

# SR-LLM: Rethinking the Structured Representation in Large Language Model

Anonymous ACL submission

## Abstract

Structured representations, exemplified by Abstract Meaning Representation (AMR), have long been pivotal in computational linguistics. However, their role remains ambiguous in the Large Language Models (LLMs) era. Initial attempts to integrate structured representation into LLMs via a zero-shot setting yielded inferior performance. We hypothesize that such a decline stems from the structure information being passed into LLMs in a code format unfamiliar to LLMs’ training corpora. Consequently, we propose SR-LLM, an innovative framework with two settings to explore a superior way of integrating structured representation with LLMs from training-free and training-dependent perspectives. The former integrates structural information through natural language descriptions in LLM prompts, whereas its counterpart augments the model’s inference capability through fine-tuning on linguistically described structured representations. Performance improvements were observed in widely downstream datasets, with particularly notable gains of 3.17% and 12.38% in PAWS. To the best of our knowledge, this work represents the pioneering demonstration that leveraging structural representations can substantially enhance LLMs’ inference capability. We hope that our work sheds light and encourages future research to enhance the reasoning and interoperability of LLMs by structure data.

## 1 Introduction

Structured representations (SR), manifested in Abstract Meaning Representation (AMR) (Damonte et al., 2016; Knight et al., 2021; Ramirez, 2024), Parse Syntax Trees (PST) (Sachan et al., 2020), and First-Order Logic (FOL) (Barwise, 1977), have been fundamental to NLP (Manning, 1999; Collobert et al., 2011), serving as

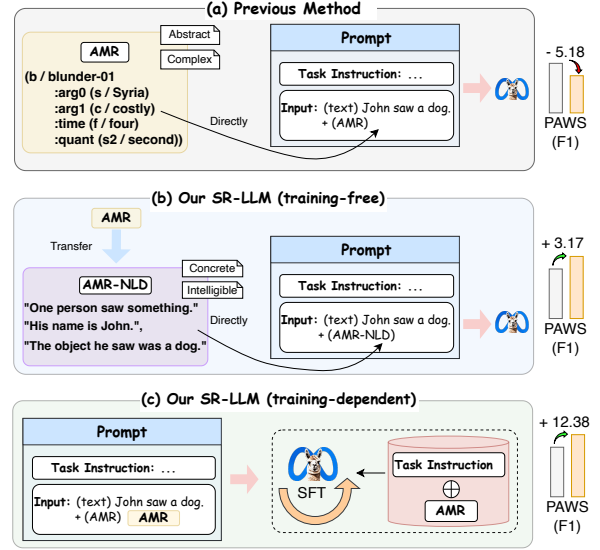


Figure 1: We propose two novel AMR integration approaches: a training-free method using natural language descriptions and a training-dependent fine-tuning paradigm. Evaluation on PAWS shows +3.17% and +12.38% improvements respectively, contrasting with the -5.18% decline in conventional code-format methods.

sophisticated frameworks for capturing semantic relationships and linguistic structures (Banarescu et al., 2013; Wang et al., 2015). An example of AMR, PST, and FOL is depicted in Figure 2.

In the era of LLMs, the paradigm for optimal SR integration remains an open research challenge. Despite LLMs’ capabilities, direct integration of SR into prompts, as illustrated in Figure 1, has proven counterproductive (Jin et al., 2024). We posit that this performance degradation stems from LLMs’ inherent limitations in processing structured representations, where direct exposure to complex linguistic structures impedes rather than enhances their reasoning process.

To address the aforementioned challenges and effectively leverage SR in LLMs, we in-




Figure 2: The AMR, PST, and FOL of the sentence “John saw a dog”.

troduce SR-LLM, a comprehensive framework with dual configurations for structural knowledge integration. The training-free approach transforms SR into natural language descriptions (SR-NLD), enhancing prompt comprehension by reformulating structured information into semantically rich, accessible formats that facilitate nuanced reasoning and reduce ambiguity. Complementarity, the training-dependent paradigm employs supervised fine-tuning on task-specific SR datasets (termed Gen-SR) to establish robust SR-task connections through iterative exposure to structured data, enabling the model to develop sophisticated internal representations and leverage deep structural knowledge during inference across diverse NLP tasks.

Our empirical evaluation encompasses a comprehensive suite of NLP benchmarks, spanning diverse linguistic phenomena from paraphrase detection (Mihalcea et al., 2006; Dolan and Brockett, 2005) and textual entailment recognition (Dagan et al., 2005; Bowman et al., 2015) to machine translation (Bahdanau, 2014; Johnson et al., 2017). This diverse benchmark selection enables rigorous evaluation of our methods across the NLP spectrum. Experimental results demonstrate the superiority of our methods over existing approaches: on PAWS, while conventional method exhibits a 5.18% performance degradation, our training-free and training-dependent approaches achieve +3.17% and +12.38% improvements respectively, which validating the efficacy of our structured information integration paradigm.

Our contributions are as follows:

- We introduce SR-LLM, a novel framework that facilitates SR integration with LLMs through dual paradigms: training-free adaptation and supervised fine-tuning.
- We provide insights into how different

types of SR (AMR, PST, FOL) impact LLMs performance across various tasks.

- To the best of our knowledge, we are the first to show that combining such SR does in fact improve LLM performance, which opens up new avenues for enhanced LLM reasoning and interoperability.

## 2 Problem Definition

This research endeavors to investigate the potential synergies between SR and LLMs, with the ultimate goal of ascertaining how their seamless integration can augment the efficacy and proficiency of LLMs in a wide array of NLP tasks.

Given a natural language input sequence  $X = (x_1, x_2, \dots, x_n)$ , where  $x_i \in V$  represents a token drawn from the vocabulary  $V$ , we also introduce the structured representation  $Z$ .  $Z$  serves as auxiliary information derived from  $X$  and can take various forms, such as AMR, PST, or FOL. These SRs capture semantic, syntactic, or logical information and provide complementary insights to natural language understanding.

The task involves generating an output sequence  $Y = (y_1, y_2, \dots, y_m)$ , where each  $y_i$  belongs to either the target vocabulary or a structured semantic output space. This transformation is performed by a model  $f$ , defined as:

$$Y = f(X, Z) \quad (1)$$

Here,  $f$  specifies how  $X$  and  $Z$  are utilized to complete a specific task by integrating natural language input with its structured representation.

The primary goal of this research is to optimize the definition of  $f$  to achieve the most effective use of  $X$  and  $Z$ , thereby maximizing task performance. Specifically, the objective is to identify the optimal model  $f^*$  that maximizes the evaluation metric  $P(\cdot)$ , such as accuracy or F1 score:

$$f^* = \arg \max_f P(f(X, Z)) \quad (2)$$

## 3 Method

This chapter introduces the SR-LLM framework, a novel paradigm designed to investigate the efficacious integration of SR into

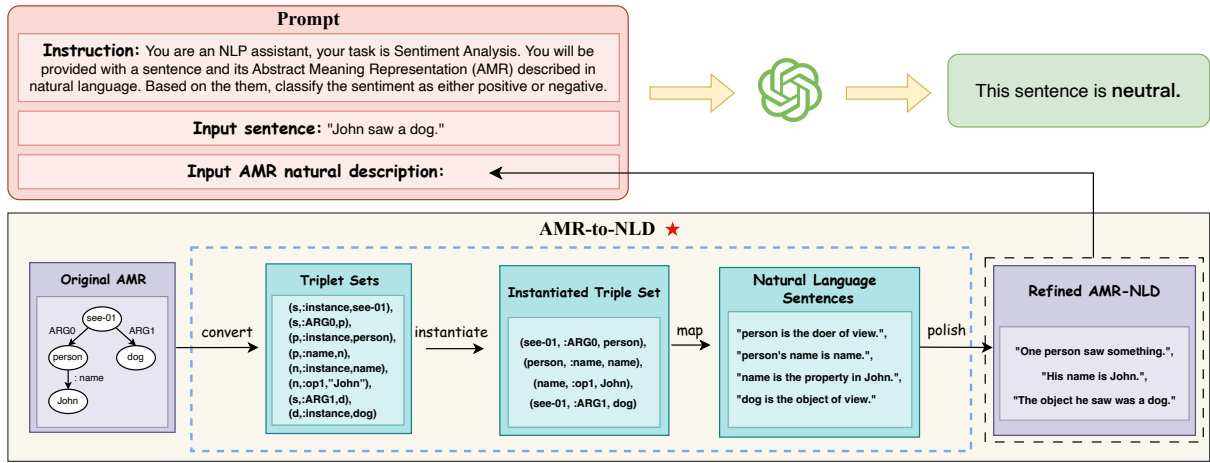


Figure 3: **The whole process of SR-LLM in training-free setting.** Initially, a task-specific prompt consists of an instruction, input sentence, and input SR structure (AMR is used here). Subsequently, the original AMR undergoes transformation via the **AMR-to-NLD** module, which employs predefined rules to map the AMR into an easily interpretable natural language description. This description is then subjected to refinement by a language model, ensuring fluency and coherence, resulting in **AMR-NLD**. Finally, the **AMR-NLD** is seamlessly integrated into the input, which is then fed into the LLM to generate the ultimate response.

LLMs. The SR-LLM framework encompasses two configurations: training-free and training-dependent. These configurations are designed to amalgamate various types of SR through differentiated methodologies, thereby enhancing the LLMs’ capability to comprehend and exploit structured information.

### 3.1 SR-LLM Training-Free

Base Prompt
<b>Instruction:</b> You are a NLP assistant whose purpose is {nlp task}... And you will be given the original text(s) ... <b>Input:</b> {Original Text} <b>Output:</b> XXX
AMRCOT Prompt
<b>Instruction:</b> You are a NLP assistant whose purpose is to {nlp task}... And you will be given the original text(s) and its(their) {Abstract Meaning Representation} to help you to handle this problem <b>Input:</b> {Original Text} and {AMR} <b>Output:</b> XXX

Figure 4: **Base prompt and AMRCOT prompt.** (Top) This is the original task prompt, with only the raw text as input, serving as the standards for performance. (Bottom) This is the AMRCOT prompt method proposed by Jin et al. (2024), serving as a baseline.

Prior approaches, exemplified by AMRCOT (Jin et al., 2024), have attempted to explicitly incorporate AMR into Chain-of-Thought (COT) prompts, as illustrated in Figure 4, have shown that this explicit approach fails to yield

performance enhancement. We hypothesize that one factor contributing to this ineffectiveness stems from the inherent difficulty LLMs face in adequately comprehending and processing abstract structures such as AMR. In view of the aforementioned challenge, as illustrated in Figure 3, we propose SR-LLM Training-Free, where the original structured representation  $Z$  is transformed into natural language descriptions termed **SR-NLD**, where SR can be instantiated with specific structured representations such as AMR, PST, and FOL. We refer to this entire transformation process as **SR-to-NLD** (Structured Representation to Natural Language Description). Specifically, the structured representations are mapped through predefined transformation rules, converting abstract symbols into easily interpretable natural language expressions. These generated natural language descriptions are then refined by a language model to ensure fluency and coherence. Finally, these descriptions are incorporated into the prompt and input into the target LLM. A pivotal advantage of this methodology lies in its training-free nature, as it does not require any additional fine-tuning or retraining of the LLM. Consequently, this technique offers remarkable flexibility, enabling rapid adaption to a diverse array of NLP tasks.

Next, we shall elucidate the SR-to-NLD process, employing AMR-NLD as our quintessential exemplar, which shown in the Algorithm 1.

### Algorithm 1 AMR-to-NLD Transformation

- 1: **Input:** AMR graph  $G = (V, E)$ , nodes collection  $V$ , edges collection  $E$ , Penman library  $\mathcal{P}$ , language model  $\theta$
- 2: **Output:** Refined natural language descriptions  $S_{\text{refined}}$
- 3: **Phase 0:** *Convert AMR to Triplets*
- 4: Convert AMR graph  $G$  into triplets  $T = \{(c_1, r, c_2) \mid c_1, c_2 \in V, r \in E\}$  using the Penman library:  $T = \mathcal{P}(G)$
- 5: **Phase 1:** *Identifier Instantiation*
- 6: **for** each triplet  $(c_1, r, c_2) \in T$  **do**
- 7:     **if**  $r = \text{:instance then}$
- 8:         Replace identifiers  $c_1, c_2$  with their corresponding concepts or instances
- 9:     **end if**
- 10: **end for**
- 11: **Phase 2:** *Mapping to Natural Language*
- 12: Convert triplets into natural language descriptions using a predefined dictionary:  $M : T' \rightarrow S$
- 13: **Phase 3:** *Refinement*
- 14: Refine the generated descriptions  $S$  using language model:  $S_{\text{refined}} = \theta(S)$
- 15: **return**  $S_{\text{refined}}$

The process first converts the AMR graph into triplets, then replaces the identifiers with actual concepts. Next, the triplets are mapped into natural language descriptions using predefined rules, and finally, the descriptions are refined by GPT-4o Mini to produce coherent AMR-NLD. To mitigate the risk of hallucination, we implemented a voting mechanism based on multiple generations. This detailed analysis forms the core of our discussion, outlining each step of the conversion process. The transformation methods for other SRs are elaborated in the Appendix A.1 for completeness. Different from traditional SR-to-Text approaches, which generate a structurally coherent and fluent text based on the SR, such as the “input sentence” in Figure 4. SR-to-NLD aims to collaboratively describe the structured information through multiple sentences, as illustrated by the Refined AMR-NLD in Figure 4.

### 3.2 SR-LLM Training-Dependent

In addition to making SRs more interpretable for LLMs, we also believe that establishing connections between tasks and structured information presents a potential opportunity. As shown in the Figure 5, in SR-LLM Training-Dependent, we constructed a task-specific hybrid dataset, named Gen-SR, where SR can be replaced by specific representations such as AMR, PST, and FOL.

The entire hybrid dataset is composed of two

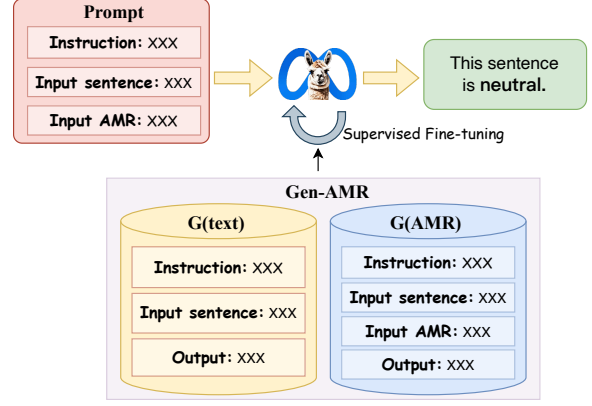


Figure 5: **The whole process of SR-LLM in training-dependent setting.** Taking AMR as an example, a dataset called **Gen-AMR**, created by combining inputs consisting of sentences and their corresponding AMR structures, is utilized for the SFT of LLM to enhance the reasoning capability.

parts: one consists of task-specific instruction pairs based on original text, while the other adds SRs in the instruction pairs based on the former. The former we mark as G(text) and the other we mark as G(SR). The complete example of these two are shown in the Appendix D. This mixed approach allows LLM to not only learn instruction-following for downstream tasks from G(text), but also to establish more robust connections between tasks and structures from G(SR), making the model achieve more effective improvements compared to learning solely from text.

## 4 Experiments

### 4.1 Datasets

To ensure comprehensive and diverse experiments, we selected 10 datasets covering various NLP tasks based on Liu et al. (2024), including five tasks from Jin et al. (2024) for result comparability. The dataset composition includes: PAWS for paraphrase detection (Zhang et al., 2019), SNLI for textual entailment recognition (Bowman et al., 2015), WMT16 for translation tasks (Bojar et al., 2016), CoNLL2003 for named entity recognition (Sang and De Meulder, 2003), Logic for logical fallacy detection (Jin et al., 2022), SST-2 for sentiment analysis (Socher et al., 2013), Pubmed45 for event extraction (Garg et al., 2016), WiC for word sense disambiguation (Pilehvar and Camacho-Collados, 2018), SPIDER for Text2SQL code generation (Yu et al., 2018),



Table 1: **Performance of SR-LLM(training-free).** In the table, a checkmark under “SR” indicates that the original SR was added to the prompt, while a checkmark under “SR-NLD” (highlighted with a gray background) represents the inclusion of SR-NLD in the prompt, which corresponds to the results of SR-LLM (training-free). No checkmarks indicate the use of the original prompt, serving as the control group for comparison. Our focus is on the performance differences between adding SR and SR-NLD, as well as their respective differences compared to the control group.

SR	SR-NLD (Ours)	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	AGNEWS (F1)	WiC (F1)	SNLI (F1)	CoNLL2003 (F1)	SST-2 (F1)	WMT16 (BLEU)	SPIDER (F1)
(a) Llama3.1- 8b-Instruct											
✓		41.59	15.48	24.35	53.88	43.99	25.81	46.28	68.72	13.16	24.80
	✓	36.41	14.20	20.69	48.17	42.05	23.17	41.75	65.66	12.34	21.53
	✓	<b>44.77</b>	<b>18.27</b>	<b>26.10</b>	<b>56.67</b>	<b>48.17</b>	<b>28.87</b>	<b>48.73</b>	<b>71.77</b>	<b>14.10</b>	<b>29.60</b>
(b) GPT 3.5-turbo											
✓		56.94	38.63	27.14	<b>85.12</b>	50.61	38.93	<b>56.52</b>	90.46	26.13	39.63
	✓	56.10	36.27	25.63	81.33	51.60	32.00	54.67	86.90	25.77	39.07
	✓	<b>57.97</b>	<b>39.40</b>	<b>28.17</b>	84.07	<b>55.27</b>	<b>41.47</b>	55.17	<b>92.60</b>	<b>27.07</b>	<b>42.27</b>
(c) GPT 4o-mini											
✓		75.80	<b>48.10</b>	<b>38.65</b>	<b>85.26</b>	<b>58.47</b>	40.59	<b>65.27</b>	91.39	<b>26.80</b>	41.55
	✓	73.50	47.32	33.11	81.62	46.65	41.30	59.21	91.01	26.21	39.33
	✓	<b>76.48</b>	47.95	36.66	83.45	56.63	<b>42.00</b>	64.12	<b>92.83</b>	26.76	<b>43.57</b>

and AGNEWS for text classification (Zhang et al., 2015).

Regarding the source of SR datasets, we used a dual-source strategy: one part includes high-quality AMR datasets from Jin (Jin et al., 2024), covering five core tasks; the other is automatically constructed using GPT-4o, comprising supplementary AMR, PST, and FOL data. The detailed collection processes and results provided in the Appendix B.1.

## 4.2 Training-Free Results

**Experimental Details.** We conducted experiments on the Llama3.1-8b-Instruct (Dubey et al., 2024), GPT-3.5-turbo, and GPT-4o-mini (Achiam et al., 2023) models, arranged from weak to strong according to their performance levels, employing two prompting strategies: Chain-of-Thought (CoT) (Wei et al., 2022) and One-Shot (Brown, 2020). CoT guides step-by-step reasoning, while One-Shot demonstrates task-solving through specific examples. All experiments were conducted independently on three types of SRs: AMR, FOL, and PST. Both PST and FOL were incorporated into the prompts using the same approach as AMR-COT (Jin et al., 2024). For brevity, the results obtained from these experiments were averaged and presented. For detailed prompts, refer to the Appendix E.

**Result Analysis.** First, as shown in Table 1, incorporating SR-NLD into the prompt consistently outperforms incorporating the origi-

nal SR. This indicates that for LLMs, transforming abstract SRs into natural language formats more familiar to the models is an effective strategy for enhancing their ability to interpret and apply structured information. Meanwhile, the comparison of the three models also reveals that the gradual decrease in the benefit of structured information as model performance increases. Specifically, for the Llama3.1-8b-Instruct model, results with SR-NLD significantly and consistently surpass those of the original prompt (i.e., without SR or SR-NLD). For GPT-3.5-turbo, most results show improvement, whereas for GPT-4o-mini, approximately half of the results demonstrate improvement, albeit with a smaller margin. This result further illustrates that weaker models benefit more from structured information as a supplement to the original text, aiding them in downstream reasoning tasks. In contrast, for stronger models, the additional structured information offers limited advantages and may even be less informative than the insights derived directly from the raw text.

## 4.3 Training-Dependent Results

**Experimental Details** We conducted experiments using the Llama3.1-8B-Instruct model to evaluate the performance of the training-dependent setting of SR-LLM, more detailed experimental parameters can be found in the Appendix A.2. The whole process of fine-tuning is a joint training across data from 10 tasks, rather than task-specific fine-tuning for any

Table 2: **Performance of SR-LLM(training-dependent)**. G(text) and G(SR) represent the types of training data, with 50% and 10% indicating their respective proportions in the total training dataset. Our focus is on the best performance of the model across various tasks under different fine-tuning strategies, as well as the performance differences between adding SR and the control group.

FT Strategy	SR	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	AGNEWS (F1)	WiC (F1)	SNLI (F1)	CoNLL2003 (F1)	SST-2 (F1)	WMT16 (BLEU)	SPIDER (EM)
-	✓	41.59 36.41	15.48 14.20	24.35 20.69	53.88 48.17	43.99 42.05	25.81 23.17	46.28 41.75	68.72 65.66	13.16 12.34	24.80 21.53
100% G(text)	✓	68.94 64.07	26.21 16.84	78.91 77.33	76.52 67.14	66.97 67.05	35.53 35.36	75.79 71.73	75.59 74.65	29.07 28.41	41.20 38.47
100% G(SR)	✓	65.34 75.39	25.23 29.89	81.13 <b>82.02</b>	75.10 81.99	66.44 70.82	36.68 <b>56.62</b>	75.40 76.27	77.49 81.62	26.93 <b>30.80</b>	37.07 40.60
50% G(SR) + 50% G(text)	✓	68.66 <b>81.04</b>	26.77 <b>36.52</b>	79.78 81.85	75.77 <b>82.63</b>	69.48 <b>74.68</b>	36.49 54.92	75.42 <b>76.67</b>	77.13 <b>83.72</b>	26.14 30.33	42.40 <b>48.93</b>

single dataset. Detailed data collection procedures and specific training data configurations are provided in the Appendix B.2. To provide a comparative analysis, we conducted three sets of experiments using the following datasets: 100%G (text), 100%G (SR), and a 50%G (text) mixed with 50% G (SR). The 50%-50% ratio was chosen because we considered it to be the most balanced approach. Further experiments, elaborated in Appendix C.2, also confirmed that this is the optimal mixing ratio. And we employed a random sampling approach. All experiments were conducted independently on three types of SRs and for brevity, the results obtained from these experiments were averaged and presented.

**Result Analysis.** As shown in the Table 2, when the fine-tuning dataset includes a certain proportion of SRs and incorporates SRs in the prompt, the model achieves superior performance in downstream tasks, consistently surpassing the case where the training data consists solely of text. Additionally, we observe that models fine-tuned with SRs data perform significantly better with prompts that include SRs, compared to the original prompts without SR. Conversely, when the training data consists entirely of text, the opposite trend is observed. These findings suggest that when a model establishes a strong association between tasks and structured representations during training, it can leverage this information more effectively during inference. Furthermore, when the training data is entirely composed of structured representations, the performance is inferior to that achieved with a balanced mix of text and structured data. This highlights the critical im-

Table 3: **Performance between different AMR Source.** Each data represents the performance difference of the model when using AMRs generated by GPT-4o versus AMRBART, calculated as the performance of AMRBART minus that of GPT-4o. As shown, the differences are almost all below 1%.

AMR	AMR (NLD)	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	WMT16 (BLEU)	SPIDER (EM)
(a) Llama3.1-8b-Instruct						
✓	✓	0.40 0.77	-0.07 -0.13	0.01 0.50	0.13 -0.02	0.28 -0.01
(b) GPT 3.5-turbo						
✓	✓	0.45 0.02	0.57 -2.40	-0.15 0.52	0.08 0.23	0.12 0.21
(c) GPT 4o-mini						
✓	✓	0.08 -0.11	0.07 0.61	0.53 0.61	0.49 -0.13	0.02 -0.13

portance of a balanced integration of raw text and structured representations in maximizing the model’s reasoning capabilities.

#### 4.4 Auxiliary Validation Experiments

##### SR from High-Quality SR-Parsing Model.

To validate the reliability of the generated SRs, we choose AMRBART (Bai et al., 2022) to generate the required AMRs, and experiments were conducted to compare the results with those generated by GPT-4o. It is a model that demonstrates exceptional performance in the AMR parsing domain with a Smatch score of 85.4 on the AMR Parsing Leaderboard, ranking among the top-performing models. As shown in the Table 3, the performance differences between AMRs and AMR-NLDs derived from these two sources were minimal, almost always within 0.5%. This indicates that the quality of the AMRs produced by AMRBART is comparable to those generated by our method.

Table 4: **Performance between different AMR Quality.** The numbers in parentheses represent the performance differences between adding AMR or AMR-NLD and the control group. ‘Flawed’ means the AMR is ambiguous or structurally flawed. ‘Gold’ means the AMR is double checked by human and LLM.

AMR Quality	AMR	AMR-NLD	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	WMT16 (BLEU)	SPIDER (EM)
(a) Llama3.1-8b-Instruct							
-	-	-	42.19	14.32	23.67	13.66	22.58
Flawed	✓	-	34.5 (-7.69)	11.52 (-2.8)	19.41 (-4.26)	11.07 (-2.59)	18.26 (-4.32)
Gold	✓	-	42.48 (+0.29)	14.7 (+0.38)	23.43 (-0.24)	14.65 (+0.99)	22.93 (+0.35)
Flawed	-	✓	32.91 (-9.29)	11.56 (-2.76)	18.39 (-5.28)	11.06 (-2.6)	18.49 (-4.09)
Gold	-	✓	<b>46.96 (+4.76)</b>	<b>18.98 (+4.66)</b>	<b>28.62 (+4.95)</b>	<b>19.13 (+5.47)</b>	<b>28.02 (+5.44)</b>
(b) GPT 3.5-turbo							
-	-	-	56.04	43.79	28.29	26.01	40.28
Flawed	✓	-	51.57 (-4.47)	41.58 (-2.21)	25.71 (-2.58)	23.79 (-2.22)	36.66 (-3.62)
Gold	✓	-	54.53 (-1.51)	44.7 (+0.91)	29.47 (+1.19)	26.17 (+0.15)	39.77 (-0.51)
Flawed	-	✓	51.33 (-4.71)	39.79 (-4.01)	26.9 (-1.38)	24.37 (-1.64)	36.74 (-3.54)
Gold	-	✓	<b>56.78 (+0.74)</b>	<b>46.49 (+2.7)</b>	<b>32.0 (+3.71)</b>	<b>28.72 (+2.71)</b>	<b>44.81 (+4.53)</b>
(c) GPT 4o-mini							
-	-	-	68.71	44.95	37.07	29.02	40.05
Flawed	✓	-	65.63 (-3.08)	42.74 (-2.2)	35.42 (-1.66)	27.31 (-1.71)	37.84 (-2.21)
Gold	✓	-	70.04 (+1.33)	45.9 (+0.96)	35.36 (-1.71)	29.62 (+0.6)	41.47 (+1.42)
Flawed	-	✓	62.63 (-6.07)	41.46 (-3.49)	34.51 (-2.56)	26.64 (-2.38)	37.30 (-2.76)
Gold	-	✓	<b>70.13 (+1.42)</b>	<b>46.18 (+1.23)</b>	<b>39.17 (+2.09)</b>	<b>30.14 (+1.12)</b>	<b>41.54 (+1.48)</b>

**Gold AMR vs Flawed AMR.** Additionally, we selected 70 AMR samples (labeled as “Flawed”) with ambiguities or structural flaws from each of the 10 datasets and refined them using a dual-process correction strategy that combined AMRBART-generated results with manual adjustments, producing high-quality AMRs (labeled “Gold”). Results in Table 4 show that AMR quality significantly impacts model performance. Using flawed AMRs led to performance declines for both direct AMR and AMR-NLD representations, with a more pronounced drop for AMR-NLD. This indirectly validates AMR-NLD’s ability to enhance LLMs’ understanding of AMR structures. In contrast, with high-quality AMRs, AMR-NLD substantially improved model performance, while direct AMR usage showed limited gains. These results demonstrate that combining high-quality AMR-NLD is more effective in helping models comprehend structured information. This effect is particularly pronounced when the quality of the AMR is high, leading to substantial performance gains.

**Fine-tuning Larger Model.** To validate the robustness of the proposed method, we selected Llama3.1-70B-Instruct and conducted training-dependent experiments, whose details were consistent with those described for the Llama3.1-8B-Instruct model above, in five tasks shown in the Table 5. The SR used in these experiments was AMR, with a 50%-50% ratio.

Table 5: **Performance of SR-LLM(training-dependent) in Llama3.1-70b-Instruct.** The numbers in parentheses represent the performance differences between adding SR and the control group. Our focus is on the performance variations across different models with different prompts.

AMR	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	WMT16 (BLEU)	SPIDER (EM)
(a) Vanilla					
✓	68.00	47.13	63.95	28.65	33.71
	60.28	43.08	48.82	27.91	29.20
	(-7.73)	(-4.04)	(-15.13)	(-0.73)	(-4.51)
(b) 50% G(AMR) + 50% G(text)					
✓	74.74	54.57	76.51	33.73	47.06
	<b>84.56</b>	<b>58.96</b>	<b>81.54</b>	<b>37.00</b>	<b>53.84</b>
	(+9.81)	(+4.39)	(+5.03)	(+3.27)	(+6.78)

We can see that, after fine-tuning, the model demonstrated improvements on all tasks, with corresponding values turning positive, more than half of which exceeded 5%. These results further validate the effectiveness of Training-Dependent method on larger-scale models.

**Experiment on Traceable LLM.** Since NLP datasets have been public for years, their role in modern LLM development is unclear. We validate our approach through experiments on OLMo (Groeneveld et al., 2024), where the training and validation data sources are clearly documented, in five tasks shown in the Table 6. The results show that when the prompt includes SR, its performance is lower than when SR is not included. However, when SR is transformed into SR-NLD, the performance improves signif-

Table 6: **Performance in OLMo.**

SR	SR_NLD (Ours)	PAWS (F1)	Logic (F1)	Pubmed45 (F1)	WMT16 (BLEU)	SPIDER (EM)
		61.52	22.96	60.27	9.67	20.44
✓		57.73	17.97	57.89	11.23	19.79
	✓	<b>65.40</b>	<b>24.89</b>	<b>65.19</b>	<b>12.44</b>	<b>22.20</b>

icantly. For instance, in the PAWS, the performance increases from 61.52% to 65.40%. This demonstrates the robustness and generalizability of our approach.

## 5 Related Work

**Structure Representations.** The SRs, including AMR, PST, and FOL, each unique advantages and applications in specific areas. AMR uses rooted, labeled graphs to abstract syntactic details, offering concise and semantically rich representations (Banarescu et al., 2013). PST, based on Chomsky’s generative grammar, employs hierarchical trees to represent sentence syntax and word dependencies (Chomsky, 2014). FOL, as a symbolic logic system, defines objects, their relations, and properties, serving as a key tool in formal logic and reasoning (Enderton, 2001; Barwise, 1977).

**Structure Representations Transformation.** The SR transformation has long been a critical area of research. Much of the existing work has focused on SR-to-Text approaches, which generate fluent text that aligns with the structure of the SR (Song et al., 2018; Ribeiro et al., 2021; Wang et al., 2020). Meanwhile, a method known as canonical expressions employs rule-based techniques to convert structures into standardized natural language representations, primarily to resolve ambiguities in non-standard sentences (Shin et al., 2021; Roy et al., 2024). Its outputs are essentially normalized texts rather than comprehensive descriptions of the SR’s full structure. In contrast, our SR-to-NLD approach preserves the integrity of structured information while enhancing its interpretability through natural language descriptions of the structure.

**Structured Representations used for NLP in LLM.** With the rise of LLM, studies like Hahn et al. (2022) showed these sequence to sequence model’s ability to generalize across formal domains, though challenges like low in-

terpretability and hallucinations persist De Belis (2023). Integrating structured representations into LLMs has improved accuracy and interpretability. Yao et al. (2024) and (Shi et al., 2024) combined AMR with LLMs for tasks like sentence simplification and Retrieval-Augmented Generation. Additionally, Hahn et al. (2022) and (Kalyanpur et al., 2024) advanced formal specification and logical reasoning in LLMs. And An et al. (2024) identified "magic prompts" that improve the performance of NLP tasks by solely focusing on semantic parsing, without the need to provide the actual parsing results. However, Jin et al. (2024) argued that simply add AMR into prompt might sometimes hinder performance in certain NLP tasks.

## 6 Conclusion

SR-LLM demonstrates significant progress in enhancing LLMs’ reasoning capabilities through structured representations. Our evaluation across diverse NLP tasks revealed SR’s potential in generating novel implicit information. We established a framework for integrating SR into LLMs, from prompt engineering to fine-tuning, providing valuable insights into structured information incorporation. These advancements led to substantial improvements in both training-free and training-dependent settings, highlighting the effectiveness of integrating semantic, syntactic, and logical features. As we refine SR-LLM, we anticipate further progress towards more interpretable, accurate, and versatile language models with enhanced reasoning capabilities in various applications.

## 7 limitations

Despite SR-NLD’s promising performance in certain tasks, its effectiveness remains inconsistent across different LLMs. The rule-based conversion method may constrain flexibility. Future research should focus on developing a more robust and adaptive structured representation, exploring task-specific optimizations, and investigating advanced conversion techniques and novel model architectures. Expanding evaluation to diverse language models and datasets will be crucial to enhance the method’s consistency, flexibility, and applicability in various NLP domains.



## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Kaikai An, Shuzheng Si, Helan Hu, Haozhe Zhao, Yuchi Wang, Qingyan Guo, and Baobao Chang. 2024. Rethinking semantic parsing for large language models: Enhancing llm performance with semantic hints. *arXiv preprint arXiv:2409.14469*.
- Dzmitry Bahdanau. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Xuefeng Bai, Yulong Chen, and Yue Zhang. 2022. Graph pre-training for amr parsing and generation. *arXiv preprint arXiv:2203.07836*.
- Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186.
- Jon Barwise. 1977. An introduction to first-order logic. In *Studies in Logic and the Foundations of Mathematics*, volume 90, pages 5–46. Elsevier.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. Findings of the 2016 conference on machine translation (wmt16). In *First conference on machine translation*, pages 131–198. Association for Computational Linguistics.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Tom B Brown. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Noam Chomsky. 2014. *Aspects of the Theory of Syntax*. 11. MIT press.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *Journal of machine learning research*, 12:2493–2537.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine learning challenges workshop*, pages 177–190. Springer.
- Marco Damonte, Shay B Cohen, and Giorgio Satta. 2016. An incremental parser for abstract meaning representation. *arXiv preprint arXiv:1608.06111*.
- Alessandro De Bellis. 2023. Structuring the unstructured: an llm-guided transition. In *DC@ISWC*.
- Bill Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Third international workshop on paraphrasing (IWP2005)*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Herbert B Enderton. 2001. *A mathematical introduction to logic*. Elsevier.
- Sahil Garg, Aram Galstyan, Ulf Hermjakob, and Daniel Marcu. 2016. Extracting biomolecular interactions using semantic parsing of biomedical text. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, A. Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Daniel Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hanna Hajishirzi. 2024. *Olmo: Accelerating the science of language models*. *arXiv preprint*.
- Christopher Hahn, Frederik Schmitt, Julia J Tillman, Niklas Metzger, Julian Siber, and Bernd Finkbeiner. 2022. Formal specifications from natural language. *arXiv preprint arXiv:2206.01962*.
- Zhijing Jin, Yuen Chen, Fernando Gonzalez, Jiarui Liu, Jiayi Zhang, Julian Michael, Bernhard Schölkopf, and Mona Diab. 2024. Analyzing the role of semantic representations in the era of large language models. *arXiv preprint arXiv:2405.01502*.
- Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea, and Bernhard Schölkopf. 2022. Logical fallacy detection. *arXiv preprint arXiv:2202.13758*.

631	Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. <i>Transactions of the Association for Computational Linguistics</i> , 5:339–351.	686
632		687
633		688
634		689
635		
636		
637		
638	Aditya Kalyanpur, Kailash Saravanakumar, Victor Barres, Jennifer Chu-Carroll, David Melville, and David Ferrucci. 2024. Llm-arc: Enhancing llms with an automated reasoning critic. <i>arXiv preprint arXiv:2406.17663</i> .	690
639		691
640		692
641		693
642		694
643		695
644	Kevin Knight, Bianca Badarau, Laura Baranescu, Claire Bonial, Madalina Bardocz, Kira Griffitt, Ulf Hermjakob, Daniel Marcu, Martha Palmer, Tim O’Gorman, et al. 2021. Abstract meaning representation (amr) annotation release 3.0.	696
645		697
646		698
647		699
648	Yang Liu, Jiahuan Cao, Chongyu Liu, Kai Ding, and Lianwen Jin. 2024. Datasets for large language models: A comprehensive survey. <i>arXiv preprint arXiv:2402.18041</i> .	700
649		701
650		702
651		
652	Christopher D Manning. 1999. <i>Foundations of statistical natural language processing</i> . The MIT Press.	703
653		704
654		705
655	Rada Mihalcea, Courtney Corley, Carlo Strapparava, et al. 2006. Corpus-based and knowledge-based measures of text semantic similarity. In <i>Aaai</i> , volume 6, pages 775–780.	706
656		
657		
658		
659	Mohammad Taher Pilehvar and Jose Camacho-Collados. 2018. Wic: the word-in-context dataset for evaluating context-sensitive meaning representations. <i>arXiv preprint arXiv:1808.09121</i> .	707
660		708
661		709
662		710
663	José Gabriel Carrasco Ramírez. 2024. Natural language processing advancements: Breaking barriers in human-computer interaction. <i>Journal of Artificial Intelligence General Science (JAIGS) ISSN: 3006-4023</i> , 3(1):31–39.	711
664		712
665		
666		
667		
668	Leonardo FR Ribeiro, Yue Zhang, and Iryna Gurevych. 2021. Structural adapters in pre-trained language models for amr-to-text generation. <i>arXiv preprint arXiv:2103.09120</i> .	713
669		714
670		715
671		716
672	Subhro Roy, Samuel Thomson, Tongfei Chen, Richard Shin, Adam Pauls, Jason Eisner, and Benjamin Van Durme. 2024. Benchclamp: A benchmark for evaluating language models on syntactic and semantic parsing. <i>Advances in Neural Information Processing Systems</i> , 36.	717
673		718
674		719
675		720
676		721
677		722
678	Devendra Singh Sachan, Yuhao Zhang, Peng Qi, and William Hamilton. 2020. Do syntax trees help pre-trained transformers extract information? <i>arXiv preprint arXiv:2008.09084</i> .	723
679		724
680		725
681		726
682	Erik F Sang and Fien De Meulder. 2003. Introduction to the conll-2003 shared task: Language-independent named entity recognition. <i>arXiv preprint cs/0306050</i> .	727
683		728
684		729
685		730
		731
		732
		733
		734
		735
		736
		737
		738
		739

## A Experimental Details

### A.1 Details of Converting SR to SR-NLD

#### A.1.1 Details of Translating AMR Triplet to Natural Sentence

According to the Figure 6, first, the triplet is converted into a sentence based on the relation mapping rules. Then, using the entity dictionary, the entities are replaced with their actual meanings to form the final sentence. Finally, the sentence is input into the LLM for refinement into a complete and coherent sentence, as shown in the Figure 7.

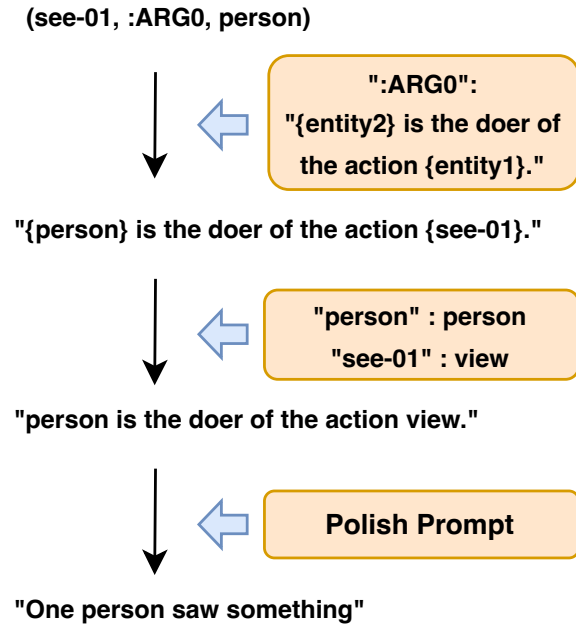


Figure 6: The process of translate entities and relationships into natural language sentences

#### A.1.2 Whole Process of Making PST-NLD

**Definition of PST.** PST is represented as a tree structure  $T = (N, E)$ . Here  $N$  denotes the set of nodes, representing the syntactic components of a sentence (e.g., part-of-speech tags and phrase labels). Node types include  $S$  (sentence),  $NP$  (noun phrase),  $VP$  (verb phrase), etc.  $E$  denotes the set of edges, representing dependencies between components. An example of the original PST structure is shown in the Figure 8.

**Conversion of PST to a Linear Structure Using Depth-First Search (DFS).** Starting from the root node (typically  $n_0$ , representing the sentence’s syntactic structure, such as

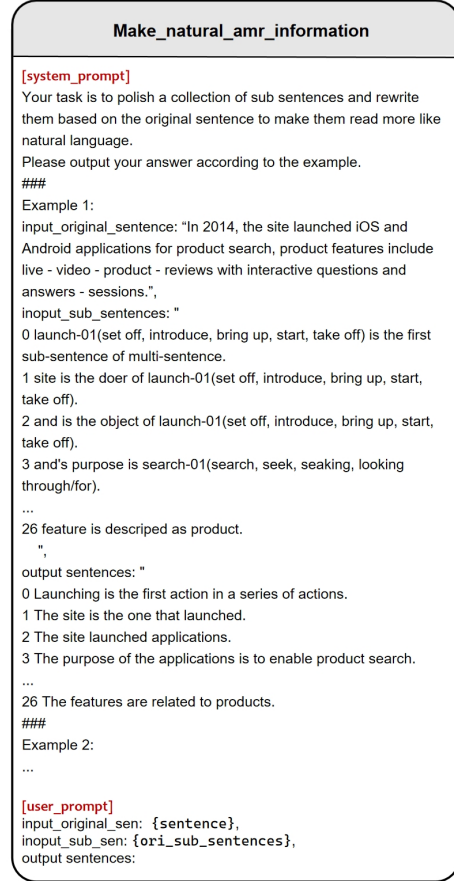


Figure 7: The prompt of polishing sentence for making AMR-NLD

$S$ ), we traverse the tree in a depth-first search (DFS) manner, converting it into a linear sequence of symbols  $P$ .

**Mapping PST Identifiers to Natural Language Descriptions.** We define a mapping function  $M$  to translate each identifier (e.g.,  $S$ ,  $NP$ ,  $VBD$ ) and its child nodes into natural language descriptions. The dictionary  $D$ , which specifies the natural language interpretation of each identifier, is detailed in the appendix. For each triplet  $(n, c_1, c_2)$ , where  $n$  is a node and  $c_1, c_2$  are its children, we apply the mapping function  $M(n) = \text{description}(n)$ . The resulting natural language description  $S$  is as shown in the Figure 8.

**Refinement of Natural Language Descriptions Using a Language Model.** To make the descriptions more natural and coherent, the generated descriptions  $S$  are refined using the language model  $F_{LM} : S \rightarrow S_{\text{refined}}$ . The specific prompt is shown in the prompt (b) of Figure 9.

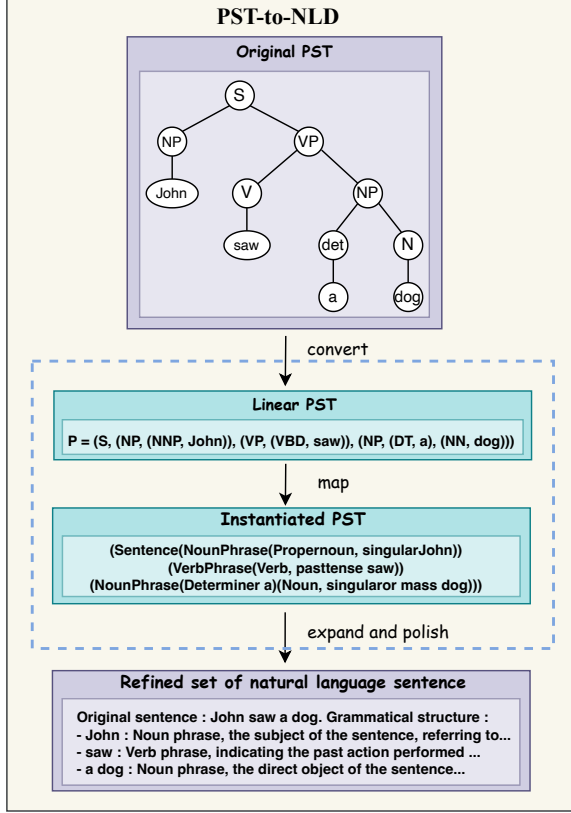


Figure 8: **The Whole process of Making PST-NLD.** The process of creating PST-NLD involves first converting the PST tree structure into a linear sequence of symbols using depth-first search (DFS). Then, a mapping function is applied to translate each node and its children into natural language descriptions. Finally, a language model is used to refine the generated descriptions, making them more natural and coherent.

### A.1.3 Whole Process of Making FOL-NLD

**Definition of FOL.** FOL is represented as  $F = (Q, V, P, C)$ , where  $Q$  denotes the set of quantifiers, used to express the existence of variables, such as  $\exists$  (exists) and  $\forall$  (for all).  $V$  represents the set of variables, representing objects in FOL, typically denoted as  $x, y, z$ .  $P$  represents the set of predicates, used to express properties of objects or relationships between multiple objects.  $C$  represents the set of logical connectives, used to connect multiple propositions, including conjunction ( $\wedge$ ), disjunction ( $\vee$ ), and negation ( $\neg$ ). An example of the original FOL structure is shown in the Figure 10.

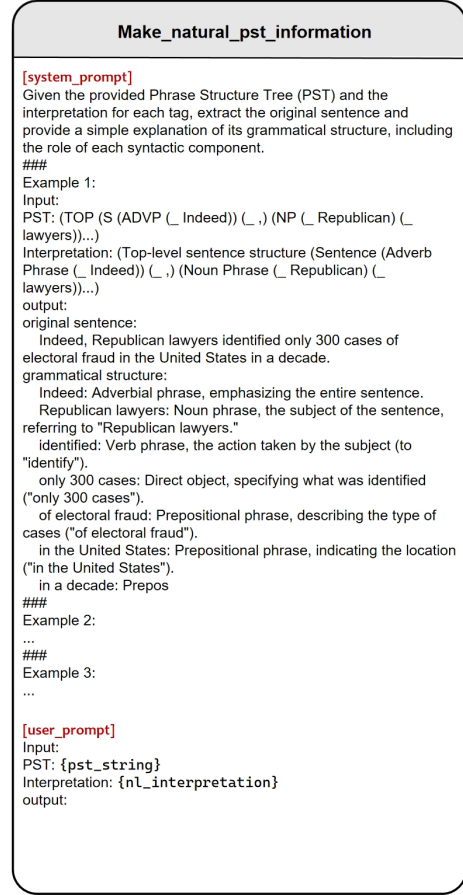


Figure 9: The prompt of polishing sentence for making PST-NLD

**Mapping FOL to Natural Language Descriptions.** We define a mapping function  $M = (D, L)$ , where  $D$  is a set of symbol mappings that translates variables, predicates, and logical operators in FOL into natural language descriptions.  $L$  is a set of logical mapping rules that transforms the logical structure of FOL into natural language syntax. By applying these mapping rules to the initial FOL expressions, we can convert logical symbols into natural language descriptions.

**Refinement of Natural Language Descriptions Using a Language Model.** To ensure that the descriptions are coherent and fluent, we refine the generated descriptions  $S$  using the language model  $F_{LM} : S \rightarrow S_{refined}$ . The specific prompt is shown in the prompt (c) of Figure 11.

## A.2 Complete Fine-tuning Details

We used Meta’s Llama-3.1-8B-Instruct as the backbone and conducted fine-tuning on 8 NVIDIA A100-80G GPUs. Optimization was



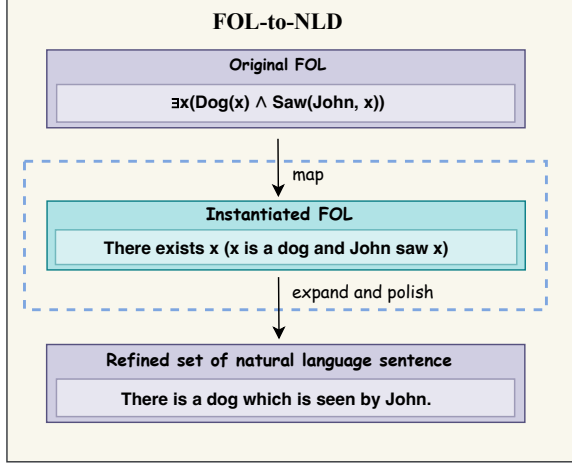


Figure 10: **The Whole process of Making FOL-NLD.** The process of converting FOL to NLD involves first mapping FOL symbols, such as variables, predicates, and logical operators, into natural language using predefined symbol mappings and logical rules. Then, the generated descriptions are refined using a language model to ensure they are coherent and fluent.

performed using the AdamW optimizer with a learning rate of  $1e-4$  and cosine learning rate decay. The training setup included a `per_device_train_batch_size` of 16 and `gradient_accumulation_steps` of 8, yielding an effective global batch size of 1024. A fixed random seed of 42 ensured reproducibility. Each dataset was fine-tuned for 10 epochs, with early stopping to prevent overfitting.

## B Data Collection

### B.1 The Process of Constructing Datasets for All Tasks of SR-LLM (training-free)

In this section, I will outline the process of collecting test data for the 10 tasks used in SR-LLM (training-free), including both the original text and three types of structured representations. The data statistics are summarized in the Table 7.

**SNLI** SNLI is a large and comprehensive dataset, with a test set containing 10,000 examples. Therefore, we directly used the test set for our experiments. The AMR, FOL, and PST data were generated using GPT-4o-turbo in a few-shots setting, with the prompt provided in

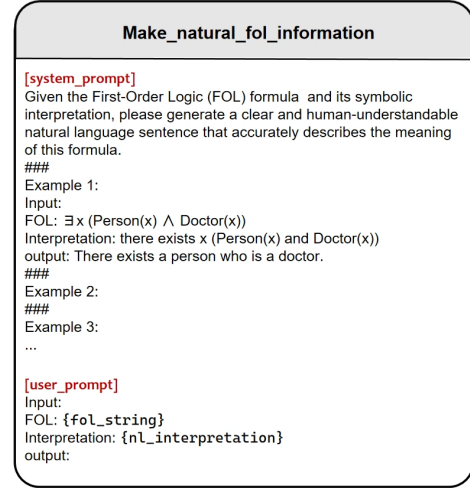


Figure 11: The prompt of polishing sentence for making FOL-NLD

Table 7: Tasks and datasets used in SR-LLM (training-free)

Dataset	Task	Test Size
PAWS	Paraphrase Detection	8000
SNLI	Recognizing Textual Entailment	10000
WMT16	Translation	5999
CoNLL2003	Named Entity Recognition	3453
LOGIC	Logical Fallacy Detection	2449
SST-2	Sentiment Analysis	872
Pubmed45	Event Extraction	5000
WiC	Lexical Disambiguation	2038
SPIDER	Text2SQL Code Generation	8034
AGNEWS	Text Classification	7600

the Figure 12, Figure 13 and Figure 14.

**CoNLL2003** CoNLL2003 is also a rich and complete dataset, with a test set of 3,453 examples, which we used directly. Structured representations were generated using the same method as described above.

**SST-2** Since the official SST-2 test set does not contain labels, we used the full validation set of 872 examples as the test set for this experiment. Structured representations were generated using the same method as described above.

**WiC** The WiC test set consists of 1,400 examples, which is relatively small. Therefore, we combined the 648 examples from the validation set to create a larger test set. Structured representations were generated using the same method as described above.

**AGNEWS** AGNEWS is another large and comprehensive dataset, with a test set of 7,600 examples, which we used directly. Structured

**Make\_AMR**

**[system\_prompt]**

You are an NLP assistant whose task is to generate an Abstract Meaning Representation (AMR) for a given sentence.

The AMR should capture the essential meaning of the sentence in a structured format, representing the relationships between concepts in the sentence.

You must accurately represent the roles of entities, actions, and other relationships in the sentence.

Ensure that the representation is semantically complete and concise.

The AMR should not contain any extra explanations or interpretations, only the AMR structure.

###

Example 1:

Input Sentence: "The boy wants to eat an apple."

Output AMR:

```
(w / want-01
 :arg0 (b / boy)
 :arg1 (e / eat-01
 :arg0 b
 :arg1 (a / apple)))
```

###

Example 2:

...

###

Example 3:

...

**[user\_prompt]**

Input Sentence: {sentence}

Output AMR:

Figure 12: The prompt of making AMR

**Make\_PST**

**[system\_prompt]**

You are an NLP assistant tasked with generating a Phrase Structure Tree (PST) for a given sentence. Your PST should illustrate the syntactic structure of the sentence, showing how words and phrases are hierarchically organized. Follow these guidelines:

**Structure:** Represent the sentence's hierarchical structure with nodes for major constituents, such as noun phrases (NP), verb phrases (VP), adjective phrases (AdjP), adverbial phrases (AdvP), and prepositional phrases (PP). The tree should have a root node representing the entire sentence, typically labeled as S (Sentence).

**Labeling:** Label each node according to its syntactic function, such as NP for noun phrases, VP for verb phrases, and PP for prepositional phrases. Ensure that terminals (leaf nodes) represent individual words or word forms.

**Hierarchy:** Show the hierarchical relationship between constituents. For example, a verb phrase (VP) might be composed of a verb and a noun phrase (NP), where the VP node branches into its subcomponents. Include intermediate nodes where applicable, such as Determiners (Det) within noun phrases.

Ensure that the PST is detailed enough to accurately reflect the syntactic structure of the sentence. Avoid including extra explanations or interpretations; focus solely on the syntactic structure.

The pst should not contain any extra explanations or interpretations, only the pst structure.

###

Example 1:

Input Sentence: "The cat slept."

Output pst: '(S (NP (Det The) (N cat)) (VP (V slept)))'

###

Example 2:

...

###

Example 3:

...

**[user\_prompt]**

Input:

FOL: {fol\_string}

Interpretation: {nl\_interpretation}

output:

Figure 13: The prompt of making PST

Table 8: Tasks and datasets used in SR-LLM (training-dependent)

Dataset	Task	Train Size	Test Size
PAWS	Paraphrase Detection	10000	8000
SNLI	Recognizing Textual Entailment	10000	10000
WMT16	Translation	10000	5999
CoNLL2003	Named Entity Recognition	10000	3453
LOGIC	Logical Fallacy Detection	10000	2449
SST-2	Sentiment Analysis	10000	872
Pubmed45	Event Extraction	10000	5000
WiC	Lexical Disambiguation	5066	1048
SPIDER	Text2SQL Code Generation	7000	1034
AGNEWS	Text Classification	10000	7600

representations were generated using the same method as described above.

**PAWS** To ensure sufficient comparability in the experiments, the original text data and AMR representations for PAWS were sourced from Jin et al. (2024). And the FOL and PST representations were generated using the same method as described above.

**WMT16, LOGIC, Pubmed45, SPIDER** The data collection for these tasks followed the same procedure as PAWS.

## B.2 The Process of Constructing Datasets for All Tasks of SR-LLM (training-dependent)

In this section, I will explain the process of collecting both training and test data for the 10 tasks used in SR-LLM (training-dependent), including the original text and three types of structured representations. Data statistics are summarized in the Table 8.

**PAWS, WMT16, Pubmed45, SNLI, CoNLL2003, SST-2, AGNEWS** These datasets contain relatively large training sets. Therefore, we randomly selected 10,000 examples from each as the training set. The structured representations were generated using GPT-4o-turbo in a few-shot setting, with sample prompts provided in the figure. The test

sets are the same as those used in the SR-LLM (training-free) experiments.

**LOGIC** Since the LOGIC dataset is relatively small, the training-free setup used all the available samples from the test, validation, and training sets combined, yielding a total of 2,449 samples as the test set. We retained these 2,449 samples for the test set in the training-dependent setting as well. To create the training set, we synthetically generated 10,000 logic examples using GPT-4o-turbo. The generation process is illustrated in the Figure 15, where a few-shot strategy was employed to guide the model to generate sentences containing different logical fallacies. The generated prompt is

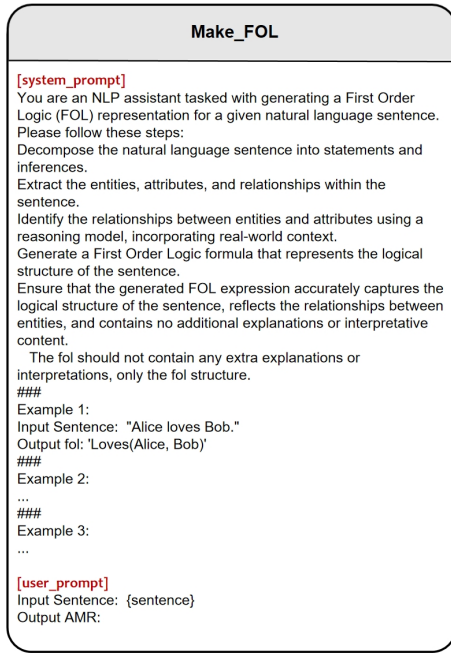



Figure 14: The prompt of making FOL

shown in Figures 16 and Figures 17. The type of logical error serves as the label, producing complete data points. Structured representations were generated in the same manner as described above.



**Figure 15: The synthetic process for LOGIC data.** Taking the “Faulty Generalization” type as an example, we employed a few-shot strategy to guide the model in generating sentences containing the logical fallacy of “Faulty Generalization” To ensure greater sentence diversity, we incorporated a thematic element during generation, such as “Sports” as shown in the figure. This thematic addition helps produce a broader variety of sentence while maintaining the specific logical error, leading to a richer and more varied dataset.

**SPIDER** Since the official SPIDER test set is not publicly available, the training-free setup used a combination of training and validation sets as the test set. However, due to the complexity of generating SPIDER-like data, we used the original 7,000 training examples for the training set in the training-dependent set-

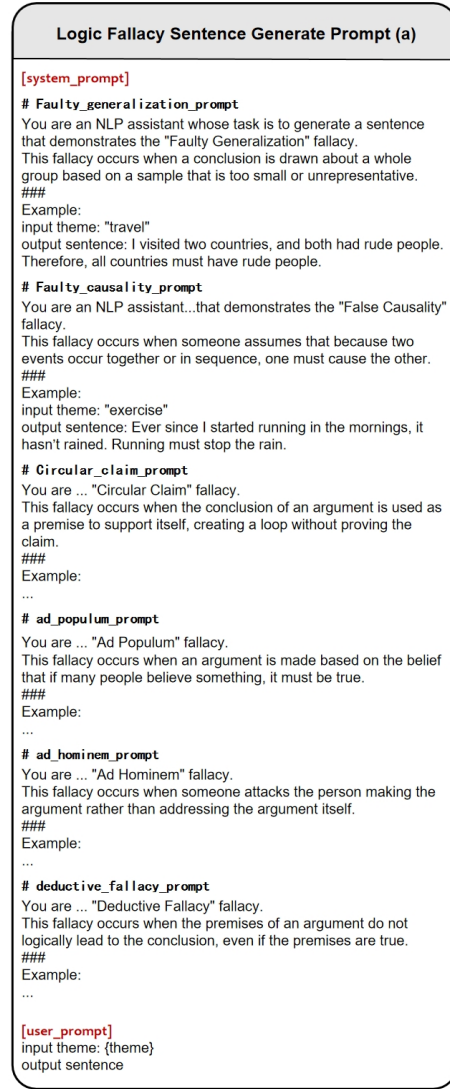


Figure 16: Logic Folly Sentence Generate Prompt (a)

ting and the 1,034 validation examples as the test set. Structured representations were generated as described above.

**WiC** As the WiC training set is relatively small, we combined the 648 validation examples with the original training set to create a total of 5,066 training samples. Structured representations were generated using the same method as described above.

## C Additional Experiments

### C.1 Comparative Analysis of Different SR Combinations and Their Impact on LLM Reasoning

We conducted an in-depth comparison of the performance of different structured representations (SR) and explored their combinations to assess whether joint usage could further en-

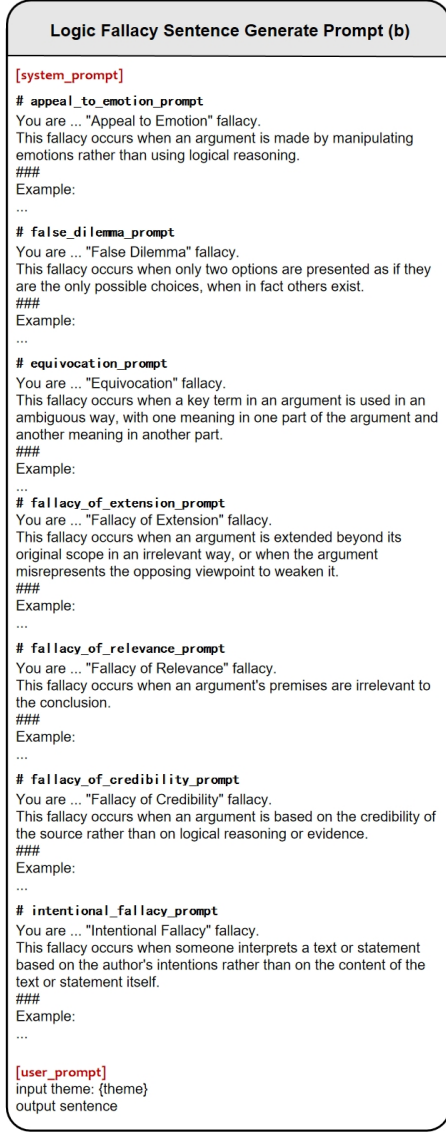




Figure 17: Logic Fallacy Generate Prompt (b)

hance LLM reasoning capabilities. Figure 18 summarizes the average performance improvements across all tasks. The results indicate that the use of individual SRs such as AMR, PST, and FOL did not lead to significant performance enhancements, which is consistent with the findings of (Jin et al., 2024). Moreover, when multiple SRs were introduced simultaneously, their combined complexity posed additional challenges for the LLMs, further dispersing the model’s attention and resulting in poorer performance compared to using a single SR. In contrast, when relatively weaker LLMs were provided with more comprehensible semantic features (AMR) and logical features (FOL), their average performance improved. The integration of these two types of features com-

plemented each other, leading to better overall results. However, the contribution of syntactic features (PST) was relatively less effective and, in some cases, even negated the positive effects of semantic and logical features.



(a) Average performance of SR Performance in different Tasks



(b) Average performance of SR-NLD Performance in different Tasks

Figure 18: **Performance comparison of different SR combinations.** (a) The average performance enhancement ( $\Delta$ ), for various SR combinations across different tasks. (b) The average performance enhancement ( $\Delta$ ), for different SR-NLD combinations across various tasks.

## C.2 Optimal Text-to-SR Ratio Analysis

To further investigate the most optimal ratio of between G(text) and G(SR), I selected five tasks, which includes PAWS, LOGIC, Pubmed45, SPIDER, WMT16 for additional experiments, adjusting the ratio of text to structured representations in the Gen-SR dataset to identify the optimal balance. The experimental results are shown in the Figure 19. As can be observed, the fluctuations in performance with different ratios are relatively small. For both AMR and PST, a 50-50 ratio between text and structured representations appears to be the most effective. However, for FOL, a 30-70 ratio (whether favoring structured representations) yields better results. This is a preliminary exploration, and I believe it repre-



sents a promising direction for further research.




Figure 19: Comparison of average performance of models at different scales in all tasks.

### C.3 Enhancing LLM’s Understanding of SR during Pretraining.

We further conducted experiments during the pretraining phase with the goal of enhancing LLM’s ability to comprehend structured representations, aiming for performance improvements in downstream tasks. Specifically, we collected 1GB of task-agnostic SR data, including AMR, PST, and FOL, following a similar procedure as in previous data collection efforts, and applied this data to the pre-training of the Llama3.1-8B-Instruct model. Building on this, we further conducted SFT, the same as SR-LLM (training-dependent), on five datasets. The final average performance results are shown in the Table 9.

The experimental results show that, compared to the vanilla model without pre-training, the pre-trained model indeed exhibited performance improvements in downstream tasks, though the improvements were relatively modest, with an average increase of less than 1%. However, after applying SFT on the pre-trained model, its performance was actually inferior to that of the vanilla model trained directly with SFT. We hypothesize that this phenomenon may be due to the model forming certain inherent understandings of structured representations during the pre-training phase, which hindered its ability to establish effective connections between structure and tasks during SFT, leading to worse performance compared to the vanilla model. This phenomenon highlights a potential conflict in how the model processes structured information during the pre-training and fine-tuning phases, which warrants further exploration and resolution in future research.

Table 9: **The SR enhancement of models with different training strategies.** These are the average SR Enhancement results across all tasks under different training strategies. Green indicates the best performance within the same SR, while red represents the worst performance.

AMR	FOL	PST	Pretrain	SFT	$\Delta$
✓					<b>-3.51%</b>
✓			✓		0.56%
✓			✓	✓	-1.16%
✓				✓	<b>11.59%</b>
	✓				<b>-2.83%</b>
	✓		✓		1.30%
	✓		✓	✓	3.10%
	✓			✓	<b>6.45%</b>
		✓			<b>-3.61%</b>
		✓	✓		-0.18%
		✓	✓	✓	1.58%
		✓		✓	<b>2.91%</b>

## D Examples of Gen-SR

We present specific examples of Gen-SR in this section. Figure 20 shows an example of G(text), Figure 21 shows an example of G(AMR), Figure 22 shows an example of G(PST), and Figure 23 shows an example of G(FOL).

**G(text)**

**[Instruction]**  
 You are an advanced natural language understanding model designed for the task of Sentiment Analysis.  
 Your goal is to classify the sentiment of a given sentence.  
 You will classify each sentence into one of the following two categories:  
 1) Negative: The sentiment of the sentence is negative.  
 2) Positive: The sentiment of the sentence is positive.  
 You will be provided with a sentence. Based on the content of the sentence, classify the sentiment as either positive or negative.  
 Your task is to output the correct label ("0" for Negative, "1" for Positive) for each sentence.  
 Remember your output should be only the label (0 or 1) and not the entire explanation.

**[Input]**  
 Input Sentence: 'It's *nothing interesting in unfaithful whatsoever*'

**[Output]**  
 Output Response: '0'

Figure 20: The Example of G(text)

## E Prompt of Testing the SR-LLM

We present the complete prompts for our experiments, including both CoT and One-shot

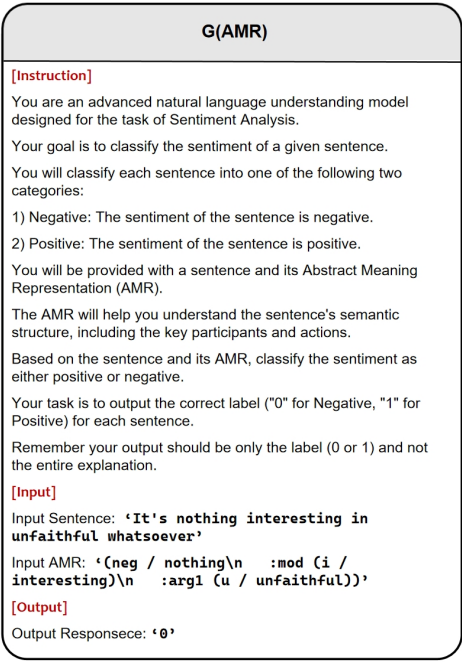


Figure 21: The Example of G(AMR)

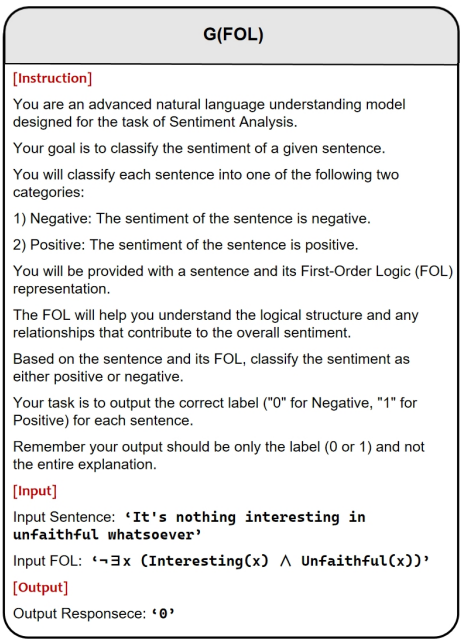


Figure 23: The Example of G(FOL)

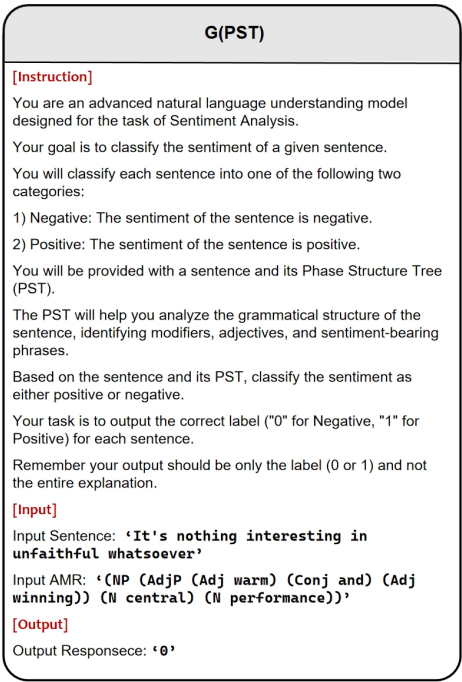


Figure 22: The Example of G(PST)

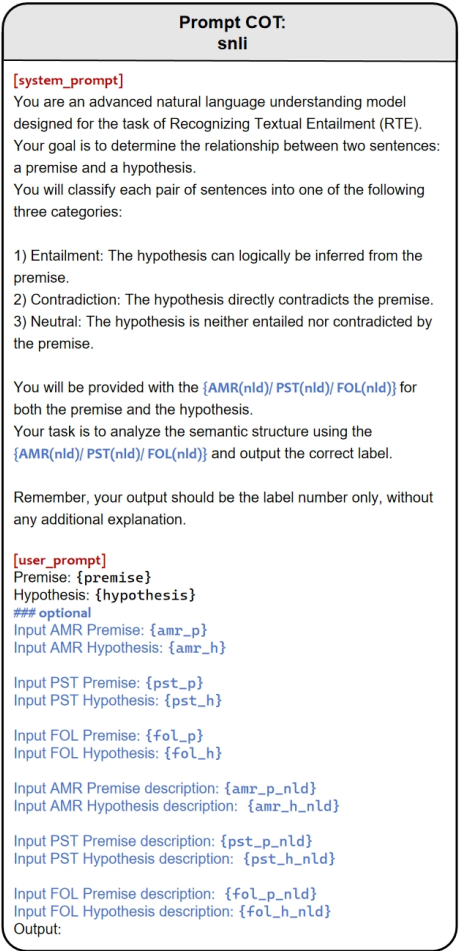


Figure 24: The COT prompt of SNLI

examples, using the SNLI dataset as an illustration in Figures 24, Figures 25 and Figures 26.

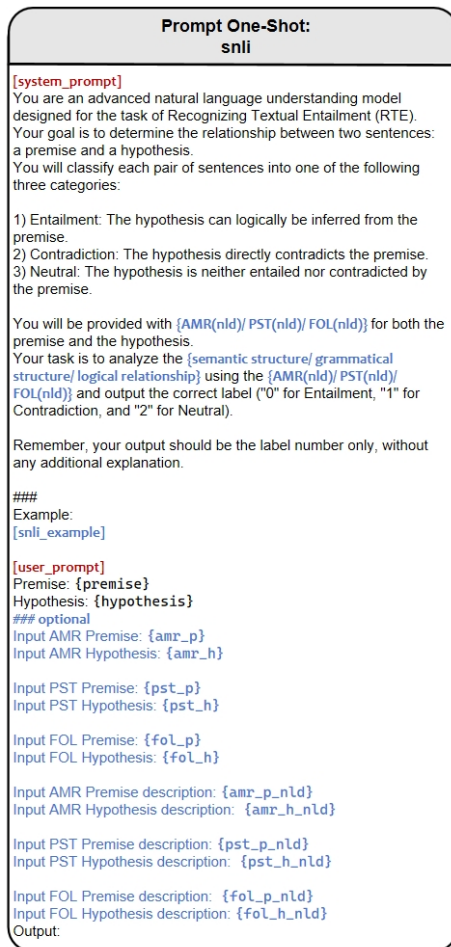


Figure 25: The One-Shot prompt of SNLI

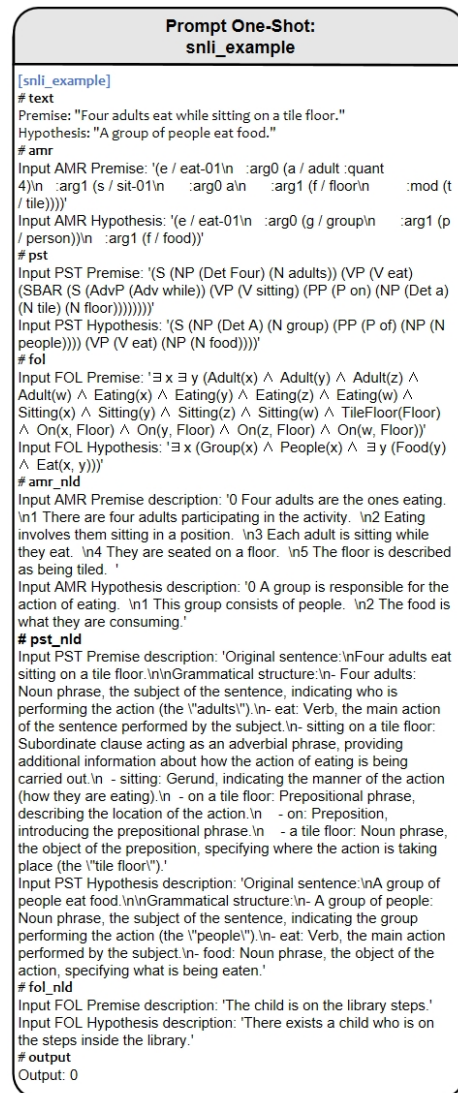


Figure 26: The One-Shot prompt of SNLI's example