
Generalised Probabilistic Modelling and Improved Uncertainty Estimation in Comparative LLM-as-a-judge

Yassir Fathullah¹

Mark J. F. Gales¹

¹Engineering Department, University of Cambridge, UK

Abstract

This paper explores generalised probabilistic modelling and uncertainty estimation in comparative LLM-as-a-judge frameworks. We show that existing Product-of-Experts methods are specific cases of a broader framework, enabling diverse modelling options. Furthermore, we propose improved uncertainty estimates for individual comparisons, enabling more efficient selection and achieving strong performance with fewer evaluations. We also introduce a method for estimating overall ranking uncertainty. Finally, we demonstrate that combining absolute and comparative scoring improves performance. Experiments show that the specific expert model has a limited impact on final rankings but our proposed uncertainty estimates, especially the probability of reordering, significantly improve the efficiency of systems reducing the number of needed comparisons by $\sim 50\%$. Furthermore, ranking-level uncertainty metrics can be used to identify low-performing predictions, where the nature of the probabilistic model has a notable impact on the quality of the overall uncertainty.

1 INTRODUCTION

Instruction-tuned Large Language Models (LLMs) have shown impressive zero-shot performance on a wide range of natural language processing and generation tasks [Ouyang et al., 2022, Zhou et al., 2023, Chung et al., 2024]. While the number of downstream applications of aligned LLMs increases [Brown et al., 2020, Achiam et al., 2023, Dubey et al., 2024], so does the need to evaluate their performance on bespoke tasks, which could lack labelled data or are costly for humans to judge at scale [Zheng et al., 2023, Wang et al., 2022, Taori et al., 2023, Fathullah et al., 2024, Lakomkin et al., 2024]. As an alternative, instruction-tuned

LLMs have increasingly been used as a replacement for humans to evaluate the quality of natural language generations that demonstrate high correlations with human judgments [Zheng et al., 2023, Liusie et al., 2024b, Bubeck et al., 2023, Wang et al., 2023b, Chiang and Lee, 2023].

There are two standard approaches to using LLM in judging responses. **Absolute scoring:** Prompt an LLM to evaluate a certain response attribute on a defined scale (e.g., 1 to 10). **Comparative scoring:** Prompt an LLM to choose which of the two responses to a given query displays higher quality for a given attribute. Absolute scoring is a straightforward and effective method for evaluating a variety of responses to a query and ranking them. However, the scores obtained are unreliable and may vary significantly between different LLM judges. Alternatively, the more expensive comparative scoring approach has consistently demonstrated higher correlations with human judgements [Zheng et al., 2023, Liusie et al., 2024b, Qin et al., 2023]. However, a drawback of this approach is that it scales quadratically with the number of response candidates, which can become prohibitively expensive due to the inference costs of LLMs.

To address the computational limitations of comparative assessment, various approaches have been proposed. Notably, it is possible to extract more information from the LLM-as-a-judge than just a binary decision. Various works have, as opposed to using simple win-ratio, resorted to using the average probability output from the LLM [Qin et al., 2023, Zheng et al., 2023, Liusie et al., 2024b, Park et al., 2024]. Building on this idea, Liusie et al. [2024c] introduced a Product-of-Experts (PoE) [Hinton, 1999, Welling, 2007] framework in modelling comparative scoring. In principle, the joint distribution of candidate scores can be broken down into arbitrarily chosen experts that model the score differences of two instances at a time, allowing for a partial set of comparisons to model the full joint distribution. This directly allows one to obtain a ranking of candidates without having to perform all possible comparisons, showing that only a fraction of the total number of comparisons is needed to obtain highly competitive performance.

Contributions: In this paper, we generalise the expert in the comparative framework and derive how there is a wide range of viable options. Starting from the Beta distribution in modelling LLM probabilities, we show that prior works are specific instances of this choice. We also propose improved estimates for the uncertainty in individual comparisons and show how these updated uncertainties allow us to make even fewer comparisons without loss of performance. In addition to the uncertainty of a comparison, we also propose an uncertainty in the overall ranking of a set of candidates. Finally, we show that the Product-of-Experts framework easily lends itself to combinations of various scoring approaches. Specifically, we show that absolute scoring can complement comparative scoring efficiently and cheaply and lead to improved overall performance.

2 BACKGROUND AND RELATED WORK

NLG Evaluation with LLMs: The extensive natural language generation (NLG) capabilities of instruction-tuned large language models [Achiam et al., 2023, Ouyang et al., 2022, Chung et al., 2024, Dubey et al., 2024] have prompted recent work on open-ended generation evaluation using LLMs. Methods such as GPTScore [Fu et al., 2023] rank responses based on the likelihood of generation and G-Eval [Liu et al., 2023] which uses chain-of-thought and form-filling to evaluate the quality of a response. Furthermore, LLM-as-a-judge [Zheng et al., 2023] approaches score responses on an absolute scale [Wang et al., 2023a, Kocmi and Federmann, 2023] or comparative manner by comparing responses against each other [Qin et al., 2023, Liusie et al., 2024b,c] and building an overall ranking through the set of pairwise comparisons.

LLM-Based Comparative Assessment: The work by Liusie et al. [2024b] showed that comparative assessment yields superior performance compared to absolute scoring methods and various custom baselines. By making all possible $N(N-1)$ pairwise comparisons of N candidate responses, and computing the win-ratio, an overall ranking can be obtained. This style of approach has found its applications in many places. Qin et al. [2023] utilised pairwise comparisons to retrieve relevant sources, using both the full set of comparisons and sorting-based algorithms. Park et al. [2024] employed comparative assessment for dialogue evaluation, calculating the average probability across a randomly sampled set of comparisons to determine score quality. Finally, Liu et al. [2024b] demonstrated the limitations of LLM scoring, and resorted to using pairwise comparisons. They introduced PAirwise-preference Search, a variation of the merge sort algorithm which utilises LLM probabilities.

Ranking from Pairwise Comparisons: The idea of generating a full ranking from pairwise comparisons has been extensively studied. Arguably the most well-known example of this is the ranking of tennis players based on the outcomes

of games and this kind of problem has applications in many different areas. Anything from sports and gaming [Beaudoin and Swartz, 2018, Csató, 2013], web search [Dwork et al., 2001], social studies [Manski, 1977, Louviere et al., 2000] to information retrieval [Cao et al., 2007, Liu et al., 2009] requires modelling through pairwise events. The most common approach to model pairwise comparisons is the Bradley-Terry (BT) model [Bradley and Terry, 1952]. By assigning each candidate a latent score, the probability that one candidate wins over another is based on the underlying score difference. The latent scores can then be obtained by maximising the log-likelihood of the data [David, 1963, Cattelan, 2012]. Finally, the TrueSkill model [Herbrich et al., 2006, Minka et al., 2018] generalises the Bradley-Terry model by incorporating uncertainties in candidate scores within a Bayesian framework in a sports context.

Product-of-Experts: Each comparison $\mathcal{C}_k = (i, j, p_{ij})$ contains the ids of the candidates being compared and the corresponding probability produced by the LLM that it think i is better than j for a certain attribute. The comparison \mathcal{C}_k between a pair of candidates then provides information about the distribution of scores $\mathbf{s} = s_{1:N}$ of the candidates.

The Product-of-Experts (PoE) [Hinton, 1999] approach presents a simple and effective way to combine the information from multiple comparisons according to:

$$p(\mathbf{s}|\mathcal{C}_{1:K}) \propto \prod_k p(\mathbf{s}|\mathcal{C}_k) = \prod_k p(s_i - s_j|\mathcal{C}_k)$$

The distribution can be simplified into a product of individual experts, and be further simplified as it involves comparisons made exclusively between pairs of candidates. The work of Liusie et al. [2024c] proposed the soft Bradley-Terry expert as an extension to the standard BT model:

$$p(s_i - s_j|\mathcal{C}_k) \propto \sigma(s_i - s_j)^{p_{ij}} (1 - \sigma(s_i - s_j))^{1-p_{ij}}$$

where the probability p_{ij} is obtained from the LLM when comparing candidate i against j . The more experts/comparisons are included, the better the resulting estimate of the scores should be. Scores can then be retrieved by optimising $p(\mathbf{s}|\mathcal{C}_{1:K})$ using iterative [Zermelo, 1929, Newman, 2023] or standard gradient-based approaches.

3 GENERALISED EXPERT MODELLING

This section will focus on the nature of the expert for modelling both absolute and pairwise comparisons. We generally have N candidates with associated scores $s_{1:N}$. Given a partial set of comparisons $\mathcal{C}_{1:K}$, the aim is to predict a set of scores $\hat{s}_{1:N}$ which ranks the candidates as closely as possible to the true ranking. The following section will then focus on how to find an optimal set of comparisons $\mathcal{C}_{1:K}$ through uncertainty estimation, aiming to achieve better correlation with the true scores using fewer comparisons.

3.1 COMPARATIVE EXPERT MODELLING

The experts in prior work all modelled the score difference of candidates $p(s_i - s_j | \mathcal{C}_k)$. Through a simple change of variables we propose a more generalised version of the expert in comparative modelling:

$$p(s_i - s_j | \mathcal{C}_k) = f'(s_i - s_j) \tilde{p}(f(s_i - s_j) | \mathcal{C}_k)$$

where the $f(\cdot)$ is a generic monotonically increasing function. Any choice of f is viable as long as the distribution \tilde{p} supports it. By letting f be the sigmoid function and using an underlying Beta distribution $\tilde{p}(f | \mathcal{C}_k) = \mathcal{B}(f; p_{ij}, 1 - p_{ij})$ we regain the soft Bradley-Terry model from Liusie et al. [2024c]. However, from this point, it is clear that there are many more viable options for modelling a pairwise event. In this work, we will investigate several other combinations starting with a sigmoid and a general:

$$\tilde{p}(f | \mathcal{C}_k) = \mathcal{B}(f; p_{ij} + \alpha, 1 - p_{ij} + \beta) \quad (1)$$

model. We will also try an unconventional choice by combining the Gaussian distribution with a sigmoid-like function:

$$p(s_i - s_j | \mathcal{C}_k) = \mathcal{N}(s_i - s_j; 0, 1) \mathcal{N}(\Phi(s_i - s_j); p_{ij}, 1) \quad (2)$$

where $\Phi(\cdot)$ is the cumulative density of the Gaussian. Abolitions will study the impact of the choice of f and the underlying distribution \tilde{p} .

3.2 EXPERT COMBINATIONS

As has been mentioned several times, absolute scoring is cheaper but worse than comparative scoring. However, it is possible that absolute scoring can provide complementary information so we propose combining the two approaches into a single model:

$$p(s | \mathcal{C}_{1:K}, \mathcal{A}_{1:N}) \propto \prod_k p(s_i - s_j | \mathcal{C}_k) \prod_n p(s_n | \mathcal{A}_n)$$

where \mathcal{A}_n contains the information from an absolute scoring prompt. When prompting an LLM on a scale of 1 to 10 it contains the probabilities of those values. Unfortunately, absolute scoring provides a discrete score and while it is possible to obtain the associated logits for each value (of 1 to 10) from the LLM (and construct a categorical distribution) it remains difficult to combine the continuous pairwise experts with the discrete absolute experts. Therefore, we also propose using moment matching to transform the categorical expert into a Gaussian $\mathcal{N}(\mu, \sigma^2)$:

$$\mu = \sum_c c p_c, \quad \sigma^2 = \sum_c (c - \mu)^2 p_c$$

where c and p_c represent the class and the associated probability that the LLM would output that class. In our running

example, we would have $c \in \{1, \dots, 10\}$. Finally, the absolute expert can be written as:

$$p(s_n | \mathcal{A}_n) \approx \mathcal{N}(s_n; \mu_n, \sigma_n^2), \quad \forall n = 1, \dots, N$$

and would allow us to operate with continuous values and optimise the skill scores using gradient-based approaches. In this work, we rely solely on simple absolute scoring but there exist many more sophisticated pointwise scoring approaches like G-Eval [Liu et al., 2023] which can provide further improvements.

3.3 HOME ADVANTAGE

A big issue plaguing LLM-based approaches is bias in the system. In our case, the probability outputs of an LLM are inconsistent $p_{ij} \neq 1 - p_{ji}$, meaning that the LLM-based judge can assign conflicting probabilities when comparing i to j as opposed to j to i . This stems from positional biases in the system [Zheng et al., 2023, Chen et al., 2024a, Liusie et al., 2024a, Wang et al., 2023c, Zhu et al., 2023, Chen et al., 2024b, Liu et al., 2024a]. To resolve such an issue, we rely on one of the two approaches. **Permutation Debiasing:** For each comparison we make two LLM calls for both i vs j and j vs i to obtain a final debiased probability $\tilde{p}_{ij} = \frac{1}{2}(p_{ij} + (1 - p_{ji}))$ which would ensure consistency: $\tilde{p}_{ij} = (1 - \tilde{p}_{ji})$. **Home Advantage:** An alternative approach is to directly incorporate the positional bias into the comparative expert model. Since we already know that a certain position will be preferred over another we can introduce a 'home advantage' [Agresti, 1990, Caron and Doucet, 2012] parameter to model the inconsistency through our function f :

$$f(s_i - s_j; \Delta) = f(s_i - s_j - \Delta)$$

While prior approaches have developed the theory for home advantage in specific use-cases such as Bradley-Terry [Caron and Doucet, 2012] and Gaussian experts [Liusie et al., 2024c], our parameterisation through the generic function f allows us to straightforwardly incorporate home advantage into any type of expert. Furthermore, while the work of Liusie et al. [2024c] estimated the advantage parameter Δ through bespoke rules for each expert, we estimate it by maximising the likelihood $p(s | \mathcal{C}_{1:K}, \Delta)$.

4 UNCERTAINTY ESTIMATION

This section will explore how to estimate uncertainty when ranking examples. Two levels of uncertainty will be explored. **Pairwise uncertainty:** The uncertainty in the score difference of a pair of candidates. By being able to identify the most uncertain pairs, we can understand what comparisons we should perform to improve the performance. **Ranking uncertainty:** The uncertainty in the overall ranking of a set of candidates, which can help us in understanding whether the predicted ranking is trustworthy or not.

4.1 LAPLACE’S APPROXIMATION

Due to the complex nature of the Product-of-Experts (PoE) distribution, $p(s|\mathcal{C}_{1:K})$, deriving analytical expressions for uncertainties is generally intractable. Therefore, following practice in Bayesian inference with complex models, we rely on Laplace’s approximation to estimate the posterior distribution and its associated uncertainties. The approach approximates the distribution $p(s|\mathcal{C}_{1:K})$ with a Gaussian:

$$p(s|\mathcal{C}_{1:K}) \approx \mathcal{N}\left(s; \boldsymbol{\mu}^{(K)}, \boldsymbol{\Sigma}^{(K)}\right)$$

where we set (and dropped the superscript for conciseness):

$$\begin{aligned} \boldsymbol{\mu} &= \arg \max_s \ln p(s|\mathcal{C}_{1:K}) \\ \boldsymbol{\Sigma}^{-1} &= -\nabla \nabla \ln p(s|\mathcal{C}_{1:K})|_{\boldsymbol{\mu}} \end{aligned}$$

There are more advanced approaches to approximating an intractable distribution but we will rely on this simple and efficient scheme in this work.

4.2 PAIRWISE UNCERTAINTY ESTIMATION

Being able to estimate the uncertainty in a pair of candidates, in their score difference, can allow us to decide which comparisons are useful and which are not. The better the quality of the uncertainty estimate, the fewer comparisons are needed to achieve good performance. Following Liusie et al. [2024c], the aim is to iteratively add additional comparisons to improve the product-of-experts model and the overall ranking of scores. Starting with a unit Gaussian expert/prior on every score we:

1. **Estimate Uncertainty:** Compute the covariance matrix $\boldsymbol{\Sigma}$ based on the currently selected set of comparisons $\mathcal{C}_{1:K}$ using Laplace approximation.
2. **Select Next Comparison:** Use an uncertainty metric and the current covariance matrix $\boldsymbol{\Sigma}$ to select the most informative comparison pair (i, j) . In large-scale tasks one can select a batch of comparisons.
3. **Perform Comparison:** Obtain the LLM probability for the selected pair (i, j) and add this new comparison to the set $\mathcal{C}_{1:K+1}$.
4. **Update Model:** Update the Product-of-Experts model to incorporate the new comparison, which implicitly updates the mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ for the next iteration.

Prior work posited that the next comparison that should be selected should induce the **minimum overall uncertainty** in the resulting distribution, giving the following selection criteria under a soft BT model:

$$\arg \max_{i,j} \sigma(\mu_i - \mu_j) \sigma(\mu_j - \mu_i) (\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}) \quad (3)$$

We propose and evaluate two additional pairwise uncertainty metrics which apply to any expert modelling choice. **Variance in Score Difference:** A straightforward measure of uncertainty in a pair of candidates i and j is the variance of their score difference under the approximated Gaussian posterior. Given the covariance matrix $\boldsymbol{\Sigma}$, the variance of the score difference $s_i - s_j$ can be directly computed as $\mathbb{V}[s_i - s_j|\mathcal{C}_{1:K}] = \Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}$. To select the next comparison, we aim to maximise this variance, choosing the pair (i, j) for which the model is most uncertain about their score difference:

$$\arg \max_{i,j} \Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj} \quad (4)$$

Probability of Reordering: While variance captures the general uncertainty in score difference, a potentially more relevant metric for ranking is the probability that the ranking between two candidates is reversed. Assuming that the current model predicts candidate i to be better than candidate j (i.e., $\mu_i > \mu_j$), the probability of reordering, i.e., the probability that candidate j is actually better than candidate i given the current comparisons $\mathcal{C}_{1:K}$, can be calculated using the Gaussian cumulative density:

$$P(s_i < s_j|\mathcal{C}_{1:K}) = \Phi\left(\frac{\mu_j - \mu_i}{\sqrt{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}}\right)$$

This metric directly quantifies the likelihood of a rank inversion between the pair. As shown in the Appendix C, this can be simplified to a similar selection form since we want to pick examples with the highest reordering probability:

$$\arg \max_{i,j} \frac{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}{(\mu_i - \mu_j)^2} \quad (5)$$

This probability of reordering metric is intuitively appealing as it directly targets the goal of improving ranking accuracy by focusing on comparisons that are most likely to correct potential ranking errors. Furthermore, it naturally encourages the exploration of comparisons between candidates with similar skill levels (high-density regions of the score distribution), as these are the pairs where rank inversions are most probable. Furthermore, since all selection mechanisms take the following form:

$$\arg \max_{i,j} w(\mu_i - \mu_j) (\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj})$$

we ablate the weight function $w(\cdot)$.

4.3 RANKING UNCERTAINTY ESTIMATION

In addition to pairwise uncertainties, it is also valuable to assess the overall uncertainty in the predicted ranking of all candidates. While a full probabilistic ranking distribution is complex to compute, we can use the entropy of the approximated Gaussian score distribution as a

proxy for ranking uncertainty. The entropy $\mathcal{H}[s|\mathcal{C}_{1:K}]$ of a multivariate Gaussian distribution $\mathcal{N}(s; \mu, \Sigma)$ is given by: $N(1 + \ln(2\pi))/2 + \ln(\det(\Sigma))/2$. Lower entropy indicates a more concentrated and certain score distribution, suggesting a more reliable overall ranking. We will use this entropy metric to assess the overall uncertainty of the rankings produced by various models.

5 EXPERIMENTAL SETUP

5.1 DATASETS

We mainly perform experiments on the summary evaluation SummEval dataset Fabbri et al. [2021] which contains 100 articles, each with 16 machine-generated summaries evaluated on four different attributes: coherency (COH), consistency (CON), fluency (FLU), and relevancy (REL). We will also use the much larger HANNA dataset Chhun et al. [2022] which has 1056 machine-generated stories annotated by humans on six different attributes. These are averaged to a single overall quality score.

5.2 METHODOLOGY

We will be relying on Flan-T5 Chung et al. [2024] and Qwen2.5-Instruct Qwen Team [2024] systems to evaluate performance on the SummEval and the larger HANNA datasets. Appendix A will detail our choices, how we structure the prompts and how the probabilities are extracted from each model.

Probabilistic Models: In almost all experiments we will rely on the soft Bradley-Terry extension as our baseline expert model. This will be compared against our proposed extensions to this approach: (1) The generalised Beta distribution in Eq. (1), (2) the extended Gaussian distribution in Eq. (2) and (3) the combination of comparative and absolute outputs in a single PoE model. We will not include simple baselines such as average win-ratio and average probability since these have been shown to be inferior on a wide range of tasks [Liusie et al., 2024c, Raina et al., 2024].

Iterative Selection: We will also compare the probabilistic models in an active learning framework where each model needs to select the comparisons that will induce the best performance. The baseline will be the **minimum uncertainty** approach given in Eq. (3) using the soft BT model. We will compare this against our proposed **variance** (Eq. (4)) and **probability of reordering** (Eq. (5)) selection mechanisms.

Ranking Uncertainties: The entropy of Product-of-Experts models will be investigated on how well they correlate with the actual performance of the predicted rankings, and whether they can identify high-performing predictions.

Evaluation Metrics: Since we are interested in predicting

the ranking of candidate responses given a context/query, the main performance metric is Spearman rank correlation between the predicted and the human labelled scores. In SummEval ($N = 16$) we perform absolute and comparative scoring and evaluate the average Spearman across all contexts. For HANNA ($N = 1056$) we rank all generated stories. Furthermore, we assess the quality of the various comparison-level uncertainties (iterative selection schemes) by the number of comparisons needed to achieve good Spearman rank correlation. This will be defined as the number of comparisons K required to reach within 90% performance when using the full set of $N(N - 1)$ comparisons. Finally, the ranking-level uncertainties are evaluated using the area under the receiver operating characteristic curve (AUROC) to detect well-performing rankings.

6 RESULTS

6.1 FORM OF DISTRIBUTION

This section investigates a wide range of PoE models conditioned on the full set of comparisons. For SummEval with $N = 16$ summaries per context, there is a total of $N(N - 1) = 240$ number of comparisons. Unless reported otherwise, all results are based on the direct biased outputs of judges. In the first block of Table 1, under a Flan-T5 judge, models showcase similar Spearman performance when evaluated on the full set of comparisons. The only model that performs noticeably better is the combined comparative-absolute model. Absolute scoring seems to extract complementary information and give the overall system a performance boost. Furthermore, improved and larger Qwen LLMs clearly display improved performance.

However, improved performance does not necessarily imply higher quality ranking-level uncertainty estimates. To benchmark, an uncertainty is predicted for each context, meaning to represent how well the predicted ranks perform. The Spearman rank performance of each context is thresholded by the median score and mapped to a binary value so that contexts are classed as '0' or '1'. This allows us to use the standard AUROC score to evaluate detection performance. The second block of Table 1 shows the AUROC performance of a range of probabilistic models. Unlike previous results, the nature of the probabilistic model seems to have a significant impact on performance. While the performance of various PoE models are similar, the predicted entropy, and by equivalence, the hessian (curvature) of the log-likelihood seems to differ between certain models. This seems to stem mainly from the choice of f but further investigations are needed to understand how to predict robust uncertainties.

The linear-Gaussian model performs significantly worse, only marginally outperforming a random classifier. Furthermore, all models that map skill differences through a sigmoid-like function f perform notably better, with the

Table 1: Spearman rank correlations (%) and AUROC (%) for SummEval.

LLM	Function f	Distribution \tilde{p}	Spearman Rank (%)					AUROC (%)				
			COH	CON	FLU	REL	Avg	COH	CON	FLU	REL	Avg
Flan-T5 (3B)	x	Gaussian	49.1	45.2	32.5	42.2	42.3	53.4	51.8	54.9	58.4	54.6
	σ	Gaussian	49.2	45.3	32.5	42.2	42.3	65.5	67.3	60.8	65.8	64.9
	Φ	Gaussian	49.2	45.3	32.5	42.3	42.3	65.0	67.3	61.0	65.9	64.8
	σ	Beta	49.2	45.2	32.5	42.2	42.3	63.8	68.3	60.6	63.8	64.1
	Φ	Beta	49.2	45.3	32.6	42.3	42.4	64.8	67.4	60.9	65.7	64.7
	Comparative (σ -Beta) + Absolute Experts			50.7	45.9	32.9	43.4	43.2	64.5	69.9	62.0	67.0
Qwen2.5 (3B)	σ	Beta	48.5	48.6	40.8	46.5	46.1	59.8	62.8	63.2	57.5	60.8
Qwen2.5 (7B)	σ	Beta	50.4	48.9	38.7	50.5	47.1	66.7	68.2	58.5	69.4	65.7
Qwen2.5 (14B)	σ	Beta	57.4	52.1	47.4	51.1	52.0	61.7	62.0	58.4	64.8	61.7

combined model being the best system similar to previous results. Even the better-performing and larger Qwen systems are unable to outperform the combined model. Furthermore, while Qwen2.5 (3B) displays better Spearman rank than the similarly sized Flan-T5, the uncertainties perform worse. While this is only a preliminary investigation into detecting well-performing examples, it is a good start into understanding the problem and what models produce better uncertainties. In Appendix B, we investigate the causes for this and pinpoint some of the problems to an overconfidence issue associated with modern language models.

6.2 ITERATIVE SELECTION

In this section, we explore the quality of various models and uncertainty metrics when iteratively selecting comparisons. Furthermore, this investigation is performed on both biased and debiased probabilities extracted from Flan-T5 (3B). Extended results for Qwen can be found in Appendix D. Four main points can be observed from the results in Figure 1:

(1) Model Convergence: As reported in the table above, all models converge towards the same final performance since the rankings are predicted from the same full set of

comparisons.

(2) Quality of Uncertainties: There is a significant gap in performance between various uncertainty metrics. Probability of reordering is shown to outperform the minimum uncertainty metric in all attributes of SummEval for both biased and debiased cases. We expect this performance difference to originate from how each metric was derived. While minimum uncertainty is simply focused on choosing the next comparison that would induce the least uncertainty, the probability of reordering is directly linked to achieving the correct ranking.

(3) Expert Model Invariance: Probability of reordering under two different models, σ -Beta (soft BT) and Φ -Gaussian (ngaussian), perform almost identically in all cases. This again reinforces the idea that the nature of the expert model does not matter as much as the uncertainty modelling used in selecting the comparisons.

(4) Performing better with less: Performance is expected to increase as one adds more and more comparisons to the PoE model. However, in many of the cases above, performance drops until the full set of comparisons is reached. This is related to an overconfidence issue plaguing the uncertainty

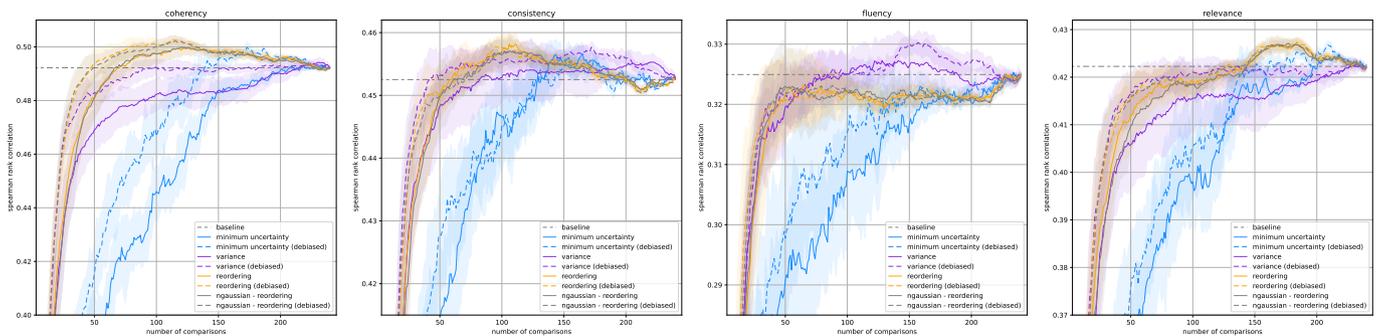


Figure 1: The Spearman Rank Correlation when iteratively selecting the next examples of lowest confidence/highest uncertainty. The baseline is the soft Bradley-Terry model with the minimum uncertainty metric. We also report the proposed variance and probability of reordering under the soft BT model. Furthermore, the Φ -Gaussian model is referred to as "ngaussian". Debiased refers to permutation debiased probabilities.

Table 2: **Measure of efficiency:** The number of comparisons needed to achieve 90% performance of the baseline soft Bradley-Terry system for each corresponding judge. The baseline is evaluated using all $N(N - 1) = 240$ comparisons.

Uncertainties	Debiased	Flan-T5 (3B)				Qwen2.5 (3B)			
		COH	CON	FLU	REL	COH	CON	FLU	REL
minimum uncertainty	✗	82.8 ± 9.8	39.4 ± 6.2	30.3 ± 9.6	61.4 ± 8.7	76.9 ± 8.4	39.9 ± 3.2	45.7 ± 6.3	71.2 ± 7.3
	✓	59.6 ± 8.7	27.1 ± 7.1	23.9 ± 6.1	46.8 ± 4.6	57.9 ± 8.3	35.9 ± 4.2	43.7 ± 6.8	63.1 ± 7.9
variance	✗	29.7 ± 6.9	17.8 ± 2.7	13.8 ± 7.6	27.5 ± 6.7	44.3 ± 7.1	24.6 ± 3.6	27.2 ± 6.6	53.9 ± 8.0
	✓	20.3 ± 2.0	14.3 ± 2.5	16.7 ± 2.3	19.7 ± 1.7	24.4 ± 4.0	16.7 ± 2.8	21.8 ± 4.2	24.5 ± 3.1
probability of reordering	✗	28.1 ± 3.4	17.7 ± 1.9	18.8 ± 8.1	22.2 ± 4.2	29.0 ± 4.5	23.3 ± 2.2	24.6 ± 2.4	34.8 ± 5.6
	✓	20.4 ± 1.7	14.2 ± 2.3	16.1 ± 3.9	18.6 ± 2.1	20.7 ± 1.8	16.1 ± 2.6	18.9 ± 3.1	24.3 ± 4.9

estimates. We show in Appendix B how the LLM-as-a-judge is miscalibrated, and how temperature annealing is not enough to calibrate and solve the overconfidence issue.

Furthermore, in Table 2, we report the number of examples required to reach 90% of the final performance when all comparisons have been selected. The baseline ‘minimum uncertainty’ metric with both biased and debiased probabilities extracted from the LLMs manages to reach the 90% threshold with at worst 35% of the number of comparisons. This already represents a significant gain in efficiency, cutting the number of examples to a fraction of the full 240. However, both proposed uncertainty metrics significantly outperform the baseline, reducing the number of required comparisons by an additional ~60% for Flan-T5 and ~50% for Qwen. Overall, we observe that the reordering metric is more efficient across datasets and models.

Finally, in Figure 2a, we compare the soft BT model with the combined comparative-absolute model. Due to the initial cost of obtaining $N = 16$ absolute experts, the combined model initially performs worse. However, both biased and debiased combined models outperform the final soft BT model performance using a fraction of comparisons.

6.2.1 Ablation Studies

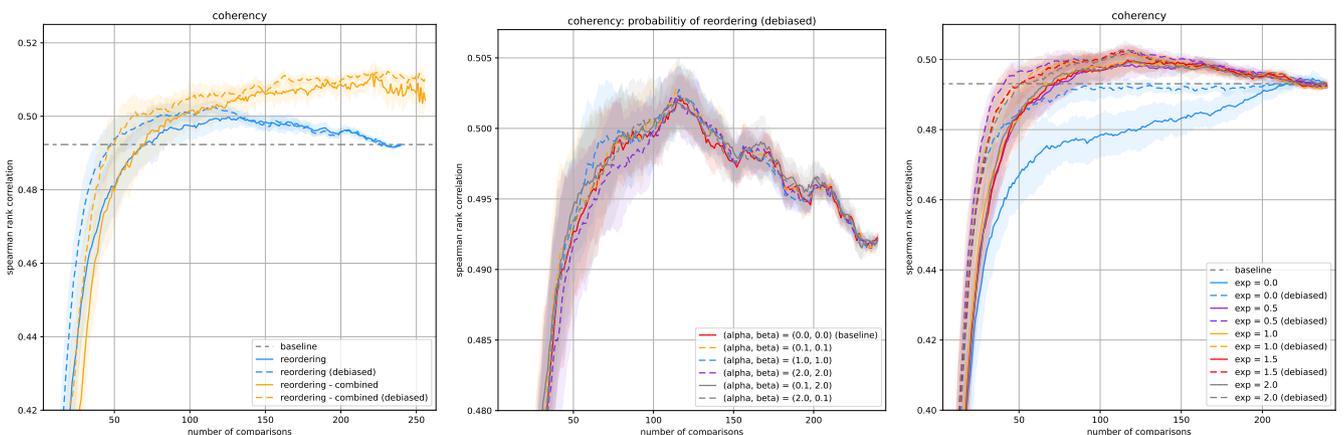
This section will explore various nuances in the modelling choices. In Figure 2b, we vary the parameters of the underlying Beta distribution in a generalised soft BT model.

Evaluated on SummEval (COH), it is clear that the underlying Beta distribution has negligible impact on both the selection process and the final performance. Furthermore, we explore generalising the selection metrics with the following:

$$\arg \max_{i,j} \frac{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}{|s_i - s_j|^\epsilon} \quad (6)$$

where we vary the exponent ϵ . Setting $\epsilon = 0$ returns variance while $\epsilon = 2$ gives probability of reordering.

The exponent is swept on the same benchmark in Figure 2c. It is evident that while $\epsilon = 0$ suffers in performance, most of the other values perform similarly. The best-performing option is $\epsilon = 0.5$ which slightly outperforms other options including the probability of reordering.

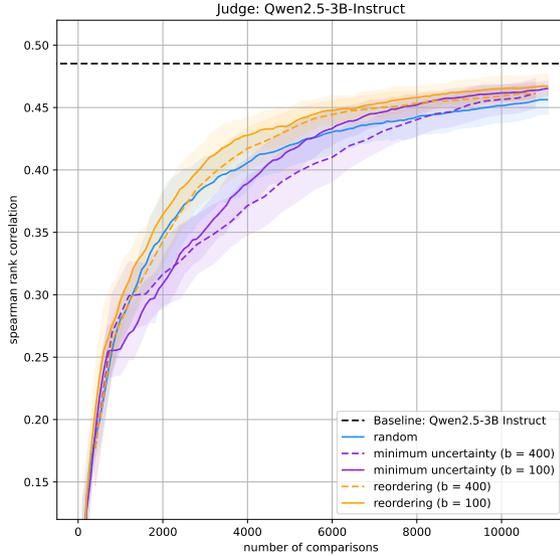


(a) Absolute and comparative combination.

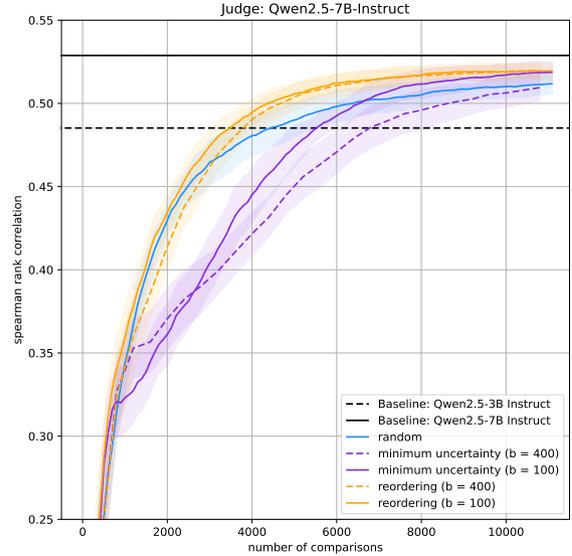
(b) Varying the Beta distribution.

(c) Varying the uncertainty estimate.

Figure 2: Extended results on Flan-T5: **(a)** Comparing the soft BT model to the combined model which relies on additional absolute experts. Both use the probability of reordering as the selection criteria. **(b)** Selection performance. Varying the parameters α and β of the Beta distribution in Eq. (1). **(c)** Selection performance. Varying the parameter ϵ in Eq. (6).



(a) Using Qwen2.5-3B-Instruct as the judge.



(b) Using Qwen2.5-7B-Instruct as the judge.

Figure 3: The Spearman Rank Correlation when iteratively selecting the next batch b of examples of lowest confidence/highest uncertainty. The baselines are Qwen2.5- $\{3B-7B\}$ models with all comparisons selected. Furthermore, the efficient baseline is set to the soft BT model with the minimum uncertainty metric.

6.3 LARGE-SCALE SELECTION

All results have been focused on ranking a small set of candidates $N = 16$ which has $K = 240$ possible comparisons. In this section we scale up to the HANNA dataset with $N = 1056$ stories and $K = 1114080$ possible comparisons. Furthermore, due to the cost of iteratively selecting a single comparison, re-estimating Laplace’s approximation, and repeating this process, we opt to perform batch acquisitions with $b = \{100, 400\}$. This will showcase in the extreme case how well our best-proposed uncertainty metric, probability of reordering, performs compared to random and minimum uncertainty selection. We will only evaluate the iterative process up until 1% of all possible comparisons to test the efficiency of approaches.

In Figure 3, various iterative schemes are reported. Furthermore, the performance of the soft BT model with all possible comparisons is reported under both backbone judges. In these results, one can observe minimum uncertainty to suffer significantly compared to the simplest baseline of random selection when using both Qwen2.5-3B and 7B as backbone judges. This is caused due to the lack of diversity when selecting a batch of comparisons. While reducing the batch size of acquisitions helps performance, it still lacks significantly compared to the probability of reordering which is far more robust to larger batch acquisitions.

7 CONCLUSIONS

This paper generalised probabilistic modelling for comparative LLM-as-a-judge, demonstrating that existing approaches

are specific instances of a broader framework. We introduced improved uncertainty estimates for individual comparisons and overall rankings, leading to more efficient iterative selection strategies. Notably, the probability of reordering proved to be a superior metric for selecting informative comparisons. We also showed the benefits of combining absolute and comparative scoring within a Product-of-Experts framework, achieving enhanced performance. While the specific expert model had limited impact on final rankings given sufficient comparisons, the choice of uncertainty estimation and the incorporation of absolute scoring significantly improved efficiency and accuracy. Our findings highlight the importance of robust uncertainty estimation in LLM-based evaluation and provide a more flexible and efficient framework for comparative assessment.

8 LIMITATIONS

The main concern lies in the quality of the estimated uncertainties, which are crucial for the efficiency of the proposed iterative selection methods. The reliance on Laplace’s approximation to derive these uncertainties introduces potential inaccuracies. This approximation assumes that the posterior distribution over model parameters is approximately Gaussian, which may not hold true in all scenarios, particularly when the true posterior is multimodal or exhibits significant skewness. Consequently, the derived uncertainty metrics, such as the variance and probability of reordering, might not perfectly reflect the true uncertainty in the model’s predictions.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Alan Agresti. *Categorical data analysis*. John Wiley & Sons, 1990.
- David Beaudoin and Tim Swartz. A computationally intensive ranking system for paired comparison data. *Operations Research Perspectives*, 5:105–112, 2018.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.
- Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. Learning to rank: from pairwise approach to list-wise approach. In *Proceedings of the 24th international conference on Machine learning*, pages 129–136, 2007.
- Francois Caron and Arnaud Doucet. Efficient bayesian inference for generalized bradley–terry models. *Journal of Computational and Graphical Statistics*, 21(1):174–196, 2012.
- Manuela Cattelan. Models for paired comparison data: A review with emphasis on dependent data. 2012.
- Guiming Hardy Chen, Shunian Chen, Ziche Liu, Feng Jiang, and Benyou Wang. Humans or llms as the judge? a study on judgement biases, 2024a.
- Xinyun Chen, Ryan A. Chi, Xuezhi Wang, and Denny Zhou. Premise order matters in reasoning with large language models, 2024b. URL <https://arxiv.org/abs/2402.08939>.
- Cyril Chhun, Pierre Colombo, Fabian Suchanek, and Chloé Clavel. Of human criteria and automatic metrics: A benchmark of the evaluation of story generation. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5794–5836, 2022.
- Cheng-Han Chiang and Hung-yi Lee. Can large language models be an alternative to human evaluations? *arXiv preprint arXiv:2305.01937*, 2023.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 2024.
- László Csató. Ranking by pairwise comparisons for swiss-system tournaments. *Central European Journal of Operations Research*, 21:783–803, 2013.
- Herbert Aron David. *The method of paired comparisons*, volume 12. London, 1963.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th international conference on World Wide Web*, pages 613–622, 2001.
- Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409, 2021.
- Yassir Fathullah and Mark Gales. Self-distribution distillation: Efficient uncertainty estimation. In *Proceedings of the 38th Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 180 of *Proceedings of Machine Learning Research*, pages 655–665. PMLR, 2022.
- Yassir Fathullah, Guoxuan Xia, and Mark J. F. Gales. Logit-based ensemble distribution distillation for robust autoregressive sequence uncertainties. In Robin J. Evans and Ilya Shpitser, editors, *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pages 582–591. PMLR, 31 Jul–04 Aug 2023. URL <https://proceedings.mlr.press/v216/fathullah23a.html>.
- Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Ke Li, Junteng Jia, Yuan Shangguan, Jay Mahadeokar, Ozlem Kalinli, Christian Fuegen, and Mike Seltzer. AudioChatLlama: Towards general-purpose speech abilities for LLMs. In Kevin Duh, Helena Gomez, and Steven Bethard, editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for*

- Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5522–5532, Mexico City, Mexico, June 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.naacl-long.309. URL <https://aclanthology.org/2024.naacl-long.309/>.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*, 2023.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.
- Ralf Herbrich, Tom Minka, and Thore Graepel. Trueskill™: a bayesian skill rating system. *Advances in neural information processing systems*, 19, 2006.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *Neural Information Processing Systems 2014 Deep Learning and Representation Learning Workshop*, 2015. arXiv preprint arXiv:1503.02531.
- Geoffrey E. Hinton. Products of experts. In *Artificial Neural Networks, 1999. ICANN 99. Ninth International Conference on (Conf. Publ. No. 470)*, volume 1, pages 1–6. IET, 1999.
- Tom Kocmi and Christian Federmann. Large language models are state-of-the-art evaluators of translation quality. *arXiv preprint arXiv:2302.14520*, 2023.
- Egor Lakomkin, Chunyang Wu, Yassir Fathullah, Ozlem Kalinli, Michael L. Seltzer, and Christian Fuegen. End-to-end speech recognition contextualization with large language models. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12406–12410, 2024. doi: 10.1109/ICASSP48485.2024.10446898.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12:157–173, 2024a. doi: 10.1162/tacl_a_00638. URL <https://aclanthology.org/2024.tacl-1.9>.
- Tie-Yan Liu et al. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval*, 3(3): 225–331, 2009.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. G-eval: NLG evaluation using gpt-4 with better human alignment. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.153. URL <https://aclanthology.org/2023.emnlp-main.153>.
- Yinhong Liu, Han Zhou, Zhijiang Guo, Ehsan Shareghi, Ivan Vulić, Anna Korhonen, and Nigel Collier. Aligning with human judgement: The role of pairwise preference in large language model evaluators, 2024b.
- Adian Liusie, Yassir Fathullah, and Mark JF Gales. Teacher-student training for debiasing: General permutation debiasing for large language models. *arXiv preprint arXiv:2403.13590*, 2024a.
- Adian Liusie, Potsawee Manakul, and Mark Gales. LLM comparative assessment: Zero-shot NLG evaluation through pairwise comparisons using large language models. In Yvette Graham and Matthew Purver, editors, *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 139–151, St. Julian’s, Malta, March 2024b. Association for Computational Linguistics. URL <https://aclanthology.org/2024.eacl-long.8>.
- Adian Liusie, Vatsal Raina, Yassir Fathullah, and Mark Gales. Efficient llm comparative assessment: a product of experts framework for pairwise comparisons. *arXiv preprint arXiv:2405.05894*, 2024c.
- Jordan J Louviere, David A Hensher, and Joffre D Swait. *Stated choice methods: analysis and applications*. Cambridge university press, 2000.
- Andrey Malinin, Bruno Mlodozieniec, and Mark Gales. Ensemble distribution distillation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1xWdeEFvr>.
- Charles F Manski. The structure of random utility models. *Theory and decision*, 8(3):229, 1977.
- Tom Minka, Ryan Cleven, and Yordan Zaykov. Trueskill 2: An improved bayesian skill rating system. *Technical Report*, 2018.
- MEJ Newman. Efficient computation of rankings from pairwise comparisons. *Journal of Machine Learning Research*, 24(238):1–25, 2023.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

- ChaeHun Park, Minseok Choi, Dohyun Lee, and Jaegul Choo. Paireval: Open-domain dialogue evaluation with pairwise comparison. *arXiv preprint arXiv:2404.01015*, 2024.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, et al. Large language models are effective text rankers with pairwise ranking prompting. *arXiv preprint arXiv:2306.17563*, 2023.
- Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.
- Vatsal Raina, Adian Liusie, and Mark Gales. Finetuning llms for comparative assessment tasks, 2024. URL <https://arxiv.org/abs/2409.15979>.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- Jiaan Wang, Yunlong Liang, Fandong Meng, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. Is chatgpt a good nlg evaluator? a preliminary study. *arXiv preprint arXiv:2303.04048*, 2023a.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhi-fang Sui. Large language models are not fair evaluators, 2023b.
- Yiwei Wang, Yujun Cai, Muhao Chen, Yuxuan Liang, and Bryan Hooi. Primacy effect of ChatGPT. In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 108–115, Singapore, December 2023c. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.8. URL <https://aclanthology.org/2023.emnlp-main.8>.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisha Liu, Noah A Smith, Daniel Khashabi, and Hananeh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*, 2022.
- M. Welling. Product of experts. *Scholarpedia*, 2(10):3879, 2007. doi: 10.4249/scholarpedia.3879. revision #137078.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=83B55k9s0h>.
- Ernst Zermelo. Die berechnung der turnier-ergebnisse als ein maximumproblem der wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 29(1):436–460, 1929.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. Lima: Less is more for alignment. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 55006–55021. Curran Associates, Inc., 2023.
- Lianghui Zhu, Xinggang Wang, and Xinlong Wang. Judgelm: Fine-tuned large language models are scalable judges. 2023.

Generalised Probabilistic Modelling and Improved Uncertainty Estimation in Comparative LLM-as-a-judge (Supplementary Material)

Yassir Fathullah¹

Mark J. F. Gales¹

¹Engineering Department, University of Cambridge, UK

A PROMPTING

When prompting an LLM for the score of a candidate, or which of two candidates is better, the important information lies in the logits of the tokens we are interested in. In this section we detail the design of our prompts for both Flan-T5 and Qwen2.5. Since the former is an encoder-decoder foundation model and the latter is a decoder-only foundation model the prompts need to be designed slightly differently.

Absolute prompting: For the Flan-T5 system, we give the encoder the following prompt:

```
Article: <context>\n\nSummary: <A>\n\nScore the response between 1 and 10 based on how coherent the summary is.
```

where we are scoring the coherency of a summary. The `<context>` and `<A>` are replaced by the article and summary. Following this we extract the logits corresponding to 1 to 10 from the decoder. The probability of each class is then:

$$p_c = \frac{\exp(z_c)}{\sum_{i=1}^{10} \exp(z_i)} \quad c = 1, \dots, 10$$

The choice of 1-10 is arbitrary and any other range could have been chosen.

Comparative prompting: For the Flan-T5 system, we give the encoder the following prompt:

```
Article: <context>\n\nSummary A: <A>\n\nSummary B: <B>\n\nWhich Summary is more coherent, Summary A or Summary B?
```

The `<context>`, `<A>` and `` are replaced by the article and two different summaries. Following this we give the following prefix to the decoder:

Summary

and extract the logits corresponding to A and B from the decoder. The prefix ensures that the probability mass of the next token is concentrated into the options "A" and "B". From these logits we extract the probability that A will win:

$$p = \frac{\exp(z_A)}{\exp(z_A) + \exp(z_B)}$$

Similarly, we prompt the Qwen2.5 system in the following matter when we want to rank various stories from HANNA:

```

{"role": "system", "content": "You
are an expert story assessor."},

{"role": "user", "content": "Story A:
<A>\n\nStory B: <B>\n\nWhich story
is better overall, Story A or B?
Answer only with Story A or Story B."}

{"role": "assistant", "content": "Story "}

```

These are then prepared by the Qwen2.5 tokenizer in the instruction following format and fed into the model. Following on, the logits corresponding to *A* and *B* are then extracted for the next token and converted into a probability.

B CALIBRATION

This section reports the calibration error and reliability diagrams for the different metrics under a biased and debiased setup. The main point is to address the overconfidence issue related to our results in Section 6.2 and why temperature annealing is not enough to solve the problem.

The calibration is based on the confidence scores of individual comparisons $\max(p, 1 - p)$. Therefore, when calibrating using temperature annealing, the resulting (binary) predictions remain the same:

$$\tilde{p} = \frac{p^{1/T}}{p^{1/T} + (1 - p)^{1/T}}$$

To understand the impact of calibration, we find the optimal temperature on the SummEval dataset by minimising the expected calibration error, see Table 3. Each attribute has its optimal temperature. We also report the corresponding reliability

Table 3: Flan-T5 (3B): Expected calibration error (%).

Method	Debiased	COH	CON	FLU	REL
-	✗	9.80	3.77	9.87	11.86
	✓	2.83	1.82	4.84	6.20
Calibrated	✗	1.02	0.68	1.28	0.98
	✓	2.58	1.72	1.08	1.07

diagrams in Figure 4. From these results, it is evident that simple temperature annealing can almost entirely resolve the miscalibration in the systems.

This next part will check how temperature annealing affects the solution of a soft Bradley-Terry model. Starting from the gradient of the log-likelihood:

$$\nabla \ln p(\mathbf{s}|\mathcal{C}_{1:K}) = \sum_{i,j \in \mathcal{C}_{1:K}} p_{ij} - \sigma(s_i - s_j) = 0$$

Looking at a single element of the sum, and under calibrated probabilities the new solution becomes:

$$\begin{aligned}
\tilde{p}_{ij} &= \sigma(\tilde{s}_i - \tilde{s}_j) \iff \\
\frac{p^{1/T}}{p^{1/T} + (1-p)^{1/T}} &= \sigma(\tilde{s}_i - \tilde{s}_j) \iff \\
\frac{1}{1 + \left(\frac{1-p}{p}\right)^{1/T}} &= \sigma(\tilde{s}_i - \tilde{s}_j) \iff \\
\frac{1}{1 + \exp\left(\frac{1}{T} \ln\left(\frac{1-p}{p}\right)\right)} &= \sigma(\tilde{s}_i - \tilde{s}_j) \iff \\
\sigma\left(\frac{1}{T} \ln\left(\frac{p}{1-p}\right)\right) &= \sigma(\tilde{s}_i - \tilde{s}_j) \iff \\
\frac{1}{T} \ln\left(\frac{p}{1-p}\right) &= \tilde{s}_i - \tilde{s}_j \iff \\
\frac{p}{1-p} &= \exp(T(\tilde{s}_i - \tilde{s}_j)) \iff \\
p &= \frac{\exp(T(\tilde{s}_i - \tilde{s}_j))}{1 + \exp(T(\tilde{s}_i - \tilde{s}_j))} \iff \\
p &= \frac{1}{1 + \exp(-T(\tilde{s}_i - \tilde{s}_j))} \iff \\
p &= \sigma(T(\tilde{s}_i - \tilde{s}_j))
\end{aligned}$$

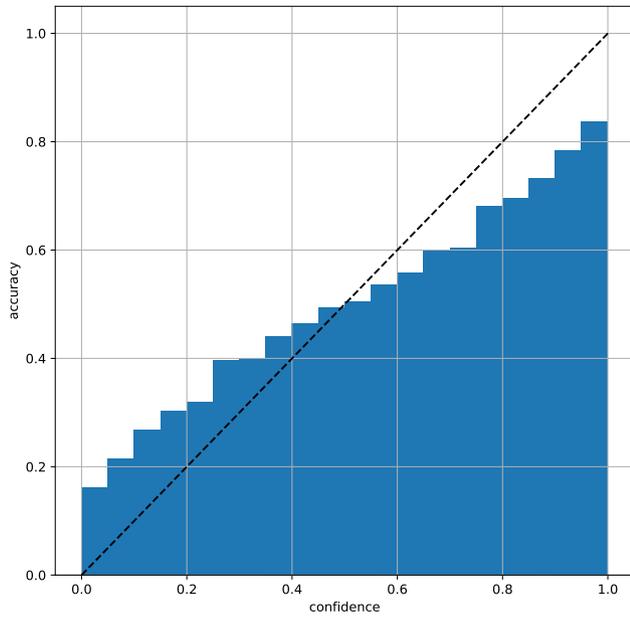
This shows that temperature annealing leads to a new solution of scores that are linearly scaled by the temperature T . Therefore, even if temperature annealing is enough to calibrate a system, it has no impact at all on the predicted rankings.

Instead, we report a different result, the diagrams in Figure 5. We obtain either the confidence of each comparison or the probability of reordering. Then the aim is to compute the accuracy of comparisons on a filtered dataset when removing the examples of lowest confidence/highest uncertainty. What one expects from high quality uncertainties is for the accuracy of the filtered dataset to improve as much as possible. While we observe that the accuracy improves as we reject samples, both metrics display a significant overconfidence issue; accuracy reduces when rejecting samples with the highest confidence and lowest uncertainty. This could partially explain why our results in Section 6.2 showcase a 'bump', where adding more comparisons decreases the system's performance. This also justifies using more advanced methods for calibrating the outputs of LLM-judges when using them to rank candidates.

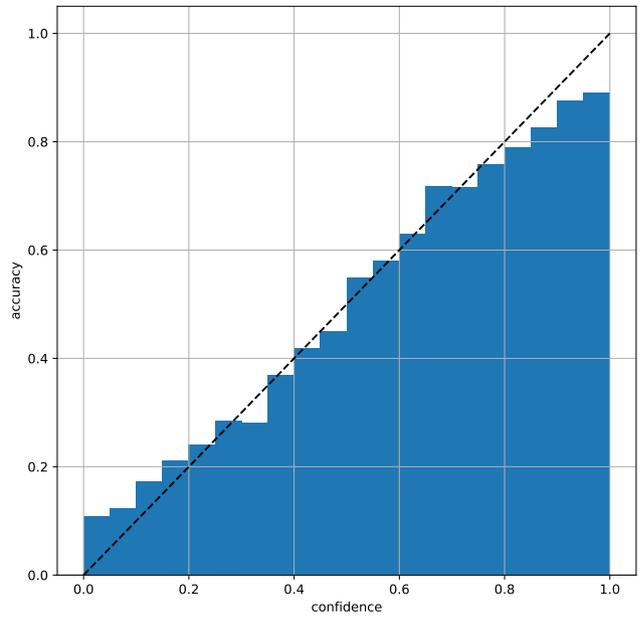
Finally, we report the reliability diagram of all language models on SummEval (COH), see Figure 6. While the Qwen2.5 family performs better in predicting the correct ranking, these systems suffer more from overconfidence as evidenced by the much higher calibration error. As discussed, temperature annealing can resolve the calibration issue without having any impact on performance. Instead, viable approaches that could resolve this include:

1. Ensembling multiple comparative LLM-as-a-judge backbones. Averaging outputs is known to reduce calibration error and the overconfidence issue Guo et al. [2017].
2. Ensembling multiple prompts using the same LLM backbone. Alternatively to above, one can design multiple prompts which achieve the same task, and average the output from each one.

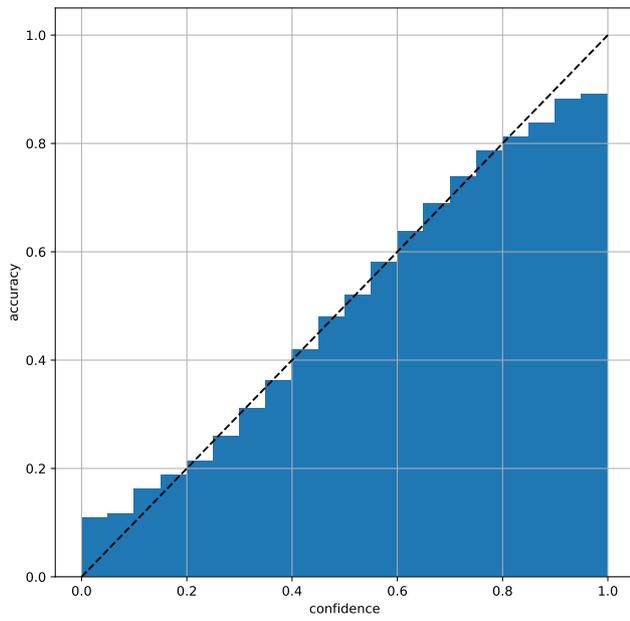
These ensembling and distillation-based strategies are well-established in the broader uncertainty estimation literature. They have been successfully applied to improve calibration and provide robust uncertainty estimates for downstream tasks like out-of-distribution detection, active learning, and grammatical error correction across various domains Hinton et al. [2015], Fathullah and Gales [2022], Fathullah et al. [2023], Malinin et al. [2020] In addition to the above, one can aim to design improved prompts which achieve the task and minimise biases in the system. We leave these approaches to future work Yang et al. [2024].



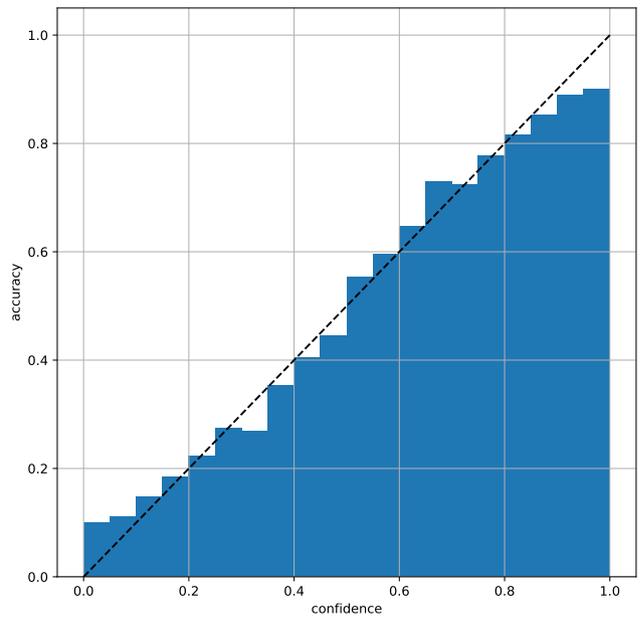
(a) Biased



(b) Debiased

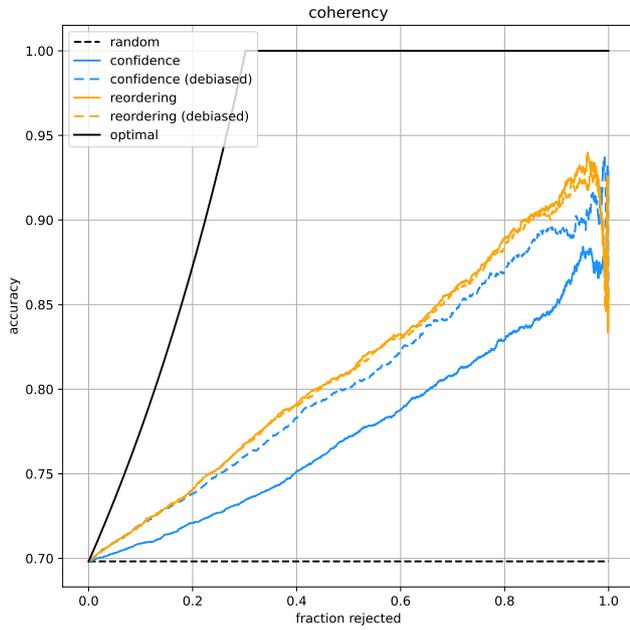


(c) Biased - calibrated

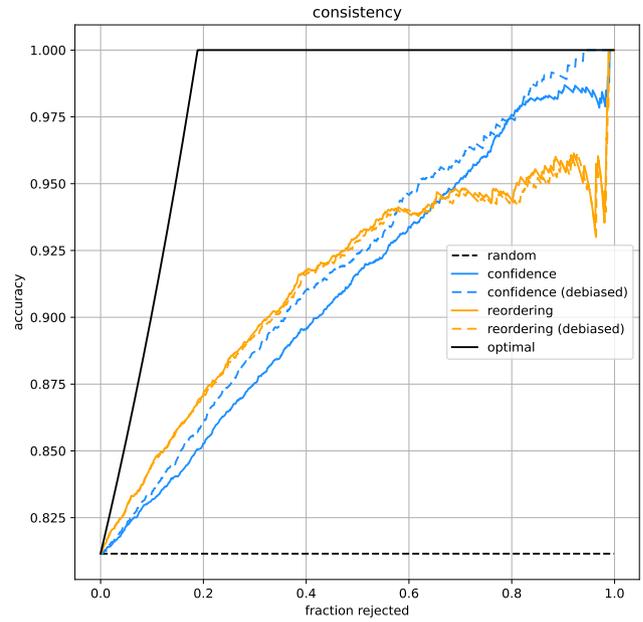


(d) Debiased - calibrated

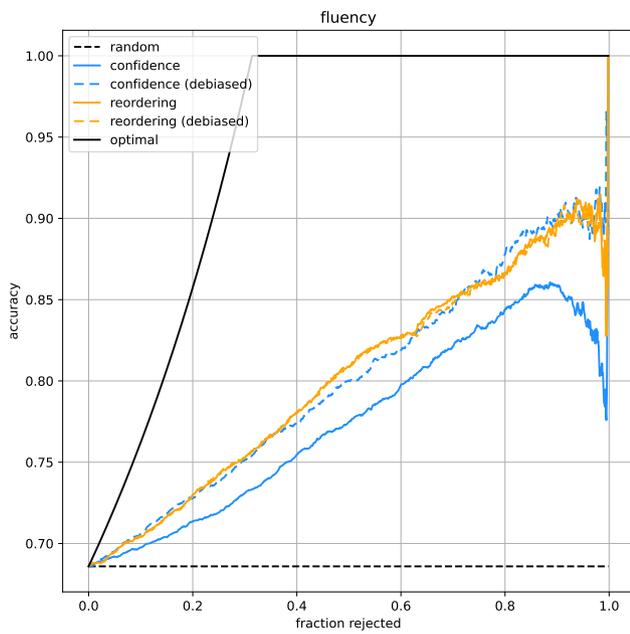
Figure 4: The reliability diagram of biased and debiased, standard and calibrated systems on the coherency metric.



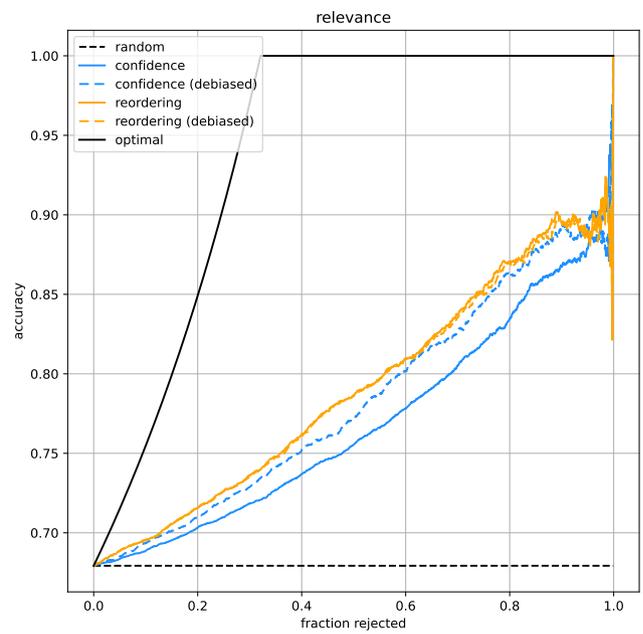
(a) COH



(b) CON

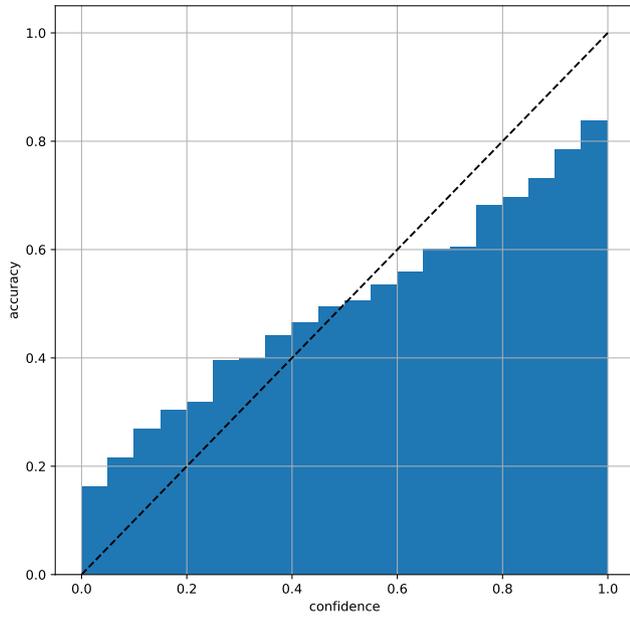


(c) FLU

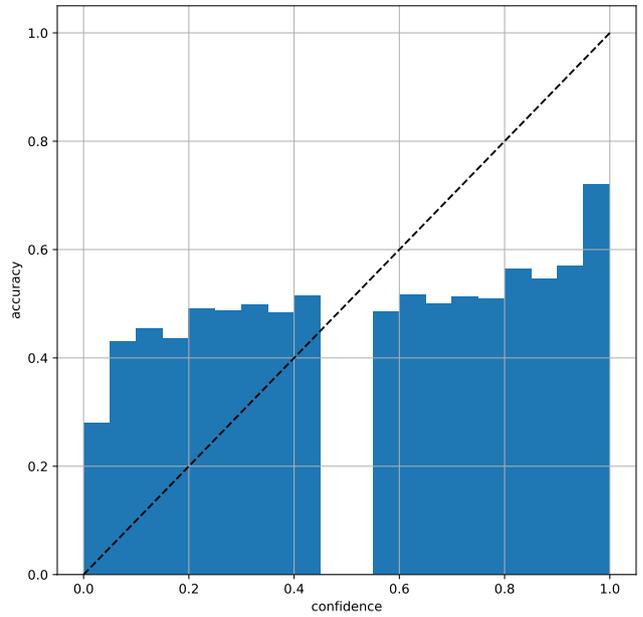


(d) REL

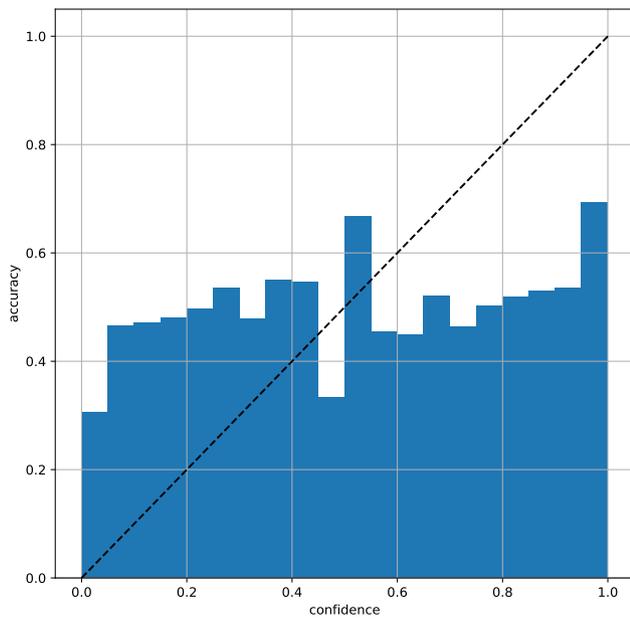
Figure 5: The accuracy at a comparison-level when the examples of lowest confidence/highest uncertainty are rejected.



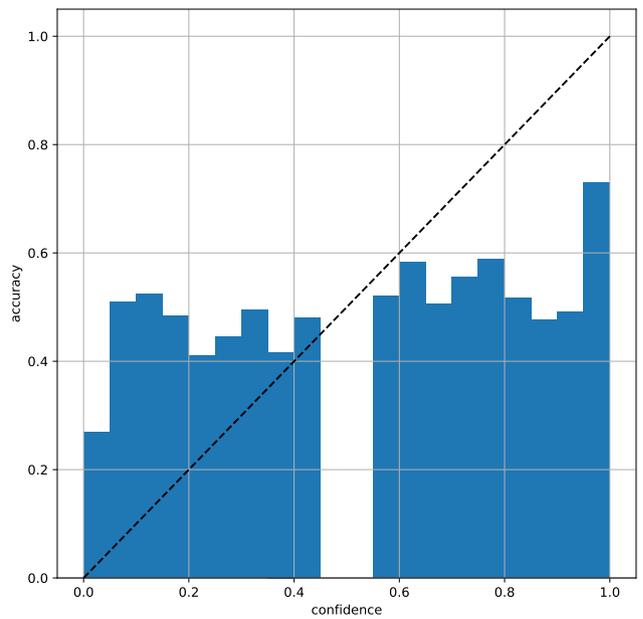
(a) Flan-T5 (3B): ECE = 9.80%



(b) Qwen2.5-Instruct (3B): ECE = 25.73%



(c) Qwen2.5-Instruct (7B): ECE = 28.17%



(d) Qwen2.5-Instruct (14B): ECE = 24.64%

Figure 6: The reliability diagram of biased systems on the coherency metric. The expected calibration error (ECE) is reported for each system.

C PROBABILITY OF REORDERING

In this section we showcase how the probability of reordering can be rephrased to a familiar form. Starting from the general Laplace's approximation:

$$p(\mathbf{s}|\mathcal{C}_{1:K}) = \mathcal{N}(\mathbf{s}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$$

the distribution of a pair of candidates becomes:

$$p(s_i, s_j|\mathcal{C}_{1:K}) = \mathcal{N}\left(\begin{bmatrix} s_i \\ s_j \end{bmatrix}; \begin{bmatrix} \mu_i \\ \mu_j \end{bmatrix}, \begin{bmatrix} \Sigma_{ii} & \Sigma_{ij} \\ \Sigma_{ji} & \Sigma_{jj} \end{bmatrix}\right)$$

We are only interested in the distribution of the difference $s_i - s_j$:

$$p(s_i - s_j|\mathcal{C}_{1:K}) = \mathcal{N}(s_i - s_j; \mu_i - \mu_j, \Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj})$$

Assuming that $s_i > s_j$ ($\mu_i > \mu_j$) the probability of reordering becomes:

$$P(s_i < s_j|\mathcal{C}_{1:K}) = \Phi\left(\frac{\mu_j - \mu_i}{\sqrt{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}}\right)$$

The selection is based on picking the examples with highest probability of reordering:

$$\begin{aligned} \hat{i}, \hat{j} &= \arg \max_{i,j} \Phi\left(\frac{\mu_j - \mu_i}{\sqrt{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}}\right) \\ &= \arg \max_{i,j} \frac{\mu_j - \mu_i}{\sqrt{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}} \\ &= \arg \min_{i,j} -\frac{\sqrt{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}}{\mu_i - \mu_j} \\ &= \arg \max_{i,j} \frac{\Sigma_{ii} - 2\Sigma_{ij} + \Sigma_{jj}}{(\mu_i - \mu_j)^2} \end{aligned}$$

Similarly, assuming $s_j > s_i$ returns the exact same expression.

D EXTENDED RESULTS FOR QWEN SYSTEMS

This section expands on the results in Section 6.2 for the Qwen family of models. Following the initial analysis, further experiments were conducted to provide a more comprehensive evaluation of the proposed methods. We structure the analysis into three parts: first, we examine the sensitivity of the final rankings to the choice of expert model; second, we assess the sensitivity of the ranking-level uncertainty to the same choices; and third, we evaluate the efficiency of different iterative selection metrics.

D.1 SENSITIVITY OF FINAL RANKINGS TO EXPERT MODEL CHOICE

Table 4 shows the final Spearman Rank correlation for various expert model configurations (f , p) across all Qwen models, using the full set of comparisons. A key finding is the relative stability of the final rankings. For a given LLM judge, the choice of the expert model has a very limited impact on the final Spearman correlation. This suggests that when a sufficient number of comparisons are available, the specific probabilistic formulation is not a critical factor in determining the final rank order of the candidates.

Table 4: **Spearman Rank Correlation (%)** on SummEval for various expert model configurations using Qwen judges. Results are shown for the full set of comparisons.

LLM	Function f	Distribution p	COH	CON	FLU	REL	Average
Qwen2.5 (3B)	x	Gaussian	48.5	48.6	40.8	46.5	46.1
	σ	Gaussian	48.4	48.6	40.9	46.4	46.1
	Φ	Gaussian	48.5	48.6	40.9	46.6	46.1
	σ	Beta	48.5	48.6	40.8	46.5	46.1
	Φ	Beta	48.6	48.6	40.9	46.5	46.2
Qwen2.5 (7B)	x	Gaussian	50.4	48.9	38.7	50.4	47.1
	σ	Gaussian	50.5	48.9	38.8	50.6	47.2
	Φ	Gaussian	50.4	48.9	38.8	50.6	47.2
	σ	Beta	50.4	48.9	38.7	50.5	47.1
	Φ	Beta	50.3	48.9	38.7	50.5	47.1
Qwen2.5 (14B)	x	Gaussian	57.3	52.1	47.4	51.1	52.0
	σ	Gaussian	57.3	52.1	47.4	51.2	52.0
	Φ	Gaussian	57.4	52.1	47.4	51.2	52.0
	σ	Beta	57.4	52.1	47.4	51.1	52.0
	Φ	Beta	57.4	52.1	47.4	51.1	52.0

D.2 SENSITIVITY OF RANKING UNCERTAINTY TO EXPERT MODEL CHOICE

In contrast to the final rankings, the quality of the ranking-level uncertainty is more sensitive to the expert model configuration. Table 5 shows the AUROC scores, which measure how well the model’s uncertainty can predict high-performing rankings. Here, we observe greater variability. Models that use a sigmoidal mapping function (σ or Φ) tend to yield better AUROC scores than the linear ‘Identity’ function. This confirms that the nature of the probabilistic model has a notable impact on the quality of the overall uncertainty estimates.

Table 5: **AUROC (%) for Ranking Uncertainty** on SummEval for various expert model configurations using Qwen judges.

LLM	Function f	Distribution p	COH	CON	FLU	REL	Average
Qwen2.5 (3B)	x	Gaussian	60.1	53.0	59.8	51.4	56.1
	σ	Gaussian	53.7	55.4	62.2	56.6	57.0
	Φ	Gaussian	50.6	54.3	65.1	58.6	57.2
	σ	Beta	59.8	62.8	63.2	57.5	60.8
	Φ	Beta	53.5	53.8	63.1	58.6	57.2
Qwen2.5 (7B)	x	Gaussian	57.7	50.6	51.9	51.2	52.9
	σ	Gaussian	63.8	61.9	55.6	70.4	62.9
	Φ	Gaussian	64.0	63.4	55.6	70.3	63.3
	σ	Beta	66.7	68.2	58.5	69.4	65.7
	Φ	Beta	64.0	63.4	55.6	70.6	63.4
Qwen2.5 (14B)	x	Gaussian	59.1	55.5	52.0	56.6	55.8
	σ	Gaussian	63.5	57.9	61.6	68.0	62.8
	Φ	Gaussian	64.6	56.9	62.2	67.3	62.8
	σ	Beta	61.7	62.0	58.4	64.8	61.7
	Φ	Beta	64.2	57.2	62.4	67.6	62.8

D.3 EFFICIENCY OF ITERATIVE SELECTION METRICS

This section evaluates the efficiency of different uncertainty metrics for iteratively selecting the most informative comparisons. Table 6 compares the $\epsilon = 0.5$ metric with the probability of reordering ($\epsilon = 2.0$) from Eq. (6). The results show that $\epsilon = 0.5$ is highly competitive and often more efficient, particularly for the largest Qwen2.5-14B model, where it consistently requires fewer comparisons to reach 90% of the final performance.

Table 6: **Efficiency comparison of uncertainty selection metrics** for Qwen models. The values represent the number of comparisons needed to achieve 90% of the final performance. Lower is better.

LLM	Metric (ϵ)	COH	CON	FLU	REL	Average
Qwen2.5 (3B)	0.5	28.2	23.1	23.6	33.5	27.1
	2.0	27.9	22.7	24.0	34.0	27.1
Qwen2.5 (7B)	0.5	32.9	25.8	29.9	27.7	29.1
	2.0	34.1	27.3	31.1	28.1	30.2
Qwen2.5 (14B)	0.5	25.7	23.5	26.0	25.6	25.2
	2.0	29.6	26.6	33.9	27.9	29.5

The broader efficiency comparison in Table 7 and visualized in Figures 7, 8, and 9 corroborates these findings. While variance only performs well under debiased outputs from an LLM, the probability of reordering is stable under both cases, significantly outperforming the baseline minimum uncertainty in all cases.

Table 7: **Measure of efficiency:** The number of comparisons needed to achieve 90% performance of the baseline soft Bradley-Terry system for each corresponding judge. The baseline is evaluated using all $N(N - 1) = 240$ comparisons.

Uncertainties	Debiased	Qwen2.5 (7B)				Qwen2.5 (14B)			
		COH	CON	FLU	REL	COH	CON	FLU	REL
minimum uncertainty	✗	73.1 ± 9.0	46.9 ± 5.6	60.8 ± 13.3	71.8 ± 8.4	75.9 ± 7.7	50.2 ± 6.0	64.2 ± 12.7	67.6 ± 9.8
	✓	61.0 ± 9.3	46.8 ± 8.3	50.9 ± 9.3	60.3 ± 11.1	67.0 ± 5.0	48.8 ± 7.4	62.0 ± 9.1	61.6 ± 9.3
variance	✗	56.1 ± 9.5	36.6 ± 9.6	52.0 ± 19.6	63.4 ± 10.7	59.7 ± 12.6	36.2 ± 8.0	61.0 ± 14.3	51.0 ± 9.6
	✓	24.3 ± 3.6	20.6 ± 4.2	23.2 ± 5.9	31.8 ± 5.9	27.6 ± 4.5	22.8 ± 2.4	31.6 ± 9.5	31.4 ± 6.0
probability of reordering	✗	35.6 ± 5.7	27.4 ± 2.7	31.6 ± 5.0	29.4 ± 4.3	28.1 ± 2.6	25.6 ± 2.5	34.2 ± 4.9	26.1 ± 2.7
	✓	21.5 ± 4.3	20.1 ± 2.5	20.5 ± 3.9	21.4 ± 2.4	20.6 ± 2.0	19.9 ± 2.1	23.1 ± 3.2	20.1 ± 1.6

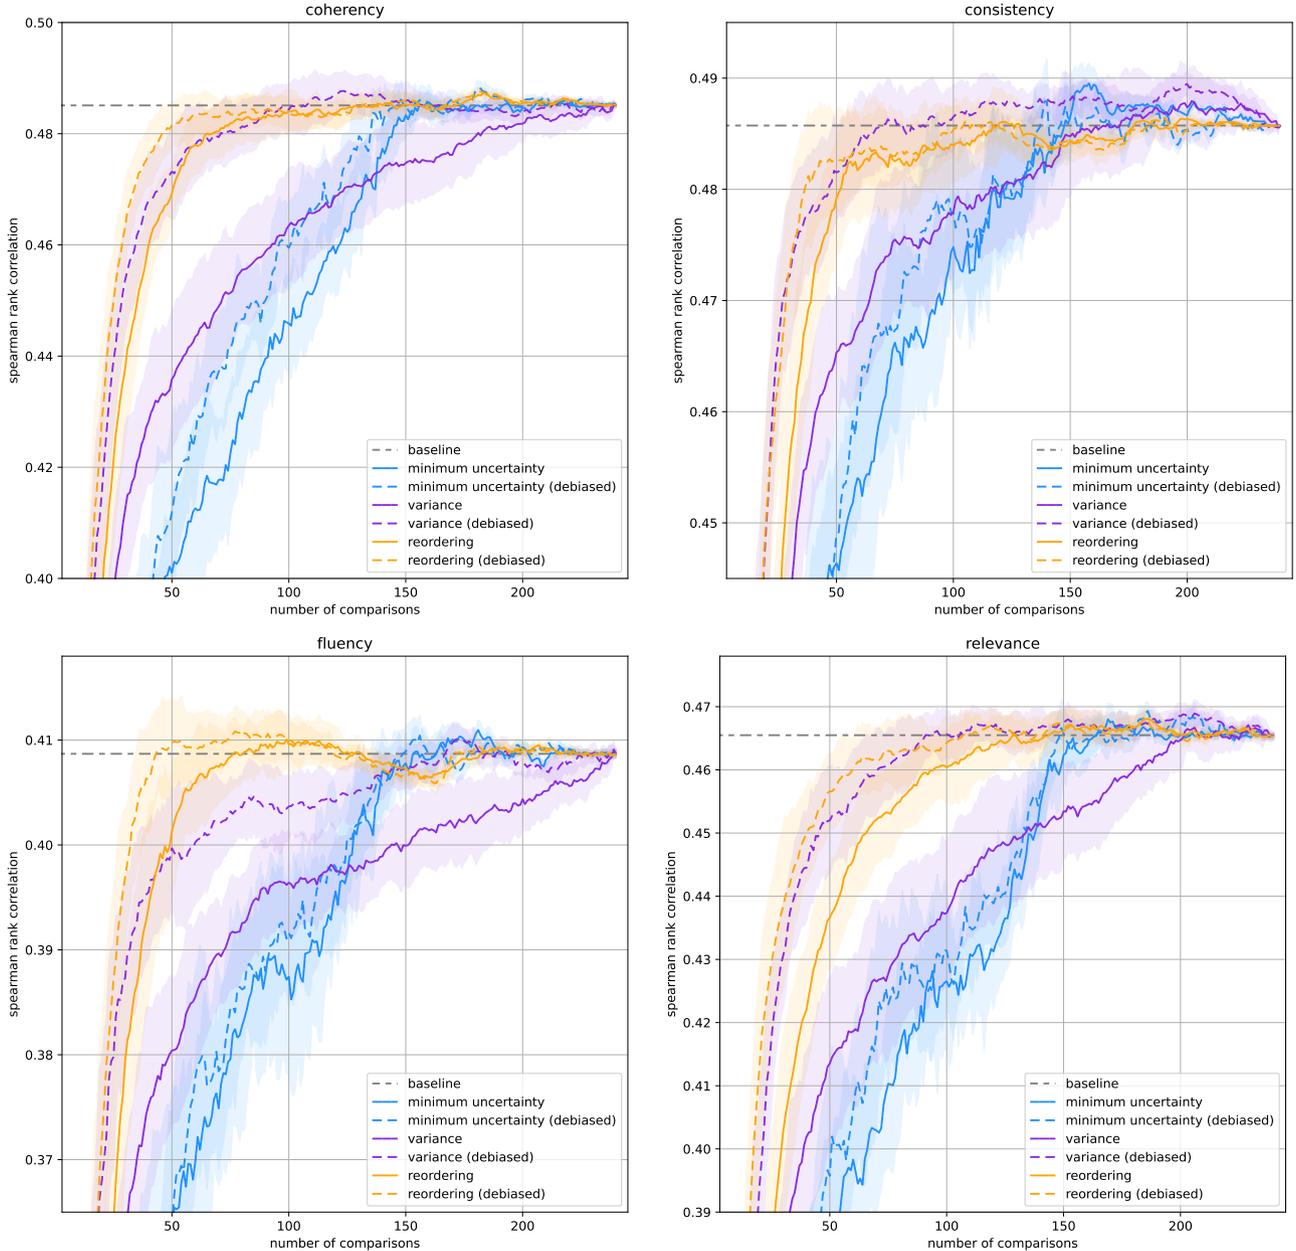


Figure 7: **Qwen2.5-3B-Instruct:** The Spearman Rank Correlation when iteratively selecting the next examples of lowest confidence/highest uncertainty. The baseline is the soft Bradley-Terry model with the minimum uncertainty metric. We also report the proposed variance and probability of reordering under the soft BT model.

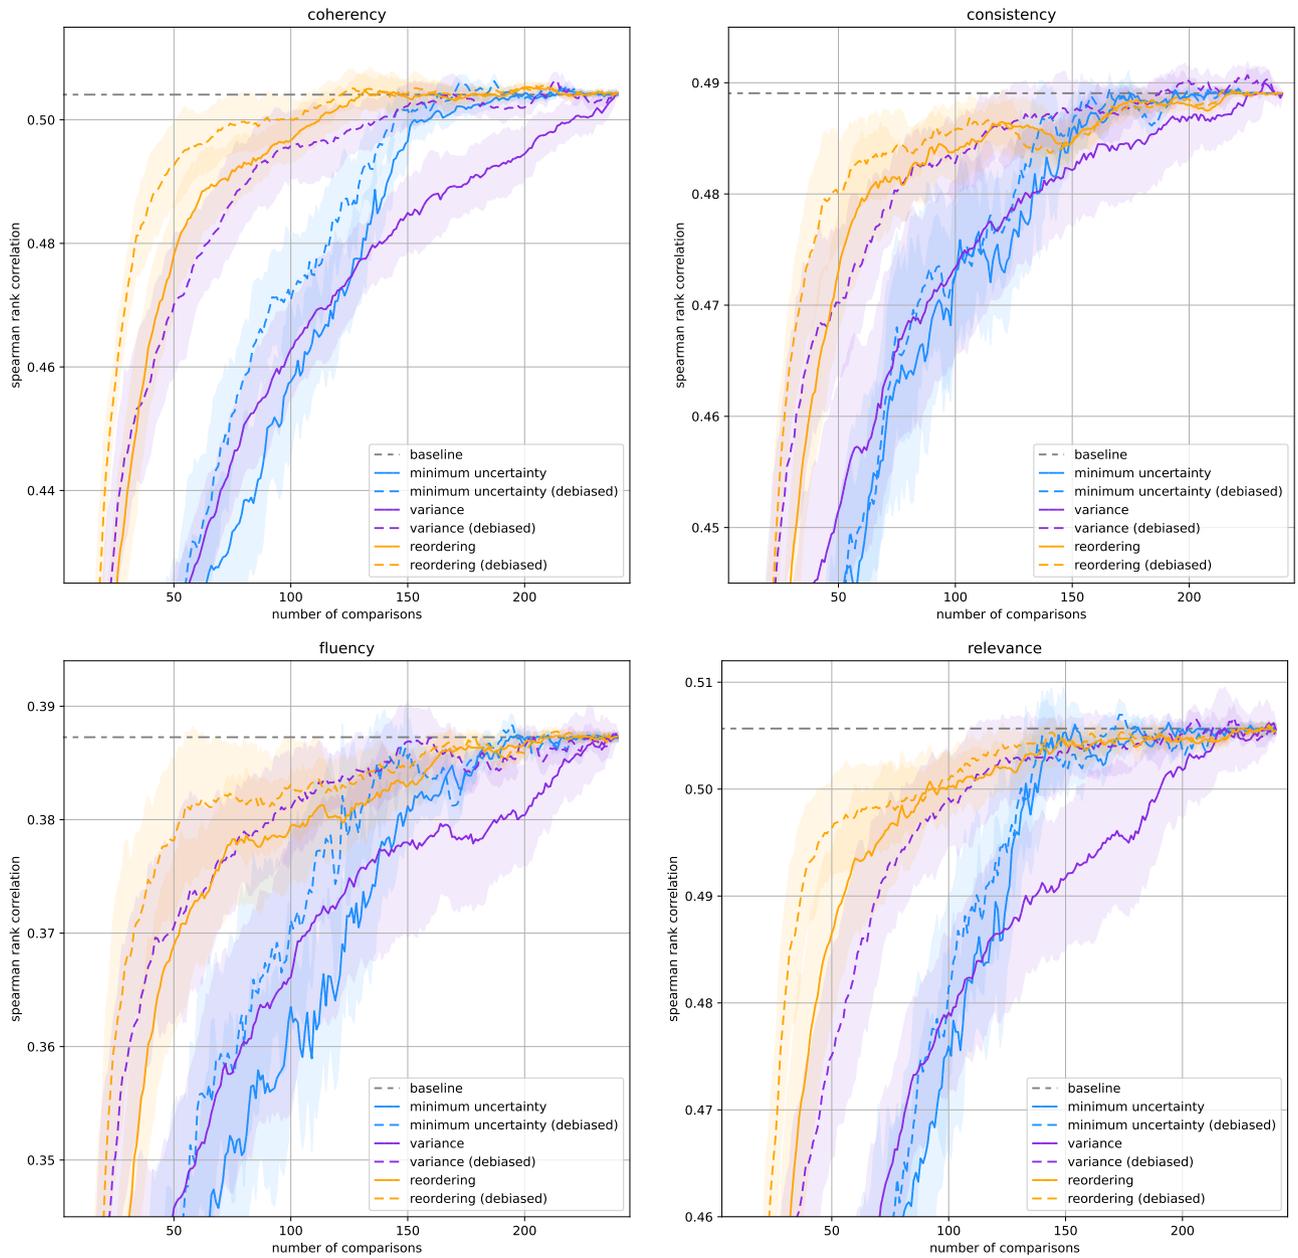


Figure 8: **Qwen2.5-7B-Instruct**: The Spearman Rank Correlation when iteratively selecting the next examples of lowest confidence/highest uncertainty. The baseline is the soft Bradley-Terry model with the minimum uncertainty metric. We also report the proposed variance and probability of reordering under the soft BT model.

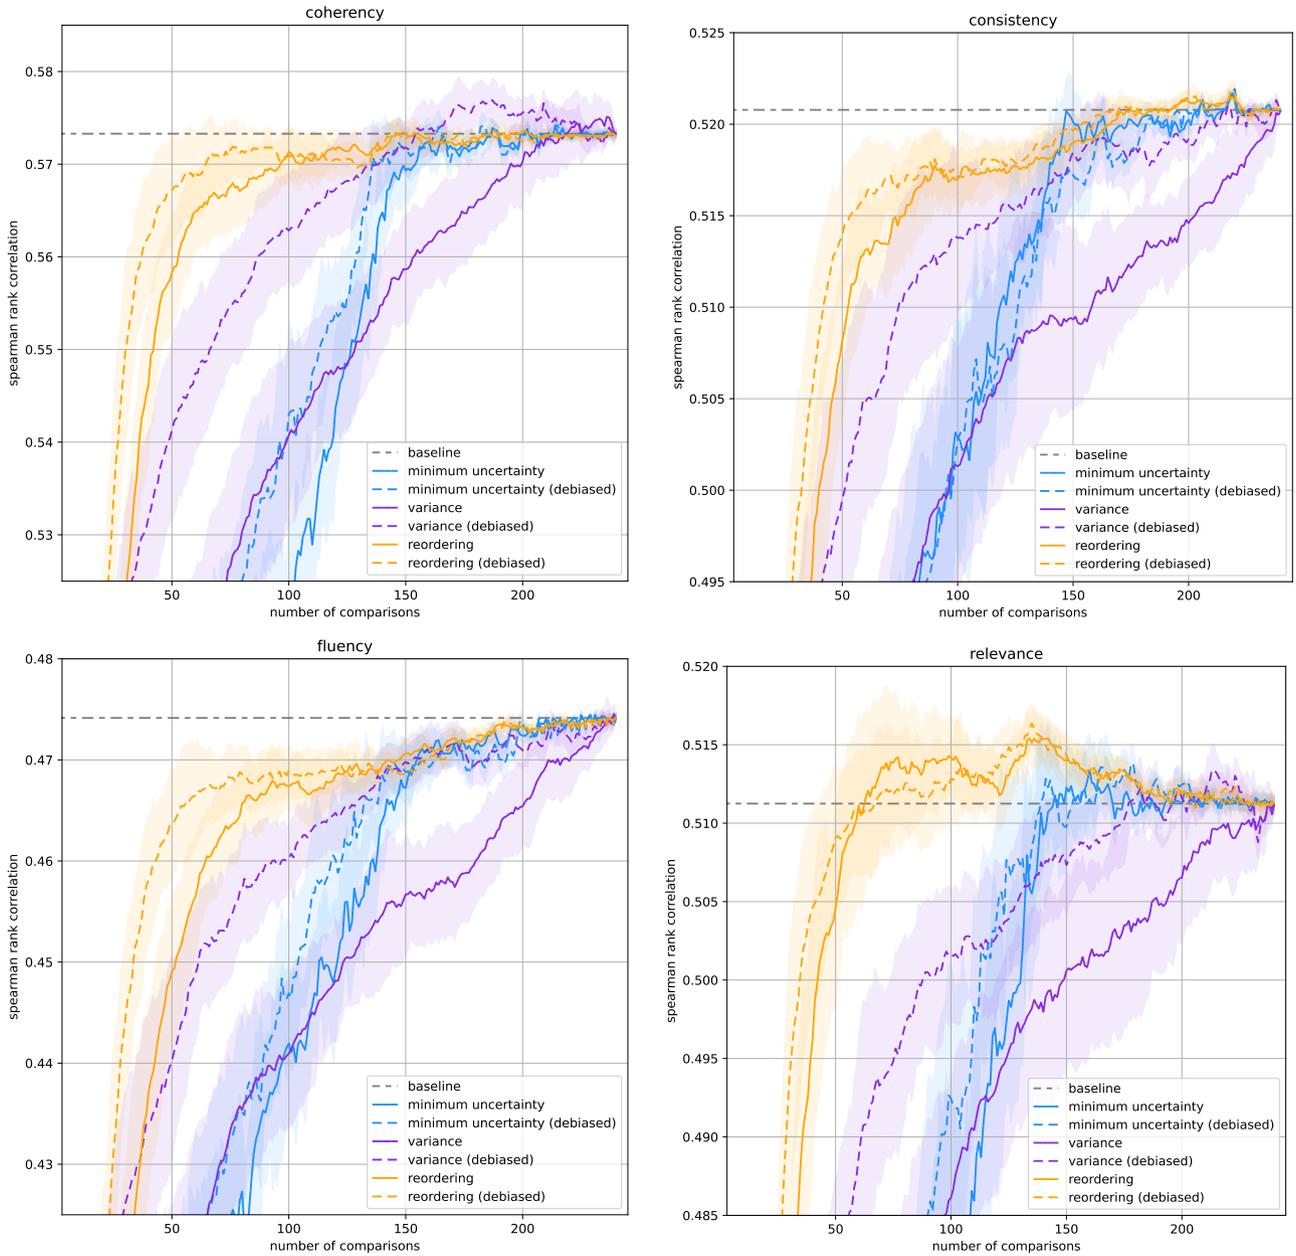


Figure 9: **Qwen2.5-14B-Instruct**: The Spearman Rank Correlation when iteratively selecting the next examples of lowest confidence/highest uncertainty. The baseline is the soft Bradley-Terry model with the minimum uncertainty metric. We also report the proposed variance and probability of reordering under the soft BT model.