

Measuring the Inconsistency of Large Language Models in Ordinal Preference Formation

Xiutian Zhao

University of Edinburgh
x.zhao-103@sms.ed.ac.uk

Ke Wang , Wei Peng

Huawei IT Innovation and Research Center
{wangke215, peng.wei1}@huawei.com

Abstract

Despite large language models' (LLMs) recent advancements, their bias and hallucination issues persist, and their ability to offer consistent and preferential rankings remains under-explored. This study investigates the capacity of LLMs to provide consistent ordinal preferences, a crucial aspect in scenarios with dense decision space or lacking absolute answers. We introduce a formalization of consistency based on order theory, outlining criteria such as transitivity, asymmetry, reversibility, and independence from irrelevant alternatives. Our diagnostic experiments on selected state-of-the-art LLMs reveal their inability to meet these criteria, indicating a strong positional bias and poor transitivity, with preferences easily swayed by irrelevant alternatives. These findings highlight a significant inconsistency in LLM-generated preferential rankings, underscoring the need for further research to address these limitations.

1 Introduction

Expressing one's preferences in an ordinal manner is a widespread and informative practice in human reasoning and communication (Arrow et al., 2010). By evaluating and comparing available options, individuals can make more informed decisions and communicate their values to others more effectively. In the domain of natural language processing (NLP), human preferential feedback serves as a valuable data type for aligning language models with human inclinations (Schulman et al., 2017; Rafailov et al., 2023).

Recent advances in large language models (LLMs) have prompted researchers to investigate the potential of LLMs in complex ranking-based tasks - such as recommendation (Li et al., 2023; Ren et al., 2024), web search (Sun et al., 2023), and text relevance comparison (Qin et al., 2023) - traditionally handled by task-specific models. Moreover, given that human annotation and evaluation

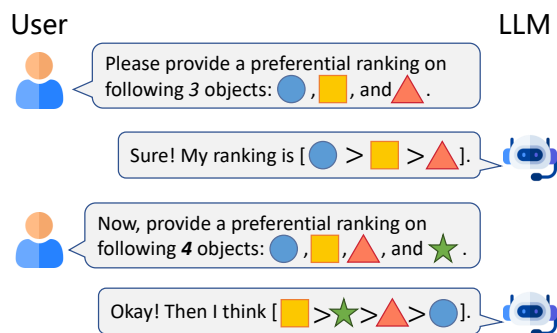


Figure 1: An example of violating *Independence from Irrelevant Alternatives (IIA)* criterion. Initially, given 3 choices, the model preferred **Circle** over **Square** over **Triangle**. However, after introducing a new choice **Star**, the **relative preferential positions** among the initial 3 choices inconsistently changed.

are resource-intensive, there is an increasing interest in augmenting or even substituting human preferential data with LLM-generated judgments to annotate, evaluate, or supplement as corpus (Wang et al., 2021; Zhao et al., 2022; Lee et al., 2023).

On the other hand, it is well-recognized that LLMs often exhibit severe bias and hallucination (Rawte et al., 2023; Zhang et al., 2023b). Specifically, prior research has identified undesirable behavioral patterns in LLMs when presented with multiple options (i.e., choices). For example, in Multiple-Choice Question Answering (MCQA), a task commonly used to benchmark LLM performance (Hendrycks et al., 2021a,b; Robinson et al., 2023), LLMs have shown a particular bias towards the position (Pezeshkpour and Hruschka, 2023) and the labeling of choices (Zheng et al., 2023).

Unlike MCQA, which requires single-selections, preferential ranking necessitates the ordinal preferences of all options, which is invaluable in scenarios lacking a definitive answer. Despite extensive research, the current literature on LLM bias has not fully addressed their behavior in preferential ranking tasks. To address this gap, our study en-

deavors to investigate a critical yet under-explored question: *To what extent can LLMs consistently and coherently provide ordinal ‘preferences’?*

This study makes an effort to measure the consistency (or more likely inconsistency) of LLMs in preferential ranking. Firstly, by incorporating order theory (Grätzer, 2002), we formalize the question and define five self-evident criteria that must be satisfied to achieve ‘consistency’ (§ 2.1). Through comprehensive diagnostic experiments on various state-of-the-art LLMs, we examine their adherence to preferential ranking criteria, namely transitivity, asymmetry, reversibility, and independence from irrelevant alternatives (IIA, as exemplified in Figure 1). We demonstrate that even the most advanced LLMs are incapable of providing consistent or coherent preferential rankings.

Specifically, we observe that: (1) The tested models generally **fail to meet the asymmetry condition** in preferential ranking (e.g., different answers for ‘compare A and B’ and ‘compare B and A’), indicating a strong positional bias (§ 3.2). (2) The preferences provided by the tested models exhibit **poor transitivity**; that is, concatenating binary preferences of choice pairs does not reliably yield an ordinal chain, and in fact, these preferences are often contested or even cyclic (§ 3.2). (3) The preferences of LLMs are **significantly influenced by the addition or removal of irrelevant alternatives** (§ 3.3). (4) When requested to provide rankings in different ordinal sequences (e.g., preferential descending and ascending), LLMs **fail to produce logically equivalent outcomes** (§ 3.4).

In summary, our contributions are threefold:

- We first formalize the measurement of consistency in LLM preferential ranking through the lens of order theory.
- We devise specific measurement metrics that align with the defined consistency conditions. A preliminary experiment not only corroborates some shared biases with the MCQA task but also highlights the unique challenges of preferential ranking.
- Through comprehensive experiments on a collection of state-of-the-art (SOTA) LLMs, we uncover a severe and widespread inconsistency in LLM preferential ranking. Our findings sound a serious alarm in related research and call for immediate mitigation efforts.

2 Experiment Setup and Preliminaries

2.1 Definition

Concretely, let $A = \{a, b, \dots, n\}$ be a finite set of n distinct alternatives, we define a preferential ranking as a *strict partial ordering* \succ of A (Grätzer, 2002; Rosen, 2007). Such ordering satisfies that, for all $a, b, c \in A$:

- **Irreflexivity**: not $a > a$.
- **Asymmetry**: if $a > b$ then not $b > a$.
- **Transitivity**: if $a > b$ and $b > c$ then $a > c$.

Besides above intrinsic conditions, In multi-round preferential ranking scenarios, we also examine following criteria

- **Independent from Irrelevant Alternative (IIA)**: if $a > [\dots] > b$ in an ordering $\succ_{original}$, given additional alternatives c, d, \dots , then $a > [\dots] > b$ in the new ordering \succ_{new} .
- **Reversibility**: if $a > [\dots] > b$ then $b < [\dots] < a$. This criterion can be regarded as a full ranking generalization of binary *Asymmetry*.

Note that there is also a *non-strict partial ordering* variation that allows $a = b$ (i.e., preferential ties). For simplicity, all experiments are conducted under *strict partial ordering* scenario.

2.2 Datasets

Following prior research that benchmarks the reasoning capabilities of LLMs (Park et al., 2022; Liu et al., 2023; Zhang et al., 2023a; Google, 2023; Jiang et al., 2023), we choose the MMLU (Hendrycks et al., 2021a) as our principal testbed. This benchmark encompasses a total of 14,079 MCQA test cases across 57 varied subject areas. Given that our study’s main focus is on preference ranking rather than choice generation, the MMLU is particularly well-suited to our research interests, as the benchmark is uniformly formatted with multiple-choice options, and the options (i.e., choices) are predefined.

It should be noted that preferential ranking is a more challenging task than MCQA because it necessitates additional ordinal information. To create a balanced test set, we curate a collection by selecting the first 20 cases from each subject, resulting in a total of 1,140 cases. In line with the original MMLU framework (Hendrycks et al., 2021a), we employ a 5-shot example prompting strategy that leverages the dataset’s fixed development set.

2.3 Evaluated Models

To investigate the potential inconsistencies in LLM preferential ranking, we have compiled a selection of open-source models, including Llama-3-70B (AI@Meta, 2024) and Qwen-1.5-72/110B (Qwen, 2024). Our selection criteria prioritized models with relatively large parameters (exceeding 70B), as smaller models are generally outperformed by their larger counterparts in text-based task performance. For proprietary models, we have included gpt-3.5 (Brown et al., 2020) and gpt-4o (OpenAI, 2023), which are among the most widely utilized closed-source models in recent times. Specifically, we adopt the snapshot models gpt-3.5-turbo-0125 and gpt-4o-2024-05-13. Sources of tested open-source models are summarized in Table 6. Detailed specifications and sources for the selected models are provided in Appendix A. To ensure reproducibility, we have set the temperature for all experiments to zero (the temperature setting ranges from 0 to 2 for OpenAI models, and from 0 to 1 for others).

2.4 Preliminary Examinations

Prior to initiating the principal experiments, it is beneficial to ascertain whether the label tokens of choices affect LLMs’s preferences and whether LLMs exhibit differential performance across single-select and preference ranking tasks. To this end, we conduct two preliminary examinations.

Alternative Label Bias Following (Zheng et al., 2023), in comparison with the original *Alphabetic* label tokens ([A, B, C, D]) of MMLU, we add *Arabic*: [(1), (2), (3), (4)], and *Roman*: [I, II, III, IV] token sets. The parentheses in *Arabic* token to reduce ambiguity for numerical questions. Few-shot examples are modified in accordance with the altered labels.

As shown in Table 1, the first-preference accuracies vary slightly for tested models. We also evaluate similarity of rankings based on minimal editing distance, and the normalized (between 0 and 1) similarities are near 0.9, suggesting a minor influence of label tokens in preferential ranking.

Question Format Sensitivity Given that the *first-preference* in a ranking context is logically congruent with a single-select choice, we juxtapose the accuracies of MCQA across these varying question formats.

Alternative Labels	Alphabetic	Arabic		Roman	
	Acc.@1	Acc.@1	Sim.	Acc.@1	Sim.
llama-3-70b	72.1	72.9	86.6	72.2	87.3
qwen-1.5-72b	72.2	73.2	89.2	70.5	89.5
qwen-1.5-110b	71.7	71.9	90.7	71.1	89.8
gpt-3.5-turbo	62.1	62.7	89.1	61.1	88.0
gpt-4o	83.7	83.1	92.5	84.6	92.5

Table 1: The accuracies and similarity scores of first-preferences among different label token sets. **Sim.** denotes the similarity score.

Question Format	Single Select	Ordinal Ranking			Cardinal Ranking		
		@1	@2	@3	@1	@2	@3
HitRate@N	-						
llama-3-70b	77.3	72.1(-5.2)	85.6	93.1	73.9(-3.4)	87.4	94.2
qwen1.5-72b	78.0	72.2(-5.8)	85.5	91.6	70.4(-7.6)	83.7	91.4
qwen1.5-110b	76.5	71.7(-4.8)	85.9	93.2	71.8(-4.7)	86.0	93.2
gpt-3.5-trubo	67.3	62.1(-5.2)	81.0	91.6	61.1(-6.2)	80.7	91.4
gpt-4o	78.1	83.7(+5.6)	92.2	96.9	82.7(+4.6)	92.4	96.8

Table 2: Full results of Question Format Sensitivity test. @N denotes that the accuracies are calculated based on the first N elements of the preferential ranking lists.

Ordinal Ranking necessitates that outputs be sequenced lists, whereas *Cardinal Ranking* obliges models to assign a numerical score to each potential answer. As evidenced in Table 2, all models, with the exception of gpt-4o, demonstrate reduced accuracies relative to the single-select format. This leads us to infer that LLMs are indeed **sensitive to the format of questions**, thereby underscoring the importance of probing into LLMs’ performance in preference-based ranking.

3 Main Experiments and Key Observations

3.1 Irreflexivity

We elect to forgo further scrutiny of this criterion, as a pilot test that we performed indicates that carefully crafted prompt instructions effectively preclude the recurrence of options within model-generated rankings, with infrequent transgressions observed (less than 1% across all evaluated models).

3.2 Asymmetry and Transitivity

Symmetry checking in LLM reasoning is fundamentally a test for positional bias. Given that prior research has identified positional bias in single-selection tasks (Pezeshkpour and Hruschka, 2023; Zheng et al., 2023), it is imperative to ascertain whether LLMs exhibit a similar propensity to modify their preference rankings when options are se-

quenced differently in questions.

Consider A as a list of n options: $[a_1, a_2, a_3, a_4]$. By soliciting the LLM to perform binary comparisons $n \times (n - 1)$ times, we can construct an $n \times n$ binary comparison matrix M . As depicted in Figure 2, we assign $m_{ij} = 1$ if the model shows a preference for a_i when presented with the ordered pair $[a_i, a_j]$ (noting that $[a_j, a_i]$ constitutes a distinct ordered pair), and $m_{ij} = -1$ if the model opts for a_j when faced with the same ordered pair.

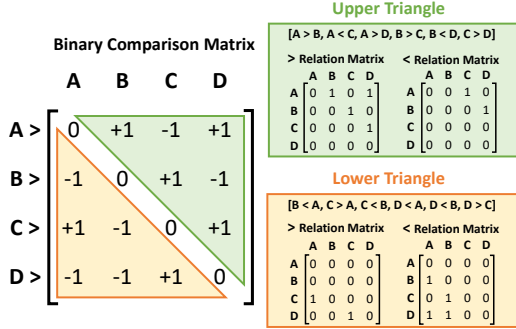


Figure 2: A 4-option binary comparison matrix (left) and a breakdown of its upper and lower triangles (right). Each triangular matrices can be transformed into a relation matrix for each relation.

Next, we calculate an asymmetry score by comparing the agreement between m_{ij} and m_{ji} :

$$\frac{2 \sum_{i,j=0,i>j}^n s_{ij}}{n(n-1)}, s_{ij} = \begin{cases} 0 & \text{if } m_{ij} \equiv m_{ji} \\ 1 & \text{if } m_{ij} \not\equiv m_{ji} \end{cases} \quad (1)$$

The average asymmetry scores, as delineated in Table 3, reveal a low degree of overall asymmetry among all models, indicative of significant positional biases in the preferential ranking task. Notably, gpt-4o, recognized as the SOTA proprietary LLM to date (Li et al., 2024), registers the lowest asymmetry score. Fundamentally, the positions of options in binary comparisons markedly affects the preferences of LLMs, culminating in a decrease in duly asymmetry. This finding also concurs prior observations that LLMs show a position bias in MCQA task (Robinson et al., 2023).

Upon identifying inconsistencies in asymmetry, we recognize that the upper and lower triangles of a binary comparison matrix do not perfectly correspond. Consequently, it is imperative to calculate transitivity separately for each triangular matrix.

Considering options as nodes and relations ($<$ and $>$) as directions within a graph, we reconceptualize the problem as one of directed reachability.

Model	Asymmetry	Transitivity		
		Upper Tri.	Lower Tri.	Avg.
random	49.9	59.4	59.4	59.4
llama-3-70b	76.6	94.5	94.7	94.6
qwen-1.5-72b	73.4	96.5	96.1	96.3
qwen-1.5-110b	82.8	97.3	96.4	96.9
gpt-3.5-turbo	73.0	94.1	94.6	94.4
gpt-4o	67.1	89.2	88.9	89.1

Table 3: Asymmetry and transitivity scores comparisons. **Upper Tri.** and **Lower Tri.** denotes the upper triangle and lower triangle results, respectively.

Removed Choice Index	Gold	Gold + 1	Gold + 2	Gold + 3	Random Non-Gold
llama-3-70b	49.7	65.3	64.9	67.1	66.3
qwen-1.5-72b	55.5	75.3	75.6	74.1	75.2
qwen-1.5-110b	57.9	76.7	76.3	76.3	74.6
gpt-3.5-turbo	62.5	71.5	71.4	70.6	69.7
gpt-4o	65.9	80.4	79.7	80.5	81.3

Table 4: Similarity scores are calculated comparing with full-option rankings. **+N** denotes the N-th option after the indices of gold answers.

For each relation, a relation matrix R can be derived from the triangular matrices, as depicted in Figure 2. Subsequently, we can compute the *transitive closure* matrix (Purdum Jr, 1970; Karp, 1990):

$$M_t = [r_{ij}]_{n \times n} = R^0 \vee R^1 \vee \dots \vee R^{n-1} \quad (2)$$

where \vee is *logical and* operation. If $r_{ij} = 1$ in M_t , then the relation has successfully transitioned; otherwise, it is deemed non-transitive.

As evidenced in Table 3, all models exhibit moderate transitivity, with the random baseline established at 59.4. Furthermore, there are subtle differences between the upper and lower triangles in all models. This observation is consistent with the positional biases identified in the asymmetry experiment. In contrast, the impact of relation symbols on transitivity is considerably less pronounced.

3.3 Independent from Irrelevant Alternative

IIA criterion assesses whether the introduction of an additional option affects the relative order of the original preference ranking. This condition is tested by calculating a normalized similarity score, $Sim = 1 - MED/2n$, where MED represents the minimum edit distance between (1) preference rankings with three options and (2) preference rankings with four options, excluding the omitted choice from the three option rankings.

As suggested by Table 4, the removal of gold answers significantly alters the LLMs’ preferences for

the remaining options. Conversely, the elimination of non-gold options results in less pronounced, yet still noticeable, impacts on the preference rankings.

3.4 Reversibility

In all preceding experiments, the models were instructed to provide preferences in a descending order, placing the most favored option first. Maintaining all other conditions constant, we now instruct the models to rank in an ascending order, positioning the least favored option at the forefront.

Table 5 encapsulates the *first-N option match rates* and the overall ranking similarities between the original rankings and the reversed sequences under the alternative output order. All models exhibit suboptimal performance on full match length (with <45% match rate), while gpt-4o outperforms other models by significant margins.

Match Length	1	2	3 (also 4)	Sim.
llama-3-70b	73.4	47.8	34.1	80.9
qwen-1.5-72b	70.1	42.0	30.3	79.9
qwen-1.5-110b	70.7	45.7	31.5	80.3
gpt-3.5-turbo	61.4	37.4	28.3	78.4
gpt-4o	85.4	61.3	44.6	84.8

Table 5: Since repetitive entries are forbidden (see § 2.1), results for match length of 3 and 4 are the same for 4-option sequences. *Sim.* denotes similarity scores.

4 Conclusion and Future Work

To conclude, we have formalized the consistency measurements in preferential ranking tasks by designing corresponding criteria and metrics. Through diagnostic experiments, we have evaluated some of the most advanced LLMs, uncovering severe inconsistencies and positional biases that are prevalent across all models, among other observations. Our study raises general awareness of discrepancies in LLMs and signals a call for future research efforts. Specifically, we highlight two areas of interest: the development of a non-MCQA benchmark for consistency measurement and the creation of mitigation methods to enhance the consistency of LLMs in ranking-based tasks.

Limitations

While the experiments on MMLU yield notable and insightful observations, we acknowledge that MCQA is not fully aligned with preferential ranking. Most QA benchmarks have predetermined ‘correct’ answers; however, preferential ranking

can also be relevant in scenarios where there is no absolute right or wrong. Therefore, an additional avenue for future work could involve constructing a benchmark that measures preference representativeness rather than one based on true-or-false judgments.

Acknowledgements

Our colleague Jack provided the deployment and inference services of the tested open-source models, thanks to him.

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Kenneth J Arrow, Amartya Sen, and Kotaro Suzumura. 2010. *Handbook of social choice and welfare*, volume 2. Elsevier.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Google. 2023. [Gemini: A family of highly capable multimodal models](#).
- George Grätzer. 2002. *General lattice theory*. Springer Science & Business Media.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#).
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the math dataset](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Richard M Karp. 1990. The transitive closure of a random digraph. *Random Structures & Algorithms*, 1(1):73–93.

- Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. 2023. [Rlaif: Scaling reinforcement learning from human feedback with ai feedback](#).
- Lei Li, Yongfeng Zhang, and Li Chen. 2023. Prompt distillation for efficient llm-based recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1348–1357.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024. [From live data to high-quality benchmarks: The arena-hard pipeline](#).
- Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2023. [Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2022. [Social simulacra: Creating populated prototypes for social computing systems](#). In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology, UIST '22*, New York, NY, USA. Association for Computing Machinery.
- Pouya Pezeshkpour and Estevam Hruschka. 2023. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*.
- Paul Purdom Jr. 1970. A transitive closure algorithm. *BIT Numerical Mathematics*, 10(1):76–94.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, et al. 2023. Large language models are effective text rankers with pairwise ranking prompting. *arXiv preprint arXiv:2306.17563*.
- Qwen. 2024. [Introducing qwen1.5](#).
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#).
- Vipula Rawte, Amit Sheth, and Amitava Das. 2023. A survey of hallucination in large foundation models. *arXiv preprint arXiv:2309.05922*.
- Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. Representation learning with large language models for recommendation. In *Proceedings of the ACM on Web Conference 2024*, pages 3464–3475.
- Joshua Robinson, Christopher Michael Rytting, and David Wingate. 2023. [Leveraging large language models for multiple choice question answering](#).
- Kenneth H Rosen. 2007. *Discrete mathematics and its applications*. The McGraw Hill Companies,.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#).
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? investigating large language models as re-ranking agents](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14918–14937, Singapore. Association for Computational Linguistics.
- Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. Want to reduce labeling cost? gpt-3 can help. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4195–4205.
- Jintian Zhang, Xin Xu, and Shumin Deng. 2023a. Exploring collaboration mechanisms for llm agents: A social psychology view. *arXiv preprint arXiv:2310.02124*.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023b. Siren’s song in the ai ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- Mengjie Zhao, Fei Mi, Yasheng Wang, Minglei Li, Xin Jiang, Qun Liu, and Hinrich Schütze. 2022. Lm-turk: Few-shot learners as crowdsourcing workers in a language-model-as-a-service framework. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 675–692.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations*.

A Reproducibility

The two proprietary models we evaluated, gpt-3.5 and gpt-4o, are commercially available close-source models. The source for open-source models are shown in Table 6.

Models	Sources
gpt-3.5-turbo	https://platform.openai.com/
gpt-4o	https://platform.openai.com/
llama-3-70b	https://huggingface.co/meta-llama/Meta-Llama-3-70B-Instruct
qwen1.5-72b	https://huggingface.co/Qwen/Qwen1.5-72B-Chat
qwen1.5-110b	https://huggingface.co/Qwen/Qwen1.5-110B-Chat

Table 6: Sources of the evaluated models.