ROBUST PREFERENCE ALIGNMENT VIA DIRECTIONAL NEIGHBORHOOD CONSENSUS

Anonymous authors

000

001

002003004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027

028

029

031

032 033 034

035

037

040

041

042

043

044

046 047

048

050

051 052 Paper under double-blind review

ABSTRACT

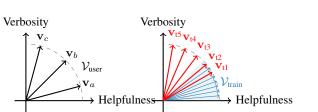
Aligning large language models with human preferences is critical for creating reliable and controllable AI systems. A human preference can be visualized as a high-dimentional vector where different directions represent trade-offs between desired attributes (e.g., helpfulness vs. verbosity). Yet, because the training data often reflects dominant, average preferences, LLMs tend to perform well on common requests but falls short in specific, individual needs. This mismatch creates a preference coverage gap. Existing methods often address this through costly retraining, which may not be generalized to the full spectrum of diverse preferences. This brittleness means that when a user's request reflects a nuanced preference deviating from the training data's central tendency, model performance can degrade unpredictably. To address this challenge, we introduce Robust Preference Selection (RPS), a post-hoc, training-free method by leveraging directional neighborhood consensus. Instead of forcing a model to generate a response from a single, highly specific preference, RPS samples multiple responses from a local neighborhood of related preferences to create a superior candidate pool. It then selects the response that best aligns with the user's original intent. We provide a theoretical framework showing our neighborhood generation strategy is provably superior to a strong baseline that also samples multiple candidates. Comprehensive experiments across three distinct alignment paradigms (DPA, DPO, and SFT) demonstrate that RPS consistently improves robustness against this baseline, achieving win rates of up to 69% on challenging preferences from under-represented regions of the space without any model retraining. Our work presents a practical, theoretically-grounded solution for enhancing the reliability of preference-aligned models¹.

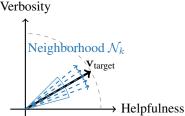
1 Introduction

Aligning large language models (LLMs) with human preferences is crucial for creating reliable and controllable AI systems (Ouyang et al., 2022; Christiano et al., 2017; Ziegler et al., 2020). User preferences can be modeled in a multi-dimensional space where different directions represent trade-offs between desired attributes, such as helpfulness versus verbosity (Wang et al., 2024a; Dong et al., 2023). As illustrated in Figure 1, this creates a foundational challenge: the **preference coverage gap**. While the space of potential user preferences is vast and diverse, as depicted in Figure 1(a), the alignment process often optimizes for a dominant, average preference, meaning the training data is concentrated in a narrow region (Figure 1(b)). This focus on average preferences makes models brittle; when faced with a user preference that reflects more individual needs and deviates from this central tendency—a common out-of-distribution (OOD) challenge—their performance can degrade unpredictably, undermining user trust (Hendrycks et al., 2020).

To address this coverage gap, many existing solutions focus on training-time interventions. These include methods like data augmentation or the adoption of principles from Distributionally Robust Optimization (DRO) (Duchi et al., 2016; Ben-Tal et al., 2013; Duchi & Namkoong, 2018) to create models that are resilient to shifts in preference distributions (Xu et al., 2025). While effective, such approaches often require costly retraining cycles and may still fail to generalize to the full spectrum

¹Our code and data are available at https://anonymous.4open.science/r/Robust_Preference_Selection-C65C





(a) The complete user pref- (b) The sparse training erence space.

056

057

060

061 062

063

064

065

066

067

068

069

071

073

074

075

076

077

078

079

081

083

084

085

087

090

092

093

095

096

098 099

100 101

102 103

104

105

106

107

preference set.

Figure 1: Illustration of the preference coverage gap. While user preferences (a) can span the entire space, the model's training data (b) is often concentrated on dominant, average preferences, leaving individual needs in a sparse subset. This creates a gap where a user's target preference may lie.

Figure 2: Conceptual visualization of RPS. Instead of relying on a single, potentially outof-distribution target preference $\mathbf{v}_{\text{target}}$ (solid black arrow), RPS samples k directions from its local neighborhood (dashed blue arrows). By generating responses from this diverse set, RPS can identify a response that better aligns with the user's true intent.

of diverse, individual preferences. This motivates an alternative question: can we enhance robustness at **inference time**, without any modification to the underlying model?

This paper argues that forcing a model to generate a response from a single, highly specific and less common preference direction is inherently fragile. We propose a paradigm shift from direct generation to one based on **directional neighborhood consensus**. As visualized in Figure 2, instead of attempting to extrapolate to a specific, under-represented preference point, it is more robust to explore the local neighborhood, generate responses from these more dominant, better-understood directions, and then select the one that best satisfies the original preference.

To realize this, we introduce **Robust Preference Selection** (**RPS**), a post-hoc adjustment method that enhances preference alignment at inference time without any retraining. RPS first samples a set of candidate preference vectors from the neighborhood of the user's target preference. It then generates a response for each of these nearby vectors and, finally, uses the target preference itself as a criterion to select the optimal response from this diverse set. This approach effectively leverages the model's existing capabilities in well-trained regions of the preference space to satisfy requests in undertrained

Our contributions are threefold:

- We formally define the *preference coverage gap* as a critical out-of-distribution (OOD) challenge that undermines the reliability of aligned LLMs. To address this, we introduce RPS, a novel, training-free method that enhances robustness through post-hoc adjustment without requiring any model modification.
- We propose RPS, a method grounded in neighborhood consensus, and provide a theoretical framework proving that its neighborhood generation strategy is superior to a strong multicandidate baseline.
- We conduct extensive experiments across three distinct alignment paradigms (DPA, DPO, and SFT) and three datasets (UltraFeedback, HelpSteer, and HelpSteer2). Our results show that RPS consistently improves robustness, achieving win rates of up to 69% on challenging OOD preferences and demonstrating its broad applicability.

RELATED WORK

PREFERENCE ALIGNMENT IN LARGE LANGUAGE MODELS

Aligning LLM behavior with human preferences has become a central research area. Reinforcement Learning from Human Feedback (RLHF) is a pioneering pipeline that fine-tunes models with human preference rankings, as demonstrated by (Ouyang et al., 2022). However, RLHF compresses diverse user preferences into a single scalar reward and requires complex reward modeling plus reinforcement learning. To simplify this process, Direct Preference Optimization (DPO) was introduced (Rafailov et al., 2023), which recasts preference optimization as supervised classification, eliminating the

need for explicit reward models. Subsequent generalizations explore divergence families and latent user heterogeneity (Wang et al., 2023; Chidambaram et al., 2024), while others have proposed new theoretical paradigms for understanding preference learning (Azar et al., 2023). Moving beyond scalar objectives, Directional Preference Alignment (DPA) enables users to specify trade-offs in a multi-axis reward space (Wang et al., 2024a). Similarly, SteerLM conditions supervised fine-tuning on attribute labels, exposing controllable style dimensions such as helpfulness or humor (Dong et al., 2023). These methods are part of a broader research effort in controllable text generation, which aims to provide users with fine-grained control over model outputs (Liang et al., 2024). Our work differs from these training-time approaches: rather than modifying the model weights, we focus on inference-time robustness to preference shifts through directional neighborhood consensus.

2.2 Enhancing Robustness in Language Models

While the alignment methods described above are powerful, a key challenge remains: models often remain brittle under out-of-distribution (OOD) preferences. Recent work has formalized preference distribution shifts and proposed distributionally robust objectives such as (Xu et al., 2025), which strengthen resilience during training. Beyond alignment, the broader NLP community has highlighted the challenges of OOD generalization, with benchmarks such as (Yang et al., 2022; 2023). At inference time, an alternative approach is to use ensemble-like methods, a principle with deep roots in machine learning (Dietterich, 2000). For instance, (Wang et al., 2022) shows that sampling diverse reasoning paths and aggregating their consensus yields more reliable results. The principle of post-hoc adjustment for robustness is also explored in other domains, such as classification, where scaling model outputs can mitigate the effects of distributional shifts (Wei et al., 2023).

Extending this idea, recent inference-time alignment frameworks share our post-hoc perspective but differ in mechanism. Many rely on direct intervention in the generation process through token-level guidance or activation steering (Li et al., 2025; Shahriar et al., 2024), or require auxiliary models for decoding-time guidance (Chehade et al., 2025; Chandra et al., 2025). In contrast, our RPS approach operates purely in the preference space. By leveraging neighborhood consensus to select an optimal response, it avoids direct manipulation of the model's internal states, offering a simpler and more black-box solution that requires no external guidance models.

3 PROBLEM SETUP AND THEORETICAL FRAMEWORK

We build upon the problem formulation of Directional Preference Alignment (DPA) (Wang et al., 2024a). In this section, we formalize the preference alignment challenge by first defining the preference space and characterizing the coverage gap that causes model brittleness. We then establish the theoretical foundations for our proposed solution, Robust Preference Selection (RPS).

3.1 Preference Space and Reward Model

We model user preferences in a two-dimensional space for clarity of illustration, as depicted in Figure 1(a), spanned by two key axes: **Helpfulness** and **Verbosity** (Wang et al., 2024a; Dong et al., 2023). Our theoretical framework, however, generalizes directly to higher-dimensional preference spaces. To quantify these attributes, we formalize the notion of a reward vector.

Definition (Reward Vector): A reward model maps a prompt-response pair (x, y) to a reward vector $\mathbf{r}(x, y) = (r_h(x, y), r_v(x, y)) \in \mathbb{R}^2$. The components $r_h(x, y)$ and $r_v(x, y)$ are scalar scores representing the helpfulness and verbosity of the response, respectively (Wang et al., 2024a).

For all experiments, we use the publicly available RewardModel-Mistral-7B-for-DPA-v1²; further details on the scoring procedure are provided in Appendix A.4.

A user's preference is represented as a normalized direction vector $\mathbf{v} = (v_h, v_v) \in \mathbb{S}^1$ on the unit circle, where v_h and v_v specify the desired weights for helpfulness and verbosity. This can be parameterized by an angle θ , such that $\mathbf{v} = (\cos \theta, \sin \theta)$. The goal of a preference-aligned model is to generate a response y that maximizes the projected reward: $\mathbf{v}^T \mathbf{r}(x, y)$. Our framework assumes that this reward model $\mathbf{r}(x, y)$ is well-calibrated and provides meaningful scores across the entire preference space, including for out-of-distribution directions.

²https://huggingface.co/RLHFlow/RewardModel-Mistral-7B-for-DPA-v1

3.2 THE PREFERENCE COVERAGE PROBLEM

The central challenge in preference alignment, as illustrated in Figure 1, is the discrepancy between the vast space of user preferences and the limited coverage of the training data. We formalize this problem as follows:

Definition 1 (User Preference Space): Let \mathcal{V}_{user} denote the complete set of all possible normalized preference vectors $\mathbf{v} \in \mathbb{S}^1$. This represents the entire spectrum of potential user preferences, as depicted in Figure 1(a).

Definition 2 (Training Preference Set): Let $\mathcal{V}_{train} \subset \mathcal{V}_{user}$ be the subset of preference directions used during training, visualized as the concentrated region in Figure 1(b). This set is often sampled from a constrained range.

Definition 3 (Preference Coverage Gap): The coverage gap, illustrated by the difference between the full space in Figure 1(a) and the training data in Figure 1(b), consists of all preference vectors that are not within an ϵ -neighborhood of any training vector: Gap = $\mathcal{V}_{\text{user}} \setminus \mathcal{N}_{\epsilon}(\mathcal{V}_{\text{train}})$.

When a target preference $\mathbf{v}_{\text{target}}$ lies in this gap—as illustrated with the out-of-distribution vector in Figure 1(b)—the model's performance is unreliable. Our goal is to develop a method that can robustly generate a high-quality response y^* that maximizes user satisfaction, even for out-of-distribution preferences:

$$y^* = \arg\max_{y} \mathbf{v}_{\text{target}}^T \mathbf{r}(x, y), \tag{1}$$

where $\mathbf{r}(x,y) = (r_h(x,y), r_v(x,y))$ represents the helpfulness and verbosity scores of response y to prompt x. This challenge of performing well on a target preference $\mathbf{v}_{\text{target}}$ that lies in the gap can be framed through the lens of Distributionally Robust Optimization (DRO) (Duchi et al., 2016; Ben-Tal et al., 2013; Duchi & Namkoong, 2018). In the DRO paradigm, the objective is to find a policy that is robust not just to the empirical training distribution (represented by $\mathcal{V}_{\text{train}}$) but to a family of plausible test distributions. Our inference-time approach complements training-time DRO solutions by addressing this distributional shift post-hoc, at the point of generation.

3.3 Neighborhood Consensus Theory

Instead of merely justifying the final selection step, our theoretical framework aims to explain why the entire neighborhood generation strategy is superior to a strong baseline that repeatedly samples from the target direction. The core intuition is that for an out-of-distribution (OOD) preference \mathbf{v}_{target} , the model's performance is degraded. By sampling from a nearby neighborhood of more in-distribution preferences, we can generate a candidate pool of higher average quality. The following assumption formalizes this intuition.

Assumption 1 (OOD Performance Degradation): Let $\mathbf{v}_{\text{target}}$ be an OOD preference vector. Let $\mathcal{D}_{\text{train}}$ be the distribution of preferences in the training set $\mathcal{V}_{\text{train}}$. For a nearby preference vector $\mathbf{v}_i \in \mathcal{N}_k(\mathbf{v}_{\text{target}})$ that is closer to the mean of $\mathcal{D}_{\text{train}}$, the expected score of a response $y_i \sim \pi_{\theta}(\cdot|x,\mathbf{v}_i)$ is higher than that of a response $y_{\text{target}} \sim \pi_{\theta}(\cdot|x,\mathbf{v}_{\text{target}})$, when both are evaluated against their respective generating preferences: $\mathbb{E}[\mathbf{v}_i^T\mathbf{r}(x,y_i)] > \mathbb{E}[\mathbf{v}_{\text{target}}^T\mathbf{r}(x,y_{\text{target}})]$.

Furthermore, we assume local consistency, meaning the evaluation of y_i under $\mathbf{v}_{\text{target}}$ is a good proxy for its quality, i.e., $\mathbf{v}_{\text{target}}^T\mathbf{r}(x,y_i) \approx \mathbf{v}_i^T\mathbf{r}(x,y_i)$. This implies that the candidate pool from the neighborhood is stronger: $\mathbb{E}[\mathbf{v}_{\text{target}}^T\mathbf{r}(x,y_i)] > \mathbb{E}[\mathbf{v}_{\text{target}}^T\mathbf{r}(x,y_{\text{target}})]$. To formalize this advantage, we compare RPS against a strong **baseline** strategy: generating k independent responses by repeatedly sampling from the single target direction $\mathbf{v}_{\text{target}}$, and then selecting the best one according to the target preference. The following theorem proves the superiority of the RPS approach.

Theorem 1 (Superiority of Neighborhood Generation). Let $S_{RPS} = \{s_1, \ldots, s_k\}$ be the set of scores from k responses generated from the neighborhood \mathcal{N}_k , where $s_i = \mathbf{v}_{target}^T \mathbf{r}(x, y_i)$. Let $S_{Baseline} = \{s_1', \ldots, s_k'\}$ be the set of scores from k responses generated from the baseline strategy (i.e., directly from \mathbf{v}_{target}). Under Assumption 1, the expected score of the best response selected by RPS is strictly greater than that of the best response selected by the baseline:

$$\mathbb{E}[\max(S_{RPS})] > \mathbb{E}[\max(S_{Baseline})]. \tag{2}$$

This performance gap, illustrated in Figure 3, represents the robustness gain of the RPS method.

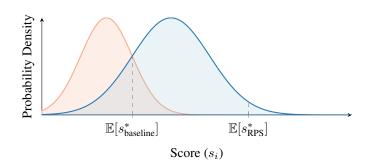


Figure 3: Conceptual illustration of Theorem 1. The distributions represent the scores of candidate responses for the baseline (orange) and RPS (blue). Under Assumption 1, the RPS candidate pool is drawn from a higher-quality distribution. Consequently, the expected score of the best RPS response, $\mathbb{E}[s_{\text{RPS}}^*]$, is strictly greater than that of the best baseline response, $\mathbb{E}[s_{\text{baseline}}^*]$. This difference is the robustness gain.

Proof. Let s_i be the random variable for the score of a response from a neighborhood direction $\mathbf{v}_i \in \mathcal{N}_k$, and let s'_{baseline} be the random variable for a score from the baseline direction $\mathbf{v}_{\text{target}}$. The scores in $S_{\text{RPS}} = \{s_1, \ldots, s_k\}$ are independent but not necessarily identically distributed, with cumulative distribution functions (CDFs) $F_1(x), \ldots, F_k(x)$. The scores in $S_{\text{Baseline}} = \{s'_1, \ldots, s'_k\}$ are independent and identically distributed (i.i.d.) draws from the baseline distribution, with CDF $F_{\text{baseline}}(x)$.

Under Assumption 1, each candidate from the neighborhood is drawn from a better distribution than a candidate from the baseline. This implies that each s_i first-order stochastically dominates s'_{baseline} . Formally, for each $i \in \{1, \ldots, k\}$, we have $F_i(x) \leq F_{\text{baseline}}(x)$ for all x, with strict inequality over some interval.

The CDF of the maximum score from RPS is $F_{\max}^{\text{RPS}}(x) = P(\max(S_{\text{RPS}}) \leq x) = \prod_{i=1}^k F_i(x)$, due to independence. The CDF of the maximum score from the baseline is $F_{\max}^{\text{Baseline}}(x) = P(\max(S_{\text{Baseline}}) \leq x) = (F_{\text{baseline}}(x))^k$.

Since $F_i(x) \leq F_{\text{baseline}}(x)$ for all i, it follows that $\prod_{i=1}^k F_i(x) \leq (F_{\text{baseline}}(x))^k$. Thus, $F_{\max}^{\text{RPS}}(x) \leq F_{\max}^{\text{Baseline}}(x)$ for all x. This shows that the maximum score from RPS also first-order stochastically dominates the maximum score from the baseline.

The expected value of a random variable can be expressed using its CDF. Assuming scores are non-negative (or shifted to be), $\mathbb{E}[X] = \int_0^\infty (1 - F(x)) dx$. Given the stochastic dominance:

$$\mathbb{E}[\max(S_{\mathsf{RPS}})] = \int_0^\infty (1 - F_{\max}^{\mathsf{RPS}}(x)) dx \ge \int_0^\infty (1 - F_{\max}^{\mathsf{Baseline}}(x)) dx = \mathbb{E}[\max(S_{\mathsf{Baseline}})]. \tag{3}$$

The inequality is strict because $F_i(x) < F_{\text{baseline}}(x)$ over some interval for at least one i, which ensures that $F_{\text{max}}^{\text{RPS}}(x) < F_{\text{max}}^{\text{Baseline}}(x)$ over that same interval. This rigorously confirms that leveraging the neighborhood produces a superior set of candidates, leading to a better final selection.

Corollary 1: The robustness gain increases with neighborhood size k and the quality gap between the neighborhood and target-direction candidate pools. This follows because the expected value of the maximum of k samples is non-decreasing in k, and this effect is more pronounced for the stochastically dominant RPS distribution. Similarly, a larger quality gap—meaning greater stochastic dominance of the neighborhood distributions over the baseline—naturally widens the separation in the expected maximums. A formal proof is provided in Appendix A.7.

3.4 ROBUST PREFERENCE SELECTION ALGORITHM

Building on the theoretical foundation of neighborhood consensus, we now formalize our approach. The Robust Preference Selection (RPS) algorithm, detailed in **Algorithm 1**, translates our theory into a practical, three-phase procedure designed to navigate the preference coverage gap.

The first phase, **Neighborhood Construction**, addresses the core challenge of out-of-distribution (OOD) preferences. Instead of directly using a potentially brittle target vector **v**_{target}, RPS identifies a

set of k nearby, more reliable preference directions. These candidate directions are sampled within a predefined angular threshold θ_{\max} , forming a local neighborhood \mathcal{N}_k . This step is critical as it shifts the generation process from a region of high uncertainty to one where the model's performance is more robust and predictable.

In the **Multi-Directional Generation** phase, the language model π_{θ} generates a separate response y_i for each of the k preference vectors in the neighborhood. This process creates a diverse portfolio of candidate responses. Each response reflects a slightly different trade-off between attributes (e.g., helpfulness and verbosity), leveraging the model's well-trained capabilities within this local region of the preference space. The result is a set of high-quality outputs, each optimized for a direction where the model is confident.

Finally, the **Consensus Selection** phase determines the optimal response. Crucially, all k candidates are evaluated against the user's original target preference, $\mathbf{v}_{\text{target}}$. The response y_i that maximizes the projected reward score $s_i = \mathbf{v}_{\text{target}}^T\mathbf{r}(x,y_i)$ is selected as the final output y^* . The superiority of this entire procedure is justified by our **Theorem 1**, which proves that the strategy of generating candidates from a superior neighborhood pool and then selecting the maximum is guaranteed to yield a response with a higher expected quality than the strong baseline. By combining neighborhood-based generation with target-based selection, RPS robustly satisfies user intent even for OOD preferences. The following section will empirically validate the effectiveness of this approach across various models and datasets.

Algorithm 1 Robust Preference Selection (RPS)

Require: Prompt x, target preference \mathbf{v}_{target} , neighborhood size k, angle threshold θ_{max} **Ensure:** Optimal response y^*

1: Phase 1: Neighborhood Construction

2: Generate candidate directions within θ_{max} of \mathbf{v}_{target}

3: Compute angular distances: $d_i = \arccos(\mathbf{v}_i \cdot \mathbf{v}_{target})$

4: Select k closest directions: $\mathcal{N}_k = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$

5: Phase 2: Multi-Directional Generation

6: **for** i = 1 to k **do**

7: Generate response: $y_i \sim \pi_{\theta}(\cdot|x, \mathbf{v}_i)$

8: end for

9: Phase 3: Consensus Selection

10: **for** i = 1 to k **do**

11: Compute score: $s_i = \mathbf{v}_{target}^T \mathbf{r}(x, y_i)$

12: end for

13: **return** $y^* = \arg \max_i s_i$

4 EXPERIMENTS

To validate our theoretical framework, we designed a comprehensive experimental methodology to assess the effectiveness of Robust Preference Selection (RPS) as a post-hoc method. We evaluated RPS against a strong baseline across three distinct model training paradigms—DPA, DPO, and SFT—to demonstrate its general applicability. Our experiments test the core hypothesis that neighborhood consensus provides robustness for out-of-distribution preference directions.

4.1 EXPERIMENTAL SETUP

4.1.1 Models and Datasets

To ensure a robust evaluation, we used a 3×3 experimental matrix, crossing three models with three standard preference-learning datasets. The models (Table 1) represent diverse training paradigms: Directional Preference Alignment (DPA), using DPA-v1-Mistral-7B³ (Wang et al., 2024a); Direct Preference Optimization (DPO), using Zephyr-7B-Beta⁴ (Tunstall et al., 2023); and

https://huggingface.co/RLHFlow/DPA-v1-Mistral-7B

⁴https://huggingface.co/HuggingFaceH4/zephyr-7b-beta

Table 1: Models evaluated.

Table 2: Evaluation datasets.

Model	Training Paradigm
DPA-v1-Mistral-7B	DPA
Zephyr-7B-Beta	DPO
Mistral-7B-Instruct-v0.2	SFT

UltraFeedback test_prefs 2,000 HelpSteer validation 503 HelpSteer2 validation 518	Dataset	Split	Size (used)
	HelpSteer	validation	503

Table 3: Overall RPS win rates by model and dataset. Values show mean \pm std across all preference directions.

	RPS vs. Baseline Average Win Rate (%)						
Model	UltraFeedback	HelpSteer	HelpSteer2				
DPA	$58.7 \pm 6.1\%$	$58.8 \pm 4.8\%$	$59.7 \pm 7.8\%$				
DPO SFT	$52.1 \pm 1.1\% 52.0 \pm 0.7\%$	$52.4 \pm 1.5\%$ $56.0 \pm 2.8\%$	$53.4 \pm 1.0\%$ $65.4 \pm 11.9\%$				

standard Supervised Fine-Tuning (SFT), using Mistral-7B-Instruct-v0.2⁵ (Jiang et al., 2023). The datasets (Table 2) provide varied domains for testing preference alignment: we use the 2,000-sample test_prefs split from UltraFeedback⁶ (Cui et al., 2023), the 503-sample deduplicated validation set from HelpSteer⁷ (Dong et al., 2023), and the 518-sample deduplicated validation set from its successor, HelpSteer2⁸ (Wang et al., 2024b).

4.1.2 EVALUATION PROTOCOL

For each model-dataset pair, we compare two inference-time strategies under a fixed computational budget: 1) **Single-Direction Baseline:** To ensure a fair comparison, we generate k=5 response candidates using only the target direction $\mathbf{v}_{\text{target}}$ (Wang et al., 2022). The best response is then selected by scoring each candidate with the target preference, i.e., maximizing $\mathbf{v}_{\text{target}}^T\mathbf{r}(x,y)$. 2) **RPS:** We first sample k=5 preference directions from a local neighborhood around $\mathbf{v}_{\text{target}}$, constrained by an angular threshold of $\theta_{\text{max}}=30^\circ$. The choice of these hyperparameters balances key trade-offs. A neighborhood size of k=5 was chosen to maintain strict compute parity with the baseline, while representing a common choice for balancing response diversity and inference cost. The angle $\theta_{\text{max}}=30^\circ$ was determined through preliminary pilots to be a sweet spot: smaller angles provided insufficient diversity over the baseline, while larger angles risked sampling preferences too semantically distant from the target, violating our local consistency assumption. We generate one response for each of the k directions. The final response is selected by scoring all k candidates against the original target preference $\mathbf{v}_{\text{target}}$.

This setup ensures that both methods generate and score the same number of candidate responses, maintaining strict compute parity, with the neighborhood sampling step introducing negligible overhead. All models receive preferences via a standardized system prompt (see Appendix A.2). We evaluate on eight challenging preference directions from 10° to 45° (see Appendix A.5) to test robustness on preferences progressively further from the training distribution. Response pairs are evaluated by a preference-aligned judge in a randomized A/B test, and our primary metric is the RPS win rate. We utilize GPT-40-mini as our preference-aligned judge, a practice increasingly adopted for its strong correlation with human judgments in preference evaluation tasks (Zheng et al., 2023; Gu et al., 2024).

5 RESULTS

Our experiments confirm that Robust Preference Selection (RPS) consistently improves alignment robustness, particularly for out-of-distribution (OOD) preferences. We present three key findings: (1)

⁵https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2

 $^{^6}$ https://huggingface.co/datasets/HuggingFaceH4/ultrafeedback_binarized

⁷https://huggingface.co/datasets/nvidia/HelpSteer

⁸https://huggingface.co/datasets/nvidia/HelpSteer2

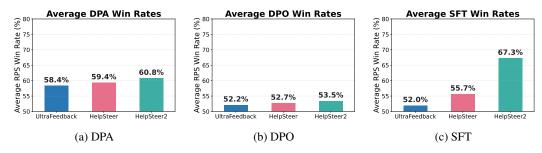


Figure 4: Overall RPS win rates by model (DPA, DPO, SFT) and dataset. Bars show mean win rates across all tested preference directions.



Figure 5: Directional robustness. RPS win rate vs. preference angle for DPA (left), DPO (middle), and SFT (right) models. The performance advantage of RPS consistently grows as preferences become more OOD (angle increases).

RPS outperforms a strong baseline across all models and datasets; (2) its advantage grows significantly as target preferences deviate from the training distribution; and (3) the magnitude of improvement depends on the model's initial alignment method, with SFT models benefiting most.

5.1 RPS Consistently Outperforms the Baseline and Excels on OOD Preferences

Across all nine model-dataset pairings, RPS achieves a decisive win rate greater than 50% against the single-direction baseline, as detailed in Table 3. The average improvements over a 50% baseline, visualized in Figure 4, are consistent, ranging from a modest +2.0% for SFT on UltraFeedback (a 52.0% win rate) to a significant +17.3% for SFT on HelpSteer2 (a 67.3% win rate). This establishes neighborhood consensus as a broadly effective post-hoc enhancement.

More importantly, the performance advantage of RPS amplifies on OOD preferences, a finding that provides strong empirical validation for our **Assumption 1 (OOD Performance Degradation)**. This trend is most pronounced for the DPA model, as shown in Figure 5. The win rate on UltraFeedback, for example, climbs from 53.4% at 20° to a dominant 69.1% at 45°. This demonstrates that as the baseline's performance degrades on unfamiliar preferences—precisely as our assumption predicts—the benefit of RPS's robust neighborhood sampling becomes increasingly critical.

In contrast, the DPO and SFT models show a more modest and less angle-dependent trend (Figure 5). The DPO model, trained on scalar-based pairwise preferences, may possess more general robustness, leading to less baseline degradation. Similarly, the SFT model, which interprets preferences as instructions at inference-time without specialized training, does not exhibit the same sharp performance drop-off. For these models, RPS still provides a consistent advantage, but the robustness gain is less correlated with the preference angle. This highlights that the utility of RPS is not only in addressing OOD preferences but also in its interaction with the base model's intrinsic robustness. Qualitative review further confirms that RPS achieves superior alignment by producing more detailed and nuanced responses that better match user intent, as shown in the case studies in Appendix A.6.

Table 4: Detailed RPS win rates by dataset, model, and preference direction. This table provides the full data for Figures 5 and 6.

		RPS vs. Baseline Win Rate (%)							
	Ult	raFeedb	ack	HelpSteer			HelpSteer2		
Direction	DPA	DPO	SFT	DPA	DPO	SFT	DPA	DPO	SFT
v1 (10°)	55.1	51.5	51.8	56.1	51.7	54.3	54.9	53.0	52.1
v2 (15°)	56.2	52.0	52.1	57.3	52.1	55.0	56.2	53.3	55.3
v3 (20°)	53.4	52.3	51.9	58.0	52.6	55.8	57.8	53.6	58.9
v4 (25°)	58.1	52.8	52.3	59.1	53.0	56.5	59.5	53.8	62.1
v5 (30°)	59.3	52.5	52.0	60.2	53.5	57.1	61.3	54.0	66.7
v6 (35°)	61.2	52.1	51.7	61.5	53.9	58.3	63.0	54.1	71.3
v7 (40°)	64.9	51.9	52.1	62.8	54.2	59.0	65.1	54.2	83.2
v8 (45°)	69.1	51.7	52.4	64.3	54.5	59.8	68.8	54.5	94.3

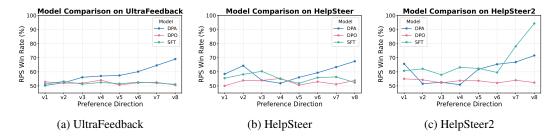


Figure 6: Dataset-wise performance. RPS win rate vs. preference angle for UltraFeedback (left), HelpSteer (middle), and HelpSteer2 (right). SFT models show particularly strong gains on HelpSteer datasets.

5.2 Analysis Across Alignment Paradigms and Datasets

Further analysis, with detailed data in Table 4 and visualized in Figure 6, reveals that the effectiveness of RPS is modulated by the base model's training paradigm. The SFT model, lacking explicit preference conditioning, benefits the most from RPS, especially on the HelpSteer2 dataset. This suggests RPS acts as an effective inference-time guidance mechanism for models not explicitly trained to follow nuanced preferences. Conversely, the DPO-tuned model, which may already possess some inherent robustness, shows more modest gains. This indicates that the utility of RPS may be inversely related to the base model's intrinsic robustness. Qualitative review further confirms that RPS achieves superior alignment by producing more detailed and nuanced responses that better match user intent, as shown in the case studies in Appendix A.6.

6 Conclusion

We have shown that the brittleness of preference-aligned models in out-of-distribution (OOD) scenarios can be effectively mitigated without retraining. Our proposed method, Robust Preference Selection (RPS), shifts from single-point generation to a more robust neighborhood consensus approach. It generates a diverse set of candidate responses from a local neighborhood of the target preference, which we show is theoretically guaranteed to produce a superior candidate pool compared to repeated sampling from the target direction itself. The optimal response is then selected using the original user preference. Extensive experiments across DPA, DPO, and SFT paradigms validate this approach, demonstrating significant robustness gains—up to a 69% win rate—for challenging OOD preferences. This work provides a practical, model-agnostic solution to the preference coverage gap and suggests that inference-time steering via neighborhood consensus is a promising path toward more adaptable and trustworthy AI systems.

ETHICS STATEMENT

This research aims to enhance the reliability and controllability of large language models, a goal with positive societal implications. Our work exclusively utilizes publicly available and widely used datasets (UltraFeedback, HelpSteer, and HelpSteer2) and open-source models. The datasets are standard benchmarks for preference alignment research and do not contain personally identifiable information. Our proposed method, RPS, is a post-hoc technique that does not involve model retraining, thereby avoiding the significant computational costs and environmental impact associated with it. We do not foresee any direct negative ethical implications arising from this work.

REPRODUCIBILITY STATEMENT

We are committed to ensuring the reproducibility of our research. All models used in our experiments (DPA-v1-Mistral-7B, Zephyr-7B-Beta, and Mistral-7B-Instruct-v0.2) are publicly available on the Hugging Face Hub, and direct links are provided in Section 4.1.1. Similarly, the datasets (UltraFeedback, HelpSteer, and HelpSteer2) are publicly accessible and cited. Our experimental setup, including the baseline and RPS configurations, is detailed in Section 4.1.2, with key hyperparameters (k=5, $\theta_{\rm max}=30^{\circ}$) specified. The Appendix provides further essential details for replication, including the exact prompts used for generation and evaluation (Appendix A.1 and A.2), the reward model scoring procedure (Appendix A.3), and the precise preference vectors used for evaluation (Appendix A.4). We believe this provides sufficient information for our results to be independently reproduced. We also provide our code and data at https://anonymous.4open.science/r/Robust_Preference_Selection-C65C.

REFERENCES

- Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. A general theoretical paradigm to understand learning from human preferences, October 2023. URL https://arxiv.org/abs/2310.12036.
- Aharon Ben-Tal, Dick den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, Feb 2013. doi: 10.1287/mnsc.1120.1641.
- Bobbili Sarat Chandra, Ujwal Dinesha, Dheeraj Narasimha, and Srinivas Shakkottai. Pita: Preference-guided inference-time alignment for llm post-training, 2025. URL https://arxiv.org/abs/2507.20067.
- Mohamad Chehade, Soumya Suvra Ghosal, Souradip Chakraborty, Avinash Reddy, Dinesh Manocha, Hao Zhu, and Amrit Singh Bedi. Bounded rationality for Ilms: Satisficing alignment at inference-time, 2025. URL https://arxiv.org/abs/2505.23729.
- Keertana Chidambaram, Karthik Vinay Seetharaman, and Vasilis Syrgkanis. Direct preference optimization with unobserved preference heterogeneity, 2024. URL https://arxiv.org/abs/2405.15065.
- Paul Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences, June 2017. URL https://arxiv.org/abs/1706.03741.
- Luzi S. Cui et al. Ultrafeedback: Boosting llms with high-quality feedback data, 2023. Dataset: HuggingFaceH4/ultrafeedback_binarized.
- Thomas G. Dietterich. *Ensemble methods in machine learning*, pp. 1–15. January 2000. doi: 10.1007/3-540-45014-9_1. URL https://link.springer.com/chapter/10.1007/3-540-45014-9_1.
- Yi Dong, Zhilin Wang, Makesh Narsimhan Sreedhar, Xianchao Wu, and Oleksii Kuchaiev. Steerlm: Attribute conditioned sft as an (user-steerable) alternative to rlhf, 2023. URL https://arxiv.org/abs/2310.05344.

- John Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization, 2018. URL https://arxiv.org/abs/1810.08750.
 - John Duchi, Peter Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach, 2016. URL https://arxiv.org/abs/1610.03425.
 - Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel Ni, and Jian Guo. A survey on llm-as-a-judge, 2024.
 - Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding, September 2020. URL https://arxiv.org/abs/2009.03300.
 - Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, et al. Mistral 7b instruct, 2023.
 - Yichen Li, Zhiting Fan, Ruizhe Chen, Xiaotang Gai, Luqi Gong, Yan Zhang, and Zuozhu Liu. Fairsteer: Inference time debiasing for llms with dynamic activation steering, 2025. URL https://arxiv.org/abs/2504.14492.
 - Xun Liang, Hanyu Wang, Yezhaohui Wang, Shichao Song, Jiawei Yang, Simin Niu, Jie Hu, Dan Liu, Shunyu Yao, Feiyu Xiong, and Zhiyu Li. Controllable text generation for large language models: a survey, August 2024. URL https://arxiv.org/abs/2408.12599.
 - Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv:2203.02155*, Mar 2022. URL https://arxiv.org/abs/2203.02155.
 - Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model, 2023. URL https://arxiv.org/abs/2305.18290.
 - Sadat Shahriar, Zheng Qi, Nikolaos Pappas, Srikanth Doss, Monica Sunkara, Kishaloy Halder, Manuel Mager, and Yassine Benajiba. Inference time Ilm alignment in single and multidomain preference spectrum, 2024. URL https://arxiv.org/abs/2410.19206.
 - Lewis Tunstall et al. Zephyr: A chaotic good model, 2023. Hugging Face blog.
 - Chaoqi Wang, Yibo Jiang, Chenghao Yang, Han Liu, and Yuxin Chen. Beyond reverse kl: Generalizing direct preference optimization with diverse divergence constraints, 2023. URL https://arxiv.org/abs/2309.16240.
 - Haoxiang Wang, Yong Lin, Wei Xiong, Rui Yang, Shizhe Diao, Shuang Qiu, Han Zhao, and Tong Zhang. Arithmetic control of llms for diverse user preferences: Directional preference alignment with multi-objective rewards, 2024a. URL https://arxiv.org/abs/2402.18571.
 - Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv*:2203.11171, Oct 2022. URL https://arxiv.org/abs/2203.11171.
 - Zhilin Wang, Yi Dong, Olivier Delalleau, Jiaqi Zeng, Gerald Shen, Daniel Egert, Jimmy J. Zhang, Makesh Narsimhan Sreedhar, and Oleksii Kuchaiev. Helpsteer2: Open-source dataset for training top-performing reward models, 2024b. URL https://arxiv.org/abs/2406.08673.
 - Jiaheng Wei, Harikrishna Narasimhan, Ehsan Amid, Wen-Sheng Chu, Yang Liu, and Abhishek Kumar. Distributionally robust post-hoc classifiers under prior shifts, 2023.
 - Zaiyan Xu, Sushil Vemuri, Kishan Panaganti, Dileep Kalathil, Rahul Jain, and Deepak Ramachandran. Robust Ilm alignment via distributionally robust direct preference optimization, 2025. URL https://arxiv.org/abs/2502.01930.

Linyi Yang, Shuibai Zhang, Libo Qin, Yafu Li, Yidong Wang, Hanmeng Liu, Jindong Wang, Xing Xie, and Yue Zhang. Glue-x: Evaluating natural language understanding models from an out-of-distribution generalization perspective, 2022. URL https://arxiv.org/abs/2211.08073.

Linyi Yang, Yaoxian Song, Xuan Ren, Chenyang Lyu, Yidong Wang, Jingming Zhuo, Lingqiao Liu, Jindong Wang, Jennifer Foster, and Yue Zhang. Out-of-distribution generalization in natural language processing: Past, present, and future. In *EMNLP*, 2023. doi: 10.18653/v1/2023. emnlp-main.276.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging Ilm-as-a-judge with mt-bench and chatbot arena, 2023.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, and OpenAI. Fine-tuning language models from human preferences. *arXiv*, January 2020. URL https://arxiv.org/abs/1909.08593v2.

A APPENDIX

A.1 USE OF LARGE LANGUAGE MODELS

This paper was prepared in accordance with ICLR's policy on Large Language Models (LLMs). The following checklist details the use of LLMs in this work:

• To aid or polish writing? Yes. LLMs were used to improve grammar, clarity, and phrasing throughout the manuscript.

A.2 RESPONSE GENERATION PROMPTS

All models (DPA, DPO, SFT) use the same system prompt format for generating responses:

```
System: "You are a helpful assistant. Your response should maximize weighted rating = helpfulness*{weight_helpfulness} + verbosity*{weight_verbosity}."

User: {original_prompt}
```

Where weight_helpfulness and weight_verbosity are integers from 0 to 100 satisfying the unit circle constraint (weight_verbosity/100) **2 + (weight_helpfulness/100) **2 == 1. This unified approach, adapted from the experimental setup in (Wang et al., 2024a), enables fair comparison across all training paradigms.

A.3 PREFERENCE-ALIGNED JUDGE PROMPTS

More aligned: A | B | Tie

We use preference-aligned A/B/TIE evaluation with randomized positioning to eliminate bias. The judge prompt template is:

```
You are a careful evaluation judge. Decide which response better
637
      aligns with the user preference that balances two criteria with
638
      given weights:
639
      - Helpfulness (weight = {wh})
640
      - Verbosity (weight = {wv})
641
642
      Safety/correctness override:
643
      - Factually wrong, unsafe, or misleading content should lose
644
      regardless of verbosity.
645
      Output format (exactly):
646
      Comparison: <one-sentence comparison and reason>
647
```

Query: {user_query}
Response A: {response_1}

Response B: {response_2}

For each target direction $\mathbf{v} = (v_h, v_v)$, we set the weight placeholders to the direction components and evaluate response pairs. A/B positions are randomized with a fixed seed for reproducibility.

A.4 REWARD MODEL SCORING

All experiments utilize the publicly available reward model Haoxiang-Wang/RewardModel-Mistral-7B-for-DPA-v1 from Wang et al. (2024a), which is trained to predict scores across multiple preference dimensions. To obtain the reward vector $\mathbf{r}(x,y)=(r_h(x,y),r_v(x,y))$ for a given prompt-response pair, we format the input according to the model's required template:

[INST] You must read the following conversation carefully and rate the assistant's response from score 0-100 in these aspects: helpfulness, correctness, coherence, honesty, complexity, verbosity

```
User: {prompt}
Assistant: {response} [/INST]
```

The model returns a vector of scores for each attribute mentioned in the prompt. For our twodimensional analysis, we extract the first score as helpfulness (r_h) and the sixth score as verbosity (r_v) to construct the reward vector used for all calculations and selection criteria in our work.

A.5 Preference Direction Specifications

Table 5 provides the specification of preference directions used in our experiments. Our evaluation focuses on directions \mathbf{v}_1 through \mathbf{v}_8 as these represent increasingly challenging preference configurations that extend beyond typical training ranges.

Table 5: Preference direction specifications with exact vector components and angles.

Direction	Vector (v_h, v_v)	Angle (°)
\mathbf{v}_1	(0.9848, 0.1736)	10.0
\mathbf{v}_2	(0.9659, 0.2588)	15.0
\mathbf{v}_3	(0.9397, 0.3420)	20.0
\mathbf{v}_4	(0.9063, 0.4226)	25.0
\mathbf{v}_5	(0.8660, 0.5000)	30.0
\mathbf{v}_6	(0.8192, 0.5736)	35.0
\mathbf{v}_7	(0.7660, 0.6428)	40.0
\mathbf{v}_8	(0.7071, 0.7071)	45.0

A.6 QUALITATIVE CASE STUDIES

To provide a more concrete understanding of our method's improvements, this section presents qualitative comparisons for each model. The examples, detailed in Table 6, are selected from the most out-of-distribution preference direction (\mathbf{v}_8 , angle 45°) to illustrate how RPS generates more nuanced and detailed responses compared to the baseline.

A.7 PROOF OF COROLLARY 1

We provide a brief justification for the two claims in Corollary 1.

Dependence on Neighborhood Size k: Let $G(k) = \mathbb{E}[\max(S_{RPS}^{(k)})] - \mathbb{E}[\max(S_{Baseline}^{(k)})]$ be the robustness gain for size k. The expected value of the maximum of a set of random variables is

Table 6: Qualitative comparisons for DPA, DPO, and SFT models on out-of-distribution preferences (\mathbf{v}_8 , angle 45°).

DPA Model

Prompt (UltraFeedback, ID: 10): "Write an investigative report on the state of homelessness in America."

Baseline Output: Title: The State of Homelessness in America: A Comprehensive Investigation... Homelessness is a pervasive issue that affects millions of people... Causes of Homelessness... Lack of affordable housing... Mental illness... Substance abuse...

RPS Output: Title: The State of Homelessness in America: A Comprehensive Investigation... The issue of homelessness has been a persistent challenge in the United States, with an estimated 580,466 people experiencing homelessness on a single night in 2019 (U.S. Department of Housing and Urban Development, 2020)... Current Trends: 1. Increasing Population... 2. Shift in Demographics... 3. Impact of COVID-19...

Judgment (RPS Win): GPT-40-mini selected the RPS response as more aligned.

Analysis: The RPS response is more helpful, including a specific statistic with a source, discussing more nuanced causes (Domestic Violence), and incorporating recent trends like the impact of COVID-19. The baseline is generic and lacks this specific, timely context.

DPO Model

Prompt (HelpSteer, ID: 5): "What are some synonyms for the word 'beautiful'?"

Baseline Output: Some synonyms for the word "beautiful" include gorgeous, lovely, pretty, stunning, radiant, enchanting, captivating, mesmerizing, breathtaking, and exquisite. All of these words convey various aspects of beauty... The choice of synonym depends on the context...

RPS Output: 1. Gorgeous 2. Stunning 3. Radiant 4. Magnificent 5. Attractive 6. Charming 7. Lovely 8. Enchanting 9. Exquisite 10. Alluring ... [list continues to 50 synonyms] ... Note: The weightage of each synonym may vary based on context and usage...

Judgment (RPS Win): GPT-40-mini selected the RPS response as more aligned.

Analysis: For a preference demanding high verbosity, the RPS response is clearly superior. It provides an extensive list of 50 synonyms, comprehensively addressing the user's request. The baseline provides only a handful of examples and a brief explanation.

SFT Model

Prompt (HelpSteer2, ID: 0): "explain master slave replication nsql"

Baseline Output: I'd be happy to help you understand master-slave replication in the context of NSQ... In master-slave replication, there is a primary node (master) that processes and writes messages, while secondary nodes (slaves) read from the master and maintain an identical copy of the data...

RPS Output: Master-slave replication is a method used in database systems... Here's how it works in the context of NSQ, a real-time distributed messaging platform: 1. Master: NSQ has a built-in database called Lookupd that stores metadata... 2. Slaves... 3. Replication lag... 4. Failover...

Judgment (RPS Win): GPT-40-mini selected the RPS response as more aligned.

Analysis: The RPS response provides a more technically accurate and structured explanation. It correctly identifies 'Lookupd' as the key component and explains concepts like replication lag and failover. The baseline's explanation is generic and less specific to NSQ's architecture.

non-decreasing with the size of the set. Therefore, both $\mathbb{E}[\max(S_{RPS}^{(k)})]$ and $\mathbb{E}[\max(S_{Baseline}^{(k)})]$ are non-decreasing in k. The gain increases because the expected improvement from adding an additional sample is greater for the RPS pool. Let $M_k^{RPS} = \max(S_{RPS}^{(k)})$. The increase in expected maximum is $\mathbb{E}[\max(M_k^{RPS}, s_{k+1})] - \mathbb{E}[M_k^{RPS}]$. Since the distribution of s_{k+1} stochastically dominates that of a baseline sample, this improvement is larger than the corresponding improvement for the baseline, causing the gap G(k) to widen.

Dependence on Quality Gap: We can formalize the "quality gap" as the degree of stochastic dominance. Let the RPS scores $\{s_i\}$ be drawn from distributions $\{F_i\}$, and consider an alternative set of "higher-quality" distributions $\{G_i\}$ such that each G_i stochastically dominates the corresponding F_i (i.e., $G_i(x) \leq F_i(x)$ for all x). Let S'_{RPS} be a set of scores drawn from $\{G_i\}$. Then $\max(S'_{RPS})$

stochastically dominates $\max(S_{RPS})$. This implies $\mathbb{E}[\max(S'_{RPS})] \geq \mathbb{E}[\max(S_{RPS})]$. The robustness gain relative to the fixed baseline therefore increases as the quality of the neighborhood candidate pool improves.