

# Uncovering the Fragility of Trustworthy LLMs through Chinese Textual Ambiguity

Xinwei Wu\*

Hongyu Liu\*

Haojie Li\*

Chalmers University of Technology  
Gothenburg, Sweden

Xinyu Ji

Boeing

Gothenburg, Sweden

Ruohan Li

Chalmers University of Technology

Gothenburg, Sweden

Yule Chen

Chalmers University of Technology  
Gothenburg, Sweden

Yigeng Zhang<sup>†</sup>

Salesforce

Burlington, USA

## Abstract

In this work, we study a critical research problem regarding the trustworthiness of large language models (LLMs): how LLMs behave when encountering ambiguous narrative text, with a particular focus on Chinese textual ambiguity. We created a benchmark dataset by collecting and generating ambiguous sentences with context and their corresponding disambiguated pairs, representing multiple possible interpretations. These annotated examples are systematically categorized into 3 main categories and 9 subcategories. Through experiments, we discovered significant fragility in LLMs when handling ambiguity, revealing behavior that differs substantially from humans. Specifically, LLMs cannot reliably distinguish ambiguous text from unambiguous text, show overconfidence in interpreting ambiguous text as having a single meaning rather than multiple meanings, and exhibit overthinking when attempting to understand the various possible meanings. Our findings highlight a fundamental limitation in current LLMs that has significant implications for their deployment in real-world applications where linguistic ambiguity is common, calling for improved approaches to handle uncertainty in language understanding. The dataset and code are publicly available at this GitHub repository<sup>1</sup>.

## CCS Concepts

• **Computing methodologies** → **Language resources**.

## Keywords

Large Language Models, AI Trustworthiness, Ambiguity Detection

## ACM Reference Format:

Xinwei Wu, Hongyu Liu, Haojie Li, Xinyu Ji, Ruohan Li, Yule Chen, and Yigeng Zhang. 2025. Uncovering the Fragility of Trustworthy LLMs through Chinese Textual Ambiguity. In *Proceedings of KDD workshop on Evaluation and Trustworthiness of Agentic and Generative AI Models (Agentic & GenAI Evaluation Workshop KDD'25)*. ACM, New York, NY, USA, 11 pages. <https://doi.org/XXXXXXX.XXXXXXX>

## 1 Introduction

Large Language Models (LLMs) have demonstrated strong language understanding capabilities and are widely deployed across a range of real-world applications [21, 29, 36]. They are used to process complex instructions in multi-turn dialogues and are integrated into various systems as agents or components of AI workflows. However, LLMs still exhibit inherent limitations in trustworthiness, such as hallucinations [2, 15], misunderstanding [22], and misalignment [9] that are particularly critical in safety-sensitive scenarios. Researchers have also invested significant effort in improving alignment, developing guardrails [1], and enhancing uncertainty understanding [13, 32] to enable more reliable use of LLMs.

In practical use cases, people typically interact with LLMs via chat interfaces using written text in a conversational or spoken style, where ambiguity frequently arises. For example, in an LLM-based e-commerce shopper agent, the instruction *return the phone and computer accessories I purchased last month* is ambiguous: it could mean returning the phone and the computer's accessories, or the accessories for both the phone and the computer. In such cases, the agent should be able to use appropriate means to resolve the ambiguity instead of proceeding with one possible interpretation, which may lead to unintended outcomes.

In this work, we focus specifically on examining how LLMs behave when faced with linguistic ambiguity, a core aspect of human language understanding. We present a new benchmark for ambiguity detection and interpretation in Chinese text. The dataset was annotated by native Chinese speakers and includes ambiguity type classification. It contains 900 ambiguous sentences sourced from real-world contexts spanning a variety of everyday scenarios. Each ambiguous sentence is annotated with all plausible interpretations and a corresponding set of disambiguated sentences, where each rewritten sentence clearly reflects one of the possible meanings.

\*These authors contributed equally to this research.

<sup>†</sup>Corresponding author. Email: yzhang168@uh.edu

<sup>1</sup><https://github.com/ictup/LLM-Chinese-Textual-Disambiguation>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*Agentic & GenAI Evaluation Workshop KDD'25, Toronto, ON, Canada*

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2025/08

<https://doi.org/XXXXXXX.XXXXXXX>

We categorize ambiguity into three types: lexical, syntactic, and semantic-pragmatic.

We conducted extensive experiments to investigate how LLMs handle ambiguity and found that they often exhibit fragile behavior in such scenarios. Our initial observation is that LLMs tend to confidently commit to one possible interpretation of an ambiguous sentence, which diverges from how humans typically respond to ambiguity. Furthermore, when explicitly asked to disambiguate, the models often assert with overconfidence that the sentence is ambiguous, even when it may not be. In some cases, the models demonstrate signs of *overthinking* when prompted to explain ambiguous content, producing unnecessarily complex or speculative explanations.

Our analysis spans a range of open-weight LLMs, including both standard and reasoning models, from small to large scales. We performed a series of experiments involving prompt engineering and retrieval-augmented generation (RAG) across different model families and sizes. The results show that even state-of-the-art open-weight models such as DeepSeek-R1 display fragile behavior when confronted with ambiguity. Our contributions lie in several dimensions:

- The study sheds light on the semantic boundaries of LLMs, demonstrating that disambiguation remains a major challenge.
- This study provides a meaningful new perspective for evaluating the trustworthiness of LLMs and related systems.
- In the discourse of NLP research, we present and open-source a new benchmark for ambiguity detection and understanding. Meanwhile, we conduct extensive experiments and analyses to investigate how LLMs behave when faced with ambiguous sentences. Furthermore, we propose a solution to improve the robustness of LLMs in such scenarios.

This work serves as a call for the community to pay closer attention to the fragility of LLMs in the face of ambiguity, and a message of caution for industry applications concerned with the trustworthiness of LLM-based AI systems to help prevent potentially catastrophic consequences.

## 2 Chinese textual ambiguity benchmark for LLMs

### 2.1 Task introduction

Ambiguity is ubiquitous and inevitable in human language, yet large language models (LLMs) rely on natural language instructions to interface with users. Given this, understanding how LLMs behave with ambiguity is essential. In this work, we focus on two core tasks: ambiguity detection and ambiguity interpretation. In the context of NLP, the first task evaluates whether an LLM can identify if a sentence is ambiguous, formulated as a binary classification problem. The second examines whether an LLM can capture latent ambiguity and generate all plausible interpretations, framed as a conditional generation task. To support this study, we introduce a new human-annotated benchmark for ambiguity detection and interpretation in Chinese. While we acknowledge that human language lacks precision and annotations may not represent absolute ground truth, our goal is to analyze the behavior of LLMs in the

face of ambiguity, highlight discrepancies with human judgments, and offer a new perspective on LLM evaluation. Through the lens of these tasks, we aim to investigate the following research questions:

- **RQ1:** To what extent does an LLM differ from human annotators in identifying ambiguous narratives?
- **RQ2:** How does an LLM perform when explaining the meaning of a sentence that contains ambiguity?
- **RQ3:** How does an LLM interpret the meaning of a sentence when it is explicitly informed that the sentence is ambiguous?

### 2.2 Dataset creation

In this work, we employ human annotators to construct ambiguous sentences along with their corresponding disambiguated versions. Annotators are also asked to provide all plausible interpretations of each ambiguous sentence. The sentence construction is grounded in real-world scenarios and everyday contexts, and the data are sourced through original writing, commonly used spoken expressions, online searches, and AI-assisted generation. The quality of the annotations is assessed by the annotators. All annotators are native Chinese speakers with qualifications sufficient for admission to graduate-level programs in science and engineering. We include only sentences that remain highly ambiguous and cannot be easily disambiguated based on human annotators' judgments.

After sentence collection and annotation, we further categorize ambiguity into three main levels: lexical, syntactic, and semantic-pragmatic, following established studies in Chinese linguistics. Within the lexical category, we further distinguish homonymy, polysemy, and part-of-speech ambiguity. Within syntactic ambiguity, we include both structural and syntax-semantic ambiguity. For semantic-pragmatic ambiguity, we identify four subtypes: speech act ambiguity, conversational implicature, deixis ambiguity, and sociocultural ambiguity. The ambiguity categories and label statistics are presented in Table 1 with examples.

## 3 Experiment and result

To evaluate the performance of different models on our three designed experimental tasks, we selected eight representative large language models, covering a range of scales and architectural characteristics:

**Qwen3 Series Models:** This includes Qwen3-4B, Qwen3-14B, Qwen3-32B, and Qwen3-235B-A22B, corresponding to 4B, 14B, 32B, and 235B parameters, respectively. Among them, Qwen3-235B-A22B is specifically optimized for reasoning-intensive tasks and shows outstanding performance in complex reasoning scenarios [41].

**Gemma2 Series Models:** Developed by Google, these instruction-tuned models include Gemma2-2B-it, Gemma2-9B-it, and Gemma2-27B-it, with 2B, 9B, and 27B parameters, respectively. These models excel in instruction following and dialogue tasks [35].

**DeepSeek-R1 Model:** A large-scale model with 671B parameters, deeply optimized for reasoning tasks, showing strong capabilities

**Table 1: Categorization of Chinese ambiguity into lexical, syntactic, and semantic-pragmatic levels, each with multiple English interpretations based on contextual usage. Each category and sub-category is accompanied by its corresponding statistics.**

Category	Sub-category	Ambiguity Example in Chinese	English Translation
Lexical (218)	Polysemy (152)	吃完饭，他冷冷地说：“这顿饭先记着，回头我们再算账。”	“Let’s keep this in mind and settle the bill later.” / “Let’s keep this in mind —I’ll get even with you later.” ( “算账” can mean settling payment or seeking revenge)
	Homonymy (27)	小明走在公园里，赞叹枝头上的杜鹃很漂亮。	Xiao Ming walked in the park and admired how beautiful the cuckoo was on the branch. / ... admired how beautiful the azalea was on the branch. ( “杜鹃” can mean a bird or a flower)
	Part-of-Speech (39)	民警来到现场勘察，发现这个门没有锁。	The police arrived and found that the door didn’t have a lock. / ... found that the door hadn’t been locked. ( “锁” as noun vs. verb)
Syntactic (327)	Structural Ambiguity (261)	接到紧急通知后，领导简单地宣布：我们需要组织人员。	Upon receiving the emergency notice, the leader briefly announced: “We need to organize the personnel.” / “We need the staff responsible for organization.” (“组织人员” can be verb-object or compound noun)
	Syntax-Semantics (66)	女儿在日记中写道：我恨她对我的刻薄不容忍。	The daughter wrote in her diary: “I hate that she cannot tolerate my harshness.” / “I hate her harshness and intolerance toward me.” (ambiguity in scope of negation and attribution)
Semantic-Pragmatic (357)	Speech Act (101)	病房里护士问病人：“你能把窗户关上吗？”	“Can you close the window?” (literal inquiry about capability) / “Please close the window.” (polite indirect request)
	Conversational Implicature (82)	聚餐时有人提议喝酒，小王说：“你们可真懂我。”	At dinner, someone suggested drinking. Xiao Wang said, “You really understand me.” (literal agreement) / “You really don’t understand me at all.” (ironic/sarcastic implication)
	Deixis Ambiguity (51)	医务室王医生突然插话说：其实开刀的是他父亲。	Dr. Wang suddenly interjected: “Actually, it was his father who had the surgery.” / “Actually, his father was the one who performed the surgery.” ( “开刀” can mean to undergo or to perform surgery)
	Sociocultural Ambiguity (123)	相亲时介绍人表示，对方孩子特别老实。	During a blind date, the matchmaker said, “Their child is very well-behaved.” (praise) / “Their child is overly obedient and lacks personality.” (implied criticism — “老实” has dual connotations in social contexts)

in mathematical reasoning and logical analysis [12].

To ensure reproducibility and comparability, we split the dataset into training, development, and test sets in a 70/15/15 ratio, using stratified sampling based on ambiguity subcategories to ensure consistent distribution of ambiguity types across subsets. All input texts were pre-processed for standardization, including removing extra spaces and unifying punctuation formats, to ensure input consistency. Model outputs were also post-processed, including formatting, answer extraction, and consistency checks. All tasks used the same data split strategy to guarantee the comparability of experimental results.

In order to investigate the models’ ability to identify and understand potential ambiguity, we design two experimental conditions.

The first is the non-explicit prompt condition (**Direct Interpretation**), where the prompt does not explicitly indicate that the input sentence may be ambiguous. The second is the explicit ambiguity prompt condition (**Prompted Disambiguation**), where the prompt explicitly states that the input sentence contains ambiguity.

### 3.1 Experimental Tasks

Based on the aforementioned research questions and task formalization, we adopted a structured experimental design and completed three experimental tasks, systematically addressing the three core issues in Chinese ambiguity processing: ambiguity detection, ambiguity understanding, and end-to-end detection and understanding. For evaluation, we used accuracy, precision, recall, and F1

score as the main metrics. Given the imbalanced distribution of ambiguous sentences in real-world corpora, we placed particular emphasis on F1 score and recall. We constructed a multi-dimensional, multi-level evaluation framework to comprehensively reflect the performance of different methods on Chinese ambiguity processing tasks.

**3.1.1 Ambiguity Detection Task.** The core goal of the ambiguity detection task is to perform binary classification on a given Chinese sentence, i.e., to determine whether the sentence contains ambiguity. In this task, the provided sentences may or may not be ambiguous, and the model needs to make its own judgment and respond with 'yes' or 'no'. This task forms the foundation of the entire ambiguity processing pipeline. The evaluation is based on standard binary classification metrics, including accuracy, precision, recall, and F1 score.

**3.1.2 Ambiguity Understanding Task.** The ambiguity understanding task is a further extension based on ambiguity detection, requiring the model to, given a Chinese sentence (with or without ambiguity), complete three sub-tasks: ambiguity source localization, multiple interpretation generation, and disambiguated sentence generation. Specifically, the model needs to identify words, phrases, or syntactic structures that may cause ambiguity and mark their positions; then, based on these sources, generate all reasonable and semantically coherent interpretations (if ambiguity exists, at least two different interpretations should be provided); finally, for each interpretation, generate a corresponding disambiguated sentence by adding context, replacing words, or adjusting structure to eliminate ambiguity. To comprehensively evaluate model capability, we designed two experimental conditions: (1) directly prompting the model to explain possible meanings without indicating whether ambiguity exists, to assess the model's overall detection and understanding ability; (2) explicitly indicating that the sentence contains ambiguity, to focus on the model's understanding and generation ability. This task places higher demands on the model's linguistic analysis, understanding, and generation capabilities. Evaluation uses exact match (EM), recall, and set F1 to assess the quality of generated interpretations, effectively reflecting model performance in multi-interpretation generation.

**3.1.3 End-to-End Detection and Understanding Task.** The end-to-end task represents the highest level of ambiguity processing. Given a raw sentence, the model must first perform ambiguity detection and ambiguity type recognition, then combine the detection results with other prompting strategies (such as chain-of-thought, RAG, etc.) to form composite prompts, guiding the large model to complete ambiguity understanding and disambiguation, and automatically output multiple interpretations and disambiguated sentences. The detection results at each stage serve as prompt conditions for subsequent reasoning, with all steps integrated into a single pipeline, achieving fully automated processing from raw input to final output without human intervention. This setup not only closely simulates real-world application scenarios, but also greatly increases task complexity. Evaluation uses joint metrics, comprehensively considering detection accuracy and understanding quality, providing a quantitative assessment of overall task performance.

We clarify that all accuracy-related metrics are used solely to measure alignment with human annotations, rather than to define any absolute ground truth. We do not claim an objective standard for determining whether a sentence is ambiguous, as natural language is inherently ambiguous and human interpretations can vary significantly. This limitation should be acknowledged when interpreting the results.

## 3.2 Detection Methods

The detection methods include both transformer-based text classifiers and large language model prompting. By observing the changes in model performance under different prompting strategies, we analyze how the design of prompts affects the model's ability to handle ambiguity.

**Table 2: Ambiguity detection performance across different LLMs. Bolded scores represent the best performance, and underlined scores indicate the second-best results. <sup>†</sup> These models are optimized for reasoning tasks and have reasoning explicitly enabled.**

Model	Params	Accuracy	Precision	Recall	Macro-F1
BERT-ft	109M	<b>94.70</b>	<b>94.16</b>	<b>89.58</b>	<b>91.81</b>
Gemma2	2B	46.06	49.60	49.57	45.77
	9B	38.19	52.24	51.18	35.85
	27B	43.20	50.53	50.50	43.14
Qwen3	4B	63.96	53.32	51.93	50.24
	14B	58.95	54.63	54.91	54.65
	32B	55.85	56.66	57.58	54.79
Qwen3 <sup>†</sup>	235B-A22B	43.68	54.97	53.91	43.08
DeepSeek-R1 <sup>†</sup>	671B-A37B	<u>65.63</u>	<u>62.41</u>	<u>63.48</u>	<u>62.62</u>

**3.2.1 Pretrained Transformer-based Text Classifier.** Transformer-based pre-trained language models, such as BERT [8], RoBERTa [23], and XLNet [42], have demonstrated strong performance across a wide range of text classification tasks. Given their effectiveness in sentence-level modeling [37, 46, 47] and passage-level discrimination [5, 16] in applied classification scenarios, we adopt Transformer-based classifiers as our foundation. As a baseline, we used the pre-trained language model hf1/chinese-roberta-wwm-ext [7] as the classifier backbone. This model is based on the RoBERTa architecture and has been specifically optimized for Chinese, achieving strong performance in various NLP tasks. We further fine-tune the model with a binary classification objective to distinguish between ambiguous and unambiguous sentences.

We added a classification head to the model and fine-tuned it for binary classification. Regarding feature engineering, in addition to textual input, we incorporated linguistic features such as sentence length, POS tag sequences, and syntactic tree depth to enhance the model's sensitivity to Chinese ambiguity. These features were fused with the main model output via additional embedding layers.

To systematically evaluate the performance of LLMs on ambiguity detection, we designed a series of experiments to investigate the impact of model scale and prompting strategy on task performance.

**3.2.2 Large Language Model Prompt Learning.** Given the strong performance of large language models in complex reasoning tasks,

**Table 3: Macro-F1 performance on ambiguity detection using different prompting strategies. Bolded scores represent the best performing model under each method, while underlined scores indicate the best performing method for each model. <sup>†</sup> Reasoning-enabled models.**

Model	Params	Direct Prompt	Few-shot	Knowledge	CoT	CoT + FS	RAG + FS
Gemma2	2B	45.77	39.58	40.40	34.50	<u>51.95</u>	46.69
	9B	35.85	32.09	38.75	37.32	41.74	<u>52.95</u>
	27B	43.14	40.75	42.65	36.32	44.61	<u>56.12</u>
Qwen3	4B	50.24	43.02	50.95	46.86	47.71	<u>58.05</u>
	14B	54.65	53.81	52.11	42.24	52.51	<u>60.83</u>
	32B	54.79	55.23	55.00	44.20	55.72	<u>69.57</u>
Qwen3 <sup>†</sup>	235B-A22B	43.08	55.46	57.68	53.25	59.35	<u>74.41</u>
DeepSeek-R1 <sup>†</sup>	671B-A37B	<b>62.62</b>	<b>63.94</b>	<b>62.63</b>	<b>55.20</b>	<b>65.16</b>	<b>87.01</b>

**Table 4: Performance on ambiguity meaning understanding task. Models are evaluated in two settings: Direct Interpretation (without disambiguation prompt) and Prompted Disambiguation (with explicit disambiguation prompt). Metrics include Set F1, Recall, and Exact Match (EM).  $\Delta$  Set F1 /  $\Delta$  Recall shows the improvement from prompting. <sup>†</sup> Reasoning-enabled models.**

Model	Params	Direct Interpretation			Prompted Disambiguation			Difference $\Delta$ Set F1 / $\Delta$ Recall
		EM	Recall	Set F1	EM	Recall	Set F1	
Gemma2	2B	0.00	26.65	40.49	0.00	27.21	41.18	0.69 / 0.55
	9B	0.00	30.33	44.71	0.00	29.78	43.92	-0.78 / -0.55
	27B	0.00	31.62	46.37	0.00	31.07	45.69	-0.69 / -0.55
Qwen3	4B	0.00	31.99	46.86	0.00	32.17	47.16	0.29 / 0.18
	14B	0.00	33.64	48.87	0.00	31.62	46.27	-2.59 / -2.02
	32B	0.00	32.17	47.16	0.00	31.07	45.88	-1.27 / -1.10
Qwen3 <sup>†</sup>	235B-A22B	0.00	36.40	51.67	0.00	37.32	52.65	0.98 / 0.92
DeepSeek-R1 <sup>†</sup>	671B-A37B	0.00	<b>39.71</b>	<b>55.49</b>	0.00	<b>40.26</b>	<b>56.18</b>	0.69 / 0.55

we designed six different prompting strategies to systematically evaluate their effectiveness in ambiguity detection:

**(1) Direct Prompting:** In the most basic method, the model receives the input sentence and directly answers *yes* or *no* to indicate whether there is ambiguity. The prompt template is concise and avoids introducing bias. For example: “Please determine whether the following sentence contains ambiguity. Just answer ‘yes’ or ‘no’: [sentence]”

**(2) Few-shot Prompting:** To leverage in-context learning, we include three carefully selected examples in the prompt that cover both ambiguous and unambiguous sentences. These examples represent different types of ambiguity, helping the model understand the task requirements. Selection follows the principles of representativeness and diversity.

**(3) Knowledge-enhanced Prompting:** We incorporate linguistic background knowledge about Chinese ambiguity into the prompt, including definitions and characteristics of lexical, syntactic, and semantic ambiguity. This approach aims to enhance the model’s theoretical understanding and improve detection accuracy and consistency.

**(4) Chain-of-Thought Prompting:** Inspired by chain-of-thought reasoning, we require the model to perform step-by-step analysis before making a final judgment. The model first analyzes the sentence structure, then identifies possible ambiguity points, and finally provides reasoning and a conclusion, improving interpretability.

**(5) Chain-of-Thought and Few-shot Combined Prompting:**


This method combines the advantages of chain-of-thought reasoning and few-shot learning, providing examples with detailed analytical processes and requiring the model to follow similar reasoning patterns for new sentences.

**(6) RAG and Few-shot Combined Prompting:** Our approach employs a RAG and few-shot prompting strategy that pre-retrieves relevant examples to construct prompt templates for guiding model reasoning. This strategy aims to address two key issues in model understanding through the guidance of semantically similar examples, thereby improving model comprehension quality: first, the tendency to select a single possible interpretation rather than all reasonable ones; second, the problem of over-interpretation and false reasoning when sufficient context is lacking.

Table 2 demonstrates that a fine-tuned BERT model can reliably distinguish ambiguous sentences from unambiguous ones with high accuracy. These results establish a strong baseline for ambiguity detection and indicate that incorporating a lightweight classifier may be a practical and effective enhancement in meaning-sensitive applications, particularly in settings where computational efficiency is a priority. In contrast, despite their strong reasoning capabilities, the reasoning LLMs exhibit poor performance in ambiguity detection, frequently misclassifying clear, unambiguous sentences (as determined by human annotators) as ambiguous. This tendency to over-predict ambiguity weakens their practical reliability in meaning-sensitive tasks.

**Table 5: Set F1 performance on ambiguity meaning understanding under different prompting strategies. Each model is evaluated in two settings: Direct Interpretation (no disambiguation prompt) and Prompted Disambiguation (with disambiguation prompt). Methods include Direct prompt, Few-shot(FS), Knowledge+Prompt, Chain-of-Thought (CoT), CoT + Few-shot, and RAG-based Few-shot. Bolded scores represent the best performing model under each method, while underlined scores indicate the best performing method for each model. <sup>†</sup> Reasoning-enabled models.**

Model	Params	Direct Interpretation (Set F1)						Prompted Disambiguation (Set F1)					
		Direct	Few-shot	Knowledge	CoT	CoT+FS	RAG-FS	Direct	Few-shot	Knowledge	CoT	CoT+FS	RAG-FS
Gemma2	2B	40.49	44.79	46.01	44.23	45.46	<u>54.67</u>	41.18	47.83	49.31	50.14	48.04	<u>55.15</u>
	9B	44.71	48.99	52.07	49.93	48.68	<u>58.46</u>	43.92	50.76	52.58	52.78	53.57	<u>62.50</u>
	27B	46.37	50.45	54.74	51.07	50.26	<u>61.47</u>	45.69	53.96	55.03	53.57	52.26	<u>63.97</u>
Qwen3	4B	46.86	51.37	53.37	48.46	48.04	<u>56.62</u>	47.16	55.13	53.76	53.17	52.77	<u>59.93</u>
	14B	48.87	52.97	54.06	51.21	53.26	<u>61.03</u>	46.27	55.70	56.84	54.35	56.46	<u>63.60</u>
	32B	47.16	55.13	54.74	51.86	53.96	<u>65.07</u>	45.88	55.51	56.53	56.17	56.54	<u>67.65</u>
Qwen3 <sup>†</sup>	235B-A22B	51.67	58.33	59.61	54.74	59.25	<u>63.70</u>	52.65	57.96	<b>61.40</b>	59.43	59.12	<u>61.76</u>
DeepSeek-R1 <sup>†</sup>	671B-A37B	<b>55.49</b>	<b>61.40</b>	<b>61.31</b>	<b>57.22</b>	<b>61.75</b>	59.85	<b>56.18</b>	<b>61.05</b>	60.69	<b>62.10</b>	<b>59.97</b>	59.63



**Figure 1: Perplexity scores of Qwen3 models for ambiguous sentences across three context types. When similar context is given for a sentence, no matter ambiguous or disambiguated, there is no observed significant difference in perplexity.**


Through experiments, as shown in Table 3, we observed that the effectiveness of prompting strategies in Chinese ambiguity detection heavily relies on models’ intrinsic reasoning capabilities and parameter scale. For small-to-medium models, the Chain-of-Thought (CoT) method demonstrated limited effectiveness. However, performance improved significantly when supplemented with few-shot examples. This suggests that small-to-medium models lack the necessary reasoning capacity to leverage CoT strategies effectively; instead, they better grasp task fundamentals through concrete examples, thereby substantially enhancing their ability to detect ambiguity. We also found that specialized reasoning models (e.g., Qwen3 and DeepSeek-R1) excelled across all prompting strategies and could more effectively unlock their potential using the CoT+FS strategy to achieve peak performance. In contrast, non-reasoning-specialized models relied more heavily on external knowledge frameworks provided by Few-shot and Knowledge strategies, with their internal reasoning processes offering limited guidance.

Table 4 presents a comprehensive evaluation comparing two strategies: Direct Interpretation (asking for the meaning directly) and Prompted Disambiguation (asking for the meaning with an

explicit cue that the sentence is ambiguous). The evaluation is conducted by comparing the predicted set of meanings with the gold-standard set, using exact match, recall, and set-level F1 score as metrics. The results indicate that models perform poorly on this task, and the inclusion of an ambiguity prompt does not yield consistent or reliable improvements. Given the prohibitive cost of human evaluation at scale, especially for tens of thousands of meaning-level sentence comparisons, we employ strong reasoning models to approximate this process by comparing the predicted and reference meaning sets and outputting the number of overlapping meanings, which is then used to compute the evaluation metrics.

The performance gap between the Direct Interpretation and the Prompted Disambiguation frameworks in Table 5 reveals the significant impact of instruction framing on model comprehension. Under the Prompted Disambiguation framework, models consistently outperformed those using Direct Interpretation across all prompting strategies, demonstrating that explicit ambiguity-specific prompting enhances models’ sensitivity to multi-interpretation scenarios. These findings provide a theoretical basis for optimizing large language models in Chinese ambiguity understanding tasks and reveal differential sensitivity to prompting strategies across models of varying scales.

Through evaluations on three tasks: ambiguity detection, ambiguity understanding, and end-to-end assessment, as shown in Table 3 and Table 5, we observe that model performance improves with increased parameter size in both ambiguity detection and meaning understanding. Reasoning models often perform better across different prompting methods. We also find that the RAG method enhances sensitivity to Chinese ambiguity, especially for medium-sized non-reasoning models, by helping them identify multiple interpretations using relevant examples. Moreover, larger models benefit more from RAG, suggesting that reasoning ability plays a key role in handling ambiguity. Specifically, in the ambiguity detection task, due to its relative simplicity, the RAG strategy shows significant improvements across all models; in ambiguity understanding and end-to-end evaluation, due to the increased task complexity, the improvement effects of the RAG strategy have upper limits, primarily constrained by the models’ inherent reasoning capabilities. As shown in Table 5, RAG provides modest improvements for non-reasoning models, while showing limited enhancement for



**Figure 2: Comparison of ambiguity probabilities (the probability that Qwen3-8B model answers YES relative to NO to the question *Is the sentence ambiguous or not?*) between ambiguous sentences and their disambiguated versions. The disambiguation does not consistently reduce the model’s perceived ambiguity.**

models with strong reasoning capabilities (such as DeepSeek-R1). This occurs because strong reasoning models rely more on internal logic rather than external prompts, and are more sensitive to retrieval noise. Although high-quality retrieval still has positive effects, low-quality retrieval may have negative impacts, exhibiting diminishing marginal returns. For medium-scale models, RAG provides additional reasoning pathways that compensate for their insufficient reasoning capabilities, while these models have sufficient capacity to process rich examples, making them the optimal range for RAG strategy application. Small models have limited capacity and difficulty fully utilizing complex examples, potentially being overwhelmed by excessive information.

## 4 Analysis and Discussion

### 4.1 Perplexity Analysis

A language model’s perplexity (PPL) on a sequence of tokens is calculated by averaging the log probability values of its predictions for each token in the sequence. Perplexity is a statistical metric that assesses a language model’s ability to predict a text sequence, reflecting the model’s uncertainty in assigning probabilities to upcoming tokens. While PPL is more considered as a measure that evaluates how well LLMs model text patterns, we also assume that it measures LLMs’ ability to understand text.

Since PPL scores are strongly affected by the model’s training data, they cannot be directly compared between different models or across different datasets. Nevertheless, if all models share the same training data, the PPL scores become more comparable. In this case, differences in perplexity can more reliably reflect variations in text understanding. Inspired by this, researchers developed log-probability-based methods to classify potentially deceptive articles [20, 48] and check AI-generated content [25, 40].

In this study, we compare the PPL scores of a set of Qwen3 models on our benchmark to evaluate their relative certainty and predictive performance. Since all these models share the same training

data and vocabulary, the perplexity values are directly comparable to some degree. This comparison can provide insights into how confidently each model handles the benchmark’s input sequences.

For each sample in the benchmark, we measure a triplet of PPL scores: (a) the PPL of the ambiguous sentence without preceding or following context; (b) the PPL of the ambiguous sentence with ambiguous context; (c) the PPL of the ambiguous sentence with disambiguated context. We filter out samples whose (b) and (c) versions differ substantially in length, ensuring that the PPL scores are more comparable.

The results are shown in Figure 1. We observe that sentences with context generally have lower perplexity than those without context. However, when the provided context is similar in both length and semantic meaning, regardless of whether they are ambiguous or disambiguated, there is no significant difference in perplexity between them. This observation suggests that PPL scores may not serve as a reliable signal for LLMs’ ambiguity understanding ability. We also note that larger models tend to have lower perplexity scores, suggesting that they are more confident in their understanding of those ambiguous sentences.

As part of analyzing the decoding dynamics of Qwen3 models under conditioned inputs, we evaluate their token-level log-probability assignments on pairs of ambiguous and disambiguated sentences. For each sentence, we explicitly ask whether it is ambiguous to assess the model’s inherent sensitivity to ambiguity. Based on prior assumptions, we hypothesized that ambiguous sentences would elicit higher probabilities for a YES response compared to their disambiguated counterparts. However, as shown in Figure 2, no clear or consistent pattern was observed. This result suggests that log-probabilities may not serve as a reliable signal for detecting ambiguity, cross-validating our earlier observation that large language models exhibit limited awareness of linguistic ambiguity in Chinese text.

### 4.2 Probing Ambiguity via Clarification Questioning

To further investigate the model’s robustness against Chinese textual ambiguity, we propose an evaluation method inspired by Natural Language Inference (NLI) framing.

Every premise contains an ambiguous expression with two possible interpretations (A and B). Then, three hypotheses are generated: **Entailment**: is inferable from the premise with interpretation A. **Neutral**: Remains ambiguous, committing to neither A nor B. **Contradiction**: Supports Interpretation B and logically contradicts A [22].

Figure 3 illustrates the step-by-step process by which an LLM addresses semantic ambiguity in an NLI scenario. When given an ambiguous premise, the model may fail to make a definitive inference judgment. It generates a clarification question to explicitly resolve the ambiguity. Once the user provides a disambiguating answer, the model determines the inference relation: entailment, contradiction, or neutral. This process provides a way to evaluate whether the model has correctly identified and understood the ambiguity. By leveraging joint reasoning to identify the minimal conditions needed for clarification or decision-making, our analysis



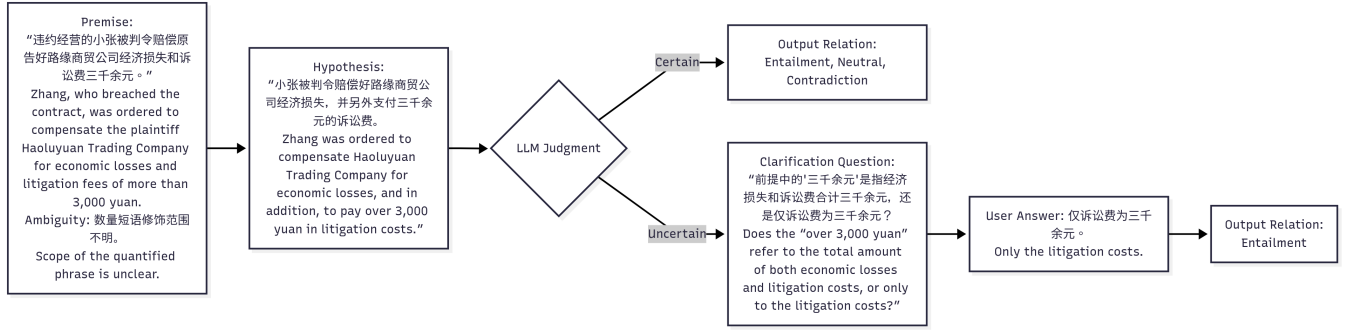


Figure 3: Example workflow for resolving ambiguity through clarification questions.




Figure 4: Illustration of an LLM’s failure to generate an effective clarification question. The model incorrectly focuses on emotional reasoning (guilt or regret) rather than resolving the core syntactic ambiguity.

method is also conceptually similar with existing work on explanation through factual and counterfactual analysis [6].

Figure 4 illustrates a case where the LLM correctly detects ambiguity, but misidentifies the source of ambiguity. Instead of focusing on the syntactic uncertainty (i.e., who cried), the model assumes the ambiguity regarding the reason for crying. As a result, it generates a clarification question that is misaligned with human intuition and fails to resolve the key ambiguity.

## 5 Related Work

Disambiguation has been an extensively studied research topic in NLP, as ambiguity is inherently present in human language and communication. Traditional machine learning-based NLP approaches have primarily focused on word sense disambiguation [4, 27], employing knowledge-based methods, vector-based 1-nn classifiers, token taggers, and sequence taggers to resolve lexical ambiguity. Ambiguity detection has also been thoroughly explored in the literature. [11] developed a taxonomy for classifying ambiguity and created POS-based and rule-based tools to detect ambiguity in requirement documents. [10] trained word embeddings on domain-specific corpora and compared cross-domain term representations to automatically identify semantic ambiguities. Large Language Models (LLMs) demonstrate exceptional capabilities in natural language understanding and reasoning tasks. Compared to early transformer models, LLMs exhibit superior performance and flexibility in comprehending and solving multiple-choice questions across diverse subjects including history, science, and mathematics, as demonstrated on benchmarks such as MMLU [14], MMLU-Pro [38], GPQA [30], and AIME [3]. LLMs also excel across multiple dimensions of language understanding, including commonsense reasoning [19] and interpretation of abstract concepts [45], and they even extend beyond the natural language domain, supporting tasks such

as coding [49], recommendation [39], forecasting, and anomaly detection [33]. However, existing reviews [17, 43] indicate that although LLMs demonstrate strong performance in language understanding tasks, they remain limited in their ability to capture fine-grained semantic nuances. Nevertheless, ambiguity remains a fundamental linguistic phenomenon that cannot be entirely overcome and has garnered significant attention from the research community. [22] presents an early work identifying limitations of LLMs in ambiguity understanding, and proposes AMBIENT, an English benchmark of ambiguous sentences. [31] specifically investigated ambiguity handling in questions to enhance LLM performance when confronted with ambiguous inputs. [18] explored improvements in LLM ambiguity handling for open-world question answering through simple prompt rewriting and context augmentation. [24] examined ambiguity detection mechanisms in LLMs. CLAMBER [44] addresses ambiguity challenges in query intention understanding and information clarification requirements for LLMs in retrieval tasks. Although these studies focus on enhancing LLM performance in specific applications, they do not examine the fundamental language understanding behaviors of LLMs when processing ambiguous content. In this work, we use Chinese as a case study to investigate how LLMs encounter and handle ambiguity with specific scenes, thereby providing meaningful insights for future research on ambiguity processing in LLMs.

## 6 Conclusion

In this work, we examine the fragility of large language models (LLMs) when handling textual ambiguity through Chinese oral input. We created a benchmark consisting of 900 ambiguous sentences with context across 9 categories, paired with corresponding disambiguation sentences. Our findings reveal that state-of-the-art



open-weight LLMs still struggle with ambiguity detection and understanding. Specifically, we observe three key issues. First, LLMs exhibit overconfidence when classifying sentences as ambiguous in detection tasks. Second, LLMs fail to effectively identify possible alternative meanings from ambiguous statements. Third, when explicitly prompted to understand ambiguous meanings, LLMs tend to overthink and generate meanings that are far-fetched compared to human interpretation. Our comprehensive experiments and analyses demonstrate several important findings. Models with more parameters perform better on these tasks, and reasoning-enhanced models show improved performance in both detection and understanding. Most notably, adding examples through retrieval-augmented generation (RAG) proves to be the most effective approach for improving both detection and understanding tasks. We also analyzed model behavior by examining perplexity differences between ambiguous and disambiguated sentences. Additionally, we explored ambiguity through probing techniques using clarification questions with case studies. This work provides a novel perspective on LLM trustworthiness and serves as a call for the community to address this inherent issue in LLMs and exercise caution in practical applications. For future work, we plan to conduct fine-grained analysis within different categories of ambiguity and develop lightweight, effective methods to mitigate these problems.

## Acknowledgements

We thank the anonymous reviewers for their valuable comments and suggestions.

## References

- [1] Suriya Ganesh Ayyamperumal and Limin Ge. 2024. Current state of LLM Risks and AI Guardrails. *arXiv preprint arXiv:2406.12934* (2024).
- [2] Razvan Azamfirei, Sapna R Kudchadkar, and James Fackler. 2023. Large language models and the perils of their hallucinations. *Critical Care* 27, 1 (2023), 120.
- [3] Mislav Balunović, Jasper Dekoninck, Ivo Petrov, Nikola Jovanović, and Martin Vechev. 2025. Matharena: Evaluating llms on uncontaminated math competitions. *arXiv preprint arXiv:2505.23281* (2025).
- [4] Michele Bevilacqua, Tommaso Pasini, Alessandro Raganato, and Roberto Navigli. 2021. Recent trends in word sense disambiguation: A survey. In *International joint conference on artificial intelligence*. International Joint Conference on Artificial Intelligence, Inc, 4330–4338.
- [5] Wei-Cheng Chang, Hsiang-Fu Yu, Kai Zhong, Yiming Yang, and Inderjit S Dhillon. 2020. Taming pretrained transformers for extreme multi-label text classification. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*. 3163–3171.
- [6] Ziheng Chen, Jin Huang, Fabrizio Silvestri, Yongfeng Zhang, Hongshik Ahn, and Gabriele Tolomei. [n. d.]. Joint Factual and Counterfactual Explanations for Top-k GNN-based Recommendations. *ACM Transactions on Recommender Systems* ([n. d.]).
- [7] Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. 2020. Revisiting Pre-Trained Models for Chinese Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*. Association for Computational Linguistics, Online, 657–668. <https://www.aclweb.org/anthology/2020.findings-emnlp.58>
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*. 4171–4186.
- [9] Leonard Dung. 2023. Current cases of AI misalignment and their implications for future risks. *Synthese* 202, 5 (2023), 138.
- [10] Alessio Ferrari and Andrea Esuli. 2019. An NLP approach for cross-domain ambiguity detection in requirements engineering. *Automated Software Engineering* 26, 3 (2019), 559–598.
- [11] Benedikt Gleich, Oliver Creighton, and Leonid Kof. 2010. Ambiguity detection: Towards a tool explaining ambiguity sources. In *Requirements Engineering: Foundation for Software Quality: 16th International Working Conference, REFSQ 2010, Essen, Germany, June 30–July 2, 2010. Proceedings* 16. Springer, 218–232.
- [12] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [13] Wenchong He, Zhe Jiang, Tingsong Xiao, Zelin Xu, and Yukun Li. 2023. A survey on uncertainty quantification methods for deep learning. *arXiv preprint arXiv:2302.13425* (2023).
- [14] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300* (2020).
- [15] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems* 43, 2 (2025), 1–55.
- [16] Tianyi Huang, Zeqiu Xu, Peiyang Yu, Jingyuan Yi, and Xiaochuan Xu. 2025. A hybrid transformer model for fake news detection: Leveraging Bayesian optimization and bidirectional recurrent unit. *2025 8th International Symposium on Big Data and Applied Statistics (ISBDAS 2025)* (2025).
- [17] Tianyi Huang, Jingyuan Yi, Peiyang Yu, and Xiaochuan Xu. 2025. Unmasking digital falsehoods: A comparative analysis of LLM-based misinformation detection strategies. In *2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE)*. IEEE, 2470–2476.
- [18] Aryan Keluskar, Amrita Bhattacharjee, and Huan Liu. 2024. Do LLMs Understand Ambiguity in Text? A Case Study in Open-world Question Answering. In *2024 IEEE International Conference on Big Data (BigData)*. IEEE, 7485–7490.
- [19] Stefanie Krause and Frieder Stolzenburg. 2023. Commonsense reasoning and explainable artificial intelligence using large language models. In *European Conference on Artificial Intelligence*. Springer, 302–319.
- [20] Nayeon Lee, Yejin Bang, Andrea Madotto, and Pascale Fung. 2021. Towards Few-shot Fact-Checking via Perplexity. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 1971–1981.
- [21] Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen. 2023. Large language models in finance: A survey. In *Proceedings of the fourth ACM international conference on AI in finance*. 374–382.
- [22] Alisa Liu, Zhaofeng Wu, Julian Michael, Alane Suhr, Peter West, Alexander Koller, Swabha Swayamdipta, Noah A Smith, and Yejin Choi. 2023. We’re Afraid Language Models Aren’t Modeling Ambiguity. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 790–807.
- [23] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [24] Behrang Mehrparvar and Sandro Pezzelle. 2024. Detecting and Translating Language Ambiguity with Multilingual LLMs. In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*. 310–323.
- [25] Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. 2023. Detectgpt: Zero-shot machine-generated text detection using probability curvature. In *International Conference on Machine Learning*. PMLR, 24950–24962.
- [26] Lingbo Mo, Boshi Wang, Muhao Chen, and Huan Sun. 2024. How Trustworthy are Open-Source LLMs? An Assessment under Malicious Demonstrations Shows their Vulnerabilities. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 2775–2792. doi:10.18653/v1/2024.naacl-long.152
- [27] Roberto Navigli. 2009. Word sense disambiguation: A survey. *ACM computing surveys (CSUR)* 41, 2 (2009), 1–69.
- [28] Wenjie Qu, Yuguang Zhou, Yongji Wu, Tingsong Xiao, Binhang Yuan, Yiming Li, and Jiaheng Zhang. 2025. Prompt inversion attack against collaborative inference of large language models. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1695–1712.
- [29] Mohaimenul Azam Khan Raiaan, Md Saddam Hossain Mukta, Kaniz Fatema, Nur Mohammad Fahad, Sadman Sakib, Most Marufatul Jannat Mim, Jubaer Ahmad, Mohammed Eunus Ali, and Sami Azam. 2024. A review on large language models: Architectures, applications, taxonomies, open issues and challenges. *IEEE access* 12 (2024), 26839–26874.
- [30] David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- [31] Zhengyan Shi, Giuseppe Castellucci, Simone Filice, Saar Kuzi, Elad Kravi, Eugene Agichtein, Oleg Rokhlenko, and Shervin Malmasi. 2025. Ambiguity Detection and Uncertainty Calibration for Question Answering with Large Language Models. In *Proceedings of the 5th Workshop on Trustworthy NLP (TrustNLP 2025)*, Trista Cao, Anubrata Das, Tharindu Kumarage, Yixin Wan, Satyapriya Krishna,

- Ninareh Mehrabi, Jwala Dhamala, Anil Ramakrishna, Aram Galystan, Anoop Kumar, Rahul Gupta, and Kai-Wei Chang (Eds.). Association for Computational Linguistics, Albuquerque, New Mexico, 41–55. doi:10.18653/v1/2025.trustnlp-main.4
- [32] Ola Shorinwa, Zhiting Mei, Justin Lidard, Allen Z Ren, and Anirudha Majumdar. [n.d.]. A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions. *Comput. Surveys* ([n.d.]).
- [33] Jing Su, Chufeng Jiang, Xin Jin, Yuxin Qiao, Tingsong Xiao, Hongda Ma, Rong Wei, Zhi Jing, Jiajun Xu, and Junhong Lin. 2024. Large language models for forecasting and anomaly detection: A systematic literature review. *arXiv preprint arXiv:2402.10350* (2024).
- [34] Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223* (2019).
- [35] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118* (2024).
- [36] Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. 2023. Large language models in medicine. *Nature medicine* 29, 8 (2023), 1930–1940.
- [37] Hao Tian, Can Gao, Xinyan Xiao, Hao Liu, Bolei He, Hua Wu, Haifeng Wang, and Feng Wu. 2020. SKEP: Sentiment Knowledge Enhanced Pre-training for Sentiment Analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 4067–4076.
- [38] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhrranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, et al. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- [39] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2024. A survey on large language models for recommendation. *World Wide Web* 27, 5 (2024), 60.
- [40] Zhenyu Xu and Victor S Sheng. 2024. Detecting AI-generated code assignments using perplexity of large language models. In *Proceedings of the aaai conference on artificial intelligence*, Vol. 38. 23155–23162.
- [41] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388* (2025).
- [42] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems* 32 (2019).
- [43] Jingyuan Yi, Zeqiu Xu, Tianyi Huang, and Peiyang Yu. 2025. Challenges and innovations in LLM-Powered fake news detection: A synthesis of approaches and future directions. In *Proceedings of the 2025 2nd International Conference on Generative Artificial Intelligence and Information Security*. 87–93.
- [44] Tong Zhang, Peixin Qin, Yang Deng, Chen Huang, Wenqiang Lei, Junhong Liu, Dingnan Jin, Hongru Liang, and Tat-Seng Chua. 2024. CLAMBER: A Benchmark of Identifying and Clarifying Ambiguous Information Needs in Large Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 10746–10766. doi:10.18653/v1/2024.acl-long.578
- [45] Yigeng Zhang, Fabio Gonzalez, and Thamar Solorio. 2024. Interpreting Themes from Educational Stories. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italia, 9190–9203. <https://aclanthology.org/2024.lrec-main.805/>
- [46] Yigeng Zhang, Mahsa Shafaei, Fabio Gonzalez, and Thamar Solorio. 2021. From None to Severe: Predicting Severity in Movie Scripts. In *Findings of the Association for Computational Linguistics: EMNLP 2021*. 3951–3956.
- [47] Yigeng Zhang, Mahsa Shafaei, Fabio Gonzalez, and Thamar Solorio. 2024. Positive and Risky Message Assessment for Music Products. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, Nicoletta Calzolari, Min-Yen Kan, Veronique Hoste, Alessandro Lenci, Sakriani Sakti, and Nianwen Xue (Eds.). ELRA and ICCL, Torino, Italia, 12893–12905. <https://aclanthology.org/2024.lrec-main.1129/>
- [48] Yigeng Zhang, Fan Yang, Yifan Zhang, and Arjun Mukherjee. 2020. Birds of a Feather Flock Together: Satirical News Detection via Language Model Differentiation. In *International Conference on Social Computing, Behavioral-Cultural Modeling & Prediction and Behavior Representation in Modeling and Simulation*.
- [49] Ziyin Zhang, Chaoyu Chen, Bingchang Liu, Cong Liao, Zi Gong, Hang Yu, Jianguo Li, and Rui Wang. 2023. Unifying the perspectives of nlp and software engineering: A survey on language models for code. *arXiv preprint arXiv:2311.07989* (2023).

- [50] Zhengyan Zhang, Xu Han, Hao Zhou, Pei Ke, Yuxian Gu, Deming Ye, Yujia Qin, Yusheng Su, Haozhe Ji, Jian Guan, et al. 2021. CPM: A large-scale generative Chinese pre-trained language model. *AI Open* 2 (2021), 93–99.

## A Implementation details

### A.1 Choice of Ambiguity Detection Model

To better handle this Chinese-specific task, we consider using language models that are pretrained on Chinese corpora and tasks [7, 34, 50]. Among them, we select hf1/chinese-roberta-wwm-ext [7], a RoBERTa-based model specifically designed for Chinese. Unlike standard BERT models that apply subword-level masking, this model adopts whole-word masking (WWM), meaning that it masks entire Chinese words during pretraining. Since Chinese words often consist of multiple characters, WWM enables the model to learn more meaningful word-level representations. This property is particularly beneficial for identifying sentence-level ambiguity, where subtle differences in phrasing can lead to different interpretations.

### A.2 Detection Model Training Procedure

To fine-tune the model for our task, we use the training set of the manually annotated samples from the dataset. Each ambiguous sentence is paired with two disambiguated versions that preserve the original meaning while removing the ambiguity. Ambiguous and disambiguated sentences are labeled as 1 and 0, respectively. This structure provides a semantic contrast between positive and negative examples.

To enhance the input representation, we added linguistic features to each sentence. Specifically, we appended a word-segmented version of the sentence using `jieba` and the corresponding part-of-speech (POS) tags. The final input format contains both lexical and syntactic cues, aiding the model in better understanding structural aspects of Chinese that are often associated with ambiguity.

For training, the model configuration included a learning rate of  $2e-5$ , a batch size of 16, 5 training epochs, Adam optimizer, and a linear learning rate decay schedule. Early stopping was used to prevent overfitting, halting training if validation performance did not improve for 3 consecutive epochs. We set a random seed for reproducibility and used CUDA to accelerate the computation. We applied gradient clipping (with `max_norm=1.0`) to avoid exploding gradients, and used a linear learning rate scheduler with warm-up steps. The early stopping strategy was implemented based on the validation F1 score, with a patience of three epochs.

To reduce performance variance and improve robustness, stratified K-fold cross-validation was used during training. Additionally, we applied automated hyperparameter tuning using the Optuna framework. The search space included batch size, learning rate, and weight decay, and each configuration was evaluated based on the cross-validation F1 score. This approach allowed us to identify a better combination of parameters with minimal manual tuning.

### A.3 Choice of LLMs

In this work, we focus exclusively on open-weight LLMs for our experiments. While proprietary models have demonstrated strong language understanding capabilities, their APIs and chat interfaces

function as black boxes, making it unclear whether additional components beyond the model weights influence the outputs. This lack of transparency may affect experimental validity and reduce reproducibility. Therefore, we selected open-weight models Gemma 2, Qwen 3, and DeepSeek R1 for our study, with Qwen and DeepSeek R1 in particular being developed by Chinese researchers and showing strong performance on tasks in Chinese. Using open-weight models for benchmarking is also a reasonable practice in the research community [26, 28].

## B Author Contributions

Xinwei Wu contributed to the creation of the benchmark dataset and conducted experiments on ambiguity detection and ambiguity understanding across multiple large language models.

Hongyu Liu contributed to a portion of the benchmark dataset development and conducted comparative experiments with result evaluation for Retrieval-Augmented Generation (RAG) methodology in Chinese ambiguity processing.

Haojie Li contributed to the creation of the benchmark dataset and implemented the full pipeline for data processing and disambiguation, including data cleaning, feature engineering, model training and optimization, and experimental evaluation.

Xinwei Wu, Hongyu Liu, and Haojie Li contributed equally to this work and share first authorship.

Xinyu Ji contributed to the creation of the benchmark dataset and conducted evaluations of ambiguity detection and ambiguity understanding across multiple large language models. All contributions were made in a personal capacity and do not reflect the views of her employer.

Ruohan Li contributed to the creation of the benchmark dataset and conducted an evaluation via clarification questioning.

Yule Chen conducted the perplexity analysis and contributed to the creation of the benchmark dataset.

Yigeng Zhang provided overall supervision and mentored the team across all stages of the project. Yigeng Zhang participated in this project as a volunteer mentor supporting early-stage researchers. All contributions were made in his personal capacity and do not represent the views of his employer.