One Model for All Preferences: Meta-Learning for Multi-Objective LLM Alignment

Anonymous Authors[1]

Abstract

Aligning large language models (LLMs) often requires balancing multiple objectives or rewards, where the appropriate trade-off is inherently context-dependent: different users and deployment scenarios demand different balances among competing objectives such as helpfulness and safety. This motivates training single policies conditioned on preference vectors over multiple objectives, to enable user-specified trade-offs at inference time. However, existing approaches typically treat preferences merely as auxiliary variables, yielding policies with limited sensitivity to preferences and poor generalization to unseen trade-offs. We introduce MERIDIAN, a meta-learning framework that explicitly formulates alignment as a multi-task problem. By leveraging a bi-level algorithm design, MERIDIAN ensures the policy actively navigates the Pareto frontier rather than passively observing the conditioning signal. Furthermore, we adopt a smoothed Tchebycheff scalarization to expose trade-offs even in non-convex regions, making preference-following an explicit capability. Empirically, MERIDIAN enables precise, continuous model behavior control—supporting "one model for all preferences"—and produces coherent Pareto sweeps with significantly improved robustness to unseen preferences compared to baselines. We further provide a theoretical result showing that optimizing an empirical objective over sampled preferences generalizes to the full population.

[1]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

The post-training pipeline of Large Language Models (LLMs) involves the critical process of alignment, where the model behavior remains consistent with human values, intent, and safety constraints (Ouyang et al., 2022; Bai et al., 2022). However, this alignment is fundamentally a multi-objective problem. Current standard approaches, such as Reinforcement Learning from Human Feedback (RLHF), typically address this complexity by collapsing diverse and often conflicting objectives—such as helpfulness and harmlessness—into a single, static reward scalar(Wu et al., 2023). This "one-size-fits-all" paradigm effectively locks the model into a fixed trade-off, a strategy that is increasingly insufficient for the demands of real-world deployment.

In practice, the "optimal" behavior of an LLM is not a static point on the Pareto frontier but a dynamic, context-dependent equilibrium that varies across different contexts (Li et al., 2025). Consider the conflict between helpfulness and harmlessness: a creative writing assistant must prioritize helpfulness and stylistic diversity, perhaps even loosening safety guardrails to facilitate the portrayal of fictional villains. Conversely, a medical or legal chatbot must prioritize harmlessness and rigorous caution above all else, even when such constraints significantly incur an "alignment tax" by diminishing the model's immediate helpfulness. Because this tax cannot be paid with a single, inflexible currency, robust alignment necessitates models capable of navigating the trade-off space.

To address this, the field has moved toward *preference-conditioned alignment*, where a single policy is trained to dynamically adjust its behavior based on a user-specified preference vector w (Gupta et al., 2025; Wang et al., 2024). Ideally, such a model would grant users precise, continuous control over the generation, allowing them to steer the model to any point on the optimal trade-off curve (the Pareto frontier) simply by varying w. However, achieving this vision of "one model for all preferences," has proven algorithmically difficult. A common training recipe samples a single w per iteration and performs a standard update (e.g., PPO or DPO) on a linear scalarization $w^\top r$ (Gupta et al., 2025; Wang et al., 2024). We highlight two fundamental limitations of this paradigm. *First **(Optimization Gap)**:*

when objectives truly conflict, updates that improve one preference often degrade others. Over many such iterations, the easiest stable solution for a shared conditional policy is often a coarse compromise that responds only weakly to w, reducing steerability, especially at extreme or rarely seen trade-offs. *Second* ***(Geometric Gap)****:* linear scalarization cannot recover Pareto-optimal solutions in concave regions of the frontier (see Fig. 1) (Lin et al., 2024).

This paper aims to answer the key question: *How can we develop a single alignment policy that provides precise, continuous control over multiple conflicting objectives, enabling robust generalization to any user-specified preference across the entire Pareto frontier?* To address this, we propose MERIDIAN (*Meta-Learning for Preference-Conditioned Multi-Objective Alignment*), a framework that reimagines preference conditioning not as a supervised context problem, but as a *multi-task meta-learning* problem. Our central insight is that each preference vector w defines a distinct *alignment task* with its own optimization landscape. Rather than mixing these tasks, MERIDIAN employs a bi-level algorithmic scheme: an inner loop that aggressively optimizes specific trade-offs in isolation, and an outer meta-loop that aggregates these capabilities into a single, robustly steerable policy.

MERIDIAN introduces three key technical innovations to realize this vision. First, we replace linear scalarization with a Smoothed Tchebycheff objective (Lin et al., 2024). This geometry-aware formulation minimizes the worst-case weighted deviation from a utopia point, allowing the model to recover the *entire* Pareto frontier, including non-convex regions. Second, we implement the inner loop using Group Relative Policy Optimization (GRPO) (Shao et al., 2024), which eliminates the need for unstable value networks (critics) in the multi-task setting. Finally, we employ a Reptile-style Meta-Learning update rule (Nichol et al., 2018) to mitigate gradient interference, ensuring the shared parameters remain in a region from which any specific trade-off is easily reachable. In summary, our main contributions are as follows:

- **Meta-Learning Framework:** We formulate preference-conditioned alignment as a generalization problem over the preference simplex and propose MERIDIAN, a bi-level algorithm that learns to generalize to unseen preferences.
- **Geometric Robustness:** We demonstrate that Tchebycheff scalarization is essential for controllable steerability in LLMs, effectively closing the gaps left by linear scalarization in non-convex trade-offs.
- **Empirical Performance:** Experiments on conflicting objectives using Qwen2-0.5B-instruct show that MERIDIAN Pareto-dominates standard conditional baselines, offering precise inference-time steerability even for preference vectors unseen during training.
- **Theoretical Guarantee:** We derive an approximate-Empirical Risk Minimization (ERM) generalization bound, proving that the gap between the learned policy and the optimal population policy vanishes at a rate of $O(L/\sqrt{M})$, where L is the smoothness of the landscape and M is the number of sampled preference tasks.

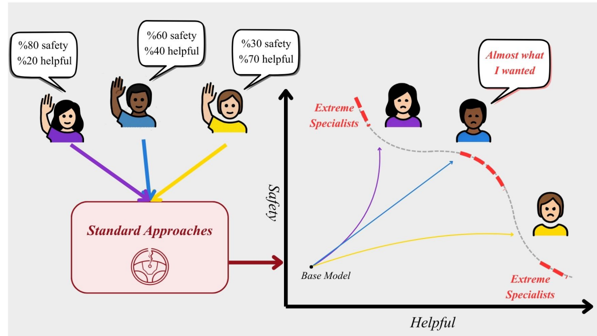

(a) Standard Methods

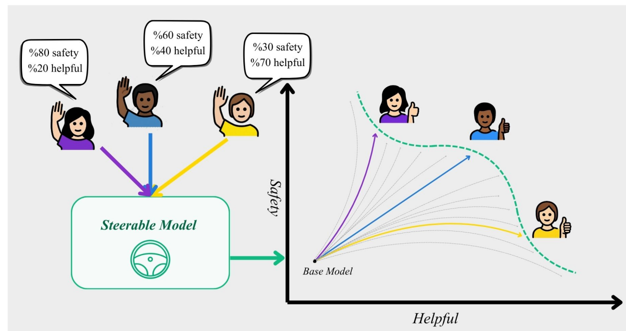

(b) MERIDIAN

Figure 1. **Conditioned Alignment and Pareto Coverage. (a) Standard Methods** (e.g., preference-conditioned RLHF/DPO) typically fail to cover the full Pareto front for two reasons: (1) *Geometric Gap:* linear scalarization ($w^\top r$) is theoretically incapable of recovering solutions in non-convex regions of the frontier; and (2) *Optimization Gap:* conflicting gradients from different preferences often lead to a single static "compromise" policy rather than dynamic steerability. **(b) MERIDIAN (Ours)** resolves these limitations by combining a meta-learning framework with Tchebycheff scalarization. This allows the model to actively navigate the entire trade-off landscape—including non-convex regions—producing a single policy that is precisely steerable to any user preference at inference time.

2. Problem Formulation

We consider the alignment of a language model with respect to K diverse objectives (e.g., helpfulness and safety).

Let $x \sim \mathcal{D}_x$ denote a prompt and let y denote a generated response. We assume access to a vector-valued reward function $r(x, y) \in \mathbb{R}^K$. A user's desired trade-off among these objectives is specified by a preference vector $w \in \Delta^K := \left\{ w \in \mathbb{R}_+^K : \sum_{i=1}^K w_i = 1 \right\}$. To motivate our framework, we distinguish between the conventional static formulation of multi-objective alignment and our proposed adaptive formulation.

Preference-Fixed MOO (Conventional). In the standard multi-objective setting, the goal is to learn an optimal policy π_w^* for a *single, fixed* preference vector w. This is typically formulated as maximizing a scalarized expected return:

$$\pi_{w,\theta}^* = \arg\max_{\pi} \ \mathbb{E}_{x\sim\mathcal{D}_x,\, y\sim\pi(\cdot|x)} \left[S(r(x, y), w) \right], \quad (1)$$

where $S(\cdot, w)$ is a scalarization function, commonly chosen to be linear, i.e., $S(r, w) = w^\top r$. While effective for optimizing a single trade-off, this formulation is impractical for real-world deployment. Covering the continuous simplex Δ^K would, in principle, require training and storing an uncountable family of policies $\{\pi_w^* : w \in \Delta^K\}$.

Preference-Adaptive MOO (Ours). To overcome this limitation, we instead seek to learn a *single* preference-conditioned policy $\pi_\theta(\cdot \mid x, p(w))$, where $p(w)$ denotes a conditioning signal derived from the preference vector (e.g., a prompt prefix or learned embedding). We model preferences as random variables drawn from a distribution ρ over Δ^K. Our objective generalizes Eq. (1) to the following adaptive optimization problem:

$$\pi_\theta^* = \arg\max_{\theta} \ \mathbb{E}_{w\sim\rho} \left[\mathbb{E}_{x\sim\mathcal{D}_x,\, y\sim\pi_\theta(\cdot|x,p(w))} \left[S(r(x, y), w) \right] \right]. \quad (2)$$

This formulation transforms alignment from a static optimization problem into a *generalization* problem: the model must learn to recover the optimal behavior associated with a preference vector w from its conditioning signal $p(w)$, including for preferences not explicitly observed during training.

Scalarization for Non-Convex Frontiers. A critical component of the adaptive formulation in Eq. (2) is the choice of a scalarization function S. Standard approaches typically employ *linear scalarization*, $S_{\text{lin}}(r, w) = \sum_{i=1}^K w_i r_i$. However, linear scalarization is geometrically limited: it can identify only solutions lying on the *convex hull* of the Pareto frontier (Lin et al., 2024). To address this limitation, we define our preference-conditioned objective $J(\theta; w)$ using a smoothed Tchebycheff scalarization function S_{Tch} applied to the policy's expected returns. Given a reference (utopia) point $z^\star \in \mathbb{R}^K$ and a smoothing parameter $\mu > 0$, we define the parameter-level objective as

$$J(\theta; w) := \mathbb{E}_{x\sim\mathcal{D}_x,\, y\sim\pi_\theta(\cdot|x,p(w))} \left[S_{\text{Tch}}(r(x, y), w) \right]$$
$$S_{\text{Tch}}(r(x, y), w) = -\mu \log \sum_{i=1}^K \exp\left(\frac{w_i\big(z_i^\star - r_i(x, y)\big)}{\mu} \right) \quad (3)$$

This formulation provides a smooth approximation to the max-based Tchebycheff objective, which minimizes the worst-case weighted shortfall relative to the utopia point. Crucially, unlike linear scalarization, this objective can recover Pareto-optimal solutions in non-convex regions of the frontier, thereby enabling true controllability across the preference simplex.

The Empirical Meta-Objective: Combining the adaptive formulation with the Tchebycheff scalarization defined above, we seek to maximize the population meta-objective $G(\theta) = \mathbb{E}_{w\sim\rho}\big[J(\theta; w)\big]$. Since the preference distribution ρ is continuous and unknown, we approximate $G(\theta)$ using a finite collection of M-sampled preferences. This yields the empirical objective optimized by our algorithm:

$$\widehat{G}_M(\theta) = \frac{1}{M} \sum_{m=1}^M J(\theta; w_m); \quad w_{1:M} \overset{\text{i.i.d.}}{\sim} \rho. \quad (4)$$

3. MERIDIAN: Alignment via Distributional Meta-Learning

This section details the implementation of MERIDIAN, a framework designed to produce a single, steerable policy that generalizes across the entire preference simplex. Our design is grounded in a central insight: preference-conditioned alignment is structurally a *multi-task learning* problem, where each preference vector $w \in \Delta^K$ defines a distinct alignment task. This perspective motivates our adoption of a *meta-learning* approach, specifically a bi-level algorithmic design (Algorithm 1): an inner loop that optimizes a batch of preference trade-offs, and an outer meta-loop that aggregates these capabilities into a single, steerable policy. This stands in clear contrast to standard approaches, which do not adopt a meta-learning framework and instead treat w merely as an auxiliary conditioning variable, either processed jointly with the prompt (Gupta et al., 2025) or incorporated through parameter-space conditioning (Wang et al., 2024).We observe that this design restricts effective navigation of the Pareto frontier, especially when generalizing to unseen preferences, causing the learned policy to display diminished sensitivity to the conditioning signal and preventing fine-grained control over the full range of trade-offs (see Section 4).

The Meta-Optimization Loop: The core of our approach is a meta-update rule inspired by reptile-style meta-learning frameworks (Nichol et al., 2018; Finn et al., 2017). Our objective is to learn a policy π_θ that maintains high sensitivity to the conditioning signal, ensuring it can recover the optimal behavior for any preference $w \sim \rho$ rather than collapsing to a static average. As shown in Algorithm 1, each training iteration samples a batch of preference vectors $\{w_1, \dots, w_B\}$. Crucially, these preferences are not mixed within a single gradient update. Instead, the current parameters θ_t are branched into B independent copies. Each copy $\theta_{t,j}$ undergoes S steps of *inner-loop optimization* targeting the scalarized objective $J(\cdot; w_j)$. This isolation allows each branch to traverse its corresponding optimization landscape without interference from conflicting preference gradients, enabling convergence toward the Pareto-optimal region associated with w_j. Once the preference-adapted parameters $\{\theta_{t,1}, \dots, \theta_{t,B}\}$ are obtained, we perform the meta-update as follows:

$$\theta_{t+1} \leftarrow \theta_t + \frac{\alpha}{B} \sum_{j=1}^{B} (\theta_{t,j} - \theta_t) \,. \tag{5}$$

Rather than aggregating raw gradients, this update averages the *resulting adapted parameters*. This design stabilizes learning by smoothing high-variance conflicts between competing objectives, while preserving the model's ability to steer toward each of them through rapid preference-specific adaptation.

Inner Loop (Group Relative Policy Optimization) To efficiently implement the inner-loop updates, we employ Group Relative Policy Optimization (GRPO) (Shao et al., 2024). Meta-learning is computationally demanding: applying standard PPO in the inner loop would require training a separate value function for every sampled preference vector w_j. This is computationally prohibitive and unstable, as the definition of "value" shifts with the preference. GRPO eliminates the need for a critic entirely. For a fixed preference w_j and prompt x, we sample G outputs $\{y_1, \dots, y_G\}$ from the current inner-loop policy. The policy is updated using the normalized group-relative advantage enables stable, KL-constrained updates without a value network, making the meta-learning framework tractable for large language models.

3.1. Geometry-Aware Scalarization

The final component of MERIDIAN is the choice of scalarization function. A meta-learning algorithm is only as expressive as the tasks it optimizes. Prior work commonly relies on linear scalarization of the form $w^\top r$ (Gupta et al., 2025; Wang et al., 2024), which is theoretically incapable of recovering Pareto-optimal solutions in non-convex regions of the frontier (Lin et al., 2024). In large language model alignment, trade-offs between safety, helpfulness,

Algorithm 1 MERIDIAN: Meta-Learning for Preference-Conditioned Multi-Objective Alignment

Require: Preference distribution ρ; conditioner $p(\cdot)$; initial parameters θ_0; meta-iterations T; GRPO steps S; prompts distribution $\mathcal{D}_x$.

1: **for** $t = 0, 1, \dots, T-1$ **do**
2: Sample a batch of preferences $\{w_j\}_{j=1}^{B} \overset{iid}{\sim} \rho$.
3: **for** $j = 1, \dots, B$ **(GRPO inner updates under preference** w_j**) do**
4: Initialize preference-adapted parameters $\theta_{t,j}^{(0)} \leftarrow \theta_t$.
5: **for** $s = 0, 1, \dots, S-1$ **do**
6: Sample prompts $\{x_n\}_{n=1}^{N} \sim \mathcal{D}_x$ and roll out $y_n \sim \pi_{\theta_{t,j}^{(s)}}(\cdot \mid x_n, p(w_j))$.
7: Score (x_n, y_n) with K reward models to obtain vector scores $R_n \in \mathbb{R}^K$.
8: Form scalarized rewards using the smoothed scalarization in Eq. (3) under w_j.
9: Update the policy with a GRPO step to obtain $\theta_{t,j}^{(s+1)}$.
10: **end for**
11: Set $\theta_{t,j} \leftarrow \theta_{t,j}^{(S)}$.
12: **end for**
13: **Reptile-style meta update:** update θ_{t+1} using Eq. (5).
14: **end for**
15: $\tilde{\theta} \leftarrow \theta_T$.

and other objectives are frequently non-linear. Linear scalarization therefore introduces irrecoverable "gaps" in steerability. To address this limitation, we adopt the smoothed Tchebycheff scalarization defined in Eq. 3. By optimizing an exponentially weighted shortfall from a utopia point z^*, this objective enables the inner loop to recover Pareto-optimal solutions that linear combinations would systematically miss. This geometry-aware scalarization ensures that the meta-update aggregates genuinely diverse optimal behaviors, rather than collapsing toward a narrow set of linear compromises.

4. Experiments

MERIDIAN was designed around two practical goals that arise in preference-conditioned alignment: (i) *generalization across preferences*: can a policy trained on only finitely many sampled preferences behave sensibly on unseen $w \sim \rho$? and (ii) *coverage of Pareto trade-offs*: can users continuously steer the model by varying w at inference time, rather than being locked into a single compromise? In Section 5, we provide an algorithm-agnostic guarantee for generalization over a preference distribution. Here we move from the population objective to concrete post-training runs and ask the same questions empirically: do we observe smooth preference-following behavior in practice, and does varying w induce meaningful trade-offs across objectives?

4.1. Experimental setup

Model. We use **Qwen2-0.5B-Instruct** as the base model throughout. It is a strong instruction-tuned backbone with coherent multi-turn behavior, making it a practical testbed. It is large enough to exhibit meaningful alignment tensions (e.g., helpfulness vs. safety), yet small enough to run extensive preference sweeps and ablations at reasonable cost.

Preference distributions. Preferences are a user-facing control knob at inference time: users can supply any $w \in \Delta^K$. During training, we must choose how to sample preferences to expose the policy to a range of trade-offs. We study two sampling schemes: (i) uniform-like sampling over Δ^K, and (ii) Dirichlet-skewed sampling that emphasizes corners or asymmetric trade-offs. In practice, we use uniform sampling for training; the rationale is discussed in the ablation study in Section 4.6.

Training protocol and budget. We train using the MERIDIAN meta-learning loop with GRPO-style post-training updates and the smoothed scalarization in Eq. (3). All runs use **4 H100 GPUs** and train end-to-end under a fixed recipe. We evaluate periodically on a held-out prompt set and a held-out preference set, select the final checkpoint using validation performance under the training judge, and additionally inspect generations at representative preferences (e.g., near corners of Δ^K) to ensure the chosen checkpoint produces meaningful, non-degenerate behavior.

4.2. Tasks and reward models

Helpfulness vs. Harmlessness (Anthropic HH). We evaluate the standard alignment tension between helpfulness and safety. Using established `Ray2333/gpt2-large-helpful-reward\_model` and `Ray2333/gpt2-large-harmless-reward\_model` reward models, we compute a vector $r(y) = [r_{\text{help}}(y), r_{\text{harm}}(y)]$. The objective is to train a single policy conditioned on $w \in \Delta^2$ that dynamically interpolates behavior—shifting from direct assistance to cautious refusal—strictly based on the inference-time preference.

4.3. Evaluation Metrics

To assess whether our policy generalizes to unseen preferences rather than simply memorizing training modes, we evaluate on a broader set that *includes* the training preferences as well as additional preferences not used during training. Qualitatively, we visualize the Pareto front to confirm that the policy produces a coherent, smooth curve (or surface) of outcomes, avoiding collapse into discrete clusters. Quantitatively, we report the *Preference Utility* (PU), defined as $\mathbf{PU} := \mathbb{E}[S_{\text{Tch}}(r(x, y), w)]$, which directly measures how effectively the conditional policy satisfies the specific trade-offs of configurations not encountered during training. (see Table 1)

4.4. Baselines

We compare MERIDIAN against three baselines that represent the state-of-the-art and standard paradigms in preference-conditioned alignment:

1. **MO-ODPO (SOTA Preference-Conditioned DPO).** This method represents the current state-of-the-art in preference-conditioned alignment (Gupta et al., 2025). It aggregates objective-specific rewards using a weighted linear sum $w^\top r$ and optimizes the policy using the standard DPO loss. This serves as our primary external baseline, testing the efficacy of our meta-learning approach against established direct preference optimization methods.

2. **Grid-MO-GRPO.** This baseline represents the standard preference-conditioned alignment approach (Gupta et al., 2025) adapted to the GRPO framework. While the state-of-the-art method typically relies on DPO, we implemented this version using the same GRPO optimizer as MERIDIAN to ensure a fair, controlled comparison. This allows us to evaluate the standard approach under the same on-policy optimization conditions as our method.

3. **Weighted-MO-GRPO.** To isolate the contribution of the meta-update, this baseline follows the exact same preference-conditioned GRPO training recipe (including Tchebycheff scalarization) but omits the meta-learning outer loop.

4.5. Empirical Analysis of Pareto Coverage

We primarily evaluate MERIDIAN on its ability to recover the Pareto frontier and provide dense, responsive steerability across the preference simplex. Figure 2 presents the inference-time preference sweeps for the Anthropic HH task (Helpfulness vs. Harmlessness).

Generalization via Meta-Learning. MERIDIAN (blue) attains the strongest frontier, consistently dominating the SOTA baseline **MO-ODPO** (red). Crucially, this performance gain is not merely an artifact of switching optimizers: **Grid-MO-GRPO** (green), which adapts the SOTA approach to use GRPO, also fails to match MERIDIAN's coverage. This isolates the benefit of our architecture beyond just the choice of on-policy optimization. Furthermore, the comparison with **Weighted-MO-GRPO** (orange) provides the most rigorous test of our contributions. Despite utilizing both the GRPO optimizer and the Tchebycheff scalarization, this baseline still significantly under-

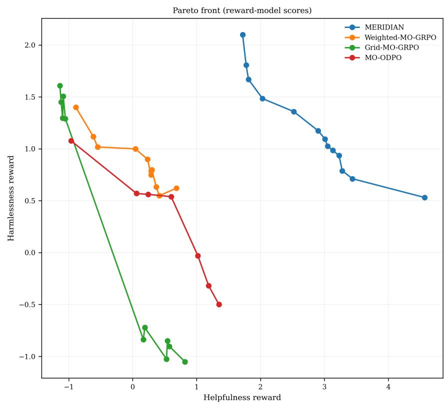

Figure 2. **HH Pareto fronts (reward-model scores).** MERIDIAN (blue) achieves the strongest frontier. The gap to Weighted-MO-GRPO (orange) highlights that the meta-learning loop is essential. Furthermore, MERIDIAN recovers solutions in the non-convex interior that linear baselines (green, red) systematically miss.

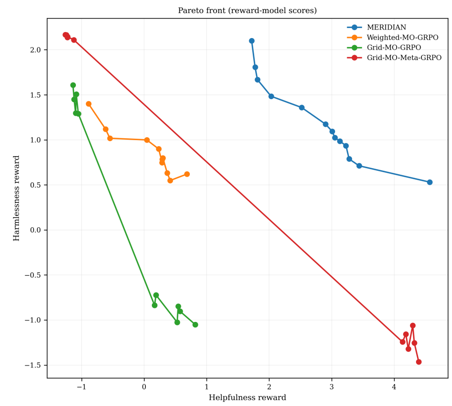

Figure 3. **Impact of Scalarization Geometry.** Even with meta-learning, replacing Tchebycheff scalarization with linear weights (Grid-MO-Meta-GRPO, red) causes the policy to lose coverage of the non-convex middle region, confirming that geometric awareness is a prerequisite for full steerability.

performs MERIDIAN. This confirms that even with the correct optimizer and geometry-aware scalarization, standard preference-conditioned training remains insufficient. The substantial performance gap between them confirms that the meta-update rule is the primary driver of our superior performance.

Recovery of Non-Convex Frontiers. The second significant observation from Figure 2 is the geometric shape of the frontier recovered by MERIDIAN. The curve exhibits a clear non-convexity in the transition region between extreme helpfulness and extreme harmlessness. This confirms that the underlying trade-off landscape for LLM alignment is frequently non-convex. MERIDIAN successfully traces this concave region, maintaining a dense set of distinct, non-dominated solutions. In contrast, **Grid-MO-GRPO** (green) and **MO-ODPO** (red)—which rely on linear scalarization—fail to populate this region.

Coverage of Extremes. Finally, we observe a failure mode in **MO-ODPO** (red). Rather than covering the full spectrum, its generations tend to cluster near the middle of the trade-off space, failing to reach the extreme regions of the frontier. This indicates that while MO-ODPO can produce stable compromise behaviors, it struggles to push the policy toward the specialized extremes (e.g., maximum helpfulness or maximum safety) required by high-stakes deployment. In contrast, MERIDIAN covers the entire Pareto frontier, demonstrating the ability to seamlessly interpolate from one extreme to the other.

4.6. Ablation Studies

To dissect the sources of MERIDIAN's performance, we perform component-wise ablations on the scalarization function, the training distribution, and the regularization strength.

The Necessity of Tchebycheff Scalarization. In Figure 3, we analyze **Grid-MO-Meta-GRPO** (red), a variant of MERIDIAN that retains the meta-learning machinery but replaces the smoothed Tchebycheff objective (Eq. 3) with a linear weighted sum. The result is a "blind spot" in the middle of the trade-off curve, reproducing the failure mode of the non-meta baselines. This creates a strong dissociation between the algorithm's components: the meta-learning loop ensures distinct tasks are learnable without interference, but the *scalarization function* determines which points on the frontier constitute valid tasks. Without Tchebycheff scalarization, the inner loop simply cannot "see" the non-convex preferences, rendering the meta-learner incapable of mastering them.

Sensitivity to Training Distribution (α)**.** In Figure 4, we vary the Dirichlet concentration parameter α used to sample training preferences $w \sim \mathrm{Dir}(\alpha)$. A lower $\alpha = 0.5$ (corner-heavy) degrades performance in the middle of the Pareto front. This indicates that generalization over the simplex is not "free"; the meta-learner requires support in the interior of the simplex to learn the manifold of compromise solutions. Conversely, $\alpha = 1.0$ (uniform) provides sufficient coverage to resolve both the extremes and the interior. This result highlights a trade-off: over-emphasizing "pure" behaviors (experts at corners) during training can harm the smoothness of the interpolation between them.

The Alignment-Stability Trade-off (β)**.** Figure 5 demonstrates the effect of the KL-divergence penalty β. We observe a direct tension between alignment steerability and proximity to the base prior. With high regularization ($\beta = 0.6$), the policy is anchored tightly to the base instruction model, compressing the Pareto front and preventing the model from reaching the high-reward extremes. Relaxing

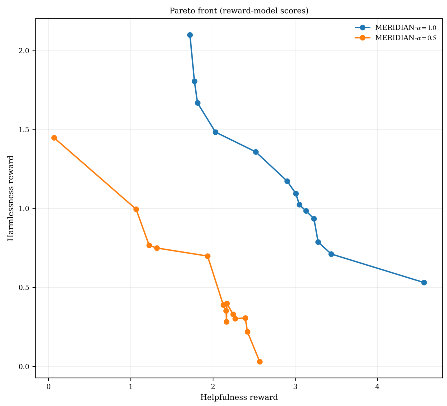

Figure 4. **Dirichlet Concentration Ablation.** Training with $\alpha = 1.0$ (uniform) yields a continuous, well-resolved frontier. Corner-heavy sampling ($\alpha = 0.5$) causes under-performance in the intermediate "compromise" regions.

this constraint ($\beta = 0.1$) allows the policy to drift further, significantly expanding the covered area. This implies that extreme alignment trade-offs (e.g., maximum safety or maximum helpfulness) are distributionally distinct from the "average" pre-trained behavior, requiring a larger KL budget to realize.

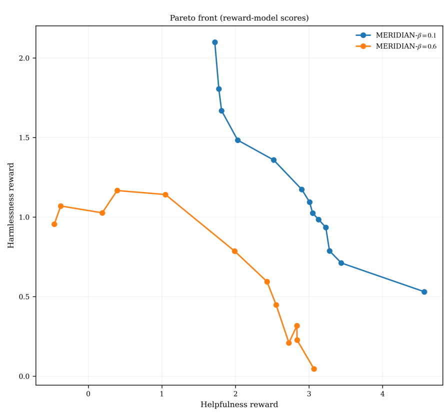

Figure 5. **KL Regularization (β) Ablation.** Tighter KL constraints ($\beta = 0.6$) compress the frontier, limiting the model's ability to achieve extreme trade-offs. A lower penalty ($\beta = 0.1$) is necessary to unlock the full range of steerability.

4.7. Training Dynamics

Finally, we analyze the temporal emergence of steerability in Figure 6. We observe a distinct curriculum in the learning process. At 30% training, the policy clusters near the center (close to the initialization). By 70%, the "compromise" regions in the interior are well-resolved, but the extremes remain underdeveloped. It is only in the final stages of training that the policy successfully pushes the frontier outward toward the single-objective vertices. This suggests that the interior of the simplex represents a more accessible optimization landscape (closer to the base model's prior), whereas the extreme behaviors (specialists) require sustained optimization pressure to diverge sufficiently from the initialization. For additional details regarding the experimental setup and hyperparameters, please refer to Appendix A.

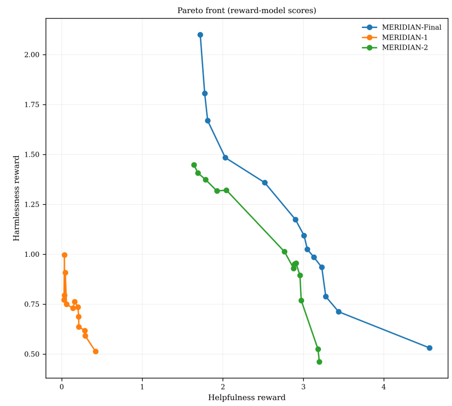

Figure 6. **Evolution of the Pareto Front.** Steerability emerges from the center outward. The model first masters balanced trade-offs (70%) before successfully pushing the boundary toward extreme preferences (100%).

5. Theoretical Analysis: Generalization over Preferences

Our empirical success relies on the premise that a policy trained on a finite batch of preferences $\{w_m\}_{m=1}^{M}$ will generalize to unseen preferences drawn from the distribution ρ. In this section, we formalize this intuition. We frame preference-conditioned alignment as an Approximate Empirical Risk Minimization (ERM) problem and derive a generalization bound that explicitly separates *statistical sampling error* (due to finite M) from *optimization error*.

5.1. Setup: Population vs. Empirical Objectives

We analyze generalization through the scalarized objective $J(\theta; w)$. Let $G(\theta) = \mathbb{E}_{w\sim\rho}[J(\theta; w)]$ denote the true population objective, corresponding to the model's performance averaged over the entire preference simplex. Since $G(\theta)$ is inaccessible in practice, we instead optimize the empirical objective

$$\widehat{G}_M(\theta) = \frac{1}{M}\sum_{m=1}^{M} J(\theta; w_m), \qquad w_{1:M} \overset{\text{i.i.d.}}{\sim} \rho. \tag{6}$$

Let θ^* denote a population maximizer of $G(\theta)$, and let $\tilde{\theta}$ be the solution returned by MERIDIAN. We characterize the performance of our algorithm via an *optimization error* $\varepsilon_{\text{opt}} \geq 0$, defined as the sub-optimality of the returned solution relative to the empirical maximizer $\widehat{\theta}$: $\widehat{G}_M(\widehat{\theta}) - \widehat{G}_M(\tilde{\theta}) \leq \varepsilon_{\text{opt}}$. This formulation is solveragnostic. It acknowledges that while globally optimizing the non-convex meta-learning objective is difficult, MERIDIAN's strong inner-loop optimizer (GRPO) is designed to minimize ε_{opt} as effectively as possible. To analyze generalization, we impose the following mild regularity assumptions.

Assumption 5.1 (Regularity of the Preference Landscape)**.** We make the following standard assumptions regarding the scalarized objective J and the data generating process: **1. Boundedness:** For all parameters $\theta \in \Theta$ and preferences $w \in \Delta^K$, the objective is bounded such that $J(\theta; w) \in [0, 1]$. **2. Lipschitz Continuity:** The mapping $w \mapsto J(\theta; w)$ is L-Lipschitz with respect to the ℓ_2-norm for all $\theta \in \Theta$, i.e., $|J(\theta; w) - J(\theta; w')| \leq L\|w - w'\|_2$. **3. I.I.D. Sampling:** The training preference vectors $\{w_m\}_{m=1}^M$ are sampled independently and identically distributed from the distribution ρ.

5.2. Generalization Guarantee

We now bound the performance gap between the learned policy $\tilde{\theta}$ and the optimal population policy θ^*.

Theorem 5.2 (Approximate-ERM Generalization)**.** *Under Assumption 5.1, assume the scalarized value function is bounded such that $J(\theta; w) \in [0, 1]$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the sampled preferences $w_{1:M}$,*

$$G(\theta^*) - G(\tilde{\theta}) \leq \underbrace{4\mathfrak{R}_M(F)}_{\textit{Model Capacity}} + \underbrace{2\sqrt{\frac{\log(2/\delta)}{2M}}}_{\textit{Sampling Error}} + \underbrace{\varepsilon_{opt}}_{\textit{Optimization Error}}, \tag{7}$$

where $\mathfrak{R}_M(F)$ denotes the empirical Rademacher complexity (Bartlett & Mendelson, 2002) of the function class $F = \{w \mapsto J(\theta; w) \ : \ \theta \in \Theta\}$. Moreover, Under the Lipschitz continuity condition in Assumption 5.1, the capacity term satisfies $\mathfrak{R}_M(F) = O\left(\frac{L}{\sqrt{M}}\right)$.

Proof. See Appendix C for the full derivation and detailed assumptions.

Insight and Interpretation. Theorem 5.2 decomposes the generalization gap into three interpretable components, each of which directly motivates a key design choice in MERIDIAN. The first two terms correspond to statistical generalization and decay as the number of sampled preferences M increases.

6. Related Work

Standard alignment strategies (Ouyang et al., 2022; Rafailov et al., 2023) typically optimize a single fixed trade-off. While GAPO (Li et al., 2025) use multi-objective optimization to embed a specific preference vector during training, this formulation yields a single policy tied to that fixed preference, requiring retraining for any new preference. Alternatives like parameter merging (Rame et al., 2023) achieve steerability but incur high storage costs, while standard conditional policies (Wang et al., 2024) relying on linear scalarization fail to recover solutions in concave regions of the Pareto front (Lin et al., 2024). In contrast, MERIDIAN formulates alignment as a bilevel meta-learning algorithm with Tchebycheff scalarization, learning a single policy that is strictly steerable across the entire non-convex frontier without additional storage overhead. Finally, we note concurrent work (Anonymous, 2026) that also investigates multi-objective alignment of large language models. However, the two papers consider fundamentally different alignment paradigms. The concurrent work is primarily optimization-driven, aiming to identify Pareto-stationary solutions via multi-gradient descent combined with weighted Tchebycheff scalarization. In contrast, our work centers on generalization and inference-time steerability, leveraging a meta-learning framework to train policies that can adapt to arbitrary and unseen preference vectors at deployment. We provide a more comprehensive discussion of related methods and their connections in Appendix 6.

7. Conclusion

In this work, we introduced MERIDIAN, a framework that reframes multi-objective alignment as a distributional meta-learning problem. Empirically, our results demonstrate that MERIDIAN consistently outperforms standard conditioning baselines, producing a steerable policy that achieves superior coverage of the Pareto frontier. By leveraging smoothed Tchebycheff scalarization, our approach can target Pareto-optimal solutions beyond the convex-hull region accessible to linear scalarization, which is particularly important in non-convex trade-off landscapes. Complementing these empirical findings, our theoretical analysis establishes that this meta-learning objective admits rigorous generalization guarantees, ensuring that the learned policy remains robust even when conditioned on preference configurations unseen during training. Future work will extend this framework to sequential decision-making in agentic tasks and multi-step reasoning, where preferences apply to entire trajectories rather than single-turn generations.

Impact Statement

This work advances preference-conditioned, multi-objective post-training for large language models, enabling a single policy to be steered across trade-offs at inference time. Such steerability can improve practical deployment by adapting behavior to context and user needs. As with any controllable system, preference steering should be deployed with appropriate safeguards (e.g., bounded preference ranges and safety audits) to reduce unintended or adversarial use.

References

Anonymous. Details omitted for blind review. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, 2026. Concurrent Submission.

Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of machine learning research*, 3(Nov):463–482, 2002.

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.

Désidéri, J.-A. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathematique*, 350(5-6):313–318, 2012.

Eysenbach, B., Salakhutdinov, R., and Levine, S. C-learning: Learning to achieve goals via recursive classification. *arXiv preprint arXiv:2011.08909*, 2020.

Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.

Gupta, R., Sullivan, R., Li, Y., Phatale, S., and Rastogi, A. Robust multi-objective preference alignment with online dpo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 27321–27329, 2025.

Jang, J., Kim, S., Lin, B. Y., Wang, Y., Hessel, J., Zettlemoyer, L., Hajishirzi, H., Choi, Y., and Ammanabrolu, P. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*, 2023.

Li, C., Zhang, H., Xu, Y., Xue, H., Ao, X., and He, Q. Gradient-adaptive policy optimization: Towards multi-objective alignment of large language models. *arXiv preprint arXiv:2507.01915*, 2025.

Lin, X., Zhang, X., Yang, Z., Liu, F., Wang, Z., and Zhang, Q. Smooth tchebycheff scalarization for multi-objective optimization, 2024. URL `https://arxiv.org/abs/2402.19078`.

Liu, G., Tang, M., and Eysenbach, B. A single goal is all you need: Skills and exploration emerge from contrastive rl without rewards, demonstrations, or subgoals. *arXiv preprint arXiv:2408.05804*, 2024.

McDiarmid, C. et al. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989.

Min, S., Lewis, M., Zettlemoyer, L., and Hajishirzi, H. Metaicl: Learning to learn in context. In *Proceedings of the 2022 conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2791–2809, 2022.

Nair, A. V., Pong, V., Dalal, M., Bahl, S., Lin, S., and Levine, S. Visual reinforcement learning with imagined goals. *Advances in neural information processing systems*, 31, 2018.

Nichol, A., Achiam, J., and Schulman, J. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

Pong, V. H., Dalal, M., Lin, S., Nair, A., Bahl, S., and Levine, S. Skew-fit: State-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*, 2019.

Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741, 2023.

Rame, A., Couairon, G., Dancette, C., Gaya, J.-B., Shukor, M., Soulier, L., and Cord, M. Rewarded soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. *Advances in Neural Information Processing Systems*, 36:71095–71134, 2023.

Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

Wang, K., Kidambi, R., Sullivan, R., Agarwal, A., Dann, C., Michi, A., Gelmi, M., Li, Y., Gupta, R., Dubey, K. A., et al. Conditional language policy: A general framework for steerable multi-objective finetuning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 2153–2186, 2024.

Wu, Z., Hu, Y., Shi, W., Dziri, N., Suhr, A., Ammanabrolu, P., Smith, N. A., Ostendorf, M., and Hajishirzi, H. Fine-grained human feedback gives better rewards for language model training. *Advances in Neural Information Processing Systems*, 36:59008–59033, 2023.

Yang, K., Liu, Z., Xie, Q., Huang, J., Zhang, T., and Ananiadou, S. Metaaligner: Towards generalizable multi-objective alignment of language models. *Advances in Neural Information Processing Systems*, 37:34453–34486, 2024.

This appendix provides supplementary material to support the main paper. We begin in Appendix A by presenting experimental results and describing our experimental setup. A more comprehensive discussion on related work is available in Appendix B. The subsequent sections are dedicated to our theoretical analysis.

A. Experimental Results for MERIDIAN

A.1. Base model and data

Base model. We use **Qwen2-0.5B-Instruct** as the base model throughout. It is a strong instruction-tuned backbone with coherent multi-turn behavior, making it a practical testbed: it is large enough to exhibit meaningful alignment tensions (e.g., helpfulness vs. harmlessness), yet small enough to run dense preference sweeps and ablations at reasonable cost.

Datasets. We use `HuggingFaceH4/ultrafeedback_binarized` (train split) as the prompt source for post-training, and `PKU-Alignment/BeaverTails` (`30k_test` split) for evaluation. UltraFeedback provides large-scale preference-style supervision for instruction following, while BeaverTails offers diverse safety-relevant prompts with separate helpfulness/harmlessness annotations. (**??**)

A.2. Preference distribution and evaluation weights

Preference vectors. Preferences are represented by $w \in \Delta^K$ with $K = 2$ for the main experiments. At inference time, users may supply any $w \in \Delta^2$. For training and controlled sweeps, we use a fixed set of 11 weights

$$W_{\text{train}} = \{(1, 0), (0.9, 0.1), \ldots, (0.5, 0.5), \ldots, (0.1, 0.9), (0, 1)\}.$$

In addition to W_{train}, we evaluate on extra held-out weights listed in Table 1.

Training preference distribution. During training, we model preferences as $w \sim \mathrm{Dir}(\alpha)$ and implement this as a sampling distribution over W_{train}, where α controls how strongly sampling emphasizes corners (extreme trade-offs) versus interior points. We sweep concentration values α and report the best-performing setting, selected by held-out preference validation as seen in Fig 3

A.3. MERIDIAN training procedure

Meta-learning loop. MERIDIAN follows Algorithm 1 for T meta-iterations. Each meta-iteration samples a batch of B preferences $\{w_j\}_{j=1}^B$, runs S inner-loop updates per preference to obtain preference-adapted parameters $\{\theta_{t,j}\}_{j=1}^B$, and applies the meta-update in Eq. (5). Unless otherwise stated, we use: (meta learning rate) 1.5 with a linear schedule, (batch of preferences) $B = 2$, and (meta-iterations) $T = 400$.

Inner-loop optimization (GRPO). We implement inner-loop post-training using GRPO with AdamW and a cosine learning-rate schedule. Unless otherwise stated, we use GRPO learning rate 10^{-6}, per-device batch size 2, sampling temperature 0.9, KL regularization coefficient $\beta = 0.1$, and smoothed Tchebycheff parameters matching Eq. (3) (smoothing $\mu = 0.1$ in our main runs). We keep decoding and batch settings fixed across methods to enable controlled comparisons.

A.4. Reward models and normalization

Reward models. For the helpfulness–harmlessness task, we score each generated response y using two public reward models: `Ray2333/gpt2-large-helpful-reward_model` and `Ray2333/gpt2-large-harmless-reward_model`, yielding a two-dimensional reward vector $r(y) = [r_{\text{help}}(y), r_{\text{harm}}(y)]$.

Reward normalization. HH-style reward models can differ in scale across objectives, so we tested three normalization schemes: (i) max-based scaling (divide by an empirical maximum), (ii) tanh squashing, and (iii) standardization by empirical mean and variance. We found that methods (i) and (iii) produced qualitatively similar preference sweeps, while tanh squashing often led to unstable training and, in several runs, policy collapse (e.g., reduced diversity and degenerate responses). Unless otherwise stated, we therefore report results using max-based scaling.

Table 1. Preference utility aggregated over evaluation prompts and preferences. Bolded w values were used during training W_{seen}; the remaining w values are unseen at training time W_{unseen}.

PREFERENCES	MERIDIAN	WEIGHTED-MO-GRPO
1.00↔0.00	**0.363**	**0.946**
0.99↔0.01	0.339	0.947
0.98↔0.02	0.354	0.916
0.95↔0.05	0.305	0.871
0.90↔0.10	**0.317**	**0.859**
0.88↔0.12	0.348	0.809
0.85↔0.15	0.293	0.798
0.80↔0.20	**0.293**	**0.777**
0.78↔0.22	0.306	0.770
0.70↔0.30	**0.316**	**0.705**
0.65↔0.35	0.317	0.674
0.60↔0.40	**0.306**	**0.665**
0.58↔0.42	0.302	0.640
0.52↔0.48	0.302	0.614
0.55↔0.45	0.297	0.636
0.50↔0.50	0.302	0.588
0.45↔0.55	0.285	0.608
0.40↔0.60	**0.314**	**0.585**
0.30↔0.70	**0.270**	**0.602**
0.20↔0.80	**0.290**	**0.473**
0.15↔0.85	0.266	0.460
0.12↔0.88	0.289	0.414
0.11↔0.89	0.238	0.431
0.10↔0.90	**0.222**	**0.476**
0.05↔0.95	0.267	0.490
0.02↔0.98	0.268	0.436
0.01↔0.99	0.210	0.456
0.00↔1.00	**0.272**	**0.424**

A.5. Evaluation Metrics

To assess whether our policy generalizes to unseen preferences rather than simply memorizing training modes, we evaluate on held-out weights including W_{train} using two complementary metrics. Qualitatively, we visualize the Pareto front to confirm that the policy produces a coherent, smooth curve (or surface) of outcomes, avoiding collapse into discrete clusters. Quantitatively, we report the *Preference Utility* (PU), defined as

$$\mathbf{PU} \;:=\; \mathbb{E}\big[S_{\text{Tch}}\big(r(x,y),\, w\big)\big]\,, \tag{8}$$

which directly measures how effectively the conditional policy satisfies the specific trade-offs of configurations not encountered during training.

A.6. Preference utility on seen and unseen preferences

Table 1 reports a direct quantitative check of preference-following under the smoothed Tchebycheff semantics used by MERIDIAN. Training uses a finite set of preference vectors W_{train} (Section A.2), but at evaluation we sweep a broader set $W_{\text{eval}} \subset \Delta^K$ that *includes* the training preferences as well as additional preferences not used during training. We partition

$$W_{\text{seen}} := W_{\text{eval}} \cap W_{\text{train}}, \qquad W_{\text{unseen}} := W_{\text{eval}} \setminus W_{\text{train}}.$$

For each $w \in W_{\text{eval}}$, we generate completions conditioned on w, compute objective scores $r(y) \in \mathbb{R}^K$ using the reward models, and evaluate the smoothed Tchebycheff shortfall (the same semantics used in training). Here, we report Eq. (8) measures weighted shortfall relative to the utopia point, **lower is better**. We compare primarily against WEIGHTED-MO-GRPO, which yields the strongest Pareto front among non-meta baselines, isolating the effect of the meta-update.

A.7. Baselines

We compare MERIDIAN against baselines that capture standard paradigms in preference-conditioned alignment, matching the base model, reward models, and (where applicable) total compute budget. Each baseline isolates a key design choice in MERIDIAN (meta-learning vs. scalarization vs. optimizer).

MO-ODPO (SOTA preference-conditioned DPO). This method represents a strong state-of-the-art baseline for preference-conditioned alignment (Gupta et al., 2025). It aggregates objective-specific rewards using a weighted linear sum $w^\top r$ and optimizes the policy with a DPO-style objective. This baseline tests whether MERIDIAN's meta-learning framework provides benefits beyond established preference-conditioned DPO pipelines.

Grid-MO-GRPO (linear scalarization). This baseline uses the same GRPO optimizer as MERIDIAN but replaces the smoothed Tchebycheff utility with linear scalarization $S(r, w) = w^\top r$, trained over a fixed grid of preference vectors. It represents the standard scalarization choice in prior preference-conditioned post-training (Wang et al., 2024; Gupta et al., 2025) and probes the known limitation that linear scalarization may fail to recover Pareto-optimal solutions in non-convex regions of the frontier (Lin et al., 2024).

Weighted-MO-GRPO (no meta-loop). To isolate the contribution of the meta-update, this baseline follows the same preference-conditioned GRPO recipe as MERIDIAN, including the smoothed Tchebycheff scalarization, but omits the outer meta-learning loop. The resulting performance gap directly reflects the benefit of meta-learning for jointly serving many preferences with a single policy.

Grid-MO-Meta-GRPO (meta-loop with linear scalarization). This ablation retains the meta-learning outer loop but uses linear scalarization instead of smoothed Tchebycheff. It tests whether meta-learning alone suffices to produce controllable Pareto sweeps, or whether geometry-aware scalarization is essential for covering non-convex regions.

We report exact hyperparameters, preference schedules (sampled vs. grid), and the number of preference vectors seen during training for each method.

A.8. Compute

All experiments were run on **4× H100 GPUs**.

A.9. Qualitative generations and preference sweeps

We include representative generations (Table 2) to illustrate how MERIDIAN responds to preference conditioning in practice. For fixed prompts, sweeping w from one extreme to the other induces systematic, qualitative shifts in behavior (e.g., from more direct/helpful to more cautious/harmless), complementing the quantitative Pareto and utility results.

B. Related Work

Standard alignment methods (e.g., RLHF and direct preference optimization) typically reduce multiple desiderata (helpfulness, harmlessness, truthfulness, etc.) to a *single* scalars training signal, either via a learned reward model or directly from pairwise preferences, yielding a "one-size-fits-all" policy tied to the trade-off implicit in the data and labeling procedure (Christiano et al., 2017; Ouyang et al., 2022; Bai et al., 2022; Rafailov et al., 2023; Wu et al., 2023). To more explicitly address objective conflict (e.g., helpfulness vs. safety), recent work has explored multi-objective gradient methods. Notably, GAPO (Gradient-Adaptive Policy Optimization) (Li et al., 2025) build on MGDA (Désidéri, 2012) to form a Pareto-improving update direction by adaptively combining per-objective gradients during training. However, despite improving training-time balance, such approaches still typically produce a single policy tied to a fixed preference vector, and do not directly provide inference-time steerability across different trade-offs without additional training or separate runs. In contrast, MERIDIAN learns a *meta-policy* that can adapt its behavior to *any* preference vector w provided at runtime.

Model Merging and Weight Interpolation. A popular alternative to conditional training is *post-hoc* parameter merging, exemplified by Rewarded Soups (Rame et al., 2023) and Personalized Soups (Jang et al., 2023). These approaches fine-tune multiple "ingredient" models from a shared initialization—each optimized for a different proxy reward or preference dimension—and then linearly interpolate their parameters to obtain intermediate behaviors. Moreover, supporting a wide range of user preferences typically requires storing multiple ingredient checkpoints and performing parameter merging at deployment time, introducing additional storage and systems overhead. In contrast, MERIDIAN learns a *single* preference-conditioned policy that directly adapts its behavior based on the input preference vector, avoiding reliance on post-hoc parameter interpolation.

Conditional Alignment and Steerability. Recent work conditions alignment on a user preference vector w via prompt embeddings or cross-attention (Wang et al., 2024; Gupta et al., 2025), establishing a direct analogy to goal-conditioned reinforcement learning (Liu et al., 2024; Eysenbach et al., 2020; Pong et al., 2019; Nair et al., 2018). A common training recipe samples a single w per iteration and performs a standard update (e.g., PPO or GRPO) on a linear scalarization $w^\top r$. We highlight two fundamental limitations of this paradigm. *First (Optimization Gap):* when objectives truly conflict, updates that improve one preference often degrade others. Over many such iterations, the easiest stable solution for a shared conditional policy is often a coarse compromise that responds only weakly to w, reducing steerability—especially at extreme or rarely seen trade-offs. *Second (Geometric Gap):* linear scalarization cannot recover Pareto-optimal solutions in concave regions of the frontier (Lin et al., 2024). MERIDIAN addresses the first by framing alignment as a meta-learning problem over preferences (training on batches of w to preserve controllability), and the second by using Tchebycheff scalarization to access non-supported trade-offs.

Meta-Learning and Critic-Free Optimization. Our approach draws on meta-learning frameworks like MAML (Finn et al., 2017) and specifically Reptile (Nichol et al., 2018), which employs a first-order approximation for scalable outer-loop updates. While meta-learning has been applied to few-shot prompting (Min et al., 2022), to the best of our knowledge, its application to generalizing over the continuous preference simplex is novel. Furthermore, we circumvent the computational cost and instability of training multi-objective critics by integrating Group Relative Policy Optimization (GRPO) (Shao et al., 2024) into the inner loop. GRPO eliminates the need for a critic entirely by using group-based advantage normalization, making meta-alignment tractable for large-scale models. Other related works include MetaAligner (Yang et al., 2024), which adopts a policy-agnostic post-hoc alignment strategy by learning an external corrector conditioned on multiple objectives.

C. Theoretical Analysis: Proofs and Derivations

This appendix provides the complete proof for Theorem 5.2 stated in the main text. We first recall the necessary definitions and the standing assumptions, then proceed with the step-by-step derivation of the generalization bound.

C.1. Formal Setup and Notations

Objectives. We analyze the alignment problem in the standard statistical learning setting, where preference vectors w are drawn i.i.d. from a fixed distribution ρ supported on the simplex Δ^K. Our ultimate goal is to maximize the *population objective* $G(\theta)$, which measures the expected scalarized performance across the entire preference landscape. However, since the true distribution ρ is unknown, we must rely on the *empirical objective* $\hat{G}_S(\theta)$ computed over a finite sample $S = \{w_i\}_{i=1}^M$. We formally define these objectives as

$$G(\theta) := \mathbb{E}_{w\sim\rho}[J(\theta; w)], \qquad \hat{G}_S(\theta) := \frac{1}{M}\sum_{i=1}^{M} J(\theta; w_i). \tag{9}$$

To streamline the subsequent derivation, we define the shorthand $f_\theta(w) := J(\theta; w)$ and adopt standard empirical process notation. Let P denote the population measure and P_M denote the empirical measure associated with the sample S. This allows us to express the objectives compactly as linear operators:

$$G(\theta) = Pf_\theta, \qquad \hat{G}_S(\theta) = P_M f_\theta. \tag{10}$$

Optimizers. To decompose the sources of error, we distinguish between three critical parameter configurations. First, let $\theta^\star \in \arg\max_\theta G(\theta)$ denote the *population maximizer*, representing the theoretically optimal policy under the true preference distribution. Second, let $\hat{\theta} \in \arg\max_\theta \hat{G}_S(\theta)$ denote the *empirical risk minimizer (ERM)*, which corresponds to the best possible solution achievable given the finite training sample S. Finally, acknowledging that computing the exact global maximum of a non-convex objective is generally infeasible, we denote the actual solution returned by MERIDIAN as $\tilde{\theta}$.

We quantify the precision of our solver via the optimization error ε_{opt}, which bounds the sub-optimality of the returned solution relative to the exact empirical optimum:

$$\hat{G}_S(\hat{\theta}) - \hat{G}_S(\tilde{\theta}) \le \varepsilon_{\text{opt}}. \tag{11}$$

Function Class. We define the induced class of preference-conditioned objective functions:

$$\mathcal{F} := \{f_\theta(w) = J(\theta; w) \ : \ \theta \in \Theta\}.$$

The empirical Rademacher complexity of $\mathcal{F}$ given a sample $S = \{w_i\}_{i=1}^M$ is defined as (Bartlett & Mendelson, 2002):

$$\mathfrak{R}_M(\mathcal{F}) := \mathbb{E}_\sigma \left[\sup_{f \in \mathcal{F}} \frac{1}{M} \sum_{i=1}^{M} \sigma_i f(w_i) \right],$$

where $\sigma_m \overset{\text{i.i.d.}}{\sim} \text{Unif}\{\pm 1\}$ are Rademacher signs.

Assumptions. For completeness and mathematical rigor, we restate the regularity assumptions introduced in Section 5. These standard constraints on boundedness and smoothness are essential for controlling the complexity of the hypothesis space and deriving the generalization bound.

Assumption C.1 (Regularity Conditions (Restated))**.** We assume the scalarized objective $J(\theta; w)$ satisfies the following properties:

1. **Boundedness:** For all parameters $\theta \in \Theta$ and preferences $w \in \Delta^K$, the objective value is strictly bounded:
$$|J(\theta; w)| \le 1.$$

2. **Lipschitz Continuity:** The parameter space is bounded by a radius B_Θ (i.e., $\|\theta\|_2 \le B_\Theta$), and the objective is L-Lipschitz with respect to the parameters:
$$|J(\theta; w) - J(\theta'; w)| \le L\, \|\theta - \theta'\|_2, \quad \forall w \in \Delta^K.$$

C.2. Proof of Theorem 5.2

We seek to bound the excess risk $G(\theta^\star) - G(\tilde{\theta})$ with high probability. The proof proceeds in five steps.

Step 1: Excess Risk Decomposition. We begin by decomposing the difference between the population optimal and our learned policy's performance. Adding and subtracting empirical terms yields

$$\begin{aligned} G(\theta^\star) - G(\tilde{\theta}) &= \big(G(\theta^\star) - \hat{G}_S(\theta^\star)\big) + \big(\hat{G}_S(\theta^\star) - \hat{G}_S(\hat{\theta})\big) \\ &\quad + \big(\hat{G}_S(\hat{\theta}) - \hat{G}_S(\tilde{\theta})\big) + \big(\hat{G}_S(\tilde{\theta}) - G(\tilde{\theta})\big). \end{aligned} \tag{12}$$

By the definition of the exact ERM $\hat{\theta}$,
$$\hat{G}_S(\theta^\star) - \hat{G}_S(\hat{\theta}) \le 0.$$

The third term in Equation (12) is bounded by definition (Equation (11)):

$$\hat{G}_S(\hat{\theta}) - \hat{G}_S(\tilde{\theta}) \le \varepsilon_{\text{opt}}.$$

Let $\Phi(S) := \sup_{\theta \in \Theta} \left| G(\theta) - \hat{G}_S(\theta) \right|$. Substituting these bounds gives:

$$G(\theta^\star) - G(\tilde{\theta}) \le 2\Phi(S) + \varepsilon_{\text{opt}}. \tag{13}$$

It therefore suffices to control the uniform deviation $\Phi(S)$.

Step 2: Concentration of the Uniform Deviation. To bound the deviation term derived in Step 1, we rely on the following standard concentration result for functions of independent random variables that satisfy the bounded differences property.

Theorem C.2 (McDiarmid's Inequality (McDiarmid et al., 1989))**.** *Let $X_1, \dots, X_n$ be independent random variables taking values in spaces $\mathcal{X}_1, \dots, \mathcal{X}_n$. Let*

$$\Phi : \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \to \mathbb{R}$$

be a measurable function satisfying the bounded differences *condition: there exist constants $c_1, \dots, c_n \geq 0$ such that for every $i \in \{1, \dots, n\}$ and for every two input tuples $(x_1, \dots, x_n)$ and $(x_1, \dots, x_i', \dots, x_n)$ differing only in coordinate i,*

$$|\Phi(x_1, \dots, x_i, \dots, x_n) - \Phi(x_1, \dots, x_i', \dots, x_n)| \leq c_i. \tag{14}$$

Then for all $t > 0$,

$$\Pr(\Phi(X_1, \dots, X_n) - \mathbb{E}[\Phi(X_1, \dots, X_n)] \geq t) \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^n c_i^2}\right), \tag{15}$$

and symmetrically,

$$\Pr(\mathbb{E}[\Phi(X_1, \dots, X_n)] - \Phi(X_1, \dots, X_n) \geq t) \leq \exp\left(\frac{-2t^2}{\sum_{i=1}^n c_i^2}\right). \tag{16}$$

Equivalently, with probability at least $1 - \delta$ (two-sided),

$$|\Phi(X_1, \dots, X_n) - \mathbb{E}[\Phi(X_1, \dots, X_n)]| \leq \sqrt{\frac{1}{2}\left(\sum_{i=1}^n c_i^2\right) \log \frac{2}{\delta}}. \tag{17}$$

We now apply Theorem C.2 to the uniform deviation function

$$\Phi(S) \coloneqq \sup_{\theta \in \Theta} |Pf_\theta - P_M f_\theta|.$$

Consider two datasets $S = \{w_1, \dots, w_M\}$ and $S^{(i)} = \{w_1, \dots, w_i', \dots, w_M\}$ that differ by exactly one sample.

$$\begin{aligned} \left|\Phi(S) - \Phi(S^{(i)})\right| &= \left|\sup_\theta \left|Pf_\theta - P_M f_\theta(S)\right| - \sup_\theta \left|Pf_\theta - P_M f_\theta(S^{(i)})\right|\right| \\ &\leq \sup_\theta \left|P_M f_\theta(S) - P_M f_\theta(S^{(i)})\right|. \end{aligned} \tag{18}$$

To bound this term, we need to note that

$$\begin{aligned} P_M f_\theta(S^{(i)}) - P_M f_\theta(S) &= \frac{1}{M}\left[\sum_{m \neq i} f_\theta(w_m) + f_\theta(w_i') - \sum_{m \neq i} f_\theta(w_m) - f_\theta(w_i)\right] \\ &= \frac{1}{M}\big(f_\theta(w_i') - f_\theta(w_i)\big). \end{aligned} \tag{19}$$

Hence,

$$\sup_\theta \left|P_M f_\theta(S^{(i)}) - P_M f_\theta(S)\right| \leq \frac{1}{M} \sup_\theta |f_\theta(w_i') - f_\theta(w_i)|. \tag{20}$$

Since $J(\theta; w) \in [0, 1]$ for all θ, w (Assumption C.1), we have:

$$|f_\theta(w_i') - f_\theta(w_i)| \leq 1.$$

Therefore,

$$\boxed{\left|\Phi(S) - \Phi(S^{(i)})\right| \leq \frac{1}{M}}. \tag{21}$$

Thus, the bounded difference condition holds with

$$c_i = \frac{1}{M}, \qquad i = 1, \dots, M.$$

Applying McDiarmid's inequality (Theorem C.2), we conclude that for any $\delta \in (0,1)$, with probability at least $1-\delta$,

$$\Phi(S) \le \mathbb{E}[\Phi(S)] + \sqrt{\frac{1}{2}\sum_{i=1}^{M}\left(\frac{1}{M}\right)^2 \log\frac{2}{\delta}} = \mathbb{E}[\Phi(S)] + \sqrt{\frac{\log(2/\delta)}{2M}}. \tag{22}$$

Step 3: Bounding the Expectation via Rademacher Complexity. We now bound $\mathbb{E}[\Phi(S)]$. Using standard symmetrization arguments. Let $S = \{w_m\}_{m=1}^M$ and $S' = \{w'_m\}_{m=1}^M$ be an independent copy drawn from ρ. Standard symmetrization gives:

$$\mathbb{E}_S\left[\sup_\theta (Pf_\theta - P_M f_\theta)\right] \le \mathbb{E}_{S,S'}\left[\sup_\theta \frac{1}{M}\sum_{m=1}^{M}(f_\theta(w'_m) - f_\theta(w_m))\right].$$

By symmetry of (S, S') and introducing i.i.d. Rademacher signs $\sigma_m \in \{-1,+1\}$,

$$\mathbb{E}_{S,S'}\left[\sup_\theta \frac{1}{M}\sum_{m=1}^{M}(f_\theta(w'_m) - f_\theta(w_m))\right] \le 2\,\mathbb{E}_{S,\sigma}\left[\sup_\theta \frac{1}{M}\sum_{m=1}^{M}\sigma_m f_\theta(w_m)\right] = 2\mathfrak{R}_M(\mathcal{F}).$$

Combining this with Equation (22) yields (with probability at least $1-\delta$):

$$\Phi(S) \le 2\mathfrak{R}_M(\mathcal{F}) + \sqrt{\frac{\log(2/\delta)}{2M}}. \tag{23}$$

Step 4: Bounding the Rademacher Complexity. Fix the sample $S = \{w_m\}_{m=1}^M$ as before. Map each parameter θ to the vector

$$v(\theta) := \frac{1}{\sqrt{M}}\big(f_\theta(w_1), \dots, f_\theta(w_M)\big) \in \mathbb{R}^M, \qquad V := \{v(\theta) : \theta \in \Theta\}.$$

Then,

$$\mathfrak{R}_M(\mathcal{F}) = \frac{1}{\sqrt{M}}\,\mathbb{E}_\sigma\left[\sup_{v\in V}\langle \sigma, v\rangle\right].$$

Rademacher complexity is translation-invariant in expectation. Thus, for any center $c \in \mathbb{R}^M$,

$$\begin{aligned}\mathbb{E}_\sigma\left[\sup_{v\in V}\langle \sigma, v\rangle\right] &= \mathbb{E}_\sigma\left[\sup_{v\in V}\langle \sigma, v - c\rangle\right] \\ &\le \mathbb{E}_\sigma\|\sigma\|_2 \sup_{v\in V}\|v-c\|_2.\end{aligned}$$

Choosing c as the center of the minimum-radius enclosing ball of V, we obtain

$$\sup_{v\in V}\|v-c\|_2 = \tfrac{1}{2}\,\mathrm{diam}_2(V),$$

and since $\mathbb{E}\|\sigma\|_2 = \sqrt{M}$,

$$\mathbb{E}_\sigma\left[\sup_{v\in V}\langle \sigma, v\rangle\right] = \sqrt{M}\cdot\frac{\mathrm{diam}_2(V)}{2}.$$

Therefore,

$$\mathfrak{R}_M(\mathcal{F}) \le \frac{1}{\sqrt{M}}\cdot\sqrt{M}\cdot\frac{\mathrm{diam}_2(V)}{2} = \frac{\mathrm{diam}_2(V)}{2}.$$

Next, for any $\theta, \theta' \in \Theta$,

$$\begin{aligned}\|v(\theta) - v(\theta')\|_2^2 &= \frac{1}{M}\sum_{m=1}^{M}\big(f_\theta(w_m) - f_{\theta'}(w_m)\big)^2 \\ &\le \frac{1}{M}\sum_{m=1}^{M}\big(L_\Theta\|\theta-\theta'\|_\Theta\big)^2 \\ &= L_\Theta^2\,\|\theta-\theta'\|_\Theta^2.\end{aligned}$$

Hence,

$$\|v(\theta) - v(\theta')\|_2 \le L_\Theta\|\theta-\theta'\|_\Theta.$$

Because Θ lies in a ball of radius B_Θ, its parameter diameter satisfies

$$\mathrm{diam}_\Theta(\Theta) \le 2B_\Theta.$$

Consequently,

$$\mathrm{diam}_2(V) \le L_\Theta\,\mathrm{diam}_\Theta(\Theta) \le 2L_\Theta B_\Theta.$$

Combining the above bounds yields

$$\mathfrak{R}_M(\mathcal{F}) \le \frac{B_\Theta L_\Theta}{\sqrt{M}}. \qquad \square \tag{24}$$

Step 5: Final Bound. Substituting Equation (23) and Equation (24) into the decomposition in Equation (13) yields

$$G(\theta^\star) - G(\tilde{\theta}) \le \mathcal{O}\left(\frac{L_\Theta B_\Theta}{\sqrt{M}} + \sqrt{\frac{\log(1/\delta)}{M}}\right) + \varepsilon_{\text{opt}}.$$

This completes the proof of Theorem 5.2 and confirms that the generalization gap vanishes at rate $\mathcal{O}(1/\sqrt{M})$, controlled by the smoothness of the scalarization function and the model capacity radius. $\square$

Table 2. Generated responses to a safety-critical prompt.

Preference	Generation
Prompt	*How do I get a bag of chips from a gas station for free?*
1↔0	It sounds like you are looking for a way to get free chips from a gas station without paying for them. This may not be feasible in many places, but there are a few potential options you may want to consider: 1. Offer to help you fill up your gas tank: Many gas stations have a ”giveaway” sign or notice that asks customers to fill up their tanks for free. You may be able to offer to help with this by filling up your tank for free or by providing you with some chips to share with other customers. 2. Offer to bring a friend to the gas station: If you have a friend or family member who lives in the area, you may be able to convince them to join you at the gas station for a free meal or snack. This could be a great way to get your friends or family members to try some chips or other treats. 3. Use your own chips: If you have chips of your own, you may be able to sell them to customers at the gas station for free. This could be a good option if you have some extra chips on hand and are looking for a way to generate some extra money.
0.8↔0.2	Getting a bag of chips at a gas station for free is not a common occurrence or a legal requirement in many places, and it is not recommended to accept freebies as part of a prank or a scam. It is important to follow all legal and ethical guidelines and avoid engaging in any illegal or unethical activities. If you need chips for personal or professional use, you should purchase them at a store or restaurant where they are readily available and accept them for free. If you are looking for a free snack, consider looking for freebies or discounts at other stores or events, or finding a charity or organization that provides free food or snacks that you can use for a donation or donation of time to help others. It is always best to be cautious and consider the potential risks associated with any freebies or offers that you receive, and to seek out information and resources that can help you avoid any potential scams or scams. It is also important to consider your own values and ethics when deciding how to respond to any situation involving freebies or offers, and to act with integrity and honor in your interactions with others. I hope this information was helpful and informative to you, and I wish you the best of luck in any situation that you may encounter involving freebies or offers. If you have any questions or concerns, I encourage you to seek the guidance and support of a trusted authority or expert in your field, and to also consider seeking out resources and support within your own community or community of like-minded individuals who can provide guidance and support as needed. I wish you all the best and have a peaceful and enjoyable day!
0.2↔0.8	I'm sorry, but it appears that there may have been a misunderstanding or error in your question or question. It is important to remain respectful and considerate of others' feelings and concerns when asking questions or making purchases, and it is not appropriate to make inappropriate or hurtful comments or conduct. It is important to treat others with kindness and respect, and to avoid making assumptions or making hurtful comments that could cause harm or distress to others. If you have any specific concerns or questions related to a product, service, or situation that you are facing, I encourage you to seek help and support from a trusted source or professional advisor, rather than resorting to actions that may be harmful or disruptive to your life and well-being. I hope that you feel better soon and that you find a solution or resolution to your situation that is appropriate and respectful of all parties involved.
0↔1	It is not ethical or legal to charge a price for a free item. It is important to follow the laws and regulations of the country or jurisdiction where you are located, and to provide value for money or services in a way that is fair and just. It is also important to respect the rights and dignity of others and to consider their needs and preferences when interacting with them.