



The Comparative Trap: Pairwise Comparisons Amplify Biased Preferences of LLM Evaluators

Hawon Jeong*¹ ChaeHun Park*¹ Jimin Hong^{1,2} Hojoon Lee¹ Jaegul Choo¹

¹ KAIST AI ² KRAFTON

{hawon, ddehun, jimmyh, joonleesky, jchoo}@kaist.ac.kr

Abstract

As large language models (LLMs) are increasingly used as evaluators for natural language generation tasks, ensuring unbiased assessments is essential. However, LLM evaluators often display biased preferences, such as favoring verbosity and authoritative tones. Our empirical analysis reveals that these biases are exacerbated in pairwise evaluation, where LLMs directly compare two outputs and easily prioritize superficial attributes. In contrast, pointwise evaluation, which assesses outputs independently, is less susceptible to such bias because each output is judged in isolation. To address the limitations of the pairwise evaluation, we introduce a novel evaluation method, PREPAIR, which integrates pointwise reasoning within a pairwise framework. PREPAIR effectively alleviates biased preference, improving performance on the adversarial benchmark (LLMBar) while outperforming pointwise evaluation on the standard benchmark (MT-Bench).

1 Introduction

Automated evaluation systems are crucial for scaling natural language generation (NLG) research, offering a more efficient and cost-effective alternative to human evaluation. Recently, large language models (LLMs) have been increasingly adopted as automatic evaluators, demonstrating promising capabilities in assessing various aspects of text quality such as coherence, relevance, and fluency (Chiang and Lee, 2023; Liu et al., 2023).

Despite their impressive performance, LLM-based evaluators often display systematic biases in their preferences—for instance, favoring verbose responses (Wang et al., 2024a), exhibiting egocentric judgment (Zheng et al., 2024), or preferring authoritative tones (Chen et al., 2024). Recent studies have begun to investigate these biases by probing the evaluator’s ability to follow explicit

evaluation instructions in controlled settings (Zeng et al., 2024; Lambert et al., 2024). In these evaluations, one output strictly follows the instructions, while the other is adversarially crafted to exploit evaluator biases—appearing superficially favorable yet failing to comply with the instructions. These studies reveal that LLM evaluators not only exhibit biased preferences but also struggle to identify these deceptive outputs effectively.

Building on these observations, we first explore how different evaluation setups influence the biased preferences in LLM evaluators. Specifically, we compare two common setups: *pairwise*, where the evaluator compares two outputs simultaneously and selects the better one, and *pointwise*, where each output is evaluated independently and then selects the one with a higher score. Our experiments demonstrate that the pairwise setup amplifies the biases of the evaluator, performing significantly worse than the pointwise setup on adversarial examples (see Fig. 1 Right). Further analysis reveals that pairwise evaluators, while recognizing weaknesses in incorrect outputs, often reach incorrect conclusions.

We hypothesize that the increased bias in pairwise evaluations stems from the structural differences between the pairwise and pointwise setups. To illustrate, consider a job interview where the interviewer must choose between Alice and Bob for a research engineer position. Although the goal is to select the candidate more suited for the role, the interviewer might favor Alice simply if she has more fluent speaking skills despite Bob possibly having stronger technical skills. Psychological studies support this phenomenon: when directly comparing options, humans often focus on superficial features rather than thoroughly evaluating each option (Wright, 1974; Payne, 1976; Olshavsky, 1979). Similarly, in a pairwise setup, LLMs may concentrate on easily exploitable, superficial attributes, causing adversarial outputs to

* Equal contribution

appear favorable even when they fail to follow the instructions.

In contrast, pointwise evaluation assesses each output independently, which allows the model to focus solely on how well the output adheres to the instructions, without being interrupted by comparative attributes. This isolation reduces susceptibility to biases and leads to a more accurate assessment of task compliance. However, pointwise evaluation is not a silver bullet for all scenarios; Suppose Alice and Bob are equally fluent speakers. In this case, evaluating both candidates simultaneously can provide additional context about who is ultimately better suited for the role. In such cases, the pairwise approach benefits from leveraging contextual information for non-adversarial examples (See Fig. 1 Left).

Motivated by our findings, we propose a simple yet effective method for reliable pairwise evaluation, namely PREPAIR, which stands for **P**ointwise **R**Easoning within a **P**AIRwise evaluation framework (Fig. 2). The key idea is to first extract unbiased insights from each output using pointwise reasoning. Then, we combine these individual explanations into a final decision process for pairwise evaluation, leveraging the strength of both pointwise and pairwise approaches. Experimental results across various models demonstrate that PREPAIR improves the performance of pairwise evaluators on adversarial datasets. Specifically, PREPAIR improves average accuracy by 24.48% across six LLMs on the adversarial dataset. Moreover, it outperforms the pointwise approach on normal datasets, demonstrating its versatility and effectiveness.

Our contributions can be summarized as follows: (1) We analyze how pointwise and pairwise evaluation setups affect LLM assessment in adversarial scenarios. (2) We identify biases in pairwise evaluation that favor superficial attributes, making adversarial responses more likely to be preferred. (3) We propose a hybrid evaluation framework by integrating pointwise reasoning into pairwise evaluation.

2 Related Work

LLMs-based Evaluation The evaluation of NLG systems has traditionally relied on human annotation and automatic metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee and Lavie, 2005). However, these

metrics often fail to capture semantic adequacy and coherence (Liu et al., 2016; Sai et al., 2022). To address these limitations, large language models (LLMs) have emerged as evaluators due to their strong instruction-following capabilities and generalization across diverse tasks (Kocmi and Federmann, 2023; Lin and Chen, 2023; Min et al., 2023; Chan et al., 2024; Wang et al., 2024b).

Among various evaluation strategies, pairwise comparison has been increasingly utilized, where an LLM ranks or selects the better response between two candidates (Chiang et al., 2024; Zeng et al., 2024; Wang et al., 2024a; Zhang et al., 2023). This method aligns with human evaluation practices and has been shown to correlate well with human judgments (Zheng et al., 2024). However, LLM evaluators exhibit systematic biases, such as favoring verbose responses (Li et al., 2023; Zheng et al., 2024) and overrating their own generations (Li et al., 2023; Koo et al., 2024). Moreover, their rankings can be inconsistent, as slight prompt variations may lead to different outcomes (Liu et al., 2023; Zhu et al., 2024). Recent studies suggest that incorporating explicit evaluation criteria and general guidelines in prompts can improve their discernment (Zeng et al., 2024). Building on these insights, we examine how different evaluation setups influence LLM-based assessment, particularly in adversarial scenarios where biases may be exacerbated.

Human Behavior with Complex Tasks Psychological research has demonstrated that decision-makers often simplify their process by focusing on limited information when handling complex tasks (Payne, 1976; Olshavsky, 1979; Einhorn and Hogarth, 1981). Moreover, Wright (1974) has indicated that cognitive overload from environmental distractions reduces decision quality, causing people to overlook important evidence or rely on irrelevant information. These findings parallel our intuitions in adversarial evaluation: assessing multiple outputs at once encourages the model to prioritize superficial aspects, similar to humans favoring simpler solutions in complex tasks.

3 Analyzing LLM Evaluators on Adversarial Evaluation Samples

We now analyze LLM-based evaluators using different datasets and prompting strategies. We introduce meta-evaluation datasets with different characteristics and then compare two representative

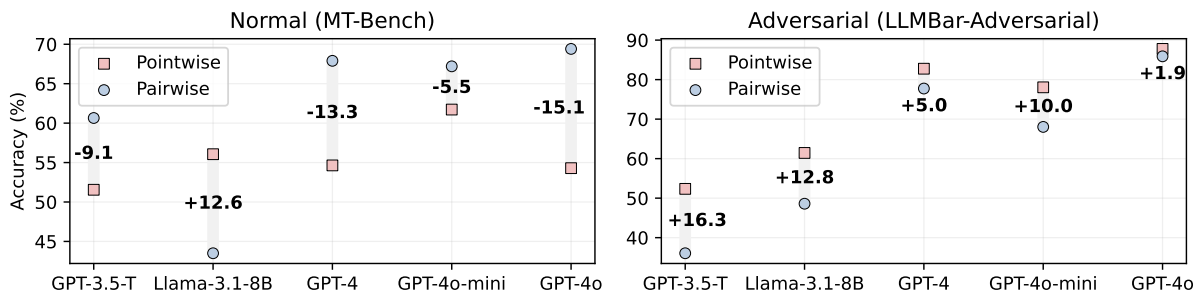


Figure 1: Results of **pointwise** and **pairwise** approaches in normal (MT-Bench) and adversarial (LLMBar-Adversarial) meta-evaluation datasets. The difference values represent the change between pointwise and pairwise approaches for each model under both normal and adversarial conditions.

evaluation approaches on these datasets. Finally, we examine the impacts of intrinsic biases in LLM evaluators by exploring their reasoning process.

3.1 Meta-Evaluation Datasets

In our experiments, we evaluate LLM evaluators using datasets with two distinct characteristics: **Normal** and **Adversarial**. The normal dataset, *MT-Bench* (Zheng et al., 2024), contains responses generated under standard conditions and reflects real-world evaluation distributions. In contrast, the adversarial dataset, *LLMBar-Adversarial* (Zeng et al., 2024), consists of responses that appear high-quality but subtly deviate from the original instruction, testing an evaluator’s ability to detect misleading outputs. Each evaluation sample in both datasets consists of (1) an input instruction, (2) a pair of two model-generated responses, and (3) a human preference label indicating which response is preferred. The datasets contain a total of 1131 evaluation samples from *MT-Bench* and 319 from *LLMBar-Adversarial*. Further details on each dataset are provided below.

MT-Bench (Zheng et al., 2024) is a meta-evaluation dataset consisting of 80 multi-turn instructions, each paired with model-generated outputs from various language models. For our experiments, we use the first-turn meta-evaluation samples. The dataset includes a total of 1131 human judgments from two groups: expert labelers (885) and the original authors (246). This dataset is widely used to assess evaluation metrics in a setting that reflects real-world model performance.

LLMBar-Adversarial (Zeng et al., 2024) is specifically designed to test the ability of evaluators to discern instruction-following responses. The dataset consists of 319 adversarial evaluation samples, where generated responses exhibit high superficial quality but fail to adhere to the original instruction. These samples are categorized into four types: Neighbor Instructions (134), GPT-4 In-

structions (92), GPT-4 Unhelpful Outputs (47), and Manual Construction (46). This dataset serves as a benchmark for evaluating whether LLM evaluators can detect subtle deviations from the expected responses.

3.2 Setup

Pointwise and Pairwise Approaches We compare the two popular approaches for pairwise text evaluation: *pairwise* and *pointwise* (Zheng et al., 2024; Liusie et al., 2024). In the pairwise approach, an LLM receives an instruction and two model-generated outputs and then chooses a higher-quality output. In the pointwise approach, an LLM assigns a single score on a Likert scale to each generated output independently. The output with the higher score is considered to be of higher quality. Please refer to Appendix A.2 for more details.

Implementation Details We use the following LLMs as evaluators to compare the approaches above (i.e., pointwise and pairwise) on *MT-Bench* and *LLMBar* datasets: Llama-3.1-8B-Instruct (Dubey et al., 2024), gpt-3.5-turbo, gpt-4, gpt-4o-mini, and gpt-4o (Achiam et al., 2023). We utilize the original prompts released by each dataset for evaluation. For *LLMBar*, we include general rules, such as prioritizing outputs that accurately follow instructions, as these guidelines generally enhance evaluation quality (Zeng et al., 2024). The performance of the evaluator is obtained by comparing its judgments to human judgments within the datasets. In *MT-Bench*, since a tie option is allowed, the pairwise evaluator is prompted to choose one of three options. The scalar value generated by the pointwise evaluator using the Likert-scale is directly used for evaluation.¹ In

¹Evaluation results of pointwise evaluators with weighted summation on a dataset that does not allow tie-option (i.e., RewardBench-Chat (Lambert et al., 2024) are also presented in Section 5.3.4.

Evaluation Sample
<p>Instruction: Create a new headline that is witty and summarizes the article. A new study has found that eating spinach can increase intelligence.</p> <p>Output A: Spinach: The New Brain Food, or Just Another Leafy Green? (<i>Human-preferred</i>)</p> <p>Output B: Spinach It To Win: Eating Spinach Can Boost Your Brain Power!</p>
<p>Explanations from LLM Evaluator</p> <p>Strengths of Output A: Clever play on words with "Brain Food" and "Leafy Green". /Succinctly summarizes the article topic.</p> <p>Weaknesses of Output A: Lacks a direct mention of the study findings. / Could be perceived as slightly less engaging compared to Output B.</p> <p>Strengths of Output B: Engaging with the phrase "Spinach It To Win." / Clearly states the benefit of eating spinach for brain power.</p> <p>Weaknesses of Output B: The use of "Win" might be seen as slightly exaggerated or sensationalized. / Could be considered a bit informal for certain audiences.</p>
<p>Final Judgement of LLM Evaluator</p> <p>Therefore, Output B is better.</p>

Table 1: Explanations and judgment of pairwise evaluators from an evaluation sample. The correct and incorrect model explanations are manually highlighted.

LLMBar, the human judgments contain only two options without tie. Following Liu et al. (2023), we use a weighted summation with token probabilities for the pointwise evaluator.

3.3 Results

Pairwise Comparisons Amplifies the Biased Preference. Fig. 1 shows the meta-evaluation results on the two datasets. The pairwise evaluator performs better than the pointwise evaluator on MT-Bench, while the pointwise evaluator outperforms the pairwise evaluator on LLMBar. These contrasting results likely stem from the differences in the evaluation structure. The pairwise approach excels with normal samples like those in MT-Bench, since it can leverage contextual information by considering two outputs simultaneously. However, the benefits of the pairwise approach turn into drawbacks when handling adversarial samples in LLMBar, as it makes the evaluators more easily susceptible to intrinsic biases. In contrast, the pointwise approach is relatively less affected by these biases, resulting in better performance on adversarial samples.

From these observations, we have questioned how the pairwise approach undermines the discerning ability of LLM evaluators on adversarial examples. Therefore, our subsequent experiments explore how intrinsic biases impact LLM evaluators using the pairwise approach in adversarial evaluations.

Pairwise Evaluators Do Correct Reasoning behind Incorrect Evaluation. To further understand why pairwise evaluation struggles with ad-

versarial samples in LLMBar, we used a Chain-of-Thought (Kojima et al., 2022; Wei et al., 2022, CoT) prompting to let LLM provide its own reasoning process before making predictions. Specifically, we encouraged the LLM to explain the strengths and weaknesses of each output to see if it can accurately identify issues in low-quality outputs. Then, we manually examine the explanation provided by LLM for 100 incorrectly predicted samples. More details on human analysis are described in Appendix A.4.

Our analysis reveals that, despite its incorrect final decisions, the pairwise evaluator often provides rational explanations for identifying flaws in low-quality outputs. We find that 64 out of 100 explanations accurately describe weaknesses that align with our manual annotations. For example, as shown in Table 1, the pairwise evaluator correctly identifies problems in a flawed output but ultimately makes the wrong judgment. This indicates that while the LLM evaluator with the pairwise approach has the potential to perform better, it is hindered by intrinsic issues within the pairwise setup. Therefore, addressing these issues is essential for reliable pairwise evaluation.

4 PREPAIR: Pointwise Reasoning for Pairwise Evaluation

Approach Our analysis in Section 3 revealed two key limitations in existing LLM-based evaluation methods. First, pairwise evaluators often struggle with adversarial samples, as they may still prefer a problematic response despite recognizing its flaws.

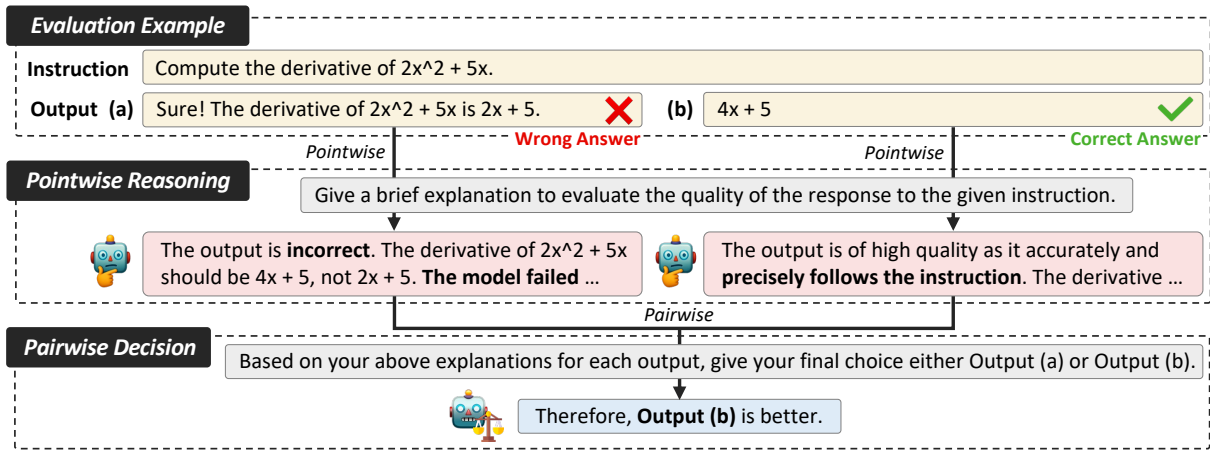


Figure 2: The overall illustration of PREPAIR.

Methods	GPT-3.5-T	Claude-3	Llama-3.1	GPT-4	GPT-4o-m	GPT-4o
Pointwise	52.35	-	61.44	82.76	78.05	87.77
Pairwise	36.05	40.75	48.59	77.74	68.03	85.89
w/ CoT	31.66	41.07	50.47	77.12	63.01	83.07
w/ PREPAIR*	<u>37.62</u>	<u>43.26</u>	<u>61.44</u>	<u>78.37</u>	<u>68.65</u>	86.21
w/ PREPAIR	54.23	53.60	70.84	80.88	78.06	86.21

Table 2: LLMBAR-Adversarial results. The pairwise approach’s highest and second-highest scores are highlighted in bold and underlined. PREPAIR* adopts pairwise reasoning but follows the generation format of PREPAIR.

Second, pointwise evaluators, which analyze responses independently, tend to perform worse on normal samples but show greater robustness against adversarial manipulations. These observations suggest the need for a method that retains the contextual advantages of pairwise evaluation while reducing biases that arise from direct comparison.

To address this, we propose PREPAIR (Pointwise Reasoning for Pairwise Evaluation), a hybrid approach that incorporates pointwise reasoning into a pairwise evaluation framework. Unlike conventional pairwise evaluation, which directly compares two responses, PREPAIR first analyzes each response independently using pointwise reasoning before making a final pairwise decision. Given an instruction and two candidate responses, the evaluator reasons about each response separately, considering its properties without direct reference to the other. These independent reasoning steps provide structured insights that are then used in the final pairwise decision, where the evaluator determines the superior response. By decoupling the reasoning process from direct comparison, PREPAIR helps mitigate biases such as favoring verbosity or model-specific stylistic preferences while maintaining the benefits of pairwise evaluation.

One trade-off of PREPAIR is the increased com-

putational cost, as it requires an additional LLM forward pass for each response before making a final comparison. However, this overhead can be reduced through caching, since the pointwise reasoning results for individual responses can be reused across multiple comparisons. This is particularly advantageous in ranking tasks, where precomputed reasoning steps allow for efficient and consistent evaluation. By combining the robustness of pointwise reasoning with the contextual awareness of pairwise evaluation, PREPAIR offers a more reliable evaluation framework, particularly in adversarial scenarios.

5 Experiments

5.1 Setup

We conduct experiments on the LLMs described in Section 3.2 and the Claude-3-Haiku (Anthropic, 2024). We compare our method for the pairwise approach against the following baselines. *w/ CoT* generates an explanation before making a prediction. PREPAIR* is an ablated version of our method that prompts pairwise evaluators to generate individual explanations for each output, using the same format as PREPAIR. This baseline aims to validate the effectiveness of the pointwise reasoning process.

Detailed setups and more results are in Appendix A and B.

Methods	GPT-3.5-T	GPT-4o-m	GPT-4	GPT-4o
Pointwise	51.55	61.72	54.64	52.29
Pairwise	60.65	67.20	67.90	69.41
w/ CoT	61.54	<u>65.87</u>	66.40	67.46
w/ PREPAIR*	60.83	64.72	<u>66.76</u>	69.14
w/ PREPAIR	<u>61.36</u>	64.10	65.34	69.50

Table 3: MT-Bench results. The indicators are the same as Table 2. Bold and underlined values indicate the highest and second-highest scores.

5.2 Results

As shown in Table 2, our method improves the accuracy of pairwise evaluators on the LLMBAR-Adversarial dataset across all tested LLMs, with the largest gain observed in Llama-3.1. By incorporating pointwise reasoning, PREPAIR mitigates biases in direct comparison and enhances robustness against adversarial responses. A key observation is that PREPAIR consistently outperforms PREPAIR*, indicating that its improvements stem from explicit pointwise reasoning rather than structural modifications alone. Additionally, PREPAIR sometimes performs even better than pointwise evaluators, suggesting that while pointwise reasoning enhances adversarial robustness, integrating it within a pairwise framework preserves the benefits of contextual evaluation. Qualitative examples in Appendix C.1 further illustrate how our method improves reliability compared to standard Chain-of-Thought (CoT) prompting.

On the MT-Bench dataset (Table 3), pairwise methods achieve similar performance and consistently outperform the pointwise approach. Since MT-Bench focuses on general response quality rather than adversarial robustness, this result suggests that PREPAIR does not compromise effectiveness in standard evaluation while offering clear benefits in adversarial settings. Overall, these results demonstrate that PREPAIR effectively balances the strengths of pointwise and pairwise evaluation, improving reliability against adversarial samples while maintaining strong performance in normal evaluation tasks.

5.3 Analysis

We conduct further analysis to understand PREPAIR comprehensively as follows.

5.3.1 Evaluating PREPAIR to Open-source LLMs

We use the following open-source LLMs to probe the effectiveness of PrePair: Llama-2-7b-chat (Touvron et al., 2023), Llama-3-8B-Instruct (Dubey et al., 2024), and Llama-3.1-8b-Instruct (Dubey et al., 2024). As shown in Table 4, PrePair consistently contributes to better performance of LLM evaluators on adversarial meta-evaluation datasets.

5.3.2 Evaluating PREPAIR to a Finetuned LLM Evaluator

We conducted experiments with Prometheus-v2-7B² (Kim et al., 2024), which supports both pointwise and pairwise evaluations. The evaluation results are shown in Table 5. We observe consistent trends with our previous results, where a pointwise evaluation usually performs better than a pairwise one in adversarial samples, and PrePair increases the performance of pairwise evaluators.

5.3.3 Extending PREPAIR with Different Reasoning Strategies

To assess the extensibility of PREPAIR, we apply it to the Self-Generated Metric approach (Saha et al., 2024; Zeng et al., 2024), where LLMs first generate instance-specific evaluation criteria before making judgments. These criteria are then incorporated into prompts to guide the evaluation process.

We compare PREPAIR with standard pairwise evaluation when generating both metrics and explanations. As shown in Table 6, PREPAIR consistently improves performance, suggesting that its reasoning mechanism enhances evaluation robustness. Moreover, its effectiveness across different reasoning setups indicates that PREPAIR can complement other evaluation strategies, making it a flexible addition to existing frameworks.

5.3.4 Impacts of Different Aggregation Strategies in Pointwise Approach

In pointwise evaluation, the final score for a response is typically derived through an aggregation strategy, which determines how individual scoring tokens contribute to the overall decision. A common approach is to compute a weighted sum over the probabilities of score tokens, which has been shown to improve performance (Liu et al., 2023). However, certain benchmarks, such as MT-Bench,

²<https://huggingface.co/prometheus-eval/prometheus-7b-v2.0>

Methods	Manual	GPTInst	GPTOut	Neighbor	Total
<i>Llama-2-13b-chat</i>					
Pairwise	28.26	28.26	51.06	29.10	31.97
w/ CoT	32.61	35.87	57.45	22.39	32.92
w/ PREPAIR	36.96	32.61	63.83	40.30	41.07
<i>Llama-3-8B-Instruct</i>					
Pairwise	45.65	48.91	48.94	32.09	41.38
w/ CoT	54.35	46.74	48.94	41.04	45.77
w/ PREPAIR	52.17	72.83	51.06	58.96	60.82
<i>Llama-3.1-8B-Instruct</i>					
Pairwise	60.87	73.91	61.70	52.99	61.44
w/ CoT	52.17	55.43	53.19	45.52	50.47
w/ PREPAIR	69.57	77.17	53.19	73.13	70.84

Table 4: Accuracies on the LLMBAR-Adversarial dataset for different Llama model variants. Each model is evaluated using Pairwise, Chain-of-Thought (CoT), and PrePair.

Methods	Manual	GPTInst	GPTOut	Neighbor	Total
Pointwise	47.83	54.35	55.32	43.28	48.90
Pairwise	50.00	33.70	57.45	32.84	39.19
w/ PREPAIR	58.70	57.61	57.45	35.82	48.59

Table 5: Accuracies on the LLMBAR-Adversarial dataset for Prometheus-v2-7b (Kim et al., 2024) model. Each model is evaluated using Pointwise, Pairwise, and PrePair.

Methods	Manual	GPTInst	GPTOut	Neighbor	Total
Pointwise	<u>56.52</u>	61.96	38.30	49.25	52.35
Pairwise	41.30	33.70	40.43	34.33	36.05
w/ CoT	28.26	39.13	46.81	22.39	31.66
w/ SGM	52.17	55.43	46.81	<u>55.97</u>	53.92
w/ PREPAIR	52.17	64.13	48.94	50.00	54.23
w/ PREPAIR + SGM	58.70	<u>63.04</u>	51.06	57.46	58.31

Table 6: Accuracies on the LLMBAR-Adversarial dataset using Pointwise, Pairwise, Chain-of-Thought (CoT), Self-Generated Metric (SGM) (Saha et al., 2024), PrePair, and PrePair with SGM.

allow tie predictions, requiring an alternative strategy. To accommodate this, our pointwise evaluator for MT-Bench directly uses the scalar value generated by the model instead of applying a weighted sum. In contrast, the RewardBench-Chat subset (Lambert et al., 2024) does not include tie options, allowing a direct comparison between aggregation strategies. Here, the pointwise evaluator follows the same setup as in LLMBAR, computing scores based on a weighted sum over Likert-scale integer tokens. The dataset, curated from AlpacaEval (Dubois et al., 2024) and MT-Bench, consists of 358 samples, each containing an instruction, two model outputs, and a preference label indicating the superior response without the possibility of a tie.

Table 7 presents the results, showing that even when using a weighted sum aggregation, pairwise evaluation still outperforms pointwise evaluation. This indicates that the aggregation strategy alone

does not bridge the performance gap between the two approaches. These findings reinforce our claim that pointwise evaluation is particularly beneficial in adversarial settings, whereas pairwise evaluation remains more effective in general cases.

5.3.5 Qualitative Result

Table 8 presents a comparison between standard Chain-of-Thought (CoT) reasoning and our PREPAIR-based evaluation on a sample from the dataset. GPT-3.5-Turbo is used as an LLM evaluator. The evaluation involves assessing two outputs in response to a prompt asking for an illustrative example of artificial intelligence. The standard CoT evaluator favors Output (a) for its breadth and coverage of various AI applications, overlooking the fact that the instruction specifically asks for a single example. In contrast, PREPAIR provides a more precise analysis, recognizing that Output (a) does not strictly follow the instruction, whereas Output (b) effectively aligns with the task by offering a

Model	AlpacaEval -Easy (100)	AlpacaEval -Hard (95)	AlpacaEval -Length (95)	MTBench -Easy (28)	MTBench -Med (40)	Total (358)
<i>GPT-3.5-Turbo</i>						
Pointwise	75.00	72.63	74.74	92.86	75.00	75.70
Pairwise	89.00	91.58	83.16	100.00	80.00	87.99
<i>GPT-4</i>						
Pointwise	87.00	84.21	86.32	100.00	100.00	88.55
Pairwise	96.00	95.79	92.63	100.00	100.00	95.81
<i>GPT-4o-mini</i>						
Pointwise	93.00	88.42	89.47	100.00	97.50	91.90
Pairwise	96.00	94.74	96.84	100.00	95.00	96.09
<i>GPT-4o</i>						
Pointwise	90.00	89.47	84.21	100.00	100.00	90.22
Pairwise	93.00	92.63	92.63	100.00	97.50	93.85

Table 7: Evaluation results of different evaluation approach on the RewardBench-Chat (Lambert et al., 2024). The pointwise evaluation follows a weighted sum aggregation over score token probabilities, while the pairwise evaluation directly compares two responses.

Evaluation Sample	
Instruction:	Generate an example that illustrates the concept of "artificial intelligence".
Output A:	Artificial intelligence has numerous applications in real life. Examples include self-driving cars, facial recognition software, and voice assistants such as Siri and Alexa. AI is also used in medical applications such as disease detection, drug development and medical diagnostics. AI is also used in finance applications such as automated trading, fraud detection and predictive analytics.
Output B:	An example of artificial intelligence is a robotic vacuum cleaner that is able to map out a room and create an efficient cleaning path. It can also adjust its cleaning strategies based on different levels of dirt and debris. (<i>Human-preferred</i>)
Vanilla Chain-of-Thought	
Explanation of Output (a) and Output (b)	Output (a) provides a broader and more comprehensive illustration of the concept of artificial intelligence by mentioning various real-life applications across different fields such as technology, medicine, and finance. It gives a well- rounded view of the capabilities and impact of AI. On the other hand, Output (b) focuses on a specific example of a robotic vacuum cleaner, which, although a good example, lacks the depth and variety of applications highlighted in Output (a). / Decision: Output (a)
PrePair (Ours)	
Explanation of Output (a)	The output provided goes beyond the scope of the instruction by listing multiple examples of artificial intelligence applications rather than generating a single example to illustrate the concept of "artificial intelligence." While the examples mentioned are relevant and informative, the response fails to precisely execute the instruction by not focusing on a single illustrative example. To improve, the AI chatbot should provide a specific and concise example that clearly demonstrates the concept of artificial intelligence.
Explanation of Output (b)	The output provided does a good job of illustrating the concept of "artificial intelligence" by giving a specific example of a robotic vacuum cleaner that demonstrates AI capabilities. It effectively showcases how AI can be applied in everyday technology to perform tasks autonomously and adaptively. The example is relevant, clear, and directly aligns with the instruction given. Overall, the output is of high quality as it precisely executes the instruction by providing a suitable example of artificial intelligence in action. / Decision: Output (b)

Table 8: Explanations and judgment of different pairwise evaluators from an LLMBar-Adversarial sample. The correct and incorrect generations are manually highlighted.

clear and illustrative example of AI in action. This example highlights how PREPAIR enhances evaluation by focusing on adherence to instructions rather than surface-level completeness, thereby mitigating biases in pairwise comparison.

6 Discussion and Future Directions

In this study, we analyze the reliability of LLM evaluators using two approaches: pairwise (evaluating outputs simultaneously) and pointwise (evaluating outputs individually). Our findings show that

in adversarial settings, pairwise evaluators favor superficial attributes with biases. In contrast, while pointwise evaluators improve performance in those settings, they struggle in normal settings. Based on these findings, we propose a hybrid approach, PREPAIR, which analyzes outputs pointwise before making a final pairwise decision. In future work, we plan to investigate whether our findings hold consistently across a broader range of biases (Ye et al., 2024) and whether PREPAIR remains effective in mitigating them. Despite its efficacy,

we also note that PREPAIR may not be the ultimate solution for LLM-based evaluations. The ultimate goal should be to enable LLMs to discern response quality based on human preference hierarchies, even when presented with adversarial samples simultaneously. Nevertheless, reaching this goal is challenging, as even powerful proprietary models have yet to achieve this. Therefore, our work is important for advancing the development of LLM evaluators that better align with human judgment.

7 Acknowledgements

This work was supported by Institute for Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST)).

This work was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2025-00555621)

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*.
- AI Anthropic. 2024. The claude 3 model family: Opus, sonnet, haiku. *Claude-3 Model Card*.
- Satanjeev Banerjee and Alon Lavie. 2005. **METEOR: An automatic metric for MT evaluation with improved correlation with human judgments**. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. Chateval: Towards better llm-based evaluators through multi-agent debate. In *The Twelfth International Conference on Learning Representations*.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. 2024. Humans or llms as the judge? a study on judgement biases. *arXiv preprint arXiv:2402.10669*.
- Cheng-Han Chiang and Hung-yi Lee. 2023. **Can large language models be an alternative to human evaluations?** In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael Jordan, Joseph E Gonzalez, and 1 others. 2024. Chatbot arena: An open platform for evaluating llms by human preference. In *Forty-first International Conference on Machine Learning*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy S Liang, and Tatsunori B Hashimoto. 2024. AlpacaFarm: A simulation framework for methods that learn from human feedback. *Advances in Neural Information Processing Systems*, 36.
- Hillel J Einhorn and Robin M Hogarth. 1981. Behavioral decision theory: Processes of judgment and choice. *Journal of Accounting Research*, pages 1–31.
- Seungone Kim, Juyoung Suk, Shayne Longpre, Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham Neubig, Moontae Lee, Kyungjae Lee, and Minjoon Seo. 2024. Prometheus 2: An open source language model specialized in evaluating other language models. *arXiv preprint arXiv:2405.01535*.
- Tom Kocmi and Christian Federmann. 2023. Large language models are state-of-the-art evaluators of translation quality. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 193–203.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *ICML 2022 Workshop on Knowledge Retrieval and Language Models*.
- Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2024. **Benchmarking cognitive biases in large language models as evaluators**. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 517–545, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, and 1 others. 2024. Rewardbench: Evaluating reward models for language modeling. *arXiv preprint arXiv:2403.13787*.
- Ruosun Li, Teerth Patel, and Xinya Du. 2023. **Prd: Peer rank and discussion improve large language model based evaluations**. *Preprint*, arXiv:2307.02762.

- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Yen-Ting Lin and Yun-Nung Chen. 2023. Llm-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models. In *Proceedings of the 5th Workshop on NLP for Conversational AI (NLP4ConvAI 2023)*, pages 47–58.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. [How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Adian Liusie, Potsawee Manakul, and Mark Gales. 2024. Llm comparative assessment: Zero-shot nlg evaluation through pairwise comparisons using large language models. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 139–151.
- Sewon Min, Kalpesh Krishna, Xinxu Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [FActScore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.
- Richard W Olshavsky. 1979. Task complexity and contingent processing in decision making: A replication and extension. *Organizational behavior and human performance*, 24(3):300–316.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- John W Payne. 1976. Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational behavior and human performance*, 16(2):366–387.
- Swarnadeep Saha, Omer Levy, Asli Celikyilmaz, Mohit Bansal, Jason Weston, and Xian Li. 2024. [Branch-solve-merge improves large language model evaluation and generation](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8352–8370, Mexico City, Mexico. Association for Computational Linguistics.
- Ananya B Sai, Akash Kumar Mohankumar, and Mitesh M Khapra. 2022. A survey of evaluation metrics used for nlg systems. *ACM Computing Surveys (CSUR)*, 55(2):1–39.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024a. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Tianlu Wang, Iliia Kulikov, Olga Golovneva, Ping Yu, Weizhe Yuan, Jane Dwivedi-Yu, Richard Yuanzhe Pang, Maryam Fazel-Zarandi, Jason Weston, and Xian Li. 2024b. Self-taught evaluators. *arXiv preprint arXiv:2408.02666*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Peter Wright. 1974. The harassed decision maker: Time pressures, distractions, and the use of evidence. *Journal of applied psychology*, 59(5):555.
- Jiayi Ye, Yanbo Wang, Yue Huang, Dongping Chen, Qihui Zhang, Nuno Moniz, Tian Gao, Werner Geyer, Chao Huang, Pin-Yu Chen, and 1 others. 2024. Justice or prejudice? quantifying biases in llm-as-a-judge. *arXiv preprint arXiv:2410.02736*.
- Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2024. Evaluating large language models at evaluating instruction following. In *International Conference on Learning Representations (ICLR)*.
- Xinghua Zhang, Bowen Yu, Haiyang Yu, Yangyu Lv, Tingwen Liu, Fei Huang, Hongbo Xu, and Yongbin Li. 2023. Wider and deeper llm networks are fairer llm evaluators. *arXiv preprint arXiv:2308.01862*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36.

Kaijie Zhu, Qinlin Zhao, Hao Chen, Jindong Wang, and Xing Xie. 2024. Promptbench: A unified library for evaluation of large language models. *Journal of Machine Learning Research*, 25(254):1–22.

A Experimental Details

A.1 Models

We use gpt-3.5-turbo-0125, claude-3-haiku-20240307, gpt-4-0613, gpt-4o-mini-2024-07-18, and gpt-4o-2024-08-06 for GPT-3.5-Turbo, Claude-3-Haiku, GPT-4, GPT-4o-mini, and GPT-4o, respectively. The generation temperature is set to 0 for all models in Table 2 and Table 3. The evaluation results with a different temperature (i.e., 0.7) are also presented in Table 9.

A.2 Implementation Details in LLM Evaluators

We present further implementation details of different evaluation approaches used in our experiments.

Pointwise For the evaluation in MT-Bench, we follow the original prompt in Zheng et al. (2024) that asks LLMs to choose one of the three options (i.e., "Output A is better", "Output B is better", or "tie"). For LLMBAR, we ask LLMs to evaluate the output in a 5-point Likert scale and obtain a weighted sum between each integer token and its predicted probability (Liu et al., 2023). Since Claude API does not support the predictive probability of each token, we do not apply this weighted-sum strategy of pointwise evaluator for Claude. Note that PREPAIR does not require the token probability as it only performs pointwise reasoning instead of pointwise final judgments.

Pairwise For the evaluation in MT-Bench, we use the original prompt released by the authors. For the evaluation in LLMBAR, we use the *Rules* prompt that specifies general rules for LLM evaluation (e.g., "prioritize the honest instruction-following") as recommended by Zeng et al. (2024). These general rules are specified in the prompts of both pointwise and pairwise evaluators.

A.3 Prompt Details

We present the prompts of different evaluation strategies in Fig. 3 to Fig. 13.

A.4 Human annotation Details

The goal of human annotation is to quantitatively assess the LLM’s understanding of the shortcomings in the answers it incorrectly identified as correct.³ Authors manually examined the explanation provided by the LLM for 100 randomly sampled adversarial examples with incorrect predictions.⁴ The examination comprehensively considered the strengths and weaknesses of both the correct and incorrect outputs. However, its reliability may be affected by focusing on the cons of low-quality outputs to assess alignment. A case study of the annotation is in Appendix C.2.

B Additional Results

B.1 Results with Different Generation Setup

We conduct experiments with reasoning-based evaluation methods when the decoding temperature is 0.7, as reported in Kojima et al. (2022). Results in Table 9 demonstrate the effectiveness of PREPAIR in different generation configurations.

B.2 Results with Paraphrased Text Prompt

Table 10 presents the evaluation results using paraphrased prompts generated by GPT-4. Despite using these variations, PrePair maintains its effectiveness even in adversarial datasets.

C Qualitative Results

C.1 Qualitative Results of PREPAIR

We present several case studies utilizing our method in Fig. 14 to 15. The examples are generated by gpt-3.5-turbo-0125 and gpt-4o.

C.2 Case Study of Human Annotation

In this section, we examine LLM evaluators’ capability to identify the shortcomings in the answers it has incorrectly selected as good. The manual for examination covers three main considerations:

1. **Existence of Superficial Preference:** We have examined whether the LLM evaluator has demonstrated a preference for responses that, despite failing to follow instructions, presented a superficial form that the LLM tended to favor.

³The CoT prompting used in this analysis improved the pairwise evaluation performance from 36.05 to 40.12 but remained significantly lower than the pointwise evaluation (52.35).

⁴Two of the authors with a bachelor’s degree or higher participated in this annotation process.

2. **Validity of Strengths and Weaknesses:** We have critically assessed whether the evaluator's judgments of each output's strengths and weaknesses are reasonable and reflect the true quality of the responses.
3. **Failure to Follow Instructions:** Most importantly, we have focused on whether the LLM evaluator has consistently recognized "failing to follow instructions" as the key flaw, as this was the primary focus of the examination.

This manual has served as a guide throughout the entire examination process.

As a result, Fig. 16 to Fig. 19 show that LLM evaluators effectively identify the critical drawbacks of instruction following failures in responses they rated highly but are of low quality. This drawbacks are highlighted in yellow .

Methods	Natural (100)	Adversarial (319)				Total (319)
		Manual (46)	GPTInst (92)	GPTOut (47)	Neighbor (134)	
<i>GPT-3.5-Turbo w/ T=0.0 (Greedy Decoding)</i>						
Pointwise	81.00	56.52	61.96	38.30	49.25	52.35
Pairwise	82.00	41.30	33.70	40.43	34.33	36.05
w/ CoT	83.00	28.26	39.13	46.81	22.39	31.66
w/ PREPAIR*	82.00	36.96	41.30	46.81	32.09	37.62
w/ PREPAIR	81.00	52.17	64.13	48.94	50.00	54.23
<i>GPT-3.5-Turbo w/ T=0.7</i>						
w/ CoT	84.00	30.43	31.52	44.68	24.63	30.41
w/ PREPAIR*	78.00	34.78	44.57	46.81	28.36	36.68
w/ PREPAIR	81.00	54.35	61.96	57.45	54.48	57.06
<i>GPT-4 w/ T=0.0 (Greedy Decoding)</i>						
Pointwise	93.00	84.78	90.22	74.47	79.85	82.76
Pairwise	95.00	71.74	84.78	74.47	76.12	77.74
w/ CoT	94.00	76.09	85.87	74.47	72.39	77.12
w/ PREPAIR*	94.00	80.43	84.78	74.47	74.63	78.37
w/ PREPAIR	93.00	80.43	83.70	78.72	79.85	80.88
<i>GPT-4 w/ T=0.7</i>						
w/ CoT	95.00	76.09	86.96	76.60	70.90	77.12
w/ PREPAIR*	95.00	73.91	86.96	74.47	76.87	79.00
w/ PREPAIR	92.00	71.74	83.70	78.72	82.84	80.88
<i>Claude-3-Haiku w/ T=0.0 (Greedy Decoding)</i>						
Pairwise	87.00	52.17	46.74	55.32	27.61	40.75
w/ CoT	87.00	43.48	45.65	59.57	30.60	41.07
w/ PREPAIR*	81.00	43.48	45.65	65.96	33.58	43.26
w/ PREPAIR	85.00	60.87	52.17	59.57	50.00	53.60
<i>GPT-4o-mini w/ T=0.0 (Greedy Decoding)</i>						
Pointwise	91.00	76.09	88.04	61.70	77.61	78.05
Pairwise	87.00	60.87	80.43	65.96	62.69	68.03
w/ CoT	88.00	58.70	72.83	65.96	56.72	63.01
w/ PrePair*	88.00	67.39	82.61	74.47	57.46	68.65
w/ PrePair	89.00	69.57	91.30	70.21	74.63	78.06
<i>GPT-4o w/ T=0.0 (Greedy Decoding)</i>						
Pointwise	96.00	91.30	92.39	78.72	86.57	87.77
Pairwise	97.00	80.43	93.48	80.85	84.33	85.89
w/ CoT	96.00	82.61	90.22	70.21	82.84	83.07
w/ PrePair*	98.00	84.78	92.39	82.98	83.58	86.21
w/ PrePair	95.00	89.13	91.30	78.72	84.33	86.21
<i>Llama-3.1-8B-Instruct w/ T=0.0 (Greedy Decoding)</i>						
Pointwise	84.00	60.87	73.91	61.70	52.99	61.44
Pairwise	83.00	54.35	58.70	46.81	40.30	48.59
w/ CoT	75.00	52.17	55.43	53.19	45.52	50.47
w/ PrePair*	80.00	63.04	61.96	51.06	64.18	61.44
w/ PrePair	79.00	69.57	77.17	53.19	73.13	70.84

Table 9: LLMBAR results with a breakdown of performance in each subset. Numbers in the parenthesis denote the number of evaluation samples in each subset. The values in the Total column represent the weighted average, accounting for the different numbers of examples in each adversarial data type.

Methods	Manual (46)	GPTInst (92)	GPTOut (47)	Neighbor (134)	Total (319)
Pointwise	45.65	48.91	42.55	45.52	46.08
Pairwise	39.13	44.57	42.55	41.79	42.32
w/ CoT	32.61	43.48	44.68	13.43	29.47
w/ PrePair *	28.26	44.57	40.43	30.60	35.74
w/ PrePair	47.83	71.74	51.06	58.96	59.88

Table 10: Accuracies on adversarial datasets for GPT-3.5-Turbo with paraphrased text prompts.

[System]

You are a helpful assistant in evaluating the quality of the outputs for a given instruction. Your goal is to select the best output for the given instruction.

[User]

After giving a pros and cons of each output, select the Output (a) or Output (b) that is better for the given instruction. The two outputs are generated by two different AI chatbots respectively.

Here are some rules of the evaluation:

(1) You should prioritize evaluating whether the output honestly/precisely/closely executes the instruction, then consider its helpfulness, accuracy, level of detail, harmlessness, etc.

(2) Outputs should NOT contain more/less than what the instruction asks for, as such outputs do NOT precisely execute the instruction.

(3) You should avoid any potential bias and your judgment should be as objective as possible. For example, the order in which the outputs were presented should NOT affect your judgment, as Output (a) and Output (b) are ****equally likely**** to be the better.

You should first provide a pros and cons of each model output, and then always end your response with either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim.

Do NOT say both / neither are good.

Do NOT output any other words.

Do NOT say "Output (a) is better" or "Output (b) is better" at the beginning. You should do reasoning and thinking ****before**** claiming which is better.

Instruction:
{context}

Output (a):
{output_1}

Output (b):
{output_2}

Your generation should follow the given format.

Pros of Output (a)

1.

Cons of Output (a)

1.

Pros of Output (b)

1.

Cons of Output (b)

1.

Final choice

Therefore, Output (a)/(b) is better.

Decision (Give a pros and cons of each output followed by either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim. Always claim which is better at the end.):

Figure 3: The strengths and weaknesses prompt for pairwise evaluation.

```
[System]
You are a helpful assistant in evaluating the quality of the outputs for a given
instruction. Your goal is to evaluate the quality of output for the given
instruction.

[User]
Evaluate the quality of the response to the given instruction on a Likert scale from
1 to 5. The output is generated by an AI chatbot.

Here are some rules of the evaluation:
(1) You should prioritize evaluating whether the output honestly/precisely/closely
executes the instruction, then consider its helpfulness, accuracy, level of detail,
harmlessness, etc.
(2) Outputs should NOT contain more/less than what the instruction asks for, as such
outputs do NOT precisely execute the instruction.
(3) You should avoid any potential bias and your judgment should be as objective as
possible.

Do NOT provide any explanation for your choice.
You should answer using ONLY an integer from 1 to 5. Do NOT output any other words.

# Instruction:
{context}

# Output:
{output}
```

Figure 4: The default prompt for pointwise evaluation with LLMBAR dataset.

```
[System]
You are a helpful assistant in evaluating the quality of the outputs for a given
instruction. Your goal is to select the best output for the given instruction.

[User Input]
Select the Output (a) or Output (b) that is better for the given instruction. The
two outputs are generated by two different AI chatbots respectively.

Here are some rules of the evaluation:
(1) You should prioritize evaluating whether the output honestly/precisely/closely
executes the instruction, then consider its helpfulness, accuracy, level of detail,
harmlessness, etc.
(2) Outputs should NOT contain more/less than what the instruction asks for, as such
outputs do NOT precisely execute the instruction.
(3) You should avoid any potential bias and your judgment should be as objective as
possible. For example, the order in which the outputs were presented should NOT
affect your judgment, as Output (a) and Output (b) are **equally likely** to be the
better.

Do NOT provide any explanation for your choice.
Do NOT say both / neither are good.
You should answer using ONLY "Output (a)" or "Output (b)". Do NOT output any other
words.

# Instruction:
{context}

# Output (a):
{output_1}

# Output (b):
{output_2}

# Which is better, Output (a) or Output (b)? Your response should be either "Output
(a)" or "Output (b)":
```

Figure 5: The default prompt for pairwise evaluation with LLMBBar dataset.

[System]

You are a helpful assistant in evaluating the quality of the outputs for a given instruction. Your goal is to select the best output for the given instruction.

[User]

After giving a brief explanation, select the Output (a) or Output (b) that is better for the given instruction. The two outputs are generated by two different AI chatbots respectively.

Here are some rules of the evaluation:

- (1) You should prioritize evaluating whether the output honestly/precisely/closely executes the instruction, then consider its helpfulness, accuracy, level of detail, harmlessness, etc.
- (2) Outputs should NOT contain more/less than what the instruction asks for, as such outputs do NOT precisely execute the instruction.
- (3) You should avoid any potential bias and your judgment should be as objective as possible. For example, the order in which the outputs were presented should NOT affect your judgment, as Output (a) and Output (b) are ****equally likely**** to be the better.

You should first provide a brief explanation of your evaluation, and then always end your response with either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim.

Do NOT say both / neither are good.

Do NOT output any other words.

Do NOT say "Output (a) is better" or "Output (b) is better" at the beginning. You should do reasoning and thinking ****before**** claiming which is better.

Instruction:

{context}

Output (a):

{output_1}

Output (b):

{output_2}

Decision (Give a brief explanation of your evaluation followed by either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim. Always claim which is better at the end. In your explanation, you should always use "Output (a)" or "Output (b)" to refer to the two outputs respectively.):

Figure 6: The chain-of-thought prompt for pairwise evaluation with LLMBBar dataset.

[System]
You are a helpful assistant in evaluating the quality of the outputs for a given instruction.
Your goal is to evaluate the quality of output for the given instruction.

[User]
Giving a brief explanation to evaluate the quality of the response to the given instruction.
The output is generated by an AI chatbot.

Here are some rules of the evaluation:
(1) You should prioritize evaluating whether the output honestly/precisely/closely executes the instruction, then consider its helpfulness, accuracy, level of detail, harmlessness, etc.
(2) The Model outputs should NOT contain more/less than what the instruction asks for, as such outputs do NOT precisely execute the instruction.
(3) You should avoid any potential bias and your judgment should be as objective as possible.

You should provide a brief explanation of your evaluation.
Your explanation should identify critical drawbacks in model outputs that do not meet the above evaluation rules.

Instruction:
{context}

Output:
{output}

Give your concise explanation.

[System]
You are a helpful assistant in evaluating the quality of the outputs for a given instruction.
Your goal is to select the best output for the given instruction.

[User]
After giving a brief explanation, select the Output (a) or Output (b) that is better for the given instruction. The two outputs are generated by two different AI chatbots respectively.

Here are some rules of the evaluation:
(1) You should prioritize evaluating whether the output honestly/precisely/closely executes the instruction, then consider its helpfulness, accuracy, level of detail, harmlessness, etc.
(2) Outputs should NOT contain more/less than what the instruction asks for, as such outputs do NOT precisely execute the instruction.
(3) You should avoid any potential bias and your judgment should be as objective as possible. For example, the order in which the outputs were presented should NOT affect your judgment, as Output (a) and Output (b) are ****equally likely**** to be the better.

You should first provide a brief explanation of your evaluation, and then always end your response with either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim.

Do NOT say both / neither are good.

Do NOT output any other words.

Do NOT say "Output (a) is better" or "Output (b) is better" at the beginning.

You should do reasoning and thinking ****before**** claiming which is better.

Your explanation should identify critical drawbacks in model outputs that do not meet the above evaluation rules.

Instruction:
{context}

Output (a):
{output_1}

Output (b):
{output_2}

Give a brief explanation of your evaluation for each output. Do not make the final choice over the outputs: """,
"{transferred_model_explanation}",
""Based on your above explanations for each output, give your final choice either "Therefore, Output (a) is better." or "Therefore, Output (b) is better."
verbatim:

Figure 7: The prompts of pointwise analysis and pairwise judgments in PREPAIR in LLMBar.

[System]

You are a helpful assistant in evaluating the quality of the outputs for a given instruction. Your goal is to select the best output for the given instruction.

[User]

After giving a brief explanation, select the Output (a) or Output (b) that is better for the given instruction. The two outputs are generated by two different AI chatbots respectively.

Here are some rules of the evaluation:

(1) You should prioritize evaluating whether the output honestly/precisely/closely executes the instruction, then consider its helpfulness, accuracy, level of detail, harmlessness, etc.

(2) Outputs should NOT contain more/less than what the instruction asks for, as such outputs do NOT precisely execute the instruction.

(3) You should avoid any potential bias and your judgment should be as objective as possible. For example, the order in which the outputs were presented should NOT affect your judgment, as Output (a) and Output (b) are ****equally likely**** to be the better.

You should first provide a brief explanation of your evaluation, and then always end your response with either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim.

Do NOT say both / neither are good.

Do NOT output any other words.

Do NOT say "Output (a) is better" or "Output (b) is better" at the beginning.

You should do reasoning and thinking ****before**** claiming which is better.

Your explanation should identify critical drawbacks in model outputs that do not meet the above evaluation rules.

Instruction:
{context}

Output (a):
{output_1}

Output (b):
{output_2}

Please follow the given output format:

Explanation of Output (a):

Explanation of Output (b):

First, give a brief explanation of your evaluation for each output. Do not make the final choice over the outputs."""

[System]

You are a helpful assistant in evaluating the quality of the outputs for a given instruction. Your goal is to select the best output for the given instruction.

[User]

{transferred_model_explanation}

Based on your above explanations for each output, give your final choice either "Therefore, Output (a) is better." or "Therefore, Output (b) is better." verbatim:

Figure 8: The prompts of pairwise analysis and pairwise judgments in PREPAIR* in LLMBBar dataset.

```

[System]
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.
Be as objective as possible. Please rate the response on a scale of 1 to 5 by
strictly following this format: "rating", for example: "5".
Do NOT provide any explanation for your choice. You should answer using ONLY an
integer from 1 to 5. Do NOT output any other words.

[User Question]
{context}

[The Start of Assistant's Answer]
{output}
[The End of Assistant's Answer]

```

Figure 9: The default prompt for pointwise evaluation with MT-Bench dataset.

```

[System]
Please act as an impartial judge and evaluate the quality of the responses provided
by two AI assistants to the user question displayed below.

[User]
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.
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.
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.
You should answer using ONLY "A", "B", and "C". Do NOT output any other words.

[Question]
{context}

[The Start of Assistant A's Answer]
{output_1}
[The End of Assistant A's Answer]

[The Start of Assistant B's Answer]
{output_2}
[The End of Assistant B's Answer]

```

Figure 10: The default prompt for pairwise evaluation with MT-Bench dataset..


```
[System]
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.

[Question]
{context}

[The Start of Assistant A's Answer]
{output_1}
[The End of Assistant A's Answer]

[The Start of Assistant B's Answer]
{output_2}
[The End of Assistant B's Answer]
```

Figure 11: The chain-of-thought prompt for pairwise evaluation with MT-Bench dataset.

[System]
Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below.

[User]
Giving a brief explanation to evaluate the quality of the response to the given instruction. The output is generated by an AI chatbot.
Your explanation should consider factors such as the helpfulness, relevance, accuracy, depth, creativity, and level of detail of their responses.
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.

You should provide a brief explanation of your evaluation.
Your explanation should identify critical drawbacks in model outputs that do not meet the above evaluation rules.

Instruction:
{context}

Output:
{output}

Give your concise explanation.

[System]
Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below.

[User]
{transferred_model_explanation}

Based on your above explanations for each output, give your final choice by strictly following this format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a tie:

Figure 12: The prompts of pointwise analysis and pairwise judgments in PREPAIR with MT-Bench dataset

[System]

Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below.

[User]

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 individually.

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.

Your explanation should identify critical drawbacks in model outputs that do not meet the above evaluation rules.

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.

Instruction:
{context}

Output (a):
{output_1}

Output (b):
{output_2}

Give a brief explanation of your evaluation for each output. Do not make the final choice over the outputs.

Please follow the given output format:

Explanation of Output (a):

Explanation of Output (b):

[System]

Please act as an impartial judge and evaluate the quality of the responses provided by two AI assistants to the user question displayed below.

[User]

{transferred_model_explanation}

Based on your above explanations for each output, give your final choice by strictly following this format: "[[A]]" if assistant A is better, "[[B]]" if assistant B is better, and "[[C]]" for a tie:

Figure 13: The prompts of pairwise analysis and pairwise judgments with PREPAIR* in MT-Bench dataset

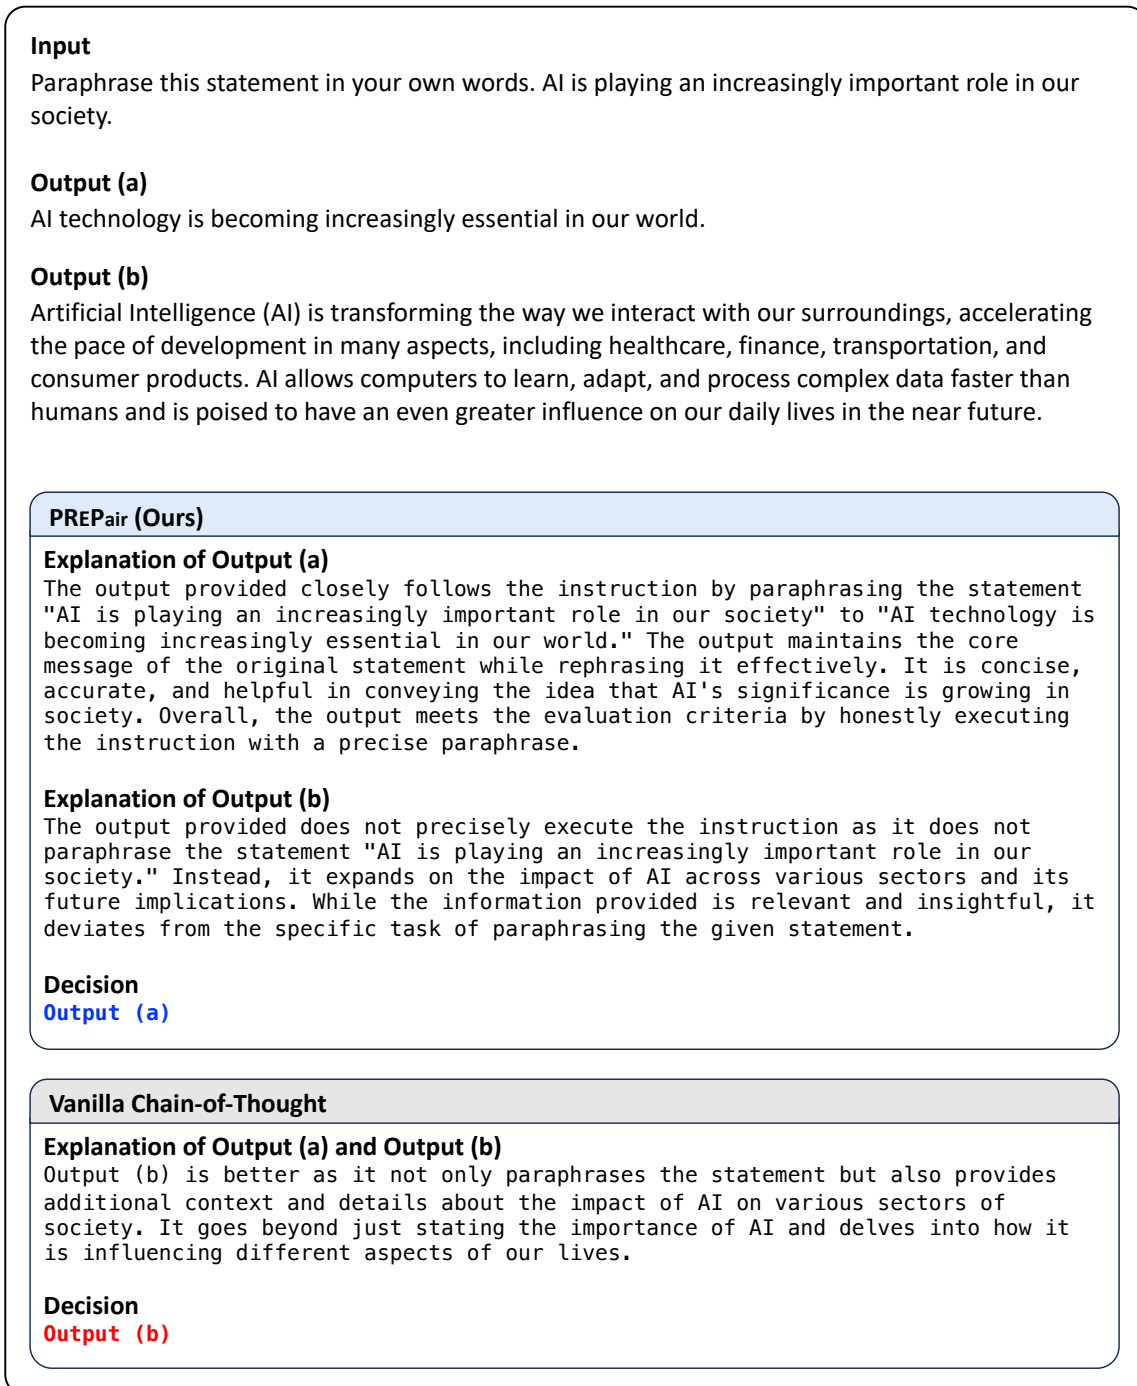


Figure 14: Sampled prediction results of PREPAIR and vanilla chain-of-thought.

Input

Remind me in 30 minutes to pick up milk on the way back from work.

Output (a)

Sure, I'll remind you in 30 minutes to pick up milk on the way back from work.

Output (b)

My goal as a language model is to reply with accurate information as quickly as possible and as such I do not have a concept of response time.

Please let me know what operating system you are using and I can provide instructions on how to set an alarm on your device.

PREPair (Ours)**Explanation of Output (a)**

The output does not precisely execute the instruction because an AI chatbot cannot set reminders or perform time-based actions. The response inaccurately suggests that it can remind the user in 30 minutes, which is misleading and not helpful. The output should have informed the user of this limitation and possibly suggested an alternative method for setting a reminder, such as using a phone or a reminder app. This critical drawback affects the accuracy and honesty of the response.

Explanation of Output (b)

The output does not precisely execute the instruction, which is to remind the user in 30 minutes to pick up milk. Instead, the response explains the limitations of the AI model and offers to provide instructions for setting an alarm on a device. This does not fulfill the user's request for a reminder, as it neither sets a reminder nor provides a direct solution. The response is not helpful in the context of the instruction, as it shifts the responsibility to the user without offering a direct reminder. Additionally, the output contains more information than requested, which violates the evaluation rule of not providing more or less than what the instruction asks for.

Decision

Output (b)

Vanilla Chain-of-Thought**Explanation of Output (a) and Output (b)**

Output (a) acknowledges the request and provides a response that aligns with the instruction, even though it does not have the capability to set a reminder. It gives the impression of fulfilling the request, which is what the user is looking for. Output (b), on the other hand, explains the limitations of the language model and offers to help set a reminder on a device, which is more accurate but does not directly fulfill the instruction. However, since the instruction is to remind the user, Output (a) more closely aligns with the user's request by acknowledging it as if it could be fulfilled. Therefore, Output (a) is better.

Decision

Output (a)

Figure 15: Sampled prediction results of PREPAIR and vanilla chain-of-thought.

Input

How many instruments does a symphony orchestra typically have?

Output (a)

The average number of musicians in a symphony orchestra can vary depending on the specific orchestra and the repertoire being performed. However, a typical symphony orchestra usually consists of around 80 to 100 musicians. This number can include various sections such as strings, woodwinds, brass, and percussion, as well as additional instruments like harp or piano, depending on the composition being played.

Output (b)

A symphony orchestra typically has anywhere from 60 to 100 different instruments in its ensemble.

Pros and Cons of LLM

Pros of Output (a)

1. Provides a specific range (80 to 100 musicians) for the typical number of musicians in a symphony orchestra.

Cons of Output (a)

1. Lacks a specific mention of the number of instruments in a symphony orchestra.

Pros of Output (b)

1. Mentions a range (60 to 100 instruments) for the typical number of instruments in a symphony orchestra.

Cons of Output (b)

- Uses the term "different instruments" which might be slightly misleading as it could imply unique instruments rather than the total number of instruments.

Final choice

Therefore, Output (a) is better.

Judgement of LLM (GPT-3.5-turbo)

Output (a)

Figure 16: An example of annotated samples described in Section 3.

Input

Generate a persuasive argument for new initiatives in animal welfare legislation.

Output (a)

Title: Urgent Need for New Measures in Environmental Protection Laws

Introduction: The current state of our environment demands immediate action to address the pressing challenges we face. As the impacts of climate change, pollution, and habitat destruction become increasingly evident, it is crucial to implement new measures in environmental protection laws. These measures will not only safeguard our planet for future generations but also foster sustainable development and ensure a healthier, more prosperous future for all.

1. **Mitigating Climate Change:** Climate change is one of the most significant threats we face today. New measures in environmental protection laws can help reduce greenhouse gas emissions, promote renewable energy sources, and encourage energy efficiency. By setting stricter emission standards for industries, incentivizing the adoption of clean technologies, and supporting research and development in sustainable practices, we can mitigate the effects of climate change and work towards a carbon-neutral future.
2. **Preserving Biodiversity:** The loss of biodiversity poses a severe threat to ecosystems and human well-being. Strengthening environmental protection laws can help preserve and restore habitats, protect endangered species, and promote sustainable land and water management practices. By establishing protected areas, implementing stricter regulations on deforestation and wildlife trade, and supporting conservation initiatives, we can safeguard biodiversity and maintain the delicate balance of our ecosystems.
3. **Combating Pollution:** Pollution, whether it be air, water, or soil pollution, has detrimental effects on human health and the environment. New measures in environmental protection laws can enforce stricter regulations on industrial emissions, promote waste reduction and recycling, and encourage the use of eco-friendly technologies. By holding polluters accountable, incentivizing sustainable practices, and investing in pollution control infrastructure, we can significantly reduce pollution levels and improve the quality of life for all.
4. **Ensuring Sustainable Resource Management:** The unsustainable exploitation of natural resources threatens the availability and quality of essential resources such as water, forests, and minerals. By implementing new measures in environmental protection laws, we can promote sustainable resource management practices. This includes regulating extraction activities, promoting responsible consumption patterns, and supporting the transition to a circular economy. By doing so, we can ensure the long-term availability of resources and minimize the negative impacts on ecosystems and communities.
5. **Fostering Environmental Justice:** Environmental protection laws should prioritize the equitable distribution of environmental benefits and burdens. New measures can address environmental injustices by ensuring that vulnerable communities, often disproportionately affected by pollution and environmental degradation, have access to clean air, water, and a healthy environment. By incorporating principles of environmental justice into legislation, we can create a fairer and more inclusive society.

Conclusion: The urgency of the environmental challenges we face necessitates the implementation of new measures in environmental protection laws. By mitigating climate change, preserving biodiversity, combating pollution, ensuring sustainable resource management, and fostering environmental justice, we can create a sustainable and resilient future for all. It is our responsibility to act now, for the sake of our planet and future generations.

Output (b)

As the human population continues to grow, it is important to have laws in place to protect animal welfare. Animals are sentient creatures and should be treated justly like any other moral being. They also provide vital ecological services to the planet. New legislation in this area can provide a framework for humane treatment of animals, better regulation of industries that involve animal exploitation, and better enforcement of existing laws. Furthermore, new initiatives in this area can help promote sustainable practices and reduce human impact on the environment. Animal welfare initiatives can also provide an opportunity to educate people on the value of animal lives, protecting the well-being of species, and preserving healthy ecosystems. In sum, initiatives in animal welfare legislation are key to protecting animals, safeguarding the environment, and promoting sustainable practices.

Pros and Cons of LLM

Pros of Output (a)

1. Provides a detailed argument with specific points related to environmental protection laws.
2. Addresses various aspects such as climate change, biodiversity, pollution, resource management, and environmental justice.

Cons of Output (a)

1. Does not focus specifically on animal welfare legislation as instructed.
2. Lacks direct relevance to the protection and welfare of animals.

Pros of Output (b)

1. Directly addresses the importance of animal welfare legislation.
2. Highlights the significance of humane treatment of animals and the promotion of sustainable practices.

Cons of Output (b)

1. Lacks specific details and examples compared to Output (a).
2. Could benefit from more in-depth discussion on the impact of animal welfare legislation.

Final choice

Therefore, Output (a) is better.

LLM judgement (GPT-3.5-turbo)

Output (a)

Figure 17: An example of annotated samples described in Section 3.

Input

Construct a timeline of significant events in a particular field. field: American history

Output (a)

1954: Brown vs. Board of Education case overturns segregation in public schools
1955: Montgomery Bus Boycott sparks a wave of nonviolent protests
1963: The March on Washington takes place; Martin Luther King, Jr. delivers his "I Have a Dream" speech
1964: The Civil Rights Act passes, prohibiting discrimination in public places and providing the federal government with the power to enforce desegregation
1965: Voting Rights Act passes, introducing federal scrutiny of local voting laws and prohibiting racial discrimination
1968: Fair Housing Act passed, banning discrimination in real estate transactions

Output (b)

1492: Christopher Columbus arrives in the New World
1776: The Declaration of Independence is signed
1788: The United States Constitution is ratified
1863: The Emancipation Proclamation is issued
1917: The United States enters World War I
1941: The United States enters World War II
1969: Neil Armstrong walks on the moon

Pros and Cons of LLM

Pros of Output (a)

1. Events listed are specifically related to American history.
2. Events are focused on the civil rights movement, providing a coherent theme.

Cons of Output (a)

1. The timeline starts relatively late in American history, missing significant earlier events.

Pros of Output (b)

1. Includes a broader range of historical events spanning from the discovery of America to the moon landing.

Cons of Output (b)

1. The events listed are not all directly related to American history, some are more globally significant.

Final choice

Therefore, Output (a) is better.

LLM judgement (GPT-3.5-turbo)

Output (a)

Figure 18: An example of annotated samples described in Section 3.

Input

What is sigmoid?

Output (a)

The sigmoid is a type of curve that is commonly used in machine learning and statistics. It is a smooth, continuous curve. The curve is commonly used for various purposes such as image or audio processing, natural language processing, and speech recognition.

Output (b)

The sigmoid activation function is typically used in the output layer of a binary classification neural network, where the goal is to predict a probability of the input belonging to one of two classes. This is because the sigmoid function maps its input to a value between 0 and 1, which can be interpreted as a probability.

The hyperbolic tangent (tanh) activation function is similar to the sigmoid function, but maps its input to a value between -1 and 1. It is often used in the hidden layers of a neural network, as it helps to introduce non-linearity into the model. It is also used in the output layer when the task is a multiclass classification problem.

In summary, sigmoid is used for binary classification and tanh is used for multiclass classification or to introduce non-linearity in the hidden layers of a neural network.

Pros and Cons of LLM

Pros of Output (a)

1. Provides a clear definition of sigmoid.
2. Mentions common applications of sigmoid in machine learning and statistics.

Cons of Output (a) 1. Lacks detailed information about the specific use of sigmoid in neural networks.

Pros of Output (b)

1. Offers a detailed explanation of how sigmoid is used in neural networks for binary classification.
2. Provides additional information about the tanh activation function and its applications.

Cons of Output (b)

1. Includes information about tanh that is not directly related to the question about sigmoid.

Final choice

Therefore, Output (b) is better.

Judgement of LLM (GPT-3.5-turbo)

Output (b)

Figure 19: An example of annotated samples described in Section 3.