

# Mitigating Boundary Ambiguity and Inherent Bias for Text Classification in the Era of Large Language Models

Zhenyi Lu<sup>1,2\*</sup>, Jie Tian<sup>1,2\*</sup>, Wei Wei<sup>1,2†</sup>, Xiaoye Qu<sup>1,3</sup>,  
Yu Cheng<sup>4</sup>, Wenfeng Xie<sup>5</sup>, Dangyang Chen<sup>5†</sup>

<sup>1</sup> Cognitive Computing and Intelligent Information Processing (CCIIP) Laboratory, School of Computer Science and Technology, Huazhong University of Science and Technology, <sup>2</sup> Joint Laboratory of HUST and Pingan Property & Casualty Research (HPL), <sup>3</sup> Shanghai AI Laboratory, <sup>4</sup> The Chinese University of Hong Kong, <sup>5</sup> Ping An Property & Casualty Insurance Company of China, Ltd.  
{luzhenyi529, jaytian.china}@gmail.com, {weiw, xiaoye}@hust.edu.cn  
chengyu@cse.cuhk.edu.hk, julian\_wind@163.com, chendangyang273@pingan.com.cn

## Abstract

Text classification is a crucial task encountered frequently in practical scenarios, yet it is still under-explored in the era of large language models (LLMs). This study shows that LLMs are vulnerable to changes in the number and arrangement of options in text classification. Our extensive empirical analyses reveal that the key bottleneck arises from ambiguous decision boundaries and inherent biases towards specific tokens and positions. To mitigate these issues, we make the first attempt and propose a novel two-stage classification framework for LLMs. Our approach is grounded in the empirical observation that pairwise comparisons can effectively alleviate boundary ambiguity and inherent bias. Specifically, we begin with a self-reduction technique to efficiently narrow down numerous options, which contributes to reduced decision space and a faster comparison process. Subsequently, pairwise contrastive comparisons are employed in a chain-of-thought manner to draw out nuances and distinguish confusable options, thus refining the ambiguous decision boundary. Extensive experiments on four datasets (Banking77, HWU64, LIU54, and Clinic150) verify the effectiveness of our framework. Furthermore, benefitting from our framework, various LLMs can achieve consistent improvements. Our code and data are available in <https://github.com/Chuge0335/PC-CoT>.

## 1 Introduction

Text classification, as one of the most fundamental and common tasks in natural language processing (NLP), has a wide spectrum of applications (Qu et al., 2019; Lei et al., 2023b; Chen et al., 2022; Zaremba and Demir, 2023; Labonne and Moran,

\* Equal contribution.

† Corresponding authors: Wei Wei and Dangyang Chen.

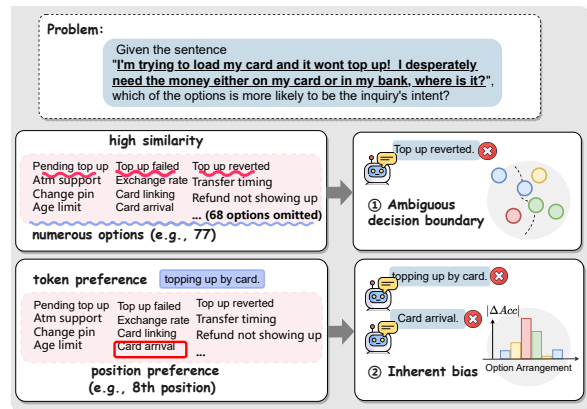


Figure 1: LLMs in text classification suffer from two main challenges: (a) ambiguous decision boundary, which arises from the explosion of options and their similarity. (b) inherent bias, favoring specific optional positions or tokens. These challenges lead to decreased classification accuracy.

2023). A common approach for text classification involves supervised fine-tuning on a specific dataset, which can be costly and may suffer from generalization issues (Sung et al., 2023; Bohra et al., 2023; Jiang et al., 2023). Recent Large Language Models (LLMs), such as ChatGPT (OpenAI, 2022), LLaMA2 (Touvron et al., 2023) and Qwen (Bai et al., 2023), have consistently exhibited remarkable language comprehension capabilities and performance emergence (Wei et al., 2022a). Such advancement has introduced a new paradigm of classification without the need for finetuning. LLMs can now make selections in a zero-shot manner, or employ techniques like few-shot demonstrations (Brown et al., 2020) and chain-of-thought (CoT) (Kojima et al., 2022).

Unfortunately, we observe that LLMs exhibit vulnerability to changes in the number and arrangement of options in text classification tasks. Our extensive empirical evaluation, involving 11 LLMs

ranging from 1B to 70B parameters and classification tasks with over 50 options, demonstrates that **boundary ambiguity** and **inherent bias** are prevalent issues in various large language models. *These issues cannot be effectively mitigated by long-context capabilities, simple prompting strategies, or single increasing the model size.* LLMs struggle with distinguishing among a growing number of options, a phenomenon called *boundary ambiguity*. For example, gpt-3.5-turbo achieves a zero-shot accuracy of only 32.51% when choosing the correct option from 60 categories, compared to 94.29% for two categories. LLMs also exhibit *inherent bias*, favoring specific positions or selectively retaining or discarding tokens in classification tasks. Placing the correct option consistently in the first position results in performance fluctuations of 35.29% for LLaMA2-70B-Chat and 38.52% for Qwen-72B-Chat. Furthermore, LLaMA2 shows an abnormal bias towards selecting “movies” and “convert” in the LIU54 benchmark.

To address these challenges for LLMs, our intuition is that LLMs often have the highest accuracy in **pairwise** classification, and focusing on two options at a time can significantly reduce the number of tokens and positions considered, potentially alleviating bias issues. This divide-and-conquer approach also mirrors human decision-making with multiple options, where we first eliminate less likely options and then carefully differentiate between similar ones. The latter process often involves juxtaposing confusable objects pairwise to assess their similarities and differences.

In view of this, we design a general two-stage LLM classification framework consisting of reduction and comparison stages. Within such a framework, we introduce two novel self-reduction techniques for the reduction process: one iteratively isolates the most probable option from candidates, and the other induces distinct option “windows” via clustering, operating within these windows to avoid similarity confusion, termed as ITR and CBWR. Additionally, we propose PC-CoT, an innovative contrastive chain-of-thought technique, to draw out nuances and distinguish features among confusable options, refining the ambiguous decision boundary. This type of detailed comparison encourages LLMs to conduct more in-depth analyses of actual content, rather than making shortcuts based solely on position or tokens. We conduct extensive experiments to demonstrate that our framework is con-

sistently effective for various LLMs to combat ambiguity and bias, even when the number of options increases to 150. Notably, our method achieves a relative improvement of 54.1% in average accuracy compared to the full-option zero-shot performance of gpt-3.5-turbo. It significantly reduces the token bias of LLaMA-70B-Chat, improving the relative performance by 36.88%.

To sum up, our contributions are as follows: (1) Our work provides extensive empirical analyses (covering 11 LLMs) to explore the vulnerability of LLMs in option number and arrangement in text classification. We identify that their ambiguous decision boundaries and inherent biases on tokens and positions remain a main challenge. (2) To address these issues, we propose utilizing pairwise comparison to combat such ambiguity and bias and develop an innovative two-stage classification framework for LLMs. Within the framework, novel self-reduction algorithms and a contrastive elimination strategy are introduced. (3) We carry out comprehensive experiments on four primary benchmarks along with diverse LLMs, demonstrating the effectiveness of our framework.

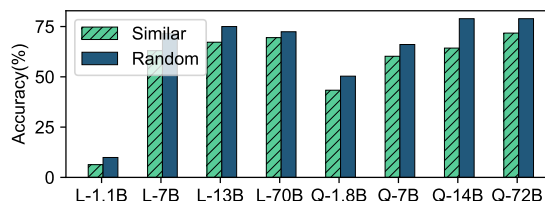
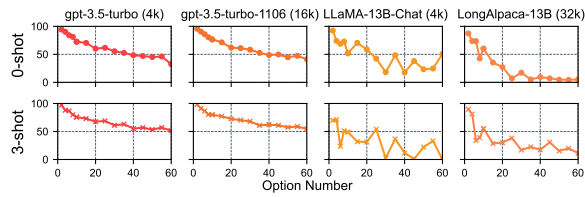


Figure 2: The classification accuracy decreases across various LLMs as the similarity among options increases. “L” denotes the LLaMA-2-Chat series, while “Q” represents the Qwen-Chat series.

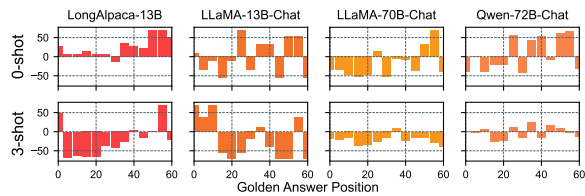
## 2 Preliminary Study

**Setup** We evaluate various LLMs to understand factors influencing classification accuracy. Our experiments include the gpt-3.5 series, LLaMA2-Chat, and Qwen-Chat models. We use the challenging portions of the Banking77 and LIU54 datasets (details in Section 4) to reduce experimental costs. Due to space constraints, the main results are summarized in the main text, with detailed dataset samples and additional results provided in Appendix A.

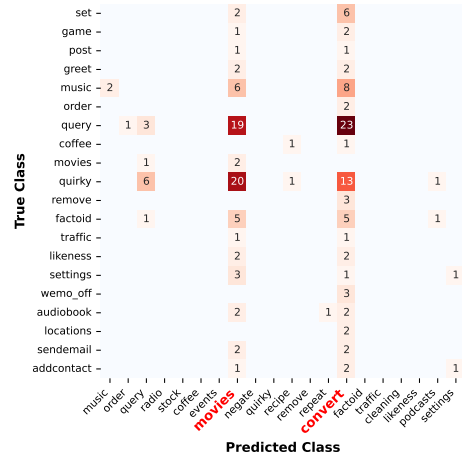
**Key Observations** Our findings, presented in Figure 2 and Figure 3, reveal the following key observations and insights:



(a) LLMs classification accuracy in relation to the number of options consistently demonstrates a decline.



(b) The relative performance change rate for fixed correct options in a position (range from 0 to 60) compared to random configuration performance. Simple option arrangement leads to dramatic performance fluctuations in both zero-shot and few-shot scenarios.



(c) The confusion matrix sheds light on the token bias of LLaMA-2 in the LIU54 dataset. Notably, the examination reveals a pronounced bias of LLaMA-2-70B-Chat towards the tokens “movies” and “convert”.

Figure 3: The influence of option number and arrangement on LLM performance. Full results are in the Appendix A.

**(1) The option similarity consistently poses a challenge.** We test randomly sampled or similarity-retrieved options in Figure 2 and observe a consistent performance drop across various LLM sizes and types. This finding indicates that LLMs encounter difficulties when faced with the inherent fuzziness of semantics.

**(2) The increased number of options consistently presents a challenge for LLMs.** Across various LLMs, including LLaMA, Qwen, and gpt-3.5-turbo, we consistently observe a decline in performance as the number of options increases, as shown in Figure 3a. In addition, this challenge persists when a longer context window is introduced, *e.g.*, the gpt-3.5-turbo-1106 with 16k context length and LongAlpaca with 32k context length. This phenomenon suggests that the noise stemming from the proliferation of options remains a significant obstacle in the landscape of LLMs.

**(3) Bias towards option positions and tokens.** In Figure 3b, we discern a significant bias caused by the position of the correct choice, especially in open-source models such as LLaMA and Qwen, resulting in performance fluctuations exceeding 50%. In contrast, gpt-3.5-turbo demonstrates minimal effects, with fluctuations below 10%. This lack of robustness demonstrates inherent position bias in LLMs. In Figure 3c, we analyze the confusion matrix of LLaMA for the LIU54 benchmark. A preference for “movies” and “convert” tokens can be clearly observed regardless of the golden label,

which indicates the existence of token bias.

**(4) The above issues cannot be effectively mitigated by long-context capabilities, simple prompting strategies, or single increasing the model size.** Even with a longer context, issues persist regarding the number and placement of options. For example, the gpt-3.5-turbo-1106 (16k context length) exhibits similar declining trends as its 4k-context counterpart with increasing option numbers. The LongAlpaca-32k has 50% performance fluctuation when the golden option is fixed at the 50<sup>th</sup> position. Furthermore, experiments with varying numbers of options show consistent trends in both few-shot and zero-shot scenarios. There are distinct bias distributions between 3-shot and zero-shot scenarios, indicating that in-context examples can alter the bias distributions of option positions but cannot effectively eliminate them. Regarding the model size, we observe that LLMs of different sizes all suffer from option similarity, as shown in Figure 2. Even 70B LLaMA and Qwen models suffer from an increase in option number and simple arrangement, as depicted in Figure 3.

**From Full Option to Pairwise Comparison** Inspired by Figure 3a, selecting the correct label from two choices is much easier for LLMs than directly selecting from all options. For instance, gpt-3.5-turbo achieves a 94.29% zero-shot accuracy when choosing between two categories. To this end, we propose a pairwise comparison framework to max-

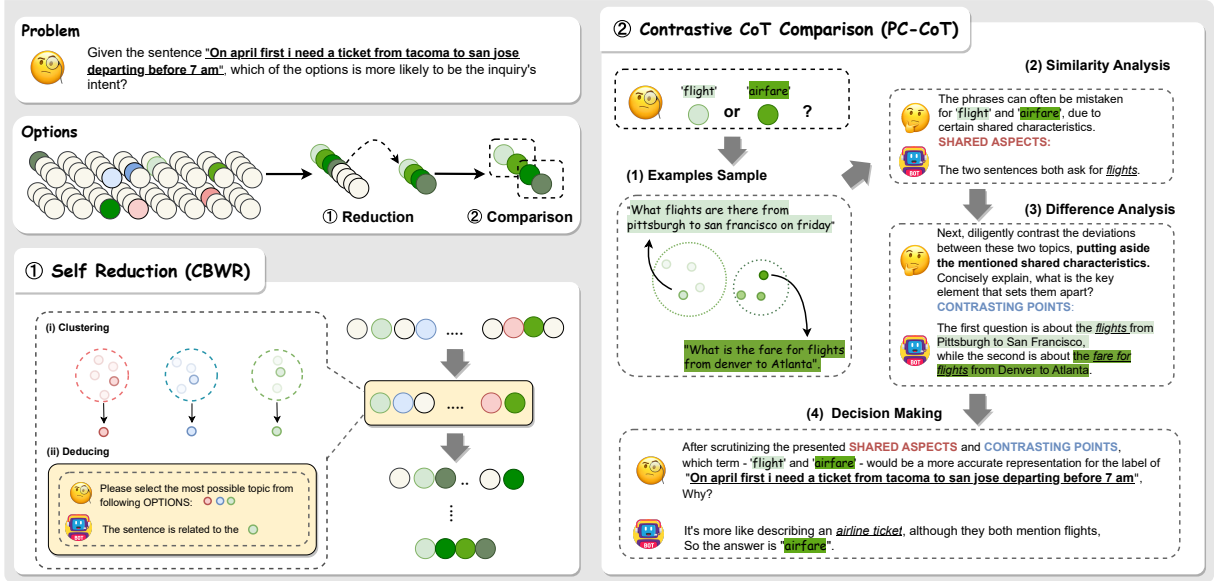


Figure 4: An overview of our framework for text classification with numerous options. It comprises two main stages: (1) Self Reduction (CBWR) which eliminates extraneous options, and (2) Contrastive CoT Comparison (PC-CoT) that conducts pairwise comparison. In CBWR, distinct option “windows” are induced via clustering to avoid confusion due to similarity. For pairwise comparison, our PC-CoT employs a contrastive chain-of-thought technique to highlight similarities and differences among the paired options for final decision making.

imize the capabilities of LLMs. Our approach focuses on comparing pairs of options, which reduces the range of positions involved in the classification process. This reduction has been shown to effectively mitigate position bias (Pezeshkpour and Hruschka, 2023). Furthermore, using pairwise comparisons allows for more nuanced comparative strategies, avoiding token shortcuts and ensuring that the most distinguish content guides decision-making.

### 3 Methodology

#### 3.1 Formulation

Formally, given a fixed LLM parameterized by  $\theta$ , we use  $x$  to denote the text classification input,  $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$  as  $N$  potential label options. Then, we utilize the prompt to instruct LLM to select the appropriate label  $y \in \mathcal{Y}$  for input  $x$ :

$$y \sim p_{\theta}(y|\text{Prompt}(x, \mathcal{Y})) \quad (1)$$

For simplicity, we will subsequently refer to  $p_{\theta}(y|\text{Prompt}(x, \mathcal{Y}))$  as  $p_{\theta}^{\text{Prompt}}(y|x, \mathcal{Y})$

The expressions for in-context learning ( $p_{\theta}^{\text{ICL}}$ ) and chain-of-thought methodology ( $p_{\theta}^{\text{CoT}}$ ) for tackling text classification problems are formulated as:

$$\begin{aligned} p_{\theta}^{\text{ICL}}(y | S_{icl}, x, \mathcal{Y}), \\ p_{\theta}^{\text{CoT}}(z, y | S_{cot}, x, \mathcal{Y}) \end{aligned} \quad (2)$$

where the set  $S$  denotes the demonstration sets, with an empty size ( $S = \emptyset$ ) in the zero-shot scenario and a non-empty size in the few-shot scenario. For the few-shot instances,  $S_{icl}$  includes pairs from  $\{x', y' | y' \in \mathcal{Y}, x' \neq x\}$ . For CoT methods,  $S_{cot}$  additionally include texts  $z'$ , viewed as the step-by-step reasoning “thoughts”. These thoughts serve as instructional support for LLMs to induce the conditional distribution of  $y'$  based on the given  $x'$ .

#### 3.2 Proposed Framework

The fundamental concept of our approach is to compare options in a pairwise manner. In cases where the options set is substantial, directly contrasting original options in pairs becomes inefficient and costly. Therefore, we propose a two-stage framework: (i) **Reduction**, strategically eliminating extraneous options  $\mathcal{Y}$  to a reduced set  $\mathcal{R} \subset \mathcal{Y}$ . (ii) **Comparison**, navigating through the interference options within  $\mathcal{R}$  via pairwise comparison in a chain-of-thought fashion. In our two-stage framework, various reduction and comparison algorithms can be seamlessly integrated. This paper further introduces two new algorithms designed for this framework: self-reduction (in Section 3.3), which utilizes LLMs to filter out irrelevant options automatically, and the contrastive chain-of-thought technique (in Section 3.4), which mimics how humans compare confusable objects. The overall ar-

chitecture of our algorithm is illustrated in Figure 4.

### 3.3 Self Reduction

Given a sentence  $x$  and option sets  $\mathcal{Y}$ , LLMs iteratively refine the options to  $\mathcal{R}$  through a multi-iteration strategy, *i.e.*,  $p_\theta(\mathcal{R} | x, \mathcal{Y})$ . Specifically, we design two types of multi-iteration strategies to improve the elimination accuracy: (1) Iterative Top Reduction (ITR): where LLMs iteratively choose the most likely options from the given candidate set and subsequently isolate them in the following iteration. This cyclic process continues until a pre-set threshold is reached, facilitating the steady refinement of labels. (2) Cluster-Based Window Reduction (CBWR): Inspired from Figure 2, a high degree of option similarity may degrade selection precision. To mitigate this, CBWR refines the full set of options for an input sentence through iterative clustering and selection within a window. The remaining options are clustered into  $K$  clusters, with samples evenly sampled from each cluster and placed into a window. Top- $N$  candidates are then selected from the window by LLM, progressively discarding less probable options. The detailed procedure can be found in Algorithm 2 in Appendix B.

### 3.4 Contrastive CoT Comparison

In this phase, the remaining candidates within the set  $\mathcal{R}$  often exhibit significant ambiguity. Each option has a close relationship with input  $x$  from the perspective of LLMs. To accurately identify the candidate, we propose a novel strategy based on a “similar-difference-decide” reasoning chain during the pairwise comparison process, which we termed PC-CoT. This approach is inspired by the concept of contrastive learning.

As specified in Algorithm 1, our approach begins by initiating a candidate pool with the reduced label set, systematically selecting two labels in each iteration, denoted as  $y_1, y_2$ . Subsequently, the demonstration set expands to include the demonstrations of both labels, *i.e.*,  $S = S_{y_1} \cup S_{y_2}$ . The LLMs are then prompted to analyze the similarity ( $z_s$ ) and differences ( $z_d$ ) of demonstration sentences associated with these labels. This step aims for a better understanding of the relationship between the two candidates and the refinement of the final decision on which label to eliminate. This iterative removal of candidates continues until we arrive at a sole label, which is regarded as the final answer. The algorithm exhibits a time complexity of  $O(|\mathcal{R}|)$ .

Model	Dataset	P	SC	ITR	CBWR
gpt-3.5-turbo	Banking77	40.48	65.35	<b>80.05</b>	<u>77.02</u>
	HWU64	49.93	76.25	<u>90.96</u>	<b>90.96</b>
	LIU54	32.17	37.99	<u>52.51</u>	<b>55.69</b>
	Clinc150	66.11	88.28	<u>93.11</u>	<b>94.37</b>
	Average	47.17	66.97	<u>79.16</u>	<b>79.51</b>
LLaMA2-70B-Chat	Banking77	33.44	37.34	<b>71.43</b>	<u>61.36</u>
	HWU64	39.09	27.27	<b>81.82</b>	<u>78.90</u>
	LIU54	35.04	27.95	<b>63.78</b>	<u>49.61</u>
	Clinc150	52.44	37.33	<b>85.78</b>	<u>73.99</u>
	Average	40.00	32.47	<b>75.70</b>	<u>65.97</u>
Qwen-72B-Chat	Banking77	42.21	71.75	<b>80.52</b>	<u>79.55</u>
	HWU64	56.36	72.73	<u>85.45</u>	<b>91.82</b>
	LIU54	31.50	53.94	<b>63.78</b>	<u>58.27</u>
	Clinc150	61.78	86.67	<u>93.33</u>	<b>94.67</b>
	Average	47.96	71.27	<u>80.77</u>	<b>81.08</b>

Table 1: HIT@5 performance evaluation for the reduction stage. We compare the effectiveness of our two self-reduction strategies, ITR and CBWR, against LLM employing either a standard reduction prompt (P) or paired with a self-consistency technique (SC). The best results are highlighted in **bold**, and the second best results are underlined.

Given that the size of  $\mathcal{R}$  is typically small, our approach ensures a cost-efficient performance.

#### Algorithm 1 PC-CoT

---

**Require:** Input  $x$ , reduced label options  $\mathcal{R} = \{y_1, y_2, \dots, y_r\}$ , language model  $p_\theta$ , training set  $\mathcal{D}$ , similar analysis prompt  $\text{Sim}$ , difference analysis prompt  $\text{Diff}$  and decision prompt  $\text{Dec}$ .

- 1: **while**  $|\mathcal{R}| > 1$  **do**
- 2:   Select two label options  $\{y_1, y_2\}$  from  $\mathcal{R}$
- 3:    $\mathcal{R} = \mathcal{R} \setminus \{y_1, y_2\}$
- 4:   Sample demonstrations pairs  $\{x, y\} \in \mathcal{D}$  for  $\{y_1, y_2\}$  as  $S_{y_1}, S_{y_2}$
- 5:    $S \leftarrow S_{y_1} \cup S_{y_2}$
- 6:   Sample  $z_s \sim p_\theta^{\text{sim}}(z | x, S, \{y_1, y_2\})$
- 7:   Sample  $z_d \sim p_\theta^{\text{diff}}(z | z_s, x, S, \{y_1, y_2\})$
- 8:   Predict  $y_p \sim p_\theta^{\text{dec}}(y | z_s, z_d, x, S, \{y_1, y_2\})$
- 9:    $\mathcal{R} = \mathcal{R} \cup \{y_p\}$
- 10: **end while**
- 11: **return**  $R$

---

## 4 Experiments

To evaluate the performance of our framework, we choose four widely recognized classification datasets, including: **Banking77** (Casanueva et al., 2020) with 77 labels in the banking services domain; **HWU64** (Casanueva et al., 2020) of 64 intents from 21 different domains; **LIU54** (Liu et al., 2021) of 54 classes gathered from Amazon Mechanical Turk; **Clinc150** (Larson et al., 2019) covering 150 classes and 10 domains. Notice that we sample challenging segments from these datasets to reduce the experiment cost. The comprehensive sampling process and supervised performance are

Model	Dataset	Full Options				Our Framework				
		ZS	ZS-CoT	FS	FS-CoT	ZS	ZS-CoT	FS	FS-CoT	PC-CoT
gpt-3.5-turbo	Banking77	30.81 $\pm$ 2.80	37.39 $\pm$ 1.97	41.91 $\pm$ 2.97	47.40 $\pm$ 0.65	44.43 $\pm$ 2.22	45.19 $\pm$ 3.09	<u>50.06</u> $\pm$ 4.33	47.53 $\pm$ 2.87	<b>61.82</b> $\pm$ 2.57
	HWU64	41.48 $\pm$ 2.88	<u>60.59</u> $\pm$ 1.99	56.05 $\pm$ 1.34	57.54 $\pm$ 4.89	52.43 $\pm$ 0.11	54.46 $\pm$ 0.25	59.05 $\pm$ 0.22	58.79 $\pm$ 0.27	<b>65.09</b> $\pm$ 0.87
	LIU54	29.30 $\pm$ 0.42	31.83 $\pm$ 1.17	36.37 $\pm$ 0.17	35.39 $\pm$ 2.76	33.95 $\pm$ 2.62	32.34 $\pm$ 1.91	<u>36.39</u> $\pm$ 3.38	31.67 $\pm$ 2.35	<b>37.35</b> $\pm$ 3.75
	Clinic150	59.03 $\pm$ 1.37	68.61 $\pm$ 0.55	78.77 $\pm$ 3.16	76.71 $\pm$ 0.47	77.68 $\pm$ 3.25	72.12 $\pm$ 0.17	78.75 $\pm$ 3.54	<u>81.29</u> $\pm$ 2.12	<b>83.78</b> $\pm$ 3.84
	Average	40.24 $\pm$ 1.24	49.60 $\pm$ 1.42	53.27 $\pm$ 1.91	54.26 $\pm$ 2.19	52.12 $\pm$ 2.05	51.03 $\pm$ 1.35	<u>56.06</u> $\pm$ 2.87	54.82 $\pm$ 2.33	<b>62.01</b> $\pm$ 2.33
LLaMA2-70B-Chat	Banking77	32.04 $\pm$ 4.55	32.14 $\pm$ 3.19	27.27 $\pm$ 3.10	35.71 $\pm$ 4.50	41.56 $\pm$ 1.26	27.93 $\pm$ 0.31	<b>50.54</b> $\pm$ 2.28	31.60 $\pm$ 1.60	<u>49.91</u> $\pm$ 1.01
	HWU64	43.94 $\pm$ 5.48	7.57 $\pm$ 0.53	31.21 $\pm$ 6.82	42.73 $\pm$ 2.73	50.91 $\pm$ 1.99	30.31 $\pm$ 0.52	61.21 $\pm$ 2.78	40.61 $\pm$ 3.20	<b>50.61</b> $\pm$ 1.39
	LIU54	29.13 $\pm$ 1.80	10.37 $\pm$ 7.04	8.66 $\pm$ 3.61	30.05 $\pm$ 2.98	33.86 $\pm$ 1.68	25.98 $\pm$ 0.68	36.22 $\pm$ 6.58	29.26 $\pm$ 0.23	<b>38.97</b> $\pm$ 2.19
	Clinic150	43.70 $\pm$ 2.72	5.19 $\pm$ 1.03	38.67 $\pm$ 9.69	52.74 $\pm$ 9.90	68.44 $\pm$ 1.77	42.07 $\pm$ 0.25	<b>82.22</b> $\pm$ 3.11	55.26 $\pm$ 2.00	<u>73.33</u> $\pm$ 2.10
	Average	37.20 $\pm$ 3.64	13.82 $\pm$ 2.95	26.45 $\pm$ 5.80	40.31 $\pm$ 5.03	48.69 $\pm$ 1.67	31.57 $\pm$ 0.44	<b>57.55</b> $\pm$ 3.69	39.18 $\pm$ 1.76	<u>53.21</u> $\pm$ 1.67
Qwen-72B-Chat	Banking77	43.83 $\pm$ 1.84	42.05 $\pm$ 0.23	48.70 $\pm$ 3.21	47.89 $\pm$ 2.06	49.68 $\pm$ 0.46	44.97 $\pm$ 0.23	53.09 $\pm$ 4.82	60.07 $\pm$ 5.06	<b>61.69</b> $\pm$ 0.92
	HWU64	54.55 $\pm$ 2.57	56.82 $\pm$ 4.50	55.00 $\pm$ 0.64	57.73 $\pm$ 0.64	43.64 $\pm$ 0.53	47.73 $\pm$ 0.64	<u>60.46</u> $\pm$ 3.22	56.82 $\pm$ 0.64	<b>63.64</b> $\pm$ 1.29
	LIU54	32.68 $\pm$ 5.56	31.10 $\pm$ 1.67	40.36 $\pm$ 4.18	36.62 $\pm$ 0.56	31.70 $\pm$ 0.28	29.53 $\pm$ 0.15	<u>39.37</u> $\pm$ 2.79	38.78 $\pm$ 0.28	<b>41.73</b> $\pm$ 1.12
	Clinic150	57.11 $\pm$ 0.95	58.00 $\pm$ 2.21	73.56 $\pm$ 2.20	73.56 $\pm$ 4.08	67.34 $\pm$ 0.32	63.12 $\pm$ 0.63	63.11 $\pm$ 3.78	<u>83.78</u> $\pm$ 0.31	<b>84.89</b> $\pm$ 0.49
	Average	47.04 $\pm$ 2.73	46.99 $\pm$ 2.15	54.40 $\pm$ 2.56	53.95 $\pm$ 1.84	48.09 $\pm$ 0.40	46.34 $\pm$ 0.41	54.01 $\pm$ 3.65	<u>59.86</u> $\pm$ 1.57	<b>62.99</b> $\pm$ 0.95

Table 2: Accuracy performance evaluation of three LLMs on four challenge benchmarks. We experiment with various LLM prompting techniques: Zero-Shot (ZS), Few-Shot (FS), Zero-Shot-CoT (ZS-CoT), and Few-Shot-CoT (FS-CoT). We apply these techniques either with full options or within our reduction-then-pairwise-comparison framework. Our comparison algorithm, PC-CoT, is used within our framework. For gpt-3.5-turbo, we employ 3-shot for fewshot demonstrations, while for LLaMA2 and Qwen, we use 5-shot. The best results are highlighted in **bold**, and the second best results are underlined.

detailed in Appendix A.2. Our experiments primarily utilized gpt-3.5-turbo alongside open-source models like LLaMA2-Chat and Qwen-Chat series. Here we use the instruction versions of models for better instruction following ability. We set the temperature to 0 for all models across different experiments.

#### 4.1 Evaluation For the Reduction Stage

**Setups.** For evaluating the performance of the reduction stage and our self-reduction algorithm, we configure the size of the reduced set as  $|\mathcal{R}| = 5$  and report the HIT@5 metric on the benchmark. This metric measures the proportion of instances where the correct label is among the top 5 labels chosen by the label reducer. We compare the effectiveness of our two techniques, ITR and CBWR, against the following baselines: (1) standard prompt, manually prompts the LLM to select the top 5 labels in a single call. (2) self-consistency (Wang et al., 2023), which invokes the aforementioned prompt multiple times, followed by a voting process to determine the final results. Appendix B provides additional details about their prompts.

**Results.** Table 1 provides a detailed comparison of HIT@5 performance. It is obvious that both self-reduction strategies ITR and CBWR consistently outperform the standard prompt and the self-consistency method by a relatively large margin. The self-consistency approach also exhibits an advantage over the standard prompt for gpt-3.5-turbo

and Qwen-72B. However, it may degrade the performance of the LLaMA2-70B model, which could result from the differences in training data. Another observation is significantly low reduction performance occurs on the LIU54 dataset. This could be attributed to a large proportion of short sentences in LIU54, making them challenging to interpret and reason effectively, more details are listed in Appendix E. Despite the above challenges, our strategies can consistently improve the reduction performance, minimizing error propagation to the next stage.

#### 4.2 Evaluation For Overall Framework

**Setups** We primarily use accuracy as the performance metric for our framework, considering both the reduction and pairwise comparison stages. The selected baselines include common prompting techniques: zero-shot (ZS), zero-shot-CoT (ZS-CoT), few-shot (FS), and few-shot-CoT (FS-CoT). We evaluate each method in two scenarios: with and without the reduction stage, to study the impact of the overall framework and compare the pairwise comparison algorithm. In the former scenario, we prompt the LLM directly with the full set of options, while in the latter, we conduct comparison experiments based on the CBWR result for gpt-3.5-turbo and Qwen, and the ITR result for LLaMA2. It is important to note that for the latter scenario, we compare four prompting methods and our proposed PC-CoT using the same reduction technique, to en-

sure a fair evaluation for the comparison stage. For the few-shot method, if the text exceeds the context length of 4096, we randomly omit demonstrations to reduce the context length. To avoid random errors, we conduct the experiments on each dataset five times and present the results in terms of mean and standard variance. Appendix C provides further details regarding the prompts.

**Results** As shown in Table 2, we have the following observations:

(1) Our framework significantly improves the stability and predictability of option selection in ZS, ZS-CoT, FS, and FS-CoT settings. Specifically, it boosts the zero-shot performance of gpt-3.5-turbo relatively by 11.88%, the few-shot performance of LLaMA2 from 26.45% to 54.01% (see Section 4.4 for detail analysis), and the FS-CoT performance of Qwen is relatively improved by 10.95%. These results demonstrate the effectiveness of our framework in mitigating ambiguity and bias.

(2) Our PC-CoT paired with our framework stands out with the best results for gpt-3.5-turbo and Qwen. Notably, it achieves a relative increase of 54.1% in average accuracy over a directly selected answer for gpt-3.5-turbo, from 40.24% to 62.01%. However, for LLaMA2-70B, our performance is slightly lower than the FS method on the Banking77 and Clinc150 benchmarks. We hypothesize this may have some relationship with LLaMA2 model training data, which makes it insensitive to the chain-of-thought guidance related to classification problems because we can observe that FS-CoT demonstrates a degradation in performance, dropping from 57.55% under FS to 39.18%.

(3) Compared to original CoT techniques, our PC-CoT demonstrates better generality across different types and sizes of LLMs. While the FS-CoT method exhibits satisfactory performance with the Qwen model, it deteriorates the performance of LLaMA and even hinders the zero-shot performance of gpt-3.5-turbo on Clinc150 and LIU54. This result suggests that the prompting method in classification scenarios requires tailored adjustments to effectively prioritize comparison for ambiguous boundaries.

The detailed cases of all nine prompting methods in Table 2 can be found in Appendix F.

### 4.3 Ablation Study of PC-CoT algorithm

As outlined in Table 3, we demonstrate the impact of removing key components from our PC-

Dataset	PC-CoT	w/o S	w/o D	w/o SD
Banking77	<b>61.82</b>	51.53	50.18	47.34
HWU64	<b>65.09</b>	52.44	61.11	59.01
LIU54	<b>37.35</b>	30.31	31.34	33.37
Clinc150	<b>83.78</b>	75.68	80.22	74.65
Average	<b>62.01</b>	52.49	55.71	52.59

Table 3: The ablation study for our PC-CoT algorithm (with gpt-3.5-turbo). We evaluate the performance under various conditions: remove the similar analysis prompt (w/o S), remove the difference analysis prompt (w/o D), and remove both similar and difference analysis prompt (w/o SD).

CoT. Without the similarity or difference analysis prompt, we observe an evident performance decline (approximately 15%). Furthermore, using only the difference analysis prompt leads to inferior results, supporting our belief that we need to discard the confusable points to better extract the distinctive features effectively. When both similarity and difference analyses are removed, our method reverts to plain FS-CoT, thus exhibiting a performance level similar to FS-CoT.

### 4.4 Debias Study

We find that our framework significantly improves the ZS-CoT and FS performance of LLaMA2 in Table 2. This enhancement is particularly evident in the reduced bias effect of certain tokens, such as “movies” and “convert” in the LIU54 dataset, and “iot/hue\_lightchange” and “recommendation/movies” in the HWU64 dataset. By visualizing the confusion matrices for both datasets (Figure 5), we can observe a noticeable shift in preferences after applying our framework, demonstrating its effectiveness. This improvement can be attributed to our reduction technique and pairwise comparison approach, which reduce the number of tokens and positions considered, thereby mitigating noise interference and enhancing the capabilities of few-shot demonstration and CoT reasoning.

### 4.5 Efficiency Analysis

Our method shows promising results but may require more LLM calls due to multiple iterations for reduction and pairwise comparisons to enhance accuracy and reduce bias. To assess the applicability of our approach, we compare our inference speed and call frequency with other methods in Table 4. Note that all baseline methods require multiple calls to LLMs, except for partial methods using the full option. However, the full option

Method	Full Options FS	P + FS	SC + FS	SC + FS-CoT	ITR + FS	CBWR + FS	CBWR+PC-CoT
#Call Avg.	1	5	12	24	9	8	12
Time / 1000 items	2min2s	11min32s	28min1s	53min1s	18min55s	14min21s	25min56s
Performance Gain	0	-0.39	+0.21	+0.19	+3.89	+5.01	+8.74

Table 4: LLM Calls Comparison

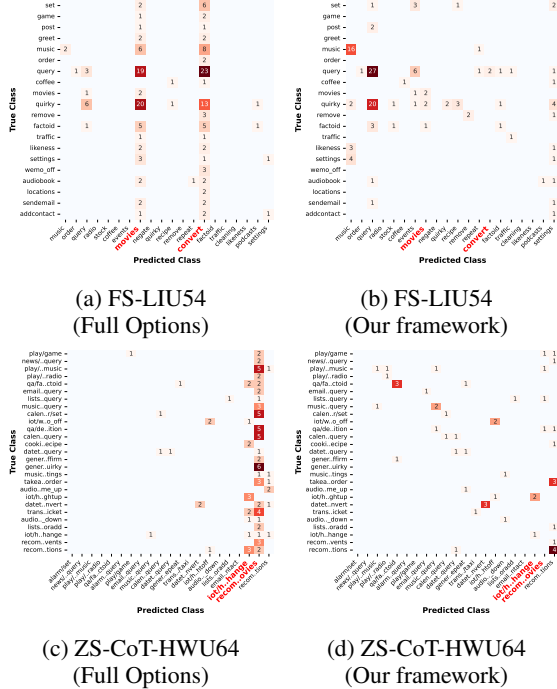


Figure 5: The confusion matrix visualization illustrates the debiasing effect of our framework on LLaMA70B. In the original few-shot setting, the model showed a preference for “movies” and “convert”, while in the ZS-CoT setting, it preferred “iot/hue\_lightchange” and “recommendation/movies”. After applying our framework, the preference bias clearly shifted.

Method	Banking77	HWU64	LIU54	Clincl50
Retrieval-Augmented ICL	56.25	52.66	35.2	78.73
Ours	61.83	65.09	37.35	83.78

Table 5: Comparison with in-context learning baselines

struggles with boundary ambiguity and inherent biases, leading to low accuracy and practical deployment difficulties. Our approach significantly outperforms other baselines by adding only a limited number of calls and achieves the highest performance gain (as shown in the last line of Table 4). In contrast, self-consistency-based methods require more calls than our approach but offer limited performance improvements. With the assistance of technologies such as flash-attn and vLLM, we have successfully reduced the computational overhead of our method when calling the 70B model. As a result, we can accurately predict 1,000 pieces of data in just half an hour, which is significantly

Method	reminder	agnews	bbcnews	restaurant
FS-CoT	78.79	81.31	92.19	86.98
Ours	79.69	82.83	92.25	87.37

Table 6: Performance on simple text classification tasks

faster than the multiple-call methods used in Self-Consistency. Moreover, the future of technology is headed towards large-scale models. It is foreseeable that with technological advancements, the costs associated with API calls and model computations will decrease further. Our method, with its enhancements addressing bias and boundary issues, will become increasingly relevant and significant in solving real-world classification problems. Therefore, we believe that the scaling trend of language models will elevate the importance of our method in the future.

#### 4.6 In Context Learning

We experiment with a similar in-context learning baseline (Milios et al., 2023) that also uses clustering algorithms in a few-shot setting. As shown in Table 5, our approach consistently outperforms the baseline. The main reason is that clustering similar demonstrations can blur the model’s decision boundary when options are highly confusable, reducing performance. In contrast, our method introduces diversity through ITR/CBWR and provides a more detailed comparison via PC-CoT.

#### 4.7 Simple Text Classification Tasks

Our framework is applicable to simple text classification tasks and has comparable performance to the few shot chain-of-thought method, as shown in Table 6. In our primary experiments, we focused on classification tasks with over 50 categories to demonstrate our framework’s ability to handle unclear decision boundaries arising from numerous options and similar, easily confused elements.

#### 4.8 Lightweight Models

As shown in Table 7, our method still applies to smaller models. While the 14B model outperforms other baselines, the Qwen 1.8B model slightly underperforms compared to the Few-Shot CoT methods. Because smaller models (<10B) may struggle



Model	Dataset	Full Options				Our Framework				
		ZS	ZS-CoT	FS	FS-CoT	ZS	ZS-CoT	FS	FS-CoT	PC-CoT
Qwen-1.8B-Chat	Banking77	18.18	17.53	4.87	24.68	26.62	23.7	19.16	<b>31.82</b>	19.16
	HWU64	21.82	22.73	4.55	15.45	24.55	26.36	20.04	<b>29.09</b>	23.64
	LIU54	17.32	17.32	0.79	18.11	23.62	<b>20.47</b>	17.72	20.08	17.32
	Clinic150	30.22	33.78	5.78	20.89	39.11	33.78	35.56	<b>54.22</b>	25.78
Qwen-14B-Chat	Banking77	33.44	33.12	38.64	46.43	35.06	26.62	51.95	46.10	<b>56.82</b>
	HWU64	40.01	41.82	47.27	50.02	37.27	38.18	55.45	53.64	<b>62.73</b>
	LIU54	29.53	27.95	31.50	27.95	31.10	25.98	37.40	32.28	<b>38.19</b>
	Clinic150	53.78	52.44	58.22	67.11	55.56	42.22	77.78	73.78	<b>82.67</b>

Table 7: Performance Comparison of LightWeight LLMs

with tasks such as instruction following and generating high-quality few-shot examples, as highlighted in prior research (Fu et al., 2023).

## 5 Related Work

**Text Classification** Traditional text classification approaches typically involve feature extraction by RNN (Xie et al., 2020), GCN (Lin et al., 2021), or Transformers (Chai et al., 2020). These features are subsequently input into a classifier head and assigned the corresponding label. To deal with the common data scarcity issue, the meta-learning paradigm emerges to identify novel categories with minimal labeled examples by leveraging prior knowledge acquired from known categories (Chen et al., 2022; Gao et al., 2021; Jiang et al., 2023; Liu et al., 2023; Sung et al., 2023). Nonetheless, these approaches require supervised fine-tuning and exhibit limited applicability. The powerful capabilities of modern LLMs enable zero-shot or few-shot classification without necessitating model update (Fei et al., 2023; Loukas et al., 2023). Despite these advancements, the exploration of LLM-based classification methodologies in addressing practical challenges, such as coping with ambiguous boundaries and inherent bias, remains underexplored.

**Chain of Thought** Wei et al. (2022b) introduces few-shot chain-of-thought, enriching LLM reasoning capabilities by incorporating intermediate reasoning steps in each demonstration. Kojima et al. (2022) highlights the efficacy of a straightforward prompt, “Let’s think step by step”, in enhancing zero-shot reasoning. While recent investigations have explored diversifying demonstrations (Zhang et al., 2022) or modeling the reasoning process using tree and graph structures (Yao et al., 2023; Lei et al., 2023a), the realm of classification problems continues to lack specialized CoT methods to com-

bat ambiguity and bias. THOR (Fei et al., 2023) exclusively addresses sentiment analysis, sequentially inducing aspects and opinions for sentiment identification. CARP (Sun et al., 2023) focuses on simple classification problems with only 2-4 options and utilizes supervised fine-tuning of the model, which is impractical in more complex classification scenarios. In contrast, we focus on handling a large number of label options and subtle semantic differences. We also address the migration of preferences for tokens and positions, enabling seamless integration into a wide range of text classification tasks.

## 6 Conclusions

In this paper, we investigate the vulnerability of LLMs in terms of option quantity and arrangement for text classification tasks. Our empirical analysis reveals that the main challenges lie in their ambiguous decision boundaries and inherent biases towards tokens and positions. To mitigate these challenges, we propose a novel two-stage classification framework for LLMs, involving reduction and comparison stages. This framework strategically employs pairwise comparisons to address ambiguity and bias. Additionally, we introduce innovative self-reduction algorithms, including ITR and CBWR strategies, and our PC-CoT, a contrastive chain-of-thought technique, to enhance the framework’s utility. Extensive experiments demonstrate our framework can bring improvement to the stability and debiasing for various prompting techniques. Our ITR and CBWR methods prove highly effective for reduction, while our PC-CoT exhibits superior performance in classification accuracy.

## Limitations

Though our method significantly outperforms other baselines by adding only a limited number of calls

and achieves the highest performance gain, it does increase the number of LLM calls. For some smaller-scale LLMs, their limited reasoning ability can impede precise reduction and comparisons, potentially diminishing the effectiveness of our method.

## Ethics Statement

In this study, the classification datasets and LLM responses used in our experiments only serve the specific purposes of evaluating our proposed approach. To address ethical considerations, we have incorporated a rule-based filter during our challenge set sampling process to exclude offensive and sensitive information. Our study investigates position bias and token bias during the classification process. Our algorithms help LLMs focus more on the actual content of options rather than making shortcut predictions based on preferences for specific tokens and positions, which is crucial for debiasing LLMs.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant No. 62276110, No. 62172039 and in part by the fund of Joint Laboratory of HUST and Pingan Property Casualty Research (HPL). There are computational tasks are conducted using the HPC platform of Huazhong University of Science and Technology. The authors would also like to thank the anonymous reviewers for their comments on improving the quality of this paper.

## References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. [Qwen technical report](#).

Arth Bohra, Govert Verkes, Artem Harutyunyan, Pascal Weinberger, and Giovanni Campagna. 2023. Byoc: Personalized few-shot classification with co-authored class descriptions. In *Findings of the Association*

*for Computational Linguistics: EMNLP 2023*, pages 13999–14015.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45.

Duo Chai, Wei Wu, Qinghong Han, Fei Wu, and Jiwei Li. 2020. [Description based text classification with reinforcement learning](#). In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1371–1382. PMLR.

Junfan Chen, Richong Zhang, Yongyi Mao, and Jie Xu. 2022. [Contrastnet: A contrastive learning framework for few-shot text classification](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 10492–10500. AAAI Press.

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

Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430. PMLR.

Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830.

- Weisen Jiang, Yu Zhang, and James Kwok. 2023. Effective structured prompting by meta-learning and representative verbalizer. In *International Conference on Machine Learning*, pages 15186–15199. PMLR.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Maxime Labonne and Sean Moran. 2023. [Spam-t5: Benchmarking large language models for few-shot email spam detection](#).
- Stefan Larson, Anish Mahendran, Joseph J Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K Kummerfeld, Kevin Leach, Michael A Laurenzano, Lingjia Tang, et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1311–1316.
- Bin Lei, pei Hung Lin, Chunhua Liao, and Caiwen Ding. 2023a. [Boosting logical reasoning in large language models through a new framework: The graph of thought](#).
- Shuo Lei, Xuchao Zhang, Jianfeng He, Fanglan Chen, and Chang-Tien Lu. 2023b. [TART: Improved few-shot text classification using task-adaptive reference transformation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11014–11026, Toronto, Canada. Association for Computational Linguistics.
- Muyang Li, Runze Wu, Haoyu Liu, Jun Yu, Xun Yang, Bo Han, and Tongliang Liu. 2023. Instant: Semi-supervised learning with instance-dependent thresholds. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Yuxiao Lin, Yuxian Meng, Xiaofei Sun, Qinghong Han, Kun Kuang, Jiwei Li, and Fei Wu. 2021. Bertgcn: Transductive text classification by combining gnn and bert. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1456–1462.
- Han Liu, Feng Zhang, Xiaotong Zhang, Siyang Zhao, Fenglong Ma, Xiao-Ming Wu, Hongyang Chen, Hong Yu, and Xianchao Zhang. 2023. [Boosting few-shot text classification via distribution estimation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):13219–13227.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. Benchmarking natural language understanding services for building conversational agents. In *Increasing Naturalness and Flexibility in Spoken Dialogue Interaction: 10th International Workshop on Spoken Dialogue Systems*, pages 165–183. Springer.
- Lefteris Loukas, Ilias Stogiannidis, Prodromos Malakiotis, and Stavros Vassos. 2023. [Breaking the Bank with ChatGPT: Few-Shot Text Classification for Finance](#).
- Aristides Milios, Siva Reddy, and Dzmitry Bahdanau. 2023. [In-context learning for text classification with many labels](#). In *Proceedings of the 1st GenBench Workshop on (Benchmarking) Generalisation in NLP*, pages 173–184, Singapore. Association for Computational Linguistics.
- OpenAI. 2022. Chatgpt.
- Pouya Pezeshkpour and Estevam Hruschka. 2023. [Large language models sensitivity to the order of options in multiple-choice questions](#).
- Xiaoye Qu, Zhikang Zou, Yu Cheng, Yang Yang, and Pan Zhou. 2019. Adversarial category alignment network for cross-domain sentiment classification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2496–2508.
- Xiaofei Sun, Xiaoya Li, Jiwei Li, Fei Wu, Shangwei Guo, Tianwei Zhang, and Guoyin Wang. 2023. [Text classification via large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8990–9005, Singapore. Association for Computational Linguistics.
- Mujeen Sung, James Gung, Elman Mansimov, Nikolaos Pappas, Raphael Shu, Salvatore Romeo, Yi Zhang, and Vittorio Castelli. 2023. [Pre-training intent-aware encoders for zero- and few-shot intent classification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10433–10442, Singapore. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas

- Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#).
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. 2020. [Unsupervised data augmentation for consistency training](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Adam Zaremba and Ender Demir. 2023. Chatgpt: Unlocking the future of nlp in finance. *Modern Finance*, 1(1):93–98.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*.

## A More Detail of Preliminary Study And Dataset

### A.1 Full result for Figure 3

In Figure 6a, we observe that across various LLM sizes and types, including LLaMA, Qwen, and gpt-3.5-turbo, we consistently observe a decline in performance as the number of options increases. Different LLMs may exhibit varying degrees of performance downgradations in response to the increased number of options, with the gpt-3.5-turbo series demonstrating heightened stability. The challenge persists even with longer contextual LLMs, though they may show a more gradual performance decline. This observation suggests that the noise stemming from the proliferation of options remains a significant obstacle in the landscape of LLMs.

In Figure 6b, we discern a significant bias caused by the position of the correct choice, especially evident in open-source models such as LLaMA and Qwen, resulting in performance fluctuations exceeding 50%. In contrast, gpt-3.5-turbo demonstrates minimal effects, with fluctuations below 10%, highlighting its superiority over open-source counterparts. This lack of robustness can be attributed to the inherent position bias inherent in LLMs, as elucidated by Pezeshkpour and Hruschka (2023). The bias in option positions becomes particularly apparent when the span of option choices is extensive, as the prolonged context amplifies the imbalance in position, presenting challenges for maintaining stable performance in classification problems.

In Figure 7, we show the full confusion matrix from Figure 3c. A preference for “movies” and “convert” tokens can be clearly observed regardless of the golden label. Similar token biases are also observed for “iot/hue\_lightchange” and “recommendation/movies” for LLaMA on the HWU64 benchmark, or neglect the “music”, “query”, and “quirky” for the Qwen model.

The long-context ability of LLM is insufficient in managing an increasing number of options or their respective positions. Despite the common intuition that extended context should aid generative models in handling larger label sets, we find that performance diminishes with a growing number of labels, coupled with a persistent bias towards specific option positions. This is exemplified by comparisons such as gpt-3.5-turbo-1106 (16k context length) versus gpt-3.5-turbo-0613 (4k context length) and LongAlpaca (32k context length) ver-

sus LLaMA-13B (4k context length). Although a longer context does impart stability and smoothness in the face of increased options, challenges persist concerning both the number and positioning of options. the LongAlpaca-32k also suffers from 50% performance fluctuation when changing the golden option to 50<sup>th</sup> position.

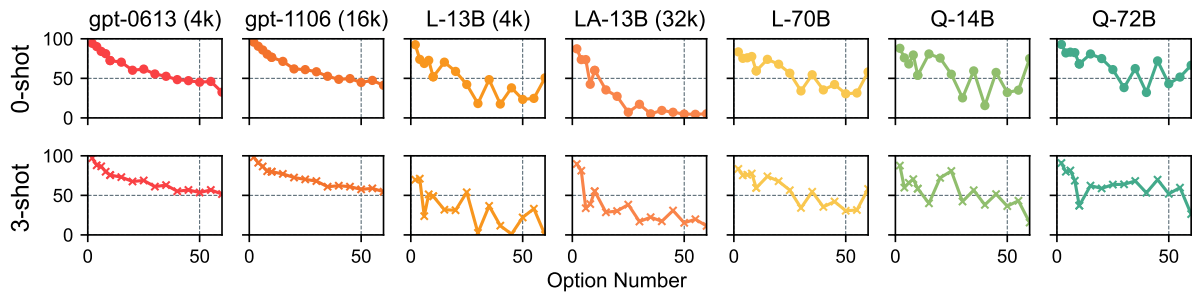
The In-context examples may reduce option bias for large models while altering the distribution. Examining in-context examples, we observe a reduction in bias for larger models like LLaMA-2-70B and Qwen-72B when moving from zero-shot to 3-shot demonstrations, while this effect is less pronounced for smaller models. Additionally, we note distinct bias distributions between 3-shot and zero-shot scenarios, suggesting that in-context examples may alter options position bias distributions. In contrast, for experiments involving varying numbers of options, the trends observed in few-shot and zero-shot scenarios remain essentially consistent.

### A.2 Challenge Set Sampling

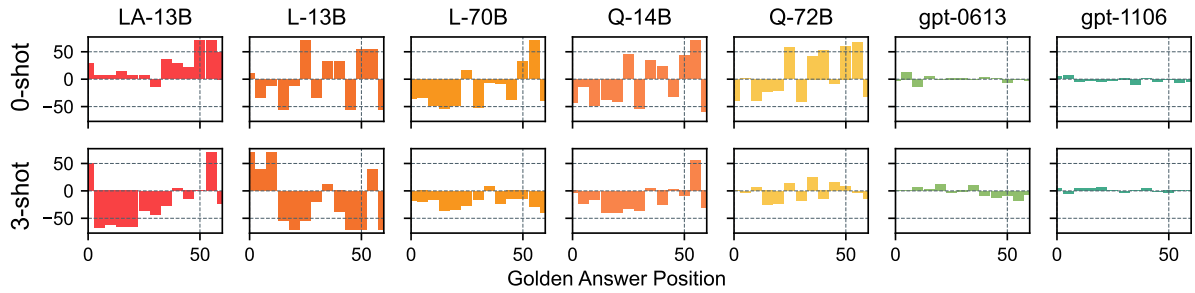
To align with real-world application scenarios, we employed a sampling strategy for a more challenging subset from these datasets (only their original test splits) by the precision margin (Li et al., 2023),  $M(f(x)) := \max(0, \max_y f_y(x) - \max_{k \neq y} f_k(x))$ , which is an indicator of the confidence of the base classifier  $f$  (the supervised finetuned BERT model). *i.e.*, the difference between the largest and second-largest label probability for every data item in the test dataset. A smaller prediction margin correlates with decreased confidence in the prediction, representing a more challenging setting. As for the training dataset, we randomly sample five items from their training splits as our few-shot demonstrations.

### A.3 Supervised Finetuning Performance on the challenge set

To better illustrate how challenging our curated dataset is, we select a series of encoder-only models and decoder-only LLMs for supervised fine-tuning (SFT) in a few-shot setting. For the encoder-only series, we fine-tune them for 30 epochs. For the decoder-only series, due to the huge SFT cost, we only select TinyLLaMA-1.1B and Qwen-1.8B. We employ an alpaca format for fine-tuning (with the following prompt). The details of the results are listed in Table 8.



(a) LLMs classification accuracy concerning the number of options consistently demonstrates a decline.



(b) The change rate of classification accuracy after moving the correct option to a fixed position compared to the original accuracy. This simple option arrangement can lead to dramatic performance fluctuations in both zero-shot and few-shot scenarios.

Figure 6: The Full result for Figure 3a and Figure 3b

Shot	Dataset	BERT-base	RoBERTa-base	BERT-Large	RoBERTa-large	Llama-1.1B	Qwen-1.8B
K=1	Banking77	1.95	1.30	2.60	4.22	28.90	53.25
	Clinic150	1.33	0.89	1.78	6.67	45.78	60.44
	LIU54	2.76	4.33	7.09	4.33	24.02	39.37
	HWU64	5.45	1.82	0.91	3.64	16.36	53.64
	Avg.	2.87	2.08	3.09	4.71	28.76	51.67
K=5	Banking77	0.32	6.49	1.30	0.32	40.58	70.45
	Clinic150	7.56	36.44	15.11	47.56	68.00	79.56
	LIU54	8.66	16.14	9.84	14.17	29.92	50.39
	HWU64	6.36	18.18	21.82	32.73	47.27	65.45
	Avg.	5.73	19.32	12.02	23.70	46.44	66.46

Table 8: Accuracy performance evaluation of encoder-only and decoder-only models.

### FineTuning Prompt

Below is an instruction that describes a task. Write a response that appropriately completes the request.

### Instruction:

Given the sentence: "{text}", please select the most possible topic from the following OPTIONS: {options}.

### Response:

and fine-tuning them for 30 epochs. With only five demonstrations, BERT achieves a mere 12.02% accuracy, while RoBERTa fares slightly better at 23.70%. In contrast, decoder-only language models, particularly Qwen-1.8B, demonstrate better performance than BERT and RoBERTa. However, this superior performance is contingent upon resource-intensive fine-tuning processes and may suffer from limited generalizability, thereby constraining their practical applicability.

Table 8 indicates that the challenge datasets chosen pose significant difficulties for both BERT and RoBERTa, even when utilizing their larger versions



## B Detail of the Reduction Stage

### Standard Reduction Prompt

Consider the sentence: "{text}"  
Please select {top\_k}most possible topic  
from following OPTIONS: {options}.  
CHOICE:

The self-consistency, our ITR, and CBWR techniques employ the same reduction prompt as exemplified above. The specific procedural steps for CBWR are delineated in Algorithm 2.

---

### Algorithm 2 Self-Reduction (CBWR)

---

**Require:** Input sentence  $x$ , full options set  $\mathcal{Y} = \{y_1, y_2, \dots, y_N\}$ , LLM  $p_\theta$ , selection prompt Sel, selection number  $N$ , step limit  $T$

- 1: Initialize the selection set  $\mathcal{S} = \mathcal{Y}$
- 2: **for**  $t = 1$  to  $T$  **do**
- 3:     Cluster the option set  $\mathcal{S}$  into  $K$  clusters
- 4:      $W \rightarrow \{\}$
- 5:     **for** each cluster  $i$  from 1 to  $K$  **do**
- 6:         Randomly select options from cluster  $i$
- 7:         add them into the window  $W$
- 8:     **end for**
- 9:     Select candidate set  $C$  from  $W$ :  
        $C \leftarrow \text{Top-N}(\{y \mid p_\theta^{\text{sel}}(y \mid x, W)\})$
- 10:     Determine the discard option set  $\mathcal{D} = W \setminus C$
- 11:     Update the remaining candidates  $\mathcal{S} \leftarrow \mathcal{S} \setminus \mathcal{D}$
- 12: **end for**
- 13: **return** the remaining selection set  $\mathcal{S}$

---

## C Detail of the Comparison Stage

### C.1 Baseline Using Full Options

In this section, we list the prompt used in baselines without label reduction, *i.e.*, the LLM is prompted to directly choose the correct label from the complete set of options.

### Zero-Shot Prompt

Given the sentence: "{text}"  
Please select the most possible topic from  
the following OPTIONS: {options}  
CHOICE:

### Zero-Shot-CoT Prompt

Given the sentence: "{text}"  
Please select the most possible topic from  
the following OPTIONS: {options}  
Let's think step by step and give your  
explanation to verify your answer:

### Few-Shot Prompt

Below is a text classification problem,  
Note that you can only select the label in  
{options}  
SENTENCE: {text1}  
LABEL: {label1}

...

SENTENCE: {text}  
LABEL:

### Few-Shot-CoT Prompt

Below is a text classification problem,  
Note that you can only select the label in  
{options}. Let's think step by step and give  
your explanation to verify the answer.  
SENTENCE: {text1}  
EXPLANATION: {explain1}  
LABEL: {label1}

...

SENTENCE: {text}  
LABEL: {label}  
EXPLANATION:

Notice that we obtain the explanation demonstrations for Few-Shot-CoT by prompting LLM to automatically generate the explanation according to the golden label without human annotation. The prompt is as follows:

### Few-Shot-CoT Explanation Generation

Below is a text classification problem.  
Let's think step by step and give your  
explanation to verify the SENTENCE label:  
SENTENCE: Fears for T N pension after talks  
Unions representing workers at Turner Newall  
say they are 'disappointed' after talks with  
stricken parent firm Federal Mogul.  
LABEL: Business  
EXPLANATION: The statement discusses talks  
between unions and a parent firm, which  
relates to business-related negotiations and  
concerns regarding pensions.

SENTENCE: {text}  
EXPLANATION:

### C.2 Baselines with Pairwise Reduction

Here, we present the pairwise prompt used in baselines with label reduction. The LLMs are instructed to compare options in pairs within the refined set obtained from the first stage.



### Zero-Shot Pairwise Prompt

Which term is more likely to represent the topic of "{text}" - "{label1}" or "{label2}"?

### Zero-Shot-CoT Pairwise Prompt

Which term is more likely to represent the topic of "{text}" - "{label1}" or "{label2}"?

Let's think step by step and give your explanation to verify your answer:

### Few-Shot Pairwise Prompt

Below is a text classification problem, please complete the sentence by "{label1}" or "{label2}":

SENTENCE: {text1}

LABEL: {label1}

...

SENTENCE: {text}

LABEL:

### Few-Shot-CoT Pairwise Prompt

Below is a text classification problem. Let's think step by step and give your explanation to verify which term is more likely to represent the label of the sentence - "{label1}" or "{label2}":

SENTENCE: {text1}

EXPLANATION: {explain1}

LABEL: {label1}

...

SENTENCE: {text}

EXPLANATION:

## C.3 Our PC-CoT Prompt

Our proposed contrastive elimination methodology, PC-CoT, consists of three key stages: similarity analysis, difference analysis, and final decision-making. Specific system instructions and prompts for the three stages are provided below.

### Instruction

Below is a text classification problem:

SENTENCE: {text1}

LABEL: {label1}

SENTENCE: {text2}

LABEL: {label2}

### Similarity Analysis

The phrases can often be mistaken for "{label1}" and "{label2}", due to certain shared characteristics.

SHARED ASPECTS:

### Difference Analysis

Next, diligently contrast the deviations between these two topics, putting aside the mentioned shared characteristics. Concisely explain, what is the key element that sets them apart?

CONTRASTING POINTS:

### DecidePrompt

After scrutinizing the presented SHARED ASPECTS and CONTRASTING POINTS, which term - "{label1}" or "{label2}" - would be a more accurate representation for the label of {text}? Provide the final label in the format "LABEL: a".

## D Detail of Decision Boundary Ambiguity

The silhouette score is a metric used to quantify the compactness and separation of clusters within a dataset. In the context of our study on text data classification, the silhouette score provides insight into the overlap between different class samples. It is calculated for each data point using the formula:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}} \quad (3)$$

Here,  $a(i)$  is the average distance of a data point  $i$  to other points within the same cluster, and  $b(i)$  is the smallest average distance to points in a different cluster. The overall silhouette score for the dataset is the average of these individual scores, denoted as  $S$ :

$$S = \frac{1}{N} \sum_{i=1}^N s(i) \quad (4)$$

where  $N$  is the total number of data points. So A higher silhouette score indicates clearer clusters and distinct separation suggesting a more evident decision boundary and facilitating classification. Conversely, a lower score suggests increased overlap between class boundaries, indicating a more ambiguous decision boundary.

Figure 9 conducts an ablation study on the number of categories. From left to right, we add a new category (in dark blue), leading to increased

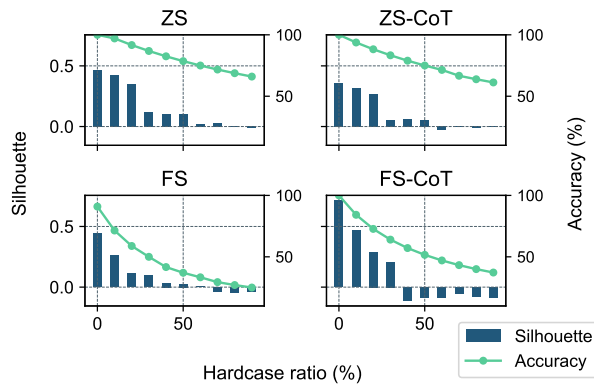
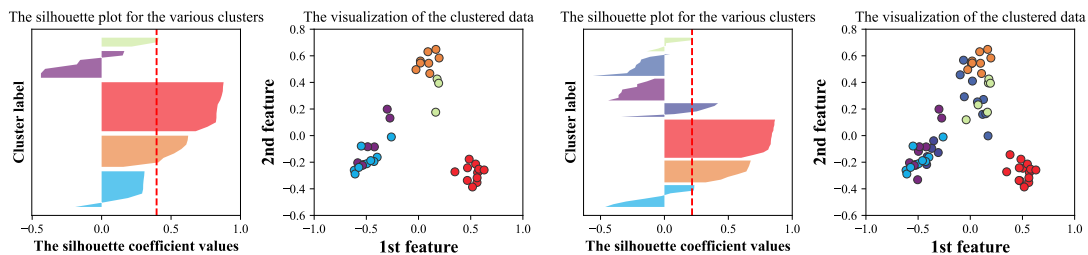


Figure 8: Demonstration sampling methods comparison.



(a) A higher silhouette score indicates clear boundaries between classes (b) A lower silhouette score suggesting overlapping classes.

Figure 9: Visualization of silhouette scores and decision boundaries

overlap with existing categories (particularly orange, green, and light blue), resulting in ambiguous boundaries. This overlap is reflected in the decreasing silhouette score, highlighted by the red line. Notably, Figure 8 demonstrates a clear positive correlation between silhouette score and accuracy, motivating our introduction of a self-reduction stage to decrease the number of categories.

## E The LIU54 Dataset

The challenge posed by the LIU54 dataset, as evidenced by the difficulty encountered by even advanced language models like gpt-3.5-turbo and Qwen-72B, underscores a significant obstacle in natural language understanding. Despite the formidable capabilities of these models, they struggle to achieve satisfactory performance, often achieving accuracy rates of only around 30% using conventional baselines. A closer examination of LIU54 reveals that its difficulty primarily stems from the brevity of its sentences, which hinders effective interpretation and reasoning. This observation underscores the importance of enhancing the capacity of large language models to capture subtle semantic nuances, particularly in the context

of classification tasks. Overall, this highlights a critical area for improvement in the development of language understanding technologies.

### Semantic incompleteness

sentence: search  
golden: factoid  
prediction: query

sentence: Investing  
golden: quirky  
prediction: stock

### Label mismatch

sentence: What's a parsec?  
golden: factoid  
prediction: definition

sentence: how we calculate  
golden: sendemail  
prediction: maths

sentence: Add that song to my playlist.  
golden: likeness  
prediction: Music

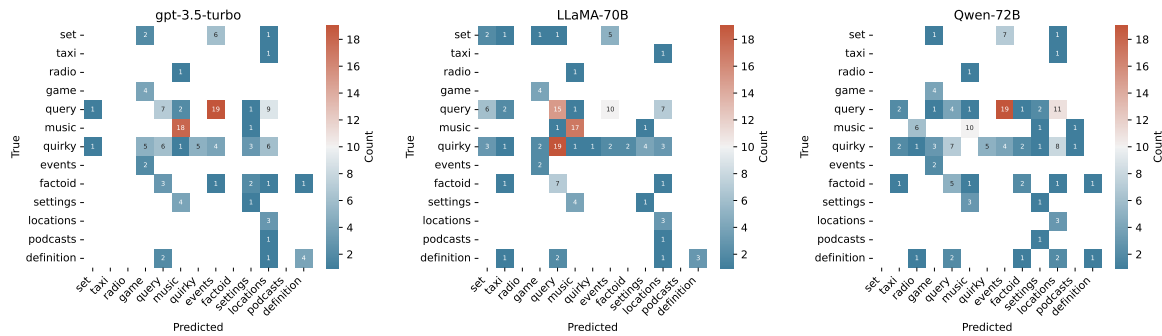


Figure 10: Visualization of bad cases of three LLMs on LIU54.

**Multi-label**

**sentence:** Set a notification for sports game

**golden:** set

**prediction:** game

**sentence:** Is the brightness of my screen running low?

**golden:** quirky

**prediction:** hue\_lighton

**sentence:** What time does the downtown parade start?

**golden:** query

**prediction:** events

We analyzed failed predictions by various Language Model (LLM) architectures using a confusion matrix (see Fig. 10) and presented selected failed cases above. Our findings reveal striking similarities in the confusing classes identified by the three LLMs. We categorize the failed cases into three overarching types: (i) Semantic Incompleteness, (ii) Label Mismatch, and (iii) Text Contents of Multiple Classes.

Semantic incompleteness refers to grammatically incomplete sentences typically comprised of a few words, such as 'search', 'Investing', 'Bingo', etc. Label mismatch occurs when the assigned label from the original dataset is inappropriate; for example, 'What's a parsec?' is more fittingly categorized as "definition" rather than "factoid". In cases of multi-label content, sentences express multiple aspects that satisfy multiple categories concurrently. For instance, 'What time does the downtown parade start?' fulfills both the 'query' and 'events' class criteria simultaneously.

Based on our analyses, we posit that confusing labels within the test set of the LIU54 Dataset may introduce bias into results. However, it also

presents an opportunity for evaluating the compatibility of different LLMs, thus serving as a challenging benchmark task.

**F Case Study**

We present the complete set of options and detailed cases for the pairwise comparison framework, ranging from Table 9 to Table 17, along with additional reduction cases outlined in Table 18.

---

**ZERO-SHOT**

---

**Q:** Given the sentence: "So I just put my top-up into the card and it hasn't changed." Please select the most possible topic from the following options.

**OPTIONS:** pending\_top\_up, top\_up\_failed, topping\_up\_by\_card, card\_not\_working, balance\_not\_updated, ...

**CHOICE:**

**A:** The most possible topic for the given sentence is **topping\_up\_by\_card**. ✘

---

Table 9: Zero-shot reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**FEW-SHOT**

---

**Q:** Below is a text classification problem, Note that you can only select labels in pending\_top\_up, top\_up\_failed, topping\_up\_by\_card, card\_not\_working, balance\_not\_updated, ...

**SENTENCE:** How long will it take for my cash withdrawal to no longer show as pending?

**LABEL:** pending\_cash\_withdrawal

...

**SENTENCE:** So I just put my top-up into the card and it hasn't changed.

**LABEL:**

**A:** card\_not\_working ✘

---

Table 10: Few-shot reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**ZERO-SHOT COT**

---

**Q:** Given the sentence: "So I just put my top-up into the card and it hasn't changed." Please select the most possible topic from the following options.

**OPTIONS:** pending\_top\_up, top\_up\_failed, topping\_up\_by\_card, card\_not\_working, balance\_not\_updated, ...

Let's think step by step and verify your answer:

**A:**Based on the given sentence "So I just put my top-up into the card and it hasn't changed," the most relevant topic seems to be **top\_up\_failed**. This is because the user mentions putting a top-up into the card, and there is an indication that something has not changed as expected. ✘

---

Table 11: Zero CoT reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**FEW-SHOT COT**

---

**Q:** Below is a text classification problem, Note that you can only select labels in pending\_top\_up, top\_up\_failed, topping\_up\_by\_card, card\_not\_working, balance\_not\_updated, ... Let's think step by step and give your explanation to verify the answer.

**SENTENCE:** I topped up but it isn't in my account

**EXPLANATION:** The statement is about a user's concern regarding a top-up that has not been reflected in their account, which is related to a pending transaction or issue with their account balance.

**LABEL:** pending\_top\_up

...

**SENTENCE:** So I just put my top-up into the card and it hasn't changed.

**EXPLANATION:**

**A:** The sentence is about a user who has topped up their card with money, but the balance has not been updated. This issue is related to the balance not being updated after a bank transfer, which falls under the label balance\_not\_updated\_after\_bank\_transfer.

**LABEL:** **balance\_not\_updated** ✗

---

Table 12: Few-shot CoT reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**ZERO-SHOT PAIRWISE**

---

**Q:** Which term is more likely to represent the topic of "So I just put my top-up into the card and it hasn't changed." - "top\_up\_failed" or "pending\_top\_up"?

**A:** The term "**top\_up\_failed**" is more likely to represent the topic of "So I just put my top-up into the card and it hasn't changed." ✗

---

Table 13: Zero-shot pairwise reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**FEW-SHOT PAIRWISE**

---

**Q:** Below is a text classification problem, please complete the sentence by "top\_up\_failed" or "pending\_top\_up":

**SENTENCE:** I topped up but it isn't in my account

**LABEL:** pending\_top\_up

...

**SENTENCE:** So I just put my top-up into the card and it hasn't changed.

**LABEL:** "

**A:** **top\_up\_failed** ✗

---

Table 14: Few-shot pairwise reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**ZERO-SHOT COT PAIRWISE**

---

**Q:** Which term is more likely to represent the topic of "So I just put my top-up into the card and it hasn't changed." - "top\_up\_failed" or "pending\_top\_up"? Let's think step by step.

**A:**Based on the given statement, the term "**top\_up\_failed**" is more likely to represent the topic. This is because the statement suggests that the top-up process did not work as expected, indicating a potential failure in the top-up transaction. **✗**

---

Table 15: Zero-shot CoT pairwise reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**ZERO-SHOT COT PAIRWISE**

---

**Q:** Below is a text classification problem. Let's think step by step and give your explanation to verify which term is more likely to represent the label of the sentence - "top\_up\_failed" or "pending\_top\_up":

SENTENCE: Why has my top up failed?

EXPLANATION: The statement directly mentions a failed top-up, indicating that there was an issue with adding additional funds or credits to something. This is a specific problem related to a top-up process, hence the label "top\_up\_failed."

LABEL: top\_up\_failed

...

SENTENCE: So I just put my top-up into the card and it hasn't changed.

EXPLANATION:

**A:** The statement mentions putting a top-up into a card, but it also mentions that it hasn't changed. This suggests that there may be an issue with the top-up process, indicating a "top\_up\_failed" label.

LABEL: **top\_up\_failed ✗**

---

Table 16: Zero-shot CoT pairwise reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

---

**PC-COT (Ours)**

---

**Q:** Below is a text classification problem:

SENTENCE: My top-up didn't go through

LABEL: top\_up\_failed

SENTENCE: why isn't my top-up going through?

LABEL: pending\_top\_up

...

The phrases can often be mistaken for "top\_up\_failed" and "pending\_top\_up", due to certain shared characteristics.

SHARED ASPECTS:

**A:** - Mention of "top-up" or "top up"

- Mention of the action not going through or failing

- Expressions of confusion or concern about the status of the top-up

**Q:** Next, diligently contrast the deviations between these two topics, putting aside the mentioned shared characteristics. Concisely explain, what is the key element that sets them apart?

CONTRASTING POINTS:

**A:** The key element that sets apart the "top\_up\_failed" and "pending\_top\_up" topics is the specific focus of the phrases.

- "top\_up\_failed" phrases primarily focus on the fact that the top-up has failed or is not working. They express frustration or confusion about why the top-up did not go through and may ask for reasons or explanations for the failure.

- On the other hand, "pending\_top\_up" phrases primarily focus on the status or confirmation of the top-up. They express uncertainty or concern about whether the top-up has gone through or not. These phrases often ask for information about why the top-up is still pending or how to determine if it has been successful. In summary, the key element that sets them apart is the specific emphasis on the failure of the top-up in "top\_up\_failed" phrases and the uncertainty about the status of the top-up in "pending\_top\_up" phrases.

**Q:** After scrutinizing the presented SHARED ASPECTS and CONTRASTING POINTS, which term - "top\_up\_failed" or "pending\_top\_up" - would be a more accurate representation for the label of So I just put my top-up into the card and it hasn't changed.? Provide the final label in the format "LABEL: a".

**A:** Based on the shared aspects and contrasting points provided, the sentence "So I just put my top-up into the card and it hasn't changed" seems to align more closely with the "pending\_top\_up" label. This is because the sentence expresses uncertainty or concern about the status of the top-up, indicating that it hasn't changed and implying that the user is unsure whether the top-up has been successful or not. Therefore, "pending\_top\_up" would be a more accurate representation of the label in this case. LABEL: **pending\_top\_up** ✓

---

Table 17: PC-CoT reasoning example for the classification problem on the BANKING77 dataset using the gpt-3.5-turbo model.

<b>Sentence</b>	So I just put my top-up into the card and it hasn't changed.	
<b>Label</b>	pending_top_up	
<b>Reduce</b>	<u>CBWR</u> verify_top_up top_up_failed balance_not_updated_after_bank_transfer pending_top_up top_up_by_card_charge	<u>ITR</u> top_up_reverted top_up_by_card_charge top_up_failed balance_not_updated_after_bank_transfer topping_up_by_card
<b>Sentence</b>	I made a transfer and am still waiting.	
<b>Label</b>	balance_not_updated_after_bank_transfer	
<b>Reduce</b>	<u>CBWR</u> failed_transfer transfer_into_account balance_not_updated_after_bank_transfer pending_transfer transfer_not_received_by_recipient	<u>ITR</u> pending_transfer transfer_not_received_by_recipient transfer_timing transfer_fee_charged failed_transfer
<b>Sentence</b>	Where can I get my card at?	
<b>Label</b>	order_physical_card	
<b>Reduce</b>	<u>CBWR</u> card_delivery_estimate card_arrival getting_spare_card order_physical_card get_physical_card	<u>ITR</u> get_physical_card card_arrival activate_my_card card_delivery_estimate getting_virtual_card

Table 18: Examples of CBWR and ITR reduction techniques. The reduction process aims to retain a subset of options with high similarity from a large pool of options as much as possible.