projection for Event Relation Extraction. In *Proceedings of ACL*.
- Wang, C.; Jian, P.; and Huang, M. 2023. Prompt-based Logical Semantics Enhancement for Implicit Discourse Relation Recognition. In *Proceedings of EMNLP*.
- Wang, Q.; Ding, K.; Liang, B.; Yang, M.; and Xu, R. 2023. Reducing Spurious Correlations in Aspect-based Sentiment Analysis with Explanation from Large Language Models. In *Proceeding of EMNLP Findings*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Ichter, B.; Xia, F.; Chi, E.; Le, Q. V.; and Zhou, D. 2022. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. In *Proceedings of NIPS*.
- Wiegrefe, S.; Marasović, A.; and Smith, N. A. 2021. Measuring Association Between Labels and Free-Text Rationales. In *Proceedings of EMNLP 2021*.
- Wu, C.; Cao, L.; Ge, Y.; Liu, Y.; Zhang, M.; and Su, J. 2022. A Label Dependence-Aware Sequence Generation Model for Multi-Level Implicit Discourse Relation Recognition. In *Proceedings of AAAI*.
- Xiang, W.; Wang, Z.; Dai, L.; and Wang, B. 2022. ConnPrompt: Connective-cloze Prompt Learning for Implicit Discourse Relation Recognition. In *Proceedings of COLING*.
- Yung, F.; Ahmad, M.; Scholman, M.; and Demberg, V. 2024. Prompting Implicit Discourse Relation Annotation. In *Proceedings of the 18th Linguistic Annotation Workshop*.
- Zeng, L.; He, R.; Sun, H.; Xu, J.; Liu, C.; and Wang, B. 2024. Global and Local Hierarchical Prompt Tuning Framework for Multi-level Implicit Discourse Relation Recognition. In *Proceedings of COLING*.
- Zhang, B.; Su, J.; Xiong, D.; Lu, Y.; Duan, H.; and Yao, J. 2015. Shallow Convolutional Neural Network for Implicit Discourse Relation Recognition. In *Proceedings of EMNLP*.
- Zhang, B.; Xiong, D.; Su, J.; and Zhang, M. 2018. Learning Better Discourse Representation for Implicit Discourse Relation Recognition via Attention Networks. *Neurocomputing*, 275: 1241–1249.
- Zhang, Y.; Meng, F.; Li, P.; Jian, P.; and Zhou, J. 2021. Context Tracking Network: Graph-based Context Modeling for Implicit Discourse Relation Recognition. In *Proceedings of NAACL*.
- Zhao, H.; He, R.; Xiao, M.; and Xu, J. 2023. Infusing Hierarchical Guidance into Prompt Tuning: A Parameter-Efficient Framework for Multi-level Implicit Discourse Relation Recognition. In *Proceedings of ACL*.
- Zhou, C.; Liang, Y.; Meng, F.; Zhou, J.; Xu, J.; Wang, H.; Zhang, M.; and Su, J. 2022a. A Multi-task Multi-stage Transitional Training Framework for Neural Chat Translation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(7): 7970–7985.
- Zhou, H.; Lan, M.; Wu, Y.; Chen, Y.; and Ma, M. 2022b. Prompt-based Connective Prediction Method for Fine-grained Implicit Discourse Relation Recognition. In *Proceedings of EMNLP Findings*.