

Can Many-Shot In-Context Learning Help LLMs as Evaluators? A Preliminary Empirical Study

Mingyang Song, Mao Zheng, Xuan Luo
Tencent Hunyuan
nickmysong@tencent.com

Abstract

Utilizing Large Language Models (LLMs) as evaluators to assess the performance of other LLMs has garnered attention. However, this evaluation approach is affected by potential biases within LLMs, raising concerns about the accuracy and reliability of the evaluation results of LLMs. To address this issue, we propose and explore two many-shot In-Context Learning (ICL) prompt templates to help LLM evaluators mitigate potential biases: Many-Shot with Reference (MSwR) and Many-Shot without Reference (MSoR). Specifically, the former utilizes in-context examples with model-generated rationales as references, while the latter does not include these references. Using these prompt designs, we investigate the impact of increasing the number of in-context examples on the consistency and quality of the evaluation results. Experimental results show that advanced LLMs, such as GPT-4o, perform better in the many-shot regime than in the zero-shot regime. Furthermore, in most cases, MSwR performs significantly better than MSoR.

1 Introduction

LLMs such as GPT-4o (OpenAI, 2023), Gemini1.5-Pro (Reid et al., 2024), and Claude3.5-Sonnet (Anthropic, 2024) have demonstrated remarkable capabilities across a wide range of Natural Language Processing (NLP) tasks, becoming integral tools in various applications. The rapid advancement of LLMs (Chowdhery et al., 2023) underscores the critical need to evaluate their alignment with human intent in generated responses. Therefore, evaluation has emerged as a crucial research area pivotal to the success of LLMs (Chang et al., 2023), especially for using LLMs as evaluators.

LLMs like GPT-4 have shown exceptional performance across various tasks, leading to their wide adoption as both evaluators (Wang et al., 2023a; Fu et al., 2023a; Wang et al., 2023c; Zheng et al., 2023;




Figure 1: Consistency between different versions of evaluation results by adopting GPT-4o as a zero-shot evaluator. v_1 and v_2 are the results based on Prompt(A) in Table 1. v_1 , v_2 , and v_3 are results based on Prompt(B) in Table 1. Prompts (A) and (B) differ in whether to output the rating first or later. The consistency evaluations show that Prompt (A) and (B) almost obtain the agreement results, but the latter is convenient for constructing many-shot in-context examples, so we adopt the latter generated rationales in this study. v_1 vs. v_2 denotes comparing the first and second versions of evaluations. v_1 vs. v_2 vs. v_3 denotes the comparison between the three versions of evaluations.

Wang et al., 2023b; Chen et al., 2024) and annotators (Peng et al., 2023). However, the robustness of LLMs as evaluators remains uncertain, given their sensitivity to textual instructions (Xu et al., 2023; Turpin et al.) and potential biases (Wang et al., 2023b; Zheng et al., 2023; Chen et al., 2024). To this end, researchers, such as Wang et al. (2023b), focus on addressing the potential biases that exist when LLMs act as evaluators.

Newly expanded context windows of LLMs allow researchers to investigate ICL with more shots than the zero-shot and few-shot regimes, namely many-shot ICL. To fully investigate the many-shot ICL, Agarwal et al. (2024) explore the impact of the number of in-context examples by scaling shots to hundreds or thousands and find that many-shot can better reduce the biases of LLMs than the few-shot regime. Specifically, they show that using

Prompt(A). The zero-shot prompt for single answer grading in Zheng et al. (2023).

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. Begin your evaluation by providing a short explanation. Be as objective as possible. After providing your explanation, please rate the response on a scale of 1 to 10 by strictly following this format: "[[rating]]", for example: "Rating: [[5]]".

Question

{question}

Response

{response}

Prompt(B). The zero-shot prompt for single answer grading in this paper.

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. Please rate the response on a scale of 1 to 10 by strictly following this JSON format: {"rating": "", "reason": ""}. The "rating" should be as objective as possible. The "reason" denotes a comprehensive explanation of your rating, which should consider factors such as helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response should also be considered. Only return the JSON results and do not give any explanation.

Question

{question}

Response

{response}

Table 1: Two versions of zero-shot prompts are used for evaluation in this paper.

the many-shot in-context examples with chain-of-thought rationales generated through the zero-shot regime is effective, and the many-shot ICL may overcome the biases of LLMs, whereas few-shot ICL struggles. Therefore, the intuitive idea is to use the many-shot ICL, allowing LLMs as evaluators to see the zero-shot evaluations of similar questions and answers first and then scoring examples before scoring. Therefore, an interesting question arises:

- *Can many-shot in-context learning help LLMs as evaluators?*

Motivated by prior findings and the above issue, we preliminary verify the consistency of the widely used prompts of using LLMs as evaluators (Zheng et al., 2023), as shown in Table 1. Concretely, the consistency experiments are based on the entire test set of GSM8K (Cobbe et al., 2021), where the inference answers are obtained based on LLaMA3-70B (The pipeline as illustrated as in Figure 2 and details in § 3). Figure 1 presents that the consistency between the two versions of evaluations is low, with nearly half of the ratings being inconsistent.

Inspired by the above question, in this paper, we investigate whether many-shot in-context learning helps LLMs as evaluators. Specifically, we introduce two versions of prompts for LLM evaluators, **Many-Shot with Reference (MSwR)** and **Many-Shot without Reference (MSoR)**. The former utilizes in-context examples with model-generated rationales as guidance, and the latter removes rationales used in the former. Meanwhile, we reveal the symbol bias in LLMs and explore a simple approach for mitigating this issue. Experiments show




Figure 2: The pipeline of the experiments.

that many-shot ICL can help GPT-4o-as-a-Judge obtain higher quality and consistent evaluation results. As the number of in-context examples increases, the quality and consistency of evaluation improves significantly. Furthermore, we further verify the effectiveness of the proposed simple yet effective approach for mitigating the symbol bias in pairwise comparison of GPT-4o as an evaluator. To the best of our knowledge, we are the first to attempt to study LLM as an evaluator using the many-shot ICL regime.

2 Methodology

2.1 Background of Many-Shot ICL

LLMs excel at few-shot in-context learning, which involves learning from a few input-output demonstrations (“shots”) provided in context at inference without weight updates (Brown et al., 2020). Newly expanded long-context LLMs allow us to investigate ICL with hundreds of in-context examples (Li et al., 2023; Agarwal et al., 2024).

2.2 Recalling Potential Biases in LLMs

LLMs as evaluators possess potential biases, which have been widely explored by Wang et al. (2023b); Wu and Aji (2023); Zheng et al. (2023); Chen et al. (2024); Zheng et al. (2024). For example,




Figure 3: Evaluate the results of LLaMA3-70b on the GSM8K dataset using the Prompt(A).




Figure 4: Results of random selection.

Do not ask why, randomly select one from the following symbols. Just return the symbol you choose, no explanation is needed.

1 2 3 4 5 6 7 8 9 10

Table 2: Prompt template for testing selection bias.

positional bias in LLMs refers to the phenomenon where, during pairwise comparisons, LLM evaluators tend to favor one side of a pair regardless of the actual quality of the answers. In this paper, we investigate whether leveraging many-shot ICL helps LLMs as better evaluators.

To select the tested LLM and reveal the selection bias, we conducted a simple symbol selection experiment (set temperature=0.7 for all tested LLMs), using the designed prompt in Table 2 to let LLMs choose a symbol. The results are shown in Figure 4, and from these results, it can be observed that different LLMs tend different symbols. Meanwhile, we also evaluate the performance of LLaMA3-70B on the GSM8K dataset using Prompt(A) in Table 1 as shown in Figure 3. From the results, we can observe that the evaluation results of GPT-4o is relatively reasonable. Moreover, previous studies have demonstrated that GPT4 series models have the highest consistency with human evaluation.

2.3 Two Prompt Templates

In this section, we introduce two prompt templates for employing LLM as an evaluator: Many-Shot with Reference (MSwR) and Many-Shot without Reference (MSoR).

2.3.1 Many-Shot with Reference

Usually, when asking LLMs to select a score from 1-10 for evaluating answers to questions, providing reference answers may improve the evaluation quality of LLMs (Wang et al., 2023b; Zheng et al., 2023). In the zero-shot ICL regime, no in-context examples are provided for the GPT-4o evaluator, which selects scores only depending on itself. In the few-shot ICL regime, a few in-context examples are provided for the GPT-4o evaluator. In the many-shot ICL regime, many in-context examples are provided for the GPT-4o evaluator. The differences among the three regimes above are as follows: (1) zero-shot, which lacks any reference information, causes LLM to score entirely based on its preferences. (2) few-shot, due to providing only a small number of examples, may lead to evaluation results that lack diversity. (3) many-shot, by offering a larger number of examples, can ensure diversity and quality in evaluations. Therefore, we present the prompt template MSwR, which uses model-generated rationales as the in-context examples. Specifically, MSwR-4 indicates the prompt template of MSwR using 4 shots with reference as a demonstration, as shown in Table 3.

2.3.2 Many-Shot without Reference

Previous studies (Agarwal et al., 2024) find that in-context examples may limit the problem-solving approach of LLMs, so we propose a prompt template that does not need a reference for each in-context example. The designed prompt contains three parts: (1) a preamble, such as, “You will be provided questions similar to the ones below:”; (2)

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. Please rate the response on a scale of 1 to 10 by strictly following this JSON format: {"rating": "", "reason": ""}. The "rating" should be as objective as possible. The "reason" denotes a comprehensive explanation of your rating, which should consider factors such as helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response should also be considered. Only return the JSON results and do not give any explanation.

Now, I am going to give you a series of demonstrations of <Question>, <Response>, and Evaluations. When you respond, respond only with the Evaluation of the final pair of <Question> and <Response>, thinking step by step.

```

<Question>
{question1}
<Response>
{response1}
Evaluation
{evaluation1}
—
...
—
<Question>
{question4}
<Response>
{response4}
Evaluation
{evaluation4}
—
<Question>
{final_question}
<Response>
{final_response}
Evaluation

```

Table 3: Example of the prompt MSwR-(K=4).

a list of unsolved inputs or problems.; (3) a few-shot prompt with outputs for the desired output format. Specifically, MSoR-(K=4)-4 indicates that using 4 shots without reference and 4 shots with reference, as shown in Table 4.

3 Experiments

3.1 Experimental Settings

In our experiments, we use LLaMA3-70B to generate answers for each question in GSM8K (Cobbe et al., 2021) with a temperature of 0.7. GSM8K¹ is previously introduced by Cobbe et al. (2021), which comprises 8.5K high-quality grade school math problems meticulously crafted by human problem writers. This dataset is divided into 7.5K training and 1K test problems. Specifically, **each problem typically requires between 2 to 8 steps to solve**, primarily involving a sequence of elementary calculations using basic arithmetic operations (addition, subtraction, multiplication, and division). The problems are designed so that a proficient middle school student can solve each one. Further-

¹<https://huggingface.co/datasets/openai/gsm8k>

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. Please rate the response on a scale of 1 to 10 by strictly following this JSON format: {"rating": "", "reason": ""}. The "rating" should be as objective as possible. The "reason" denotes a comprehensive explanation of your rating, which should consider factors such as helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response should also be considered. Only return the JSON results and do not give any explanation.

Now, I am going to give you a series of demonstrations of <Question> and <Response>.

```

<Question>
{question1}
<Response>
{response1}
—
...
—
<Question>
{question4}
<Response>
{response4}
—
Now, I am going to give you a series of demonstrations of <Question>,
<Response>, and Evaluations. When you respond, respond only with the
Evaluation of the final pair of <Question> and <Response>, thinking step
by step.
—
<Question>
{question1}
<Response>
{response1}
Evaluation
{evaluation1}
—
...
—
<Question>
{question4}
<Response>
{response4}
Evaluation
{evaluation4}
—
<Question>
{final_question}
<Response>
{final_response}
Evaluation

```

Table 4: Example of the prompt MSoR-(K=4)-4.

more, the problem-solving task requires models to solve problems with model-generated rationales, which may be challenging to evaluate. Then, we use LLaMA3-70B² to create model-generated rationales for GSM8K. There are two reasons for using LLaMA3 to infer the answers to the problems in GSM8K. First, the quality of answers obtained by LLaMA3 may be high or low, which makes it beneficial to use these answers to analyze the GPT-4o as an evaluator. Second, if GPT-4o is adopted to infer the answers, using GPT-4o to evaluate again has potential biases. We use the training set of

²<https://github.com/meta-llama/llama3>




Figure 5: Consistency between two versions of evaluation results. Concretely, the bar corresponding to "0" on the x-axis represents the number of samples with consistent and inconsistent ratings in comparing evaluation results obtained twice using GPT-4o as the evaluator in the zero-shot regime. In addition, the zero-shot generated rationales are used for MSwR and MSoR. The bar corresponding to " 2^7 " on the x-axis represents the consistency of using the GPT-4o as an evaluator in MSwR.




Figure 6: Compare the consistency of the results from the two evaluations.

GSM8K as the sampling pool of in-context samples and use the first 200 samples of the test set in GSM8K for experiments in this paper. Specifically, we randomly sample in-context examples for each K -shot prompt in each test data for reliable results. To ensure that using more shots provides additional information, any K -shot ICL prompt in our experiments includes all in-context examples from prompts with less than K examples (all examples in the training set).

Inspired by Counting-Stars (Song et al., 2024), even if too many shots are provided, the GPT-4o may not be able to utilize all of them. Because, in the Counting-Stars benchmark, when the num-

ber of pieces of evidence reaches 32, LLMs (e.g., GPT-4 Turbo, Gemini 1.5 Pro, and Claude3 Opus) may no longer accurately obtain all of them. Therefore, adding more many-shot in-context examples is probably not captured by LLMs for learning as a reference to evaluate. However, the many-shot regime and Counting-Stars are substantially different, so there is no noise from "haystack". Hence, we set the maximum number of the in-context examples to 128, i.e., 128-shots.

In the experiments, we use GPT-4o with public API access, and the specific endpoint is "gpt-4o-2024-05-13". We use Claude3.5-Sonnet with public API access, and the specific endpoint is "claude-

	2^0	2^1	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9
Average Context Length in Character of A Single Test Sample										
MSwR-K	0.7K	1.1K	2.2K	3.6K	6.5K	13.6K	26.3K	50.7K	98.3K	198.5K
MSoR-K-4	2.2K	2.6K	3.1K	4.6K	6.9K	11.7K	20.9K	40.8K	81.3K	159.7K

Table 5: Context length in tokens for K-shot MSwR-K and MSoR-K-4 ($K = 2^0, 2^1, 2^2, 2^3, 2^4, 2^5, 2^6, 2^7, 2^8, 2^9$).

Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below. You should choose the assistant that follows the user’s instructions and answers the user’s question better. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of their responses. Begin your evaluation by comparing the two responses and provide a short explanation. Avoid any position biases and ensure that the order in which the responses were presented does not influence your decision. Do not allow the length of the responses to influence your evaluation. Do not favor certain names of the assistants. Be as objective as possible. After providing your explanation, output your final verdict by strictly following this format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a tie.

Problem

Please act as an impartial judge and evaluate the quality of the response provided by an AI assistant to the user question displayed below. Your evaluation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response. Please rate the response on a scale of 1 to 10 by strictly following this JSON format: {"rating": "", "reason": ""}. The "rating" should be as objective as possible. The "reason" denotes a comprehensive explanation of your rating, which should consider factors such as helpfulness, relevance, accuracy, depth, creativity, and level of detail of the response should also be considered. Only return the JSON results and do not give any explanation.

Question

{question}

Response

{response}

The Assistant A’s Answer {Answer-A}	The Assistant A’s Answer {Answer-B}	The Assistant B’s Answer {Answer-A}	The Assistant B’s Answer {Answer-B}
The Assistant B’s Answer {Answer-B}	The Assistant B’s Answer {Answer-A}	The Assistant A’s Answer {Answer-B}	The Assistant A’s Answer {Answer-A}
<i>Compare(A, B)</i>	<i>Compare(B, A)</i>	<i>Compare(A[†], B[†])</i>	<i>Compare(B[†], A[†])</i>

Table 6: The prompts used for pairwise comparison.

3-5-sonnet-20240620”. For the convenience of introduction, when comparing zero-shot and many-shot regimes, we uniformly refer to the few-shot and many-shot regimes as the many-shot regime. The context length in character-level of a single test sample with K -shot in-context examples is presented in Table 5.

3.2 Consistency Evaluation

We investigate the consistency between different versions of evaluation results generated by the GPT-4o evaluator. Here, the single answer grading evaluation results can be used to compare the consistency between different versions. As shown in Figure 5, the bar corresponding to " 2^n " on the x-axis represents the number of samples with consistent and inconsistent ratings in comparing evaluation results obtained twice using GPT-4o as an evaluator in MSwR. From the results, consistency improves as we increase the number of shots provided as in-context examples during inference.

Recent studies (Wang et al., 2023b; Zheng et al., 2023; Chen et al., 2024) have shown that the performance of GPT-4 as an evaluator is highly in

agreement with those of humans. However, both human and LLM evaluators are subject to potential biases (Chen et al., 2024). By analyzing prior studies, we suppose it unnecessary to obtain utterly accurate evaluation results (because this is difficult) by using LLMs as evaluators. It is only required to ensure that the evaluation results are highly consistent multiple times so that the single answer grading of GPT-4o as an evaluator may be effective. From all results in this work, we find that the many-shot ICL examples help the evaluation of LLMs more consistently, which is essential. We consider that the main reason may be that the many-shot in-context examples mitigate the potential biases of the GPT-4o evaluator.

Meanwhile, we also implemented experiments via the prompt template MSoR, but the results show that this regime may be unsuitable in the scenario of acting as an evaluator because the problem of each sample is the same, but the scored questions and answers are different. From the results in Figure 6, it can be seen that the consistency corresponds to the number of appended in-context samples with model-generated rationales but nearly

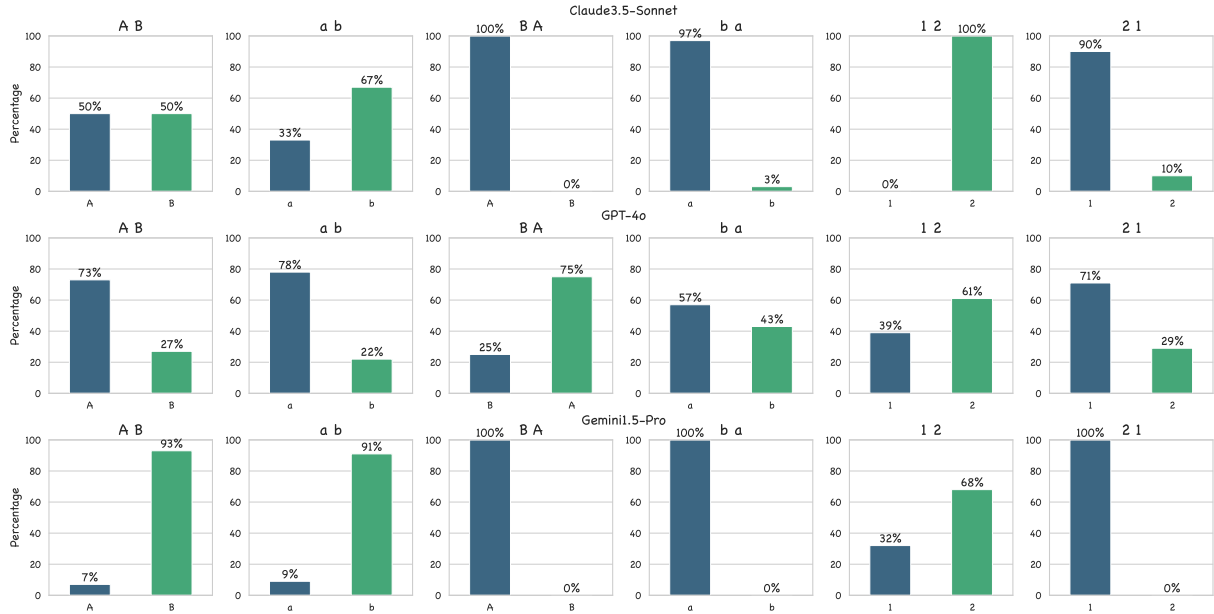


Figure 7: Results of randomly selecting different symbols.

not to the number of in-context examples without model-generated rationales.

In addition, both the few-shot and many-shot regimes are sensitive to the selection and order of in-context examples. In the experimental setting of this paper, almost no test data will have the same in-context examples. Even under this condition, the evaluation results also show a high consistency, demonstrating the effectiveness of many-shot ICL in helping GPT-4o evaluator.

3.3 Quality Evaluation

After the consistency evaluation, an important question arises about:

- *Does a high consistency refer to high-quality evaluations?*

Around this issue, we pairwise compare model-generated evaluation rationales between zero-shot and many-shot regimes using the designed prompts shown in Table 6. In pairwise comparison, the GPT-4o evaluator is presented with a question and two answers and tasked with *determining which is better or declaring a tie*. From Figure 8 *Compare(A, B)*, we can find that the evaluations obtained in the many-shot regime are significantly better than those in the zero-shot regime. However, as mentioned before, the positional bias of GPT-4o may cause these results, so we performed a second comparison by swapping the positions. As shown in Figure 8 *Compare(B, A)*, we observe that as in-context examples increase, the evaluation results in

the many-shot regime gradually turn around, that is, a higher winning rate. To obtain fairness results, we integrate the above two results (*Compare(A, B)* and *Compare(B, A)*), as shown in Figure 9. It can be seen that after mitigating the positional bias, the evaluation quality in the many-shot regime is still better than that in the zero-shot regime.

3.4 Revealing Symbol Bias

Zheng et al. (2024) discover that these LLMs face selection bias, which means they prefer to select specific option IDs as answers (like "Option A"). Meanwhile, Song et al. (2024) find that when an LLM is stress-tested, it is easy to output some wrong information, which may be an increasing array related to the test data or an English alphabet sequence starting with "A". This phenomenon shows that LLMs favor answers with the symbol "A" rather than the symbol "B" (or the symbol "1" instead of the symbol "2"). Therefore, an interesting question arises about:

- *Do LLMs prefer to choose the answer with the symbol A or B?*

To answer this question, we utilize the prompt template in Table 2 and replaced the symbols to be selected, conducting multiple experiments where LLMs chose symbols (set temperature=0.7 for all tested LLMs), as shown in Figure 7. It can be seen that the bias actually exists in LLMs when facing selection toward different symbols. Comparing the




Figure 8: Win rate compared between zero-shot and many-shot regimes.

three LLMs at the same time, we find that the symbol bias of GPT-4o is relatively small, so in the pairwise comparison, we leverage GPT-4o as the evaluator. In addition, we find an interesting phenomenon that the position bias in the Claude3.5-sonnet and Gemini1.5-Pro are stronger than the symbol bias, and these LLMs tend to choose the options where the position is at the back. In contrast, GPT-4o tends to choose the symbols at the front. In other words, for Claude3.5-sonnet and Gemini1.5-Pro, the effect of position bias is heavier, and for GPT-4o the symbol bias impacts more.

Inspired by our experiments, we swap the answers corresponding to symbols A and B, as shown in Figure 6. From Figure 8 $Compare(A^\dagger, B^\dagger)$ and $Compare(B^\dagger, A^\dagger)$, it can be seen that the evaluation results are different from the experiments. Actually, the results of $Compare(A, B)$ and $Compare(A^\dagger, B^\dagger)$ should be similar, and the results of $Compare(B, A)$ and $Compare(B^\dagger, A^\dagger)$ should be similar. This phenomenon shows that symbol bias does exist when adopting GPT-4o as an evaluator.

Recent research (Wang et al., 2023b) integrates the evaluation results of $Compare(A, B)$ and $Compare(B, A)$ to mitigate the positional bias, which motivates us to incorporate the evaluation results of $Compare(A^\dagger, B^\dagger)$ and $Compare(B^\dagger, A^\dagger)$ to reduce symbol bias. As presented in Figure 9, it can be seen that as in-context examples increase, the higher the win rate of the many-shot regime, which further verifies the effectiveness of the many-shot regime in helping GPT-4o as an evaluator.

4 Related Work

LLMs have exhibited remarkable general generation capabilities, positioning themselves as powerful assistants (Zhao et al., 2023; OpenAI, 2023). With the rapid progression of LLMs, evaluating their proficiency in adhering to human instructions is imperative. Given the advanced capabilities of




Figure 9: Comparison evaluation results after mitigating the biases via different approaches.

LLMs, researchers have begun adopting these models to evaluate the performance of LLMs in following human instructions (Koo et al., 2023; Liusie et al., 2023; Liu et al., 2023; Zhu et al., 2023; Lu et al., 2023; Fu et al., 2023b; Zheng et al., 2023; Wang et al., 2023b; Chen et al., 2024). Notably, the evaluation paradigm introduced by Zheng et al. (2023) has gained widespread adoption. However, LLMs as evaluators are revealed to have potential biases (Wang et al., 2023b; Chen et al., 2024), which leads to higher uncertainty and inconsistency during the evaluation using LLMs, questioning the validity of LLM evaluators.

5 Conclusion

In this work, we investigate and explore whether many-shot ICL helps LLMs as evaluators, such as GPT-4o. To this end, we designed two prompt templates, e.g., **Many-Shot with Reference (MSwR)** and **Many-Shot without Reference (MSoR)**. Experiments show that many-shot ICL can help the GPT-4o evaluator improve the consistency and quality of

evaluation. Meanwhile, we also revealed symbol bias in LLMs when LLMs act as evaluators, and further proposed a simple yet effective approach to mitigate the symbol bias.

6 Limitations

Considering the trade-off between costs and benefits, we do not verify too many in-context examples in the experiments, such as thousands of examples. Combining Figures 5 and Figure 9, it is not difficult to see that when the number of in-context examples increases to 256 and 512, although the consistency no longer improves, the evaluation quality has increased significantly. In addition, we consider that using GPT-4o as an evaluator in the many-shot regime is another evolution of the weak-to-strong strategy (Burns et al., 2023), which uses many zero-shot evaluation results to generate a better one. As the long-context capabilities of LLMs improve, adding more in-context examples may reveal more valuable phenomena for studying LLMs as evaluators in the future. There are still many unexplored features of many-shot in-context learning. One possibility is that it can enable LLMs to simulate data constructed from shots, thereby increasing the diversity of synthetic data. Another possibility is to design benchmarks based on many-shot in-context examples and to assess the long-context LLMs based on the capabilities of understanding and imitating many shots.

Acknowledgments

We thank the three anonymous reviewers for carefully reading our paper and their insightful comments and suggestions.

References

- Rishabh Agarwal, Avi Singh, Lei M. Zhang, Bernd Bohnet, Stephanie C. Y. Chan, Ankesh Anand, Zabeer Abbas, Azade Nova, John D. Co-Reyes, Eric Chu, Feryal M. P. Behbahani, Aleksandra Faust, and Hugo Larochelle. 2024. [Many-shot in-context learning](#). *CoRR*, abs/2404.11018.
- Anthropic. 2024. The claude 3 model family: Opus, sonnet, haiku. *Technical Report*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems*.
- Collin Burns, Pavel Izmailov, Jan Hendrik Kirchner, Bowen Baker, Leo Gao, Leopold Aschenbrenner, Yining Chen, Adrien Ecoffet, Manas Joglekar, Jan Leike, Ilya Sutskever, and Jeff Wu. 2023. [Weak-to-strong generalization: Eliciting strong capabilities with weak supervision](#). *CoRR*, abs/2312.09390.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Kaijie Zhu, Hao Chen, Linyi Yang, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2023. [A survey on evaluation of large language models](#). *CoRR*, abs/2307.03109.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. [Humans or llms as the judge? A study on judgement biases](#). *CoRR*, abs/2402.10669.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2023. [Palm: Scaling language modeling with pathways](#). *J. Mach. Learn. Res.*, 24:240:1–240:113.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *CoRR*, abs/2110.14168.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023a. [Gptscore: Evaluate as you desire](#). *CoRR*, abs/2302.04166.
- Xue-Yong Fu, Md. Tahmid Rahman Laskar, Cheng Chen, and Shashi Bhushan TN. 2023b. [Are large](#)

- language models reliable judges? A study on the factuality evaluation capabilities of llms. *CoRR*, abs/2311.00681.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2023. Benchmarking cognitive biases in large language models as evaluators. *CoRR*, abs/2309.17012.
- Mukai Li, Shansan Gong, Jiangtao Feng, Yiheng Xu, Jun Zhang, Zhiyong Wu, and Lingpeng Kong. 2023. In-context learning with many demonstration examples. *CoRR*, abs/2302.04931.
- Yiqi Liu, Nafise Sadat Moosavi, and Chenghua Lin. 2023. Llms as narcissistic evaluators: When ego inflates evaluation scores. *CoRR*, abs/2311.09766.
- Adian Liusie, Potsawee Manakul, and Mark J. F. Gales. 2023. Zero-shot NLG evaluation through pairwise comparisons with llms. *CoRR*, abs/2307.07889.
- Qingyu Lu, Baopu Qiu, Liang Ding, Liping Xie, and Dacheng Tao. 2023. Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt. *CoRR*, abs/2303.13809.
- OpenAI. 2023. GPT-4 technical report. *CoRR*, abs/2303.08774.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with GPT-4. *CoRR*, abs/2304.03277.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, and et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *CoRR*, abs/2403.05530.
- Mingyang Song, Mao Zheng, and Xuan Luo. 2024. Counting-stars: A multi-evidence, position-aware, and scalable benchmark for evaluating long-context large language models. *Preprint*, arXiv:2403.11802.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023*.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023a. Is chatgpt a good NLG evaluator? A preliminary study. *CoRR*, abs/2303.04048.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023b. Large language models are not fair evaluators. *CoRR*, abs/2305.17926.
- Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. 2023c. Pandalm: An automatic evaluation benchmark for LLM instruction tuning optimization. *CoRR*, abs/2306.05087.
- Minghao Wu and Alham Fikri Aji. 2023. Style over substance: Evaluation biases for large language models. *CoRR*, abs/2307.03025.
- Canwen Xu, Daya Guo, Nan Duan, and Julian J. McAuley. 2023. Baize: An open-source chat model with parameter-efficient tuning on self-chat data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 6268–6278. Association for Computational Linguistics.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A survey of large language models. *CoRR*, abs/2303.18223.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. Large language models are not robust multiple choice selectors. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems*, volume 36, pages 46595–46623.
- Lianghui Zhu, Xinggang Wang, and Xinlong Wang. 2023. Judgelm: Fine-tuned large language models are scalable judges. *CoRR*, abs/2310.17631.