

# UDA: Unsupervised Debiasing Alignment for Pair-wise LLM-as-a-Judge

Yang Zhang<sup>1,2,\*†</sup>, Cunxiang Wang<sup>1,3,†‡</sup>, Lindong Wu<sup>1</sup>, Wenbo Yu<sup>1</sup>,  
Yidong Wang<sup>1,2</sup>, Guangsheng Bao<sup>4</sup>, Jie Tang<sup>3‡</sup>

<sup>1</sup>ZhipuAI

<sup>2</sup>Peking University

<sup>3</sup>Tsinghua University

<sup>4</sup>Westlake University

zhang360428@stu.pku.edu.cn, wangcunxiang303@gmail.com

## Abstract

Pairwise evaluation of Large Language Models (LLMs) is a common paradigm, but it is prone to preference bias, where judges systematically favor certain outputs, such as their own. This bias leads to inconsistent and skewed rankings across different judges. To address this, we first empirically demonstrate significant and heterogeneous biases in cross-model evaluations. We then propose UDA (Unsupervised Debiasing Alignment), a framework that reduces inter-judge disagreement by dynamically adjusting the Elo rating system. For each pairwise comparison, a compact neural network learns to adaptively set the K-factor and refine win probabilities. Crucially, UDA operates in a fully unsupervised manner, guided solely by the objective of minimizing the dispersion among the Elo trajectories of all judges. This forces an alignment towards a collective consensus, which serves as an unsupervised proxy for a more stable and reproducible evaluation. In addition, we provide theoretical motivation demonstrating how alignment towards a consensus can reduce aggregate system bias. Experiments show that UDA significantly reduces the inter-judge rating standard deviation by up to 63.4% and improves the average correlation with human judgments by 24.7%. Notably, UDA elevates the performance of poorly performing judges to achieve parity with high-quality ones, fostering a more robust and reliable evaluation ecosystem.

**Code** — <https://github.com/zhang360428/Debias>

## Introduction

Large Language Models (LLMs) like GPT-4 and Claude-3 now drive progress across Natural Language Processing (NLP), while their daily use in commerce, education, and entertainment seamlessly integrates them into everyday life (Zhao et al. 2025; Minaee et al. 2025). Rigorous evaluation is therefore indispensable: it guides model optimization and equips users with an empirical basis for selection, ensuring robust deployment in the wild (Chang et al. 2023).

The paradigm of ‘LLM-as-a-judge’ (Zheng et al. 2023) delegates the evaluation to a second large model that acts as

\*Work was done at ZhipuAI intern.

†These authors contributed equally.

‡Corresponding author.

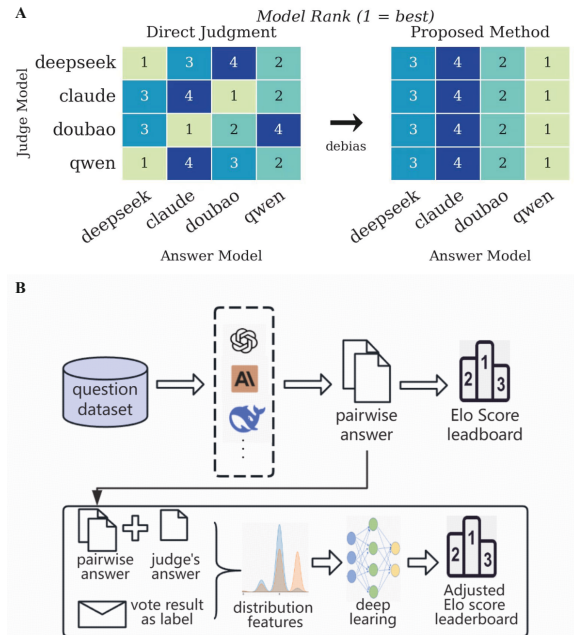


Figure 1: **A.** Some models prefer their own answers, and mitigating this bias from using different LLMs can make the results more accurate. **B.** Dynamically recalibrate, rather than naively adopts, LLM-judged scores, using consensus among judges as the supervisory label in lieu of human annotation.

an autonomous judge: the judge is prompted to score, rank or critique the outputs of the target model without human labels. Yet this very mechanism introduces systematic bias (Wang et al. 2023; Thakur et al. 2025). In this paradigm, pairwise judgments consistently outperform pointwise scoring in robustness and are widely adopted, for example, Zheng et al. first established this in the LMSys Chatbot Arena (Chiang et al. 2024): GPT-4 achieved markedly higher agreement with human preferences when ranking anonymized responses head-to-head than when assigning absolute scores. Subsequent works (Wang, Zhang, and Choi 2025) corroborate that pairwise ranking reduces intra-model

variance; on RewardBench, GPT-4o attains 90.5% accuracy in pairwise mode versus 88.0% in pointwise mode, while smaller models trend in the opposite direction. Different LLMs are built and trained in unique ways, causing them to have specific biases, such as a preference for wordy or stylistically similar answers. This results in varied scores and rankings. Prior works show that score disparities exceeding 30% among LLMs on identical benchmarks (Ye et al. 2024).

Most researches (Chang et al. 2023) mainly focus on the improvement of evaluation indicators and methods, but the studies on how to reduce the impact of model preferences during the evaluation process are not in-depth enough. Although recent work calibrates proprietary models on-the-fly via prompt engineering (Zhou et al. 2024), the practice remains ad-hoc and theoretically thin. Ensemble schemes (Wataoka, Takahashi, and Ri 2025) mitigate bias but explode in cost, and Nash-based re-ranking (Liu et al. 2025a) scales poorly. In addition, there is relatively little research on how to use deep learning techniques to dynamically adjust the evaluation process, rather than just directly using the evaluation results provided by large models, in order to reduce the bias from different judge models.

To address these limitations, we propose UDA (Unsupervised Debiasing Alignment), a fully automated, annotation-free framework that reconciles heterogeneous LLM judges. UDA retains the anonymity and simplicity of pairwise Elo tournaments but injects a learned, instance-level correction mechanism. Concretely, for each comparison, we (i) calculate the distributional similarity between the two answers, (ii) quantify the judge’s preference by contrasting its own generated response with the two candidates, and (iii) feed these continuous signals into a lightweight neural adapter that outputs an adaptive K-factor and refined win probability expectations. The adapter is trained to pull each judge’s scoring trajectory towards the collective consensus of all judges. This consensus, while not a perfect ground truth, serves as a pragmatic, unsupervised signal to mitigate extreme idiosyncratic biases. Because these adjustments are conditioned solely on the response distributions, UDA is model-agnostic and requires no additional fine-tuning.

Bias usually shifts across topics: a judge fair on one task may favor itself on another. To verify generalisation beyond public benchmarks, we evaluate UDA zero-shot on a newly curated HUMAN-ANNOTATED TRANSFER SET that contains 100 open-ended prompts distinct from the training corpus. Without any retraining, UDA still shrinks the inter-judge standard deviation by 63.4% and boosts correlation with human preferences by 24.7%, demonstrating robust transfer to unseen domains and judges.

We first quantify the tendency of each judge to over- or under-rate its own response. On our curated dataset, baseline Elo produces self-scores that deviate from the average score assigned by the remaining judges by as much as  $-21\%$  (underrating) to  $+56\%$  (overrating). After applying UDA, the same deviations narrow to only  $-15\%$  to  $+4\%$ , indicating that extreme self-preference is largely suppressed. Under mild assumptions, UDA minimises the *aggregate* bias across all judges (Theorem 1). However, because the procedure pulls every judge toward the consensus, exceptionally well-

calibrated judges may experience a slight drop in individual accuracy; the collective variance reduction dominates this per-model cost. Despite the small per-model trade-off, UDA improves the average Pearson correlation with human rankings from 0.65 to 0.81 (+25%), confirming that shrinking inter-judge disagreement also aligns the final scores more closely with human preferences.

To the best of our knowledge, we are the first to reduce pairwise LLM-as-a-judge preference bias in an unsupervised manner.

## Related Work

### LLM-as-a-Judge

The LLM-as-a-judge approach has been widely adopted for evaluating LLMs due to its low cost and high efficiency (Gu et al. 2025). This method leverages one LLM to assess the outputs of another, providing fine-grained evaluations. For instance, OpenAI’s GPT and Anthropic’s Claude are commonly utilized as judge models to evaluate the performance of other LLMs. However, this approach may introduce significant bias. Different judge models can produce varying results and rankings, as their specific traits, differences in training data, and subjective evaluation criteria can influence the evaluation results.

### Elo Rating System

The Elo rating system (Hunter 2004), traditionally used in chess, has been adapted for evaluating large language models (LLMs). Chatbot Arena (Chiang et al. 2024) introduced an Elo system that assesses chatbots through pairwise comparisons based on user votes. This method offers a dynamic and interactive approach to evaluation, providing a degree of objectivity and comparability. However, it may be influenced by user preferences and the diversity of evaluation questions. Yet, this system still relies on user feedback and faces challenges in scalability and scenario diversity. Gray et al. (2022) combined the ELO system with traditional metric-based evaluations, integrating model rankings with specific metric analyses for a more comprehensive assessment. Despite these advancements, existing ELO evaluation systems in LLMs still struggle with balancing the weight of different evaluation dimensions and criteria.

### Bias in LLM-as-a-Judge

Preference bias in LLMs is a critical evaluation issue (Wataoka, Takahashi, and Ri 2025). It causes LLMs to overrate certain outputs, especially their own output. Studies (Zheng et al. 2023; Xu et al. 2024) reveal that some LLMs favor their own generations. This bias may stem from the model’s training data and architecture. Larger and more capable models often show stronger self-preference. The model’s internal probabilities may favor familiar generation styles, and alignment training artifacts could also contribute. Panickssery, Bowman, and Feng explored the link between self-preference bias and self-recognition ability. As shown in Figure 2, it reveals heterogeneous LLM preferences.

Other biases also exist. Some models show position bias, preferring the first or last option in pairwise comparisons.




Figure 2: Pairwise answer evaluation experiments were conducted using different large models as judges on the same Arena Hard dataset. The figure shows that LLMs exhibit self-preferential bias: some over-rate their own answers relative to other judges, others under-rate them. On ArenaHard (the depicted dataset) this bias ranges from  $-38\%$  to  $+90\%$ , while on our dataset it spans  $-21\%$  to  $+56\%$ .

Model performance is influenced by whether real names or anonymous aliases are used (Koo et al. 2024). Large models tend to favor their own responses under anonymity, but this self-preference decreases when real names are shown, while smaller models’ self-preference increases.

### Mitigating LLM Judges’ Bias

Zhou et al. (2024) tackled the bias issue in the “LLM-as-a-judge” approach by quantifying the bias in large models’ predictive distributions and then subtracting it to reduce the bias. However, this method has limitations when evaluating close-source llm. Due to the inaccessibility of predictive logic, it is difficult to calculate the distributional differences between pre-trained and SFT models. Consequently, it resorts to prompt engineering to model surface quality and mitigate bias. Liu et al. (2025a) adopted a game-theoretic perspective, framing evaluation as a three-player game and proposing a Nash equilibrium solution. Yet, the equilibrium ratings lack scalability and cannot be directly applied to many real-world large models. Liu et al. (2025b) tried to optimize ELo rating systems in arena settings, they proposed the am-ELO framework, which uses maximum likelihood estimation (MLE) to adjust the Elo probability function by incorporating annotator abilities, enhancing stability and accuracy from an annotation accuracy perspective.

However, these methods either rely on inaccessible model internals, resort to prompt engineering, or heuristics-based reranking that scales poorly. Moreover, these methods overlook the unequal impact of individual comparisons in scoring systems: a weaker model’s loss to a near-equal rival is weighted identically to its loss to a markedly stronger one, an evident flaw that undermines power when sample sizes are small. They also ignore judge-induced distributional biases—e.g., inflated scores arising simply from a preference for responses resembling their own. In contrast, our approach dynamically reweights each pairwise outcome and incorporates distributional features to correct for biases, using the principle of consensus alignment as a supervisory

signal in the absence of human labels.

## Method

UDA (Unsupervised Debiasing Alignment) is an annotation-free method that reconciles a set of  $M$  heterogeneous LLM judges,  $\mathcal{J}$ , by dynamically adjusting the Elo update rule for a set of  $N$  candidate models,  $\mathcal{M}$ . Given a prompt  $x$ , a judge  $k \in \mathcal{J}$  expresses its raw binary preference  $r_{ij}^{(k)} \in \{0, 1\}$  for paired responses  $(a_i, a_j)$  from models  $i, j \in \mathcal{M}$ , which in turn determines the Elo score  $R_i^{(k)} \in \mathbb{R}$  for each model under that judge. The key idea is to treat the **collective consensus ranking** of all judges as a pragmatic, unsupervised target without human labels.

### Baseline Elo Updating Rule

Baseline Elo updates judge  $k$ ’s score via

$$R_i^{(k)} \leftarrow R_i^{(k)} + K(r_{ij}^{(k)} - \hat{p}_{ij}^{(k)}), \quad \hat{p}_{ij}^{(k)} = \sigma(R_i^{(k)} - R_j^{(k)}), \quad (1)$$

where  $\sigma(z) = 1/(1 + e^{-z})$  and  $K$  is a global constant.

### Adaptive Debiasing Network

For every pair of answers  $(a_i, a_j)$  and a judge  $k$ , we extract a feature vector  $\phi_{ij}^{(k)}$  based on the semantic embeddings of the answers  $(e_i, e_j)$  and the judge’s own generated response  $(e_k)$ . These features, constructed without human labels, capture two key aspects: (i) the semantic and stylistic relationship between the two answers being compared, and (ii) the similarity of each answer to the judge’s own output, which serves as a signal for potential self-preference bias. The feature set includes vector operations (e.g., element-wise differences) and scalar metrics (e.g., cosine similarity, KL-divergence). The precise construction of the full feature vector is detailed in Algorithm 1.

A small Multilayer Perceptron (MLP)  $f_\theta : \mathbb{R}^{6D+9} \rightarrow \mathbb{R}^3$  then maps this feature vector to the adaptive parameters:

$$(K_{ij}, s_i, s_j) = f_\theta(\phi_{ij}^{(k)}), \quad (2)$$

where  $K_{ij}$  is the *instance-level K-factor*, and  $s_i, s_j \in (0, 1)$  are soft labels that replace the hard  $r_{ij}^{(k)}$  to mitigate preference noise ( $s_i + s_j = 1$  after softmax normalization). All parameters  $\theta$  are shared across all judges and prompts.

### The Consensus Anchor Principle

A core challenge in unsupervised debiasing is the absence of a “golden” reference. Our methodology is built upon a foundational principle: while individual LLM judges are fraught with idiosyncratic biases, the consensus derived from a diverse pool of judges can serve as a powerful proxy for a more stable and reproducible evaluation.

This principle is motivated by our empirical findings (Figure 2). The biases are not monolithic; they are **heterogeneous**. For instance, on ArenaHard, a judge like glm-4-flash exhibits strong self-preference ( $+90\%$ ), while a top-tier model like gpt-4o can paradoxically underrate its own responses ( $-38\%$ ). We hypothesize that these opposing biases tend to partially cancel each other out when aggregated.

---

**Algorithm 1: Feature Vector Construction ( $\phi_{ij}^{(k)}$ ) for UDA**


---

**Require:** Answer embeddings  $\mathbf{e}_i, \mathbf{e}_j \in \mathbb{R}^D$ ; Judge’s own embedding  $\mathbf{e}_k \in \mathbb{R}^D$

- 1: Let  $\odot$  denote the element-wise product.
  - 2: **{Part 1: High-Dimensional Features (6 vectors)}**
  - 3:  $\mathbf{v}_{\text{diff},ij} \leftarrow |\mathbf{e}_i - \mathbf{e}_j|$
  - 4:  $\mathbf{v}_{\text{diff},ik} \leftarrow |\mathbf{e}_i - \mathbf{e}_k|$  {Judge self-bias feature}
  - 5:  $\mathbf{v}_{\text{diff},jk} \leftarrow |\mathbf{e}_j - \mathbf{e}_k|$  {Judge self-bias feature}
  - 6:  $n_i \leftarrow \|\mathbf{e}_i\|_2, n_j \leftarrow \|\mathbf{e}_j\|_2, n_k \leftarrow \|\mathbf{e}_k\|_2$
  - 7:  $\mathbf{v}_{\text{norm-prod},ij} \leftarrow (\mathbf{e}_i \odot \mathbf{e}_j)/(n_i n_j + \epsilon)$
  - 8:  $\mathbf{v}_{\text{norm-prod},ik} \leftarrow (\mathbf{e}_i \odot \mathbf{e}_k)/(n_i n_k + \epsilon)$  {Judge self-bias feature}
  - 9:  $\mathbf{v}_{\text{norm-prod},jk} \leftarrow (\mathbf{e}_j \odot \mathbf{e}_k)/(n_j n_k + \epsilon)$  {Judge self-bias feature}
  - 10: **{Part 2: Scalar Features (9 scalars)}**
  - 11:  $s_{ij} \leftarrow \cos(\mathbf{e}_i, \mathbf{e}_j), s_{ik} \leftarrow \cos(\mathbf{e}_i, \mathbf{e}_k), s_{jk} \leftarrow \cos(\mathbf{e}_j, \mathbf{e}_k)$
  - 12:  $\mathbf{p}_i \leftarrow \text{softmax}(\mathbf{e}_i), \mathbf{p}_j \leftarrow \text{softmax}(\mathbf{e}_j), \mathbf{p}_k \leftarrow \text{softmax}(\mathbf{e}_k)$
  - 13:  $d_{\text{KL}}^{(j,i)} \leftarrow \text{KL}(\mathbf{p}_j \parallel \mathbf{p}_i), d_{\text{KL}}^{(i,k)} \leftarrow \text{KL}(\mathbf{p}_i \parallel \mathbf{p}_k), d_{\text{KL}}^{(j,k)} \leftarrow \text{KL}(\mathbf{p}_j \parallel \mathbf{p}_k)$
  - 14:  $\ell_{\text{diff},ij} \leftarrow |n_i - n_j|, \ell_{\text{diff},ik} \leftarrow |n_i - n_k|, \ell_{\text{diff},jk} \leftarrow |n_j - n_k|$
  - 15: **{Part 3: Concatenation}**
  - 16:  $\phi_{ij}^{(k)} \leftarrow \text{concat}[\mathbf{v}_{\text{diff},ij}, \mathbf{v}_{\text{diff},ik}, \mathbf{v}_{\text{diff},jk}, \mathbf{v}_{\text{norm-prod},ij}, \mathbf{v}_{\text{norm-prod},ik}, \mathbf{v}_{\text{norm-prod},jk}, s_{ij}, s_{ik}, s_{jk}, d_{\text{KL}}^{(j,i)}, d_{\text{KL}}^{(i,k)}, d_{\text{KL}}^{(j,k)}, \ell_{\text{diff},ij}, \ell_{\text{diff},ik}, \ell_{\text{diff},jk}]$
  - 17:  $s_{ij}, s_{ik}, s_{jk}, d_{\text{KL}}^{(j,i)}, d_{\text{KL}}^{(i,k)}, d_{\text{KL}}^{(j,k)}, \ell_{\text{diff},ij}, \ell_{\text{diff},ik}, \ell_{\text{diff},jk}$
  - 18:  $\ell_{\text{diff},ij}, \ell_{\text{diff},ik}, \ell_{\text{diff},jk}$
  - 19:  $\ell_{\text{diff},ij}, \ell_{\text{diff},ik}, \ell_{\text{diff},jk}$
  - 20: **return**  $\phi_{ij}^{(k)} \in \mathbb{R}^{6D+9}$
- 

However, we acknowledge a critical limitation of this assumption: if a majority of judges share a systemic bias (e.g., a preference for verbosity), their consensus will reflect a “bias average” rather than a debiased center. In such a scenario, UDA might inadvertently reinforce this majority bias.

Therefore, UDA does not treat the consensus score, denoted as  $\mathbf{G}$ , as an infallible ground truth. Instead, we define it as a **stabilizing proxy target**. The objective of UDA is not to perfectly replicate this consensus, but to use it as a supervisory signal to **reduce system variance**. By training each judge to align with the collective agreement, UDA effectively mitigates the extreme, idiosyncratic preferences of any single judge. The underlying hypothesis is that this variance reduction will enhance the overall reproducibility and alignment with true quality, a hypothesis we validate empirically against human judgments in our experiments.

### Consensus-Driven Training

Let the mini-batch  $\mathcal{B}_b$  contain *all* pairwise comparisons issued by a single judge  $k_b$  in one forward pass. After applying the adaptive update rule, we obtain a vector of per-model Elo scores  $\mathbf{R}^{(b)} \in \mathbb{R}^N$  for the batch. Across  $B$  such judge-homogeneous batches, we first compute the consensus anchor using  $\bar{\mathbf{R}} = \frac{1}{B} \sum_{b=1}^B \mathbf{R}^{(b)}$  and  $\mathbf{G} = \frac{1}{B} \sum_{b=1}^B \mathbf{R}_{(b)}^{\text{base}}$ .

where  $\mathbf{R}_{(b)}^{\text{base}}$  are the scores produced by the baseline Elo

rule with fixed  $K = 32$ . We use  $\mathbf{G}$  as our optimization target. It is important to clarify that we are aware of the potential biases within the baseline Elo system. However, in a fully unsupervised setting, the consensus score  $\mathbf{G}$  represents the most practical and accessible proxy for a stable, judge-agnostic ranking. Our framework is designed to learn a corrective function that maps raw comparisons to debiased outcomes, using the noisy but stable signal from  $\mathbf{G}$  as its guide.

The training objective is:

$$\mathcal{L}(\theta) = \alpha \sum_b \|\mathbf{R}^{(b)} - \mathbf{G}\|_2^2 + \sigma \|\bar{\mathbf{R}} - \mathbf{G}\|_2^2 + \beta \sum_b (1 - \rho_p(\mathbf{R}^{(b)}, \mathbf{G})) \quad (3)$$

with  $\rho_p(\cdot, \cdot)$  the Pearson correlation. Minimising Eq. (3) simultaneously (i) pulls each judge-specific trajectory toward the baseline consensus, (ii) maximises the linear rank correlation with that consensus, and (iii) enforces collective agreement across all judges. The ultimate success of this approach is evaluated by its ability to improve correlation with external, human-annotated ground truth.

### Adaptive Elo Updating Rule

After training, judge  $k$  updates its scores with

---

**Algorithm 2: UDA Adaptive Elo Update**


---

- Require:** responses  $a_i, a_j$ , judge  $k$ , current scores  $R_i^{(k)}, R_j^{(k)}$
- compute features  $\phi_{ij}$  (Algorithm 1)
- $(K_{ij}, s_i, s_j) \leftarrow f_\theta(\phi_{ij})$  {forward pass}
- $\hat{p}_{ij}^{(k)} \leftarrow \sigma(R_i^{(k)} - R_j^{(k)})$
- update
- $R_i^{(k)} \leftarrow R_i^{(k)} + K_{ij}(s_i - \hat{p}_{ij}^{(k)})$
- $R_j^{(k)} \leftarrow R_j^{(k)} + K_{ij}(s_j - (1 - \hat{p}_{ij}^{(k)}))$
- 

### Theoretical Motivation for Consensus Alignment

Our core claim is that aligning diverse judges towards a consensus reduces aggregate system bias. To provide theoretical motivation for this approach, we analyze a simplified, idealized model of the alignment process.

**Theorem 0.1** (Principle of Aggregate Bias Reduction). *Let  $R_i^*$  be the unknown true Elo score for model  $i$ . Let  $R_i^{(k)}$  be the score assigned by judge  $k$ , which includes a bias term  $\epsilon_i^{(k)} = R_i^{(k)} - R_i^*$ . The UDA procedure, by optimizing each judge’s score towards the consensus, is motivated by the principle that this reduces the total expected absolute bias across all judges. Under an idealized linear shrinkage model of the consensus alignment, for any model  $i$ , the expected aggregate bias after alignment satisfies:  $\sum_{k=1}^M \mathbb{E}[|\epsilon_i^{(k)}(\text{UDA})|] \leq \sum_{k=1}^M \mathbb{E}[|\epsilon_i^{(k)}(\text{base})|]$ .*

*Proof Sketch (Illustrative).* We model the alignment as a direct shrinkage towards the mean, serving as a theoretical mo-

Model	gpt-4o	claude-3.5	glm-4-air	glm-4-plus	glm-4-flash	doubao-1.5pro	qwen-max	deepseek-v3	gemini-2.0-flash	deepseek-r1
Baseline Elo	162.0	154.5	122.6	99.7	118.2	170.1	109.0	149.6	<b>341.9</b>	157.5
UDA Method	<b>71.5</b>	<b>44.8</b>	<b>47.8</b>	<b>42.1</b>	<b>50.3</b>	<b>72.6</b>	<b>58.1</b>	<b>61.1</b>	<b>128.8</b>	<b>71.1</b>
Reduction%	55.8	71.0	61.1	57.8	57.4	57.3	46.7	59.2	62.3	54.9

Table 1: Inter-LLM judge score standard deviation ( $\downarrow$ ) on ArenaHard Dataset.

Model	gpt-4o	claude-3.5	glm-4-air	glm-4-plus	glm-4-flash	doubao-1.5pro	qwen-max	deepseek-v3	gemini-2.0-flash	deepseek-r1
Baseline Elo	310.9	119.3	91.3	98.6	161.5	250.1	182.0	165.6	131.1	429.1
UDA Method	<b>125.0</b>	<b>28.3</b>	<b>56.3</b>	<b>37.2</b>	<b>80.4</b>	<b>64.8</b>	<b>61.8</b>	<b>77.4</b>	<b>52.7</b>	<b>125.8</b>
Reduction%	59.8	76.3	38.3	62.2	50.2	74.1	66.0	53.3	59.8	70.7

Table 2: Inter-LLM judge score standard deviation ( $\downarrow$ ) on HUMAN-ANNOTATED TRANSFER SET.

tivation for the consensus-driven loss in Eq. (3). Let the updated bias be a convex combination of the original bias and the average bias:  $\epsilon_i^{(k)}(\text{UDA}) = \alpha \epsilon_i^{(k)} + (1 - \alpha) \bar{\epsilon}_i$ .

By applying the triangle inequality and Jensen’s inequality, we can show that the sum of absolute biases after this alignment step is bounded by the original sum:

$$\sum_{k=1}^M |\epsilon_i^{(k)}(\text{UDA})| \leq \sum_{k=1}^M \left( \alpha |\epsilon_i^{(k)}| + (1 - \alpha) |\bar{\epsilon}_i| \right) \leq \sum_{k=1}^M |\epsilon_i^{(k)}|$$

This idealized model illustrates the underlying principle: while an individual, well-calibrated judge’s accuracy might slightly decrease (a trade-off we observe empirically), the collective variance and aggregate bias across all judges are provably reduced. We emphasize that this is a simplified linear model intended to provide intuition, while the full UDA framework employs a learned non-linear function  $f_\theta$  to achieve this alignment adaptively. The full algebraic details are provided in Appendix. ■

## Experiments

### Experimental Setup

**Metrics. Inter-Judge Standard Deviation:** The standard deviation of a model’s scores across all LLM judges. Lower values signify stronger inter-judge agreement and reduced bias, measuring UDA’s success in curbing score dispersion.

**Pearson Correlation with Human Judgments:** The correlation between a method’s generated ranking and the human-annotated ground-truth. Higher values indicate closer alignment with human preferences. **Alignment with Baseline Consensus (for analysis):** The Pearson correlation and MSE between a judge’s UDA score and the baseline Elo consensus (mean across judges). These internal metrics verify that our training objective successfully pulls individual judges toward a stable collective anchor.

**Datasets.** We evaluate on two datasets: **ArenaHard:** A subset of 500 prompts from the ArenaHard benchmark, with responses from 10 diverse LLMs (e.g., GPT-4o, DeepSeek-V3). The same 10 models act as judges, yielding 450,000

pairwise judgments. **Human-Annotated Transfer Set:** A new dataset of 100 open-ended prompts with human preference labels, designed to test zero-shot generalization. The data pipeline is identical to ArenaHard.

**Training Details.** The adapter is a 3-layer MLP trained with AdamW. We use a learning rate scheduler and select the best model based on validation loss. Full hyperparameters are in the Appendix.

**Baselines.** The main baseline is the **Baseline Elo** with a fixed K-factor ( $K = 32$ ). We compare it against our full **UDA** method. The **Consensus** score, our optimization target, is evaluated for its correlation with human judgments.

### Experiment Results

**Cross-Model Variance Reduction on ArenaHard.** Table 1 reports the per-model standard deviation of Elo scores across ten distinct judge LLMs on ArenaHard. UDA significantly reduces the average inter-judge standard deviation from 158.5 to 64.8, a **relative decrease of 59.1%**. The most volatile model, *gemini-2.0-flash*, exhibits the largest absolute shrinkage ( $341.9 \rightarrow 128.8$ ), confirming UDA’s efficacy against extreme judge-specific bias.

**Visual Consistency on ArenaHard.** Figure 3 displays the score trajectories under both methods. With baseline Elo (left), divergent priors yield highly inconsistent curves with score envelopes exceeding 340 points. UDA (right) compresses this envelope to within 130 points, demonstrating the emergence of a stable consensus. This consistent reduction in variance verifies that UDA successfully suppresses judge-specific priors, yielding more stable rankings across evaluators.

**Zero-Shot Transfer to Human-Annotated Dataset.** We applied the UDA adapter trained solely on ArenaHard directly to the HUMAN-ANNOTATED TRANSFER SET without any retraining to evaluate its zero-shot generalization.

**Variance Reduction and Correlation with Human Judgments.** Even on unseen prompts, UDA proves highly effective. As shown in Table 2, it reduces the average




Figure 3: Score stability across ten judge llms. Left: Baseline Elo. Right: UDA. Our method markedly aligns scores across diverse LLM judges, yielding significantly lower inter-judge variance.

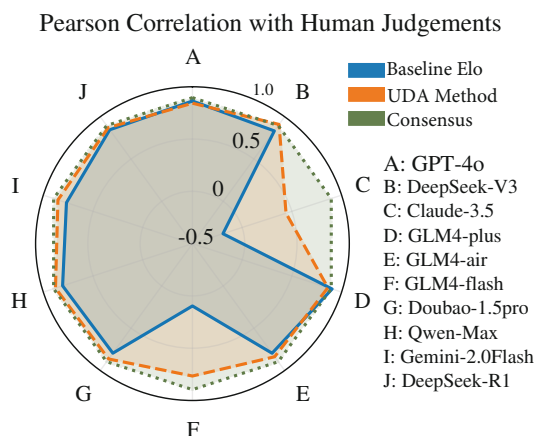


Figure 4: Per-model Pearson correlation ( $\uparrow$ ) between the judge scores computed with the baseline method, UDA method and consensus result, respectively, and the human-annotated ground-truth scores. Notably, consensus scores correlate best with human judgments; the uniform correlation of 0.89 arises because the consensus is shared across all judges, irrespective of individual model participation. This validates the consensus as a robust optimization target.

inter-judge standard deviation from 193.95 to 70.97, a **63.4% relative decrease**. More critically, Figure 4 shows that this increased agreement translates to better alignment with human preferences. UDA (orange markers) consistently outperforms the baseline (blue markers), raising the average Pearson correlation from 0.651 to 0.812 (**+24.7%**). Notably, it elevates even low-capacity judges (e.g. `glm-4-flash`) to reach near-parity with high-quality ones (e.g. `deepseek-r1`) after debiasing. Furthermore, the consensus result exhibits the strongest human correlation, validating its use as an effective optimization objective.

**Qualitative Heatmap Analysis.** Figure 5 offers a qualitative view of the alignment by heatmaps of judge scores. The **Baseline Elo** map (left) shows pronounced vertical stripes,

indicating strong, judge-induced distortions where a single judge rates models differently from its peers. In contrast, the **UDA Method** (right) yields a much smoother pattern where the columnar structure aligns more closely with human annotations, confirming effective cross-judge calibration and alignment with a more objective standard.

Method	Avg. Inter-Judge Std. Dev. ( $\downarrow$ )	Human Pearson Corr. ( $\uparrow$ )
Baseline Elo	193.95	0.651
UDA (Full)	70.97	<b>0.812</b>
UDA (Ablated)	<b>65.64</b>	0.510
Reduction (Full UDA)	63.4%	<b>+24.7%</b>
Reduction (Ablated)	<b>66.2%</b>	-21.6%

Table 3: Ablation study results on the HUMAN-ANNOTATED TRANSFER SET. Removing self-awareness features (UDA (Ablated)) further reduces score variance but critically degrades correlation with human judgments, highlighting their essential role in achieving meaningful alignment.

**Ablation Study: The Critical Role of Self-Awareness Features** To isolate the contribution of our key innovation—features derived from the judge’s own response—we conducted a critical ablation study. We trained and evaluated a variant of UDA, hereafter referred to as **UDA (Ablated)**, where all features related to the judge’s own answer embedding ( $\mathbf{e}_k$ ) were removed from the input vector  $\phi_{ij}^{(k)}$ . This reduced the feature space dimensionality from 4617 to 1539 and effectively rendered the model blind to the judge’s stylistic priors. The model was trained using the same consensus-driven objective.

The results, summarized in Table 3, reveal a fascinating trade-off. On one hand, the ablated model achieved an even greater reduction in inter-judge standard deviation, cutting it by **66.2%** (from 193.95 to 65.64). This indicates that, when stripped of self-awareness features, the model becomes exceptionally effective at its direct optimization goal: forcing all judges to agree with each other.




Figure 5: Judge-score heat-maps on HUMAN-ANNOTATED TRANSFER SET. After refinement with UDA, the scores judged by different LLM judges converge markedly, yielding a visibly narrower chromatic variance within each column of the heatmap.

However, this enhanced consensus came at a steep price. The average Pearson correlation with human judgments plummeted from **0.651** (Baseline) to **0.510** for the ablated model, a significant decrease. In contrast, the full UDA model, equipped with self-awareness features, boosted this correlation to **0.812**.

This finding provides powerful evidence that the self-awareness features are not merely beneficial; they are **essential for grounding the debiasing process**. Without them, the consensus-driven optimization risks converging to a "false consensus"—a state where judges agree with each other more strongly, but their collective judgment diverges from a human-aligned ground truth. In essence, knowing its own stylistic preferences allows a judge (and by extension, the UDA framework) to properly discount them, thereby steering the collective evaluation towards a more objective and meaningful alignment.

## Analysis and Discussion

In this section, we analyze **UDA** in light of recent advances in LLM-as-a-judge calibration and contrast its empirical behavior with the theoretical claims made in Section .

### Methodological Novelty in Context

**UDA** revises the canonical Elo update in two orthogonal directions that jointly attack *judge-induced inconsistency*. First, baseline Elo applies a *global, static*  $K$ , implicitly assuming homoscedastic noise. In practice, this is violated by LLM judges whose priors correlate with their own outputs. **UDA** replaces the constant  $K$  with an *instance-adaptive*  $K_{ij}$  learned from human-free features that capture both inter-answer relationships and self-preference signals. Consequently, likely biased comparisons are down-weighted.

Second, **UDA** replaces the raw win label  $r_{ij}^{(k)}$  with learned *soft* labels  $(s_i, s_j)$  that represent a bias-corrected win probability. This entire procedure is trained unsupervised, guided only by minimizing dispersion among all judges' Elo trajectories to distill a robust latent consensus.

### Empirical Fidelity

Our experiments confirm that UDA's consensus-driven objective successfully translates to empirical gains. The training loss (Eq. 3) guides the model to reduce inter-judge variance while maintaining rank fidelity to a stable consensus anchor. As evidenced by zero-shot transfer results (detailed in Appendix), UDA significantly improves alignment with the baseline consensus across both Pearson correlation and MSE metrics, confirming that the optimization effectively suppresses judge-specific noise.

Crucially, this reduction in variance corresponds to a greater alignment with human evaluators. UDA boosts the average Pearson correlation with human judgments by 24.7%, democratizing evaluation by elevating weaker judges to parity with stronger ones (Figure 4). The dramatic compression of score variance, as visualized in Figure 3, empirically validates the principle motivating Theorem 1: pulling scores toward a consensus reduces system-wide variance and produces more reproducible, judge-invariant rankings.

## Conclusion

We introduced **UDA**, an unsupervised, plug-and-play framework that mitigates self-preference bias in LLM-as-a-judge by dynamically calibrating the Elo update rule. By learning an instance-level K-factor and soft win-probabilities from lightweight, human-free features, UDA distills a consensus ranking that is simultaneously *judge-invariant* and *human-aligned*. Extensive experiments on two benchmarks confirm that UDA (i) shrinks inter-judge score dispersion by 59–63%, (ii) raises average Pearson correlation with human labels by 25%, and (iii) elevates low-capacity judges to parity with state-of-the-art models without retraining. Theoretical analysis supports the bias reduction under mild assumptions. Future work will explore richer feature spaces, weighted consensus objectives, and sub-quadratic approximations to scale beyond hundreds of judges. Code and data are publicly released to foster reproducible LLM evaluation.

## Acknowledgments

This work was completed during Yang Zhang’s internship at ZhipuAI. We would like to express sincere gratitude to the ZhipuAI data annotation team, including Zikang Wang, Xiaohan Zhang and others, for their diligent efforts in curating the human-annotated datasets that made this research possible.

## References

- Chang, Y.; Wang, X.; Wang, J.; Wu, Y.; Yang, L.; Zhu, K.; Chen, H.; Yi, X.; Wang, C.; Wang, Y.; Ye, W.; Zhang, Y.; Chang, Y.; Yu, P. S.; Yang, Q.; and Xie, X. 2023. A Survey on Evaluation of Large Language Models. arXiv:2307.03109.
- Chiang, W.-L.; Zheng, L.; Sheng, Y.; Angelopoulos, A. N.; Li, T.; Li, D.; Zhang, H.; Zhu, B.; Jordan, M.; Gonzalez, J. E.; and Stoica, I. 2024. Chatbot Arena: An Open Platform for Evaluating LLMs by Human Preference. arXiv:2403.04132.
- Gray, A.; Rahat, A.; Crick, T.; Lindsay, S.; and Wallace, D. 2022. Using Elo Rating as a Metric for Comparative Judgment in Educational Assessment. arXiv:2204.01805.
- Gu, J.; Jiang, X.; Shi, Z.; Tan, H.; Zhai, X.; Xu, C.; Li, W.; Shen, Y.; Ma, S.; Liu, H.; Wang, S.; Zhang, K.; Wang, Y.; Gao, W.; Ni, L.; and Guo, J. 2025. A Survey on LLM-as-a-Judge. arXiv:2411.15594.
- Hunter, D. R. 2004. MM algorithms for generalized Bradley-Terry models. *The Annals of Statistics*, 32(1): 384 – 406.
- Koo, R.; Lee, M.; Raheja, V.; Park, J. I.; Kim, Z. M.; and Kang, D. 2024. Benchmarking Cognitive Biases in Large Language Models as Evaluators. arXiv:2309.17012.
- Liu, S.; Gemp, I.; Marris, L.; Piliouras, G.; Heess, N.; and Lanctot, M. 2025a. Re-evaluating Open-ended Evaluation of Large Language Models. arXiv:2502.20170.
- Liu, Z.; Li, J.; Zhuang, Y.; Liu, Q.; Shen, S.; Ouyang, J.; Cheng, M.; and Wang, S. 2025b. am-ELO: A Stable Framework for Arena-based LLM Evaluation. arXiv:2505.03475.
- Minaee, S.; Mikolov, T.; Nikzad, N.; Chenaghlu, M.; Socher, R.; Amatriain, X.; and Gao, J. 2025. Large Language Models: A Survey. arXiv:2402.06196.
- Panickssery, A.; Bowman, S. R.; and Feng, S. 2024. LLM Evaluators Recognize and Favor Their Own Generations. arXiv:2404.13076.
- Thakur, A. S.; Choudhary, K.; Ramayapally, V. S.; Vaidyanathan, S.; and Hupkes, D. 2025. Judging the Judges: Evaluating Alignment and Vulnerabilities in LLMs-as-Judges. arXiv:2406.12624.
- Wang, P.; Li, L.; Chen, L.; Cai, Z.; Zhu, D.; Lin, B.; Cao, Y.; Liu, Q.; Liu, T.; and Sui, Z. 2023. Large Language Models are not Fair Evaluators. arXiv:2305.17926.
- Wang, V.; Zhang, M. J. Q.; and Choi, E. 2025. Improving LLM-as-a-Judge Inference with the Judgment Distribution. arXiv:2503.03064.
- Wataoka, K.; Takahashi, T.; and Ri, R. 2025. Self-Preference Bias in LLM-as-a-Judge. arXiv:2410.21819.
- Xu, W.; Zhu, G.; Zhao, X.; Pan, L.; Li, L.; and Wang, W. Y. 2024. Pride and Prejudice: LLM Amplifies Self-Bias in Self-Refinement. arXiv:2402.11436.
- Ye, J.; Wang, Y.; Huang, Y.; Chen, D.; Zhang, Q.; Moniz, N.; Gao, T.; Geyer, W.; Huang, C.; Chen, P.-Y.; Chawla, N. V.; and Zhang, X. 2024. Justice or Prejudice? Quantifying Biases in LLM-as-a-Judge. arXiv:2410.02736.
- Zhao, W. X.; Zhou, K.; Li, J.; Tang, T.; Wang, X.; Hou, Y.; Min, Y.; Zhang, B.; Zhang, J.; Dong, Z.; Du, Y.; Yang, C.; Chen, Y.; Chen, Z.; Jiang, J.; Ren, R.; Li, Y.; Tang, X.; Liu, Z.; Liu, P.; Nie, J.-Y.; and Wen, J.-R. 2025. A Survey of Large Language Models. arXiv:2303.18223.
- Zheng, L.; Chiang, W.-L.; Sheng, Y.; Zhuang, S.; Wu, Z.; Zhuang, Y.; Lin, Z.; Li, Z.; Li, D.; Xing, E. P.; Zhang, H.; Gonzalez, J. E.; and Stoica, I. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. arXiv:2306.05685.
- Zhou, H.; Huang, H.; Long, Y.; Xu, B.; Zhu, C.; Cao, H.; Yang, M.; and Zhao, T. 2024. Mitigating the Bias of Large Language Model Evaluation. arXiv:2409.16788.