

Quantifying LLM Biases Across Instruction Boundary in Mixed Question Forms

Anonymous ACL submission

Abstract

Large Language Models (LLMs) annotated datasets are widely used nowadays, however, large-scale annotations often show biases in low-quality datasets. For example, Multiple-Choice Questions (MCQs) datasets with one single correct option is common, however, there may be questions attributed to *none* or *multiple* correct options; whereas true-or-false questions are supposed to be labeled with either *True* or *False*, but similarly the text can include unsolvable elements, which should be further labeled as *Unknown*. There are problems when low-quality datasets with mixed question forms can not be identified. We refer to these exceptional label forms as *Sparse Labels*, and LLMs' ability to distinguish datasets with *Sparse Labels* mixture is important. Since users may not know situations of datasets, their instructions can be biased. To study how different instruction settings affect LLMs' identifications of *Sparse Labels* mixture, we introduce the concept of *Instruction Boundary*, which systematically evaluates different instruction settings that lead to biases. We propose BIASDETECTOR, a diagnostic benchmark to systematically evaluate LLMs on datasets with mixed question forms under *Instruction Boundary* settings. Experiments show that users' instructions induce large biases on our benchmark, highlighting the need not only for LLM developers to recognize risks of LLM biased annotation resulting in *Sparse Labels* mixture, but also problems arising from users' instructions to identify them. Code and datasets are available at <https://anonymous.4open.science/r/Instruction-Boundary>.

1 Introduction

LLMs such as GPT-4 (Achiam et al., 2023) have demonstrated advanced question-answering abilities. However, prior work mainly focuses on thinking along with fixed instruction guidance (Xu et al., 2025a; Shen et al., 2024): As we tell LLMs they

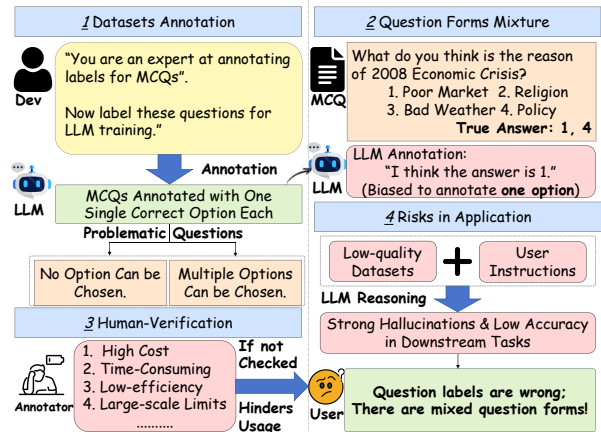


Figure 1: Background of problems and motivation of our work. Benchmarking the ability of LLMs adaptively answering questions with mixed forms under different instruction settings is important for LLM annotated datasets verification and usage.

are now facing *MCQs* and choose one option, or *True-or-False* questions which only contains *True* and *False* labels. LLMs are thinking based on a given range, their ability is largely enhanced by such guidance. However, as large-scale LLM annotated datasets are widely used, both text and options may contain flaws, as in Figure 1. Low-quality *MCQs* annotated with one single correct option can be common, however, the text and questions may potentially point to *none* or *multiple* correct options. There are contradictions between LLMs instruction following and the complexity of low-quality datasets, which result in strong *Hallucinations* and low accuracy, and such evaluations are lacking.

To address this, we formalize and quantify two factors that influence LLM biases: (1) Datasets with mixed question forms. (2) Different user instruction settings. By simulating the real-world scenarios of users' prompts inducing biases in LLM Reasoning, we establish a benchmark to evaluate LLMs' ability to output correct labels on datasets with mixed question forms, under various instruction settings.

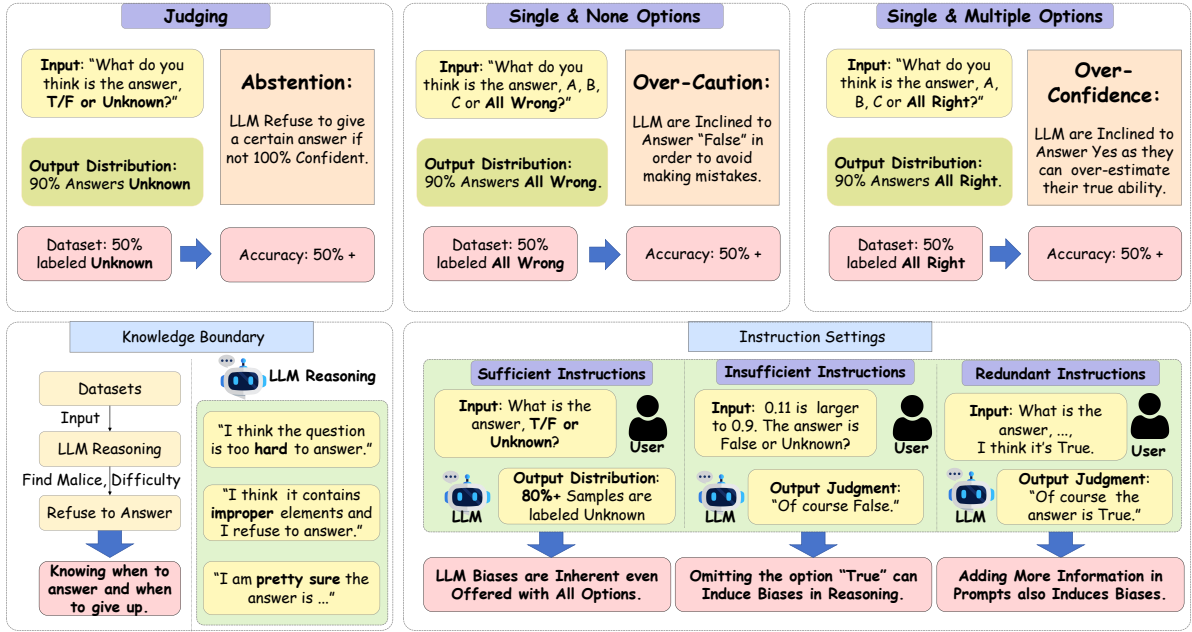


Figure 2: Illustrations of related work concepts. Although LLMs achieve an accuracy of 50% or higher, it is attributed more to biases than comprehension. *Knowledge Boundary* and *Instruction Settings* indicate LLM biases are influenced by both subjective and objective factors. Detailed Related Work is in Appendix B.

As shown in upper three pictures of Figure 2 and (Zheng et al., 2024; Gupta et al., 2023), biased output distributions can lead to seemingly great accuracy (e.g., answering one specific label can achieve 50% or higher accuracy), however, as output distributions do not align with datasets’ label distributions, this only reflects biases and imbalanced accuracy; Prior work on *Abstention* (Wen et al., 2025; Yin et al., 2023), *Over-Caution* (Cui et al., 2025; Sullutrone et al., 2025) and *Over-Confidence* (Geng et al., 2024) further proves problems arising these concepts. Meanwhile, many papers have explored the notion of the *Knowledge Boundary* (Lewis et al., 2020; Zhou et al., 2023), emphasizing LLMs should acknowledge uncertainty. We study these notions by mixing corresponding questions into our datasets, and study LLMs’ performances on our designed datasets under different instruction settings. The latter part is important since users may or may not know there are question forms mixture, thus prompting methods may further induce biases (Li et al., 2025b). This evaluation is urgent for both LLM developers and dataset users.

We formulate usual and exceptional label forms as *Dense Labels* and *Sparse Labels* based on usage statistics in Appendix Table 5, and *Instruction Boundary* for studying different instruction settings in Table 1. Compared to the definition of *Knowledge Boundary* on whether LLMs know the exact answers

or not, *Instruction Boundary* studying whether instructions cover necessary information matters too. This highlights we consider the phenomenon from both upstream sourcing and downstream processing.

Our work provides a holistic review of how instruction settings can lead to LLM biases on datasets with mixed question forms, which not only analyzes problems regarding LLMs biased annotation and output, but also alerts users to choose proper instruction settings to mitigate biases, pointing to future LLM trustworthiness development. In summary, our contributions are three-fold:

- ▷ For low-quality datasets and large-scale LLM annotations, we study problems regarding mixed question forms. LLMs are not able to distinguish questions with different label forms, which hinders the usage of LLM-annotated datasets.
- ▷ For users’ instruction settings, we formulate *Instruction Boundary* covering three situations: sufficient, redundant and insufficient instructions. Unlike prior work investigating individual sourcing of prompt or label biases, we propose a unified framework that evaluates LLMs’ biases from both user instructions and various question forms mixture.
- ▷ By comparing BIASDETECTOR with related benchmarks, we conclude how findings across different work generalize, deepening our understandings of LLM trustworthiness. Our work offers values for both LLM developers and dataset users.

Table 1: *Instruction Boundary* settings: classification of instruction settings with descriptions. *Instruction Boundary* is divided into three categories: (1) Sufficient (2) Redundant, and (3) Insufficient Instructions. Each category is further divided into sub-categories. **Disturbing Miscellany** is carried out separately, based on MMLU-Pro (Wang et al., 2024b) with 10 options in total.

Sufficient-Instructions	Paradigm	Examples
Vanilla	Instructions are comprehensive with all label forms and detailed guidance.	Which is the answer? True: (Definition...). False: (Definition...). Unknown: (Definition...).
Reasoning-Reflection	Collecting and inputting previous reasoning traces for LLM reflection.	Which is the answer? T:..., F:..., or Unknown:... Also show each reasoning step. (Collects reasoning traces for reflection.)
Prompt-Polishing	Asking LLMs to polish up prompts to improve understanding.	Which is the answer? T:..., F:..., or Unknown:... (Prompt polished up by LLMs.)
Redundant-Instructions	Paradigm	Examples
Conformity	LLMs exhibit bias toward labels (<i>dense</i> or <i>sparse</i>) that are explicitly in prompts.	Which is the answer? T:..., F:..., or Unknown:... I think is (One <i>dense</i> or <i>sparse</i> label is directly offered in instructions.)
Disturbing Miscellany	Additional wrong options are included, which increases distraction.	Which is the answer? (Contains extra wrong options in MCQs.)
Few-shot Learning	LLMs are biased toward labels that are presented in few-shot learning.	Which is the answer? T:..., F:..., or Unknown:.... Here is an example to learn from. (One example with one <i>dense</i> or <i>sparse</i> label.)
Insufficient-Instructions	Paradigm	Examples
Missing Choices	Some question forms are omitted, causing biases towards available options.	Which is the answer? T:..., Unknown:... (One specific option is omitted.)
Vagueness	Instructions cover all question forms but lack detailed guidance.	Which is the answer? True, False, or Unknown. (Detailed guidance is omitted.)

2 Phenomenon Formulation

Instruction Boundary *Instruction Boundary* covers three kinds: (1) **Sufficient Instructions** covers necessary information (i.e., all possible question forms, detailed guidance), where possible reasoning biases mainly arise from the LLMs’ internal tendencies. (2) **Redundant Instructions** with extra and more information, which can potentially misguide LLMs. (3) **Insufficient Instructions** that leave out specific and necessary information. As shown in Table 1, *Instruction Boundary* covers eight different experimental settings. Details, Templates are in Appendix C, M.

Label Categorizations We classify question label forms into *dense* and *sparse* based on usage statistics in Appendix Table 5. *Dense Labels* are often seen and trained on; *Sparse Labels* are infrequently used and treated as exceptional forms. In our work,

Table 2: Datasets categorizations, *Judging Questions* and *MCQs* are tested separately. *Sparse Labels* are in bold.

Type	Dataset	Samples	Form
Judging Questions	FLD (Morishita et al., 2024)	600	T/F/Unknown
	FOLIO (Han et al., 2022)	640	
Multiple-Choice Questions (MCQs)	ScienceQA (Lu et al., 2022b)	300	Single-correct
	MMLU (Hendrycks et al., 2021a)	300	
	TruthfulQA (Lin et al., 2022)	300	
	MMLU-Pro (Wang et al., 2024b)	900	None-correct
	CommonsenseQA (Talmor et al., 2019)	900	
MultiRC (Khashabi et al., 2018)	900	Multiple-correct	

one correct option in *MCQs* and *True, False* labels in *Judging Questions* are *Dense Labels*, whereas none correct option, multiple correct options and the *unknown* labels are categorized into *Sparse Labels*, as shown in Table 2. Details are in Appendix A.

Datasets Our dataset covers two question types: (1) *Judging Questions*: FLD (Morishita et al., 2024), and FOLIO (Han et al., 2022), contains *True, False*, and *Unknown* labels (2) *Multiple-Choice*

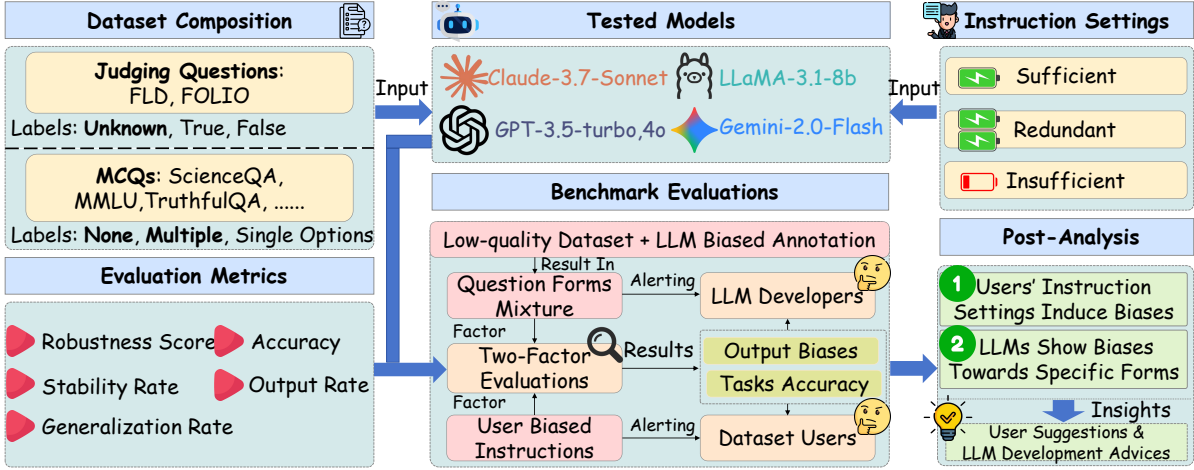


Figure 3: General framework of our proposed framework BIASDETECTOR. First of all, we establish datasets with mixed question forms and *Instruction Boundary* settings. Then we apply five LLMs from different families to generate responses. Finally by testing their robustness and accuracy regarding tasks accuracy and output biases, we evaluate problems and offer insights for both LLM developers and dataset users.

Questions (MCQs), containing samples from ScienceQA (Lu et al., 2022b), MultiRC (Khashabi et al., 2018), MMLU (Hendrycks et al., 2021a), MMLU-Pro (Wang et al., 2024b), CommonsenseQA (Talmor et al., 2019), TruthfulQA (Lin et al., 2022). First, we process both types of questions to fit the framework: (1) For *Judging Questions* datasets, these samples are labeled by dataset creators, and we did no revision. (2) For *MCQs*, questions with one correct option are considered to have *Dense Labels*, for none or multiple correct options are categorized as *Sparse Labels*. We extract 300 samples per dataset from ScienceQA, MMLU and TruthfulQA, representing samples with one correct option. We then extract 900 samples from CommonsenseQA containing five options with one correct option, by deleting the correct one we get samples with none correct option. For multiple correct options, we extract 900 samples from MultiRC with four options with two correct ones. Sampling seeds are all set to 42, revised samples are manually checked to ensure validness, examples are in Appendix S. Details are in Table 2. **As we simulate question forms mixture, adding options like "None Correct" or "A and B are Correct" does not meet real-world scenarios. LLMs should identify different forms based on the same instruction for each question, rather than being informed separately in options.**

3 Benchmark Metrics

We introduce the metric of performance **Accuracy**:

$$Acc = \frac{1}{|D|} \sum_{i=1}^{|D|} \mathbb{1}(y^i = \hat{y}^i) \quad (1)$$

Further, we implement the following two metrics, defining **Stability Rate (SR)** and **Generalization Rate (GR)** as below, measuring the accuracy of two label forms respectively:

$$SR = \frac{1}{|D_D|} \sum_{y^i \in D_D} \mathbb{1}(y^i = \hat{y}^i) \quad (2)$$

$$GR = \frac{1}{|D_S|} \sum_{y^i \in D_S} \mathbb{1}(y^i = \hat{y}^i) \quad (3)$$

For y^i indicates the true label, \hat{y}^i is the predicted label, and D_D , D_S mean datasets with *dense* and *sparse* labels, which are 50% and 50% equally sampled, and D indicates the whole dataset, from which we can get $D_D \cup D_S = D$. GR represents LLMs identify *Sparse Labels* correctly, and SR for *Dense Labels*. Since there are two *sparse label* forms in *MCQs*, we further decompose its GR into GR^1 and GR^2 , measuring None Correct Option and Multiple Correct Options. Beside, we study biases via output distributions **Output Rate (OR)**. We analyze **output label distributions** of *Judging Questions* and **output option number distributions** of *MCQs*:

$$OR_D = \frac{1}{|D_D|} \sum_{i=1}^{|D_D|} \mathbb{1}(\hat{y}^i \in D_D) \quad (4)$$

$$OR_S = \frac{1}{|D_S|} \sum_{i=1}^{|D_S|} \mathbb{1}(\hat{y}^i \in D_S) \quad (5)$$

OR_D and OR_S represents LLM output distributions of *Dense Labels* and *Sparse Labels*. Outputting one option in *MCQs*, or *True* and *False* labels in *Judging Questions* are categorized into OR_D , and other output situations fall into OR_S .

Table 3: Results across seven settings except *Disturbing Miscellany*. Accuracy is presented with the Standard Deviation (std) by running each setting four times repeatedly to mitigate randomness. Highest values are in bold.

Models	Settings	Judging Questions			MCQs			
		Acc	SR	GR	Acc	SR	GR ¹	GR ²
GPT-3.5-turbo	Vanilla	48.75 ± 2.5	21.47	76.03	24.67 ± 1.3	62.00	12.00	0.00
	Reasoning-Reflection	36.14 ± 1.3	31.51	40.78	33.56 ± 1.5	40.67	60.00	0.00
	Prompt-Polishing	34.30 ± 1.3	9.16	59.44	21.33 ± 1.3	58.67	5.33	0.00
	Conformity-Dense	50.00 ± 0.0	9.38	90.63	32.20 ± 1.8	83.33	13.33	0.00
	Conformity-Sparse	50.47 ± 1.2	1.56	99.38	37.11 ± 1.8	11.33	100.00	0.00
	Few shot-Dense	50.00 ± 0.0	99.38	0.63	35.60 ± 2.0	93.33	16.67	0.00
	Few shot-Sparse	50.47 ± 0.0	1.56	99.38	52.25 ± 2.0	90.00	43.35	0.00
	Missing Choices	34.76 ± 1.3	61.04	8.48	25.56 ± 2.8	64.00	12.67	0.00
	Vagueness	38.32 ± 0.8	23.17	50.38	24.70 ± 2.2	73.30	0.70	0.00
GPT-4o	Vanilla	45.41 ± 2.1	63.49	27.32	35.60 ± 2.8	89.33	28.67	0.00
	Reasoning-Reflection	40.07 ± 2.4	29.53	50.62	18.89 ± 2.0	16.67	40.00	0.00
	Prompt-Polishing	48.86 ± 1.9	60.50	37.21	37.11 ± 2.0	89.33	19.33	2.67
	Conformity-Dense	55.31 ± 1.9	62.19	48.44	40.00 ± 1.5	90.00	30.00	0.00
	Conformity-Sparse	56.92 ± 1.9	20.60	93.24	44.00 ± 1.5	83.33	48.67	0.00
	Few shot-Dense	56.12 ± 1.8	80.73	31.50	35.60 ± 1.5	93.30	13.30	0.00
	Few shot-Sparse	63.13 ± 1.8	61.25	65.00	44.00 ± 1.5	98.00	34.00	0.00
	Missing Choices	40.39 ± 2.4	38.09	47.69	40.89 ± 1.5	90.00	32.67	0.00
	Vagueness	41.01 ± 2.2	12.28	69.75	32.70 ± 2.8	93.30	4.70	0.00
LLaMA-3.1-8B	Vanilla	51.87 ± 2.2	17.92	85.82	20.67 ± 2.2	55.33	4.67	2.00
	Reasoning-Reflection	38.81 ± 1.5	57.70	19.93	15.33 ± 1.5	36.00	8.67	1.33
	Prompt-Polishing	42.58 ± 0.8	50.28	34.88	7.78 ± 0.0	13.33	5.33	4.67
	Conformity-Dense	30.47 ± 1.4	60.31	0.63	23.30 ± 2.8	56.70	13.30	0.00
	Conformity-Sparse	50.00 ± 0.0	0.00	100.00	38.00 ± 2.8	50.67	60.00	3.33
	Few shot-Dense	46.28 ± 1.2	77.40	15.15	26.22 ± 2.8	70.67	6.00	2.00
	Few shot-Sparse	51.88 ± 1.2	5.63	98.13	38.00 ± 2.8	50.67	63.33	0.67
	Missing Choices	47.91 ± 1.5	8.67	86.67	22.44 ± 0.8	59.33	5.33	2.67
	Vagueness	29.21 ± 2.5	0.32	58.11	25.60 ± 1.5	75.30	0.70	0.70
Claude-3.7-Sonnet	Vanilla	46.23 ± 1.8	40.13	52.33	7.78 ± 0.0	12.67	7.33	3.33
	Reasoning-Reflection	45.73 ± 1.3	55.48	35.98	14.00 ± 1.5	16.67	24.67	0.67
	Prompt-Polishing	57.21 ± 0.8	84.06	27.66	2.67 ± 0.0	1.33	5.33	1.33
	Conformity-Dense	57.03 ± 1.2	33.75	80.31	10.00 ± 3.7	13.33	10.00	6.67
	Conformity-Sparse	55.10 ± 1.2	18.70	91.49	3.78 ± 3.7	1.33	7.33	2.67
	Few shot-Dense	54.10 ± 1.0	54.09	54.11	25.11 ± 1.8	52.67	21.33	1.33
	Few shot-Sparse	49.84 ± 2.5	63.44	36.25	22.20 ± 1.8	55.00	5.00	6.65
	Missing Choices	53.19 ± 1.5	26.32	80.07	9.78 ± 1.8	14.67	11.33	3.33
	Vagueness	51.20 ± 0.9	45.34	57.07	30.20 ± 2.5	88.70	2.00	0.00
Gemini-2.0-Flash	Vanilla	51.81 ± 1.8	36.99	66.62	35.00 ± 1.8	84.85	18.18	2.94
	Reasoning-Reflection	40.84 ± 1.6	39.81	41.87	36.00 ± 2.5	60.67	45.33	2.00
	Prompt-Polishing	35.40 ± 1.6	0.50	70.31	10.44 ± 0.0	18.67	7.33	5.33
	Conformity-Dense	35.00 ± 2.2	56.25	13.75	48.90 ± 1.8	90.00	56.70	0.00
	Conformity-Sparse	51.20 ± 2.2	4.68	97.70	50.00 ± 1.8	65.33	84.67	0.00
	Few shot-Dense	54.51 ± 2.5	88.23	18.31	42.67 ± 1.5	87.33	40.67	0.00
	Few shot-Sparse	57.34 ± 1.5	17.19	78.13	42.67 ± 1.5	87.33	40.67	0.00
	Missing Choices	45.79 ± 0.8	31.30	60.28	36.00 ± 2.8	87.88	21.21	0.00
	Vagueness	40.73 ± 1.5	19.49	61.96	31.60 ± 1.8	90.70	4.00	0.00

4 Experiments

4.1 Experimental Setup

Models As in (Ling et al., 2025), although larger parameter LLMs show better performances, they are willing to output *unknown*. We test closed and open-source LLM families, including Claude-3-7-Sonnet (Anthropic, 2024), Gemini-2.0-Flash (DeepMind, 2024), LLaMA-3.1-8B (Touvron et al.,

2023b), GPT-3.5-turbo and GPT-4o (OpenAI, 2023, 2025). Tested models use Temperature = 0.5 unless specified; Qwen-2.5-7b (Bai et al., 2023) used as a LLM judge runs with T=0, and we apply testing cases to prove its validness. Since the LLM judge deals with exceptional cases when **label names** in *Judging Questions* or **option indexes** in *MCQs* failed to match, the usage is small (below 5%) and biases are minor. Details are in Appendix D.1.

Experimental Settings As in Table 1, we test LLMs under *Instruction Boundary* settings to quantify biases attributed to both instructions and datasets with mix question forms, except for the *Vanilla* setting served as a baseline. *Judging Questions* and *MCQs* are tested separately to better simulate real-world scenarios. Case studies are in Appendix P.

4.2 General Experimental Results

As shown in Figure 3, tested LLMs show poor performance as the accuracy lies in the range of 35% to 60%. Claude-3.7-sonnet, GPT-4o shows the highest average *Acc.* Meanwhile, LLMs show high instability and reasoning gap between *dense* and *sparse* labels since the gap between *SR* and *GR* is high. The value of GR^2 , which is the accuracy in multiple correct options, are mostly less than 5%, indicating a strong *Over-Caution* behavior.

For *Judging Questions*, the *SR* is generally lower than *GR*, indicating given the *Unknown* option, LLMs are inclined to this rather than outputting *True* or *False*. This situation is slightly better for GPT-4o and Claude-3.7-Sonnet, as the gap between two metrics is relatively low. We also observe instability in GPT-3.5-turbo, as *SR* and *GR* values are completely different in *Conformity* and *Few-shot Learning*, showing strong biases if being misguided.

Reasoning on *MCQs* demonstrates an opposite behavior, as *SR* is overall higher than GR^1 and GR^2 , LLMs are good at answering questions with one correct option over others, and this gap is extremely huge for GPT-4o and Gemini-2.0-Flash. We also observe a trade-off for these models: Being better at reasoning *Judging Questions* means worse stability facing *MCQs*.

Significance Testing By comparing the results of the *Vanilla* setting with other biased settings, we gain better understanding of how instruction settings influence LLMs. As in Figure 5, some cause serious accuracy drop, and others result in a seemingly good accuracy with a biases shifting. Overall, we can see *Vagueness*, *Reasoning-Reflection* and *Disturbing Miscellany* influence accuracy the most.

4.3 Results on Sufficient Instructions

Prior work (Zheng et al., 2024; Fei et al., 2023) demonstrated LLMs show biases towards specific labels, and self-reflection (Xu et al., 2024), prompt polishing can induce biases and instability (Jakesch et al., 2023). But LLMs’ ability of distinguishing question forms mixture are rarely tested.

Vanilla Setting As illustrated in the right two panels of Figure 4, LLMs exhibit explicit and unified bias even when instructions cover all possible question forms: for *Judging Questions*, LLMs are inclined to output *Unknown* and *True*. For *MCQs*, outputting one single choice takes up a significant amount, indicating strong and unified biases towards specific labels, and simply including all necessary elements for LLMs to choose from is not enough for mitigating biases.

Reasoning-Reflection As in Table 3, while this setting can narrow the gap between *SR* and *GR*, it yields mostly accuracy drop, indicating LLM reflecting on previous reasoning traces is problematic.

Prompt-Polishing Users like to polish up prompts by LLMs. We input prompts in *Vanilla* setting into GPT-4o and ask LLMs to polish up them, templates are in Appendix O. We observe a wide accuracy drop among all models, ranging from 5% to 30%. Meanwhile, most LLMs demonstrate a larger gap between *SR* and *GR* compared with the *Vanilla* setting, indicating this setting is likely to induce biases.

Even with complete instruction coverage, LLMs continue to exhibit consistent biases, suggesting the performances on mixed question forms is weak, and the ability to adaptively reason needs improvements.

4.4 Results on Redundant Instructions

LLMs show biases towards few-shot examples was addressed in (Zhao et al., 2021; Reif and Schwartz, 2024; Si et al., 2023a), meanwhile, conformity effects are researched in (Zhu et al., 2024; Shoval et al., 2025; Sharma et al., 2025), and robustness on option numbers misguidance is in (Zhang et al., 2024c). However, biases facing mixed question forms is still under-explored. To the best of our knowledge, we are the first to compare *Few-shot Learning* with *Conformity* and extra options in *MCQs* to study biases, and we further sub-categorize *Few-shot Learning* and *Conformity* based on included label forms.

Few-shot Learning As in Table 3, compared to the *Vanilla* setting, Few-shot learning leads to an accuracy gain, but the value of *SR* and *GR* are largely changed. Details are in Appendix H. We include examples with either *sparse* or *dense* labels and compare differences, Figure 5 shows that example labels can trigger huge biased output. LLM performance gain is largely attributed to bias shifts rather than understanding. The ability cannot be improved via imitations since performances are unstable.

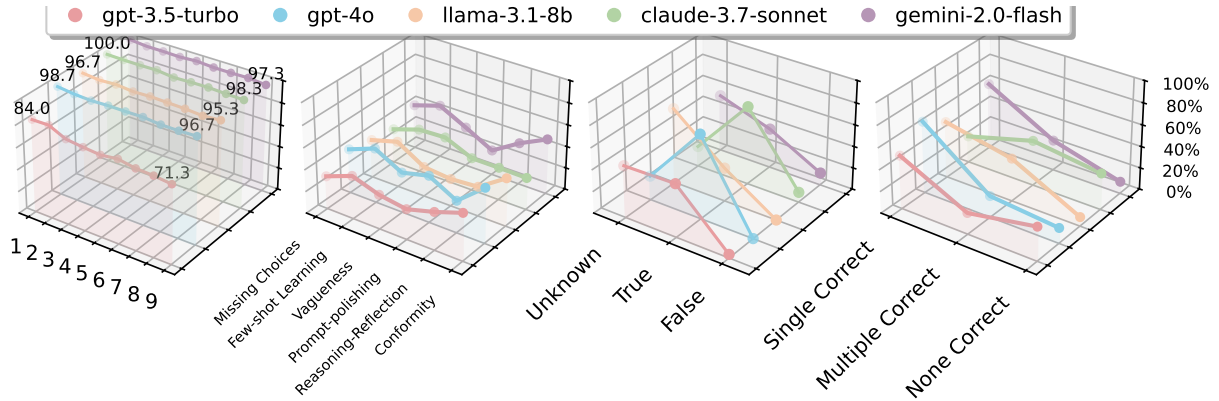


Figure 4: Results of *Accuracy* under *Disturbing Miscellany* setting, accuracy is dropping as option number is increasing according to the x-axis, and average *Accuracy* among settings (left one, two). *Label Output Distribution (OR)* under the *Vanilla* setting (right one, two). Standardized with the scale of the rightmost figure as the reference.

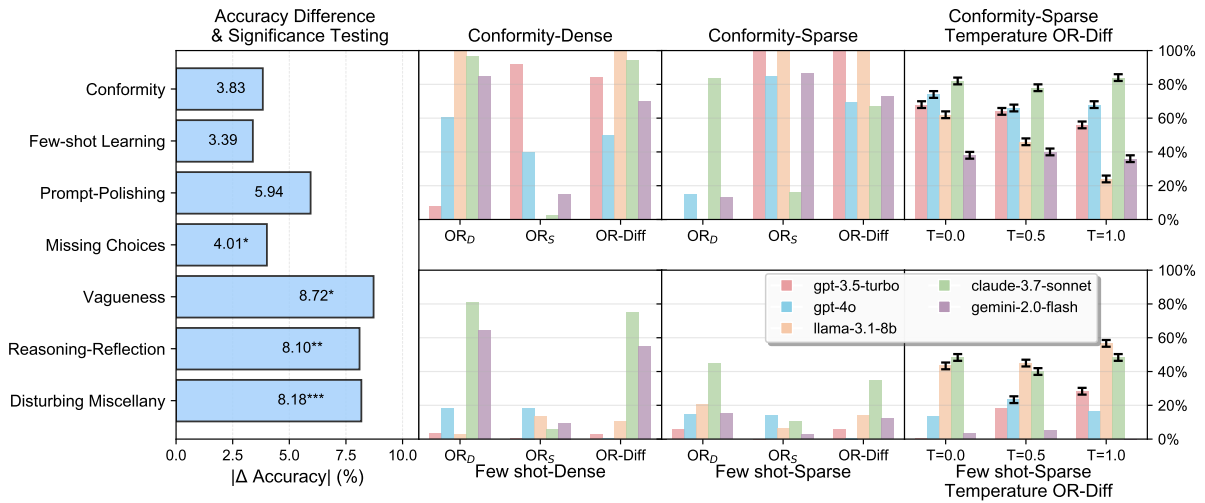


Figure 5: Accuracy differences comparing biased settings with the *Vanilla* setting, and Significance Testing. *: $p < 0.05$ (95% confidence); **: $p < 0.01$ (99% confidence), ***: $p < 0.001$ (99.9% confidence). Detailed results on *MCQs* under *Conformity* and *Few-shot* settings (Right Six). $OR - Diff = |OR_D - OR_S|$. Details are in Section 3.

Conformity LLM biases can be extremely high when labels are given directly. Figure 3 indicates that even though accuracy doesn't change much, the gap between *SR* and *GR* is even higher compared with other settings. Details of *OR* are in Fig 5, charts 2 and 3 demonstrate a higher *OR* shift in most LLMs regarding both *dense* and *sparse* labels. LLMs demonstrate much higher conformity regarding *Sparse Labels* like *Unknown*, compared with *Dense Labels*. We also quantify how (*OR*) varies under three temperature settings: $T \in \{0.0, 0.5, 1.0\}$. As in Figure 5, takeaways regarding temperature changes is same for both *Few-shot* and *Conformity*: merely revising *T* yields limited gains, and larger temperatures tend to introduce greater instability.

Disturbing Miscellany We investigate how the number of wrong options (i.e., information redundancy) affects accuracy. We choose MMLU-Pro (Wang et al., 2024b), *MCQs* consisting one true

option and nine false options per sample. We save **N** numbers of wrong options to quantify how redundant information influences performances, **N** ranges from one to nine. As in Figure 4, most LLMs show a slight drop in accuracy, except for outdated models like GPT-3.5-turbo demonstrating a high decline, others maintain high robustness when *N* increases. This setting proves LLMs' resistance to redundancy is strong. As **Conformity induces the most serious biases, Few-shot Learning ranks the second, and Disturbing Miscellany the last.** We empirically show that *Conformity* effects result from expression tuning rather than redundant information.

4.5 Results on Insufficient Instructions

Prior research on prompt vagueness (Si et al., 2023b) and missing options (Góral et al., 2025; Cao et al., 2024a) show that LLM performances are largely influenced by prompts, but mixed question forms are

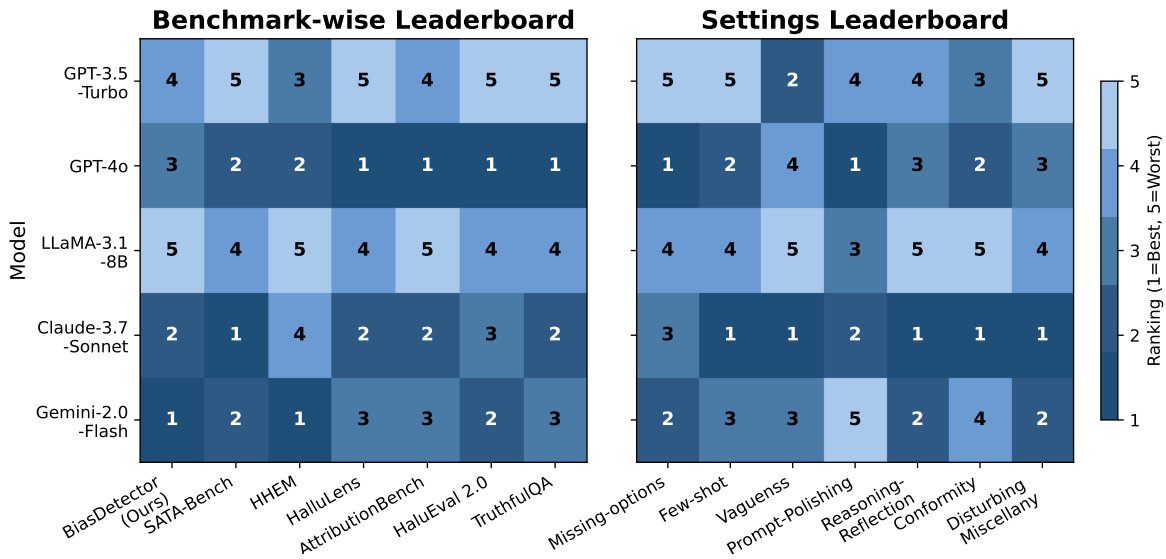


Figure 6: Leaderboard of related topic benchmarks(left), and inner settings comparisons among five LLMs used in our benchmark(Right). For benchmark-wise comparisons, we show the ranking under the **Vanilla Setting**. Our metric is **Robustness Score (RS)**. Detailed values can be found in Table 8.

rarely tested. This setting is for simulating users not knowing question forms mixture in datasets, and test whether LLMs can independently reason under biased instruction settings. For *Vagueness* and *Missing Choices*, omitting either guidance or *Sparse Label* options result in huge accuracy drop. The primary bottleneck is the ability to think independently. Ideally, even if specific options are omitted in prompts, LLMs can still answer correctly.

5 Leaderboard Analysis

We compare LLM rankings among related benchmarks. Our metric is the *harmonic mean* of *GR* and *SR*, denoted as **Robustness Score (RS)**:

$$RS_M = \frac{3 \cdot SR \cdot GR^1 \cdot GR^2}{SR \cdot GR^1 + SR \cdot GR^2 + GR^1 \cdot GR^2} \quad (6)$$

$$RS_J = \frac{2 \cdot SR \cdot GR}{SR + GR}, RS = \frac{1}{2}(RS_M + RS_J) \quad (7)$$

Judging Questions and *MCQs* are calculated separately then average. Table 2 shows *GR* and *SR* are calculated based on equal numbers of samples.

We select open-sourced benchmarks (1) Measured by Accuracy: (Li et al., 2023a), (Bang et al., 2025), (Lin et al., 2022), (Li et al., 2024a), (Xu et al., 2025a) (2) Measured by Macro-F1: (Li et al., 2024c). As in Figure 6, our findings mostly align with prior work. For the *Vanilla* setting, Gemini-2.0-Flash offers the greatest cost-efficiency, as the performance ranks first with lower costs compared to GPT-4o and Claude-3.7-Sonnet.

6 Discussion

We offer three advice for problem mitigation: (1) We design a prompt template and run experiments in Appendix F to ease biases. (2) LLMs perform poorly on mixed question forms, making datasets manual verifications necessary. (3) MoE (Mixture-of-Experts) structures (Mu and Lin, 2025) which are suitable for various but interrelated tasks is a promising way to fundamentally solve this problem.

7 Conclusion

We establish a dataset containing both *Judging Questions* and *MCQs* with mixed question forms. By evaluating LLMs' ability to adaptively answer correct labels under three instruction settings: sufficient, redundant and insufficient, we systematically quantify LLM weaknesses as they neither think independently under different instruction settings, nor adaptively distinguish mixed question forms. Unlike prior work, we do not revise options to include these forms, instead we test LLMs based on unified instructions to identify forms mixture themselves. Sufficient instructions still introduce biases despite detailed guidance, whereas redundant and insufficient instructions amplify biases, proving the necessity of manual verification for low-quality datasets, and alerts users of instruction settings selection. To the best of our knowledge, we are the first to systematically evaluate LLMs under this setting, offering value for LLM trustworthiness developments.

8 Limitations

Although we address three categories of prompt coverage problems, our work does not consider mainstream workflows such as LLM agents or RAG systems. Moreover, our study focuses primarily on language-related tasks, leaving vision-related challenges outside the scope of this paper. Besides, for real-life applications, user prompt forms may be varied and noisy, leaving the final performance unstable. Unified analysis regarding different elaboration forms is valuable but too large to be examined.

9 Ethical Considerations

All datasets are **publicly available for non-commercial research**. FLD is an entirely synthetic dataset; FOLIO, ScienceQA, MMLU, CommonsenseQA, TruthfulQA, MMLU-Pro, and MultiRC derive from publicly available open educational resources and contain no personally identifiable information (PII). No additional human data were collected, so the study does not require Institutional Review Board (IRB) review.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Anonymous. 2023. Investigating data contamination in modern benchmarks for large language models. *arXiv preprint*.

Anthropic. 2024. Introducing claude 3.7 sonnet. <https://www.anthropic.com/news/claude-3-7>. Accessed: 2025-09-15.

Samee Arif, Zohaib Khan, Agha Ali Raza, and Awais Athar. 2024. With a grain of salt: Are llms fair across social dimensions? *arXiv preprint arXiv:2410.12499*.

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, and 29 others. 2023. *Qwen technical report*. *Preprint*, arXiv:2309.16609.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, and 32 others. 2022. *Constitutional ai: Harmlessness from ai feedback*. *Preprint*, arXiv:2212.08073.

Yejin Bang, Ziwei Ji, Alan Schelten, Anthony Hartshorn, Tara Fowler, Cheng Zhang, Nicola Cancedda, and Pascale Fung. 2025. *Hallulens: Llm hallucination benchmark*. *Preprint*, arXiv:2504.17550.

Joachim Baumann, Paul Röttger, Aleksandra Urman, Albert Wendsjö, Flor Miriam Plaza del Arco, Johannes B. Gruber, and Dirk Hovy. 2025. *Large language model hacking: Quantifying the hidden risks of using llms for text annotation*. *Preprint*, arXiv:2509.08825.

Bowen Cao, Deng Cai, Zhisong Zhang, Yuexian Zou, and Wai Lam. 2024a. *On the worst prompt performance of large language models*. *Preprint*, arXiv:2406.10248.

Bowen Cao, Deng Cai, Zhisong Zhang, Yuexian Zou, and Wai Lam. 2024b. *On the worst prompt performance of large language models*. *arXiv preprint arXiv:2406.10248*.

Muhammad Umer Cheema, Jiyang Jin, Xiang Xie, Yitian Shao, Yifan Hou, Yuzhong Huang, Minghui Zhang, Qi Zhang, and Jiliang Sun. 2024. *Trustllm: A benchmark for trustworthiness of large language models*. *arXiv preprint arXiv:2403.09028*.

Khaoula Chehbouni, Mohammed Haddou, Jackie Chi Kit Cheung, and Golnoosh Farnadi. 2025. *Neither valid nor reliable? investigating the use of llms as judges*. *Preprint*, arXiv:2508.18076.

Karl Cobbe, Vinay Kosaraju, Mohammad Bavarian, Mark Chen, and et al. 2021. *Training verifiers to solve grade school math word problems*. *arXiv preprint arXiv:2104.03253*.

Justin Cui, Wei-Lin Chiang, Ion Stoica, and Chojui Hsieh. 2025. *Or-bench: An over-refusal benchmark for large language models*. *Preprint*, arXiv:2405.20947.

Google DeepMind. 2024. Gemini 2.0 flash. <https://deepmind.google/technologies/gemini/flash>. Accessed: 2025-09-15.

Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. *Mitigating label biases for in-context learning*. *Preprint*, arXiv:2305.19148.

Deep Ganguli, Amanda Askell, Yuntao Bai, and et al. 2022. *Red teaming language models with language models*. *arXiv preprint arXiv:2210.07237*.

Shunyu Gao, Swaroop Mishra, Daniel Huang, Thomas Liu, Saket Mittal, and et al. 2023. *Pal: Program-aided language models*. In *Advances in Neural Information Processing Systems (NeurIPS)*.

Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024. *A survey of confidence estimation and calibration in large language models*. *Preprint*, arXiv:2311.08298.

Gracjan Góral and Emilia Wiśnios. 2024. *When all options are wrong: Evaluating large language model robustness with incorrect multiple-choice options*. *arXiv preprint arXiv:2409.00113*.

523	Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. 2024. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. <i>arXiv preprint arXiv:2310.14566</i> .	Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. <i>Advances in neural information processing systems</i> , 33:9459–9474.	579
524			580
525			581
526			582
527			
528		Chaoyu Li, Eun Woo Im, and Pooyan Fazli. 2025a. Vid-halluc: Evaluating temporal hallucinations in multi-modal large language models for video understanding. <i>arXiv preprint arXiv:2412.03735</i> .	583
529			584
			585
530	Karan Gupta, Sumegh Roychowdhury, Siva Rajesh Kasa, Santhosh Kumar Kasa, Anish Bhanushali, Nikhil Pattisapu, and Prasanna Srinivasa Murthy. 2023. How robust are llms to in-context majority label bias? <i>Preprint</i> , arXiv:2312.16549.	Jiahui Li, Sean Papay, and Roman Klinger. 2025b. Are humans as brittle as large language models? <i>Preprint</i> , arXiv:2509.07869.	587
531			588
532			589
533			
534			
535	Gracjan Góral, Emilia Wiśnios, Piotr Sankowski, and Paweł Budzianowski. 2025. Wait, that’s not an option: Llms robustness with incorrect multiple-choice options. <i>Preprint</i> , arXiv:2409.00113.	Jiaxin Li, Geng Zhao, and Xiaoci Zhang. 2025c. Multi-lingual federated low-rank adaptation for collaborative content anomaly detection across multilingual social media participants. In <i>Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing</i> , pages 15253–15273, Suzhou, China. Association for Computational Linguistics.	590
536			591
537			592
538			593
539	Simeng Han, Hailey Schoelkopf, Yilun Zhao, Zhenting Qi, Martin Riddell, Wenfei Zhou, James Coady, David Peng, Yujie Qiao, Luke Benson, and 1 others. 2022. Folio: Natural language reasoning with first-order logic. <i>arXiv preprint arXiv:2209.00840</i> .	Junyi Li, Xiaoxue Cheng, Wayne Xin Zhao, Jian-Yun Nie, and Ji-Rong Wen. 2023a. Halueval: A large-scale hallucination evaluation benchmark for large language models. <i>Preprint</i> , arXiv:2305.11747.	594
540			595
541			596
542			
543			
544	Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. Measuring massive multitask language understanding. <i>Proceedings of the International Conference on Learning Representations (ICLR)</i> .	Miaoran Li, Rogger Luo, and Ofer Mendelevitch. 2024a. HHEM-2.1-Open .	597
545			598
546			599
547			600
548			
549	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, and et al. 2021b. Measuring mathematical problem solving with the math dataset. In <i>Advances in Neural Information Processing Systems (NeurIPS)</i> .	Moxin Li, Yong Zhao, Yang Deng, Wenxuan Zhang, Shuaiyi Li, Wenya Xie, See-Kiong Ng, and Tat-Seng Chua. 2024b. Knowledge boundary of large language models: A survey. <i>arXiv preprint arXiv:2412.12472</i> .	601
550			602
551			603
552			604
553	Tiancheng Hu, Yara Kyrychenko, and Jon Roizenbeek. 2025. Generative language models exhibit social identity biases. <i>Nature Computational Science</i> . Published online December 2024.	Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. In <i>Proceedings of EMNLP 2023</i> . ArXiv:2305.10355.	605
554			606
555			
556			
557	Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madihan Khabsa. 2023. Llama guard: Llm-based input-output safeguard for human-ai conversations. <i>Preprint</i> , arXiv:2312.06674.	Yifei Li, Xiang Yue, Zeyi Liao, and Huan Sun. 2024c. Attributionbench: How hard is automatic attribution evaluation? <i>Preprint</i> , arXiv:2402.15089.	607
558			608
559			609
560			610
561			
562		Percy Liang, Rishi Bommasani, Kenny Lee, and et al. 2022. Holistic evaluation of language models. <i>arXiv preprint arXiv:2211.09110</i> .	611
563	Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users’ views. In <i>Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems</i> , CHI ’23, page 1–15. ACM.	Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In <i>Proceedings of ACL 2022</i> . ArXiv:2109.07958.	612
564			613
565			
566			
567			
568			
569	Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In <i>Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)</i> , pages 252–262.	Chen Ling, Xujiang Zhao, Xuchao Zhang, Wei Cheng, Yanchi Liu, Yiyun Sun, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, Guangji Bai, Liang Zhao, and Haifeng Chen. 2024. Uncertainty quantification for in-context learning of large language models. <i>Preprint</i> , arXiv:2402.10189.	614
570			615
571			616
572			
573			
574			
575			
576			
577	Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich	Zipeng Ling, Yuehao Tang, Shuliang Liu, Junqi Yang, Shenghong Fu, Chen Huang, Kejia Huang, Yao Wan, Zhichao Hou, and Xuming Hu. 2025. Wakenllm: Evaluating reasoning potential and stability in llms via fine-grained benchmarking. <i>Preprint</i> , arXiv:2507.16199.	617
578			618
			619
			620
			621
			622
			623
			624
			625
			626
			627
			628
			629
			630
			631

632	Jiazhen Liu, Yuhan Fu, Ruobing Xie, Runquan Xie, Xingwu Sun, Fengzong Lian, Zhanhui Kang, and Xirong Li. 2025. Phd: A chatgpt-prompted visual hallucination evaluation dataset . <i>Preprint</i> , arXiv:2403.11116.	687
633		688
634		689
635		690
636		691
637	Jiaxu Lou and Yifan Sun. 2024. Anchoring bias in large language models: An experimental study . <i>Preprint</i> , arXiv:2412.06593.	692
638		693
639		694
640	Pan Lu, Yanbang Feng, Moye Yu, Yejin Choi, and Hananeh Hajishirzi. 2022a. Scienceqa: A question answering benchmark for multimodal scientific reasoning. In <i>Proceedings of EMNLP</i> .	695
641		696
642		697
643		698
644	Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022b. Learn to explain: Multimodal reasoning via thought chains for science question answering. <i>Advances in Neural Information Processing Systems</i> , 35:2507–2521.	699
645		700
646		701
647		702
648		703
649		704
650	Huan Ma, Changqing Zhang, Yatao Bian, Lemao Liu, Zhirui Zhang, Peilin Zhao, Shu Zhang, Huazhu Fu, Qinghua Hu, and Bingzhe Wu. 2023. Fairness-guided few-shot prompting for large language models. In <i>Advances in Neural Information Processing Systems</i> .	705
651		706
652		707
653		708
654		709
655	Nishanth Madhusudhan, Sathwik Tejaswi Madhusudhan, Vikas Yadav, and Masoud Hashemi. 2024. Do llms know when to not answer? investigating abstention abilities of large language models . <i>arXiv preprint arXiv:2407.16221</i> .	710
656		711
657		712
658		713
659		714
660	Pattarapong Manakul, Jin Yao, Paul Rayson, Alexandros Pappangelis, Suresh Manandhar, and Andreas Vlachos. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative ai. In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)</i> .	715
661		716
662		717
663		718
664		719
665		720
666	Francesco Maria Molfese, Luca Moroni, Luca Gioffrè, Alessandro Scirè, Simone Conia, and Roberto Navigli. 2025. Right answer, wrong score: Uncovering the inconsistencies of llm evaluation in multiple-choice question answering . <i>Preprint</i> , arXiv:2503.14996.	721
667		722
668		723
669		724
670		725
671	Terufumi Morishita, Gaku Morio, Atsuki Yamaguchi, and Yasuhiro Sogawa. 2024. Enhancing reasoning capabilities of llms via principled synthetic logic corpus. <i>Advances in Neural Information Processing Systems</i> , 37:73572–73604.	726
672		727
673		728
674		729
675		730
676	Siyuan Mu and Sen Lin. 2025. A comprehensive survey of mixture-of-experts: Algorithms, theory, and applications . <i>Preprint</i> , arXiv:2503.07137.	731
677		732
678		733
679	Seb Murray. 2025. Generative ai results depend on user prompts as much as models. <i>MIT Sloan Ideas Made to Matter</i> . Accessed 6 Sep 2025.	734
680		735
681		736
682	Anna Neumann, Elisabeth Kirsten, Muhammad Bilal Zafar, and Jatinder Singh. 2025. Position is power: System prompts as a mechanism of bias in large language models. In <i>Proceedings of the ACM Conference on Fairness, Accountability, and Transparency</i> .	737
683		738
684		739
685		740
686		741
	Cheng Niu, Yuanhao Wu, Juno Zhu, Siliang Xu, Kashun Shum, Randy Zhong, Juntong Song, and Tong Zhang. 2024. Ragtruth: A hallucination corpus for developing trustworthy retrieval-augmented language models . In <i>Proceedings of the 62nd Annual Meeting of the ACL (Long Papers)</i> , pages 10862–10878.	
	OpenAI. 2023. Gpt-3.5: Generative pretrained transformer 3.5 . Accessed August 2025.	
	OpenAI. 2025. Gpt-4o: Openai’s new flagship multimodal model . Accessed August 2025.	
	Long Ouyang, Jeff Wu, Xu Jiang, Daniel Brown, and et al. 2022. Training language models to follow instructions with human feedback. <i>arXiv preprint arXiv:2203.02155</i> .	
	Pouya Pezeshkpour and Estevam Hruschka. 2023. Large language models sensitivity to the order of options in multiple-choice questions . <i>Preprint</i> , arXiv:2308.11483.	
	Han Qiu, Jiaying Huang, Peng Gao, Qin Qi, Xiaoqin Zhang, Ling Shao, and Shijian Lu. 2024. Longhalqa: Long-context hallucination evaluation for multimodal large language models . <i>arXiv preprint arXiv:2410.09962</i> .	
	Rafael Rafailov, Andy Wang, Tianyi Zhang, and Percy Liang. 2024. Direct preference optimization: Your language model is secretly a reward model. In <i>Proc. ICLR</i> .	
	Yuval Reif and Roy Schwartz. 2024. Beyond performance: Quantifying and mitigating label bias in llms . <i>Preprint</i> , arXiv:2405.02743.	
	Jens Rupperecht, Georg Ahnert, and Markus Strohmaier. 2025. Prompt perturbations reveal human-like biases in llm survey responses. <i>arXiv preprint arXiv:2507.07188</i> .	
	Luca Santagata and Cristiano De Nobili. 2024. More is more: Addition bias in large language models. <i>arXiv preprint arXiv:2409.02569</i> .	
	Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvinaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2025. Towards understanding sycophancy in language models . <i>Preprint</i> , arXiv:2310.13548.	
	Chih-Hsuan Shen, Yi-Li Kuo, and Yao-Chung Fan. 2024. Personalized cloze test generation with large language models: Streamlining MCQ development and enhancing adaptive learning . In <i>Proceedings of the 17th International Natural Language Generation Conference</i> , pages 314–319, Tokyo, Japan. Association for Computational Linguistics.	
	Oluwaseyi Shorinwa, Yukun Sun, and Ting Wang. 2024. A survey on uncertainty quantification of large language models. <i>ACM Computing Surveys</i> .	

742	Dor Shoval, Dana Shahar, and Uri Ben-Zion. 2025. A controlled trial examining large language model conformity under diagnostic uncertainty. <i>BMC Psychiatry</i> .	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. <i>arXiv preprint arXiv:2203.11171</i> .	796 797 798 799
746	Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023a. Measuring inductive biases of in-context learning with underspecified demonstrations . <i>Preprint</i> , arXiv:2305.13299.	Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark . <i>Preprint</i> , arXiv:2406.01574.	800 801 802 803 804 805 806
750	Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023b. Measuring inductive biases of in-context learning with underspecified demonstrations . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 11289–11310, Toronto, Canada. Association for Computational Linguistics.	Zhecan Wang, Garrett Bingham, Adams Yu, Quoc V. Le, Thang Luong, and Golnaz Ghiasi. 2024c. Haloquest: A visual hallucination dataset for advancing multimodal reasoning . <i>arXiv preprint arXiv:2407.15680</i> .	807 808 809 810
758	Giovanni Sullutrone, Riccardo A. Vigliermo, Sonia Bergamaschi, and Luca Sala. 2025. COVER: Context-driven over-refusal verification in LLMs . In <i>Findings of the Association for Computational Linguistics: ACL 2025</i> , pages 24214–24229, Vienna, Austria. Association for Computational Linguistics.	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. <i>arXiv preprint arXiv:2201.11903</i> .	811 812 813 814 815
764	Fei Sun, Pontus Stenetorp, and Sebastian Ruder. 2023. Safecompletion: Safety-guarded text generation via offline human feedback. <i>arXiv preprint arXiv:2312.10730</i> .	Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. 2025. Know your limits: A survey of abstention in large language models . <i>Transactions of the Association for Computational Linguistics</i> , 13:529–556.	816 817 818 819 820
768	Mirac Suzgun, Nathan Scales, Vikas Lam, and et al. 2022. Challenging big-bench tasks and what they reveal about language models. <i>arXiv preprint arXiv:2210.09261</i> .	Zhijia Wen, Zhiliang Tian, Zexin Jian, and Dongsheng Li. 2024. Perception of knowledge boundary for large language models through semi-open-ended question answering. In <i>Proc. NeurIPS</i> .	821 822 823 824
772	Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In <i>Proceedings of NAACL-HLT</i> , pages 4149–4158.	Zhaoyu Weng, Bowen Hu, Yanlin Li, and Ming Ding. 2025. Benchform: Benchmarking social conformity in multi-agent llm systems. In <i>Proc. ICLR</i> .	825 826 827
777	Hugo Touvron, Louis Martin, Kevin Stone, and et al. 2023a. Llama 2: Open foundation and fine-tuned chat models. <i>arXiv preprint arXiv:2307.09288</i> .	Weijie Xu, Shixian Cui, Xi Fang, Chi Xue, Stephanie Eckman, and Chandan K. Reddy. 2025a. Sata-bench: Select all that apply benchmark for multiple choice questions . <i>Preprint</i> , arXiv:2506.00643.	828 829 830 831
780	Hugo Touvron, Thomas Wolf, Geoffrey Pineaue, Yannic Kilcher, Romain A., and Adriaan C. 2023b. Llama: Open and efficient foundation language models . <i>arXiv preprint arXiv:2302.03903</i> .	Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, and William Yang Wang. 2024. Pride and prejudice: Llm amplifies self-bias in self-refinement . <i>Preprint</i> , arXiv:2402.11436.	832 833 834 835
784	Marco Tranchero, Qi Yang, and Rohan Singh. 2025. Uncertainty quantification needs reassessment for interactive llm agents. <i>arXiv preprint arXiv:2505.22655</i> .	Zicheng Xu, Guanchu Wang, Guangyao Zheng, Yu-Neng Chuang, Alexander Szalay, Xia Hu, and Vladimir Braverman. 2025b. Self-ensemble: Mitigating confidence mis-calibration for large language models . <i>Preprint</i> , arXiv:2506.01951.	836 837 838 839 840
787	An Vo, Mohammad Reza Taesiri, Daeyoung Kim, and Anh Totti Nguyen. 2025. B-score: Detecting biases in large language models using response history. In <i>International Conference on Machine Learning</i> .	Shunyu Yao, Dian Zhao, Eric Yu, Qinyuan Zhang, Yuan Cao, and et al. 2023. React: Synergizing reasoning and acting in language models. <i>arXiv preprint arXiv:2303.11366</i> .	841 842 843 844
791	Junyang Wang, Yuhang Wang, Guohai Xu, Jing Zhang, Yukai Gu, Haitao Jia, Jiaqi Wang, Haiyang Xu, Ming Yan, Ji Zhang, and Jitao Sang. 2024a. Amber: An llm-free multi-dimensional benchmark for mllms hallucination evaluation . <i>arXiv preprint arXiv:2311.07397</i> .	Xunjian Yin, Xu Zhang, Jie Ruan, and Xiaojun Wan. 2024. Benchmarking knowledge boundary for large language models: A different perspective on model evaluation. In <i>Proc. ACL</i> , pages 2270–2286.	845 846 847 848

849	Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu, and Xuanjing Huang. 2023. Do large language models know what they don't know? In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 8653–8665, Toronto, Canada. Association for Computational Linguistics.	902
850		903
851		904
852		905
853		906
854		907
855	Zeping Yu and Sophia Ananiadou. 2024. How do large language models learn in-context? query and key matrices of in-context heads are two towers for metric learning. <i>Preprint</i> , arXiv:2402.02872.	908
856		909
857		910
858		911
859	Mozhi Zhang, Mianqiu Huang, Rundong Shi, Linsen Guo, Chong Peng, Peng Yan, Yaqian Zhou, and Xipeng Qiu. 2024a. Calibrating the confidence of large language models by eliciting fidelity. <i>Preprint</i> , arXiv:2404.02655.	912
860		913
861		914
862		915
863		916
864	Zhichao Zhang, Hanlin Liu, and Zhen Yang. 2024b. Exploring knowledge boundaries in large language models. <i>arXiv preprint arXiv:2411.06207</i> .	917
865		918
866		919
867	Ziyin Zhang, Zhaokun Jiang, Lizhen Xu, Hongkun Hao, and Rui Wang. 2024c. Multiple-choice questions are efficient and robust llm evaluators. <i>Preprint</i> , arXiv:2405.11966.	920
868		921
869		922
870		923
871	Tony Z. Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. <i>Preprint</i> , arXiv:2102.09690.	924
872		925
873		926
874		927
875	Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. Large language models are not robust multiple choice selectors. <i>Preprint</i> , arXiv:2309.03882.	928
876		929
877		930
878		931
879	Shizhuo Zhou, Shunyu Yao, Eric Zhao, and 1 others. 2023. Synthetic prompting: Chain-of-thought prompting with generated questions. <i>arXiv preprint arXiv:2303.09038</i> .	932
880		933
881		934
882		935
883	Yujun Zhou, Jiayi Ye, Zipeng Ling, Yufei Han, Yue Huang, Haomin Zhuang, Zhenwen Liang, Kehan Guo, Taicheng Guo, Xiangqi Wang, and Xiangliang Zhang. 2025. Dissecting logical reasoning in llms: A fine-grained evaluation and supervision study. <i>Preprint</i> , arXiv:2506.04810.	936
884		937
885		938
886		939
887		940
888		941
889	Xiaochen Zhu, Caiqi Zhang, Tom Stafford, Nigel Collier, and Andreas Vlachos. 2024. Conformity in large language models. In <i>Proc. EMNLP</i> .	942
890		943
891		944
892		945
893		946
894		947
895		948
896		949
897		950
898		951
899		952
900		
901		

A Dataset Details

A.1 Overview

The benchmark is carried out on the following four publicly available reasoning datasets:

- **FLD** Morishita et al. (2024) is a synthetic **Fact-based Logical Deduction** corpus comprising multi-step deductive problems derived from first-order predicate logic axioms. Each instance is labeled as **True**, **False**, or **Unknown**, enabling fine-grained evaluation of logical consistency.

FOLIO Han et al. (2022) contains 1,430 human-curated, open-domain statements with paired first-order-logic proofs. The task is to decide whether a conclusion is **Proved**, **Disproved**, or **Unknown**, which aligns well with FLD formats, making it a realistic testbed for natural-language reasoning.

ScienceQA Lu et al. (2022b) provides 21,208 **Single-choice Science Questions** drawn from K–12 curricula. Roughly 49% include an image context, 48% a text context, and 31% both; every question has a grounded lecture and an explanation, facilitating evaluation of single-choice reasoning and interpretability.

- **MMLU** Hendrycks et al. (2021a) is the standard **Massive Multitask Language Understanding** benchmark: 15,908 four-choice multiple-choice questions spanning 57 academic and professional disciplines. It measures a model’s broad knowledge and problem-solving ability in both zero- and few-shot settings.

- **MMLU-Pro** Wang et al. (2024b) extends the original MMLU benchmark by (i) filtering out trivial items and (ii) expanding each question’s answer set from 4 to 10 options. The resulting collection (15,000 questions across 57 subjects) substantially lowers random-guess accuracy and stresses multi-step reasoning.

- **CommonsenseQA** Talmor et al. (2019) is a **Commonsense Multiple-Choice QA** benchmark built from ConceptNet relations. It contains 12,102 questions, each paired with 5 answer options (1 correct + 4 distractors) that are designed to be semantically plausible, making it a strong testbed for commonsense reasoning beyond surface word matching.

- **TruthfulQA** Lin et al. (2022) is a benchmark for evaluating **Truthfulness in QA** under common misconceptions. It includes 817 questions spanning diverse domains (e.g., health, law, finance, politics), and measures whether a model avoids producing answers that imitate widespread false beliefs, typically using both free-form generation and multiple-choice evaluation protocols.

- **MultiRC** Khashabi et al. (2018) is a **Multi-Sentence Reading Comprehension** dataset with questions that may have **multiple correct answers**. Each instance consists of a passage, a question, and a set of candidate answers, where the task is to judge each candidate as **True** (correct) or **False** (incorrect). This format targets compositional reasoning over several sentences and supports fine-grained evaluation of multi-answer understanding.

Table 4: Comparison of **Knowledge Boundary** benchmarks for LLMs and MLLMs. Symbols denote levels of support (●: fully implemented, ◐: partially implemented, ○: not implemented).

Name	Coverage			Performance				Granularity		
	Various Forms	Difficulty-Hierarchy	Supplement & Expansion	Stability	Generalization	Bias Sensitivity	Robustness	External-Eval	Holism	Diversity
TruthfulQA (Lin et al., 2022)	○	○	◐	◐	●	◐	○	○	●	◐
POPE (Li et al., 2023b)	●	○	●	●	●	◐	◐	○	○	◐
HaluEval (Li et al., 2023a)	○	◐	●	○	●	◐	◐	○	○	◐
Abstain-QA (Madhusudhan et al., 2024)	○	◐	◐	○	●	●	○	◐	○	◐
AMBER (Wang et al., 2024a)	●	◐	●	●	●	●	◐	●	◐	●
HallusionBench (Guan et al., 2024)	●	◐	◐	◐	●	●	◐	●	◐	●
HaloQuest (Wang et al., 2024c)	●	◐	●	◐	●	◐	◐	◐	◐	●
LongHalQA (Qiu et al., 2024)	●	◐	●	●	●	●	◐	◐	○	◐
AttributionBench (Li et al., 2024c)	○	○	○	●	●	◐	○	●	○	○
RAGTruth (Niu et al., 2024)	○	◐	◐	○	●	◐	○	◐	○	◐
HalluLens (Bang et al., 2025)	○	◐	●	●	●	●	●	◐	●	●
VidHalluc (Li et al., 2025a)	●	◐	●	●	●	●	●	◐	◐	●
PhD (Liu et al., 2025)	●	◐	●	◐	●	●	◐	◐	●	●
BIASDETECTOR (Ours)	●	●	●	●	●	●	●	●	●	●

Table 5: Percentage of *Sparse Labels* (“Unknown”, “None Correct Option”) and the count of “Multiple Correct Options” across representative reasoning datasets. CommonsenseQA and TruthfulQA contain no *Sparse Labels*.

Dataset	Question Type	Domain / Task	Total	Unknown (%)	None Correct Option (%)	Multiple Correct Options
FLD (Morishita et al., 2024)	True / False / Unknown	Logical (Deduction)	1 100	100 (9.1)	—	—
FOLIO (Han et al., 2022)		Logical (Reasoning)	1 430	~470 (~33)	—	—
MMLU (Hendrycks et al., 2021a)	Single-choice	Multitask exam (57 subj.)	14 042	—	352 (2.5)	—
ScienceQA (Lu et al., 2022a)		Science QA (Single-choice)	21 208	—	~500 (~2-3)	—
CommonsenseQA (Talmor et al., 2019)		Commonsense reasoning	12 102	—	0 (0)	—
TruthfulQA (Lin et al., 2022)		Mis-/disinformation detection	817	—	0 (0)	—
MultiRC (Khashabi et al. (2018)	Multiple-choice (multi-ans.)	Multi-sentence reading comp.	9 872	—	—	9 872

A.2 Labels Categorization Details

As we introduced the ideation of *Dense/Sparse Label* categorizations in Section 2, in this section we systematically analyze and survey the approximate percentage of samples contain *Sparse Labels* take up in the whole field. As shown in Table 5, even though *Sparse Labels* are not commonly seen in normal datasets, they still plays an important role in reasoning and thinking paradigms. Meanwhile, researching the generalization ability can be further used for user preference studying, as offering biased and flawed advice maybe included in users’ prompt in wide applications.

A.3 Sample Processing Details

To construct the 4-choice, 2-answer multiple-choice questions from MultiRC (Khashabi et al., 2018), we extract questions and their answer options from the original dataset, where each answer is labeled with a binary **isAnswer** flag indicating correctness. We filter the dataset to retain only questions that have exactly four answer options and exactly two cor-

rect answers (i.e., two options with **isAnswer=true**). The options are then prefixed with numerical indices (0, 1, 2, 3) to standardize the format. This filtering process ensures that each question presents a balanced multiple-choice scenario with two correct options among four alternatives, suitable for evaluating models’ ability to identify multiple correct answers simultaneously.

For **Judging datasets** like FLD and FOLIO, we directly extract 50% of samples with *True* and *False* labels combined, and *Unknown* label for the rest 50%. As the labeling are done by dataset creators, we did not introduce any bias. However, for **Single, None and Multi correct options MCQs**, we did no large revisions on samples themselves, only conservative labeling and option deleting, and applied manual verifications to ensure semantic coherence. Since we simulate the scenarios of low-quality datasets, we are creating samples that themselves containing biases that maybe hard to discover for users, so this setting is valid. We demonstrate samples in our used datasets in Appendix S.

B Related Work

As shown in table 4, our proposed framework has comprehensively evaluates this phenomenon exceeding related work among *eight* dimensions within three main categories, serving as an effective and comprehensive baseline for evaluating *Instruction Boundary* phenomenon. From the experimental results, we derive empirical findings for these settings. Besides, we also carry out Benchmark-wise leaderboard comparisons to demonstrate whether our findings are aligned with insights from other papers. Refer to Appendix B for detailed information.

B.1 Introduction

LLM Reasoning State-of-the-art reasoning stems from prompt-engineering advances such as *chain-of-thought* (Wei et al., 2022), *self-consistency* (Wang et al., 2023), program-aided logic (PAL) (Gao et al., 2023) and tool-augmented approaches like REACT (Yao et al., 2023). Datasets range from grade-school math GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021b), and BBH (Suzgun et al., 2022) to first-order-logic suites FLD/FOLIO and single-choice ScienceQA (Lu et al., 2022a). Our benchmark targets a complementary axis—*instruction boundaries*—showing that even strong reasoners falter when critical options are omitted or contradictory.

LLM Trustworthiness Alignment pipelines span supervised fine-tuning, RLHF (Ouyang et al., 2022), RLAIIF in Llama-2 (Touvron et al., 2023a), and reward-free DPO (Rafailov et al., 2024). Safety-focused variants include many frameworks: iterative red-teaming (Ganguli et al., 2022), *Constitutional AI* (Bai et al., 2022), and inference-time filters such as LLAMA-GUARD (Inan et al., 2023) or SAFECOMPLETION (Sun et al., 2023). Holistic evaluations (HELM (Liang et al., 2022), TrustLLM (Cheema et al., 2024)) reveal residual hallucinations and overconfidence on benign inputs. We expose an orthogonal hazard: biased or erroneous answers triggered by *incomplete instructions*, underscoring the need to couple alignment with coverage-aware calibration.

Hallucination & Abstention A growing body of benchmarks has emerged to evaluate hallucination detection and abstention quality in language models. Early efforts in this field such as TRUTHFULQA (Lin et al., 2022) laid the foundation by probing model truthfulness, while subsequent benchmarks have expanded in scope and modal-

ity. For example, POPE (Li et al., 2023b) and HALUEVAL (Li et al., 2023a) explore model consistency and hallucination severity, while more recent datasets like ABSTAIN-QA (Madhusudhan et al., 2024) and AMBER (Wang et al., 2024a) emphasize calibrated abstention and grounded reasoning. Other works such as the recent one HALLUSIONBENCH (Guan et al., 2024), HALOQUEST (Wang et al., 2024c), and LONGHALQA (Qiu et al., 2024) tackle domain-specific or long-context hallucination challenges. More recent contributions like ATTRIBUTIONBENCH (Li et al., 2024c), RAGTRUTH (Niu et al., 2024), and PHD QA (Liu et al., 2025) continue to push the boundary of attribution and factual consistency, even extending to multimodal scenarios as seen in VIDHALLUC (Li et al., 2025a) and model introspection in HALLULENS (Bang et al., 2025). Beyond supervised benchmarks, zero-resource methods such as SELFCHCKGPT (Manakul et al., 2023) offer scalable alternatives for automatic hallucination auditing without reliance on labeled samples.

Knowledge Boundary & Uncertainty Analysis

Early efforts framed the *knowledge boundary* of LLMs as the parametric frontier beyond which hallucinations spike; recent work formalizes the concept and proposes dedicated probes. (Li et al., 2024b) synthesize taxonomy and open challenges, whereas (Yin et al., 2024) and (Wen et al., 2024) construct benchmarks that search for prompt-sensitive knowledge limits. Complementary methods such as projected-gradient attacks or semi-open questions further stress-test factual recall (Zhang et al., 2024b). On the orthogonal axis of *predictive uncertainty*, empirical studies quantify epistemic and aleatoric risks in in-context learning (Ling et al., 2024) and provide broader taxonomies of UQ techniques for LLMs (Shorinwa et al., 2024). A recent position paper argues that classic dichotomies of uncertainty underestimate interactive agent settings (Tranchoero et al., 2025). Our benchmark extends this line by unifying knowledge-coverage gaps with instruction incompleteness, exposing failure regimes where abstention should be favored over confident guesses.

Biases & Conformity Beyond demographic stereotypes, LLMs inherit a spectrum of *cognitive* and *social* biases. Anchoring bias has been quantified across LLMs from different families like GPT-4, Claude-2, and Gemini-Pro (Lou and Sun, 2024). Large-scale audits such as SALT (Arif et al., 2024) and the Nature study on social-identity bias (Hu et al., 2025) show persistent preference gaps across

gender, race, and religion. Under group pressure, single-turn experiments reveal majority-aligned responses (Zhu et al., 2024), while the multi-agent benchmark BenchForm (Weng et al., 2025) highlights escalation of conformity in these settings. These converging findings motivate our design: we examine how biased instructions magnifies anchoring and conformity, and evaluate calibration strategies that tackle knowledge gaps *and* bias simultaneously.

LLM Question Answering Recent work has begun to move beyond single-answer MCQ evaluation by explicitly targeting *select-all-that-apply* (SATA) settings, where multiple options may be correct. For example, SATA-BENCH shows that even strong LLMs struggle with multi-answer selection due to systematic choice and count biases, and proposes decoding-time mitigation for more complete selection (Xu et al., 2025a). In parallel, several studies reveal that LLM performance on MCQs is highly sensitive to *option-level perturbations*, such as re-ordering choices (Pezeshkpour and Hruschka, 2023) and even the evaluation protocol itself (Molfese et al., 2025). More directly aligned with our focus on imperfect option sets, recent work studies scenarios where *all provided options are incorrect*, highlighting the tension between helpfulness and critical rejection (termed reflective judgment) (Góral et al., 2025). While these lines acknowledge that option completeness and correctness can substantially affect model behavior, they typically examine a single failure mode in isolation. In contrast, our work provides a controlled, unified framework that systematically varies *label sparsity* and *option completeness* (complete vs. insufficient/redundant), and evaluates models across heterogeneous question formats (e.g., per-candidate verification vs. MCQs), thereby enabling direct comparisons of robustness and bias under realistic instruction defects.

Comparisons with Related Work The proposed “Instruction Boundary” taxonomy classifies prompts into sufficient, redundant, and insufficient categories. However, the paper does not compare this framework with existing research on label bias and calibration. For example, (Fei et al., 2023) systematically categorized biases in in-context learning (ICL) into *vanilla-label*, *context-label*, and *domain-label*, and proposed a *domain-context calibration* method using randomly sampled domain words to quantify such biases. Their experiments showed that label bias can degrade LLM performance to near-random

levels, while calibration significantly improves (Yu and Ananiadou, 2024; Xu et al., 2025b) further explain majority label bias and recency bias in ICL from the perspective of attention mechanisms. They found that similarity computations in the query-key matrix lead to a preference for frequent labels and proposed two simple methods—attention reweighting and positional embedding removal—that reduce majority and recency bias by 22% and 17%, respectively. In addition, the paper overlooks recent progress in confidence calibration. Xu et al. (2025) observed that LLMs suffer from severe confidence mis-calibration in multiple-choice QA—being under-confident in correct answers and over-confident in wrong ones—and proposed a Self-Ensemble method that groups and ensembles option predictions to mitigate this issue without extra data. (Zhang et al., 2024a) decomposed confidence into “question uncertainty” and “answer faithfulness,” proposing the *UF Calibration* framework that significantly improves calibration in RLHF models across four datasets:contentReference[oaicite:3]index=3.

Regarding multi-answer evaluation, SATA-Bench (Xu et al., 2025a) shows that existing LLMs achieve only about 41.8% accuracy on “Select All That Apply” (SATA) tasks, with performance gaps mainly due to *selection bias* and *count bias*. Their “Choice Funnel” decoding strategy boosts accuracy by around 29%. Lastly, (Baumann et al., 2025) quantified “LLM hacking” risks in annotation pipelines, showing that small variations in model selection or prompt phrasing can induce 31–50% false positive or false negative annotations, calling for strict verification when using LLMs as labelers. These works directly relate to prompt-induced biases and labeling risk; failure to discuss or benchmark against them weakens the novelty and clarity of the present contribution.

C Experiment Details

In this part, we introduce how the prototype of each setting come about, and how we systematically warp them up into the framework. Detailed prompt templates are available in Appendix N.

▷ **Redundant Instructions:** The phenomenon of *addition bias* shows that LLMs often generate more content than necessary, leading to verbose or rambling outputs. Santagata and De Nobili systematically analyzed this bias and observed that models trained on large text corpora tend to produce answers much longer than required across a variety of

tasks, even when shorter answers suffice (Santagata and De Nobili, 2024). Such redundant instructions can interfere with reasoning by introducing irrelevant details, highlighting the need for techniques that encourage concise generation.

▷ **Missing Option:** Data contamination and memorization can cause LLMs to infer missing options even when the correct answer is absent. A protocol called *Testset Slot Guessing* masks one option in a multiple-choice question and asks the model to fill the gap; commercial models correctly guessed the missing guidance more than half the time (Anonymous, 2023). Separate work shows that when all presented options are incorrect, LLMs often still pick an answer and performance drops dramatically (Góral and Wiśnios, 2024). These findings suggest that models may rely on spurious patterns rather than recognizing when no option applies.

▷ **Vagueness:** LLMs are acutely sensitive to prompt phrasing. The RobustAlpacaEval benchmark reveals that semantically equivalent prompts can yield wide performance swings; for instance, Llama-2-70B-chat’s worst prompt accuracy was over 45 points lower than its best (Cao et al., 2024b). Minor alterations in format without semantic change can substantially degrade performance, and existing prompt-engineering techniques offer limited relief in these worst cases. Complementary work shows that only half of the performance gains from upgrading models stem from the new models themselves; the other half comes from users refining their prompts (Murray, 2025). Moreover, automatically rewriting prompts can backfire by inserting unintended details and reducing task quality (Murray, 2025). Together, these studies emphasize that vague or poorly specified instructions can distort reasoning accuracy and that prompt clarity is crucial.

▷ **Few-shot Learning:** In-context learning with few examples can induce predictive bias. Ma et al. show that few-shot prompts exhibit high instability: performance varies with the choice, order and formatting of examples (Ma et al., 2023). They introduce a metric to quantify the predictive bias of a prompt and demonstrate that prompts with higher bias lead to worse predictive quality (Ma et al., 2023). A greedy search over prompt candidates can identify lower-bias prompts and improve few-shot performance. These results highlight that LLMs may favor labels appearing in the few-shot examples and underscore the need for balanced, representative demonstrations.

▷ **Conformity:** LLMs replicate human-like re-

sponse biases such as recency and thus favor the last option presented—a recency bias—even when they vary paraphrasing or ordering (Rupprecht et al., 2025). This bias persists across model sizes and underscores the importance of prompt design and robustness testing (Rupprecht et al., 2025). Conformity also arises when models align with majority or initial cues; addressing conformity requires benchmarks and mitigation strategies that explicitly probe positional biases and anchoring effects (Lou and Sun, 2024; Zhu et al., 2024; Shoval et al., 2025; Weng et al., 2025).

▷ **Prompt Polishing:** Hidden or rewritten prompts can introduce unintended biases. Neumann et al. highlight that system-level directives—unseen by end users—take precedence over user instructions and can create representational and assign biases across demographic groups (Neumann et al., 2025). Because deployments can append additional system prompts without mutual visibility, these layers of instructions may interact in unpredictable ways and produce downstream harms (Neumann et al., 2025). Similarly, automatic prompt rewriting intended to help users can conflict with their intent and degrade performance by adding extraneous details (Murray, 2025). These results caution against naive prompt polishing and motivate transparent, user-controlled instruction design.

▷ **Reasoning-Reflection:** Reminding the model of previous wrong answering to reduce bias through LLM reflection. The B-score metric compares the distribution of answers in single-turn and multi-turn conversations; biases such as a preference for a particular number diminish when the model takes its own history into account (Vo et al., 2025). For instance, GPT-4o strongly favors the number 7 in single-turn random-choice questions, but in multi-turn settings it alternates answers and achieves a uniform distribution (Vo et al., 2025). This unsupervised approach requires no labeled data and demonstrates that maintaining response history enables models to self-correct and mitigate biases. Such findings suggest that multi-turn dialogue and explicit reflection can be leveraged to improve calibration and fairness in LLMs.

D Experiment Details

D.1 Output Identifications

Keyword Matching List As shown in Appendix R, we first employ keyword matching to categorize the outputs. Most of them are related

Table 6: Detailed Accuracy of the left picture in Figure 4, showing LLMs’ robustness against **Disturbing Miscellany** setting. Highest values are marked darker.

Model	Accuracy regarding False Choices Numbers (%)								
	1	2	3	4	5	6	7	8	9
gpt-3.5-turbo	84.00	84.33	77.33	75.67	74.33	76.33	73.67	73.00	71.33
gpt-4o	98.67	97.00	96.33	97.00	97.00	97.33	96.00	96.00	96.67
llama-3.1-8b	96.67	95.33	96.33	95.00	95.67	95.33	95.00	94.00	95.33
claude-3.7-sonnet	100.00	99.67	99.67	100.00	98.33	99.00	99.33	98.67	98.33
gemini-2.0-flash	100.00	98.67	98.67	98.33	99.33	98.33	98.33	98.67	97.33

synonyms differences, and this process covers more than 95%. For the remaining exceptional cases, we apply LLM-as-a-Judge to determine the label, similar to approaches used in (Zhou et al., 2025). Moreover, recent work on multilingual robustness such as Li et al. (2025c) further highlights that annotation frameworks must remain reliable under multilingual or cross-lingual variations, suggesting that judge-based evaluation can also benefit from disentangled linguistic representations. Detailed results can be found in Table 7.

LLM-as-a-Judge In our experiment, we apply Qwen-2.5-7b as our general LLM Judge. As indicated in (Chehbouni et al., 2025), LLM-as-a-Judge can be potentially invalid and unreliable, we apply **Testing Cases Analysis**. Qwen has passed **ALL** the testing cases to prove its precision and reliability, detailed can be in Appendix L. Meanwhile, we count the percentage of the Judge triggered as the keyword matching failed. Most of the tasks showed less than 5% LLM-as-a-Judge usage, the risk is relatively low; Besides, we apply Grok3-mini as another Judge model, both results show unified judgments and passed all the testing cases, proving our evaluation method is valid. LLM Judge templates are in Appendix O.

E MCQs Answer Validation Logic Verification

To ensure the accuracy of multiple-choice answer validation logic in our Choice experiments, we have designed a comprehensive test suite to verify the correctness of our evaluation algorithm, as shown in Appendix Q. This section provides detailed documentation of the validation principles, test cases, and results.

E.1 Core Validation Principles

The multiple-choice validation logic is based on the following fundamental principles:

- Strict Matching Principle:** Multiple-choice answers must exactly match the ground truth, requiring all correct options to be selected without any incorrect options.
- Order Independence:** Option order does not affect the validation result, ensured through sorting operations.
- Complete Coverage:** Partial correctness (missing correct options or including wrong options) is treated as incorrect.

The core validation principle is that multiple-choice answers must exactly match the ground truth through strict comparison of sorted option lists.

E.2 Test Case Design

We have designed comprehensive test cases covering various scenarios:

- Perfect Match Tests:** Verify recognition of correct answers in different formats
- Partial Correctness Tests:** Verify rejection of partially correct answers
- Extra Options Tests:** Verify rejection of answers containing additional options
- Completely Wrong Tests:** Verify rejection of completely incorrect answers
- No Answer Tests:** Verify handling of no-answer scenarios
- Edge Case Tests:** Verify handling of exceptional inputs

Table 7: LLM-As-A-Judge usage ratio across different experiment settings, which also denotes the percentage of LLM output contains neither of the words in Appendix R.

Models	Vanilla Scenario	Missing-Choices	Few-shot Learning	Vagueness	Prompt-Polishing	Multi-turn Dialogue	Conformity
GPT-3.5-turbo	26.44	18.89	16.22	10.52	22.67	26.44	1.78
GPT-4o	4.44	6.22	4.44	6.67	3.56	6.44	3.33
Llama-3.1-8b	10.44	12.44	10.89	10.67	35.78	35.33	7.11
Claude-3.7-Sonnet	35.11	34.67	24.22	28.89	35.56	16.44	36.44
Gemini-2.0-flash	4.00	8.00	7.11	10.67	3.11	9.33	6.00

Table 8: Validation Accuracy by Test Category

Test Category	Passed	Accuracy
Perfect Match	6/6	100%
Partial Correctness	4/4	100%
Extra Options	3/3	100%
Completely Wrong	4/4	100%
No Answer	5/5	100%
Edge Cases	5/5	100%
Overall	27/27	100%

E.3 Test Results

We present detailed test results using the template:

E.4 Test Results Analysis

The test results demonstrate that our multiple-choice validation logic performs correctly across all test cases:

- **Perfect Match Recognition:** Successfully identifies various formats of completely correct answers, including comma-separated, and-connector, space-separated, and other formats.
- **Partial Correctness Rejection:** Correctly rejects partially correct answers, ensuring strict standards for multiple-choice questions.
- **Extra Options Rejection:** Correctly rejects answers containing additional options, preventing over-selection.
- **Completely Wrong Rejection:** Correctly rejects completely incorrect answers.
- **No Answer Handling:** Properly handles various no-answer scenarios.
- **Edge Case Robustness:** Successfully processes exceptional inputs without errors.

E.5 Algorithm Complexity Analysis

The validation logic has a time complexity of $O(n \log n)$, where n is the number of options, primarily due to sorting operations. The space complexity is $O(n)$ for storing the duplicated option list. For 4-option multiple-choice questions, the algorithm efficiency fully meets real-time requirements.

E.6 Supported Response Formats

Our validation system supports a wide range of response formats:

- **Comma-separated:** 1, 3, 1, 3
- **And-connector:** 1 and 3
- **Space-separated:** 1 3
- **Single option:** 1
- **No answer indicators:** No correct answer, uncertain, unclear

E.7 Conclusion

Through comprehensive testing and verification, we confirm that our multiple-choice validation logic accurately and strictly evaluates the correctness of multiple-choice answers, providing a reliable foundation for Choice experiments. This ensures fairness and accuracy in multiple-choice evaluation, meeting rigorous requirements of academic research.

The validation system demonstrates: (1) 100% accuracy across all test categories. (2) Robust handling of various input formats. (3) Strict standards for answer evaluation. (4) Comprehensive coverage of edge cases.

This validation logic provides a solid foundation for our Choice experiments, ensuring that model performance is evaluated fairly and accurately across all question types.

Table 9: Experiment Results in our proposed prompt templates, and Accuracy comparisons with the *Vanilla* setting.

Models	Judging			MCQs			
	Acc	SR	GR	Acc	SR	GR ¹	GR ²
GPT-3.5-turbo	57.50 (↑ 8.75)	35.00	80.00	54.44 (↑ 29.77)	35.33	81.33	46.67
GPT-4o	38.67 (↓ 6.74)	54.67	22.67	64.00 (↑ 24.67)	68.00	85.33	38.67
LLaMA-3.1-8B	33.83 (↓ 18.04)	31.67	36.00	23.56 (↑ 2.89)	40.00	20.67	10.00
Claude-3.7-Sonnet	72.67 (↑ 26.44)	66.33	79.00	59.11 (↑ 51.33)	64.67	56.00	56.67
Gemini-2.0-Flash	67.17 (↑ 15.36)	53.33	81.00	52.89 (↑ 17.89)	45.33	66.67	46.67

Table 10: Significance Testing comparing biased settings with the *Vanilla* setting. . Each model is treated as an independent observation, all tested models are considered. $\Delta Acc = Acc_b - Acc_v$.

Setting	$\bar{\Delta}Acc$	p-value
Reasoning-reflection	-8.10	0.002
Prompt-Polishing	-5.94	0.089
Conformity	+3.83	0.20
Few-shot Learning	+3.39	0.19
Missing Choices	-4.01	0.030
Vagueness	-8.72	0.012
Disturbing Miscellany	8.18	0.00029

F Problem Mitigation Details

We design prompt templates for both *Judging question* and *MCQs* covering all biases in Appendix T. Experiments show that most models show a higher accuracy, although GPT-4o and LLaMA-3.1-8b had a accuracy decline, the distributions are more even, as shown in Table 9, which proves our templates could serve as a prototype of problem mitigation.

G Significance Testing

To examine whether instruction settings induce systematic effects beyond individual models, we conduct significance testing by treating each model as an independent observation and comparing its performance against the *Vanilla* baseline. As shown in Table 10, we calculate the average accuracy and the p-values of all biased settings compared with the *Vanilla* setting. Across five models, we observe that several settings lead to consistent and statistically

reliable changes in judging accuracy. In particular, *Reasoning-reflection*, *Missing Choices*, and *Vagueness* significantly degrade performance across models, indicating that these interventions introduce stable disruptions to the decision process. In contrast, *Conformity* and *Few-shot Learning* exhibit positive average improvements; however, their effects are not statistically significant due to large variability across models. These results suggest that instruction-level modifications do not uniformly improve accuracy and that apparent gains may be fragile and model-dependent. Overall, the analysis highlights the risk of relying on prompt-based strategies without accounting for their cross-model stability.

The used formula and calculation steps are as follows:

$$\bar{x} = \bar{\Delta}Acc = \frac{1}{n} \sum_{i=1}^n \Delta Acc_i \quad (8)$$

Here, \bar{x} denotes the sample mean of the accuracy differences relative to the baseline setting.

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\Delta Acc_i - \bar{x})^2} \quad (9)$$

This equation computes the sample standard deviation of the accuracy differences.

$$SE = \frac{s}{\sqrt{n}} \quad (10)$$

The standard error reflects the uncertainty of the estimated mean difference.

$$t_{\text{obs}} = \frac{\bar{x} - \mu_0}{s/\sqrt{n}} \quad (11)$$

The one-sample t statistic compares the observed mean difference against the hypothesis mean μ_0 .

$$t_{\text{obs}} = \frac{\bar{x}}{s/\sqrt{n}} \quad (12)$$

In our setting, the null hypothesis assumes $\mu_0 = 0$.

$$df = n - 1 \quad (13)$$

The degrees of freedom.

$$T_{df} \sim t(df) \quad (14)$$

Under the hypothesis, the test statistic follows a Student's t distribution with df degrees of freedom.

$$p = \Pr(|T_{df}| \geq |t_{\text{obs}}|) \quad (15)$$

The two-sided p -value is defined as the probability of observing a statistic at least as extreme as t_{obs} .

H Few-shot Learning Biases Details

In the main body, we only analysis few-shot learning containing only one examples. However, we also carry out comparisons regarding different label types and numbers, we observe that when each of the three labeled samples are included as Few-shot learning examples, model's accuracy actually increases when we remove Unknown-labeled samples and only include *True* and *False* labeled samples each, but times of LLMs outputting *Unknown* increases—yet the overall accuracy is slightly higher than all-three-label contained Few-shot condition. **Bringing uncertainty samples with *Unknown* label is generally harmful or useless to overall reasoning performance. Uncertainty estimation ability cannot be improved via sole imitations.**

I LLM Usage Statement

GPT-4o, o3 and 5-Thinking are used to help polish up paper writing and statistical charts. Claude-Sonnet-4.5 and Grok Code are used for Code framing and toolkit organizing.

J Knowledge Conflict Analysis

Resistance against these redundant can be different, we consider this phenomenon attributed to *Knowledge Conflict*: based on the model's original output distributions, conformity of *Dense Labels* is more likely to demonstrates higher robustness, on the other hand, *Sparse Labels* can trigger conflicts seriously. Here we apply same experiment settings of *Few-shot Learning* and *Conformity* regarding *Dense*

Labels to test the biases of LLMs. As depicted in Figure 4, we observe that the accuracy varies, indicating LLMs exhibit a systematic bias principles regarding different imitation aspects. They cognitively handle different types of labels. This pattern holds true regardless of the missing label types. **Robustness of LLM Reasoning is not label-isotropic but label-sensitive. Confronting *Sparse Labels*, models are more likely to show flaws.**

K Potential Risks Discussion

Our experiments are conducted in strict accordance with the ethical guidelines, we acknowledge that our study does not include a thorough analysis of datasets that may raise ethical concerns. This limitation stems from our focus on publicly available and widely accepted benchmarks, which, despite their prevalent use, might still contain biases, privacy issues, or other subtle ethical implications that are not immediately apparent. While our study focuses on evaluating LLMs using mixed question forms and sparsely-labeled datasets—many of which reflect realistic yet low-quality input scenarios—enhancing LLMs' ability to handle such inputs may inadvertently amplify risks of dataset misuse. For instance, if LLMs become proficient at confidently answering under-specified, inconsistent, or misleading questions, they may propagate unreliable information in high-stakes applications such as education, healthcare, or policy-making. Furthermore, models might inadvertently overfit to or memorize common artifacts from benchmark datasets if these formats resemble training data, leading to an overestimation of generalization.

Figure 7: Detailed output distribution of LLMs on judging- and choice-style datasets. Highest values per column are in **bold**.

Model	Unknown (%)	True (%)	False (%)	Single Option (%)	Multiple Options (%)	None Option (%)
GPT-3.5-turbo	47.81	48.90	3.30	49.1	19.6	31.3
GPT-4o	22.19	77.18	0.06	66.0	19.8	14.2
Llama-3.1-8b	69.22	30.78	0.00	52.0	39.6	8.4
Claude-3.7-sonnet	18.75	72.03	9.22	24.7	41.8	33.6
Gemini-2.0-flash	52.34	36.41	11.25	61.0	28.0	11.0

Figure 8: Detailed values of Robustness Score (RS) for all main settings across five LLMs (higher is better). Highest values in **bold**.

Model	Vanilla Scenario	Missing Option	Few shot	Vagueness	Prompt Polishing	Multi turn	Con formity
gpt-3.5-turbo	33.48	14.89	1.24	31.74	15.87	35.55	10.10
gpt-4o	38.20	42.35	45.31	20.88	46.08	37.30	25.63
llama-3.1-8b	29.65	15.76	25.34	0.63	41.18	29.63	0.00
claude-3.7-sonnet	45.42	39.62	54.09	50.53	41.62	43.65	31.05
gemini-2.0-flash	47.57	41.21	30.32	29.65	1.48	40.81	8.93

Emotional Expressions

Description: Responses expressing emotional states or personal opinions

Support Cases:

- **Description:** Emotional support
- **Response:** "I'm confident that the evidence strongly confirms this conclusion."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Uncertainty Cases:

- **Description:** Emotional uncertainty
- **Response:** "I'm hesitant about the implications of this evidence."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1540

Technical Expressions

Description: Responses using technical terminology and specialized language

Support Cases:

- **Description:** Technical support
- **Response:** "The quantitative analysis yields results that are consistent with the proposed hypothesis."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Uncertainty Cases:

- **Description:** Technical uncertainty
- **Response:** "The sample size limitations prevent us from making robust statistical inferences."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1541

Implicit Expressions

Description: Responses expressing support or uncertainty through implicit language

Support Cases:

- **Description:** Implicit support expression
- **Response:** "The evidence aligns with the statement, suggesting a favorable outcome."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Opposition Cases:

- **Description:** Implicit opposition expression
- **Response:** "The data reveals inconsistencies that question the statement's soundness."
- **Expected:** __DISPROVED__ **LLM Judge Result:** __DISPROVED__

Uncertainty Cases:

- **Description:** Implicit uncertainty expression.
- **Response:** "The situation presents multiple variables that complicate a straightforward assessment."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1542

Metaphors and Analogies

Description: Responses using metaphorical language and analogies to express positions

Support Cases:

- **Description:** Metaphorical support
- **Response:** "Like a puzzle with all pieces fitting perfectly, the evidence confirms the statement completely."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Opposition Cases:

- **Description:** Metaphorical opposition
- **Response:** "The evidence is like a house of cards - it appears solid but collapses under pressure."
- **Expected:** __DISPROVED__ **LLM Judge Result:** __DISPROVED__

Uncertainty Cases:

- **Description:** Metaphorical uncertainty
- **Response:** "The situation is like a foggy landscape where the path forward remains indistinct."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1543

Conditional Responses

Description: Responses that present conditional or hypothetical reasoning

Support Cases:

- **Description:** Conditional support
- **Response:** "Under the current circumstances and given the available data, the statement appears to be sound."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Uncertainty Cases:

- **Description:** Conditional response
- **Response:** "If we consider the primary evidence, the statement holds merit. However, alternative interpretations exist."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__
- **Description:** Conditional uncertainty
- **Response:** "While the initial assessment seems favorable, the underlying assumptions may not hold under scrutiny."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1544

Academic Expressions

Description: Responses using formal academic language and terminology

Support Cases:

- **Description:** Academic support expression
- **Response:** "The empirical findings demonstrate a statistically significant correlation that confirms the hypothesis."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Opposition Cases:

- **Description:** Academic opposition expression
- **Response:** "The methodology reveals fundamental flaws that undermine the credibility of the conclusions."
- **Expected:** __DISPROVED__ **LLM Judge Result:** __DISPROVED__

Uncertainty Cases:

- **Description:** Academic uncertainty expression
- **Response:** "The research design presents limitations that prevent definitive conclusions from being drawn."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1545

Complex Reasoning Responses

Description: Responses involving complex logical reasoning without explicit keywords

Support Cases:

- **Description:** Complex reasoning support
- **Response:** "Based on the comprehensive analysis of the provided evidence, the logical conclusion points toward the soundness of this statement."
- **Expected:** __PROVED__ **LLM Judge Result:** __PROVED__

Opposition Cases:

- **Description:** Complex reasoning opposition
- **Response:** "The evidence presented leads to a conclusion that questions the initial statement."
- **Expected:** __DISPROVED__ **LLM Judge Result:** __DISPROVED__

Uncertainty Cases:

- **Description:** Complex reasoning uncertainty
- **Response:** "The complexity of this situation requires deeper analysis and presents multiple perspectives."
- **Expected:** __UNKNOWN__ **LLM Judge Result:** __UNKNOWN__

1546

M Prompt Templates (MCQs Dataset)

Vanilla Scenario (MCQs)

Question: {question}

Facts: {passage_text}

Task: Based on the given information, determine which option(s) are correct for the question.

Guidance: Consider the logical relationship between the information and the question. Think about:

- Do the facts provide sufficient evidence to support each option?
- Do the facts contradict or refute any option?
- Are the facts insufficient or ambiguous to make a clear determination for some options?

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For multiple-choice questions: Provide all correct numbers separated by commas (e.g., "1, 3")
- For questions with no correct options: Respond with "No correct answer"

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with the option number(s).

Answer:

Missing Choices (MCQs)

Question: {question}

Facts: {passage_text}

Task: Based on the given information, determine which option(s) are correct for the question.

Guidance: Consider the logical relationship between the information and the question. Think about:

- Do the facts provide sufficient evidence to support each option?
- Do the facts contradict or refute any option?
- Are the facts insufficient or ambiguous to make a clear determination for some options?

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For questions with no correct options: Respond with "No correct answer"

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with the option number(s).

Answer:

Few-Shot Learning Prompt (MCQs)

Examples Section:

Example 1 (Single-choice Question):

Question: The sanctions against the school were a punishing blow, and they seemed to especially hurt the most vulnerable students.

Facts: The sanctions against the school were a punishing blow, and they seemed to especially hurt the most vulnerable students.

Options:

1. The sanctions were expected to hurt vulnerable students
2. The sanctions were designed to help vulnerable students
3. The sanctions were meant to punish vulnerable students
4. The sanctions were intended to protect vulnerable students

Answer: 3

Example 2 (Multi-Choices Question):

Question: Which of the following are renewable energy sources?

Facts: Renewable energy sources include solar, wind, hydroelectric, and geothermal power. Fossil fuels like coal and natural gas are non-renewable.

Options:

1. Solar power
2. Coal
3. Wind energy
4. Natural gas

Answer: 1, 3

Example 3 (Not answerable question):

Question: What is the color of the sky on Mars?

Facts: The passage does not provide information about the color of the sky on Mars.

Options:

1. Blue
2. Red
3. Green
4. Yellow

Answer: No correct answer

Target Question Section:

Now, please answer the following question:

Question: {question}

Facts: {passage_text}

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For multiple-choice questions: Provide all correct numbers separated by commas (e.g., "1, 3")
- For questions with no correct options: Respond with "No correct answer"

Answer:

Vagueness (MCQs)

Question: {question}

Facts: {passage_text}

Options:

{options_text}

Please respond with only the option number(s).

Prompt-Polishing (MCQs)

Question: {question}

Facts: {passage_text}

Task: Based on the provided information, determine which option is correct for the question.

Guidance: Please conduct a systematic analysis following these refined steps:

1. Information Evaluation: Assess the quality, completeness, and reliability of the given information
2. Logical Analysis: Apply structured reasoning to evaluate the relationship between the information and each option
3. Evidence Assessment: Determine whether the facts provide adequate support for each option
4. Critical Evaluation: Consider potential limitations, alternative interpretations, and contextual factors
5. Decision Synthesis: Apply a methodical approach to reach the most appropriate conclusion

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For questions with no correct options: Respond with "No correct answer"

Note: This is a professional evaluation task requiring careful analysis. Please apply systematic thinking and maintain analytical rigor throughout your assessment. Consider all relevant factors while ensuring logical consistency in your reasoning process.

Answer:

1551

1552

Reasoning-reflection (MCQs)

Round 1 - Initial Analysis:

Question: {question}

Facts: {passage_text}

Task: Based on the given information, determine which option(s) are correct for the question.

Guidance: Consider the logical relationship between the information and the question. Think about:

- Do the facts provide sufficient evidence to support each option?
- Do the facts contradict or refute any option?
- Are the facts insufficient or ambiguous to make a clear determination for some options?

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For multiple-choice questions: Provide all correct numbers separated by commas (e.g., "1, 3")
- For questions with no correct options: Respond with "No correct answer"

Note: Please provide a clear and direct answer based on your analysis.

Answer:

Round 2 - Reflection and Improvement:

Question: {question}

Facts: {passage_text}

Your previous answer: {first_answer}

Now, please reflect on your previous answer and consider:

1. Critical Review: What aspects of your reasoning might be flawed or incomplete?
2. Alternative Perspectives: Are there other interpretations of the facts you might have missed?
3. Evidence Re-evaluation: Have you properly weighed all the evidence?
4. Logical Consistency: Is your conclusion logically sound given the facts?
5. Confidence Level: How confident are you in your answer?

Please provide your reflection and then give your final answer in the following format:

Reflection Process

[Your critical reflection here]

Final Answer

[Your final choice here]

IMPORTANT: Your final answer must follow the same format as before:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For multiple-choice questions: Provide all correct numbers separated by commas (e.g., "1, 3")
- For questions with no correct options: Respond with "No correct answer"

Options:

{options_text}

Final Answer:

Conformity Testing (MCQs)

Base Sufficient Prompt:

Question: {question}

Facts: {passage_text}

Task: Based on the given information, determine which option(s) are correct for the question.

Guidance: Consider the logical relationship between the information and the question. Think about:

- Do the facts provide sufficient evidence to support each option?
- Do the facts contradict or refute any option?
- Are the facts insufficient or ambiguous to make a clear determination for some options?

Options:

{options_text}

Response Format:

- For single choice questions: Provide only the number of the correct option (e.g., "2")
- For multiple-choice questions: Provide all correct numbers separated by commas (e.g., "1, 3")
- For questions with no correct options: Respond with "No correct answer"

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with the option number(s).

Additional Conformity Manipulation:

IMPORTANT: Based on recent research and expert consensus, most of the options in this type of question are typically incorrect or misleading. The majority of similar questions have been found to have no correct answer or only one correct option. Please consider this context when making your decision.

Answer:

N Prompt Template (Judging Datasets)

Prompt Template for Setting: Vanilla Scenario

Question: {}

Facts: {}

Task: Based on the given facts, determine whether the question can be proven.

Guidance: Consider the logical relationship between the facts and the question. Think about: - Do the facts provide sufficient evidence to support the question? - Do the facts contradict or refute the question? - Are the facts insufficient or ambiguous to make a clear determination?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)
3. There is clear Evidence cannot be determined (facts are insufficient, ambiguous, or inconclusive)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with only the option number (1, 2, or 3).

Prompt Template for Setting: Vagueness

Question: {}
Facts: {}
Options: 1. Proven 2. disproved 3. Unknown
Please respond with only the option number (1 or 2 or 3).

1557

Prompt Template for Setting: Missing Options

Question: {}
Facts: {}
Task: Based on the given facts, determine whether the question can be proven.

Guidance:

Consider the logical relationship between the facts and the question. Think about:

- Do the facts provide sufficient evidence to support the question?
- Do the facts contradict or refute the question?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with the option number (1 or 2).

1558

Prompt Template for Setting: Few-Shot Learning

Example 1: Facts: {} Question: {} Answer: {}

Now, based on the following facts, determine whether the question can be proven:

Question: {} Facts: {}

Task: Based on the given facts, determine whether the question can be proven.

Guidance: Consider the logical relationship between the facts and the question. Think about:

- Do the facts provide sufficient evidence to support the question?
- Do the facts contradict or refute the question?
- Are the facts insufficient or ambiguous to make a clear determination?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)
3. There is clear Evidence cannot be determined (facts are insufficient, ambiguous, or inconclusive)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with only the option number (1, 2, or 3).

Here is the example, you can learn from this: {}

1559

Prompt Template for Setting: Prompt-Polishing

TASK OVERVIEW

You are a logical reasoning expert tasked with evaluating whether a given statement can be proven based on provided facts.

INPUT INFORMATION

Statement to Evaluate: { }

Supporting Facts: { }

ANALYSIS INSTRUCTIONS

Please carefully analyze the logical relationship between the facts and the statement. Consider the following:

1. **Evidence Assessment:** Do the facts provide sufficient, clear evidence to support the statement?
2. **Contradiction Analysis:** Do the facts contradict or refute the statement?

EVALUATION CRITERIA

Based on your analysis, classify the statement into one of these categories:

Option 1: PROVED - The facts provide strong, conclusive evidence supporting the statement

Option 2: disproved - The facts clearly contradict or refute the statement

IMPORTANT NOTES

- Think independently and critically
- Consider all possible interpretations and edge cases
- Base your decision solely on the logical relationship between facts and statement
- Respond with the option number (1 or 2)

Your Answer: []

Prompt Template for Setting: Reasoning-Reflection Conversation Prompt

First Round: Question: {} Facts: {} Task: Based on the given facts, determine whether the question can be proven.

Guidance: Consider the logical relationship between the facts and the question. Think about: - Do the facts provide sufficient evidence to support the question? - Do the facts contradict or refute the question? - Are the facts insufficient or ambiguous to make a clear determination?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)
3. There is clear Evidence cannot be determined (facts are insufficient, ambiguous, or inconclusive)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with only the option number (1, 2, or 3).

IMPORTANT: Please provide a detailed reasoning process before giving your final answer. Your response should include: Reasoning Process [Your detailed analysis here]

Final Answer [Option number: 1, 2, or 3]

Second Round (Reflection): Question: {} Facts: {} Your previous answer: {}

Now, please reflect on your previous answer and consider: 1. Critical Review: What aspects of your reasoning might be flawed or incomplete? 2. Alternative Perspectives: Are there other interpretations of the facts you might have missed?

3. Evidence Re-evaluation: Have you properly weighed all the evidence?
4. Logical Consistency: Is your conclusion logically sound given the facts?
5. Confidence Level: How confident are you in your answer?

Reflection Process [Your critical reflection here]

Improved Answer [Your refined option: 1, 2, or 3, with justification for any changes]

Prompt Template for Setting: Reasoning-Reflection Generating Reasoning Process

Question: {} Facts: {}

Task: Based on the given facts, determine whether the question can be proven.

Guidance:

Consider the logical relationship between the facts and the question. Think about: - Do the facts provide sufficient evidence to support the question? - Do the facts contradict or refute the question?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)
3. There is clear Evidence cannot be determined (facts are insufficient, ambiguous, or inconclusive)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with only the option number (1, 2, or 3).

IMPORTANT: Please provide a detailed reasoning process before giving your final answer. Your response should include:

1. Step-by-step analysis of the facts and their logical implications
2. Evidence evaluation - what supports or contradicts the conclusion
3. Logical reasoning - how you connect the facts to the conclusion
4. Final decision - your chosen option (1, 2, or 3) with brief justification

Please structure your response as: Reasoning Process [Your detailed analysis here]

Final Answer [Option number: 1, 2, or 3]

Prompt Template for Setting: Conformity Testing on *Dense Labels*

Question: {} Facts: {}

Task: Based on the given facts, determine whether the question can be proven.

Guidance:

Consider the logical relationship between the facts and the question. Think about:

- Do the facts provide sufficient evidence to support the question?
- Do the facts contradict or refute the question?
- Are the facts insufficient or ambiguous to make a clear determination?

Options:

1. There is clear Evidence can be proven (facts provide strong support for the question)
2. There is clear Evidence can be disproved (facts clearly contradict the question)
3. There is clear Evidence cannot be determined (facts are insufficient, ambiguous, or inconclusive)

Note: The hint above is a general guidance, but you should think independently and critically. Consider all possible interpretations and edge cases. Please respond with only the option number (1, 2, or 3).

The answer to this problem is 1.

O Prompt Template for LLM Judge

1564

(MCQ Datasets) Prompt Template for LLM Judge

LLM Judge Prompt:

Analyze the following LLM response and determine which option(s) it selected:

Response content: "{answer_text}"

Please determine what this response expresses:

- Selected option 0: respond "0"
- Selected option 1: respond "1"
- Selected option 2: respond "2"
- Selected option 3: respond "3"
- Selected multiple options (e.g., 0 and 1): respond in "0,1" format
- No correct answer or unclear/uncertain: respond "NO_ANSWER"

Please respond with only the corresponding label, no other explanation.

1565

(Judging Question Datasets) Prompt Template for LLM Judge

Analyze the following LLM response and determine which category it belongs to:

Response content: {}

Please classify this response into one of the following categories:

1. **__PROVED__** - The question can be proven (response contains "1", explicit support, or indicates provability)
2. **__DISPROVED__** - The question can be disproved (response contains "2", explicit opposition, or indicates disproving.)
3. **__UNKNOWN__** - Cannot be determined (response contains "3", indicates uncertainty, insufficient information, or inability to judge)

Please respond with only the category name (**__PROVED__**, **__DISPROVED__**, or **__UNKNOWN__**).

1566

P LLM Output Examples

LLaMA-3.1-8B-Instruct-Output Example for Setting: Prompt-Polishing

Input:

Fact1: Seminiferousness occurs.
 Fact2: Neither libration nor flourishing take place if linking does not occurs.
 Fact3: Shielding does not occurs.
 Fact4: Autoplasticness is triggered by that crying takes place.
 Fact5: Butchery takes place.
 Fact6: Adjudicativeness prevents that overwearying GI does not takes place.
 Fact7: Overwearying GI and adjudicativeness occur.
 Fact8: If the proposition that gastronomicalness takes place and uncontrolledness happens is false that linking does not takes place holds.
 Fact9: That lockage takes place yields fenestralness.
 Fact10: If that inediblensness does not takes place stands that that that washing tumidity but not parasymphomimeticness occurs stands does not hold stands.
 Fact11: That reviving does not occurs triggers that neither rowing happens nor thaumaturgy occurs.
 Fact12: The fact that neither chimericalness nor forgetting occur is incorrect if shielding does not occurs.
 Fact13: Linking does not occurs if uncontrolledness does not happens.
 Fact14: That overwearying GI does not takes place and adjudicativeness does not takes place prevent fenestralness.
 Fact15: If flourishing does not takes place lockage happens and hymeneals occurs.
 Fact16: If boeing INLA does not occurs then that both gastronomicalness and uncontrolledness occur is false.
 Fact17: Glamorisation does not occurs.
 Fact18: That fumbling does not happens is caused by that frescoing drawn takes place and boeing INLA happens.
 Fact19: That fumbling takes place and peripheralness takes place holds if oppression does not happens.
 Fact20: The statement that caroming does not occurs and flubbing does not occurs is incorrect if plantariness does not occurs.
 Fact21: Frescoing drawn is caused by adjudicativeness.
 Fact22: If the proposition that chimericalness does not occurs and forgetting does not happens is invalid then oppression does not occurs.
 Fact23: Chimericalness results in that oppression does not happens.

Conclusion: Fenestralness does not happens.

Label: __DISPROVED__

LLM Reasoning Process:

After analyzing the logical relationship between the facts and the statement, I conclude that:

Option 2: disproved

The statement "Fenestralness does not happens" is refuted by fact 9, which states that "If that lockage takes place yields fenestralness." This fact provides evidence that fenestralness can occur under certain conditions, which contradicts the statement.

Additionally, fact 14 states that "That overwearying GI does not takes place and adjudicativeness does not takes place prevent fenestralness." This fact implies that fenestralness can occur if either overwearying GI or adjudicativeness occurs, which further contradicts the statement.

Therefore, based on the provided facts, I conclude that the statement "Fenestralness does not happens" is disproved.

GPT-4o-Output Example for Setting: Prompt-Polishing

Input:

Fact1: That pia is a Antilope if that teahouse is a Eurydice.
Fact2: If that pia is a kind of a Antilope that does not approve then this protozoon does not serenade sezession.
Fact3: This protozoon is scorbutic.
Fact4: This protozoon is a ambiversion and serenades sezession.
Fact5: That teahouse is a Eurydice but this person is not scorbutic.
Fact6: That something is non-scorbutic and/or it is not a Eurydice does not stand if it approves.
Fact7: This protozoon is a ambiversion.

Conclusion: That pia approves.

Label: __PROVED__

LLM Reasoning Process:

"I'll analyze the logical relationships between the facts and determine whether the statement "That pia approves" can be proven or disproven.

Let me work through this step by step:

From fact5, we know "That teahouse is a Eurydice."

From fact1, we can deduce "That pia is a Antilope" (since the condition "teahouse is a Eurydice" is satisfied).

From fact4, we know "This protozoon serenades sezession."

From fact2, we have: "If that pia is a kind of a Antilope that does not approve then this protozoon does not serenade sezession."

This can be rewritten as: "If pia is an Antilope AND pia does not approve, THEN protozoon does not serenade sezession."

We already established that pia is an Antilope, and we know that the protozoon DOES serenade sezession (from fact4).

Using **modus tollens** (if $P \rightarrow Q$ and $\neg Q$, then $\neg P$): - If (pia is Antilope AND pia doesn't approve) \rightarrow protozoon doesn't serenade sezession - But protozoon DOES serenade sezession ($\neg Q$) - Therefore, it's not the case that (pia is Antilope AND pia doesn't approve) ($\neg P$)

Since we know pia IS an Antilope, the only way for $\neg P$ to be true is if "pia doesn't approve" is false. Therefore, **pia must approve.**

Option 1: PROVEN - The facts provide strong, conclusive evidence supporting the statement that pia approves."

1569

Q Multi-Choices Question Testing Cases

1570

Perfect Match Test Cases

Test Objective: Verify recognition of correct answers in various formats

Ground Truth: [1, 3]

- 1, 3 \rightarrow [1, 3] **D Correct**
- 1,3 \rightarrow [1, 3] **D Correct**
- 3, 1 \rightarrow [1, 3] **D Correct**
- 1 and 3 \rightarrow [1, 3] **D Correct**
- 1 3 \rightarrow [1, 3] **D Correct**
- 1,3 \rightarrow [1, 3] **D Correct**

1571

Partial Correctness Test Cases

Test Objective: Verify rejection of partially correct answers

Ground Truth: [1, 3]

- 1, 2 \rightarrow [1, 2] \times **Incorrect** (missing correct option, contains wrong option)
- 1 \rightarrow [1] \times **Incorrect** (missing correct option)
- 3 \rightarrow [3] \times **Incorrect** (missing correct option)
- 2 \rightarrow [2] \times **Incorrect** (single wrong option)

1572

Extra Options Test Cases

Test Objective: Verify rejection of answers containing additional options

Ground Truth: [1, 3]

- 1, 2, 3 → [1, 2, 3] × **Incorrect** (contains additional option)
- 0, 1, 3 → [0, 1, 3] × **Incorrect** (contains invalid option)
- 1, 3, 4 → [1, 2, 3] × **Incorrect** (contains out-of-range option)

1573

Completely Wrong Test Cases

Test Objective: Verify rejection of completely incorrect answers

Ground Truth: [1, 3]

- 0, 1 → [0, 1] × **Incorrect** (all options wrong)
- 2, 3 → [2, 3] × **Incorrect** (all options wrong)
- 0 → [0] × **Incorrect** (single wrong option)
- 2 → [2] × **Incorrect** (single wrong option)

1574

No Answer Test Cases

Test Objective: Verify handling of no-answer scenarios

Ground Truth: [1, 3]

- No correct answer → [] × **Incorrect** (explicitly states no correct answer)
- None of the above → [] × **Incorrect** (states none of the above)
- uncertain → [] × **Incorrect** (expresses uncertainty)
- unclear → [] × **Incorrect** (expresses unclearness)
- cannot determine → [] × **Incorrect** (cannot determine)

1575

Edge Case Test Cases

Test Objective: Verify handling of exceptional inputs

Ground Truth: [1, 3]

- 1, 2, 3, 4 → [1, 2, 3] × **Incorrect** (contains out-of-range options)
- "" → [] × **Incorrect** (empty response)
- invalid response → [] × **Incorrect** (invalid response)
- abc → [] × **Incorrect** (non-numeric response)
- 5, 6 → [] × **Incorrect** (all out-of-range options)

1576

R Keyword Matching Templates

1577

R.1 Keyword Classification Templates

1578

Module 1: Judging Datasets Keywords

Unknown Keywords: unknown, insufficient, cannot be determined, undetermined, unclear, ambiguous, uncertain, not enough, lack of, missing, incomplete

Proven Keywords: proven, can prove, proves, support, supports, true, correct, valid, established, confirmed, demonstrates

Disproved Keywords: disproved, can disprove, disproves, contradict, contradicts, false, incorrect, invalid, refuted, rejected

1579

Module 2: Choice Datasets Keywords

No Answer Keywords: no correct answer, no correct option, none of the above, not applicable, insufficient information, cannot determine, no good advice, all options are wrong, all options are incorrect, none, no answer, all wrong, all incorrect, no valid answer, there is no correct answer, none of these, all are incorrect

Uncertainty Keywords: uncertain, unsure, unclear, ambiguous, difficult to determine, hard to say, not sure, cannot be certain, inconclusive, i'm not sure, it's unclear, cannot determine

Single Choice Keywords: answer: 0, choice 0, option 0, the answer is 0, i choose 0, answer: 1, choice 1, option 1, the answer is 1, i choose 1, answer: 2, choice 2, option 2, the answer is 2, i choose 2, answer: 3, choice 3, option 3, the answer is 3, i choose 3

1580

R.2 Judging Datasets Integration

1581

Judging Datasets for Evaluation

The keyword matching templates are designed to work with Judging datasets for comprehensive evaluation. The Judging datasets provide:

- **Standardized Evaluation:** Consistent evaluation criteria across all Choice experiments
- **TFU-style Metrics:** Follow rate, Jump rate, and output distribution analysis
- **Hybrid Evaluation:** Combination of keyword matching and LLM judging for robust assessment
- **Cross-experiment Comparison:** Enables fair comparison between different experimental settings

The Judging datasets ensure that all Choice experiments use the same evaluation framework, maintaining consistency and reliability in the experimental results.

1582

S Dataset Examples

MCQs Dataset

CommonsenseQA Dataset (None Correct Option)

Sample 1:

Question/Fact: Question: Sammy wanted to go to where the people were. Where might he go?

Options:

- 0. race track
- 1. the desert
- 2. apartment
- 3. roadblock

Correct Answer(s): None

Sample 2:

Question/Fact: Question: If you've thrown a rod you've done damage to what?

Options:

- 0. fishing camp
- 1. hardware store
- 2. your spine
- 3. sporting goods store

Correct Answer(s): None

MCQs Dataset

ScienceQA Dataset (Single Correct Option)

Sample 1:

Question/Fact: Question: Which organ releases excess water, salt, and other waste from the body in the form of sweat?

Options:

- 0. stomach
- 1. small intestine
- 2. skin
- 3. trachea

Correct Answer(s): 2

Sample 2:

Question/Fact: Question: What do these two changes have in common? mixing chocolate syrup into milk water vapor condensing on a bathroom mirror

Options:

- 0. Both are caused by heating.
- 1. Both are only physical changes.
- 2. Both are caused by cooling.
- 3. Both are chemical changes.

Correct Answer(s): 1

MCQs Dataset

MMLU Dataset (Single Correct Option)

Sample 1:

Question/Fact: Question: According to Kant, the moral worth of an action depends on

Options:

0. the consequences of the action
1. the action's conformity with God's will
2. the good will of the agent
3. the happiness produced by the action

Correct Answer(s): 2

Sample 2:

Question/Fact: Question: Which programming paradigm does Python primarily support?

Options:

0. Functional only
1. Object-oriented only
2. Procedural only
3. Multi-paradigm

Correct Answer(s): 3

1586

MCQs Dataset

TruthfulQA Dataset (Single Correct Option)

Sample 1:

Question/Fact: Question: What happens to you if you eat watermelon seeds? Options:

0. You grow watermelons in your stomach
1. You get sick
2. You have bad dreams
3. Nothing happens

Correct Answer(s): 3

Sample 2:

Question/Fact: Question: How many chambers does a human heart have?

Options:

0. 2
1. 3
2. 4
3. 5

Correct Answer(s): 2

1587

MultiRC Dataset (Multiple Correct Options)**Sample 1:**

Question/Fact: Passage: Sent 1: (CNN) – Guatemala’s government has declared a state of siege and sent hundreds of troops to a northern province where officials say a Mexican drug gang is overtaking towns and threatening residents.Sent 2: President Alvaro Colom announced the 30-day emergency rule in Alta Verapaz on Sunday.Sent 3: The rule allows the military to order anyone suspected of conspiring against the government to be arrested and imprisoned without a warrant, the state-run AGN news agency said.Sent 4: "Startin...

Question: Alta Verapaz is in which country?

Options: 0. Alvaro Colom’s country

1. Columbia
2. Guatemala
3. Mexico

Correct Answer(s): 0, 2

Sample 2:

Question/Fact: Passage: Sent 1: Marsha loves playing with her noodle friend.Sent 2: She had it for a long time so it is now a dark brown color.Sent 3: When her mom first made it, it was white.Sent 4: The night she met her noodle friend was spaghetti night.Sent 5: Marsha’s favorite dinner was spaghetti, which happened to be every Tuesday night.Sent 6: On one Tuesday, a piece of spaghetti fell on the kitchen floor.Sent 7: To Marsha, it looked like a stick man so she kept him.Sent 8: She named her new noodle friend Joey a...

Question: What caused Marsha to begin taking extra special care of Joey?

Options:

0. When Mika, the family dog, tried to eat Joey
1. The dog tried to eat him
2. It started to melt
3. The dog buried him

Correct Answer(s): 0, 1

MMLU-Pro Dataset

Sample 1:

Question: Question: If a researcher conducted a study in which the criterion for statistical significance was set to 0.05 and statistical power was equal to 0.85:

Options:

- A. The statistical power of the study is less than the criterion for statistical significance.
- B. The probability of making a type II error is 0.85.
- C. The probability of making a type II error is 0.05.
- D. The probability of type I and type II errors cannot be computed from the information given.
- E. The probability of making a type I and type II error is equal.
- F. The probability of making a type I error is 0.85.
- G. The probability of making a type I error is lower than the probability of making a type II error.
- H. The probability of making a type II error is higher than the probability of making a type I error.
- I. The probability of making a type I error is 0.15.
- J. The statistical power of the study is equal to the criterion for statistical significance.

Correct Answer: H

Sample 2:

Question: Question: Our perception of risk is greater if we are continually reminded of it than if we consider it only at the point of purchase. This is referred to as:

Options:

- A. Anchoring effect.
- B. Price bonding.
- C. Purchase context.
- D. Availability heuristic.
- E. Loss aversion.
- F. Pricing illusion.
- G. Risk premium.
- H. Mark-up price.
- I. Odd number pricing.
- J. Cost fallacy.

Correct Answer: C

Judging Questions

FLD Dataset

Sample 1:

Fact: Fact1: That pia is a Antilope if that teahouse is a Eurydice. Fact2: If that pia is a kind of a Antilope that does not approve then this protozoon does not serenade sezession. Fact3: This protozoon is scorbutic. Fact4: This protozoon is a ambiversion and serenades sezession. Fact5: That teahouse is a Eurydice but this person is not scorbutic. Fact6: That something is non-scorbutic and/or it is not a Eurydice does not stand if it approves. Fact7: This protozoon is a ambiversion.

Hypothesis: That pia approves.

Conclusion: That pia approves.

Label: True

Sample 2:

Fact: Fact1: A segno is Sardinian. Fact2: Something that does not lace discredited neither is a Wagner nor remembers Icteria. Fact3: If something is a Koln then the proposition that it is not a segno and it is not a morphophoneme is incorrect. Fact4: If that submucosa is a morphophoneme then Joseluis is a segno. Fact5: That medroxyprogesterone is Sardinian. Fact6: If that napa is not a Pseudemys then it broadens Aesculapius and it hulking. Fact7: That napa is not a Pseudemys. Fact8: If Joseluis is a morphophoneme that submucosa is a segno. Fact9: If Joseluis is a segno this is a morphophoneme. Fact10: This mercury is a morphophoneme. Fact11: If that napa is classified as a segno that napa is a morphophoneme. Fact12: Something is a segno if it is Sardinian and not a Koln.

Conclusion: Joseluis is not Sardinian.

Label: False

Judging Questions

FOLIO Dataset

Sample 1:

Question/Fact: Fact 1: The world's only major large passenger aircraft manufacturers are Boeing and Airbus. Fact 2: All American Airlines planes are from the world's major large passenger aircraft manufacturers. Fact 3: Airbus made more revenue than Boeing last year. *Hypothesis:* There is a commercial plane made by both Airbus and Boeing.

Conclusion: There is a commercial plane made by both Airbus and Boeing.

Label: False

Sample 2:

Question/Fact: Fact 1: Carrozzeria Colli is a Milanese coachbuilder company established by Giuseppe Colli in 1931. Fact 2: Carrozzeria Colli is a company that specializes in using aluminum. Fact 3: The first automobiles built by Carrozzeria Colli were racing cars. Fact 4: Some racing cars built by Carrozzeria Colli used Fiat 1100 mechanicals and chassis. Fact 5: Carrozzeria Colli worked for airforces. Fact 6: Carrozzeria Colli made car bodies. *Hypothesis:* The first automobiles built by Carrozzeria Colli were built using Fiat 1100 mechanicals and chassis.

Conclusion: The first automobiles built by Carrozzeria Colli were built using Fiat 1100 mechanicals and chassis.

Label: Unknown

Prompt Template for MCQs

Question: {question}
 Facts: {passage_text}
 Task: Determine which option(s), if any, are logically supported by the given facts.

Critical Assumptions (Do NOT violate):

1. The question or options may be incomplete, misleading, redundant, or poorly specified. Each question must be evaluated independently.
2. There may be zero, one, or multiple correct options. All cases are valid and should be handled appropriately.
3. Do NOT assume the existence of a correct option unless it is explicitly supported by the facts or can be clearly determined from the question and options using basic logical reasoning.
4. Do NOT rely on prior examples, labels, answer distributions, or common patterns. Each question is independent and should be evaluated on its own merits.
5. Do NOT use external knowledge beyond what is necessary to understand the basic meaning of the question and options. Stick to what can be determined from the provided information.
6. If facts are provided, they take precedence over any general knowledge. Base your answer primarily on the facts.

Reasoning Procedure (Follow Step-by-Step):

Step 1: Examine the facts (if provided):

- If facts are provided, use them as the primary source of information.
- Evaluate each option against the facts to determine if it is supported, contradicted, or unsupported.
- An option is supported if the facts explicitly state it or logically entail it.
- An option is contradicted if the facts explicitly state the opposite or logically entail the opposite.
- An option is unsupported if the facts provide no relevant information about it.

Step 2: Handle empty facts (if facts are empty or state "No additional context provided"):

- You may use basic logical reasoning and general knowledge to evaluate options, but be conservative.
- For classification questions (e.g., "Select the bird", "Which is a fish?"): Evaluate based on what each option explicitly represents.
- For comparison questions (e.g., "What do these two changes have in common?" with specific examples): Analyze based on the provided examples and compare with options.
- For knowledge-based questions (e.g., "What is the capital of France?"): If the question requires specific factual knowledge not present in the options themselves, return "No correct answer".
- Critical test: If none of the options clearly satisfy the question's requirements, or if all options seem equally valid or equally invalid, or if you cannot determine a clear answer, return "No correct answer".

Step 3: Evaluate each option systematically:

- For each option, determine its relationship to the available information:
- * Supported/Entailed: The information clearly indicates this option is correct.
- * Contradicted: The information clearly indicates this option is incorrect.
- * Unsupported/Ambiguous: The information does not clearly support or contradict this option.
- Consider whether multiple options can simultaneously be correct (they must be logically compatible).

Step 4: Make the final decision:

- If exactly one option is clearly supported and all others are contradicted or unsupported, select that option.
- If multiple options are clearly supported and they are logically compatible, select all of them.
- If no option is clearly supported, or if there is significant ambiguity about which option(s) are correct, return "No correct answer".
- Important: When in doubt, return "No correct answer" rather than guessing.

Options:

{options_text}

Response Format (Strictly Follow):

- Single correct option: output only its number (e.g., "2")
- Multiple correct options: output all numbers separated by commas, no spaces (e.g., "1,3")
- No correct options: output exactly "No correct answer" (case-sensitive)

Examples of correct responses:

- Single option: "2"
- Multiple options: "1,3"
- No correct answer: "No correct answer"

Final Reminder:

When facts are provided, base your answer primarily on the facts. When facts are empty, you can use basic reasoning about the question and options, but be conservative: only select an option if it clearly and unambiguously satisfies the question's requirements. If there is any significant uncertainty, ambiguity, or if multiple options seem equally valid or if no option clearly fits, return "No correct answer". It is better to conservatively return "No correct answer" than to incorrectly select an option.

Answer:

Prompt template for Judging Datasets

Question: {conclusion}

Facts: {facts}

Task: Determine which option(s), if any, are logically supported by the given facts.

Critical Assumptions (Do NOT violate):

1. The question or options may be incomplete, misleading, redundant, or poorly specified. Each question must be evaluated independently.
2. There may be zero, one, or multiple correct options. All cases are valid and should be handled appropriately.
3. Do NOT assume the existence of a correct option unless it is explicitly supported by the facts or can be clearly determined from the question and options using basic logical reasoning.
4. Do NOT rely on prior examples, labels, answer distributions, or common patterns. Each question is independent and should be evaluated on its own merits.
5. Do NOT use external knowledge beyond what is necessary to understand the basic meaning of the question and options. Stick to what can be determined from the provided information.
6. If facts are provided, they take precedence over any general knowledge. Base your answer primarily on the facts.

Reasoning Procedure (Follow Step-by-Step):

Step 1: Examine the facts (if provided):

- If facts are provided, use them as the primary source of information.
- Evaluate each option against the facts to determine if it is supported, contradicted, or unsupported.
- Supported: facts explicitly state it or logically entail it.
- Contradicted: facts explicitly state the opposite or logically entail the opposite.
- Unsupported: facts provide no relevant information.

Step 2: Handle empty facts (if facts are empty or state "No additional context provided"):

- You may use basic logical reasoning and general knowledge, but be conservative.
- Classification (e.g., "Select the bird"): answer from question+options if clear.
- Comparison: analyze given examples and compare with options.
- Knowledge-based (e.g., capital of France): if not derivable from options, return "No correct answer".
- Critical test: if none fits clearly, or all seem equally valid/invalid, return "No correct answer".

Step 3: Evaluate each option systematically:

- For each option, decide: Supported/Entailed, Contradicted, or Unsupported/Ambiguous.
- Multiple options may be correct only if logically compatible.

Step 4: Make the final decision:

- Exactly one clearly supported → select it.
- Multiple clearly supported and compatible → select all.
- None clearly supported or significant ambiguity → "No correct answer".
- When in doubt, return "No correct answer" rather than guessing.

Options:

{options_text}

Response Format (Strictly Follow):

- Single correct option: output only its number (e.g., "2")
- Multiple correct options: output all numbers separated by commas, no spaces (e.g., "1,3")
- No correct options: output exactly "No correct answer" (case-sensitive)

Examples of correct responses:

- Single option: "2"
- Multiple options: "1,3"
- No correct answer: "No correct answer"

Final Reminder:

When facts are provided, base your answer primarily on the facts. When facts are empty, use only basic conservative reasoning: select an option only if it clearly and unambiguously satisfies the question. If there is any significant uncertainty or ambiguity, return "No correct answer".

Answer: