

Rethinking Diverse Human Preference Learning through Principal Component Analysis

Feng Luo^{1*}, Rui Yang^{2*}, Hao Sun³, Chunyuan Deng¹, Jiarui Yao², Jingyan Shen⁴,
Huan Zhang², Hanjie Chen¹

¹Rice University, ²University of Illinois at Urbana-Champaign,

³University of Cambridge, ⁴ Columbia University

f38@rice.edu, ry21@illinois.edu, hanjie@rice.edu

Abstract

Understanding human preferences is crucial for improving foundation models and building personalized AI systems. However, preferences are inherently diverse and complex, making it difficult for traditional reward models to capture their full range. While fine-grained preference data can help, collecting it is expensive and hard to scale. In this paper, we introduce Decomposed Reward Models (DRMs), a novel approach that extracts diverse human preferences from binary comparisons without requiring fine-grained annotations. Our key insight is to represent human preferences as vectors and analyze them using Principal Component Analysis (PCA). By constructing a dataset of embedding differences between preferred and rejected responses, DRMs identify orthogonal basis vectors that capture distinct aspects of preference. These decomposed rewards can be flexibly combined to align with different user needs, offering an interpretable and scalable alternative to traditional reward models. We demonstrate that DRMs effectively extract meaningful preference dimensions (e.g., helpfulness, safety, humor) and adapt to new users without additional training. Our results highlight DRMs as a powerful framework for personalized and interpretable LLM alignment. Our code is available at <https://github.com/amandaluof/DRMs>

1 Introduction

Reinforcement Learning from Human Feedback (RLHF) (Stiennon et al., 2020; Ouyang et al., 2022; Bai et al., 2022; Team et al., 2023; Grattafiori et al., 2024; Liu et al., 2024a) has proven to be a powerful approach for fine-tuning large language models (LLMs) to serve as better general assistants (Team et al., 2023; Achiam et al., 2023), AI agents (Nakano et al., 2021; Zhao et al., 2024; Yang et al., 2025), and clinical decision support systems

(Tian et al., 2023; Wang et al., 2024a; Chen et al., 2025). Typically, LLMs are optimized using a scalar reward model trained on human preference data, acting as a proxy for overall user satisfaction. However, this approach has two key limitations: (1) it often reflects the preferences of the majority, potentially marginalizing underrepresented groups (Chakraborty et al., 2024; Chidambaram et al., 2024) and failing to capture the full diversity of human preferences, and (2) it struggles to represent the complex, multifaceted, and sometimes conflicting nature of human preferences with a single scalar reward (Jang et al., 2023; Rame et al., 2024; Yang et al., 2024c; Zhou et al., 2024).

As demand grows for more personalized LLMs (Zhang et al., 2024b), researchers have explored ways to capture fine-grained, multidimensional human preferences. Some studies have introduced datasets that evaluate multiple aspects such as relevance, correctness, completeness, helpfulness, and harmlessness (Wu et al., 2023; Wang et al., 2023; Cui et al., 2023; Pitis et al., 2024). Building on these datasets, others have proposed multi-objective optimization methods (Qiu et al., 2024) to accommodate diverse user needs (Yang et al., 2024c,a; Wang et al., 2024c). However, these approaches face significant challenges: collecting fine-grained human annotations is expensive, and using GPT-generated labels can introduce biases. Considering that large-scale binary preference datasets—where users simply compare two responses—are easier to collect and more widely available, we ask the question: *Can we infer multidimensional human preferences directly from large-scale binary comparisons?*

To address this question, we propose Decomposed Reward Models (DRMs), a framework that extracts fine-grained human preferences using binary comparison data. Unlike traditional probabilistic models such as Bradley-Terry (BT) (Bradley and Terry, 1952) or pairwise preference

*Equal contribution.

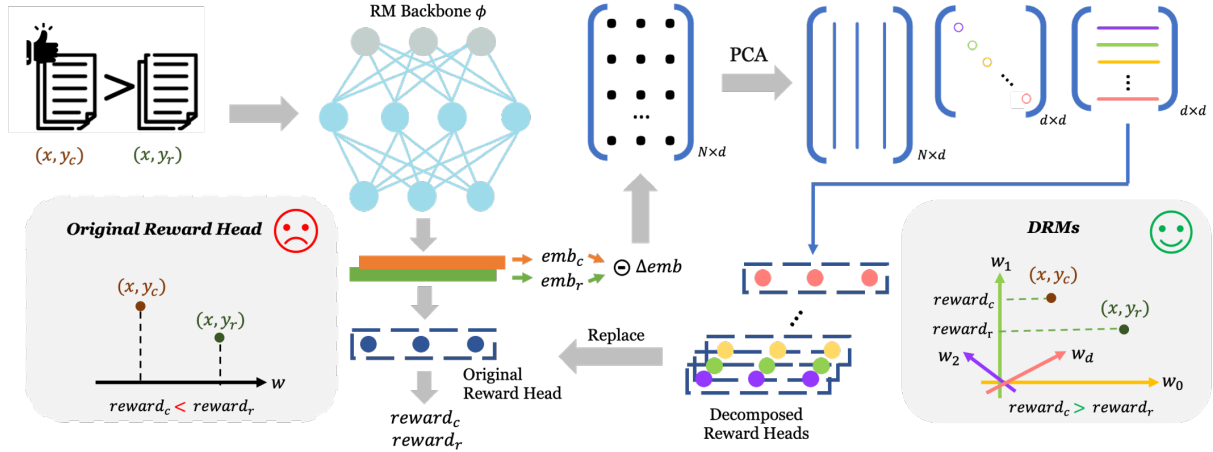


Figure 1: Illustration of the decomposition pipeline in DRMs. In the original single-dimensional head, a prompt–response pair can be predicted incorrectly. In contrast, DRMs capture preferences along multiple dimensions, aligning more effectively with the complex and multifaceted nature of human preferences.

models (Jiang et al., 2023), our approach represents human preferences as d -dimensional vectors, corresponding to the learned weights of the final linear layer in a reward model. We show that this vector-based representation establishes a connection between Principal Component Analysis (PCA) and preference learning. Our method constructs a dataset of embeddings for all prompt–response pairs and applies PCA to the differences between chosen and rejected response embeddings. This process identifies orthogonal basis vectors, each capturing a distinct human preference direction. These basis vectors can be combined with their corresponding feature extractor to construct reward models. The framework is shown in Figure 1. DRMs are particularly effective for personalized preference learning. Given a small adaptation dataset from a new user, we compute coefficients for the basis vectors and form a linear combination that best aligns with the user’s preferences. This approach allows for flexible, scalable, and interpretable preference modeling without requiring additional training.

In our experiments, we first show that the learned DRM basis vectors capture diverse human preference attributes. For instance, one vector may strongly correlate with safety, while another may reflect humor. Next, we evaluate DRM’s effectiveness in adapting to user preferences at test time. Our results show that DRMs outperform both the single-head reward model and previous test-time alignment methods based on reward ensembles. These findings emphasize DRM’s advantages in modeling and adapting to diverse human prefer-

ences. In addition, we also provide an explainability study to better understand human preferences.

In summary, our work makes three key contributions. (1) We propose a vector-based representation of human preferences, providing a more structured and interpretable approach to preference learning. (2) We establish a connection between preference learning and PCA, enabling the extraction of diverse and meaningful reward components without additional training. (3) We empirically validate DRMs, demonstrating their effectiveness in capturing diverse human preferences and adapting to new users at test time. These findings highlight DRMs as a scalable and flexible alternative to traditional reward models, with promising applications in interpretable AI and LLM personalization.

2 Preliminary

The Bradley-Terry (BT) model (Bradley and Terry, 1952) is a probabilistic framework commonly used in preference learning. Given a prompt x and two responses, y_1 and y_2 , the BT model defines the probability of y_1 being preferred over y_2 as:

$$P(y_1 \succ y_2 | x) = \frac{\exp(r(x, y_1))}{\exp(r(x, y_1)) + \exp(r(x, y_2))}$$

where $r(x, y)$ represents the reward model, which assigns a score to a given prompt–response pair. Given a dataset of comparisons $\mathcal{D} = \{(x_i, y_i^c, y_i^r)\}_{i=1}^N$, where y_i^c is the preferred response and y_i^r is the rejected one, the reward model is trained by maximizing the likelihood of human

preferences. This results in the following objective:

$$\max_{\theta} \mathbb{E}_{(x_i, y_i^c, y_i^r) \sim \mathcal{D}} [\log \sigma (r_{\theta}(x_i, y_i^c) - r_{\theta}(x_i, y_i^r))] \quad (1)$$

where $r_{\theta}(x, y)$ is the reward score parameterized by θ , and $\sigma(\cdot)$ denotes the sigmoid function. By minimizing this loss, the reward model learns to assign higher scores to human-preferred responses, making it a useful proxy for human preferences.

3 Methodology

The standard approach to preference learning relies on a scalar-valued reward model, which may not adequately capture the full complexity of human preferences. To address this limitation, we first introduce a vector representation of human preference and establish its connection to PCA. Building on this, we propose a PCA-based method that decomposes human preference into multiple basis vectors, allowing any preference to be represented as a linear combination of these vectors. This approach enables a novel paradigm for diverse reward modeling without additional training.

3.1 Vector Representation of Preferences

In conventional reward modeling approaches, human preferences are transferred into *scalar scores* using the BT models, as illustrated in Eq. (1). In practice, modern reward models are typically fine-tuned from pretrained or instruction-fine-tuned language models, with a reward head that maps hidden states to a scalar reward prediction.

In this paper, we consider a **vector representation** of human preference, similar to recent theoretical work (Sun et al., 2024; Zhang et al., 2024a; Shen et al., 2025) and empirical studies (Sun et al., 2023, 2025). We denote $\phi(x, y) \in \mathbb{R}^d$ as a d -dimensional feature extractor (e.g., the output from the penultimate layer), and apply a final linear layer w in the reward model, where the reward is computed as $r_{\theta}(x, y) = w^{\top} \phi(x, y)$. Under this formulation, the BT objective can be rewritten as:

$$\begin{aligned} \max_w \mathbb{E}_i [\log \sigma (w^{\top} \phi(x_i, y_i^c) - w^{\top} \phi(x_i, y_i^r))] \\ = \max_w \mathbb{E}_i [\log \sigma (w^{\top} (\phi(x_i, y_i^c) - \phi(x_i, y_i^r)))] \end{aligned}$$

where $(x_i, y_i^c), (x_i, y_i^r)$ denote the chosen and rejected samples, respectively.

This reformulation is particularly interesting because it allows human preference to be captured using a vector $w \in \mathbb{R}^d$ instead of relying on a

large set of model parameters. Consequently, human preference can be interpreted as a direction in the d -dimensional space. When the feature space difference $\phi(x_i, y_i^c) - \phi(x_i, y_i^r)$ aligns with such a direction, it means the response pair is aligned with human preference (i.e., $y_i^c \succ y_i^r$); otherwise, it indicates a contradiction. Since the distances between preference vectors provide a natural way of comparing preferences, *such a vector-based representation lays a foundation for a more interpretable framework for understanding human preference.*

3.2 Rethinking Preference Learning with Principal Component Analysis (PCA)

The next step is selecting a *good basis* to represent any vector preferences. The basis should be universal, interpretable, and have a strong connection with human preferences. Our choice of the basis is motivated by the following observations.

Formally, we start by defining $z_i = \phi(x_i, y_i^c) - \phi(x_i, y_i^r)$ and letting it be $\{z_i\}_{i=1}^N$ zero-centered, (i.e., $\mathbb{E}_i[z_i] = 0$, which can be ensured via normalization). Without loss of generality, we consider the unified preference vector $\|w\|_2^2 = 1$. Then the objective of reward modeling becomes

$$\max_w \mathbb{E}_i [\log \sigma (w^{\top} z_i)], \quad \text{s.t.} \quad \|w\|_2^2 = 1. \quad (2)$$

Here, w is optimized to find a direction that best distinguishes the preference data, which bears a resemblance to PCA: both methods seek a meaningful projection of the data onto a dimension. To further explore this connection, consider the following regularized objective ($\lambda > 0$):

$$J(w) = \mathbb{E}_i [\log \sigma (w^{\top} z_i)] - \lambda \|w\|_2^2. \quad (3)$$

Taking the gradient, we obtain:

$$\nabla_w J(w) = \mathbb{E}_i [(1 - \sigma(w^{\top} z_i)) z_i] - 2\lambda w. \quad (4)$$

For small $w^{\top} z_i$, we have $1 - \sigma(x) \approx \frac{1}{2} + cx$ according to Taylor expansion, where c is a constant. Therefore, we have

$$\begin{aligned} \nabla_w J(w) &\approx c \mathbb{E}_i [(z_i z_i^{\top}) w] - 2\lambda w \\ &= c \Sigma w - 2\lambda w, \end{aligned} \quad (5)$$

where $\Sigma = \mathbb{E}_i [z_i z_i^{\top}]$ is the covariance matrix of $\{z_i\}$. Setting the gradient to zero gives:

$$c \Sigma w = 2\lambda w \quad \Rightarrow \quad \Sigma w = \frac{2\lambda}{c} w. \quad (6)$$

Discussion. This equation suggests that under certain conditions, the learned preference direction w aligns with an eigenvector of the covariance matrix Σ . While this does not imply a direct equivalence between preference learning and PCA, it highlights an interesting connection: both methods extract a principal direction from the data. Unlike PCA, which maximizes variance in an unsupervised manner, preference learning optimizes a supervised ranking objective, making the relationship approximate rather than exact.

Additionally, while PCA eigenvectors are direction-agnostic—meaning their signs are arbitrary—human preference is inherently directional. As a result, when deriving a preference vector w from PCA, both w and $-w$ need to be considered to ensure an accurate representation.

3.3 From Scalar Reward to Diverse Rewards

Our analysis above suggests a connection between the eigenvectors of the covariance matrix and human preferences. A key observation is that the covariance matrix has d eigenvectors—for example, $d = 2048$ for gemma-2B (Team et al., 2024) and $d = 4096$ for Llama-3.1-8B (Grattafiori et al., 2024). This means we can extract a large number of meaningful preference vectors w from PCA applied to the embedding dataset $\{z_i\}_{i=1}^N$.

These eigenvectors form an orthonormal basis in the d -dimensional space, meaning they are mutually orthogonal and span the entire space. Mathematically, the eigendecomposition of the covariance matrix is given by: $\Sigma = W\Lambda W^T$ where $W = [w_1, w_2, \dots, w_d]$ is an orthonormal matrix whose columns are the eigenvectors of Σ , $\Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_d)$ is a diagonal matrix of eigenvalues. This ensures that w_1, w_2, \dots, w_d represent diverse preference directions. Any human preference can then be expressed as a linear combination of these basis vectors:

$$w' = \sum_{j=1}^d k_j w_j,$$

where k_1, \dots, k_d are weight parameters. This formulation enables a flexible and expressive human preference representation as a combination of these decomposed rewards.

3.4 Decomposed Reward Models (DRMs)

As illustrated in Figure 1, we propose Decomposed Reward Models (DRMs) by applying PCA to a

human preference dataset $\mathcal{D} = \{(x_i, y_i^c, y_i^r)\}_{i=1}^N$ using an embedding extractor $\phi(x, y)$. ϕ can be any language models (pretrained or instruction-tuned) or reward models that produce hidden states of dimension d . DRMs consist of two main steps: (1) *Embedding Extraction*: We run inference over the preference dataset \mathcal{D} using ϕ to obtain a dataset of embedding differences:

$$\mathcal{D}_e = \{z_i\}_{i=1}^N, \quad \text{where } z_i = \phi(x_i, y_i^c) - \phi(x_i, y_i^r).$$

(2) *Principal Decomposition*: Given the dataset \mathcal{D}_e of shape $N \times d$, we perform PCA to obtain a set of basis vectors $W = [w_1, w_2, \dots, w_d]$. Each eigenvector w_j captures a distinct directional human preference and can be combined with the feature extractor $\phi(x, y)$ to construct a reward model, leading to a total of d rewards.

To fully utilize DRMs, we aim to represent any human preference as a linear combination of these decomposed rewards. Various methods can achieve this, such as mixture-of-experts (Quan, 2024; Wang et al., 2024b) and test-time alignment (Lee et al., 2024). Since both approaches are fundamentally similar, and test-time alignment can better leverage our method without additional fine-tuning, we adopt it as our implementation. However, *DRMs are compatible with any downstream method that requires multiple diverse reward heads*.

We implement test-time alignment based on HyRe (Lee et al., 2024), which adapts to diverse human preferences by optimizing reward weights given a small adaptation set $\mathcal{D}_{\text{adapt}}$. This is practical in real-world scenarios where new users’ data helps refine the strategy to better align with individual preferences. The weights k_m are computed based on the loss function over the adaptation set:

$$k_m = \frac{\exp(-\mathcal{L}(w_m, \mathcal{D}_{\text{adapt}}))}{\sum_j \exp(-\mathcal{L}(w_j, \mathcal{D}_{\text{adapt}}))} \quad (7)$$

where

$$\mathcal{L}(w_m, \mathcal{D}_{\text{adapt}}) = -\mathbb{E}_{(x_i, y_i^c, y_i^r) \sim \mathcal{D}_{\text{adapt}}} [\log \sigma(w_m^T (\phi(x_i, y_i^c) - \phi(x_i, y_i^r)))] \quad (8)$$

Here, we also normalize the features based on $\mathcal{D}_{\text{adapt}}$ to prevent the loss \mathcal{L} from being dominated by samples with large-scale features. This formulation ensures that weights k_m dynamically adjust based on the dataset, assigning higher importance to directions w_m with lower loss and reducing the influence of those with higher loss. As a result, we can obtain a soft, weighted sum of DRMs that better aligns with the desired human preference.

Advantages of DRMs. (1) *Simplicity*: Our proposed DRMs offer a simple yet effective approach to reward modeling without requiring additional training. (2) *Diversity*: Unlike traditional scalar reward models that struggle to represent heterogeneous human preferences, DRMs leverage a diverse set of basis vectors to capture a wide range of preferences. (3) *Adaptivity*: By decomposing human preference data through PCA, DRMs naturally extract meaningful directional preferences, allowing them to adapt flexibly to different user needs. This adaptability can be leveraged by test-time alignment, which dynamically adjusts reward weights to better suit specific preferences. As a result, DRMs provide a scalable and interpretable solution for modeling complex and diverse human preferences.

4 Experiment

In this section, we conduct extensive experiments to evaluate the effectiveness of DRMs, focusing on the *diversity and interpretability* of the decomposed heads, as well as their *adaptivity* to downstream human preferences.

4.1 Experimental Setup

Dataset. We choose the *mixture2_and_safe_pku* dataset¹ (Dong et al., 2024), a collection of 550k pairwise preference samples. This dataset combines diverse data sources, including human-labeled preferences from HH-RLHF (Bai et al., 2022) and GPT-labeled preferences from Ultra-Feedback (Cui et al., 2023), making it well-suited for studying diverse preference decomposition. To evaluate the effectiveness of the decomposed reward heads, we test models on two unseen benchmarks with multiple attributes: (1) RewardBench (Lambert et al., 2024), a dataset designed to evaluate reward models across various dimensions, including chat quality, safety, and reasoning. (2) Reasonable Preference Reversal (RPR) test set (Pitis et al., 2024) focuses on personalized context-aware preference evaluation. From RPR, we sample five fine-grained categories (i.e., User-Friendliness, Narrative Quality, Linguistic Creativity, Scientific Rigor, and Humor), each with over 80 annotated samples. Other categories were excluded due to insufficient data for reliable evaluation.

Base Model. For experiments on decomposed reward heads and test-time adaptation, we use two

open-source reward models, *Gemma-2B-RM* and *Llama3-8B-RM* (Yang et al., 2024b), along with an instruction-tuned language model, *gemma-2-9b-it* (Team et al., 2024), as our base models. To analyze their performance, we keep the backbone fixed as our feature extractors while generating multiple new reward heads for them.

Baselines. We compare the performance of DRMs against several baselines. (1) *Single-Head RM* fine-tunes a single reward head on top of a fixed backbone using the same dataset as DRMs. (2) *Share-Base RM* (Lee et al., 2024) is a training-based method that employs multiple learnable reward heads while incorporating a frozen prior network (Osband et al., 2023) to maintain diversity, with the final output derived from their combination. (3) *Random Head* uses multiple reward heads with randomly initialized weights to capture diverse but largely unstructured preferences. Specifically, we experiment with both uniform and Gaussian initialization, followed by L2 normalization to ensure consistency with DRM vectors.

4.2 What information is Captured by DRMs?

We aim to better understand the decomposed reward heads in DRMs. To achieve this, we evaluate the performance of the top 100 reward vectors, ranked by eigenvalue, on both RewardBench and RPR’s fine-grained subsets. Table 1 reports the scores of a trained single-head baseline. The “Max Value” column shows the highest score achieved for each attribute, while the “Max Head” column indicates which reward head achieves this score. We also compare the results with the single-head baseline. The results reveal the following findings: (1) **Diversity and Interpretability**: DRMs effectively capture diverse human preferences, with different reward heads excelling at different attributes. For instance, in *Gemma-2B-RM*, head_9 performs best on “User-Friendliness” (accuracy: 0.798) and “Humor and Entertainment” (0.964), while head_12 excels in “Narrative and Storytelling” (0.825) and “Linguistic Creativity” (0.885). In contrast, the single-head RM fails to capture this diversity, yielding suboptimal performance on attributes such as “User-Friendliness” (0.506) and “Humor and Entertainment” (0.690) for *Gemma-2B-RM*. These results indicate that DRMs not only capture a broader range of human preferences but also provide interpretable representations that align well with certain preference attributes.

(2) **The first head is the most informative**: An

¹https://huggingface.co/datasets/weqweasdas/preference_dataset_mixture2_and_safe_pku

Benchmark	Attribute	Gemma-2B-RM			Llama3-8B-RM		
		Single Head	Max Value	Max Head	Single Head	Max Value	Max Head
RewardBench	Overall	0.733	0.735	head_0	0.862	0.869	head_0
	Chat	0.944	0.950	head_0	0.983	0.986	head_0
	Chat Hard	0.467	0.660	head_3	0.684	0.695	head_3
	Safety	0.759	0.745	head_0, head_8	0.868	0.886	head_0
	Reasoning	0.759	0.821	head_32	0.912	0.923	head_0
RPR	Overall	0.714	0.735	head_0	0.853	0.839	head_0
	User-Friendliness	0.506	0.798	head_9, head_26	0.719	0.899	head_10
	Narrative & Storytelling	0.662	0.825	head_12	0.838	0.912	head_5
	Linguistic Creativity	0.817	0.885	head_12	0.875	0.981	head_37
	Scientific Rigor	0.881	0.881	head_34	0.940	0.964	head_0
	Humor & Entertainment	0.690	0.964	head_9	0.893	0.952	head_37, head_74

Table 1: Performance of top 100 decomposed reward heads. “Single Head” is the trained single-head baseline. “Max Value” refers to the highest score achieved for each attribute, while “Max Head” indicates which specific head attains this maximum score. “Overall” represents the average accuracy of a single head across all attributes.

Base Model	Method	User Friendliness	Narrative & Storytelling	Linguistic Creativity	Scientific Rigor	Humor & Entertainment	Overall
Gemma-2B-RM	Single Head	0.506	0.662	0.817	0.881	0.690	0.714
	Shared-Base	0.517(0.000)	0.688(0.000)	0.817(0.000)	0.881(0.000)	0.690(0.000)	0.721(0.000)
	Random (Uniform)	0.713(0.062)	0.782(0.068)	0.920(0.045)	0.907(0.043)	0.907(0.026)	0.848(0.024)
	Random (Gaussian)	0.582(0.039)	0.760(0.060)	0.823(0.063)	0.873(0.025)	0.817(0.053)	0.771(0.022)
	DRM (Ours)	0.789(0.062)	0.871(0.033)	0.953(0.034)	0.907(0.019)	0.975(0.017)	0.900(0.017)
Llama3-8B-RM	Single Head	0.685	0.825	0.846	0.964	0.905	0.844
	Shared-Base	0.674(0.000)	0.825(0.000)	0.827(0.000)	0.964(0.000)	0.881(0.000)	0.832(0.000)
	Random (Uniform)	0.616(0.104)	0.860(0.037)	0.798(0.103)	0.958(0.007)	0.906(0.031)	0.823(0.041)
	Random (Gaussian)	0.730(0.100)	0.892(0.027)	0.887(0.055)	0.956(0.008)	0.919(0.032)	0.875(0.028)
	DRM (Ours)	0.812(0.063)	0.946(0.029)	0.945(0.015)	0.969(0.010)	0.991(0.011)	0.931(0.016)

Table 2: Evaluation Results on RPR ($n = 5$). We compare DRMs with trained baselines (“Single Head” and “Shared-Base”), and randomly generated multi-head baselines (“Random”). Except for single-head baseline, other methods use HyRe for test-time adaptation. Standard deviation over 20 sampled adaptation sets are reported.

interesting observation is that the head head_0 consistently achieves the highest overall accuracy for both models. This aligns with expectations, as head_0 corresponds to the eigenvector with the largest variance, i.e., the most informative direction. Furthermore, among the top 100 heads, most of the high-performing heads appear before index 40, which aligns with PCA’s property that the explained variance decreases as the head index increases. This finding further supports our argument that PCA can approximate preference learning.

In summary, our results show that a single reward head is insufficient to represent the full spectrum of human preferences. Instead, DRMs provide high-quality and interpretable estimations of diverse human preferences, supporting our analysis in Section 3.2.

4.3 Test-time Preference Adaptation

A natural application of DRMs is test-time adaptation—deriving linear combinations to match new

user preferences. Following the adaptation method in Section 3.4, we use a small subset of test data for each attribute (e.g., $n = 15$ for RewardBench and $n = 5$ for RPR), which corresponds to less than 4% of the available data per attribute in RewardBench and less than 6% for RPR. We compare DRMs against several baselines, including the single-head and ensemble-head baselines trained on the same dataset, and two random-head baselines that sample random heads. To ensure efficiency, we limit all models, including DRMs, to using only 100 heads. The results for the two benchmarks, shown in Table 2 and Table 3, include the standard deviation over 20 repetitions of sampled adaptation sets.

The results demonstrate that DRMs achieve the best overall performance across different base models and test sets. The improvement is particularly significant for Gemma-2B-RM, where DRMs improve the single-head baseline from 0.733 to 0.814 on RewardBench and from 0.714 to 0.90 on RPR.

Base Model	Method	Chat	Chat Hard	Safety	Reasoning	Overall
Gemma-2B-RM	Single Head	0.944	0.467	0.759	0.759	0.733
	Shared-Base	0.947(0.000)	0.476(0.000)	0.765(0.000)	0.774(0.000)	0.740(0.000)
	Random (Uniform)	0.940(0.005)	0.567(0.029)	0.800(0.010)	0.843(0.019)	0.787(0.009)
	Random (Gaussian)	0.951(0.005)	0.573(0.033)	0.781(0.015)	0.839(0.021)	0.786(0.008)
	DRMs(Ours)	0.953(0.003)	0.650(0.028)	0.783(0.030)	0.872(0.025)	0.814(0.013)
Llama3-8B-RM	Single Head	0.989	0.684	0.891	0.920	0.871
	Shared-Base	0.986(0.000)	0.684(0.000)	0.895(0.000)	0.927(0.000)	0.873(0.000)
	Random (Uniform)	0.985(0.003)	0.623(0.089)	0.903(0.010)	0.915(0.014)	0.857(0.023)
	Random (Gaussian)	0.982(0.004)	0.663(0.096)	0.889(0.009)	0.936(0.011)	0.868(0.024)
	DRMs(Ours)	0.986(0.002)	0.755(0.032)	0.885(0.036)	0.914(0.036)	0.885(0.012)

Table 3: Evaluation Results ($n = 15$), comparing different methods across two base models.

Benchmark	Attributes	Single Head	Random (Uniform)	DRMs (Ours)
RewardBench	Overall	0.759	0.770	0.830
	Chat	0.905	0.897	0.920
	Chat Hard	0.621	0.600	0.692
	Safety	0.699	0.753	0.786
	Reasoning	0.813	0.832	0.920
RPR	Overall	0.746	0.630	0.796
	User-Friendliness	0.640	0.555	0.657
	Narrative & Storytelling	0.713	0.610	0.763
	Linguistic Creativity	0.808	0.595	0.843
	Scientific Rigor	0.762	0.661	0.806
	Humor & Entertainment	0.798	0.744	0.905

Table 4: Performance of DRMs on **instruction-tuned model** Gemma-2-9b-it. Single head baseline is trained with the same dataset used for DRMs. Aligned with the previous setting, we use $n = 15$ and $n = 5$ for RewardBench and RPR respectively.

Interestingly, the ‘‘Shared-Base’’ ensemble baseline does not outperform the single-head baseline, suggesting that it lacks diversity in its learned reward heads. In contrast, the random-head baseline offers some improvement over the single-head baseline but remains inferior to DRMs. This confirms that DRMs provide a diverse and well-structured set of basis vectors that enable efficient test-time adaptation to user preferences.

Another important question is whether a language model can serve as a feature extractor instead of a reward model. To explore this, we conduct experiments using Gemma-2-9B-it as the feature extractor. The results, presented in Table 4, show that DRMs can successfully exploit a language model for this purpose, outperforming the single-head trained baseline by 7.8% on RewardBench and 26% on RPR. Furthermore, comparing results across models reveals that while DRMs perform well with language models as feature extractors, reward models remain a more effective choice.




Figure 2: Weight distributions of the top 100 decomposed reward heads on RewardBench for DRMs using Gemma-2B-RM as the backbone.

4.4 Quantitative Attribute Explainability

Beyond its adaptability, DRMs offer a significant advantage in interpretability, helping not only to understand human preferences but also to analyze the multiple attributes present in current test sets. Specifically, for each attribute subset, we can obtain the weight parameters $\mathbf{k} = [k_1, \dots, k_d]$ corresponding to each basis vector. Figure 2 visualizes some of these weight vectors in RewardBench, revealing distinct patterns. The *Chat* subset primarily relies on the first few basis vectors, which capture the most data variance. In contrast, the *Chat Hard* and *Safety* subsets exhibit a more uniform weight distribution, while the *Reasoning* subset depends more heavily on basis vectors indexed between 1 and 50. These variations highlight the differing preference requirements across subsets.

Using the weight vectors \mathbf{k} for all attributes, we compute their Pearson correlation, providing a quantitative explanation of the dataset’s attributes. The resulting correlation matrix, shown in Figure 3, reveals meaningful relationships between attributes.

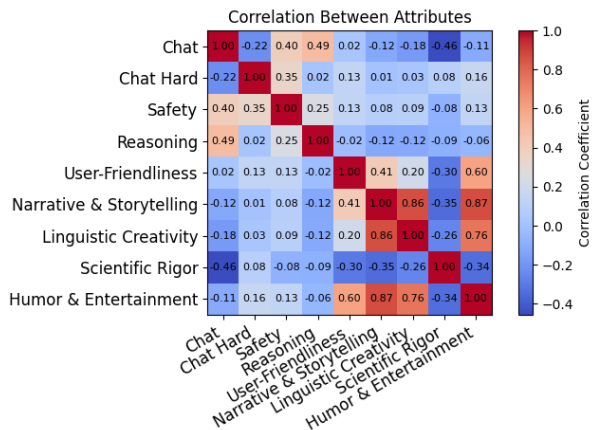


Figure 3: Correlation among attributes’ weight vectors for DRMs. The feature extractor is Gemma-2B-RM.

For instance, “*Narrative & Storytelling*” strongly correlates with “*Humor & Entertainment*” and “*Linguistic Creativity*” (with a correlation of approximately 0.87), which aligns with the idea that humor and creativity enhance storytelling. This correlation suggests that some attributes may be redundant. On the other hand, “*Scientific Rigor*” negatively correlates with several attributes, including “*Chat*” and “*Narrative & Storytelling*” (Pearson $r = -0.46$ and -0.35 , respectively), suggesting that scientific rigor may conflict with other human preferences. Additionally, many attributes show weak or negligible correlations (with absolute Pearson $r < 0.1$). Overall, DRMs provide a structured framework for quantitatively explaining attribute relationships, offering deeper insights into benchmark design and multi-attribute evaluation.

4.5 Ablation Study

We analyze two key factors affecting test-time adaptation: adaptation set size and the number of DRM heads used. Using Gemma-2B-RM as the feature extractor, we present results on RPR and RewardBench in Figure 4. Our findings show that performance improves with a larger adaptation set, converging on RewardBench at $n \geq 15$. Similarly, increasing the number of heads enhances performance but saturates beyond 100, likely because the most meaningful PCA directions lie within the first 100 heads. When the adaptation set is small (e.g., $n = 3$), performance is unstable, and fewer heads can yield better results. This may be due to difficulty in correctly weighting heads with limited data, whereas using more heads increases the risk of assigning incorrect weights. However, with sufficient data, more heads eventually lead to better performance. These results suggest that a slightly




Figure 4: Ablations on the adaptation set size and number of reward heads for test-time adaptation based on Gemma-2B-RM.

larger adaptation set and a carefully chosen number of heads are key to optimizing performance.

5 Related work

The Heterogeneity of Reward Modeling. Reward models (Lambert et al., 2024; Liu et al., 2024b) are typically trained using preference annotations from the Bradley-Terry model (Christiano et al., 2017; Bradley and Terry, 1952) or demonstration data (Wulfmeier et al., 2024; Xiao et al., 2024). However, human preferences are diverse and complex, making it difficult to capture all relevant attributes with a single objective (Yang et al., 2024c; Rame et al., 2024; Chakraborty et al., 2024). To address this, researchers are exploring multi-objective preference learning. This includes collecting datasets that assess multiple attributes (Wu et al., 2023; Wang et al., 2023; Cui et al., 2023; Pitis et al., 2024) and developing multi-head reward models that learn diverse user preferences (Quan, 2024; Wang et al., 2024b).

Embedding-based Reward Model. Recent advances in reward modeling have highlighted embedding-based approaches for their efficiency and scalability (Ahmed et al., 2024; Sun et al., 2023; Zhang et al., 2024a). These models are backed by strong theoretical foundations (Sun et al., 2024) and demonstrate high flexibility with competitive or superior performance (Li et al., 2024b; Tennenholtz et al., 2024). Additionally, they integrate well with established statistical learning tools (Dykstra, 1960; Springall, 1973; Han et al., 2020) and offer greater transparency through statistical insights (Shen et al., 2025; Feng et al., 2025). Unlike prior methods, our approach represents human preferences through the final linear layer, enhancing interpretability and enabling the decomposition of preference components.

Dimensionality Reduction and Embedding Anal-

ysis. Linear dimensionality reduction techniques like PCA have proven effective in extracting latent dimensions that capture key human preferences for model alignment (Freire et al., 2024). Beyond alignment, broader studies have explored structuring high-dimensional data into more interpretable embeddings (Huertas-García et al., 2023; Kanerva et al., 2000). Methods such as Q-Probe (Li et al., 2024b) and DeepMDP (Gelada et al., 2019) further enhance model alignment by efficiently exploring embedding spaces. Recent work also shows that improving embedding quality leads to better model performance and generalization (Li et al., 2024a; Yang et al., 2024b).

6 Conclusion

In this paper, we establish the connection between preference learning and PCA, introducing Decomposed Reward Models (DRMs). DRMs represent diverse human preferences as a set of orthogonal basis vectors using a novel vector-based formulation of preference. This approach enables efficient test-time adaptation to user preferences without requiring additional training, making it both scalable and practical. Beyond the efficiency, DRMs provide a structured way to understand human preferences. By decomposing complex preferences into interpretable components, they reveal how preferences are formed and interact. We hope this work inspires further research into the fundamentals of human preference learning while promoting more transparent and personalized AI systems.

7 Limitations

In this paper, we obtain a large number of decomposed rewards from DRMs. However, due to the large scale (e.g., 2048 or 4096 reward heads), we did not manually examine each head to identify its corresponding preference attribute. Future work could focus on developing automated methods to analyze these rewards, such as recognizing patterns in the first 100 reward heads and assessing whether the last 100 primarily capture noise or meaningful subtleties. Additionally, we did not incorporate interdisciplinary study by collaborating with psychology experts to explore human preferences in depth. Future research could benefit from such collaboration to bridge the gap between computational models and cognitive science.

8 Ethics Statement

This paper introduces Decomposed Reward Models (DRMs), a step toward improving multi-objective alignment in LLMs. Here, we discuss the potential benefits of our approach while acknowledging the associated risks.

Our method enhances LLM alignment with diverse human preferences. DRMs are lightweight, flexible, and easily adaptable to new users and evolving preferences. Their efficiency reduces resource demands and broadens accessibility, paving the way for personalized preference learning and scalable LLM alignment. By offering an interpretable framework, DRMs promote greater transparency and customization in human preference modeling.

We carefully follow the license for all datasets and models used in our paper. Human preference datasets often contain biases, reflecting the perspectives and prejudices of their sources. If not properly managed, these biases could propagate through the model, influencing decomposition and principal components, potentially leading to unintended consequences. Mitigating this risk requires careful curation, filtering, and bias reduction before large-scale deployment. Additionally, our method does not inherently control the meaning of each reward head, which could unintentionally capture harmful human preferences. Therefore, thorough evaluation is necessary before deployment to ensure ethical and responsible use.

9 Acknowledgements

We thank the anonymous reviewers for their valuable comments. We thank the Chili Lab at Rice for helpful discussions and suggestions throughout this research.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Ahmed M Ahmed, Rafael Rafailov, Stepan Sharkov, Xuechen Li, and Sanmi Koyejo. 2024. Scalable ensembling for mitigating reward overoptimisation. *arXiv preprint arXiv:2406.01013*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al.

2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Ralph Allan Bradley and Milton E Terry. 1952. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Souradip Chakraborty, Jiahao Qiu, Hui Yuan, Alec Koppel, Furong Huang, Dinesh Manocha, Amrit Bedi, and Mengdi Wang. 2024. Maxmin-rlhf: Towards equitable alignment of large language models with diverse human preferences. In *ICML 2024 Workshop on Models of Human Feedback for AI Alignment*.
- Hanjie Chen, Zhouxiang Fang, Yash Singla, and Mark Dredze. 2025. Benchmarking large language models on answering and explaining challenging medical questions. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3563–3599.
- Keertana Chidambaram, Karthik Vinay Seetharaman, and Vasilis Syrgkanis. 2024. Direct preference optimization with unobserved preference heterogeneity. *arXiv preprint arXiv:2405.15065*.
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Ganqu Cui, Lifan Yuan, Ning Ding, Guanming Yao, Wei Zhu, Yuan Ni, Guotong Xie, Zhiyuan Liu, and Maosong Sun. 2023. Ultrafeedback: Boosting language models with high-quality feedback. *arXiv preprint arXiv:2310.01377*.
- Hanze Dong, Wei Xiong, Bo Pang, Haoxiang Wang, Han Zhao, Yingbo Zhou, Nan Jiang, Doyen Sahoo, Caiming Xiong, and Tong Zhang. 2024. Rlhf workflow: From reward modeling to online rlhf. *arXiv preprint arXiv:2405.07863*.
- Otto Dykstra. 1960. Rank analysis of incomplete block designs: A method of paired comparisons employing unequal repetitions on pairs. *Biometrics*, 16(2):176–188.
- Yunzhen Feng, Ariel Kwiatkowski, Kunhao Zheng, Julia Kempe, and Yaqi Duan. 2025. Pilaf: Optimal human preference sampling for reward modeling. *arXiv preprint arXiv:2502.04270*.
- Pedro Freire, ChengCheng Tan, Adam Gleave, Dan Hendrycks, and Scott Emmons. 2024. Uncovering latent human wellbeing in language model embeddings. *arXiv preprint arXiv:2402.11777*.
- Carles Gelada, Saurabh Kumar, Jacob Buckman, Ofir Nachum, and Marc G Bellemare. 2019. Deepmdp: Learning continuous latent space models for representation learning. In *International conference on machine learning*, pages 2170–2179. PMLR.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, and et al. 2024. *The llama 3 herd of models*. *Preprint*, arXiv:2407.21783.
- Ruijian Han, Rougang Ye, Chunxi Tan, and Kani Chen. 2020. Asymptotic theory of sparse bradley–terry model. *The Annals of Applied Probability*, 30(5):2491–2515.
- Álvaro Huertas-García, Alejandro Martín, Javier Huertas-Tato, and David Camacho. 2023. Exploring dimensionality reduction techniques in multilingual transformers. *Cognitive Computation*, 15(2):590–612.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*.
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*.
- Pentti Kanerva, Jan Kristoferson, and Anders Holst. 2000. Random indexing of text samples for latent semantic analysis. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 22.
- Nathan Lambert, Valentina Pyatkin, Jacob Morrison, LJ Miranda, Bill Yuchen Lin, Khyathi Chandu, Nouha Dziri, Sachin Kumar, Tom Zick, Yejin Choi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. *Rewardbench: Evaluating reward models for language modeling*. *Preprint*, arXiv:2403.13787.
- Yoonho Lee, Jonathan Williams, Henrik Marklund, Archit Sharma, Eric Mitchell, Anikait Singh, and Chelsea Finn. 2024. Test-time alignment via hypothesis reweighting. *arXiv preprint arXiv:2412.08812*.
- Jiachen Li, Weixi Feng, Wenhui Chen, and William Yang Wang. 2024a. Reward guided latent consistency distillation. *arXiv preprint arXiv:2403.11027*.
- Kenneth Li, Samy Jelassi, Hugh Zhang, Sham Kakade, Martin Wattenberg, and David Brandfonbrener. 2024b. Q-probe: A lightweight approach to reward maximization for language models. *arXiv preprint arXiv:2402.14688*.

- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Chris Yuhao Liu, Liang Zeng, Jiakai Liu, Rui Yan, Jujie He, Chaojie Wang, Shuicheng Yan, Yang Liu, and Yahui Zhou. 2024b. Skywork-reward: Bag of tricks for reward modeling in llms. *arXiv preprint arXiv:2410.18451*.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.
- Ian Osband, Zheng Wen, Seyed Mohammad Asghari, Vikranth Dwaracherla, Morteza Ibrahimi, Xiuyuan Lu, and Benjamin Van Roy. 2023. Epistemic neural networks. *Advances in Neural Information Processing Systems*, 36:2795–2823.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Silviu Pitis, Ziang Xiao, Nicolas Le Roux, and Alessandro Sordani. 2024. Improving context-aware preference modeling for language models. *arXiv preprint arXiv:2407.14916*.
- Shuang Qiu, Dake Zhang, Rui Yang, Boxiang Lyu, and Tong Zhang. 2024. Traversing pareto optimal policies: Provably efficient multi-objective reinforcement learning. *arXiv preprint arXiv:2407.17466*.
- Shanghaoran Quan. 2024. Dmoerm: Recipes of mixture-of-experts for effective reward modeling. *arXiv preprint arXiv:2403.01197*.
- Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa Shukor, Laure Soulier, and Matthieu Cord. 2024. Rewarded soups: towards pareto-optimal alignment by interpolating weights fine-tuned on diverse rewards. *Advances in Neural Information Processing Systems*, 36.
- David A Ross, Jongwoo Lim, Ruei-Sung Lin, and Ming-Hsuan Yang. 2008. Incremental learning for robust visual tracking. *International journal of computer vision*, 77:125–141.
- Yunyi Shen, Hao Sun, and Jean-François Ton. 2025. Reviving the classics: Active reward modeling in large language model alignment. *arXiv preprint arXiv:2502.04354*.
- A Springall. 1973. Response surface fitting using a generalization of the bradley-terry paired comparison model. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 22(1):59–68.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. Learning to summarize with human feedback. *Advances in Neural Information Processing Systems*, 33:3008–3021.
- Hao Sun, Alihan Hüyük, and Mihaela van der Schaar. 2023. Query-dependent prompt evaluation and optimization with offline inverse rl. In *The Twelfth International Conference on Learning Representations*.
- Hao Sun, Yunyi Shen, and Jean-Francois Ton. 2024. Rethinking bradley-terry models in preference-based reward modeling: Foundations, theory, and alternatives. *arXiv preprint arXiv:2411.04991*.
- Hao Sun, Yunyi Shen, Jean-Francois Ton, and Mihaela van der Schaar. 2025. Reusing embeddings: Reproducible reward model research in large language model alignment without gpus. *arXiv preprint arXiv:2502.04357*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Guy Tennenholtz, Yinlam Chow, Chih-Wei Hsu, Lior Shani, Ethan Liang, and Craig Boutilier. 2024. Embedding-aligned language models. *arXiv preprint arXiv:2406.00024*.
- Yuanhe Tian, Ruyi Gan, Yan Song, Jiaying Zhang, and Yongdong Zhang. 2023. Chimed-gpt: A chinese medical large language model with full training regime and better alignment to human preferences. *arXiv preprint arXiv:2311.06025*.
- Guanchu Wang, Junhao Ran, Ruixiang Tang, Chia-Yuan Chang, Yu-Neng Chuang, Zirui Liu, Vladimir Braverman, Zhandong Liu, and Xia Hu. 2024a. Assessing and enhancing large language models in rare disease question-answering. *arXiv preprint arXiv:2408.08422*.
- Haoliang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. 2024b. Interpretable preferences via multi-objective reward modeling and mixture-of-experts. *arXiv preprint arXiv:2406.12845*.

- Kaiwen Wang, Rahul Kidambi, Ryan Sullivan, Alekh Agarwal, Christoph Dann, Andrea Michi, Marco Gelmi, Yunxuan Li, Raghav Gupta, Avinava Dubey, et al. 2024c. Conditional language policy: A general framework for steerable multi-objective finetuning. *arXiv preprint arXiv:2407.15762*.
- Zhilin Wang, Yi Dong, Jiaqi Zeng, Virginia Adams, Makeesh Narsimhan Sreedhar, Daniel Egert, Olivier Delalleau, Jane Polak Scowcroft, Neel Kant, Aidan Swope, et al. 2023. Helpsteer: Multi-attribute helpfulness dataset for steerlm. *arXiv preprint arXiv:2311.09528*.
- Zequ Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. 2023. Fine-grained human feedback gives better rewards for language model training. *Advances in Neural Information Processing Systems*, 36:59008–59033.
- Markus Wulfmeier, Michael Bloesch, Nino Vieillard, Arun Ahuja, Jorg Bornschein, Sandy Huang, Artem Sokolov, Matt Barnes, Guillaume Desjardins, Alex Bewley, et al. 2024. Imitating language via scalable inverse reinforcement learning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Teng Xiao, Mingxiao Li, Yige Yuan, Huaisheng Zhu, Chao Cui, and Vasant G Honavar. 2024. How to leverage demonstration data in alignment for large language model? a self-imitation learning perspective. *arXiv preprint arXiv:2410.10093*.
- Kailai Yang, Zhiwei Liu, Qianqian Xie, Tianlin Zhang, Nirui Song, Jimin Huang, Ziyang Kuang, and Sophia Ananiadou. 2024a. Metaaligner: Conditional weak-to-strong correction for generalizable multi-objective alignment of language models. *arXiv preprint arXiv:2403.17141*.
- Rui Yang, Hanyang Chen, Junyu Zhang, Mark Zhao, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, et al. 2025. Embodiedbench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents. *arXiv preprint arXiv:2502.09560*.
- Rui Yang, Ruomeng Ding, Yong Lin, Huan Zhang, and Tong Zhang. 2024b. Regularizing hidden states enables learning generalizable reward model for llms. *arXiv preprint arXiv:2406.10216*.
- Rui Yang, Xiaoman Pan, Feng Luo, Shuang Qiu, Han Zhong, Dong Yu, and Jianshu Chen. 2024c. Rewards-in-context: Multi-objective alignment of foundation models with dynamic preference adjustment. *arXiv preprint arXiv:2402.10207*.
- Yifan Zhang, Ge Zhang, Yue Wu, Kangping Xu, and Quanquan Gu. 2024a. General preference modeling with preference representations for aligning language models. *arXiv preprint arXiv:2410.02197*.
- Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. 2024b. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*.
- Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. 2024. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19632–19642.
- Zhanhui Zhou, Jie Liu, Jing Shao, Xiangyu Yue, Chao Yang, Wanli Ouyang, and Yu Qiao. 2024. Beyond one-preference-fits-all alignment: Multi-objective direct preference optimization. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 10586–10613.

A Appendix

A.1 Ablations on Llama3-8B-RM

We add the ablation results on Llama3-8B-RM in Figure 5. The trend is similar to the ablations in our main paper.




Figure 5: Ablations on the adaptation set size and number of reward heads for test-time adaptation on Llama3-8B-RM.

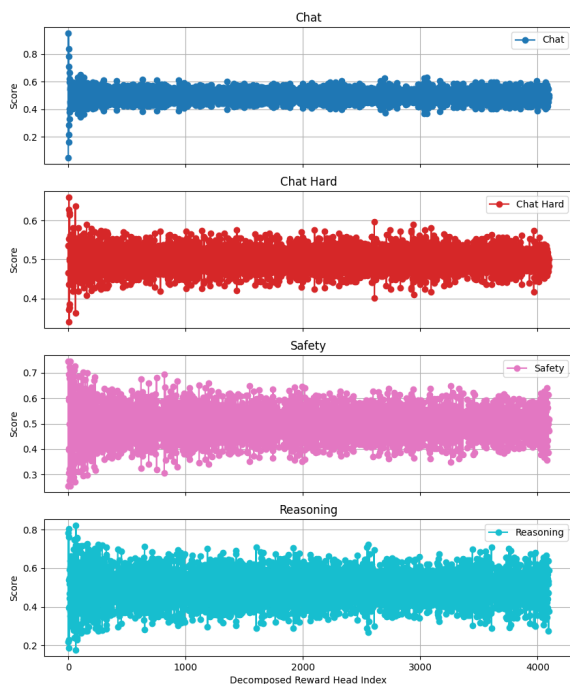


Figure 6: Reward scores of all decomposed reward heads on RewardBench for DRMs using Gemma-2B-RM as the backbone.

A.2 Reward Scores on decomposed reward heads

As shown in Figure 6 and 7, we visualize the reward scores of individual decomposed reward heads on both RewardBench and RPR, respectively. With a hidden dimension of 2048 in Gemma-2B-RM, the total number of reward heads is 4096. While

most heads’ scores fall within a certain range, a few outliers are identified, which will be utilized during test-time preference adaptation for specific tasks.




Figure 7: Reward scores of all decomposed reward heads on RPR for DRMs using Gemma-2B-RM as the backbone.

A.3 Implementation Details

For all training-based reward head models, including both the Single Head and Share-Base variants, we train them on the *mixture2_and_safe_pku* dataset² (Dong et al., 2024) for one epoch with a batch size of 16.

For our proposed Decomposed Reward Models (DRMs), we apply Principal Component Analysis (PCA) using scikit-learn’s default settings to get the component vectors. We experiment with 100 heads for all methods. Note that DRMs utilize 50 distinct reward heads, and including their negative counterparts results in a total of 100 heads.

²https://huggingface.co/datasets/weqweasdas/preference_dataset_mixture2_and_safe_pku

A.4 Scalability of DRMs

DRMs are already highly efficient and require significantly lower computational cost than training-based methods. In our experiments, we applied DRMs to a large preference dataset of size 0.5M, one of the largest open-source preference datasets. For comparison, training-based reward models using LLaMA3-8B (4096 hidden dimension) require 1-2 hours on a 48GB A6000 GPU. In contrast, the PCA step in DRMs runs in less than 1 minutes on the same server using only CPUs.

While PCA involves covariance matrix computations that can become challenging at extreme scales, this is not a bottleneck for current dataset sizes (<1M). For larger datasets, scalable PCA solvers can be applied, eg. Incremental PCA (Ross et al., 2008), Distributed PCA and Randomized PCA. These methods are well-supported in libraries like scikit-learn and PySpark, enabling DRMs to scale to millions of samples efficiently. To validate this, we conducted a scalability test using Incremental PCA across varying dataset sizes. As shown in Table 5, Incremental PCA exhibits approximately linear time complexity, demonstrating its practical scalability.

Dataset size	0.5 M	1 M	10 M
Running Time	45.59 s	97.09 s	1117.13 s

Table 5: Running time of Incremental PCA with different dataset sizes.