

Re-evaluating Automatic LLM System Ranking for Alignment with Human Preference

Mingqi Gao^{*1,2} Yixin Liu^{*2} Xinyu Hu¹
Xiaojun Wan¹ Jonathan Bragg³ Arman Cohan^{2,3}

¹Peking University ²Yale University ³Allen Institute for AI
{mingqi.gao,yixin.liu,arman.cohan}@yale.edu

Abstract

Evaluating and ranking the capabilities of different LLMs is crucial for understanding their performance and alignment with human preferences. Due to the high cost and time-consuming nature of human evaluations, an automatic LLM *bencher* (i.e., an automatic evaluation framework that aims to rank LLMs based on their alignment with human preferences) is indispensable. An automatic LLM *bencher* consists of four components: the input set (e.g., a user instruction), the evaluation model (e.g., an LLM), the evaluation type (e.g., pairwise comparison), and the aggregation method (e.g., the ELO rating system). However, previous work has not thoroughly explored how to select these components or how their different combinations influence the results. In this work, through controlled experiments, we provide a series of recommendations on how to choose each component to better automate the evaluation of LLMs. Furthermore, we discovered that when evaluating LLMs with similar performance, the performance of the automatic LLM *bencher* declines sharply, underscoring the limitations of current *benchers* and calling for future work. Lastly, we found that the evaluation models’ performance at the instance level (e.g., the accuracy of selecting the best output) does not always align with their effectiveness when used as a component of a *bencher*, highlighting the importance of dedicated system-level evaluation of *benchers*.

1 Introduction

Recently, various large language models (LLMs) (Ouyang et al., 2022) and LLM-based agents (Wang et al., 2024a) have been released, demonstrating strong capabilities across various tasks. These systems enable efficient interactions with human users and can be instructed to perform various complex activities, making their performance

in such tasks an important aspect of evaluation. To benchmark the capabilities of these systems, human judgments of output quality remain indispensable as the gold standard, as many tasks do not have standard answers and are inherently open-ended (Zheng et al., 2023; Dubois et al., 2023).

Chatbot Arena (Chiang et al., 2024) embodies this concept. It is a real-time evaluation platform aimed at a large user base, where users can freely provide input, select any two hosted LLMs to generate responses, and indicate which one they prefer. Chatbot Arena periodically derives a leaderboard of various LLMs by aggregating instance-level pairwise human evaluations. To date, Chatbot Arena has collected over 1.5 million human judgments, involving more than 100 systems. Due to its substantial size and the comprehensiveness of the systems included, its LLM ranking has been widely regarded as a trustworthy indicator of an LLM’s general capabilities (Li et al., 2023, 2024b; Lin et al., 2024; Zhao et al., 2024).

However, automatic LLM system rankers (i.e. *benchers*) are often needed because of the expensive and time-consuming nature of human evaluation.¹ As a result, several widely used automatic *benchers* have been proposed, such as Alpaca Eval (Li et al., 2023; Dubois et al., 2024) and Arena Hard (Li et al., 2024b), which use a strong LLM such as GPT-4 (OpenAI, 2023) to compare various systems on a carefully designed input set. These automatic *benchers* can produce LLM rankings that correlate well with Chatbot Arena, achieving Spearman’s ρ of around 0.95. Therefore, they have been widely used to develop and evaluate LLMs for tasks like alignment fine-tuning (Tunstall et al., 2024; Yang et al., 2024).

¹LLMs used to evaluate the output of other LLMs are also referred as “evaluators” or “LLM-as-a-judge” in the literature. We use the word “*bencher*” to explicitly emphasize the broader automatic evaluation frameworks used to derive rankings of models in a benchmark (in which an LLM evaluator is often used as part of the framework).

*Equal contribution.

Formally, an automatic LLM bench for evaluating instruction following capabilities consists of the following components:

- *Input Set*: A set of instructions that each LLM is tasked with to generate responses.
- *Evaluation Model*: An (often strong) LLM (e.g. GPT-4o) that serves as a proxy for humans in evaluating a model’s responses.
- *Evaluation Type*: The type of instance-level evaluation generated by the evaluation model, such as pointwise scores or pairwise comparisons, usually determined by the prompt.²
- *Aggregation Method*: The method by which instance-level evaluations are converted into system (or model) scores and rankings.

While these automatic benchers have shown a high level of alignment with human evaluations in Chatbot Arena, we note a lack of a systematic evaluation across all the above-mentioned components involved with the benchers. For example, while Arena Hard claims a higher alignment with human evaluation than Alpaca Eval and attributes its success to the careful selection of its input set, they also adopt different evaluation types for the evaluation: Arena Hard performs five-point pairwise comparisons, while Alpaca Eval produces binary comparisons. Therefore, the exact source of its superior performance remains unclear because of the lack of controlled comparisons.

Therefore, we aim to conduct a more rigorous examination of these automatic benchers. To this end, the first research question we explore is **RQ1: how to choose the appropriate components for building an effective automatic LLM bench?** Specifically, we perform controlled comparisons of various input sets, evaluation models, evaluation types, and aggregation methods to investigate which choices of each component maximize the benchers’ performance. Our key findings are:

(1) *Input Set*: Compared to the earlier released Alpaca Eval, using Arena Hard as the input set, which is filtered from crowdsourced user prompts through an LLM-powered pipeline, always yields higher correlations with the system rankings of Chatbot Arena.

(2) *Evaluation Model*: While proprietary models like GPT-4-turbo generally perform best as evaluation models, the open-source Llama-3.1-

²Following the literature, throughout this paper, “Instance-level” refers to an evaluation instance, which could consist of a pair of model outputs (in pairwise evaluation) or a single model output (in pointwise evaluation).

70B (Dubey et al., 2024) also delivers strong performance, particularly when paired with a suitable combination of other components in automatic benchers.

(3) *Evaluation Type*: With a fixed number of LLM queries, pointwise evaluation is slightly better than base pairwise evaluation³ for strong LLMs, but pairwise evaluation considerably outperforms pointwise with less capable LLMs. Notably, in almost all cases, the widely used 5-point pairwise evaluation performs *worse* than the base pairwise approach. Additionally, using reference systems is an efficient strategy.

(4) *Aggregation Method*: The system rankings produced by the Bradley-Terry model (BRADLEY and TERRY, 1952) and win ratio are mostly the same when the evaluation type is base pairwise. Interestingly, when the evaluation type is pointwise, the arithmetic mean, the most common aggregation method, is not always optimal.

Having provided a thorough evaluation of the automatic benchers, we take a step forward to explore a more challenging setting. A critical question regarding the evaluation of the LLM benchers is its reliability when the task is to evaluate similar-performing LLMs, which is an important use case. For example, these benchers are widely used to compare different model checkpoints fine-tuned from the same pre-trained model, which have similar performance and behavior. Finally, based on the above experimental results, it seems that we can easily build an effective automatic bench using GPT-4-turbo as the evaluation model, achieving a Spearman’s rho of over 0.90 with Chatbot Arena. However, **does the performance of the bench decline when systems with smaller performance differences are selected for evaluation (RQ2)?**

One aim of our above analyses was to identify suitable evaluation models (see Finding (2) above). Related work like LLMBAR (Zeng et al., 2024b) and RewardBench (Lambert et al., 2024) has taken a different approach to evaluating evaluation models, by measuring correlation or accuracy with ground truth instance-level human annotations rather than aggregate system rankings from Chatbot Arena. Compared to the crowd-sourced ChatBot Arena, human annotations collected for instance-level evaluations are usually of higher quality but lack the scale in terms of the number

³We refer to binary pairwise (win/lose) comparisons as the base pairwise evaluation method.

of systems used to generate evaluation examples. Sometimes system information is even unavailable for instance-level evaluation datasets. Given the significant efforts made in this direction, it would be desirable for instance-level evaluation results to also generalize to system-level evaluation, which leads to our **RQ3: can we use instance-level rankings of evaluation models as a good reference to select evaluation models for LLM benchers?**

Our investigation of the three research questions underscores several key findings for the users and developers of LLM benchers: (1) RQ1: It is important to choose the suitable components to build a benchner. For example, we find that using Llama-3.1-70B with a pairwise comparison protocol, the BT model for aggregation on Arena Hard can yield a comparable Spearman correlation of 0.93 with human annotators compared to the standard AlpacaEval2 evaluation pipeline where GPT-4 is used as the evaluator. (2) RQ2: Despite their promising performance in general, the benchers' alignment with humans degrades quickly when they are used to rank close-performing systems, calling for further work to enhance their robustness and accuracy. (3) RQ3: The evaluation results of LLM-based evaluators on the instance level do not always generalize to the system-level performance of the corresponding benchers. This highlights the need for dedicated evaluations at the system level.

2 Problem Formalization

In this section, we formalize some crucial concepts and our research questions.

2.1 Automatic LLM Benchner

Given a set of systems ($\mathcal{S} = \{s_1, s_2, \dots, s_n\}$) to be evaluated/ranked, an automatic benchner \mathcal{J}_A is composed of the following four components:

Input Set $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$. Each system s_i generates a response o_{ij} for each input x_j .

Evaluation Model \mathcal{E} : This is an LLM (e.g., GPT-4-turbo) that acts as a proxy for human evaluators. It generates instance-level evaluations for each system. The evaluation model \mathcal{E} outputs instance-level evaluation for each response o_{ij} .

Evaluation Type \mathcal{T} : This refers to the type of evaluation generated at the instance level. There are two main types⁴:

⁴There is also the "list-wise" approach, where a list of model outputs is provided, and the LLM is tasked with ranking them at once. However, this type of evaluation is not commonly employed in LLM-based judges or benchmarks.

(1) **Pointwise Evaluation**: The evaluation model assigns a score to a system's response. The pointwise score of the i -th system's response to the j -th input is denoted as $\mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{ij}) \in \mathbb{R}$. The total number of instance-level evaluations is mn .

(2) **Pairwise Evaluation**: the evaluation model compares the quality of two responses o_{ij}, o_{kj} generated by two systems s_i and s_k on the same input x_j . The base comparison result is binary, indicating whether s_i performed better than s_k or vice versa. This can be expressed as: $\mathcal{T}_{\text{base.pairwise}}(\mathcal{E}, x_j, o_{ij}, o_{kj}) \in \{o_{ij} > o_{kj}, o_{ij} < o_{kj}\}$. In 5-point pairwise comparisons, the evaluation model provides more detailed feedback about the difference between two systems' performances, capturing both the magnitude and direction of the preference. The 5 possible comparison outcomes for systems s_i and s_k on input x_j are: $\mathcal{T}_{5\text{-point.pairwise}}(\mathcal{E}, x_j, o_{ij}, o_{kj}) \in \{o_{ij} \ll o_{kj}, o_{ij} < o_{kj}, o_{ij} = o_{kj}, o_{ij} > o_{kj}, o_{ij} \gg o_{kj}\}$. The total number of instance-level pairwise evaluations is $mn(n-1)$ rather than $mn(n-1)/2$ because we always need to swap the presentation order of the responses s_i and s_j to address the position bias of LLM-based evaluations. In addition, due to the high computational cost of performing full-scale comparisons, some automatic LLM benchers (Li et al., 2023, 2024b) introduce a **reference system** in pairwise evaluation. In this approach, all systems under evaluation only need to be compared with a reference system (e.g. GPT-4), which reduces the complexity of pairwise evaluation to a level similar to that of pointwise evaluation.

Aggregation Method \mathcal{G} : This will create a mapping from systems to scores ($\mathcal{G}(s_i) \in \mathbb{R}$), depending on the evaluation type.

For pointwise evaluations, aggregation can be done via the mean or median of the instance-level scores across all inputs:

$$\mathcal{G}_{\text{mean}}(s_i) = \frac{1}{m} \sum_{j=1}^m \mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{ij})$$

$$\mathcal{G}_{\text{median}}(s_i) = \text{median}(\{\mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{ij})\}_{j=1}^m)$$

For pairwise evaluations, we consider the Bradley-Terry model (BRADLEY and TERRY, 1952) and win ratio⁵. They can only be applied to base pairwise instance-level evaluations, so other

⁵We do not consider the online ELO rating system (Elo, 1978) because the rankings it produces are influenced by the input order of instance-level evaluations.

types of instance-level evaluations need to be converted to the base pairwise type. The conversion rules are described in Appendix A. The Bradley-Terry model \mathcal{G}_{BT} computes the strength of each system via maximum likelihood estimate, and the win ratio is the proportion of pairwise wins for a system over others:

$$\mathcal{G}_{\text{win_ratio}}(s_i) = \frac{\sum_{j=1}^m \sum_{k \neq i} \mathbb{1}(\mathcal{T}_{\text{base,pairwise}}(\mathcal{E}, x_j, o_{ij}, o_{kj}) = o_{ij} > o_{kj})}{\sum_{j=1}^m \sum_{k \neq i} 1}$$

where m is the total number of inputs and the denominator counts the total number of pairwise comparisons involving s_i . Since pointwise data can be converted into pairwise data, all aggregation methods for pairwise evaluations can also be applied to pointwise evaluations after conversion.

2.2 How to Evaluate a Bencher

Ground Truth. From Chatbot Arena, we obtain system ratings, denoted as $r(s_i)$, for each system $s_i \in S$, where $r(s_i) \in \mathbb{R}$ represents the overall performance of the system as derived from human judgments. The ranking of the systems is obtained by sorting them based on their ratings: $R_H : S \rightarrow \{1, 2, \dots, n\}$, where $R_H(s_i)$ represents the rank of system s_i based on the ratings $r(s_i)$.

Evaluation Measure. The performance of an automatic bencher \mathcal{J}_A is evaluated by comparing its ranking R_A of systems with the ground truth ranking R_H , using Spearman’s rank correlation coefficient ($\rho(R_A, R_H)$) and Kendall’s tau ($\tau(R_A, R_H)$). ρ is more widely used.

Controllable Kendall’s Tau (τ_u): Inspired by Deutsch et al. (2022), we propose controllable Kendall’s tau that evaluates the agreement between rankings by focusing only on system pairs where the performance difference is small, i.e., where the absolute difference in the human-provided scores for two systems is less than a specified threshold u . This is useful for evaluating the bencher’s ability to distinguish between closely matched systems.

Let $\Delta r(s_i, s_k) = |r(s_i) - r(s_k)|$ represent the performance difference between systems s_i and s_k , where $r(s_i)$ is the system ratings derived from Chatbot Arena. For a given threshold u , we define the set of system pairs that are considered for evaluation as: $P_u = \{(s_i, s_k) \mid s_i, s_k \in S, \Delta r(s_i, s_k) \leq u, i \neq k\}$. To prevent the performance gap between two systems from being

so small that humans also find it difficult to distinguish between them, we only consider system pairs whose 95% confidence intervals in the Chatbot Arena ratings do not overlap: $Q = \{(s_i, s_k) \mid s_i, s_k \in S, \text{CI}(r(s_i)) \cap \text{CI}(r(s_k)) = \emptyset, i \neq k\}$. The controllable Kendall’s tau τ_u is then calculated only using the system pairs in $P_u \cap Q$:

$$\tau_u(R_A, R_H) = \frac{C_u - D_u}{\sqrt{(C_u + D_u + T_{A,u})(C_u + D_u + T_{H,u})}}$$

where C_u is the number of concordant pairs (i.e., pairs of systems (s_i, s_k) where the rank order between s_i and s_k is the same in both R_A and R_H). D_u is the number of discordant pairs. $T_{A,u}$ and $T_{H,u}$ are the number of ties in the automatic bencher’s ranking R_A and the human-provided ranking R_H . Only the system pairs within the set P_u are used.

2.3 Different Ways of Selecting Evaluation Models

We are comparing different evaluation models \mathcal{E} based on their ability to rank systems or predict instance-level human preferences. There are three key settings for evaluating evaluation models:

Setting 1: Standard Meta-Evaluation with System-Level Rankings ($R_{\mathcal{E}}^{(1)}$). The idea of Setting 1 is to rank evaluation models based on how well their corresponding automatic benchers align with the system rankings derived from Chatbot Arena. It is the standard way benchmarks such as Arena Hard are designed.

An automatic bencher $\mathcal{J}_A^{(i)}$ is composed of an evaluation model \mathcal{E}_i and **fixed** input set \mathcal{X} , evaluation type \mathcal{T} , and aggregation method \mathcal{G} , producing a system ranking $R_A^{(i)}$. The performance of evaluation model \mathcal{E}_i is measured by comparing $R_A^{(i)}$ with R_H using either Spearman’s $\rho(R_A^{(i)}, R_H)$, or Kendall’s $\tau(R_A^{(i)}, R_H)$ mentioned above. Evaluation models are ranked based on their performance, leading to a ranking for evaluation models: $R_{\mathcal{E}}^{(1)} : \{\mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_K\} \rightarrow \{1, 2, \dots, K\}$.

Setting 2: Instance-Level Human Judgments as Ground Truth ($R_{\mathcal{E}}^{(2)}$). We consider instance-level human judgments as the ground truth, without system-level aggregation or knowledge of which systems generated the responses.

Let $\mathcal{D} = \{(x_j, o_j^{(1)}, o_j^{(2)}, h_j)\}_{j=1}^m$ represent the dataset, where each entry consists of an input

x_j , two responses $o_j^{(1)}, o_j^{(2)}$, and human preference $h_j \in \{o_j^{(1)} > o_j^{(2)}, o_j^{(1)} < o_j^{(2)}\}$, indicating which response was preferred.

The evaluation model \mathcal{E}_i generates predictions $\hat{h}_j^{(i)} = \mathcal{T}_{\text{base.pairwise}}(\mathcal{E}_i, x_j, o_j^{(1)}, o_j^{(2)})$, and the instance-level accuracy is defined as:

$$\text{Accuracy}(\mathcal{E}_i) = \frac{1}{m} \sum_{j=1}^m \mathbb{1}(\hat{h}_j^{(i)} = h_j)$$

Evaluation models are ranked ($R_{\mathcal{E}}^{(2)}$) based on their accuracy in predicting human preferences.

Setting 3: Instance-Level Human Judgments with System Information and Aggregation ($R_{\mathcal{E}}^{(3)}$). We design this new setting to help us compare **Setting 1** and **Setting 2**. If we have access to instance-level human judgments as well as system information, we can aggregate both human judgments and evaluation model predictions into system-level rankings.

Dataset with system information: Each entry in the dataset $\mathcal{D}' = \{(x_j, o_j^{(1)}, o_j^{(2)}, s_j^{(1)}, s_j^{(2)}, h_j)\}_{j=1}^m$ includes: an input x_j , two responses $o_j^{(1)}, o_j^{(2)}$ generated by systems $s_j^{(1)}, s_j^{(2)}$, and human preference $h_j \in \{o_j^{(1)} > o_j^{(2)}, o_j^{(1)} < o_j^{(2)}\}$, indicating which response was preferred⁶. The steps to producing evaluation model rankings are as follows: (1) Human preferences are aggregated to produce a system ranking R'_H with an aggregate method \mathcal{G} . (2) The evaluation model \mathcal{E}_i provides instance-level predictions with \mathcal{T} , which are aggregated using the same method \mathcal{G} to produce a system ranking $R_A^{(i)}$. (3) Compare the two system rankings R'_H and $R_A^{(i)}$ using Spearman’s $\rho(R_A^{(i)}, R'_H)$ or Kendall’s $\tau(R_A^{(i)}, R'_H)$. (4) The evaluation models are then ranked ($R_{\mathcal{E}}^{(3)}$) based on the correlation between their aggregated rankings and the aggregated human judgments.

Setting 1 and Setting 3 are system-level evaluation, because aggregation methods are applied and a correlation value is calculated between two sets of system rankings, while **Setting 2 is the instance-level evaluation**. Note that we set the evaluation

⁶Please note that the instance-level human judgments in Chatbot Arena cannot be utilized in Setting 3, as they are not fully disclosed. Specifically, $x_j, o_j^{(1)}$, and $o_j^{(2)}$ are unavailable, and only $s_j^{(1)}, s_j^{(2)}$, and h_j are provided. Consequently, only the aggregated system rankings from Chatbot Arena (as used in Setting 1) are typically used.

type to $\mathcal{T}_{\text{base.pairwise}}$ and aggregation method to \mathcal{G}_{BT} in Setting 1 and Setting 3 to compare the results of the three settings more fairly.

2.4 Research Questions

Our research questions can be formalized as:

- RQ1: How to choose the appropriate components for building an effective automatic LLM benchmer?

$$\Rightarrow \mathcal{X}, \mathcal{E}, \mathcal{T}, \mathcal{A} = \arg \max_{\mathcal{X}, \mathcal{E}, \mathcal{T}, \mathcal{A}} \rho(R_A, R_H)$$

- RQ2: Does the performance of the benchmer decline when systems with smaller performance differences are selected for evaluation?

$$\Rightarrow \text{Does } \tau_u(R_A, R_H) \text{ decrease as } u \text{ decrease?}$$

- RQ3: Can we use instance-level rankings of evaluation models as a good reference to select evaluation models for LLM benchmers?

$$\Rightarrow \rho(R_{\mathcal{E}}^{(1)}, R_{\mathcal{E}}^{(2)}) \approx 1 \text{ and } \rho(R_{\mathcal{E}}^{(1)}, R_{\mathcal{E}}^{(3)}) \approx 1 \text{ and } \rho(R_{\mathcal{E}}^{(2)}, R_{\mathcal{E}}^{(3)}) \approx 1?$$

3 Experiments and Analyses

3.1 Setups

We selected two input sets, Arena Hard (Li et al., 2024b) and Alpaca Eval (Li et al., 2023), 18 LLMs as the systems to be evaluated, and 12 LLMs, including four proprietary OpenAI models and six open-source models, to serve as evaluation models. Please see Appendix C for detailed information.

3.2 RQ1: How to choose $\mathcal{X}, \mathcal{E}, \mathcal{T}, \mathcal{A}$ to maximize the benchmer’s performance?

This section discusses component selection for the benchmer and includes a cost analysis.

Input Set. Table 1 confirms that using Arena Hard as input set always yields higher correlations with the system rankings of Chatbot Arena compared to Alpaca Eval through different combinations of $\mathcal{E}, \mathcal{T}, \mathcal{A}$.

Evaluation Type. Table 1 shows the performance of a benchmer under different combinations of evaluation types and aggregation methods, with several representative LLMs as evaluation models. For many open-source models with fewer parameters, base pairwise is the optimal evaluation type, offering a significant advantage over other evaluation types. For strong proprietary models, the difference between base pairwise and point-wise evaluation types is small. Across all models

Alpaca Eval as the input set									
Evaluation Model		llama-3.1-70b		mixtral-8x7b		gpt-4o		gpt-4-turbo	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_tery	0.8844	0.7255	0.4757	0.3333	0.8782	0.7255	0.8611	0.6951
pointwise	win_ratio	0.8844	0.7255	0.4757	0.3333	0.8782	0.7255	0.8611	0.6951
pointwise	mean	0.5604	0.3987	0.2054	0.1373	0.8535	0.6993	0.8493	0.6863
pointwise	median	0.1806	0.1373	0.3953	0.2810	0.8762	0.7386	0.8803	0.7255
pairwise_base	bradley_tery	0.9009	0.7647	0.5253	0.3856	0.8838	0.7475	0.9112	0.7778
pairwise_base	win_ratio	0.9009	0.7647	0.5253	0.3856	0.8838	0.7475	0.9112	0.7778
pairwise_5point	bradley_tery	0.5026	0.3595	0.4510	0.3203	0.6615	0.5033	0.8184	0.6471
pairwise_5point	win_ratio	0.5046	0.3725	0.3251	0.2549	0.6037	0.4379	0.7977	0.6471
pairwise_base_ref	bradley_tery	0.8596	0.7124	0.4530	0.3333	0.8803	0.7124	0.9071	0.7908
pairwise_base_ref	win_ratio	0.8617	0.6993	0.4530	0.3333	0.8741	0.6993	0.9092	0.7778
pairwise_5point_ref	bradley_tery	0.5005	0.3856	0.4551	0.3333	0.6636	0.5163	0.8658	0.7124
pairwise_5point_ref	win_ratio	0.4881	0.3725	0.3457	0.2680	0.5707	0.4510	0.7090	0.5556

Arena Hard as the input set									
Evaluation Model		llama-3.1-70b		mixtral-8x7b		gpt-4o		gpt-4-turbo	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_tery	0.7207	0.5639	0.2693	0.2026	0.9401	0.8039	0.9092	0.7778
pointwise	win_ratio	0.7207	0.5639	0.2693	0.2026	0.9401	0.8039	0.9092	0.7778
pointwise	mean	0.6863	0.4771	0.6739	0.5163	0.9628	0.8693	0.9463	0.8562
pointwise	median	0.5439	0.4248	0.1538	0.1111	0.9587	0.8562	0.9319	0.7908
pairwise_base	bradley_tery	0.9340	0.8170	0.7441	0.5948	0.9443	0.8039	0.9381	0.8431
pairwise_base	win_ratio	0.9340	0.8170	0.7441	0.5948	0.9443	0.8039	0.9381	0.8431
pairwise_5point	bradley_tery	0.8411	0.7124	0.7296	0.6078	0.9216	0.7778	0.9463	0.8431
pairwise_5point	win_ratio	0.8142	0.6732	0.6223	0.4902	0.9360	0.8039	0.9133	0.8039
pairwise_base_ref	bradley_tery	0.8989	0.7647	0.7372	0.6033	0.9530	0.8262	0.9278	0.8170
pairwise_base_ref	win_ratio	0.8968	0.7516	0.7372	0.6033	0.9427	0.8000	0.9257	0.8039
pairwise_5point_ref	bradley_tery	0.8101	0.6732	0.7523	0.6340	0.9154	0.7516	0.8885	0.7778
pairwise_5point_ref	win_ratio	0.7131	0.5817	0.6367	0.4771	0.8473	0.6863	0.7915	0.6209

Table 1: Standard meta-evaluation results of different combinations of input sets, evaluation models, evaluation types, and aggregation methods (the correlation between the automatic benchers’ and ChatBot Arena’s evaluations). For the pairwise evaluation using a reference system, we show the results with gpt-4-0314 as the reference system. For the results with other LLMs as evaluation models, please refer to Table 4 and Table 5 in appendix.

used as evaluation models, 5-point pairwise generally performs worse than base pairwise, with the decline being particularly pronounced for open-source models, possibly due to their weaker ability to follow complex instructions. **Therefore, using 5-point pairwise as the evaluation type, as in Arena Hard, may be problematic.** In addition, by comparing the results of pairwise evaluation using the full data with those using GPT-4 as the reference system, we found that for most strong evaluation models, the performance of the benchers is similar when using the reference system, indicating that this is an efficient strategy. Further discussion on the reference system is in Appendix D.

Aggregation Method. As shown in Table 1, when the evaluation type is base pairwise, the system rankings produced by the Bradley-Terry model

and win ratio aggregation methods are mostly the same. Interestingly, when the evaluation type is pointwise, mean, the most common aggregation method, is not always optimal. Moreover, when the evaluation model is a smaller open-source model, pointwise combined with the Bradley-Terry model significantly outperforms its combination with mean. After examination, we found that this is because the majority of instance-level scores produced by these evaluation models are very similar, and aggregating them using mean or median across all samples leads to system scores that are hardly effective. In contrast, converting these scores to base pairwise retains the slight differences in scores for different responses to the same input. On the other hand, for OpenAI models, the performance of aggregating pointwise type data using either mean or the Bradley-Terry model is comparable.

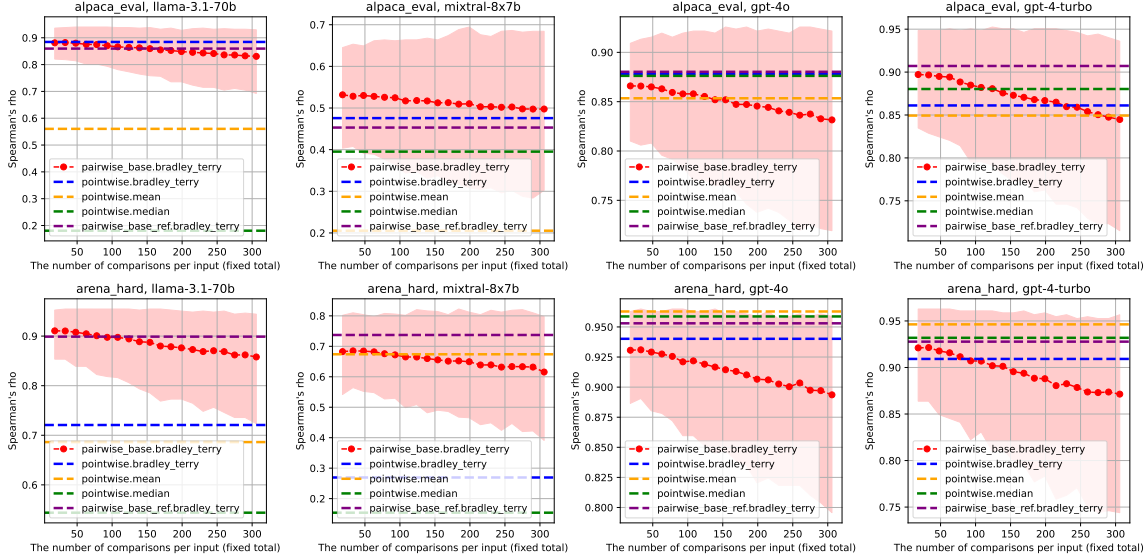


Figure 1: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic pairwise evaluation and human judgment when the total number of examples is set equal to that of pointwise evaluation. Results for pointwise and pairwise evaluations using gpt-4-0314 as the reference system are shown as horizontal lines. Both input sets and comparisons are sampled with replacement, with 1000 iterations. See Figure 7 and 8 in the appendix for other evaluation models.

Evaluation Model	Alpaca Eval		Arena Hard	
	ρ	Ranking	ρ	Ranking
gpt-4-turbo	0.9112	1	0.9381	2
gpt-4o	0.8838	3	0.9443	1
gpt-4o-mini	0.6821	4	0.8101	5
gpt-3.5-turbo	0.4923	7	0.6636	8
llama-3.1-70b	0.9009	2	0.9340	3
qwen-2.5-72b	0.4861	8	0.8411	4
mixtral-8x7b	0.5253	6	0.7441	6
llama-3.1-8b	0.6347	5	0.6656	7
glm-4-9b	0.2755	9	0.6367	9
mistral-7b-v0.3	0.2549	10	0.4076	10
mistral-7b-v0.1	0.1476	11	0.2755	12
llama-2-7b	-0.1291	12	0.2839	11

Table 2: Evaluation model rankings under Setting 1. The evaluation type and aggregation method are fixed to base pairwise and the Bradley-Terry model for all evaluation models.

Evaluation Models. Table 2 displays the performance of benchers with different LLMs as evaluation models. We found that gpt-4-turbo and gpt-4o are the strongest evaluation models. The rankings of evaluation models obtained from the Alpaca Eval and Arena Hard input sets show slight differences, but the correlation between them is high, with $\rho = 0.98$. Additionally, most of the smaller open-source models perform unsatisfactorily as evaluation models. However, llama-3.1-70b is an exception, performing even better than gpt-3.5-turbo and gpt-4o-mini.

Cost Analysis. In the previous analysis, we did not take into account the total number of instance-level evaluations produced by the evaluation models. As mentioned earlier, the total number of instance-level evaluations of pointwise evaluation and pairwise evaluation are different: pointwise evaluations amount to mn , whereas pairwise evaluations amount to $mn(n-1)$. This implies that using pairwise as the evaluation type will incur significantly higher costs, especially when using proprietary models as evaluation models. Therefore, we consider it necessary to limit the number of instance-level evaluations for pairwise to be equal to that of pointwise, i.e., mn . In this scenario, when using pairwise as the evaluation type, we need to balance the **number of inputs** with the **number of comparisons per input**. Figure 1 illustrates the bootstrapping mean and confidence intervals for base pairwise as the evaluation type under different numbers of comparisons per input. We found that:

1. When some open-source models are used as evaluation models, base pairwise still has a certain advantage even when the total number of instance-level evaluations for base pairwise is limited to that of the pointwise type.
2. As the number of comparisons per input increases (i.e., the number of inputs decreases), both the bootstrapping mean and the lower bound of the confidence interval for bencher performance decline. **This suggests that we should prioritize**

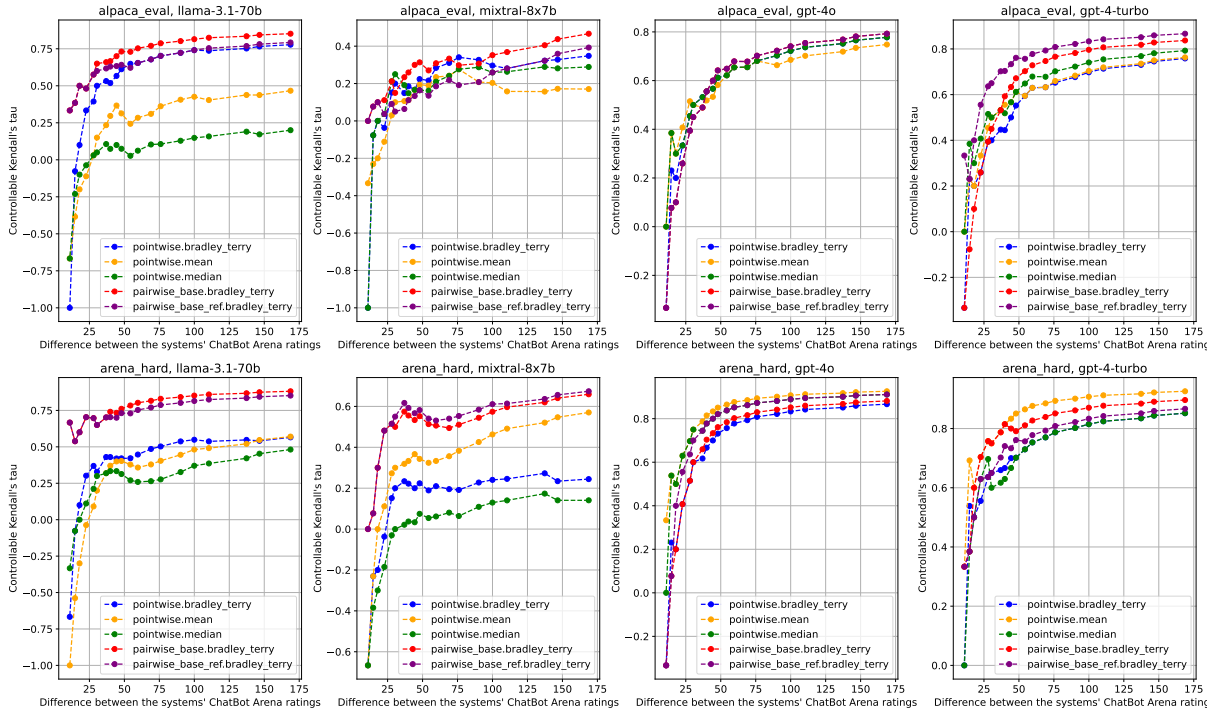


Figure 2: Controllable Kendall’s tau (τ_u) between the system rankings from automatic LLM benchers and human judgment when only partial system pairs are used. The X-axis denotes the value of threshold u , which controls the maximum difference between the systems’ ChatBot Arena ratings. For pairwise_base_ref, we show the results with gpt-4-0314 as the reference system. Across all settings, we found that the benchers’ performance degrades when they evaluate close-performing systems. See Figure 15 and 16 in the appendix for other evaluation models.

using more inputs rather than conducting more comparisons per input. As mentioned earlier, using a reference system is a practical strategy. For a more detailed cost analysis, please see appendix E.

3.3 RQ2: Does the performance of automatic LLM benchers degrade when evaluating LLMs with similar performance?

As shown in Figure 14 in the appendix, the performance differences among the selected LLM systems are uneven, so the difficulty of distinguishing between different LLMs varies. We introduced a threshold u so that only system pairs with performance differences smaller than it are used to calculate τ_u . Since fewer system pairs meet the requirement as u decreases, we can control the threshold u to select a specific proportion of system pairs. Specifically, we selected 5%, 10%, ..., and 100% of system pairs and observed the changes in bencher performance. Figure 2 shows that as u decreases (i.e., the performance differences between systems become smaller), the performance of almost all benchers declines sharply. For example, when Alpaca Eval is used as the input set and gpt-4o is used as the evaluation model, we found

that the Bencher’s performance degrades drastically by 25 points when comparing system pairs whose ChatBot Arena ratings differ by approximately 40 points, which is similar to the difference between qwen-1.5-72B and gpt-4-0314.

3.4 RQ3: Can we use instance-level rankings of evaluation models as a good reference to select evaluation models for LLM benchers?

Figure 3 displays the Spearman’s ρ between evaluation model rankings obtained from different datasets and settings. We observed the following key findings: First, the evaluation model rankings vary considerably across different settings, as evidenced by the fact that none of the Spearman’s ρ values are found to be 1.0. Second, system-level evaluation results are more consistent with each other. The evaluation model ranking obtained from Setting 3 is very similar to that from Setting 1, despite the systems to be evaluated and input sets in Alpaca Farm being different from those in Setting 1. In contrast, although the input set and involved systems in alpaca_farm_instance (Setting 2) and alpaca_farm_system (Setting 3)

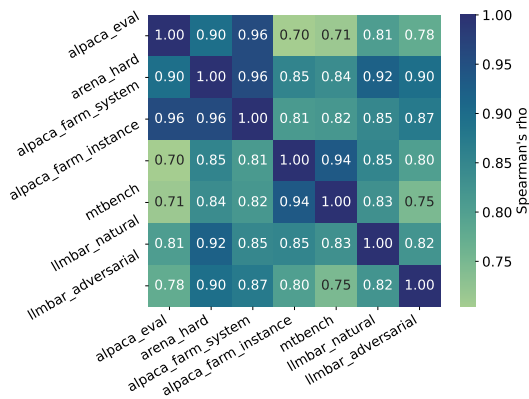


Figure 3: Spearman’s ρ between the evaluation model rankings under different meta-evaluation settings. **alpaca_eval** and **arena_hard** denote the evaluation model rankings are produced under **Setting 1**. **alpaca_farm_system** denotes that evaluation model rankings are produced under **Setting 3** where the instance-level human judgment of Alpaca Farm is used as ground truth, the Bradley-Terry model is used to aggregate both automatic evaluation and human evaluation. **alpaca_farm_instance**, **mtbench**, **llmbar_natural**, and **llmbar_adversarial** refer to that the evaluation model rankings are produced under **Setting 2** where instance-level human judgment of Alpaca Farm, MT-Bench, LLMBAR are used.

are exact the same, the evaluation model rankings they produce differ substantially. This suggests that whether the aggregation methods are applied has a great impact on the evaluation results. In addition, compared to `llmbar_adversarial`, the evaluation model ranking from `llmbar_natural` is more closely aligned with that from the system-level evaluation results. This is probably due to the responses in `llmbar_adversarial` are not from real systems, which makes it more different from system-level evaluation setting.

4 Related Work

Automatic LLM benchers. This kind of research always create a automatic leaderboard for various LLMs. As mentioned above, the main contribution of Arena Hard (Li et al., 2024b) lies in designing a pipeline for automatically constructing a challenging input set, also for Wildbench (Lin et al., 2024). In contrast, Alpaca Eval 2 (Dubois et al., 2024), focus on designing better evaluation models and use existing instruction datasets as the input set. Besides, MixEval (Ni et al., 2024) build LLM benchmarks by strategically mixing off-the-shelf ground-truth-based benchmarks that match real-

world user queries. Zhao et al. (2024) automate the entire LLM evaluation process using LLM-powered agents. There is another line of research that specifically focuses on how build LLM-based evaluation models (Gao et al., 2024).

Comparison between different evaluation types and aggregation methods. Liusie et al. (2024) compared the effectiveness of LLM-based evaluation models with absolute score prediction and pairwise comparison on several NLG tasks. Liu et al. (2024b) investigated which evaluation type is better under the instance-level evaluation settings of instruction following thoroughly. However, these studies are not oriented to system-level evaluation and aggregation methods are not involved. Liu et al. (2024a) proposed a more efficient rank aggregation method, but it is also used for ranking different instances rather than systems. Peyrard et al. (2021) showed the advantages of the Bradley-Terry model over mean and median on aggregating instance-level scores of NLP systems but this study focuses on what aggregation method makes more sense to choose as part of an evaluation metric. Daynauth et al. (2024) proposed three desirable properties of aggregation methods and evaluated the robustness of several aggregation methods. It is worth mentioning that a concurrent work also focused on automatic system-level LLM judges and analyzed their decisiveness and bias (Gera et al., 2024).

Instance-level evaluation of evaluation models. A lot of studies have focused on collecting instance-level human judgments to evaluate LLM-based evaluation models (Zeng et al., 2024b; Zheng et al., 2023; Dubois et al., 2023; Zhang et al., 2023; Wang et al., 2024b; Li et al., 2024a; Ye et al., 2024). However, to the best of our knowledge, no work has examined whether the evaluation results of them are consistent with the system-level settings.

5 Conclusions

We have explored how to construct effective LLM benchers and evaluate them. We provide several recommendations on how to combine the four components of an LLM bencher, analyze the current evaluation settings for LLM benchers and the evaluation models that serve as their components, and highlight the shortcomings in these settings. Our code and data will be available at <https://github.com/yale-nlp/RealRank>.

Limitations

The observations and conclusions of this paper are dependent on the datasets and selected LLMs, and cannot be guaranteed to apply under all circumstances. Furthermore, the open-source models included as evaluation models are relatively small in scale, which may not reflect the evaluation capabilities of larger open-source models.

References

- RALPH ALLAN BRADLEY and MILTON E. TERRY. 1952. [RANK ANALYSIS OF INCOMPLETE BLOCK DESIGNS: THE METHOD OF PAIRED COMPARISONS](#). *Biometrika*, 39(3-4):324–345.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael I. Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. [Chatbot arena: An open platform for evaluating llms by human preference](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Roland Daynauth, Christopher Clarke, Krisztian Flautner, Lingjia Tang, and Jason Mars. 2024. [Ranking unraveled: Recipes for llm rankings in head-to-head ai combat](#). *CoRR*, abs/2411.14483.
- Daniel Deutsch, Rotem Dror, and Dan Roth. 2022. [Re-examining system-level correlations of automatic summarization evaluation metrics](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 6038–6052. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurélien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Rozière, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Grégoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel M. Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, and et al. 2024. [The llama 3 herd of models](#). *CoRR*, abs/2407.21783.
- Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B. Hashimoto. 2024. [Length-controlled alpacaeval: A simple way to debias automatic evaluators](#). *CoRR*, abs/2404.04475.
- Yann Dubois, Chen Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- A.E. Elo. 1978. [The Rating of Chessplayers, Past and Present](#). Arco Pub.
- Mingqi Gao, Xinyu Hu, Jie Ruan, Xiao Pu, and Xiaojun Wan. 2024. [Llm-based NLG evaluation: Current status and challenges](#). *CoRR*, abs/2402.01383.
- Ariel Gera, Odellia Boni, Yotam Perlitz, Roy Bar-Haim, Lilach Eden, and Asaf Yehudai. 2024. [Justrank: Benchmarking llm judges for system ranking](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *CoRR*, abs/2310.06825.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Raghavi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Rewardbench: Evaluating reward models for language modeling](#). *CoRR*, abs/2403.13787.
- Junlong Li, Shichao Sun, Weizhe Yuan, Run-Ze Fan, Hai Zhao, and Pengfei Liu. 2024a. [Generative judge for evaluating alignment](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E. Gonzalez, and Ion Stoica. 2024b. [From crowdsourced data to high-quality benchmarks: Arena-hard and benchbuilder pipeline](#). *CoRR*, abs/2406.11939.

- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [AlpacaEval: An automatic evaluator of instruction-following models](https://github.com/tatsu-lab/alpaca_eval). https://github.com/tatsu-lab/alpaca_eval.
- Bill Yuchen Lin, Yuntian Deng, Khyathi Raghavi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, and Yejin Choi. 2024. [Wildbench: Benchmarking llms with challenging tasks from real users in the wild](#). *CoRR*, abs/2406.04770.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen 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, EMNLP 2023, Singapore, December 6-10, 2023*, pages 2511–2522. Association for Computational Linguistics.
- Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulic, Anna Korhonen, and Nigel Collier. 2024a. [Aligning with human judgement: The role of pairwise preference in large language model evaluators](#). *CoRR*, abs/2403.16950.
- Yixin Liu, Kejian Shi, Alexander R. Fabbri, Yilun Zhao, Peifeng Wang, Chien-Sheng Wu, Shafiq Joty, and Arman Cohan. 2024b. [Reife: Re-evaluating instruction-following evaluation](#). *CoRR*, abs/2410.07069.
- Adian Liusie, Potsawee Manakul, and Mark J. F. 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, EACL 2024 - Volume 1: Long Papers, St. Julian's, Malta, March 17-22, 2024*, pages 139–151. Association for Computational Linguistics.
- Jinjie Ni, Fuzhao Xue, Xiang Yue, Yuntian Deng, Mahir Shah, Kabir Jain, Graham Neubig, and Yang You. 2024. [Mixeval: Deriving wisdom of the crowd from LLM benchmark mixtures](#). *CoRR*, abs/2406.06565.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Maxime Peyrard, Wei Zhao, Steffen Eger, and Robert West. 2021. [Better than average: Paired evaluation of NLP systems](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2301–2315. Association for Computational Linguistics.
- Lewis Tunstall, Edward Emanuel Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro Von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M Rush, and Thomas Wolf. 2024. [Zephyr: Direct distillation of LM alignment](#). In *First Conference on Language Modeling*.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. 2024a. [A survey on large language model based autonomous agents](#). *Frontiers Comput. Sci.*, 18(6):186345.
- Yidong Wang, Zhuohao Yu, Wenjin Yao, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. 2024b. [Pandalm: An automatic evaluation benchmark for LLM instruction tuning optimization](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. 2024. [Qwen2 technical report](#). *CoRR*, abs/2407.10671.
- Seonghyeon Ye, Doyoung Kim, Sungdong Kim, Hyeonbin Hwang, Seungone Kim, Yongrae Jo, James Thorne, Juho Kim, and Minjoon Seo. 2024. [FLASK: fine-grained language model evaluation based on alignment skill sets](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang,

Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024a. [Chatglm: A family of large language models from GLM-130B to GLM-4 all tools](#). *CoRR*, abs/2406.12793.

Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. 2024b. [Evaluating large language models at evaluating instruction following](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

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](#). *CoRR*, abs/2308.01862.

Ruochen Zhao, Wenxuan Zhang, Yew Ken Chia, Deli Zhao, and Lidong Bing. 2024. [Auto arena of llms: Automating LLM evaluations with agent peer-battles and committee discussions](#). *CoRR*, abs/2405.20267.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. 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 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

A Conversion of 5-Point Pairwise and Pointwise Evaluation Type to Base Pairwise Type

To convert 5-point pairwise comparisons into base pairwise comparisons, we apply the following rules as Li et al. (2024b):

- $o_{ij} \gg o_{kj}$ or $o_{ij} \ll o_{kj}$ converts to 6 instances of $o_{ij} > o_{kj}$ or $o_{ij} < o_{kj}$
- $o_{ij} > o_{kj}$ or $o_{ij} < o_{kj}$ converts to 2 instances of $o_{ij} > o_{kj}$ or $o_{ij} < o_{kj}$
- $o_{ij} = o_{kj}$ converts to 1 instance of $o_{ij} > o_{kj}$ and 1 instance of $o_{ij} < o_{kj}$

For pointwise evaluation type, the conversion follows this rule, and the total number of its instances is expanded to $mn(n-1)/2$.

$$\mathcal{T}_{\text{base.pairwise}}(\mathcal{E}, x_j, o_{ij}, o_{kj}) = \begin{cases} o_{ij} > o_{kj} & \text{if } \mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{ij}) > \mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{kj}) \\ o_{ij} < o_{kj} & \text{if } \mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{ij}) < \mathcal{T}_{\text{pointwise}}(\mathcal{E}, x_j, o_{kj}) \end{cases}$$

B Definitions of Spearman’s ρ and Kendall’s τ

$$\rho(R_A, R_H) = 1 - \frac{6 \sum_{i=1}^n (R_A(s_i) - R_H(s_i))^2}{n(n^2 - 1)}$$

Where $R_A(s_i)$ is the rank of system s_i produced by the automatic bench and n is the total number of systems being ranked.

$$\tau(R_A, R_H) = \frac{C - D}{\sqrt{(C + D + T_A)(C + D + T_H)}}$$

Where C is the number of concordant pairs (i.e., pairs of systems (s_i, s_k) where the rank order between s_i and s_k is the same in both R_A and R_H). D is the number of discordant pairs (i.e., pairs where the rank order differs between R_A and R_H). T_A is the number of ties in the automatic bench’s ranking R_A (i.e., where two systems have the same rank). T_H is the number of ties in the human-provided ranking R_H .

C Experimental Setup

Input Set. For RQ1 and RQ2, we selected two input sets, Arena Hard (Li et al., 2024b) and Alpaca Eval (Li et al., 2023), containing 500 and 805 inputs, respectively.

Systems. 18 LLMs that ranked highly on Chatbot Arena were chosen as the systems to be evaluated⁷. The data from Chatbot Arena as of July 30, 2024, was used to compute the scores and rankings, serving as the ground truth. We reused the responses from the systems on Arena Hard, collected by Li et al. (2024b), and the responses on Alpaca Eval, collected by Li et al. (2023).

Evaluation Models. We selected 10 LLMs, including four proprietary OpenAI models and six open-source models, to serve as evaluation models. These include models gpt-4-turbo (gpt-4-turbo-2024-04-09 (OpenAI, 2023)), gpt-4o (gpt-4o-2024-08-06), gpt-4o-mini (gpt-4o-mini-2024-07-18), gpt-3.5-turbo (gpt-3.5-turbo-0125), llama-3.1-70b (meta-llama/Llama-3.1-70B (Dubey et al., 2024)), qwen-2.5-72b (Qwen/Qwen2.5-72B-Instruct (Yang et al., 2024)), mixtral-8x7b (mistralai/Mixtral-8x7B-Instruct-v0.1), llama-3.1-8b (meta-llama/Llama-3.1-8B), glm-4-9b

⁷gpt-4-turbo-2024-04-09, claude-3-opus-20240229, claude-3-sonnet-20240229, gpt-4-0314, gpt-4-0613, mistral-large-2402, mistral-medium, qwen1.5-72b-chat, claude-2.0, claude-2.1, gpt-3.5-turbo-0613, mixtral-8x7b-instruct-v0.1, yi-34b-chat, gemini-pro, dbrx-instruct-preview, tulu-2-dpo-70b, vicuna-33b, starling-lm-7b-alpha

(THUDM/glm-4-9b-chat (Zeng et al., 2024a)), mistral-7b-v0.3 (mistralai/Mistral-7B-Instruct-v0.3), mistral-7b-v0.1 (mistralai/Mistral-7B-Instruct-v0.1 (Jiang et al., 2023)), llama-2-7b (meta-llama/Llama-2-7b-chat-hf).

Evaluation Types. The evaluation types were implemented as follows. As for pointwise, given an input and a response, evaluation models were asked to score the response on a 0-9 scale. We followed the method of Liu et al. (2023), using token probability to compute a weighted sum of the model’s output for the 0-9 scale as the final instance-level score. For base Pairwise or 5-Point Pairwise, given an input and two responses, evaluation models were asked to either choose the better response or make a five-category comparison. The results of the automatic evaluations were extracted using regular expressions. If the evaluation model did not provide the requested judgment, a random evaluation result was assigned to the instance. The specific prompt designs can be found in the appendix.

Aggregation Methods. We adopted Chatbot Arena’s implementation of the Bradley-Terry model and implemented other aggregation methods ourselves.

Additional Setup for RQ3. We selected three datasets containing instance-level human judgments: Alpaca Farm (Dubois et al., 2023), LLMBAR (Zeng et al., 2024b), and MT-bench (Zheng et al., 2023). Among these, Alpaca Farm has a larger data scale, and 22 systems are involved, making it suitable for both Setting 2 and Setting 3. The number of systems in MT-bench and LLMBAR is too small, or system information is unavailable, so they can only be used for Setting 2. Additionally, LLMBAR contains two subsets: natural and adversarial. The natural subset contains responses from real systems, while the responses of the adversarial subset were generated using adversarial techniques.

D Discussions on Reference Systems

Although using a strong system as a reference system is a common practice, its rationale has not been thoroughly examined. Therefore, we calculated the performance of the benchner when each of the 18 evaluated systems is used as the reference system and demonstrated the relationship between the benchner’s performance and the Chatbot Arena rating of the LLM used as the reference system, as shown in Figures 4, 5, and 6. We found that, for

most strong evaluation models, the performance of the benchner is negatively correlated with the level of the LLM used as the reference system. Therefore, we believe this issue warrants further investigation, and at the very least, we should be more cautious in selecting the reference system.

E Cost Analysis

To perform a more granular cost analysis of the benchner, we measure cost in terms of the total number of tokens consumed during automatic evaluation and the evaluation model’s cost per million tokens. For OpenAI models, we use the pricing for the Batch API⁸ provided on their official website. For open-source models, we estimate costs based on electricity consumption during inference on GPU servers, ignoring the expenses of purchasing and maintaining GPUs.

Specifically, we estimate the number of tokens processed per second by open-source models based on throughput data from the vLLM blog⁹:

- **8B Model on 1xH100:** Approximately 30 queries per second (QPS) with an average of 179 output tokens per query. Tokens per second: $30 \times 179 = 5370$ tokens/second.
- **70B Model on 4xH100:** Approximately 15 QPS with an average of 179 output tokens per query. Tokens per second: $15 \times 179 = 2685$ tokens/second.

Next, we calculate the time required to process 1 million tokens:

- **8B Model:** $\frac{1000000}{5370} \approx 186.2$ seconds.
- **70B Model:** $\frac{1000000}{2685} \approx 372.4$ seconds.

We then estimate the energy consumption for processing 1 million tokens:

- **8B Model on 1xH100:** Energy consumed = Power (700W) \times Time (186.2 seconds) = 130340 watt-seconds \approx 36.21 watt-hours (Wh).
- **70B Model on 4xH100:** Energy consumed per GPU = Power (700W) \times Time (372.4 seconds) = 260680 watt-seconds \approx 72.41 Wh, Total energy for 4 GPUs = 72.41 Wh \times 4 = 289.6 Wh.

⁸<https://openai.com/api/pricing/>

⁹<https://blog.vllm.ai/2024/09/05/perf-update.html>

Model	\$/1M tokens
llama-3.1-8b	0.006
llama-3.1-70b	0.05
qwen-2.5-72b	0.05
gpt-4o	1.25
gpt-4-turbo	5.00
gpt-3.5-turbo	0.25
gpt-4o-mini	0.075

Table 3: Financial cost of different evaluation models during inference. The prices of open-source models are estimated by electricity rate. We use the cost of input tokens because our output length is negligible.

Finally, assuming an electricity rate of \$0.17 per kilowatt-hour (kWh), which is the average rate in the US¹⁰:

- **8B Model:** Cost = $36.21 \text{ Wh} \times \frac{\$0.17}{1000 \text{ Wh}} = \0.00577 ($\sim \$0.006$) per 1M tokens.
- **70B Model:** Cost = $289.6 \text{ Wh} \times \frac{\$0.17}{1000 \text{ Wh}} = \0.0492 ($\sim \$0.05$) per 1M tokens.

Thus, we derived the cost per 1M tokens for each model, as shown in Table 3. Furthermore, we performed sampling with replacement on the input set at proportions of 0.02, 0.03, 0.04, 0.05, 0.10, and up to 1.00 to generate input sets of varying sizes¹¹. Consequently, the cost of different combinations of evaluation models, evaluation types, and aggregation methods across input sets of different sizes can be estimated. Figures 9 and 10 illustrate the best-performing combinations under varying budget constraints. From these results, we observe that open-source models demonstrate a significant advantage when the budget is highly constrained.

F Confidence Interval Analysis

In Table 1, we computed the performance of benchers of different combinations of $\mathcal{X}, \mathcal{E}, \mathcal{T}, \mathcal{G}$, but these are merely point estimates. To assess the accuracy of our estimates, we obtained the 95% confidence intervals for bencher performance through bootstrapping. As shown in Figure 11, 12, and 13, with an increasing number of inputs, the confidence intervals for most benchers become narrower. However, for some combinations of evaluation type and aggregation method, the confidence

intervals remain relatively wide when the sample ratio equals 1.0, indicating a high degree of uncertainty in the performance estimates. In contrast, the combination of base pairwise and the Bradley-Terry model has narrower confidence intervals (the red section), suggesting that the performance estimates for this combination are more precise.

G Other Figures and Tables

¹⁰https://www.bls.gov/regions/midwest/data/averageenergyprices_selectedareas_table.htm

¹¹Repeating 100 times and taking the average performance.

Alpaca Eval as the input set									
Evaluation Model		qwen-2.5-72b		llama-3.1-8b		gpt-4o-mini		gpt-3.5-turbo	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_terry	0.8101	0.6471	0.2982	0.1503	0.8452	0.6863	0.5707	0.4510
pointwise	win_ratio	0.8101	0.6471	0.2982	0.1503	0.8452	0.6863	0.5707	0.4510
pointwise	mean	0.8576	0.7124	0.0423	0.0719	0.8700	0.7124	0.6574	0.4902
pointwise	median	0.7895	0.6209	0.2239	0.0980	0.8493	0.6863	0.5315	0.4248
pairwise_base	bradley_terry	0.4592	0.3203	0.6347	0.4902	0.6821	0.5033	0.4923	0.3464
pairwise_base	win_ratio	0.4592	0.3203	0.6347	0.4902	0.6821	0.5033	0.4923	0.3464
pairwise_5point	bradley_terry	0.4221	0.3072	0.4241	0.2810	0.4799	0.3333	0.4510	0.3725
pairwise_5point	win_ratio	0.2611	0.1895	0.2033	0.1503	0.4551	0.3203	0.3354	0.2549
pairwise_base_ref	bradley_terry	0.4014	0.2680	0.5996	0.4510	0.5397	0.4248	0.5459	0.3987
pairwise_base_ref	win_ratio	0.4014	0.2680	0.5865	0.4328	0.5335	0.3987	0.5459	0.3987
pairwise_5point_ref	bradley_terry	0.6140	0.4641	0.3808	0.2680	0.5191	0.3464	0.1930	0.1242
pairwise_5point_ref	win_ratio	0.3395	0.2418	0.3457	0.2680	0.4324	0.3072	0.0815	0.0588

Arena Hard as the input set									
Evaluation Model		qwen-2.5-72b		llama-3.1-8b		gpt-4o-mini		gpt-3.5-turbo	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_terry	0.6739	0.4641	0.5232	0.4379	0.8824	0.7255	0.5212	0.3987
pointwise	win_ratio	0.6739	0.4641	0.5232	0.4379	0.8824	0.7255	0.5212	0.3987
pointwise	mean	0.8658	0.6993	0.1744	0.0980	0.9174	0.7908	0.5480	0.3856
pointwise	median	0.7069	0.5294	0.3333	0.2941	0.9112	0.7778	0.4840	0.3464
pairwise_base	bradley_terry	0.8411	0.6863	0.6656	0.5556	0.8101	0.6732	0.6636	0.5686
pairwise_base	win_ratio	0.8411	0.6863	0.6656	0.5556	0.8101	0.6732	0.6636	0.5686
pairwise_5point	bradley_terry	0.9236	0.7647	0.3086	0.2288	0.6491	0.4771	0.6656	0.5033
pairwise_5point	win_ratio	0.8989	0.7386	0.2797	0.2157	0.6491	0.4771	0.6120	0.4510
pairwise_base_ref	bradley_terry	0.7207	0.5771	0.6267	0.5246	0.7152	0.5817	0.5772	0.4721
pairwise_base_ref	win_ratio	0.7186	0.5639	0.6371	0.5377	0.7214	0.5817	0.5772	0.4721
pairwise_5point_ref	bradley_terry	0.9030	0.7386	0.2549	0.1895	0.6347	0.4771	0.6161	0.4641
pairwise_5point_ref	win_ratio	0.6801	0.5033	0.1786	0.1311	0.6285	0.4641	0.5872	0.4379

Table 4: Standard meta-evaluation results of different combinations of input sets, evaluation models, evaluation types, and aggregation methods (the correlation between the automatic benchers’ and ChatBot Arena’s evaluations). For the pairwise evaluation using a reference system, we show the results with gpt-4-0314 as the reference system.

Alpaca Eval as the input set									
Evaluation Model		glm-4-9b		mistral-7b-v0.3		mistral-7b-v0.1		llama-2-7b	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_tery	0.2693	0.2026	0.3354	0.2288	-0.0898	-0.0327	-0.0072	0.0196
pointwise	win_ratio	0.2693	0.2026	0.3354	0.2288	-0.0898	-0.0327	-0.0072	0.0196
pointwise	mean	0.1063	0.0588	0.2219	0.1373	-0.1930	-0.1373	-0.0506	-0.0065
pointwise	median	0.2900	0.2157	0.2755	0.1895	-0.0093	0.0327	-0.0114	0.0065
pairwise_base	bradley_tery	0.2755	0.1765	0.2549	0.1765	0.1476	0.1111	-0.1291	-0.0724
pairwise_base	win_ratio	0.2755	0.1765	0.2549	0.1765	0.1476	0.1111	-0.1291	-0.0724
pairwise_5point	bradley_tery	0.2425	0.1503	-0.0279	0.0196	-0.1414	-0.085	0.1269	0.0719
pairwise_5point	win_ratio	0.1765	0.1111	-0.1517	-0.0588	-0.1414	-0.085	0.1311	0.0588
pairwise_base_ref	bradley_tery	0.3065	0.2157	0.3397	0.2098	0.2095	0.1503	-0.2183	-0.1258
pairwise_base_ref	win_ratio	0.2466	0.1765	0.2571	0.1574	0.1744	0.1242	-0.2257	-0.1329
pairwise_5point_ref	bradley_tery	0.2776	0.1634	-0.1703	-0.1242	-0.1176	-0.0821	0.0795	0.0588
pairwise_5point_ref	win_ratio	0.2095	0.1242	-0.2611	-0.1765	-0.0836	-0.0525	0.0609	0.0588

Arena Hard as the input set									
Evaluation Model		glm-4-9b		mistral-7b-v0.3		mistral-7b-v0.1		llama-2-7b	
Type	Aggregation	ρ	τ	ρ	τ	ρ	τ	ρ	τ
pointwise	bradley_tery	0.4696	0.3464	-0.0918	-0.0196	0.0237	0.0458	0.1723	0.0980
pointwise	win_ratio	0.4696	0.3464	-0.0918	-0.0196	0.0237	0.0458	0.1723	0.0980
pointwise	mean	0.4345	0.3203	0.2755	0.1765	0.0918	0.0327	0.2879	0.1895
pointwise	median	0.5005	0.3595	-0.1662	-0.0719	-0.0299	0.0065	0.2054	0.1373
pairwise_base	bradley_tery	0.6367	0.4902	0.4076	0.2941	0.2755	0.1765	0.2839	0.1836
pairwise_base	win_ratio	0.6367	0.4902	0.4076	0.2941	0.2755	0.1765	0.2839	0.1836
pairwise_5point	bradley_tery	0.5315	0.4118	0.1538	0.0850	0.2487	0.1503	-0.1930	-0.1634
pairwise_5point	win_ratio	0.4964	0.3725	0.2012	0.1242	0.2487	0.1503	-0.1414	-0.0980
pairwise_base_ref	bradley_tery	0.6536	0.5115	0.2690	0.1728	0.2654	0.1967	0.1998	0.1595
pairwise_base_ref	win_ratio	0.6536	0.5115	0.2690	0.1728	0.2241	0.1705	0.1595	0.1267
pairwise_5point_ref	bradley_tery	0.5418	0.4248	0.2136	0.1242	0.0974	0.0490	-0.4200	-0.3203
pairwise_5point_ref	win_ratio	0.5624	0.4510	0.1765	0.0850	0.2239	0.1503	-0.4407	-0.3464

Table 5: Standard meta-evaluation results of different combinations of input sets, evaluation models, evaluation types, and aggregation methods (the correlation between the automatic benchers’ and ChatBot Arena’s evaluations). For the pairwise evaluation using a reference system, we show the results with gpt-4-0314 as the reference system.

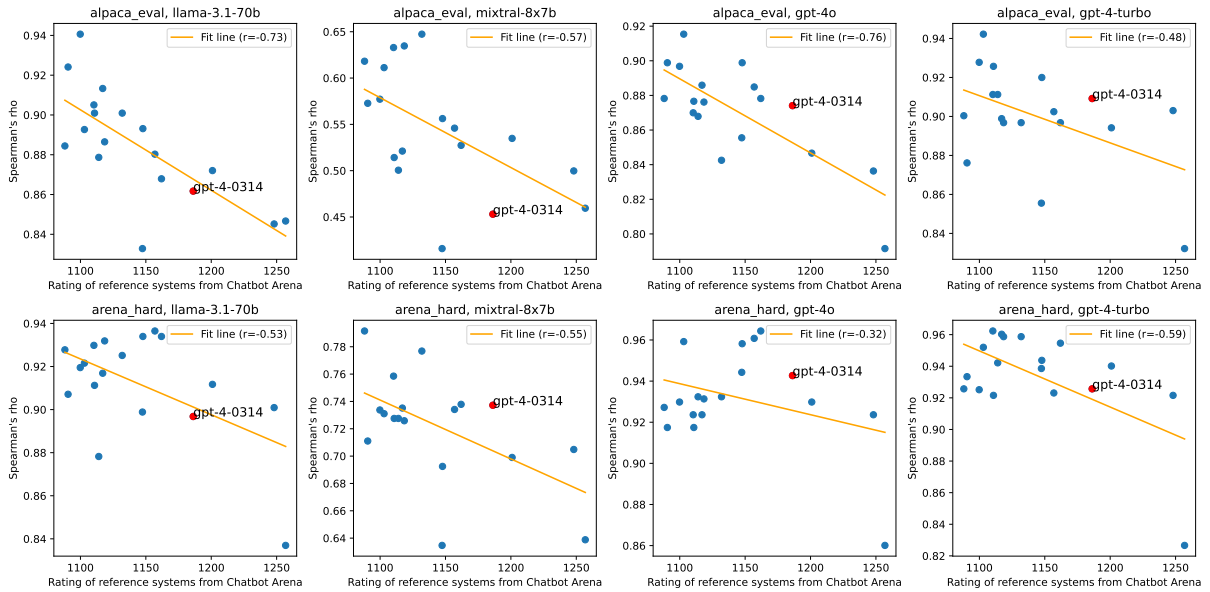


Figure 4: The relationships between the performance of the bencher and the Chatbot Arena ratings of the LLM used as the reference systems. The evaluation type is fixed as base pairwise.

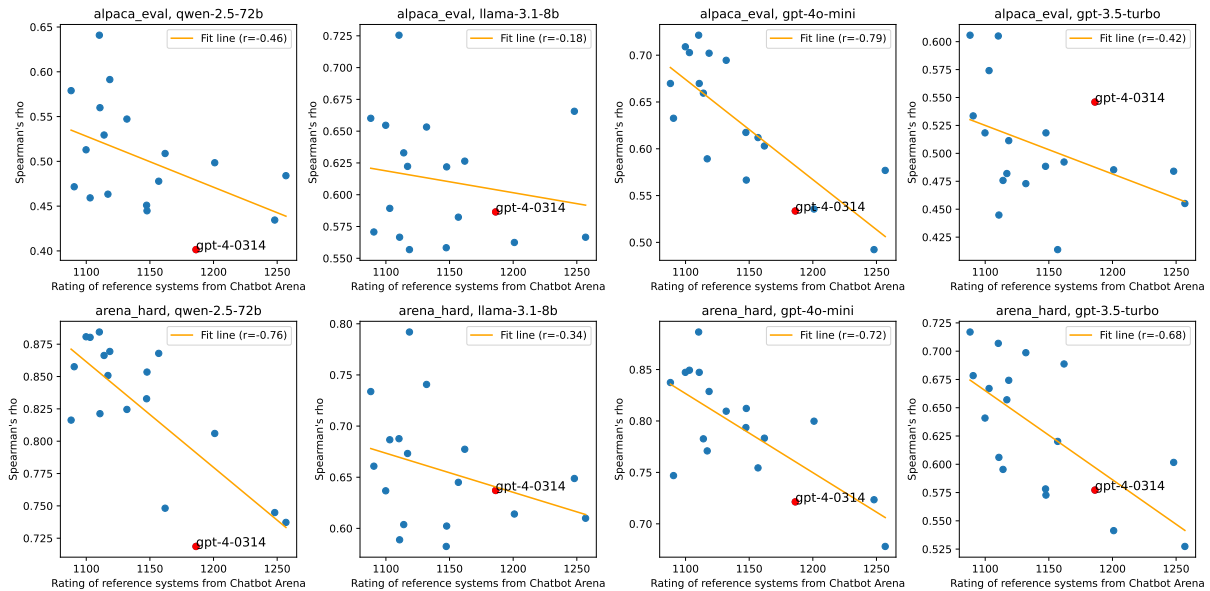


Figure 5: The relationships between the performance of the bencher and the Chatbot Arena ratings of the LLM used as the reference systems. The evaluation type is fixed as base pairwise.

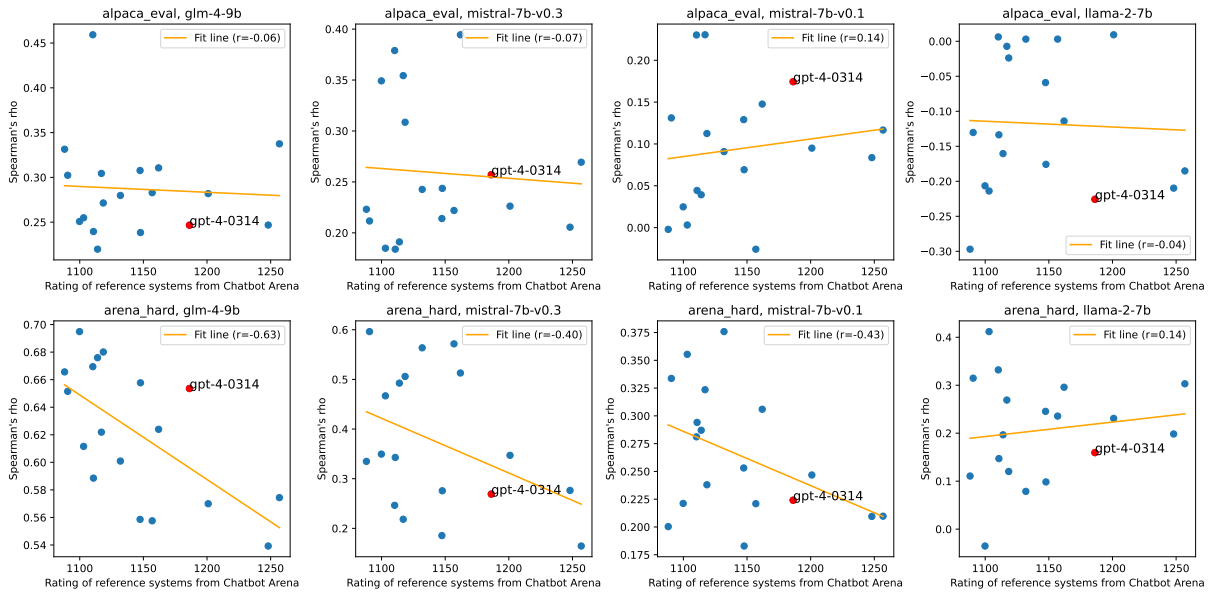


Figure 6: The relationships between the performance of the bencher and the Chatbot Arena ratings of the LLM used as the reference systems. The evaluation type is fixed as base pairwise.

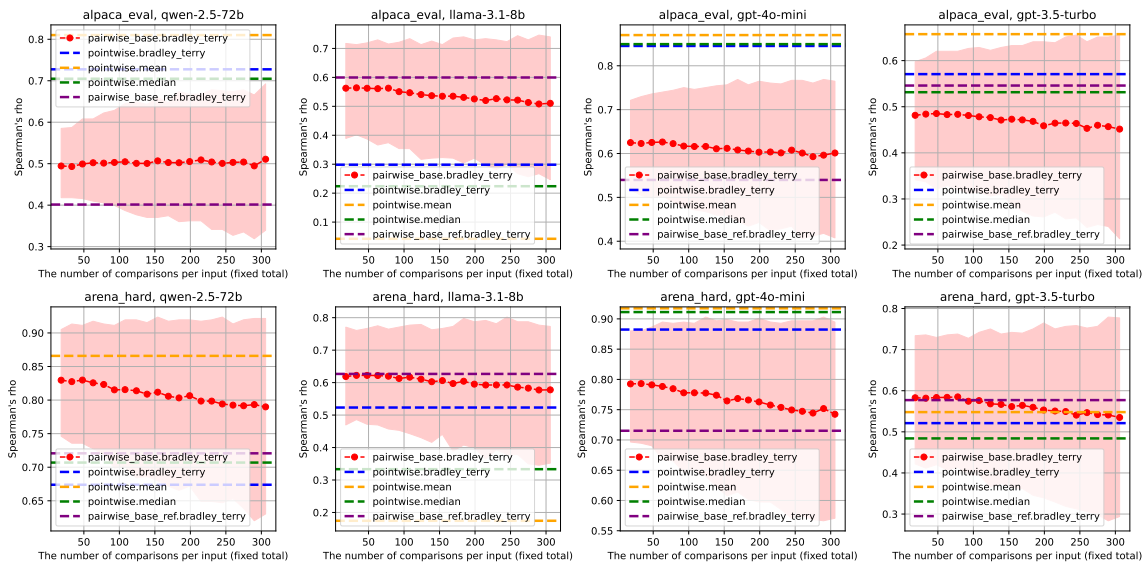


Figure 7: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic pairwise evaluation and human judgment when the total number of examples is set equal to that of pointwise evaluation. Results for pointwise and pairwise evaluations using gpt-4-0314 as the reference system are shown as horizontal lines. Both input sets and comparisons are sampled with replacement, with 1000 iterations.

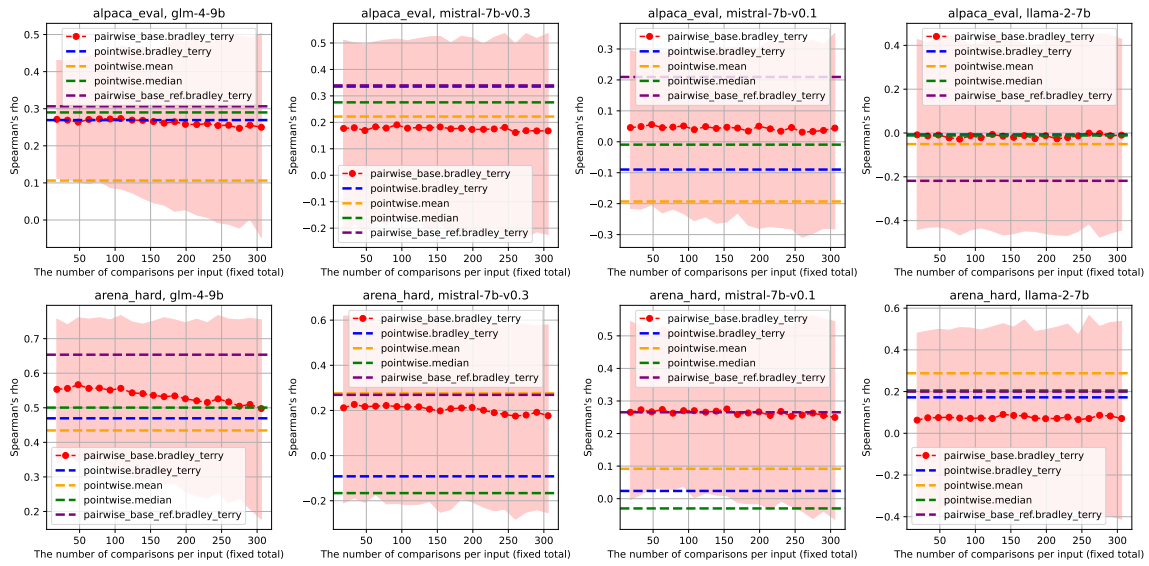


Figure 8: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic pairwise evaluation and human judgment when the total number of examples is set equal to that of pointwise evaluation. Results for pointwise and pairwise evaluations using gpt-4-0314 as the reference system are shown as horizontal lines. Both input sets and comparisons are sampled with replacement, with 1000 iterations.




Figure 9: The best-performing combinations under varying budgets on Alpaca Eval.




Figure 10: The best-performing combinations under varying budgets on Arena Hard. Using OpenAI models as evaluation models and employing pairwise evaluation incurs higher costs but does not outperform pointwise evaluation.

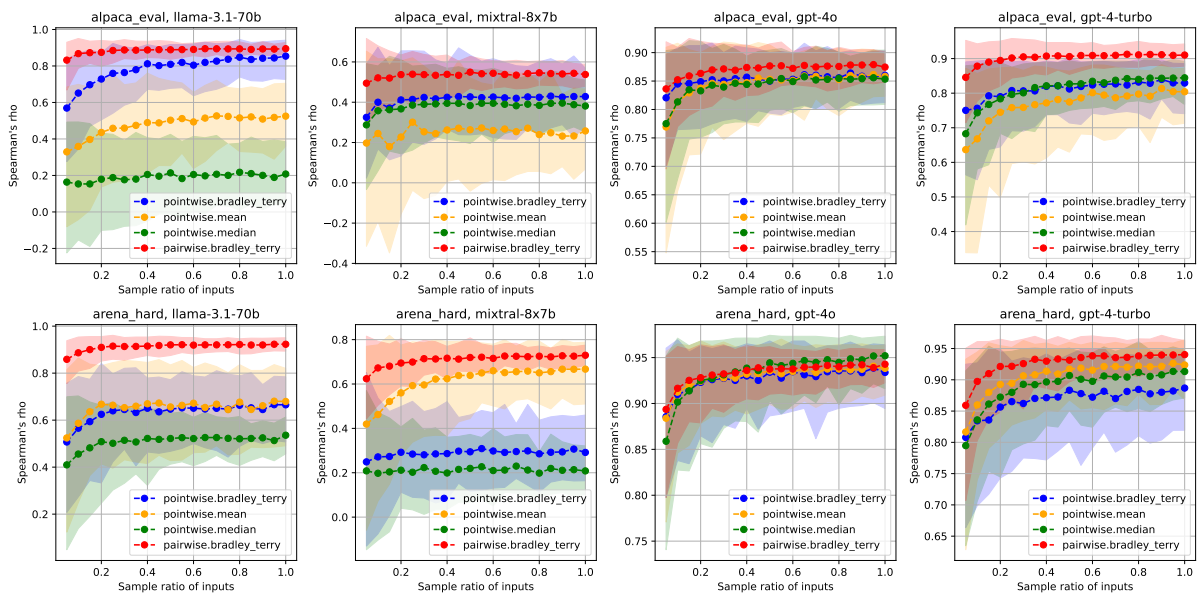


Figure 11: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic evaluation metrics and human judgment. The input set is sampled with replacement. The number of iterations is set to 100.

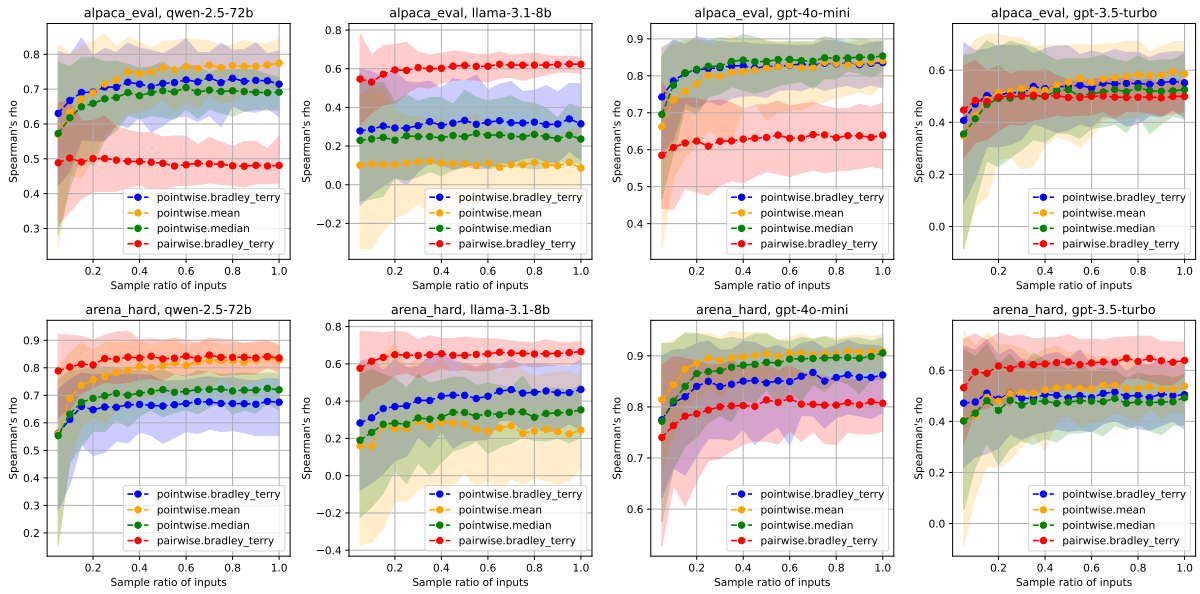


Figure 12: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic evaluation metrics and human judgment. The input set is sampled with replacement. The number of iterations is set to 100.

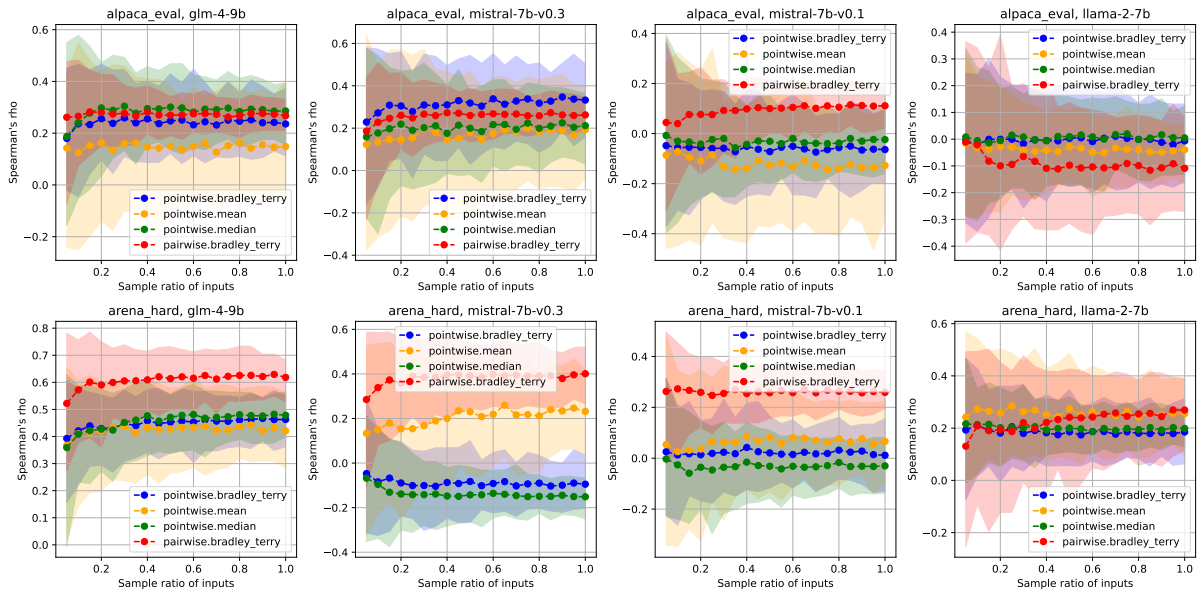


Figure 13: Bootstrapping mean and 95% confidence interval of the correlations between the system rankings of automatic evaluation metrics and human judgment. The input set is sampled with replacement. The number of iterations is set to 100.




Figure 14: System ratings of Chatbot Arena with bootstrapping mean and 95% confidence interval.

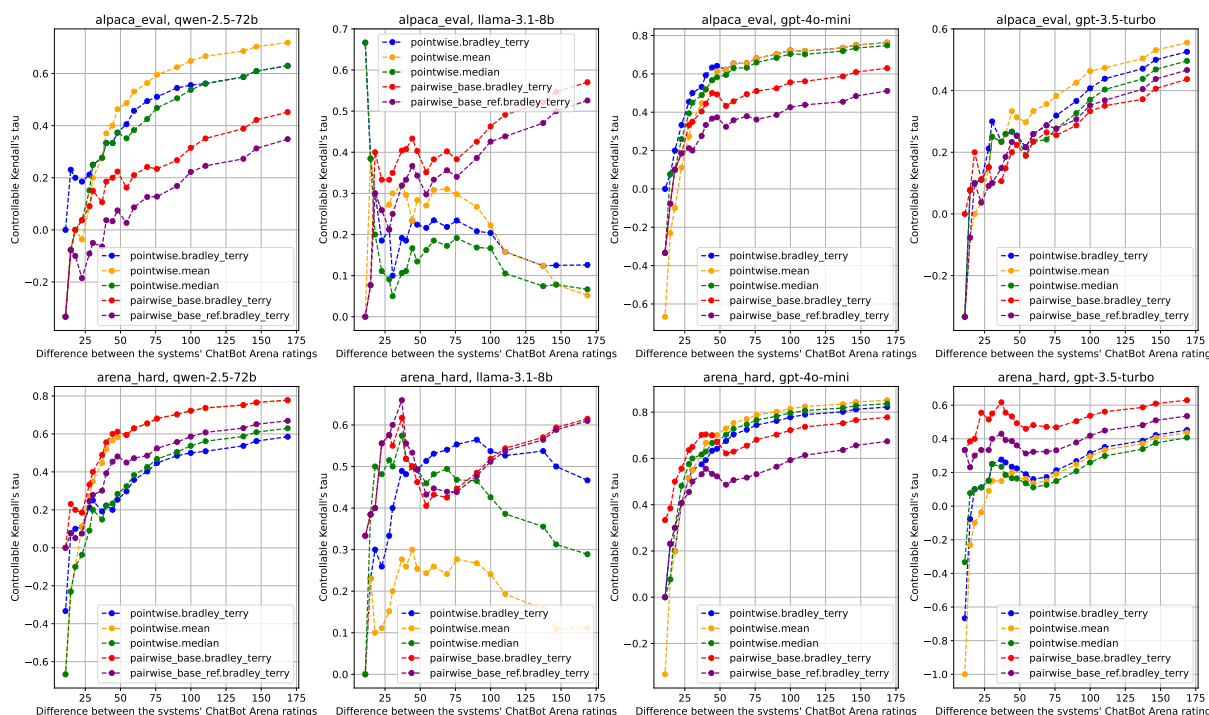


Figure 15: Controllable Kendall's tau (τ_u) between the system rankings from automatic benchers and human judgment when only partial system pairs are used. The X-axis denotes the value of threshold u , which controls the maximum difference between the systems' ChatBot Arena ratings. Across all settings, we found that the LLM benchers' performance degrades when they evaluate close-performing systems.

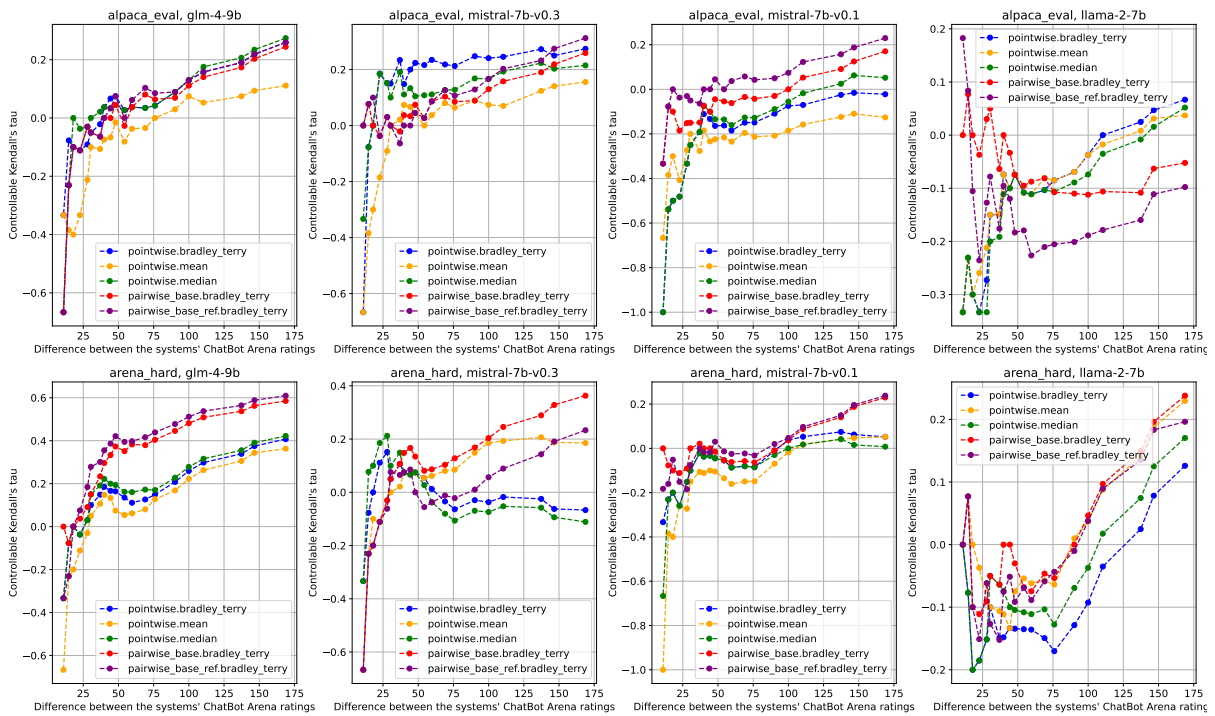


Figure 16: Controllable Kendall’s tau (τ_u) between the system rankings from automatic benchers and human judgment when only partial system pairs are used. The X-axis denotes the value of threshold u , which controls the maximum difference between the systems’ ChatBot Arena ratings. Across all settings, we found that the LLM benchers’ performance degrades when they evaluate close-performing systems.

Prompts and Instructions

<lim_start> system
 You are a helpful assistant in evaluating the quality of the outputs for a given instruction. Your goal is to score the output for the given instruction.
 <lim_end!>
 <lim_start> user
 Score the Output on a scale from 0 to 9 for the given instruction, where a score of zero means "poor quality" and score of nine means "perfect quality". 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 one number from 0 to 9. Do NOT output any other words.

Instruction:
 {INSTRUCTION}

Output:
 {OUTPUT}

What is your rating for the Output?
 <lim_end!>

Table 6: Prompt template for pointwise evaluation.

Prompts and Instructions

<lim_start> 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.

<lim_endl>

<lim_start> user

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:

{INSTRUCTION}

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)":

<lim_endl>

Table 7: Prompt template for base pairwise evaluation.

Prompts and Instructions

<lim_start> 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.

<lim_endl>

<lim_start> user

Output (a) and Output (b) are generated by two different AI chatbots for the given instruction respectively. Output one of the following choices as your verdict:

1. Output (a) is significantly better.
2. Output (a) is slightly better.
3. Tie, relatively the same.
4. Output (b) is slightly better.
5. Output (b) is significantly better.

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 1, 2, 3, 4, or 5. Do NOT output any other words.

Instruction:
{INSTRUCTION}

Output (a):
{OUTPUT_1}

Output (b):
{OUTPUT_2}

What is your verdict? Your response should be 1, 2, 3, 4, or 5:
<lim_endl>

Table 8: Prompt template for 5-point pairwise evaluation.