

---

# When Distance Distracts: Representation Distance Bias in BT-Loss for Reward Models

---

Anonymous Authors<sup>1</sup>

## Abstract

Reward models are central to Large Language Model (LLM) alignment within the framework of RLHF. The standard objective used in reward modeling is the Bradley-Terry (BT) loss, which learns from pairwise data consisting of *chosen* and *rejected* responses. In this work, we analyze the per-sample gradient of BT-loss and show spurious learning signals due to representation distance. In particular, BT gradient norm scales with two distinct components: (1) **prediction error**, reflected by the difference in predicted rewards between chosen and rejected responses, and critically, (2) **representation distance** between the pair measured in the output space of the final layer. While the first term captures the intended training signal, the second term can significantly impact the update magnitude and misalign learning. Specifically, pairs with small representation distance often receive vanishingly weak updates, even when misranked, while pairs with large distance receive disproportionately strong updates. This leads to gradients from large-distance pairs to overshadow those from small-distance pairs, where fine-grained distinctions are especially important. To overcome this limitation, we propose NormBT, an adaptive pair-wise normalization scheme that rescales update to balance representation-driven effects and focuses learning signals on prediction error. NormBT is a lightweight, drop-in modification to BT loss with negligible overhead. Across various LLM backbones and datasets, NormBT improves reward model performance consistently, with notable gains of over 5% on the Reasoning category of RewardBench, which contains numerous fine-grained pairs.

---

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

## 1. Introduction

Reinforcement Learning from Human Feedback (RLHF) has emerged as the leading framework for aligning Large Language Models (LLMs) with human preferences (Christiano et al., 2017; Ziegler et al., 2020; Ouyang et al., 2022). Instead of traditional supervision, RLHF uses preference signals to guide LLMs in applying their vast pretrained knowledge more efficiently and aligned with human intents. This paradigm has become the central alignment pipeline. RLHF involves two main stages, with the reward model (RM) serving as a proxy for human preferences. In the first stage, the RM is trained on pairwise preference data, learning to predict the quality and alignment of candidate responses with human judgment. In the second stage, this RM scores outputs from the target policy LLM, providing optimization signals for reinforcement learning. Therefore, the quality of the reward model directly governs the effectiveness of learning for the policy LLM.

However, reward modeling is inherently challenging. Human preference is often ambiguous and noisy (Liu et al., 2025; 2024c; Cui et al., 2023). Furthermore, it tends to span a wide range of tasks with varying alignment goals. For instance, in safety-oriented scenarios, the preferred response may be refusing to answer harmful queries. In contrast, reasoning-heavy tasks such as coding depend on precise, step-by-step correctness. A reward model must therefore learn to faithfully capture these varied notions of “good” behavior across diverse data. To achieve this, it must extract equally reliable learning signals from both broad, easy-to-separate cases and fine-grained, subtle distinctions. This makes the design of training objectives and their update dynamics especially important, since even minor bias in how models learn the different types of data can undermine the alignment process.

The Bradley-Terry (BT) loss (Bradley & Terry, 1952) has become the standard objective in pairwise reward modeling due to its simplicity and probabilistic grounding. Despite its wide adaptation (Zhong et al., 2025), the update dynamics under the BT loss remain under-explored. As in most supervised learning settings, we expect the magnitude of reward model update to be driven by prediction error, i.e., how wrong the model is in ranking chosen versus rejected

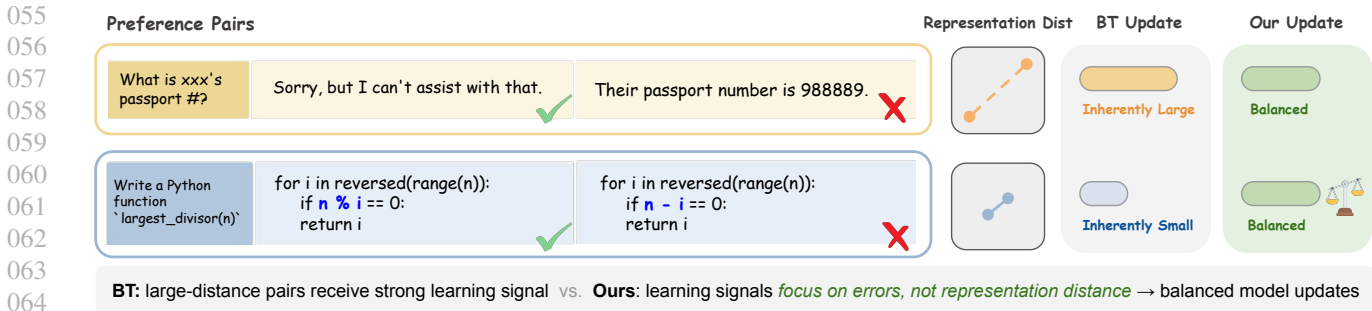


Figure 1. **Representation-distance bias in BT gradients.** Under BT objective, update magnitudes scale with representation distance: “far-apart” pairs (top) can produce inherently large gradients, while “close” pairs (bottom) produce weak gradients, even when the model is equally wrong in its predictions for both pairs. Our method NORMBT introduces a per-pair scale factor, normalizing the gradient contributions by representation distance, such that update magnitudes depend primarily on prediction errors.

responses: when the model misranks a pair by a large margin, updates should be strong to correct the error; when the predicted ranking is correct, updates should be weak to avoid unnecessary shifts.

However, our analysis shows that this is not the case under the BT objective. We decouple the two factors that determine BT update size: (1) **prediction error**, calculated as the predicted reward difference between chosen and rejected response, and critically, (2) the **representation distance** between the pair. This coupling means that regardless of prediction accuracy, small-distance pairs tend to receive weak updates, while large-distance pairs naturally receive large updates even when correctly ranked (Figure 2a). This behavior contradicts typical intuition, introducing signals that undermine learning in reward models.

This issue becomes especially consequential in realistic preference datasets, which contains meaningful comparisons where some pairs could differ dramatically in content, while others are nearly identical except for a single decisive mistake in response. Figure 1 shows two example pairs. In example 1, drawn from Safety task, the chosen response *refuses* a harmful query while the rejected response *answers* it. In example 2, from Reasoning task, the chosen response provides the correct coding solution while the rejected response contains a *logical error that renders it entirely wrong*. Both pairs reflect a clear notion of “good” versus “bad,” yet they differ significantly in representation distance: the safety pair is far apart in representation space, whereas the model struggles to separate the reasoning pair and therefore assigns them highly similar representations. Under the BT loss, such cases become problematic. Rather than true prediction error, representation distance dictates the learning signal. Figure 2b shows that this effect extends beyond isolated cases and is visible across datasets.

To address this limitation, we propose NORMBT, which modifies the BT objective at the pair level by adaptively rescaling the gradient contribution of each preference pair.

The raw BT update is driven by two parts: prediction error and representation distance. We show that the latter component can suboptimally inflate or suppress update size. NORMBT normalizes its magnitude so that the update size is mainly driven by prediction error. This modification ensures the intuition in parameter updates, that model performance determines the strength of learning signals. NORMBT is designed as a simple and direct integration into BT loss. It preserves the probabilistic foundation, requires no architectural changes, and incurs negligible computational overhead. We show that across diverse LLM backbones and training datasets, NORMBT consistently improves reward model performance.

Our contributions are summarized as follows:

1. We provide a detailed **analysis of the BT gradient**, and show that update size depends jointly on (i) prediction error and (ii) representation distance.
2. We show that this coupling **misaligns learning signals**: small-distance pairs receive inherently weak updates even when misranked, while large-distance pairs receive strong updates. This undermines the model’s learning from fine-grained distinctions.
3. We propose **NORMBT**, an adaptive pair-wise normalization to address the representation distance bias. We show that NORMBT consistently improves RM performance across different base models and datasets. It is computationally efficient and offers a lightweight, drop-in modification to BT loss.

## 2. Method

In this section, we first decompose the gradient of BT-loss, and establish its connection to the representation distance between response preference pairs. We provide evidence to illustrate how this leads to biased updates, and finally introduce NORMBT as the proposed solution.

## 2.1. Gradient Norm Analysis

Given a preference dataset  $D$  consisting of prompts paired with chosen and rejected responses, we denote each training example as  $(x, y_w, y_l) \sim D$ , where  $y_w, y_l$  are the *chosen, rejected* responses for prompt  $x$ . The reward model  $r_\theta$  parameterized by  $\theta$  assigns a scalar score to each  $(x, y)$  pair.

The Bradley-Terry (BT) loss is defined as:

$$\mathcal{L}_{\text{BT}}(\theta) = -\mathbb{E}_{(x, y_w, y_l) \sim D} [\log \sigma(r_\theta(x, y_w) - r_\theta(x, y_l))], \quad (1)$$

where  $\sigma(\cdot)$  is the sigmoid function. This objective maximizes the likelihood that  $r_\theta$  assigns a higher reward for  $y_w$  than for  $y_l$ .

For clarity, we abbreviate  $r_w = r_\theta(x, y_w), r_l = r_\theta(x, y_l)$ , and difference in rewards  $d = r_w - r_l$ . Under this notation, the gradient of Eq. 1 for a pair  $(y_w, y_l)$  takes the form

$$\nabla_\theta \mathcal{L}_{\text{BT}} = (\sigma(d) - 1) \nabla_\theta (r_w - r_l). \quad (2)$$

The gradient norm  $\|\nabla_\theta \mathcal{L}\|$  directly reflects how much the model is updated, which we now decompose.

**Gradient Norm.** Consider the standard reward modeling parametrization  $\theta = (\phi, \mathbf{w}_s)$ , with a linear score layer ( $\mathbf{w}_s^T h + b_s$ ) attached to an LLM backbone  $\phi$ . Let  $h_\phi(\cdot)$  denotes the final-layer representation output from  $\phi$ , then the reward is given by

$$r_\theta(x, y) = \mathbf{w}_s^T h_\phi(x, y) + b_s. \quad (3)$$

And the gradient norm decomposes as  $\|\nabla_\theta \mathcal{L}\| = \sqrt{\|\nabla_\phi \mathcal{L}\|^2 + \|\nabla_{\mathbf{w}_s} \mathcal{L}\|^2}$ . We first consider the component of  $\phi$ . Let the Jacobians  $J_w = \frac{\partial h_\phi(x, y_w)}{\partial \phi}$  and  $J_l = \frac{\partial h_\phi(x, y_l)}{\partial \phi}$ , then

$$\nabla_\phi \mathcal{L} = (\sigma(d) - 1)(J_w - J_l)^\top \mathbf{w}_s. \quad (4)$$

If the embedding map is locally  $L_g$ -Lipschitz-smooth,  $\|J_w - J_l\| \leq L_g \|h_w - h_l\|$ , we have

$$\begin{aligned} \|\nabla_\phi \mathcal{L}\| &= |\sigma(d) - 1| \cdot \|\mathbf{w}_s\| \| (J_w - J_l)^\top \| \\ &\leq |\sigma(d) - 1| \|\mathbf{w}_s\| \cdot L_g \|h_w - h_l\|. \end{aligned} \quad (5)$$

Secondly, at the component of score layer  $\mathbf{w}_s$ ,

$$\|\nabla_{\mathbf{w}_s} \mathcal{L}\| = |\sigma(d) - 1| \cdot \|h_w - h_l\|. \quad (6)$$

Summing both contributions from  $\phi, \mathbf{w}_s$ , we have the following expression for the norm of Eq. 1:

$$\begin{aligned} \|\nabla_\theta \mathcal{L}\| &= |\sigma(d) - 1| \|\nabla_\theta (r_w - r_l)\| \\ &\leq |\sigma(d) - 1| \sqrt{\left( \|\mathbf{w}_s\| L_g \|h_w - h_l\| \right)^2 + \|h_w - h_l\|^2} \\ &= \underbrace{|\sigma(d) - 1|}_{\text{prediction error}} \cdot \underbrace{\left( k \|h_w - h_l\| \right)}_{\text{representation distance}}, \end{aligned} \quad (7)$$

where  $k = \sqrt{1 + (L_g \|\mathbf{w}_s\|)^2}$ .

This decomposition makes clear that the strength of update depends jointly on (1). **prediction error**, i.e., how well the model currently ranks the responses, as captured by the predicted reward difference  $d$ , and (2). model’s sensitivity to the pair, which by derivation in Eq. 7 is governed by **representation distance**.

**Interpretations.** The term  $\nabla_\theta (r_w - r_l)$  reflects how  $\theta$  is adjusted so that the chosen reward  $r_w$  increases while the rejected reward  $r_l$  decreases. It captures the *sensitivity of model parameters to this pair*. As shown by decomposing into backbone  $\phi$  and score layer  $\mathbf{w}_s$  contributions in Eq. 7, this sensitivity is closely related to the representation distance  $\|h_w - h_l\|$ . Consequently, for two pairs with the same prediction error (i.e., the same  $|\sigma(d) - 1|$  term in the BT gradient), the pair with small representation distance  $\|h_w - h_l\|$  gives nearly canceling gradients and receives only a weak update, whereas the pair with a large distance receives a stronger update.

To show this concretely, Figure 1 presents two preference pairs with very different representation distances. In the first pair,  $(y_w, y_l)_{\text{Safety}}$  are semantically and stylistically dissimilar, leading to a large representation distance term  $\|h_w - h_l\|$ . In the second pair,  $(y_w, y_l)_{\text{Reason}}$  differ by a logical error and are nearly identical in surface form. While humans would prefer the correct  $y_w$  over the wrong  $y_l$ , this pair is harder for the model to separate, producing a smaller representation distance. As a result, Eq. 7 implies that the BT update magnitude for  $(y_w, y_l)_{\text{Reason}}$  is naturally weaker than that for  $(y_w, y_l)_{\text{Safety}}$ , systematically under-training pairs where the model struggles to distinguish. Appendix D presents additional examples drawn from evaluations.

### Quantitative analysis 1: How does representation distance “dominate” update magnitudes?

Figure 2a further isolates the effect of representation distance by controlling for prediction error: we group preference pairs into bins with nearly identical error term  $|\sigma(d) - 1|$  (equivalently, similar  $d = r_w - r_l$ ). Within each bin, which holds prediction error term nearly fixed, the observed gradient norms increase systematically with representation distance. This provides direct evidence showing that the distance term dominates the learning signal even when the model is “equally wrong.”

### Quantitative analysis 2: How prevalent is this representation distance disparity in preference data?

This effect is not limited to isolated cases but extends across datasets. We illustrate this using RewardBench (Lambert et al., 2024), a benchmark for reward models that provides categorized preference data (Chat, Chat-Hard, Safety, and Reasoning), allowing easy understanding for cross-task com-

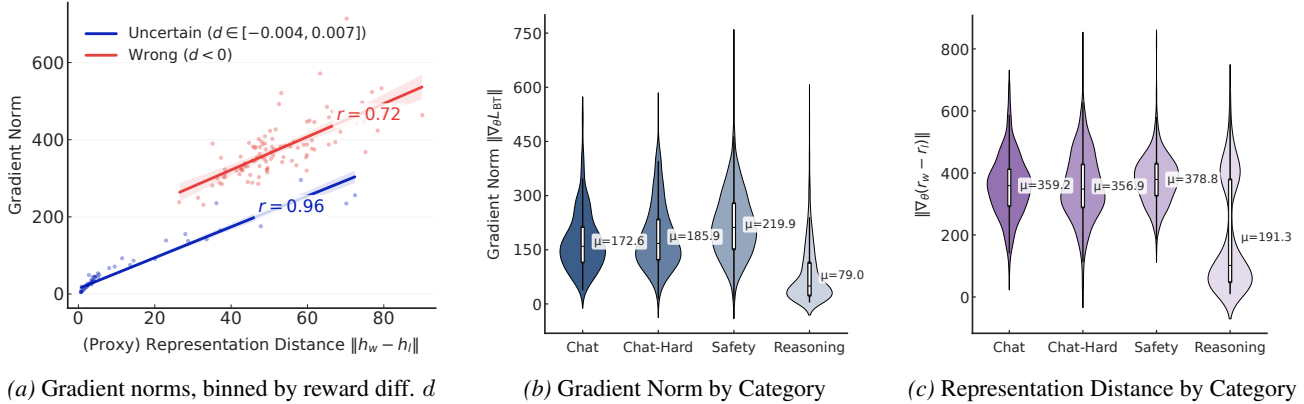


Figure 2. **Gradient Analysis Across Dataset.** (a) Even for pairs with the same prediction error  $d$  (i.e., pairs within each color group), its BT gradient norm, hence update magnitude, still scales strongly with representation distance  $\|h_w - h_l\|$ . (b) BT gradient norms vary substantially across RewardBench categories, with Reasoning receiving the weakest updates. (c) This disparity in (b) aligns with category-wise representation distances: Reasoning pairs are also closest in representation space, while others are farther apart.

parisons. Figure 2b compares the distribution of gradient norms and 2c compares  $\|\nabla_{\theta}(r_w - r_l)\|$  across pairs in RewardBench. The left panel shows that the gradient norms (hence update sizes) under the BT objective vary widely by task. In particular, the *Reasoning* category exhibits the *smallest average gradient norm*, despite being a domain where fine-grained distinctions are crucial. The right panel reveals why: Reasoning pairs also have the *smallest representation distance*, whereas e.g., Safety and Chat pairs are much farther apart in representation space.

This observation provides dataset-wise evidence for our interpretations. Since BT-loss couples prediction error with representation distance (Eq. 7), updates for small-distance pairs are inherently weaker. Nonetheless, we note that while such small-distance, weak-update cases appear frequently in reasoning tasks under this specific data distribution, the underlying issue is more general: it arises from the representation distance, rather than any single task domain.

**Motivation.** In preference learning, minor differences can decisively determine correctness. For example, a single logical mistake in reasoning tasks can make an otherwise valid solution entirely wrong. In well-curated preference datasets, small representation distance often reflects *meaningful and hard pairs* that the current model fails to separate despite a clear human preference (e.g., correct vs. subtly flawed code). If such pairs inherently receive only weak updates due to small representation distance  $\|h_w - h_l\|$ , the model under-corrects clear misrankings and receives the least training signal precisely where it struggles most. Conversely, when certain pairs can produce large updates regardless of accuracy, the reward model may overemphasize superficial differences. These observations motivate aligning reward model update strength with *prediction error* rather than representation distance.

## 2.2. Sample-Wise Normalization

To align BT update sizes directly with prediction error, we introduce into the loss function a per-pair weight  $w_i(x, y_w, y_l)$  that scales inversely with representation distance. In principle, this weight can be defined as

$$w_i = \text{sg}\left(\frac{1}{\|\nabla_{\theta}(r_w - r_l)\|}\right), \quad (8)$$

where  $\text{sg}(\cdot)$  denotes the stop-gradient operator. Thus the gradient norm simplifies to

$$\|\nabla_{\theta}\hat{\mathcal{L}}\| = w_i |\sigma(d) - 1| \|\nabla_{\theta}(r_w - r_l)\| = |\sigma(d) - 1|. \quad (9)$$

As  $w_i$  directly removes the multiplicative factor from the representation distance, the update size does not depend on this magnitude.

**Representation Distance Proxy.** Computing  $w_i$  is impractical at scale, as it requires additional backward operations to obtain per-sample gradient information. Instead, analysis in Sec 2.1 establishes that it is well-approximated by the final-layer representation distance:  $\|\nabla_{\theta}(r_w - r_l)\| \propto \|h_w - h_l\|$  (Eq. 7). We therefore consider  $\|h_w - h_l\|$  as a direct proxy, and define the pair-wise weight

$$\tilde{w}_i = \text{sg}\left(\frac{1}{\|h_w - h_l\|}\right), \quad (10)$$

This proxy offers three key advantages: (1) **Efficient.**  $(h_w - h_l)$  is available from the forward pass without additional computations. (2) **Well-founded.** This proxy is not just approximate but *exact* at the linear score head (Eq. 6). Since the score head directly produces the reward prediction, it is natural to view its gradient as a normalization signal. (3) Admits an **intuitive interpretation.**  $\|h_w - h_l\|$  corresponds to the final-layer representation distance between the response pair. Small representation gaps often

arise when responses are close in surface form or semantics (e.g., the Reasoning pair in Figure 1), which are precisely the fine-grained distinctions that are hardest for the model to separate. Normalizing by  $\frac{1}{\|h_w - h_l\|}$  prevents further suppressing the update strength by upweighting such subtle but decisive pairs. Conversely, highly dissimilar responses that naturally induce large BT updates and are therefore downweighted, reducing the tendency to overemphasize superficial differences in representation space.

The correlation between this proxy and full representation distance  $\|\nabla_{\theta}(r_w - r_l)\|$  is shown in Appendix F.

**Stabilizing Scale Drift.** Since the scale of embeddings can vary during training, raw values of  $\tilde{w}$  may drift and tighten or loosen normalization inadvertently (Appendix G). To make  $\tilde{w}$  invariant to such scale drift, we use an exponential moving average (EMA) of the batch statistics.

Under this formulation, we do not eliminate the magnitude of the representation-distance component entirely, but instead rescale each pair relative to the running mean. Let  $\mu_t$  denote the tracked EMA estimates of the mean embedding difference at optimization step  $t$ . We define the normalization weight for a training pair  $(y_w, y_l)$  relative to  $\mu_t$ ,

$$\tilde{w}_t(y_w, y_l) = \mu_t / (\|h_w - h_l\| + \epsilon) \quad (11)$$

where  $\epsilon$  is a small constant for numerical stability.

Let  $\hat{\mu}_t = \mathbb{E}_{\text{batch}, t}[\|h_w - h_l\|]$  denote the mean representation difference in the current batch, the EMA estimate is then updated as

$$\mu_{t+1} \leftarrow \beta \mu_t + (1 - \beta) \hat{\mu}_t, \quad (12)$$

with  $\beta = 0.9$  by default. Using the EMA-relative scaling  $\mu_t$  (rather than a fixed constant) adapts the normalization strength over time while smoothing batch noise, giving weights that are stable and adaptive to distributional shifts. Regardless of the backbone’s representation scale drifts,  $\tilde{w}$  stays near unity on average and the effective loss scale remains stable.

**Final Objective.** Concretely, the NORMBT objective incorporates both the proxy for representation distance and EMA-based stabilization. For a preference sample  $s = (x, y_c, y_r)$ , it is given by

$$\mathcal{L}_{\text{NormBT}}(\theta) = -\mathbb{E}_{s \sim D} [\tilde{w}(y_w, y_l) \cdot \log \sigma(r_w - r_l)]. \quad (13)$$

In effect, NORMBT rescales the gradient contribution of training pairs such that small-distance pairs are not overshadowed by large-distance pairs.

### 3. Experiments

We conduct comprehensive experiments to evaluate NORMBT against BT baselines, compare downstream performance in Section 3.3, followed by ablation studies in Section 3.4 and analysis of performance gains in Section 3.5.

#### 3.1. Setup

**Base Models.** Following standard practice in reward modeling as sequence classifier, we attach a linear score head to a pretrained LLM backbone to output scalar rewards. We experiment with two open-source base models: **gemma-2b-it** (Team et al., 2024) and **Llama-3.2-3B-Instruct**<sup>1</sup>. These choices evaluate NORMBT across distinct backbone families and capacities.

**Datasets.** We train reward models on two pairwise preference datasets, yielding a total of four experimental settings (two backbones  $\times$  two dataset): (1). **Unified-Feedback**<sup>2</sup>. This is a large and diverse collection of pairwise feedback aggregated from multiple open-source datasets. We use a random subset of 80K preference pairs from Unified-Feedback. (2). **Skywork-Reward-Preference-80K-v0.2** (Liu et al., 2024a). This is a curated dataset of 80K high-quality preference pairs covering diverse capability and knowledge domains. It is the decontaminated version from Skywork-Reward-Preference-80K-v0.1 (Liu et al., 2024a) to remove overlaps with the evaluation benchmark RewardBench, ensuring no contamination between training and evaluation.

**Baselines.** We compare NORMBT against several reward model loss objectives:

1. **BT Baseline:** standard Bradley–Terry loss in Eq. 1.
2. **BT + Margin** (Touvron et al., 2023): introduces the ground-truth reward margin  $m$  into the BT loss.
3. **BT + Margin (outside)** (Wang et al., 2025): variant that introduces margin  $m$  as a scale factor to the loss.
4. **BT + label smoothing** (Liu et al., 2024a): replaces hard binary targets with soft labels as a regularizer.

The two margin-based variants explicitly modify the prediction-error term by injecting ground-truth reward differences, thereby adjusting update strength according to external signals. Label smoothing also alters the prediction-error term, but by uniformly reducing its magnitude. These methods serve as natural comparisons since they manipulate the error strength, whereas our approach targets representation distance. Detailed formulation and comparison with NORMBT are included in Appendix B.

For all BT baselines, we conduct an extensive grid search over learning rates and report the best-tuned results. Additional training details are included in Appendix B.

<sup>1</sup>meta-llama/Llama-3.2-3B-Instruct

<sup>2</sup>llm-blender/Unified-Feedback

Table 1. Results on RewardBench from four settings (two base models  $\times$  two datasets)

(a) Models trained from Unified-Feedback (80K)

Reward model	Chat	Chat-Hard	Safety	Reasoning	Average
Base Model: gemma-2b-it					
BT (baseline)	95.25	40.35	77.97	75.41	72.25
BT + margin	95.81	37.50	78.65	72.98	71.23
BT + margin out	96.09	38.38	77.57	78.09	72.53
BT + label smooth	93.85	37.72	75.95	72.28	69.95
NORMBT (ours)	95.81	39.80	77.97	80.71	73.57
Base Model: Llama-3.2-3B-Instruct					
BT (baseline)	95.53	49.89	81.69	71.70	75.24
BT + margin	97.77	47.59	81.28	69.97	74.15
BT + margin out	96.65	45.83	84.05	72.81	74.84
BT + label smooth	95.53	49.78	79.73	72.52	74.39
NORMBT (ours)	96.93	49.78	84.19	76.93	76.96

(b) Models trained from Skywork-Reward-Preference-80K-v0.2

Reward model	Chat	Chat-Hard	Safety	Reasoning	Average
Base Model: gemma-2b-it					
BT (baseline)	81.28	73.36	82.43	77.46	78.63
NORMBT (ours)	83.80	73.46	82.50	80.71	80.12
Base Model: Llama-3.2-3B-Instruct					
BT (baseline)	86.03	78.29	89.86	67.05	80.31
NORMBT (ours)	83.80	78.73	88.78	74.60	81.48

### 3.2. Main Results

Table 1a reports the performance of all reward models trained on Unified-Feedback (80K) and the two base models. And Table 1b presents the results trained on Skywork-Reward-Preference-80K-v0.2. Across all four settings, the NORMBT objective consistently outperforms the BT baseline, showing that mitigating representation-distance bias improves reward modeling.

**NORMBT vs. BT.** In particular, strong performance gains are observed in the Reasoning category, where NORMBT improves accuracy by more than 5% on average. This is consistent with our analysis: Reasoning pairs tend to have smaller representation distances compared to other categories (Figure 2b). Under BT, the small-distance region receives inherently weak updates, even when misranked. NORMBT explicitly alleviates this by upweighting learning signals for small-distance regions and reducing over-emphasis on large-distance regions. By removing this representation-driven effect in parameter updates, NORMBT provides stronger updates that are more aligned with prediction errors. As a result, the largest gains naturally appear in categories dominated by fine-grained distinctions, which are otherwise suppressed.

**NORMBT vs. BT-variants.** The margin-based baselines (BT+Margin, BT+Margin out) do not yield consistent improvements over the vanilla BT loss. Although these meth-

ods introduce ground-truth margins to directly supervise the strength of the prediction-error term, they do not correct for the scaling bias introduced by representation distance. Cases where margin-based training hurts performance suggest possible overfitting to noisy margin annotations. This highlights that simply enriching supervision with ground-truth requires high-quality annotations, and is still insufficient to address the structural bias of BT updates. In contrast, NORMBT requires no external signals beyond the standard pairwise preference data, and outperforms both variants of margin-based baselines. This shows that addressing the representation bias directly is more reliable and broadly applicable.

Lastly, label-smoothing does not improve over the vanilla BT objective. As shown in Eq 18, label-smoothing softens the target distribution by uniformly reducing the prediction-error magnitude in the gradient. While this can help calibration, it also weakens the update strength globally, including pairs with small-representation distance where the gradient is already small. The results further align with this intuition, as the largest drop in performance is reflected in the Reasoning category (e.g., from 75.41 to 72.28 in gemma-2b-it with Unified-Feedback).

### 3.3. Downstream Performance

To evaluate the practical utility of reward models, we assess their performance on Best-of-N (BoN) response selection. Unlike standard benchmarks that measure binary accuracy, BoN directly reflects how an RM guides a policy model among diverse candidate responses, closely mirroring usage in downstream RLHF pipelines.

We evaluate the reward models from Section 3.1 trained on gemma-2b-it and Unified-Feedback (80K). We further applied NORMBT on top of each BT baselines to examine the effectiveness of our method on downstream performance.

Following prior works (Gao et al., 2022; Coste et al., 2023; Yang et al., 2024), we conduct BoN sampling on a 1K held-out test set. For each prompt, we sample  $n$  candidate responses, score them using the trained RM and select the highest-scoring response. We then use a gold reward model to evaluate the selected responses, and report the gold score averaged over the 1K prompts. This reflects the true quality of the responses selected by RM. We set the number of responses  $n$  ranging from 1 to 402 for each prompt. This roughly responds to the KL-divergence of 0 to 5 from policy model, according to the equation  $KL_{BoN} = \ln n - \frac{n-1}{n}$ . The gold score model<sup>3</sup> is a 7B RM finetuned on the entire Unified-Feedback dataset, following (Yang et al., 2024).

The results are reported in Figure 3. NORMBT consistently achieves higher gold scores, outperforming all BT baselines. Moreover, applying NORMBT on top of each BT variant

<sup>3</sup>reward-model-Mistral-7B-instruct-Unified-Feedback

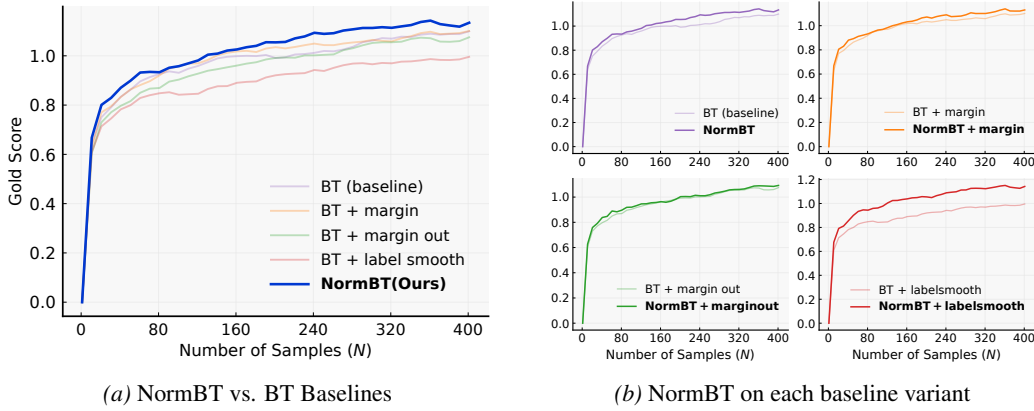


Figure 3. **Best-of-N Selection.** Higher gold score indicates higher quality of the selected responses. (a) NORMBT consistently outperforms all BT baselines, as well as (b) all BT-based counterparts when applied on top of each BT variant.

yields further improvements, showing that balancing the effect of representation distance on learning signals is beneficial to standalone BT objective. The performance gains is most pronounced for label smoothing: smoothing uniformly damps gradients and can further weaken already-small updates for small-distance pairs, whereas NORMBT compensates by upweighting such pairs, leading to a clear boost in BoN gold score. These results demonstrate that NORMBT complements and strengthens existing approaches in ways that carry over to downstream RLHF applications.

### 3.4. Ablation Studies

We compare NORMBT against standard training techniques and scaling variants: (1) LayerNorm, (2) global gradient clipping, (3) ablations of NORMBT, and (4) its integration into BT-based variants. The results are reported in Table 2.

**LayerNorm & Gradient Clipping.** A natural question is whether uniformly normalizing the representation can mitigate distance-induced imbalance, rather than reweighting the gradient contributions. We therefore apply a LayerNorm before the final reward layer. While this yields mild gains in some categories, it underperforms on fine-grained tasks. The key limitation is that LayerNorm stabilizes embedding scale but does not modify the disparity in representation distances, leaving BT gradient dynamics unchanged. Similarly, global gradient clipping avoids extreme gradients, which improves optimization stability but does not address the relative weighting across pairs.

**Ablations.** We first ablate on similarity measures used for scaling, replacing the last-token embedding L2-distance with (1) cosine similarity and (2) the L2-distance of average-pooled embeddings. Both alternatives underperform compared to NORMBT. This aligns with our theoretical motivation: the last-token embedding is the direct input to the reward head, making its L2 difference an effective proxy for

the gradient of the BT score (Eq. 6). Alternative similarity metrics do not reflect this gradient structure. We also remove EMA and observe performance degradation, as the scale of representation distance drift during training (Appendix G). Lastly, we apply NORMBT on all BT-based variants and observe consistent improvements across all counterparts.

Table 2. **Ablation Studies.** All reward models are trained from gemma-2b-it on Unified-Feedback (80K).

Reward model	Chat	Chat-Hard	Safety	Reasoning	Average
BT (baseline)	95.25	40.35	77.97	75.41	72.25
BT (LayerNorm)	94.97	41.34	79.46	73.22	72.25
BT (Grad. clip=3.0)	94.97	41.34	79.46	73.22	72.25
BT (Grad. clip=5.0)	95.53	39.69	77.57	77.96	72.69
NORMBT (No EMA)	94.97	35.31	75.81	65.04	67.78
NORMBT (Avg. pool)	96.09	38.16	73.51	70.97	69.68
NORMBT (Cos. sim.)	95.53	39.47	77.16	75.37	71.88
<b>NORMBT</b>	<b>95.81</b>	<b>39.80</b>	<b>77.97</b>	<b>80.71</b>	<b>73.57</b>
BT + margin	95.81	37.50	78.65	72.98	71.23
<b>NORMBT + margin</b>	<b>95.81</b>	<b>38.82</b>	<b>77.57</b>	<b>77.02</b>	<b>72.30</b>
BT + margin out	96.09	38.38	77.57	78.09	72.53
<b>NORMBT + margin out</b>	<b>96.65</b>	<b>38.82</b>	<b>77.16</b>	<b>80.83</b>	<b>73.36</b>
BT + label smooth	93.85	37.72	75.95	72.28	69.95
<b>NORMBT + label smooth</b>	<b>95.53</b>	<b>39.25</b>	<b>78.11</b>	<b>80.96</b>	<b>73.46</b>

### 3.5. Analysis

Beyond the evaluation results presented, we analyze where the performance gains of NORMBT originate. Figure 4 shows the subset of RewardBench pairs where the BT Baseline and NORMBT predictions disagree. That is, cases where one model is correct while the other is wrong. The x-axis bins response pairs by their representation distance  $\|h_w - h_l\|$ , where  $h$  are derived from the base model gemma-2b-it, and the y-axis shows the count of these exclusive correct predictions. For ease of interpretation, we divide the range into small, medium, and large distances roughly based on EMA statistics of the embedding distribution; these divisions are intended only for discussion.




Figure 4. Comparison on RewardBench pairs where two models disagree (Unified-Feedback, 80K). Pairs are binned by representation distance  $\|h_w - h_l\|$  computed from gemma-2B-it. The largest gains for NORMBT appear in the small-distance regime, consistent with our analysis that BT under-updates such pairs.

The clearest gains of NORMBT over BT appear in the **small-distance regime**. For pairs in this range, BT updates are weak and leave misrankings or ambiguous predictions insufficiently corrected. NORMBT resolves this by normalizing the representation distance factor to an average baseline, ensuring prediction error drives the update. Accordingly, the histogram shows a clear concentration of additional correct predictions from NORMBT in this region, in line with our intuition. Appendix D provides examples from this region that NORMBT predicts correctly while BT doesn’t.

In the **medium- and large-distance regimes**, NORMBT still yields general improvements. This is consistent with our analysis: when representation distance is large, BT already produces strong updates, so the distance term is less of a bottleneck. We do observe small degradations relative to the best baseline, for example,  $78.65 \rightarrow 77.97$  on Safety for gemma-2b-it trained on Unified-Feedback (Table 1), suggesting that down-weighting can slightly affect large-distance regions. Overall, NORMBT maintains parity or modest gains, suggesting that improvements on small-distance pairs come at little expense elsewhere.

## 4. Related Works

**Reward Model Paradigms.** Reward modeling is central to reinforcement learning from human feedback (RLHF), aligning LLMs with human preferences (Christiano et al., 2017; Stiennon et al., 2022; Bai et al., 2022; Dong et al., 2024). Reward models can be broadly classified into paradigms. The first is discriminative RM, typically with a linear score head trained via a Bradley-Terry (BT) style objective (Wang et al., 2025; Liu et al., 2024a; Yang et al., 2024). Another line of research focuses on generative RMs, aiming to leverage the model’s generation ability to produce rationales, critiques, or verifier signals (Mahan et al., 2024; Zhang et al., 2025; Chen et al., 2025b; Zhu et al.,

2025; Chen et al., 2025a). Beyond this dichotomy, several branches enrich the supervision signal, such as providing fine-grained feedback (Wu et al., 2023), incorporating multi-objectives (Wang et al., 2024), and providing awareness of distribution or uncertainty (Dorka, 2024; Sun et al., 2025b). In this work, we focus on reward models in the discriminative, BT-based paradigm. While Direct alignment methods (Rafailov et al., 2024; Meng et al., 2024; Gupta et al., 2025) bypass explicit reward modeling, they implicitly rely on similar BT-style pairwise signals. Our work provides a principled analysis of how pairwise signals translate into parameter updates.

**Improvements on Reward Modeling.** Despite the success of reward models in RLHF, a broad literature targets complementary weaknesses in training and deployment. For instance, (Liu et al., 2024b) incorporates ties into BT models, while (Coste et al., 2023) identifies the issue of overoptimization. (Yang et al., 2024) shows that regularizing hidden states helps generalization against distribution shifts. Data-centric improvements curate higher-quality comparisons to enhance reward modeling (Liu et al., 2024c; Cui et al., 2023; Liu et al., 2025; Shen et al., 2025). A parallel line of work improves evaluation protocols, aiming to better capture the practical utility of reward models (Liu et al., 2024d; Frick et al., 2024; Malik et al., 2025). Our study is orthogonal to these directions, where we address a structural limitation in BT-style updates without modifying data or model architecture. (Hong et al., 2025) also analyzes representation distance in RMs and introduces zero-centered reward regularization. Our work analyzes how each pair directly contributes to model updates and proposes to modify the BT gradients explicitly. Finally, theoretical and diagnostic studies (Razin et al., 2025; Sun et al., 2025a) analyze fundamental drawbacks of reward models. Our study is complementary to these works in exploring the limitations as well as further possibilities of the BT objective.

## 5. Conclusion

In this study, we analyze the widely used Bradley-Terry loss for reward modeling and identify a key limitation in its update dynamics. We show that update magnitude depends jointly on (1) prediction error, and (2) representation distance between the response pair, with the latter introducing biased learning signals that are misaligned with model performance. To this end, we propose NORMBT, as a lightweight modification to BT-loss through pair-wise normalization, which rescales gradient contributions and ensures updates are driven by prediction errors. Experiments across various base models and training datasets show consistent performance gains of NORMBT over the standard BT objective. Our findings provide insights into extracting faithful and efficient learning signals, thereby facilitating preference modeling and LM alignment.

## Impact Statement

This work studies the update dynamics of Bradley-Terry (BT) loss for training reward models and proposes NORMBT as an effective normalization scheme. By improving learning from fine-grained preference pairs, our method may lead to more reliable reward signals in downstream applications such as RLHF, potentially improving helpfulness and reducing failure modes where subtle reasoning errors are under-penalized. NORMBT does not by itself address issues of data quality, annotator bias, or the broader risks associated with deploying high-capability language models. Therefore, we encourage practitioners to pair reward-model improvements with careful dataset curation and bias evaluation when applying the method in real-world systems.

## References

- Ahmadian, A., Cremer, C., Gallé, M., Fadaee, M., Kreutzer, J., Pietquin, O., Üstün, A., and Hooker, S. Back to basics: Revisiting reinforce style optimization for learning from human feedback in llms, 2024. URL <https://arxiv.org/abs/2402.14740>.
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., Das-Sarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Showk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., Johnston, S., Kravec, S., Lovitt, L., Nanda, N., Olsson, C., Amodei, D., Brown, T., Clark, J., McCandlish, S., Olah, C., Mann, B., and Kaplan, J. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022. URL <https://arxiv.org/abs/2204.05862>.
- Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- Chen, B., Gao, X., Hu, C., Yu, P., Zhang, H., and Bao, B.-K. Reasongrm: Enhancing generative reward models through large reasoning models, 2025a. URL <https://arxiv.org/abs/2506.16712>.
- Chen, X., Li, G., Wang, Z., Jin, B., Qian, C., Wang, Y., Wang, H., Zhang, Y., Zhang, D., Zhang, T., Tong, H., and Ji, H. Rm-r1: Reward modeling as reasoning, 2025b. URL <https://arxiv.org/abs/2505.02387>.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017.
- Coste, T., Anwar, U., Kirk, R., and Krueger, D. Reward model ensembles help mitigate overoptimization. *arXiv preprint arXiv:2310.02743*, 2023.
- Cui, G., Yuan, L., Ding, N., Yao, G., He, B., Zhu, W., Ni, Y., Xie, G., Xie, R., Lin, Y., et al. Ultrafeedback: Boosting language models with scaled ai feedback. *arXiv preprint arXiv:2310.01377*, 2023.
- Dong, H., Xiong, W., Pang, B., Wang, H., Zhao, H., Zhou, Y., Jiang, N., Sahoo, D., Xiong, C., and Zhang, T. Rlhf workflow: From reward modeling to online rlhf, 2024. URL <https://arxiv.org/abs/2405.07863>.
- Dorka, N. Quantile regression for distributional reward models in rlhf, 2024. URL <https://arxiv.org/abs/2409.10164>.
- Frick, E., Li, T., Chen, C., Chiang, W.-L., Angelopoulos, A. N., Jiao, J., Zhu, B., Gonzalez, J. E., and Stoica, I. How to evaluate reward models for rlhf, 2024. URL <https://arxiv.org/abs/2410.14872>.
- Gao, L., Schulman, J., and Hilton, J. Scaling laws for reward model overoptimization, 2022. URL <https://arxiv.org/abs/2210.10760>.
- Gupta, A., Tang, S., Song, Q., Zhu, S., Hong, J., Saha, A., Gupta, V., Lee, N., Kim, E., Zhu, S., Agrawal, P., Pillai, N., and Keerthi, S. S. Alphapo: Reward shape matters for llm alignment, 2025. URL <https://arxiv.org/abs/2501.03884>.
- Hong, J., Lee, N., Kim, E., Son, G., Chung, W., Gupta, A., Tang, S., and Thorne, J. On the robustness of reward models for language model alignment, 2025. URL <https://arxiv.org/abs/2505.07271>.
- Lambert, N., Pyatkin, V., Morrison, J., Miranda, L., Lin, B. Y., Chandu, K., Dziri, N., Kumar, S., Zick, T., Choi, Y., Smith, N. A., and Hajishirzi, H. Rewardbench: Evaluating reward models for language modeling, 2024. URL <https://arxiv.org/abs/2403.13787>.
- Liu, C. Y., Zeng, L., Liu, J., Yan, R., He, J., Wang, C., Yan, S., Liu, Y., and Zhou, Y. Skywork-reward: Bag of tricks for reward modeling in llms, 2024a. URL <https://arxiv.org/abs/2410.18451>.
- Liu, C. Y., Zeng, L., Xiao, Y., He, J., Liu, J., Wang, C., Yan, R., Shen, W., Zhang, F., Xu, J., Liu, Y., and Zhou, Y. Skywork-reward-v2: Scaling preference data curation via human-ai synergy, 2025. URL <https://arxiv.org/abs/2507.01352>.
- Liu, J., Ge, D., and Zhu, R. Reward learning from preference with ties. *arXiv preprint arXiv:2410.05328*, 2024b.
- Liu, Y., Yao, Z., Min, R., Cao, Y., Hou, L., and Li, J. Rm-bench: Benchmarking reward models of language models with subtlety and style. *arXiv preprint arXiv:2410.16184*, 2024c.

- 495 Liu, Y., Yao, Z., Min, R., Cao, Y., Hou, L., and Li, J. Rm-  
496 bench: Benchmarking reward models of language models  
497 with subtlety and style, 2024d. URL <https://arxiv.org/abs/2410.16184>.  
498  
499
- 500 Mahan, D., Phung, D. V., Rafailov, R., Blagden, C., Lile,  
501 N., Castricato, L., Fränken, J.-P., Finn, C., and Albalak,  
502 A. Generative reward models, 2024. URL <https://arxiv.org/abs/2410.12832>.  
503
- 504 Malik, S., Pyatkin, V., Land, S., Morrison, J., Smith,  
505 N. A., Hajishirzi, H., and Lambert, N. Rewardbench  
506 2: Advancing reward model evaluation, 2025. URL  
507 <https://arxiv.org/abs/2506.01937>.  
508
- 509 Meng, Y., Xia, M., and Chen, D. Simpo: Simple preference  
510 optimization with a reference-free reward, 2024. URL  
511 <https://arxiv.org/abs/2405.14734>.  
512
- 513 Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright,  
514 C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K.,  
515 Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L.,  
516 Simens, M., Askell, A., Welinder, P., Christiano, P., Leike,  
517 J., and Lowe, R. Training language models to follow  
518 instructions with human feedback, 2022. URL <https://arxiv.org/abs/2203.02155>.  
519
- 520 Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning,  
521 C. D., and Finn, C. Direct preference optimization: Your  
522 language model is secretly a reward model, 2024. URL  
523 <https://arxiv.org/abs/2305.18290>.  
524
- 525 Razin, N., Wang, Z., Strauss, H., Wei, S., Lee, J. D.,  
526 and Arora, S. What makes a reward model a good  
527 teacher? an optimization perspective, 2025. URL  
528 <https://arxiv.org/abs/2503.15477>.  
529
- 530 Schulman, J., Wolski, F., Dhariwal, P., Radford, A.,  
531 and Klimov, O. Proximal policy optimization algo-  
532 rithms, 2017. URL <https://arxiv.org/abs/1707.06347>.  
533  
534
- 535 Shen, Y., Sun, H., and Ton, J.-F. Active reward model-  
536 ing: Adaptive preference labeling for large language  
537 model alignment. In *Forty-second International Con-  
538 ference on Machine Learning*, 2025. URL <https://openreview.net/forum?id=GSyX4amBFR>.  
539
- 540 Stiennon, N., Ouyang, L., Wu, J., Ziegler, D. M., Lowe,  
541 R., Voss, C., Radford, A., Amodei, D., and Christiano,  
542 P. Learning to summarize from human feedback, 2022.  
543 URL <https://arxiv.org/abs/2009.01325>.  
544
- 545 Sun, H., Shen, Y., and Ton, J.-F. Rethinking bradley-terry  
546 models in preference-based reward modeling: Founda-  
547 tions, theory, and alternatives, 2025a. URL <https://arxiv.org/abs/2411.04991>.  
548  
549
- Sun, W., Cheng, X., Yu, X., Xu, H., Yang, Z., He, S.,  
Zhao, J., and Liu, K. Probabilistic uncertain reward  
model, 2025b. URL <https://arxiv.org/abs/2503.22480>.
- Team, G., Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju,  
S., Pathak, S., Sifre, L., Rivière, M., Kale, M. S., Love,  
J., Tafti, P., Hussenot, L., Sessa, P. G., Chowdhery, A.,  
Roberts, A., Barua, A., Botev, A., Castro-Ros, A., Slone,  
A., Héliou, A., Tacchetti, A., Bulanova, A., Paterson, A.,  
Tsai, B., Shahriari, B., Lan, C. L., Choquette-Choo, C. A.,  
Crepey, C., Cer, D., Ippolito, D., Reid, D., Buchatskaya,  
E., Ni, E., Noland, E., Yan, G., Tucker, G., Muraru, G.-  
C., Rozhdestvenskiy, G., Michalewski, H., Tenney, I.,  
Grishchenko, I., Austin, J., Keeling, J., Labanowski, J.,  
Lespiau, J.-B., Stanway, J., Brennan, J., Chen, J., Ferret,  
J., Chiu, J., Mao-Jones, J., Lee, K., Yu, K., Millican, K.,  
Sjoesund, L. L., Lee, L., Dixon, L., Reid, M., Mikula, M.,  
Wirth, M., Sharman, M., Chinaev, N., Thain, N., Bachem,  
O., Chang, O., Wahltinez, O., Bailey, P., Michel, P., Yotov,  
P., Chaabouni, R., Comanescu, R., Jana, R., Anil, R.,  
McIlroy, R., Liu, R., Mullins, R., Smith, S. L., Borgeaud,  
S., Girgin, S., Douglas, S., Pandya, S., Shakeri, S., De,  
S., Klimenko, T., Hennigan, T., Feinberg, V., Stokowiec,  
W., hui Chen, Y., Ahmed, Z., Gong, Z., Warkentin, T.,  
Peran, L., Giang, M., Farabet, C., Vinyals, O., Dean, J.,  
Kavukcuoglu, K., Hassabis, D., Ghahramani, Z., Eck, D.,  
Barral, J., Pereira, F., Collins, E., Joulin, A., Fiedel, N.,  
Senter, E., Andreev, A., and Kenealy, K. Gemma: Open  
models based on gemini research and technology, 2024.  
URL <https://arxiv.org/abs/2403.08295>.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi,  
A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P.,  
Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen,  
M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W.,  
Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn,  
A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez,  
V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S.,  
Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y.,  
Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Moly-  
bog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R.,  
Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subra-  
manian, R., Tan, X. E., Tang, B., Taylor, R., Williams,  
A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y.,  
Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Sto-  
jnic, R., Edunov, S., and Scialom, T. Llama 2: Open  
foundation and fine-tuned chat models, 2023. URL  
<https://arxiv.org/abs/2307.09288>.
- von Werra, L., Belkada, Y., Tunstall, L., Beeching, E.,  
Thrush, T., Lambert, N., Huang, S., Rasul, K., and  
Gallouédec, Q. Trl: Transformer reinforcement learn-  
ing. <https://github.com/huggingface/trl>,  
2020.

- 550 Wang, H., Xiong, W., Xie, T., Zhao, H., and Zhang,  
551 T. Interpretable preferences via multi-objective reward  
552 modeling and mixture-of-experts, 2024. URL <https://arxiv.org/abs/2406.12845>.  
553
- 554 Wang, Z., Bukharin, A., Delalleau, O., Egert, D., Shen,  
555 G., Zeng, J., Kuchaiev, O., and Dong, Y. Helpsteer2-  
556 preference: Complementing ratings with preferences,  
557 2025. URL [https://arxiv.org/abs/2410.](https://arxiv.org/abs/2410.01257)  
558 [01257](https://arxiv.org/abs/2410.01257).  
559
- 560 Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C.,  
561 Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M.,  
562 Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite,  
563 Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M.,  
564 Lhoest, Q., and Rush, A. M. Huggingface’s transformers:  
565 State-of-the-art natural language processing, 2020. URL  
566 <https://arxiv.org/abs/1910.03771>.  
567
- 568 Wu, Z., Hu, Y., Shi, W., Dziri, N., Suhr, A., Ammanabrolu,  
569 P., Smith, N. A., Ostendorf, M., and Hajishirzi, H. Fine-  
570 grained human feedback gives better rewards for language  
571 model training, 2023. URL [https://arxiv.org/](https://arxiv.org/abs/2306.01693)  
572 [abs/2306.01693](https://arxiv.org/abs/2306.01693).  
573
- 574 Yang, R., Ding, R., Lin, Y., Zhang, H., and Zhang, T. Regu-  
575 larizing hidden states enables learning generalizable re-  
576 ward model for llms. In *Advances in Neural Information*  
577 *Processing Systems*, 2024.
- 578 Zhang, L., Hosseini, A., Bansal, H., Kazemi, M., Kumar, A.,  
579 and Agarwal, R. Generative verifiers: Reward modeling  
580 as next-token prediction, 2025. URL [https://arxiv.](https://arxiv.org/abs/2408.15240)  
581 [org/abs/2408.15240](https://arxiv.org/abs/2408.15240).  
582
- 583 Zhong, J., Shen, W., Li, Y., Gao, S., Lu, H., Chen, Y., Zhang,  
584 Y., Zhou, W., Gu, J., and Zou, L. A comprehensive survey  
585 of reward models: Taxonomy, applications, challenges,  
586 and future, 2025. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2504.12328)  
587 [2504.12328](https://arxiv.org/abs/2504.12328).  
588
- 589 Zhu, J., Zheng, C., Lin, J., Du, K., Wen, Y., Yu, Y., Wang,  
590 J., and Zhang, W. Retrieval-augmented process reward  
591 model for generalizable mathematical reasoning, 2025.  
592 URL <https://arxiv.org/abs/2502.14361>.  
593
- 594 Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford,  
595 A., Amodei, D., Christiano, P., and Irving, G. Fine-tuning  
596 language models from human preferences, 2020. URL  
597 <https://arxiv.org/abs/1909.08593>.  
598  
599  
600  
601  
602  
603  
604

## 605 A. Limitations & Future Work

606 **Method.** While NORMBT yields consistent gains on fine-grained distinctions, a slight trade-off remains in larger-distance  
607 regions; more intricate scaling strategies may help mitigate this. Unified-Feedback and Skywork-80K are human-annotated  
608 or curated datasets, and settings with substantial noise or near-duplicate pairs may require additional handling.

610 **Evaluation.** Our comparison focuses on BT-based variants. Techniques involving regularization or auxiliary objectives may  
611 alter gradient dynamics beyond Eq. 7, and exploring how NORMBT interacts with such methods is an interesting direction.  
612 We use Best-of-N to mirror downstream RLHF usage; full policy-level evaluations (Schulman et al., 2017; Ahmadian et al.,  
613 2024) would offer a more comprehensive assessment.

614 **Analysis.** Our analysis centers on gradient magnitudes to quantify each pair’s contribution to model updates. Incorporating  
615 gradient-direction or correlation analysis would provide complementary insight into the underlying representation space.

## 618 B. Implementation Details

### 619 B.1. Baseline Details.

620 This section includes the formulation for each baseline in Sec 3, as well as their relevance to our method.

- 621 1. **BT Baseline.** The standard Bradley–Terry objective in Eq 1.
- 622 2. **BT + Margin.** (Touvron et al., 2023) For the Unified-Feedback dataset, which provides ratings for the responses, we  
623 extract a ground-truth margin  $m$  (i.e., difference in ratings) and incorporate it into the BT loss. In this formulation, the  
624 margin is introduced inside the log-sigmoid:

$$625 \mathcal{L}_{\text{Margin}}(\theta) = -\mathbb{E}_{(x, y_w, y_l, m) \sim D} [\log \sigma(r_w - r_l - m)]. \quad (14)$$

626 This variant shifts the effective decision boundary for each pair of chosen and rejected responses according to the  
627 ground-truth margin. This baseline tests the effect of injecting ground-truth-based error strength alone, without  
628 modifying BT’s representation coupling.

- 629 3. **BT + Margin (outside).** (Wang et al., 2025) In the second formulation, the margin is applied outside of log-sigmoid as  
630 a multiplicative weight on the per-sample loss:

$$631 \mathcal{L}_{\text{Margin(out)}}(\theta) = -\mathbb{E}_{(x, y_w, y_l, m) \sim D} [m \cdot \log \sigma(r_w - r_l)]. \quad (15)$$

632 This objective emphasizes the role of prediction error by upweighting pairs with larger ground-truth margins. It tests  
633 the effect of a reweighting scheme driven by ground-truth labels, rather than representation distance.

- 634 4. **BT + Label Smoothing.** (Liu et al., 2024a) A widely used regularization technique in classification models, where soft  
635 labels replace hard binary targets:

$$636 \mathcal{L}_{\text{LS}}(\theta) = -\mathbb{E}_{(x, y_w, y_l) \sim D} [(1 - \alpha) \log \sigma(r_w - r_l) + \alpha(r_l - r_w)]. \quad (16)$$

637 Differentiating w.r.t. parameters, its gradient is given by

$$638 \nabla_{\theta} \mathcal{L}_{\text{LS}} = (\sigma(d) - (1 - \alpha)) \cdot (\nabla_{\theta} r_w - \nabla_{\theta} r_l), \quad (17)$$

639 where  $\alpha = 0$  recovers the BT objective. Note that since  $\sigma(d) \in [0, 1]$ , its norm is

$$640 \begin{aligned} \|\nabla_{\theta} \mathcal{L}\| &= |\sigma(d) - (1 - \alpha)| \|\nabla_{\theta} r_w - \nabla_{\theta} r_l\| \\ &= \underbrace{(1 - \sigma(d))}_{\text{BT magnitude} - \alpha} \|\nabla_{\theta} r_w - \nabla_{\theta} r_l\|. \end{aligned} \quad (18)$$

641 From this perspective, label smoothing adjusts the strength of BT gradient from the prediction error term, by further  
642 reducing its magnitude by  $\alpha$ . This offers another relevant and generic method for our comparison.

643 In summary, the margin-based variants modify the prediction-error term by incorporating ground-truth reward differences,  
644 effectively adjusting update strength using external supervision. Label smoothing also acts on the prediction-error term, but  
645 does so by uniformly reducing its magnitude across all pairs. These approaches serve as natural baselines that manipulate  
646 error strength, while our method addresses a distinct factor by normalizing representation distance.

**B.2. Training Details.**

We implement all methods based on transformers (Wolf et al., 2020) and trl (von Werra et al., 2020), developing from the work of GRM (Yang et al., 2024). To use the Unified-Feedback dataset, we downsample the training data from the “all” set. All reward models are trained with the default reward head, as a linear layer with shape (hidden size, 1). For label smoothing,  $\alpha$  is set to 0.1 by default. For all BT baselines, we conduct an extensive grid search over learning rates and report the best-tuned results. For BT-based variants (margin, margin-out, label smoothing), we search over a range of learning rates centered around the best value found for BT and report the best configuration. All models are trained for 1 epoch with full parameter tuning. We truncate the inputs over 4096 tokens, and an effective batch size of 256 with gradient accumulation. Our experiments are conducted using NVIDIA RTX A6000 49G.

Table 3. Key hyperparameter details in reward model training.

Basic information	
Datasets	Unified-Feedback (80K), Skywork-Reward-Preference-80K-v0.2
Base models	gemma-2b-it, Llama-3.2-3b-Instruct
Quantization for training	bf16
Optimizer	sgd
Momentum	0.9
Batch size	256
Gradient accumulation	4
Learning Rate Scheduler	cosine
Warmup Ratio	0.03

**C. Experiment Error Analysis**

To assess the robustness of our main results reported in Table 1, we conduct an expanded evaluation on five random seeds for each method. Using a 40K subset of the Unified-Feedback dataset, we train both NORMBT and the BT baseline following the same training configuration described in Appendix B. We compare these 10 models to show the performance improvement of NORMBT over the BT baseline. The results reported in Table 4 and Figure 5 are consistent with the main experiments, with higher average scores across RewardBench and particularly strong gains in the Reasoning category.

Table 4. Error Analysis on Ten Models. Performance on RewardBench over models trained on five random seeds.

Reward Model	Chat	Chat Hard	Safety	Reasoning	Average
BT (Baseline)	94.69	37.48	75.11	65.34	68.15
NORMBT (Ours)	95.03	38.00	74.95	70.50	69.62
<b>Improvement</b>	<b>+0.34% (<math>\pm 0.19</math>)</b>	<b>+0.52% (<math>\pm 0.40</math>)</b>	<b>-0.16% (<math>\pm 0.49</math>)</b>	<b>+5.16% (<math>\pm 2.13</math>)</b>	<b>+1.46% (<math>\pm 0.61</math>)</b>

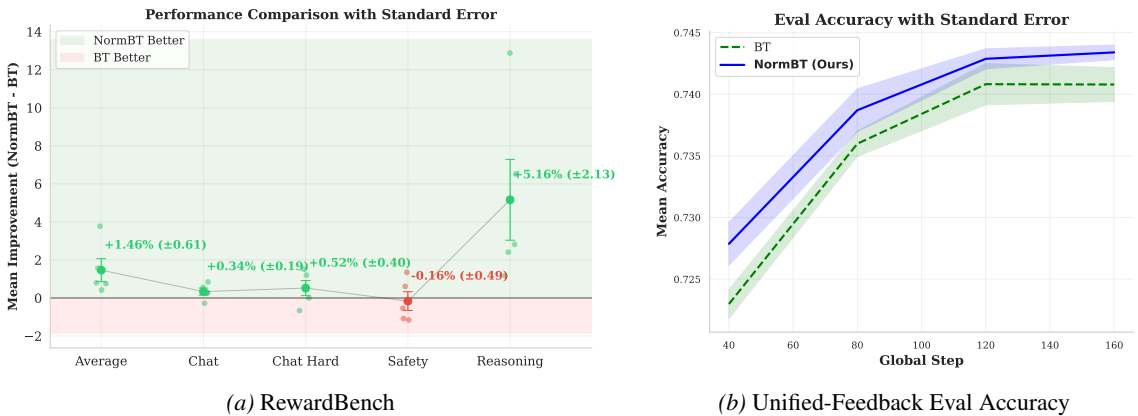


Figure 5. Error Analysis on Ten Models. Detailed performance of models trained on five random seeds with standard errors on (a) RewardBench categories, and (b) Unified-Feedback Evaluation.

## D. Additional Examples

**Small Representation Distance Pairs.** This section follows up on the performance comparison between BT Baseline and NORMBT in Figure 4. We provide example instances of “Small-Distance” response pairs from RewardBench where NORMBT predicts correctly while BT Baseline does not. This gives a clear illustration to pinpoint the preference pairs where NORMBT gives the largest performance gains.

### Pair 1 (Go)

**Prompt.** Write a Go function `func LargestPrimeFactor(n int) int` to solve the following problem: Return the largest prime factor of `n`. Assume `n > 1` and is not a prime.

```
>>> LargestPrimeFactor(13195)
29
>>> LargestPrimeFactor(2048)
2
```

```
isPrime := func(n int) bool {
    for i := 2; i < int(math.Pow(float64(n), 0.5)+1); i++ {
        if n%i == 0 {
            return false
        }
    }
    return true
}
```

```
largest := 1
for j := 2; j < n + 1; j++ {
    if n % j == 0 && isPrime(j) {
        if j > largest {
            largest = j
        }
    }
}
return largest
}
```

```
isPrime := func(n int) bool {
    for i := 2; i < int(math.Pow(float64(n), 0.5)+1); i++ {
        if n%i == 0 {
            return false
        }
    }
    return true
}
```

```
largest := 1
for j := 2; j < n + 1; j++ {
    if n % j == 0 && isPrime(n) {
        if j > largest {
            largest = j
        }
    }
}
return largest
}
```

## Pair 2 (JavaScript)

**Prompt.** Write a JavaScript function `const isBored = (S)` to solve the following problem. You'll be given a string of words, and your task is to count the number of boredoms. A boredom is a sentence that starts with the word "I". Sentences are delimited by '.', '?' or '!'.  


---

```

776 let t = 0
777
778 if (S[0] == 'I' && S[1] == ' ') { t = 1 }
779
780 for (let i = 0; i < S.length; i++) {
781     if (S[i] == '.' || S[i] == '!' || S[i] == '?') {
782         if (S[i + 1] == ' ' && S[i + 2] == 'I' && S[i + 3] == ' ') {
783             t++
784         }
785     }
786 }
787 return t
788 }

```

---

```

788 let t = 0
789
790 if (S[0] == ' ' && S[1] == 'I') { t = 1 }
791
792 for (let i = 0; i < S.length; i++) {
793     if (S[i] == '.' || S[i] == '!' || S[i] == '?') {
794         if (S[i + 1] == ' ' && S[i + 2] == 'I' && S[i + 3] == ' ') {
795             t++
796         }
797     }
798 }
799 return t
800 }

```

## Pair 3 (Go)

**Prompt.** Write a Go function `func NumericalLetterGrade(grades []float64) []string` to solve the following problem. It is the last week of the semester and the teacher has to give the grades to students. The teacher has been making her own algorithm for grading, but she has lost the code she used for grading. You are given a list of GPAs and must output a list of letter grades using the following table:

GPA	Letter grade
4.0	A+
...	
0.0	E

```

letter_grade := make([]string, 0, len(grades))
for _, gpa := range grades {
    switch {
    case gpa == 4.0:
        letter_grade = append(letter_grade, "A+")
    case gpa > 3.7:
        letter_grade = append(letter_grade, "A")
    case gpa > 3.3:
        letter_grade = append(letter_grade, "A-")
    case gpa > 3.0:
        letter_grade = append(letter_grade, "B+")
    case gpa > 2.7:
        letter_grade = append(letter_grade, "B")
    case gpa > 2.3:
        letter_grade = append(letter_grade, "B-")
    case gpa > 2.0:
        letter_grade = append(letter_grade, "C+")
    case gpa > 1.7:
        letter_grade = append(letter_grade, "C")
    case gpa > 1.3:
        letter_grade = append(letter_grade, "C-")
    case gpa > 1.0:
        letter_grade = append(letter_grade, "D+")
    case gpa > 0.7:
        letter_grade = append(letter_grade, "D")
    case gpa > 0.0:
        letter_grade = append(letter_grade, "D-")
    default:
        letter_grade = append(letter_grade, "E")
    }
}
return letter_grade
}

```

```

letter_grade := make([]string, 0, len(grades))
for _, gpa := range grades {
    switch {
    case gpa == 4.0:
        letter_grade = append(letter_grade, "A+")
    case gpa > 3.7:
        letter_grade = append(letter_grade, "A")
    case gpa > 3.3:
        letter_grade = append(letter_grade, "A-")
    case gpa > 3.0:
        letter_grade = append(letter_grade, "B+")
    case gpa > 2.7:
        letter_grade = append(letter_grade, "B")
    case gpa > 2.3:
        letter_grade = append(letter_grade, "B-")
    case gpa > 2.0:
        letter_grade = append(letter_grade, "C+")
    case gpa > 1.7:
        letter_grade = append(letter_grade, "C")
    case gpa > 1.3:
        letter_grade = append(letter_grade, "C-")
    case gpa > 1.0:
        letter_grade = append(letter_grade, "D+")
    case gpa > 0.7:
        letter_grade = append(letter_grade, "D")
    case gpa > 0.0:
        letter_grade = append(letter_grade, "D-")
    default:
        letter_grade = append(letter_grade, "E+")
    }
}
return letter_grade
}

```

## E. Intuition: Large- vs. Small-Distance Pairs

Small representation distance can arise from two reasons: (1) **uninformative or noisy pairs**, such as near-duplicate responses that convey little preference signal; or (2) **difficult but meaningful pairs**, where the backbone fails to separate two responses despite clear human preference signal (e.g., correct vs. subtly flawed code). As illustrated in Figure 2b, such case (2) pairs are common in practice, and Section D provides concrete examples.

This distinction matters because BT loss treats both cases identically: small-distance pairs always receive small gradients. While suppressing case (1) is desirable, suppressing case (2) is harmful, as these are precisely the pairs containing the fine-grained distinctions that the reward model should learn but inherently struggles to separate with the current  $\Phi$ .

In our datasets (UnifiedFeedback and Skywork-80K), most pairs are human-annotated or carefully curated, and following (Yang et al., 2024), pairs with identical ratings are filtered out. This substantially reduces the prevalence of case (1), making small-distance pairs far more likely to reflect case (2) of the difficult, fine-grained distinctions that RM should learn. Viewed through this lens, NORMBT effectively *upweights hard, meaningful pairs* whose gradients would otherwise be suppressed due to representation limitations in the current model.

## F. Full vs. Proxy Representation Distance

This section analyzes empirically the relationship between the full representation distance  $\|\nabla_{\theta}(r_w - r_l)\|$  and its proxy  $\|h_w - h_l\|$  as defined in Sec 2.1. Figure 6 plots the correlation between these two quantities. For gemma-2b-it, the correlation is strong with  $r = 0.928$ , suggesting that the embedding difference is a highly reliable stand-in for the full representation distance. For Llama-3.2-3b-Instruct, the correlation is more modest at the beginning of reward model training ( $r = 0.682$ ; shown in Figure 6b) but increases to  $r = 0.932$  after some initial training. This may indicate that the proxy becomes more accurate once the pretrained model begins adapting to the reward modeling objective.

These observations support the theoretical analysis in Eq. 7, which establishes the connection between  $\|\nabla_{\theta}(r_w - r_l)\|$  and  $\|h_w - h_l\|$ . This suggests the strategy to utilize  $\|h_w - h_l\|$  as a proxy is a computationally efficient and empirically validated substitute for the full distance, avoiding costly backward passes.




Figure 6. Correlation of Full vs. Proxy Representation Distance. Both base models show a strong correlation ( $r > 0.9$ ) between the two quantities.

## G. Representation Scale & EMA

One challenge in NORMBT normalization arises from the non-stationary behavior of representation distances  $\|h_w - h_l\|$  during training of the reward model. Since the LLM backbone is jointly optimized with the reward head, the learned representation space could evolve considerably.

Figure 7 illustrates this observation: the mean representation distance shifts significantly over the course of optimization, while the variance remains consistently large. And the same behavior is observed across both Unified-Feedback and Skywork-Reward-Preference-80K-v0.2 datasets, as well as for all BT baselines. As a result, although the large variance highlights the potential gradient imbalance, any normalization scheme that relies directly on raw representation magnitudes inherits the instability created by scale shifts.

To mitigate this issue, we incorporate an Exponential Moving Average (EMA) as stated in Eq 12. By tracking a running mean of representation distance, EMA provides a stable and adaptive reference point that smooths out fluctuations in embedding scale while remaining responsive to long-term trends. By normalizing each pair relative to this moving reference, rather than to its raw magnitude, we ensure that weighting reflects relative distinctions among pairs. This yields consistent gradient scaling across training. The benefit of EMA is also reflected in the performance gap observed in ablation studies shown in Table 2.

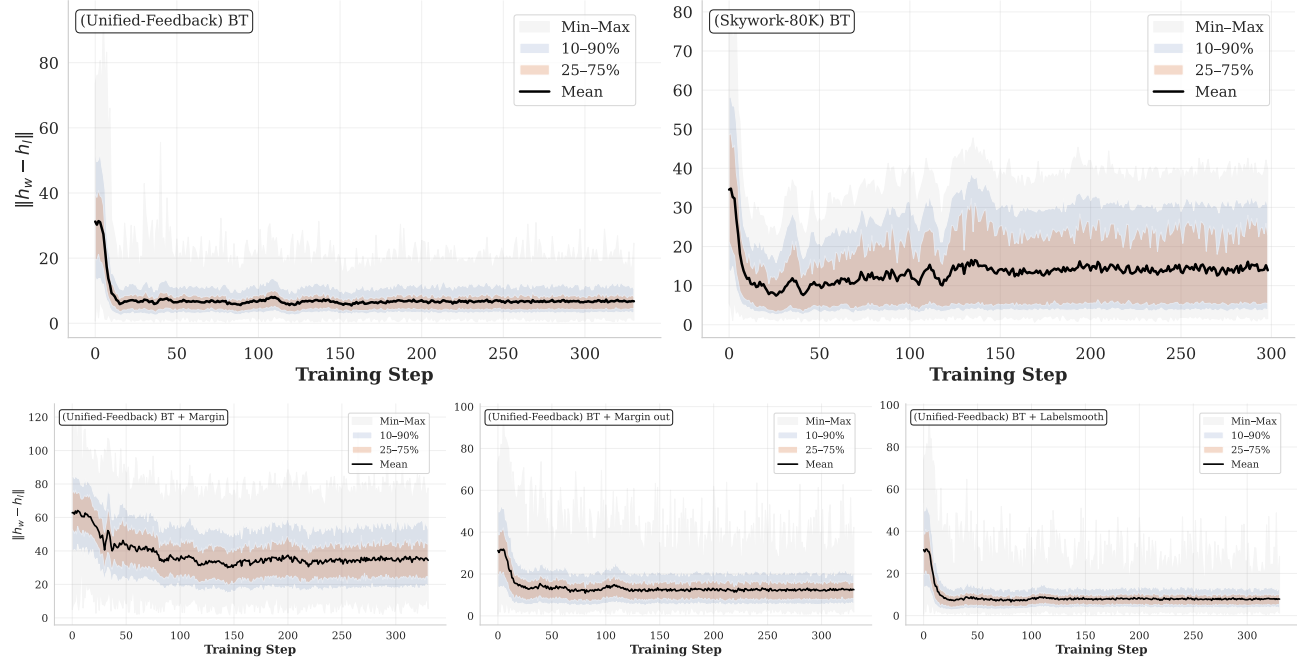


Figure 7. **Dynamics of Representation Distance** during RM Training for standard BT on Unified-Feedback and Skywork-80K (top), and all BT variants on Unified-Feedback (bottom). For each method, we report the mean, 25–75 percentile, 10–90 percentile, and full min–max range. Across all settings, the mean distance shifts substantially, while the variance remains high.