GRADIENT-BASED DIVERSITY OPTIMIZATION WITH DIFFERENTIABLE TOP-k OBJECTIVE

Anonymous authors

000

001

002003004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

024

025

026027028

029

031

033

034

037

038

040

041

042

043

044

046

047

048

051

052

Paper under double-blind review

ABSTRACT

Predicting relevance is a pervasive problem across digital platforms, covering social media, entertainment, and commerce. However, when optimized solely for relevance and engagement, many machine-learning models amplify data biases and produce homogeneous outputs, reinforcing filter bubbles and content uniformity. To address this issue, we introduce a pairwise top-k diversity objective with a differentiable smooth-ranking approximation, providing a model-agnostic way to incorporate diversity optimization directly into standard gradient-based learning. Building on this objective, we cast relevance and diversity as a joint optimization problem, we analyze the resulting gradient trade-offs, and propose two complementary strategies: direct optimization, which modifies the learning objective, and indirect optimization, which reweights training data. Both strategies can be applied either when training models from scratch or when fine-tuning existing relevanceoptimized models. We use recommendation as a natural evaluation setting where scalability and diversity are critical, and show through extensive experiments that our methods consistently improve diversity with negligible accuracy loss. Notably, fine-tuning with our objective is especially efficient, requiring only a few gradient steps to encode diversity at scale.

1 Introduction

Modern prediction models are typically evaluated by their ability to maximize accuracy, that is, to rank or classify items in line with ground-truth labels. However, optimizing exclusively for accuracy often yields homogeneous outputs: models repeatedly select similar items, overlook long-tail instances, and fail to provide novel or diverse options. This challenge is particularly evident in recommender systems, which influence decisions in shopping, entertainment, and news (Ricci et al., 2010). However, it also arises broadly in ranking and selection tasks across machine learning (Wang et al., 2023). When diversity is neglected, models risk reinforcing bias, amplifying popularity effects, and reducing the utility of top-k prediction sets (Nguyen et al., 2014; Lambrecht & Tucker, 2019).

Diversity has thus emerged as an important complementary criterion. Recent user studies confirm that diverse outputs can improve satisfaction and engagement when achieved without major loss in relevance (Chen et al., 2018a; Holtz et al., 2020; Anderson et al., 2020). The trade-off between relevance and diversity is mainly tackled by either *post-hoc re-ranking methods* or *learning-based methods*. *Post-hoc re-ranking methods* (e.g., MMR, DPP) modify top-k sets to improve diversity (Carbonell & Goldstein, 1998; Chen et al., 2018b), but typically suffer from degrading accuracy when diversity increases (Chen et al., 2017). *Model-specific learning-based methods* integrate diversity into training objectives (Borodin et al., 2017; Hurley, 2013; Wang et al., 2023), achieving strong gains with modest accuracy loss. However, these methods rely on opaque models that obscure the source of diversity gains, exhibit slow convergence, and are sensitive to the choice of trade-off parameter balancing relevance and diversity. Additionally, *data-centric approaches* such as augmentation, reweighting, and debiasing (Wang et al., 2021; Lai et al., 2023; Ren et al., 2018; Rastegarpanah et al., 2019) address bias in data distributions, but do not explicitly target diversity. Despite these efforts, we still lack a *unified, differentiable, and model-agnostic framework* for optimizing both relevance and diversity directly during training.

To address these challenges, we propose a unified framework that leverages differentiable ranking to optimize diversity in top-k prediction sets in a scalable and model-agnostic way. At the core, we use

an effective diversity objective that can be integrated into the gradient-based training without requiring architecture changes or post-processing. Building on this objective, we introduce two diversification methods. (i) *direct diversity-guided tuning* (DDT), which augments the loss with a joint relevance–diversity term, and (ii) *meta-diversity reweighting* (MDR), which preserves relevance-only training while reweighting data points using the joint loss as a meta-objective. Our approach offers a flexible alternative to post-hoc or model-specific diversification without compromising efficiency.

Our contributions are threefold: (1) We propose a unified differentiable framework for optimizing relevance and diversity in top-k prediction sets, applicable to both end-to-end training and fine-tuning; (2) we provide a theoretical analysis of gradient conflicts, deriving feasible intervals for the trade-off parameter β and showing that an adaptive update coincides with the two-objective solution of multi-gradient descent algorithm (MGDA), guaranteeing convergence to Pareto-stationary points; (3) we empirically validate the framework on five benchmark datasets and two model architectures, demonstrating that DDT and MDR achieve substantial diversity improvements with minimal relevance loss, outperforming strong baselines. Notably, the diversity gains extend beyond the explicitly optimized top-k range, reshaping subsequent predictions as well.

2 RELATED WORK

Our work is related to diverse recommender systems and multi-objective learning.

Diversity in recommender systems. Among the vast literature on recommender systems, the closest are post-hoc and learning-based diversification methods (Zhao et al., 2025); see the survey for a broader overview. *Post-hoc* methods re-rank the output of a relevance-only model to balance relevance and diversity. Representative approaches include maximal marginal relevance (MMR) (Carbonell & Goldstein, 1998), diversity-weighted utility maximization (DUM) (Ashkan et al., 2015), and determinantal point processes (DPP) (Chen et al., 2018b). These methods are model-agnostic and easy to implement, but their performance is limited by the quality of the initial relevance ranking, and diversity gain usually comes at a cost of reduced accuracy (Chen et al., 2017).

Learning-based approaches incorporate diversity objectives directly into training, including penalties for similarity among recommended items (Hurley, 2013; Wasilewski & Hurley, 2016), formulations that optimize relevance—diversity trade-offs (Wang et al., 2023) list-wise, and graph-based models that encourage coverage of item categories or long-tail exposure (Zheng et al., 2021; Yang et al., 2023). While they often outperform post-hoc re-ranking, they require architectural modifications or adversarial training, making them model-specific and computationally heavy. In contrast, our framework is differentiable and model-agnostic: it can be integrated into standard training pipelines without altering architectures or adding inference overhead.

Multi-objective learning. Related is the study of multi-objective optimization for balancing goals such as accuracy, fairness, and revenue (Zheng & Wang, 2022). Classical approaches include scalarization (Paul et al., 2022; Di Noia et al., 2017), which reduces multiple objectives to a single weighted loss, and population-based heuristics such as evolutionary algorithms (Cai et al., 2020), which approximate the Pareto front. While effective in some cases, these approaches either rely on carefully tuned weights or suffer from high computational cost.

More recently, gradient-based methods such as the multi-gradient descent algorithm (MGDA) (Désidéri, 2012) have been applied to recommendation. For instance, MGDA has been used to balance accuracy with fairness (Du et al., 2025; Wu et al., 2022) and with revenue (Milojkovic et al., 2019). These methods guarantee convergence to Pareto-stationary solutions, but their application has so far been limited to objectives other than diversity.

3 Problem Formulation

In this section, we introduce top-k diversity and our objectives, starting with notation. Given a set of candidate items $\mathcal{I} = \{i_1, i_2, \dots, i_m\}$ and a collection of users $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$, the goal is to identify those items in \mathcal{I} that are most relevant for each $u \in \mathcal{U}$. We assume access to a partially-observed supervision matrix $\mathbf{R} \in \mathbb{R}^{n \times m}$, where entries $\mathbf{R}_{u,i}$ represent relevance scores (e.g., rating, label, or interaction). The set of all *observed* scores is $\Omega = \{(u, i, \mathbf{R}_{u,i})\}^N$, where $N \ll nm$ is the number of observations. We use a relevance prediction model \mathcal{F}_{Θ} with parameters

 Θ to estimate the remaining scores $\tilde{\mathbf{R}}_{u,i} = \mathcal{F}_{\Theta}(u,i)$. To learn such a model, we consider two widely-used approaches. First, we consider *matrix factorization* (MF) (Koren et al., 2009) which predicts ratings $\tilde{\mathbf{R}}_{u,i} = \mathbf{x}_u^{\top} \mathbf{y}_i$, for user u and item i using the embeddings $\mathbf{x}_u, \mathbf{y}_i \in \mathbb{R}^d$. Second, we consider *neural network models* (He et al., 2017), which predict ratings $\tilde{\mathbf{R}}_{u,i} = \text{MLP}([\mathbf{x}_u; \mathbf{y}_i])$ with a multi-layer perceptron (MLP), allowing non-linear interactions.

In both cases, the models are trained to minimize the regularized mean squared error (MSE)

$$\mathbf{\Theta}^* = \arg\min_{\mathbf{\Theta}} \sum_{(u,i) \in \Omega_{\mathbf{T}}} (\mathbf{R}_{u,i} - \tilde{\mathbf{R}}_{u,i})^2 + \lambda \|\mathbf{\Theta}\|_2^2 , \qquad (1)$$

between observed relevance scores and predictions. While this approach yields models that predict accurately, neither MF nor MLP is optimized for diversity. Top-k diversity seeks to predict relevance scores $\tilde{\mathbf{r}}_u \in \mathbb{R}^m$ of a user u, so that the top-k scores correspond to a set of diverse items $Z_u(k) = \{z_1, \ldots, z_k\}$, indicated by $\mathbf{l}_u = \mathrm{top}_k(\tilde{\mathbf{r}}_u)$. To measure the diversity of the top-k (highest scoring) items, we take a *distance-based* approach (Hassin et al., 1997). In particular, for a given pairwise item-item affinity matrix $\mathbf{S} \in \mathbb{R}^{m \times m}$, we define diversity as the *average pairwise dissimilarity*

$$D_{\mathbf{S}}(Z_u(k)) = \frac{2}{k(k-1)} \sum_{i=1}^{m} \sum_{j=1}^{m} \mathbf{l}_u(i) \mathbf{l}_u(j) (1 - \mathbf{S}_{i,j})$$
(2)

of items in $Z_u(k)$. Then, the average top-k diversity is simply

$$\mathcal{L}_{DRO}(k) = \frac{1}{n} \sum_{u=1}^{n} D_{\mathbf{S}}(Z_u(k)) . \tag{3}$$

While this serves as a natural diversity objective, it cannot be directly optimized with gradients as $top_k(\cdot)$ involves non-differentiable operations. Our goal, however, is to optimize for relevance and top-k diversity *simultaneously* as part of gradient-based optimization. To achieve this, we relax the non-differentiable diversity reward objective (DRO) using a *differentiable surrogate* (DDRO) next.

Differentiable diversity. To overcome the non-differentiability challenge in $top_k(\cdot)$, we adopt differentiable ranking (Blondel et al., 2020), which is a continuous relaxation of sorting. The key idea is to replace the discrete permutation \mathbf{z}_u with a soft ranking vector $\tilde{\mathbf{z}}_u^{(\varepsilon)} \in \mathbb{R}^m$, obtained by projecting the predicted scores $\tilde{\mathbf{r}}_u$ onto the permutahedron \mathcal{P}_m —the convex hull of all permutations of $(1, 2, \ldots, m)$ embedded in an m-dimensional space. This projection is computed by solving the following entropy-regularized optimization problem (Blondel et al., 2020):

$$\tilde{\mathbf{z}}_{u}^{(\varepsilon)} = \operatorname{softrank}(\tilde{\mathbf{r}}_{u}) := \arg\min_{r \in \mathcal{P}_{m}} \left\{ \frac{1}{\varepsilon} \langle \tilde{\mathbf{r}}_{u}, r \rangle + H(r) \right\}, \tag{4}$$

where H(r) is the entropy regularizer and $\varepsilon>0$ controls the approximation smoothness. This definition enables the top-k soft indicator $\tilde{\mathbf{I}}_u(i)=\sigma_\tau(k-\tilde{\mathbf{z}}_u(i))$ using a scaled sigmoid function $\sigma_\tau(x)=[1+\exp(-x/\tau)]^{-1}$ for user-defined smoothness-sharpness τ , often set to 1 in our experiments. In turn, this makes it possible to train the prediction model end-to-end using diversity-aware gradient updates. Like the discrete counterpart, soft ranking operates with $O(n\log n)$ time and O(n) space complexity. The obvious question is, does soft ranking lead to sufficiently accurate top-k recommendations? The better soft ranking approximates the hard ranking, the more reliable are our top-k recommendations, formally summarized in Lemma. 1.

Lemma 1 (Soft rank approximation (Blondel et al., 2020)). Given a rating vector $\tilde{\mathbf{r}}_u$, let $\tilde{\mathbf{z}}_u^{(\varepsilon)} \in \mathbb{R}^n$ be the soft rank vector obtained from optimizing (4). Then, as $\varepsilon \to 0$, the soft ranks converge to the true ranks of $\tilde{\mathbf{r}}_u \lim_{\varepsilon \to 0} \tilde{\mathbf{z}}_u^{(\varepsilon)} = \operatorname{rank}(\tilde{\mathbf{r}}_u)$, where $\operatorname{rank}(\tilde{\mathbf{r}}_u) \in \{1, \dots, n\}^n$ denotes the discrete ranks (breaking ties arbitrarily).

Replacing indicator in Equation (3) with a the soft-ranking-derived indicator \tilde{l}_u yields the differentiable diversity reward objective

$$\mathcal{L}_{\text{DDRO}} = \frac{1}{n \cdot N} \sum_{u}^{n} \sum_{i=1}^{m} \sum_{j=1}^{m} \tilde{\mathbf{I}}_{u}(i) \tilde{\mathbf{I}}_{u}(j) (1 - \mathbf{S}_{i,j}) . \tag{5}$$

To achieve both a high relevance and a high diversity in top-k outputs, we balance both the relevance objective \mathcal{L}_{rel} and diversity objective \mathcal{L}_{div} introducing our problem below.

Problem 2. For a given model class \mathcal{F}_{Θ} , an item affinity matrix \mathbf{S} , a user-defined relevance-diversity trade-off $\beta \in [0,1]$; find parameters $\mathbf{\Theta}$ that minimize the joint loss

 $\mathcal{L}_{\text{JOINT}}(\beta, \mathbf{\Theta}) = \beta \, \mathcal{L}_{\text{rel}}(\mathbf{\Theta}) + (1 - \beta) \, \mathcal{L}_{\text{div}}(\mathbf{\Theta}) \,. \tag{6}$

In practice, we take relevance objective $\mathcal{L}_{\mathrm{rel}} = \mathcal{L}_{\mathrm{MSE}}$ and diversity objective $\mathcal{L}_{\mathrm{div}} = -\mathcal{L}_{\mathrm{DDRO}}$. While leading to an efficiently optimizable objective (6), joining them combines two diametrically opposed goals: diversity and relevance. Diverse outputs are not necessarily the most "relevant" ones, and vice versa. We discuss how to deal with this balance in the following section.

4 Joint Gradient-Based Training

In this section, we discuss how this objective can be incorporated into practical training. We propose two strategies: one *direct*, by optimizing the joint loss explicitly, and one *indirect*, by using the joint loss as a meta-objective to reweight training samples, starting with the direct approach.

4.1 BALANCING RELEVANCE AND DIVERSITY

Our first approach is straightforward: we directly optimize the joint loss (6) using gradient-based updates, named as direct diversity tuning (DDT). This loss balances relevance and diversity through a trade-off parameter $\beta \in [0,1]$. If both objectives have gradients that point in similar directions, we say that they are 'aligned'. If the two gradient directions are aligned, any β decreases both terms. However, when the gradients $g_{\rm rel} = \nabla \mathcal{L}_{\rm rel}$ and $g_{\rm div} = \nabla \mathcal{L}_{\rm div}$ are misaligned, any linear combination will necessarily favor one objective at the expense of the other. We want to ensure a descent that ensures a good diversity and accuracy balance for which we realign the gradients. For this, we adaptively compute the optimal balance parameter β^* that ensures simultaneous descent during optimization. Formally, for gradient norms $a = \|g_{\rm rel}\|$ and $b = \|g_{\rm div}\|$, we denote the cosine similarity by $\rho = \langle g_{\rm rel}, g_{\rm div} \rangle/(ab)$. The combined gradient is $g_{\beta} = \beta g_{\rm rel} + (1-\beta)g_{\rm div}$. For a step along $-g_{\beta}$ to decrease both objectives simultaneously, the following conditions

$$-\delta_{\rm rel}(\beta) > 0 \iff \langle g_{\rm rel}, g_{\beta} \rangle > 0, \quad -\delta_{\rm div}(\beta) > 0 \iff \langle g_{\rm div}, g_{\beta} \rangle > 0$$
 (A-B)

must be satisfied (Désidéri, 2012). Equivalently, the projections of the combined gradient g_{β} onto $g_{\rm rel}$ and $g_{\rm div}$ should be positive. This guarantees that a step along $-g_{\beta}$ decreases both losses at once. We give the feasible region of β that satisfies (A-B):

Lemma 3 (Common descent (Désidéri, 2012)). For any a, b > 0 and $\rho \in [-1, 1]$, the feasible region of $\beta \in [0, 1]$ satisfying (A–B) is

- 1. If $\rho > 0$ (aligned), all $\beta \in [0,1]$ are feasible.
- 2. If $\rho = 0$ (orthogonal), the feasible set is $\beta \in (0, 1)$.
- 3. If $\rho < 0$ (opposing), the feasible set is $\beta \in \left(\frac{b|\rho|}{a+b|\rho|}, \frac{b}{b+a|\rho|}\right)$.

Proof sketch. The result follows from expanding the directional derivatives $\delta_{\rm rel}(\beta) = -\beta a^2 - (1-\beta)ab\rho$ and $\delta_{\rm div}(\beta) = -(1-\beta)b^2 - \beta ab\rho$, and solving the inequalities (A–B) in the three cases $\rho > 0$, $\rho = 0$, and $\rho < 0$. Full details are provided in Appendix A.

Even when β lies in the feasible interval, different values may still lead to unbalanced overall progress, with one loss improving much more than the other. To avoid this, we select the parameter

$$\beta^{\star} = \arg \max_{\beta \in [0,1]} \min\{-\delta_{\rm rel}(\beta), \ -\delta_{\rm div}(\beta)\} \ ,$$

that maximizes the minimum per-step decrease, yielding a *multiple gradient descent algorithm* (MGDA) (Désidéri, 2012) specialized to two objectives. MGDA finds a convex combination of gradients that minimizes the maximum directional derivative across tasks using the closed-form

$$\beta^* = \frac{b(b - a\rho)}{a^2 + b^2 - 2ab\rho} \ . \tag{7}$$

When $\rho \leq 0$, the ideal balance β^* lies in the feasible interval described in Lemma 3, ensuring valid common descent. Intuitively, β^* equalizes the first-order decreases of relevance and diversity, yielding a balanced update. Finally, when β_t is chosen adaptively by the MGDA rule, we obtain convergence to Pareto-stationary solutions.

216217

Table 1: Statistics of datasets and diversity metrics.

224 225 226

228 229 230

227

231232233

234235236237

238239240241

242

243

244245246

247 248

249250251252253

256 257 258

254

255

259260261

262263264265266

267

268

269

 $DRO(\mathcal{I})$ **Dataset** $|\mathcal{C}|$ $|\mathcal{U}|$ $|\mathcal{I}|$ $|\Omega|$ Coat 290 300 6960 0.73 33 3 3 2 7 KuaiRec 1411 4676570 0.91 31 0.83 4999 27 Netflix 1112 557 176 684 782 58 Yahoo-R2 4050 5 000 0.26 6040 18 MovieLens 3 706 1 000 208 0.83

Corollary 4 (Pareto-stationarity with adaptive β_t (Sener & Koltun, 2018)). *Under diminishing step sizes, if each* β_t *is chosen by the minimax rule (or projected variant), then every accumulation point of* $\{\Theta^t\}$ *is Pareto-stationary for* $(\mathcal{L}_{rel}, \mathcal{L}_{div})$.

While the adaptive choice provides balanced progress at each step, a fixed β remains useful, as it avoids per-iteration computation. To contextualize our results, we introduce an additional alternative approach which operates with a fixed β next.

4.2 Optimizing Diversity by Example Reweighting.

A complementary line of work on machine learning fairness and robustness has shown that *reweighting training examples* can effectively mitigate inherent data bias. The central idea is to retain the standard prediction objective, but assign *adaptive weights* to individual samples so that the resulting model better aligns with a criterion (Ren et al., 2018; Rastegarpanah et al., 2019). This approach is naturally formulated as meta-learning: the inner loop minimizes a weighted relevance loss, while the outer loop adjusts weights using a meta-objective that encodes the criterion.

Inspired by this paradigm, we aim to improve diversity by reducing bias at the *data level*. To this end, we propose *meta-diversity reweighting* (MDR). Specifically, we introduce a per-sample weight $w_{u,i} \in [0,1]$ for each interaction (u,i) in a mini-batch \mathcal{B} , and optimize the *reweighted relevance loss*

$$\mathcal{L}_{\text{MSE}}^{w} = \sum_{(u,i)\in\mathcal{B}} w_{u,i} \left(\mathbf{R}_{u,i} - \tilde{\mathbf{R}}_{u,i}\right)^{2}. \tag{8}$$

The joint relevance–diversity objective is used only as a meta-loss to update the weights w, and is never applied directly to model parameters.

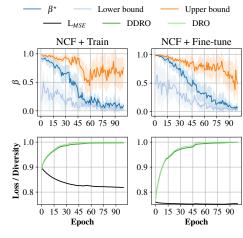
In contrast to direct diversity-guided tuning, which modifies the training objective itself, MDR preserves the standard relevance-oriented loop while implicitly reshaping the effective data distribution. The meta-objective encourages weights that downplay biased interactions and upweight samples that contribute to both accuracy and diversity. This reweighting perspective allows us to test whether diversity gains can be obtained not only by altering the optimization objective, but also by correcting data imbalance through implicit meta-optimization.

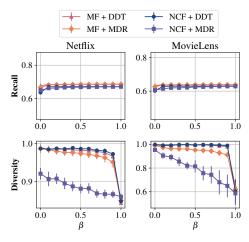
We evaluate both DDT and MDR in two settings: *from-scratch training*, where relevance and diversity are optimized jointly, and *fine-tuning*, where a relevance-trained model is adapted for diversity.

5 EXPERIMENTAL EVALUATION

We describe the experimental setup to evaluate the effectiveness of our proposed solutions, introducing the datasets and evaluation criteria in this section, and provide further details in Appendix C and D.

Datasets. We evaluate our methods across three domains: *entertainment*, *product*, and *social* recommendations. For entertainment, we use MovieLens, Netflix, and Yahoo-R2, which contain user ratings on movies or music with genre/category annotations. For product recommendation we use the Coat dataset, which includes user ratings and product attributes. For social recommendation we use KuaiRec, a large-scale mobile video dataset with watch-time based ratings. The item affinity matrices **S** are pre-computed by Jaccard similarity scores based on genre/category information of items. Basic statistics and detailed preprocessing are given in the Appendix C.





- (a) Optimization with adaptive β (MovieLens).
- (b) Optimization with varying 'static' β s.

Figure 1: Optimization of relevance-diversity joint objective. Figure 1a shows results of training and fine-tuning of NCF models with adaptive β on the MovieLens dataset. Figure 1b shows the performance comparison of DDT and MDR applied to two models (NMF and NCF) across two datasets with varying fixed $\beta \in [0,1]$.

Baselines. We compare against relevance- and diversity-aware approaches: (i) *Non-negative Matrix Factorization* (NMF) (Lee & Seung, 1999) as a relevance-only baseline; (ii) greedy post-processing methods such as *Maximal Marginal Relevance* (MMR) (Carbonell & Goldstein, 1998) and *Diversity-weighted Utility Maximization* (DUM) (Ashkan et al., 2015); (iii) probabilistic diversification via *Determinantal Point Processes* (DPP) (Chen et al., 2018b); and (iv) the graph-based *Diversified GNN Recommender* (DGRec) (Yang et al., 2023). We use the authors' publicly-available implementations.

Evaluation criteria. We evaluate accuracy using *hit rate*, *precision*, and *recall*. For each user u, let \mathcal{R}_u denote the top-k recommended items and \mathcal{T}_u the set of ground-truth relevant items (i.e., rated above a threshold), the *Hit rate* measures whether at least one relevant item appears in \mathcal{R}_u ; *precision* is the fraction of items in \mathcal{R}_u that are in \mathcal{T}_u ; and *recall* is the fraction of relevant items in \mathcal{T}_u that are retrieved in \mathcal{R}_u . When ground-truth is unknown, we report the estimated user satisfaction using preference likelihoods $p(\tilde{\mathbf{R}}_{u,i}) \in [0,1]$ and threshold $\tau = 0.8$ as the *relevance score*

$$Relevance(u) = \frac{1}{k} \sum_{i \in \mathcal{R}_u} \mathbb{I}[p(\tilde{\mathbf{R}}_{u,i}) > \tau] . \tag{9}$$

5.1 EXPERIMENTAL ANALYSIS

Having introduced our setup, we now introduce our research questions and experimental analysis.

- **Q1** Does adaptive β behave as predicted by our theory?
- **Q2** How does the choice of β affect the accuracy–diversity trade-off?
- **Q3** How do our approaches compare with established diversification methods?
- **Q4** How do diversity gains evolve with varying k?

In the following, we experimentally answer all these questions in detail.

Evaluating adaptive β **optimization.** We first study the direct optimization approach using both from-scratch training and fine-tuning of a relevance-pretrained model. At each step, we adopt the adaptive β rule, computing both the optimal coefficient β^* and the feasible interval from Lemma 3. Figure 1a shows from-scratch training and fine-tuning for 100 epochs on MovieLens. We employ the neural collaborative filtering (NCF) model with an MLP architecture, and consider top-10 diversity.

In the top row of Figure 1a, we observe that the adaptive coefficient β^* remains consistently within the feasible interval defined by Lemma 3. We see how β^* gradually shifts from favoring relevance

Table 2: Performance of DDT, MDR, and five competitors in terms of diversity and relevance across Coat, Yahoo-R2, Netflix, MovieLens, and KuaiRec. We highlight the best results with **bold**, and underline the second best, reporting mean and standard deviation over 10 trials.

Algorithm	Coat		Yahoo-R2		Netflix		MovieLens		KuaiRec	
_	Diversity	Relevance	Diversity	Relevance	Diversity	Relevance	Diversity	Relevance	Diversity	Relevance
NMF	0.77 (0.02)	0.41 (0.04)	0.09 (0.05)	0.76 (0.02)	0.84 (0.01)	0.88 (0.03)	0.62 (0.08)	0.98 (0.00)	0.89 (0.01)	0.84 (0.01)
MMR	0.80 (0.01)	0.40(0.04)	0.80 (0.06)	0.68 (0.03)	0.93 (0.01)	0.85 (0.04)	0.94 (0.03)	0.95 (0.01)	0.99 (0.00)	0.76 (0.01)
DUM	0.81 (0.01)	0.31 (0.04)	0.98 (0.02)	0.60 (0.02)	0.93 (0.00)	0.71 (0.03)	0.93 (0.01)	0.91 (0.01)	0.98 (0.00)	0.42 (0.01)
DPP	0.81 (0.01)	0.39 (0.04)	1.00 (0.00)	0.58 (0.01)	0.96 (0.00)	0.84 (0.04)	0.98 (0.01)	0.95 (0.01)	1.00 (0.00)	0.75 (0.01)
DGRec	0.71 (0.01)	0.69 (0.02)	0.33 (0.01)	0.83 (0.02)	0.76 (0.00)	0.83 (0.01)	0.73 (0.01)	0.47 (0.02)	0.91 (0.02)	0.18 (0.04)
DDT	0.83 (0.01)	0.50 (0.05)	0.98 (0.01)	$\overline{0.85\ (0.01)}$	0.98 (0.00)	0.97 (0.01)	1.00 (0.00)	1.00 (0.00)	0.98 (0.02)	0.95 (0.00)
MDR	0.82 (0.01)	0.47 (0.03)	0.86 (0.09)	0.82 (0.02)	0.98 (0.01)	0.93 (0.02)	0.97 (0.02)	0.99 (0.00)	0.97 (0.01)	0.85 (0.01)

toward favoring diversity, confirming that the adaptive rule preserves a common descent direction. The bottom row reports the optimization trajectories. The relevance loss ($\mathcal{L}_{\mathrm{MSE}}$) decreases steadily while diversity (DRO and DDRO) increases, showing that adaptive β achieves a balanced trade-off in both training objectives. We see that fine-tuning preserves the relevance of the pretrained model, whereas from-scratch converges to a worse stationary point.

We further observe that fine-tuning preserves the relevance performance of the pretrained model, whereas training from scratch reduces $\mathcal{L}_{\mathrm{MSE}}$ but converges to a worse stationary point. In contrast, both approaches reach a similar level of diversity, suggesting that regularization may hinder the convergence towards more 'relevant' models during from-scratch optimization. Fine-tuning, on the other hand, starts with a relevance-optimized model, which we effectively tune for diversity. Finally, because DRO and DDRO are close, we empirically validated the accuracy of our relaxation. We discuss the impact of the parameter ϵ on approximation accuracy later. Overall, we saw that adaptive β effectively delivers simultaneous improvements in relevance and diversity.

Relevance-diversity trade-off. We next examine how the choice of static β affects relevance and diversity. Figure 1b reports recall and diversity for DDT and MDR with MF and NCF on Netflix and MovieLens. Starting from a relevance-trained model, we fine-tune with the joint objective (6) while varying $\beta \in [0,1]$, ranging from $\beta = 1$ (pure relevance) to $\beta = 0$ (pure diversity). We run 10 random trials for each experiment and report the mean and variance. We observe a stable low-variance recall across a wide range of β , showing that introducing diversity does not substantially compromise accuracy with fixed β . A noticeable drop occurs only when β approaches zero, where the objective focuses almost exclusively on diversity. In contrast, diversity improves as β decreases. Moreover, DDT consistently outperforms MDR in terms of diversity with a neural network model, while the two strategies behave more similarly with a matrix factorization model, suggesting that reweighting is particularly beneficial for simpler low-rank models. These findings demonstrate that fixed- β optimization offers a practical means to explore the relevance—diversity trade-off, enabling substantial diversity gains with minimal loss of accuracy. Similar patterns are observed for MSE loss, hit rate, and precision, with detailed results on additional datasets provided in Appendix D.

Table 3: Diversity gain achieved by Direct Diversity Tuning (DDT) on top-1 $\sim k$ and $k+1 \sim 2k$ recommendations across five datasets Coat, Yahoo-R2, Netflix, MovieLens, and KuaiRec. We vary the diversity reward parameter $k \in \{5, 10, 20, 30, 40\}$ and apply DDT to pre-trained NMF and NCF, reporting diversity gain. We report the mean and standard deviation across 10 runs.

		Diversity gain (%)										
	Dataset			$1 \sim k$			$k+1 \sim 2k$					
		k=5	k=10	k=20	k=30	k=40	k=5	k=10	k=20	k=30	k=40	
	Coat	10.9 (4.3)	5.4 (1.8)	3.5 (1.7)	2.9 (1.1)	2.6 (0.8)	-0.5 (4.7)	-2.5 (4.4)	-1.8 (1.3)	-1.5 (1.7)	-1.9 (0.8)	
1	Yahoo-R2	75.0 (14.3)	77.4 (8.7)	78.5 (3.8)	78.4 (2.4)	77.2 (1.6)	54.9 (11.2)	50.8 (9.9)	47.1 (2.7)	45.2 (2.1)	41.9 (1.3)	
N	Netflix	14.7 (2.3)	13.8 (1.0)	11.3 (0.7)	10.2 (0.6)	9.6 (0.5)	7.8 (2.6)	4.2(1.7)	3.1(2.1)	2.4(2.2)	2.2(1.8)	
Z	MovieLens	41.4 (13.1)	37.2 (8.3)	28.4 (4.6)	23.8 (3.2)	21.6 (2.6)	25.6 (10.4)	16.2 (5.1)	10.1 (2.8)	8.3 (2.3)	6.8 (1.6)	
	KuaiRec	9.2 (1.7)	7.2 (0.5)	9.4 (0.8)	10.6 (0.7)	10.9 (0.8)	5.5 (1.5)	10.3 (2.1)	6.0 (2.0)	3.3 (1.0)	1.4 (1.4)	
	Coat	9.7 (2.4)	1.2 (1.8)	6.9 (1.3)	0.3 (0.8)	4.8 (0.5)	-0.9 (0.6)	3.5 (0.4)	-1.1 (0.7)	2.7 (0.5)	-1.0 (0.5)	
124	Yahoo-R2	79.5 (12.8)	9.4 (17.5)	79.3 (7.6)	4.6 (6.0)	78.2 (8.7)	4.9 (13.7)	77.5 (4.6)	0.6(9.0)	77.0 (6.3)	1.7 (7.3)	
NC	Netflix	13.7 (2.0)	8.2 (1.9)	13.3 (0.9)	4.5 (1.4)	10.8 (0.7)	3.9 (1.2)	9.5 (0.5)	3.6(1.1)	9.0(0.5)	3.2(1.1)	
_	MovieLens	46.8 (13.0)	19.6 (6.6)	36.2 (7.8)	12.7 (3.0)	26.9 (4.5)	10.9 (1.9)	22.9 (3.0)	9.6 (1.5)	20.9 (2.5)	9.0 (1.5)	
	KuaiRec	8.8 (1.3)	5.4 (1.8)	6.5 (0.3)	10.7 (2.6)	8.7 (0.7)	8.0 (1.2)	10.6 (0.7)	4.2 (0.9)	11.1 (0.5)	2.3 (0.9)	

Diversity-relevance performance of all approaches. Next, we compare the proposed methods against alternative diversification approaches across all datasets. Fixing k=10 and $\beta=0.2$, we fine-tune a pre-trained relevance-optimized matrix factorization model for 100 epochs and select the most diverse checkpoint. In Table 2 we report the mean and standard deviation of diversity and relevance scores for all methods from 10 trials.

We see that our direct diversity tuning approach (DDT) demonstrates strong performance in both relevance and diversity, achieving either the best or second-best results across nearly all metrics and datasets, showing the effectiveness of jointly optimizing. Similarly, we see that the implicit data reweight approach (MDR) shows highly competitive performance. Its performance is particularly notable on Netflix, MovieLens and KuaiRec, where it approaches or matches the performance of DDT. This supports our earlier observation that implicit, data-driven reweighting can offer strong benefits. On the other hand, while post-hoc diversification of NMF using MMR, DUM, and DPP, considerably improve top-k diversity, they often notably reduce relevance, aligning with the results reported in previous studies (Chen et al., 2017). For example, DPP occasionally achieves the highest diversity—particularly on Yahoo-R2 and KuaiRec—but at a cost of low relevance. Greedy methods like MMR and DUM yield moderate diversity improvements but underperform in relevance. DGRec excels in diversity on certain datasets but suffers from severe relevance degradation, especially on KuaiRec, with a diversity of 0.91 and a relevance of 0.18. In contrast, DDT and MDR maintain a significantly better trade-off. In summary, fine-tuning with direct and implicit methods both achieve superior relevance-diversity trade-offs across the board, outperforming the competitors.

Diversity gain with varying k. We examine how diversity changes across the growing number of recommended items. Starting with relevance-optimized matrix factorization and neural network models, we fine-tune with the joint objective $\mathcal{L}_{\text{JOINT}}$ with DDT for 10 epochs at a fixed $\beta=0.5$, and vary $k\in\{5,10,20,30,40\}$, selecting the best results. Since larger top-k sets approach coverage of the entire item space, it becomes impossible to improve diversity beyond the dataset's average. To account for this, we report relative diversity gains over the relevance-only baseline, normalized by the maximum achievable score. Specifically, we measure DRO gains as the difference between fine-tuned and pre-trained models, normalized to lie in [0,1] and reported as percentages. We distinguish between in-objective gains (top-k) and out-of-objective gains (the subsequent $k+1 \sim 2k$ items).

In Table 3, we observe positive *in-objective* diversity gains across all datasets. In Coat, Netflix, and MovieLens, we see a decreasing gain as k increases. The improvement is especially pronounced on Yahoo-R2, with diversity gains exceeding $78.5 \pm 3.8\%$ in NMF and 79.3 ± 7.6 in NCF at k=20, due to its low initial diversity (e.g., 0.09 in the base model). Substantial gains are also observed on MovieLens (e.g., $41.4 \pm 13.1\%$ and $46.8 \pm 13\%$ at k=5) and Netflix, demonstrating the effectiveness of our method across both sparse and dense recommendation scenarios.

Generality beyond the optimized range. We next ask whether the benefits of diversity optimization extend beyond the explicitly encoded objective. To this end, we evaluate diversity not only at the target top-k set but also in the subsequent $k+1 \sim 2k$ items in Table 3. We observe broadly positive $out\text{-}of\text{-}objective}$ gains, which again diminish as k increases. This effect is especially strong on Yahoo-R2 and MovieLens, showing that optimization reshapes the ranking itself: improvements are not limited to the items directly optimized, but generalize to deeper parts of the recommendation list. In contrast, results on Coat fluctuate around zero, likely due to its small scale and limited item pool. These findings highlight the generality of our diversity-guided optimization: it improves diversity not only within the objective's explicit target (top-k) but also beyond, demonstrating that the learned updates capture a broader notion of diversity than what is directly optimized.

Approximation error with vary ϵ . In Figure 2, we report the approximation error between the exact diversity DRO and its differentiable surrogate DDRO as the top-k size varies. We define the error as $|\mathcal{L}_{\mathrm{DRO}} - \mathcal{L}_{\mathrm{DDRO}}|/\mathcal{L}_{\mathrm{DRO}}$. For small k (e.g., k < 100), setting $\epsilon < 10^{-4}$ already yields a close approximation. As k increases, however, a smaller ϵ is required to maintain the same accuracy. This highlights that the choice of ϵ is critical: an improper value can lead to arbitrarily large errors, rendering the diversity measure unreliable for optimization.

Case study. To demonstrate the utility of our approaches, we examine the recommendations generated from an NMF model and its diversity-optimized variant. In Table 4, we observe that NMF produces homogeneous items where six out of ten items are labeled as *drama*, with remaining entries only marginally extending into *romance*, *documentary*, *comedy* or *War*. In contrast, the DDT-

Table 4: Comparison of recommendations lists for a user from DDT and NMF. Movies appearing in both lists are **bold**.

	Rank	Genres	Movie Title
	1	Horror	Vampyros Lesbos (Las Vampiras)
	2	Thriller	The Spiral Staircase
	3	War	Prisoner of the Mountains
	4	Animation, Musical	Melody Time
Ĭ	5	Documentary	The Gate of Heavenly Peace
Ī	6	Crime	Lured
	7	Drama	Mamma Roma
	8	Comedy	Smashing Time
	9	Adventure	Ulysses (Ulisse)
	10	Romance	Persuasion
	1	Drama	Mamma Roma
	2	Drama	Foreign Student
	3	Drama	The Apple
	4	Drama, Romance	Leather Jacket Love Story
₽	5	Comedy	Smashing Time
É	6	Documentary	The Gate of Heavenly Peace
	7	Documentary	Modulations
	8	Comedy, Romance, War	Forrest Gump
	9	Drama	Schlafes Bruder (Brother of Sleep)

Schindler's List

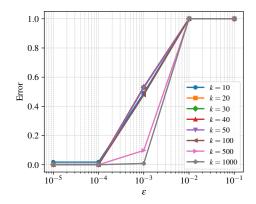


Figure 2: Approximation error of DDRO with different ϵ and k.

generated list spans a much broader range of genres, including *horror*, *thriller*, *animation*, *crime*, and *adventure*. Despite diversity gains, DDT retains three of the top items in relevance-optimized recommendation (*Mamma Roma*, *Smashing Time*, and *Gate of Heavenly Peace*), which collectively represent the relevant genre themes (*drama*, *comedy*, and *documentary*, respectively), suggesting that DDT successfully preserves relevant while enhancing diversity.

6 Conclusion

Drama, War

We addressed the issue of limited top-k diversity of relevance prediction models, which contributes to echo chambers, reduced novelty, and social polarization. By integrating diversity into a gradient-based optimization, we presented a unified framework for diversity-aware recommendation by introducing a differentiable diversity objective that enables end-to-end optimization of both relevance and diversity. We proposed two complementary, model-agnostic algorithms to support explicit and implicit integration of diversity into standard recommender systems. With extensive experiments on real-world datasets we demonstrated that our methods consistently improve diversity, converge efficiently, and introduce minimal computational overhead.

Limitations and future work. While our framework effectively promotes both relevance and diversity in top-k recommendations, several challenges remain. First, the adaptive β is applicable only for a direct joint objective optimization, while it is optimal in each step, but the optimization trajectory may converge to a compromised stationary point, which is not globally optimal. The mechanism that leads to a global optimal solution remains unexplored. Future work includes developing new ways to balance our competing goals via multi-objective formulations. Second, we primarily relies on categorical similarity to quantify diversity. While this provides interpretability, it does not capture more nuanced relationships or learned semantic embeddings. Future work extends the research into more expressive and context-sensitive diversity metrics. Finally, as diversity bias stems from the training data—implicit and explicit approaches often achieve similar performance levels—we see significant potential in the data-driven diversification, such as counterfactual data augmentation or diversity-aware sampling.

Despite these research opportunities, we observe that our methods significantly increase diversity with only an imperceptible decrease in relevance.

ETHICS STATEMENT

This work is primarily theoretical and focuses on incorporating diversity into machine-learning models for ranking. All datasets used in our experiments are publicly available, and no personally identifiable or sensitive information was collected or processed. No user studies or interventions involving human participants were conducted.

While our approach is motivated by the goal of improving diversity in ranking outcomes, we acknowledge that any deployment of such methods may have broader societal implications. Potential concerns include including fairness considerations, potential biases in the underlying data, or unintended effects depending on the application context. We leave a deeper exploration of these implications to future research.

REPRODUCIBILITY STATEMENT

Our code is available for review in an online anonymous repository. It will also be publicly available upon acceptance.

AI USAGE DISCLOSURE

Large language models (LLMs) were employed during the research phase to assist with surveying related literature, including identifying and summarizing relevant papers and methods, as well as sketching and testing proposed approaches described in prior work. Generative models were also used to draft scripts for data processing (e.g., formatting and visualizing data, preliminary experimental validation). These scripts were only used for exploration and are not part of the final experimental pipeline or released codebase. In addition, ChatGPT and Grammarly were used to assist with grammar and phrasing in the manuscript. All outputs from these tools were reviewed, edited, and verified by the authors, who take full responsibility for the final content.

REFERENCES

Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of The Web Conference*, pp. 2155–2165. ACM, 2020. URL https://dl.acm.org/doi/10.1145/3366423.3380281.

Azin Ashkan, Branislav Kveton, Shlomo Berkovsky, and Zheng Wen. Optimal greedy diversity for recommendation. In *IJCAI*, volume 15, pp. 1742–1748, 2015.

James Bennett and Stan Lanning. The Netflix prize, 2007.

Mathieu Blondel, Olivier Teboul, Quentin Berthet, and Josip Djolonga. Fast differentiable sorting and ranking. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, Proceedings of Machine Learning Research, pp. 950–959. PMLR, 2020.

Allan Borodin, Aadhar Jain, Hyun Chul Lee, and Yuli Ye. Max-sum diversification, monotone submodular functions, and dynamic updates. *ACM Transactions on Algorithms (TALG)*, pp. 1–25, 2017.

Xingjuan Cai, Zhaoming Hu, Peng Zhao, Wensheng Zhang, and Jinjun Chen. A hybrid recommendation system with many-objective evolutionary algorithm. *Expert Systems with Applications*, 159: 113648, 2020.

Jaime Carbonell and Jade Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 335–336, 1998.

¹https://anonymous.4open.science/r/DDRO-B975

- Laming Chen, Guoxin Zhang, and Hanning Zhou. Improving the diversity of top-n recommendation via determinantal point process. In *Large Scale Recommendation Systems Workshop*, pp. 3, 2017.
- Laming Chen, Guoxin Zhang, and Eric Zhou. Fast greedy map inference for determinantal point process to improve recommendation diversity. *Advances in Neural Information Processing Systems*, 31, 2018a.
 - Laming Chen, Guoxin Zhang, and Eric Zhou. Fast greedy map inference for determinantal point process to improve recommendation diversity. *Advances in Neural Information Processing Systems*, 2018b.
 - Jean-Antoine Désidéri. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathematique*, 350(5-6):313–318, 2012.
 - Tommaso Di Noia, Jessica Rosati, Paolo Tomeo, and Eugenio Di Sciascio. Adaptive multi-attribute diversity for recommender systems. *Information Sciences*, 382:234–253, 2017.
 - Gideon Dror, Noam Koenigstein, Yehuda Koren, and Markus Weimer. The yahoo! music dataset and kdd-cup'11. In *Proceedings of KDD Cup 2011*, pp. 3–18. PMLR, 2012.
 - Yongkang Du, Jieyu Zhao, Yijun Yang, and Tianyi Zhou. Controllable pareto trade-off between fairness and accuracy. *arXiv preprint arXiv:2509.13651*, 2025.
 - Chongming Gao, Shijun Li, Wenqiang Lei, Jiawei Chen, Biao Li, Peng Jiang, Xiangnan He, Jiaxin Mao, and Tat-Seng Chua. Kuairec: A fully-observed dataset and insights for evaluating recommender systems. In *CIKM*, pp. 540–550, 2022.
 - F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *ACM transactions on interactive intelligent systems (TIIS)*, 2015.
 - Refael Hassin, Shlomi Rubinstein, and Arie Tamir. Approximation algorithms for maximum dispersion. *Operations Research Letters*, pp. 133–137, 1997.
 - Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 173–182, 2017.
 - David Holtz, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. The engagement-diversity connection: Evidence from a field experiment on spotify. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 75–76, 2020.
 - Neil J Hurley. Personalised ranking with diversity. In *Proceedings of the 7th ACM Conference on Recommender Systems*, pp. 379–382, 2013.
 - Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, pp. 30–37, 2009.
 - Riwei Lai, Li Chen, Yuhan Zhao, Rui Chen, and Qilong Han. Disentangled negative sampling for collaborative filtering. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pp. 96–104, 2023.
 - Anja Lambrecht and Catherine Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management science*, pp. 2966–2981, 2019.
 - Daniel D Lee and H Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *nature*, 1999.
 - Nikola Milojkovic, Diego Antognini, Giancarlo Bergamin, Boi Faltings, and Claudiu Musat. Multigradient descent for multi-objective recommender systems. *arXiv preprint arXiv:2001.00846*, 2019.
 - Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, pp. 677–686, 2014.

- Agyemang Paul, Zhefu Wu, Kai Liu, and Shufeng Gong. Robust multi-objective visual bayesian personalized ranking for multimedia recommendation. *Applied Intelligence*, 52(4):3499–3510, 2022.
- Bashir Rastegarpanah, Krishna P. Gummadi, and Mark Crovella. Fighting fire with fire: Using antidote data to improve polarization and fairness of recommender systems. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, WSDM '19, pp. 231–239, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450359405. doi: 10.1145/3289600.3291002. URL https://doi.org/10.1145/3289600.3291002.
- Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to Reweight Examples for Robust Deep Learning. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 4334–4343. PMLR, 2018.
- Francesco Ricci, Lior Rokach, and Bracha Shapira. Introduction to recommender systems handbook. In *Recommender systems handbook*, pp. 1–35. Springer, 2010.
- Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *international conference on machine learning*. PMLR, 2016.
- Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Proceedings* of the 32nd International Conference on Neural Information Processing Systems, NIPS'18, pp. 525–536, Red Hook, NY, USA, 2018. Curran Associates Inc.
- Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. Denoising implicit feedback for recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pp. 373–381, 2021.
- Zihong Wang, Yingxia Shao, Jiyuan He, Jinbao Liu, Shitao Xiao, Tao Feng, and Ming Liu. Diversity-aware deep ranking network for recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pp. 2564–2573, 2023.
- Jacek Wasilewski and Neil Hurley. Incorporating diversity in a learning to rank recommender system. In *FLAIRS*, pp. 572–578, 2016.
- Haolun Wu, Chen Ma, Bhaskar Mitra, Fernando Diaz, and Xue Liu. A multi-objective optimization framework for multi-stakeholder fairness-aware recommendation. *ACM Transactions on Information Systems*, 41(2):1–29, 2022.
- Liangwei Yang, Shengjie Wang, Yunzhe Tao, Jiankai Sun, Xiaolong Liu, Philip S Yu, and Taiqing Wang. Dgrec: Graph neural network for recommendation with diversified embedding generation. In *Proceedings of the sixteenth ACM international conference on web search and data mining*, pp. 661–669, 2023.
- Yuying Zhao, Yu Wang, Yunchao Liu, Xueqi Cheng, Charu C Aggarwal, and Tyler Derr. Fairness and diversity in recommender systems: a survey. *ACM Transactions on Intelligent Systems and Technology*, pp. 1–28, 2025.
- Yong Zheng and David Xuejun Wang. A survey of recommender systems with multi-objective optimization. *Neurocomputing*, 474:141–153, 2022.
- Yu Zheng, Chen Gao, Liang Chen, Depeng Jin, and Yong Li. Dgcn: Diversified recommendation with graph convolutional networks. In *Proceedings of the web conference 2021*, pp. 401–412, 2021.

649 650

651 652 **PROOFS**

Inequality (B).

696

697

699 700

A.1 PROOF OF LEMMA 3 653 654 We minimize both 655 $\mathcal{L}_{ ext{rel}} := \mathcal{L}_{ ext{MSE}},$ $\mathcal{L}_{\mathrm{div}} := -\mathcal{L}_{\mathrm{DDRO}},$ 656 so decreasing $\mathcal{L}_{\mathrm{div}}$ increases diversity. Denote the gradients 657 $g_{\beta} = \beta g_{\text{rel}} + (1 - \beta)g_{\text{div}}, \quad \beta \in [0, 1],$ 658 $g_{\rm rel} = \nabla \mathcal{L}_{\rm rel}, \qquad g_{\rm div} = \nabla \mathcal{L}_{\rm div},$ 659 so that direct diversity-guided tuning (DDT) performs 660 661 $\Theta^+ = \Theta - \eta g_{\beta}$. 662 Let 663 $a = ||g_{\text{rel}}|| > 0, \qquad b = ||g_{\text{div}}|| > 0, \qquad \rho = \frac{\langle g_{\text{rel}}, g_{\text{div}} \rangle}{ab} \in [-1, 1].$ 664 665 666 1: Exact first-order changes. For an infinitesimal step along $-g_{\beta}$, the first-order changes in each 667 objective are 668 $\delta_{\rm rel} = \frac{d}{dt} \mathcal{L}_{\rm rel}(\Theta - tg_{\beta})\Big|_{t=0} = -\langle g_{\rm rel}, g_{\beta} \rangle,$ 669 670 $\delta_{\text{div}} = \frac{d}{dt} \mathcal{L}_{\text{div}}(\Theta - tg_{\beta}) \Big|_{t=0} = -\langle g_{\text{div}}, g_{\beta} \rangle.$ 671 672 Substituting $g_{\beta} = \beta g_{\rm rel} + (1 - \beta) g_{\rm div}$ gives 673 674 $\delta_{\rm rel} = -\beta \|q_{\rm rel}\|^2 - (1-\beta) \langle q_{\rm rel}, q_{\rm div} \rangle = -\beta a^2 - (1-\beta) ab\rho,$ 675 676 $\delta_{\text{div}} = -(1-\beta)\|q_{\text{div}}\|^2 - \beta\langle q_{\text{div}}, q_{\text{rel}}\rangle = -(1-\beta)b^2 - \beta ab\rho.$ 677 678 **2: Common descent conditions.** We require both decreases: 679 $\delta_{\rm rel} < 0, \quad \delta_{\rm div} < 0.$ 680 681 Rearranging each inequality yields the following necessary and sufficient conditions: 682 $\beta a + (1 - \beta)b\rho > 0$. (A) 683 684 $(1 - \beta)b + \beta a\rho > 0.$ (B) 685 686 **3: Case analysis of inequalities.** We now analyze (A) and (B) systematically. 687 688 *Inequality* (A). $\beta a + (1 - \beta)b\rho > 0 \iff \beta(a - b\rho) > -b\rho.$ 689 690 • If $a - b\rho > 0$, then $\beta > \frac{-b\rho}{a - b\rho}$, we need further analyze ρ . 691 692 • If $a - b\rho < 0$, then $\beta < \frac{-b\rho}{a - b\rho}$, we need further analyze ρ . 693 • If $a - b\rho = 0$, then the inequality holds for all β . 694

In this appendix, we provide detailed proofs for the results on the trade-off parameter β .

 $(1-\beta)b + \beta a\rho > 0 \iff b + \beta(a\rho - b) > 0.$

• If $a\rho-b>0$, then $\beta>\frac{-b}{a\rho-b}$ (RHS negative, vacuous for $\beta\in[0,1]$).

• If $a\rho - b < 0$, then $\beta < \frac{b}{b - a\rho}$, we need further analyze ρ .

• If $a\rho = b$, then the inequality holds for all β .

4: Regime specialization. We specialize the above formulas to the three relevant regimes of ρ .

Case I: Aligned gradients ($\rho > 0$). Here the two gradients point in similar directions. We revisit inequalities (A) and (B).

Condition (A).

$$\beta a + (1 - \beta)b\rho > 0 \iff \beta(a - b\rho) > -b\rho.$$

- If $a-b\rho \geq 0$, then the right-hand side is nonpositive $(-b\rho \leq 0 \text{ since } \rho > 0)$. Thus the inequality is automatically satisfied for all $\beta \in [0, 1]$; no restriction is imposed.
- If $a-b\rho<0$, then the inequality gives $\beta<\frac{-b\rho}{a-b\rho}$. But since $a-b\rho<0$ and $-b\rho<0$, the right-hand side is positive. In fact, one can check that $\frac{-b\rho}{a-b\rho}>1$. Therefore any $\beta\in[0,1]$ still satisfies the inequality.

Condition (B).

$$(1-\beta)b + \beta a\rho > 0 \iff b + \beta(a\rho - b) > 0.$$

- If $a\rho b \ge 0$, then the right-hand side is increasing in β and at $\beta = 0$ equals b > 0, so the inequality holds for all $\beta \in [0, 1]$.
- If $a\rho b < 0$, then the inequality becomes $\beta < \frac{b}{b-a\rho}$. Since $a\rho b < 0$, the denominator $b-a\rho>0$ and thus the right-hand side exceeds 1. Hence the condition imposes no restriction within $\beta \in [0, 1]$.

Conclusion. Both inequalities (A) and (B) are therefore automatically satisfied when $\rho > 0$. Thus

$$\delta_{\mathrm{rel}} < 0 \quad \text{and} \quad \delta_{\mathrm{div}} < 0 \qquad \text{for every } \beta \in [0,1].$$

In other words, when gradients are aligned, any convex combination yields a valid common descent direction. At $\beta = 0$ or $\beta = 1$, one loss still strictly decreases while the other is nonincreasing; for $0 < \beta < 1$, both decrease strictly.

Case II: Orthogonal gradients ($\rho = 0$). Plugging $\rho = 0$ into (A) and (B) gives

$$\beta a > 0 \iff \beta > 0, \qquad (1 - \beta)b > 0 \iff \beta < 1.$$

$$\implies \delta_{\text{rel}} < 0, \ \delta_{\text{div}} < 0 \quad \text{for all } \beta \in (0, 1).$$

Case III: Opposing gradients ($\rho < 0$). Let $\rho = -|\rho| < 0$. Then

$$\text{From (A): } \beta > \frac{b|\rho|}{a+b|\rho|} =: L, \qquad \text{From (B): } \beta < \frac{b}{b+a|\rho|} =: U.$$

Thus the feasible interval is

$$\beta \in (L, U) = \left(\frac{b|\rho|}{a + b|\rho|}, \frac{b}{b + a|\rho|}\right).$$

This interval is nonempty whenever $|\rho| < 1$, since

$$U - L = \frac{ab(1 - |\rho|^2)}{(b + a|\rho|)(a + b|\rho|)} > 0.$$

If $|\rho| = 1$ (exactly opposite gradients), then any $\beta \in (0,1)$ still makes both $\delta_{\rm rel}, \delta_{\rm div} < 0$.

Conclusion. We have characterized the feasible region of β :

- If $\rho > 0$: every $\beta \in [0,1]$ yields common descent.
- If $\rho = 0$: every $\beta \in (0,1)$ yields common descent.
- If $\rho < 0$: feasible region is (L, U) as defined above.

This completes the proof of Lemma 3.

A.2 PROOF OF β^*

We seek β^* such that $-\delta_{\rm rel}(\beta) = -\delta_{\rm div}(\beta)$. Expanding:

$$\beta a^2 + (1 - \beta)ab\rho = (1 - \beta)b^2 + \beta ab\rho.$$

Rearranging,

$$(a^2 + b^2 - 2ab\rho)\beta = b^2 - ab\rho.$$

Thus

$$\beta^* = \frac{b(b - a\rho)}{a^2 + b^2 - 2ab\rho}.$$

When $\rho \leq 0$, substitution verifies $L < \beta^* < U$, where (L, U) is the interval from Lemma 3. Hence β^* is feasible and maximizes the minimum first-order decrease.

B ALGORITHMS

Direct diversity-guided tuning To obtain top-k relevant and diverse recommendations—without relying on post-hoc re-ranking or heuristic interventions—we propose an integrated approach. Our approach DDT integrates a β -balanced relevance-diversity loss into a standard end-to-end training pipeline. Specifically, we use gradient-based joint optimization, which promotes relevance and diversity. The overall procedure is outlined in Algorithm 1. This method preserves the simplicity and efficiency of conventional end-to-end training while explicitly aligning the optimization process with our dual objective. By embedding both goals within a single differentiable framework, DDT eliminates the need for separate re-ranking stages.

Meta diversity-guided reweighting Our second approach is inspired by meta-learning and bilevel optimization, where hyperparameters or per-example weights are optimized with respect to a validation objective (Ren et al., 2018). Here we ask "Which training examples, if included in the update, will reduce both training loss and improve diversity?" To this end, we estimate and assign importance weights $w_{u,i} \in [0,1]$ to training samples (u,i,r) for each batch $\mathcal{B} \subseteq \Omega_T$, yielding

$$\mathcal{L}_{\text{MSE}}^{w} = \sum_{(u,i)\in\mathcal{B}} w_{u,i} \left(\mathbf{R}_{u,i} - \tilde{\mathbf{R}}_{u,i}\right)^{2}.$$
 (10)

As shown in Algorithm 2, we begin each mini-batch update (line 4) by initializing the per-example weights $w_{u,i}=0$, effectively ignoring all samples. We then perform a one-step inner update to obtain a temporary model Θ' using the weighted training loss $\mathcal{L}_{\text{MSE}}^w$ (line 5–6), where w is currently all zero. Next, we re-evaluate predictions using the updated parameters (line 7) and compute the joint meta-loss $\mathcal{L}_{\text{JOINT}}$ (line 8), combining relevance and diversity. For this, we compute the gradient of $\mathcal{L}_{\text{JOINT}}$ with respect to w, producing a utility score for each sample. We rectify it via $\tilde{\mathbf{w}} = \max(-\nabla_{\mathbf{w}}\mathcal{L}_{\text{JOINT}},0)$ (line 9), and normalize (line 10) to obtain the final per-example weights \mathbf{w} , enforcing $\sum_{(u,i)\in\mathcal{B}} w_{u,i}=1$. Normalization ensures that the overall gradient magnitude—and thus the effective learning rate—remains consistent across training steps, similar to standard Stochastic Gradient Descent (SGD), which averages over the batch (Ren et al., 2018). The final model update is then performed using this reweighted loss (line 11). This implicit diversity-guide optimization in MDR allows us to explore whether modifying the data sampling distribution alone can help achieve a decent relevance-diversity trade-off. We provide an empirical answer in Section 5, demonstrating that such implicit reweighting can be a practical and effective alternative to explicit multi-objective optimization.

C DETAILED EXPERIMENT SETTING.

Datasets. To evaluate our methods across different recommendation scenarios, we consider datasets from three domains: entertainment, product, and social recommendations, as detailed in Tab. 1. To cover the *entertainment* domain, we use Netflix (Bennett & Lanning, 2007) and MovieLens (Harper & Konstan, 2015) for user—movie recommendations, as well as Yahoo-R2 (Dror et al., 2012) for

Algorithm 1 Direct Diversity-Guided Tuning (DDT)

```
Require: Initialization model parameter Θ, training data \Omega_{\rm T}, item distance matrix D, trade-off \beta, learning rate \eta
1: for each epoch from 1 to T do
2: for each mini-batch \mathcal{B} \in \Omega_{\rm T} do
3: \tilde{\mathbf{R}}_{u,i} \leftarrow \mathcal{F}_{\Theta}(u,i), for all (u,i,\mathbf{R}_{u,i}) \in \mathcal{B}
4: \mathcal{L}_{\rm MSE} \leftarrow {\rm Compute} \ \mathcal{L}_{\rm MSE}(\tilde{\mathbf{R}}_{u,i},\mathbf{R}_{u,i};\mathcal{B}) with (1)
5: \mathcal{L}_{\rm DDRO} \leftarrow {\rm compute} \ {\rm DDRO} \ {\rm with} \ (5)
6: \mathcal{L}_{\rm JOINT} \leftarrow \beta \cdot \mathcal{L}_{\rm MSE} - (1 - \beta) \cdot \mathcal{L}_{\rm DDRO}
7: \Theta \leftarrow \Theta - \eta \cdot \nabla_{\Theta} \mathcal{L}_{\rm JOINT}
8: return Θ
```

Algorithm 2 Meta Diversity Reweighting (MDR)

Require: Initialization model parameter Θ , training data Ω_T , item distance matrix \mathbf{D} , trade-off β , learning rate η

```
1: for each epoch from 1 to T do
                    for mini-batches \mathcal{B} \in \Omega_T do
  3:
                            \hat{\mathbf{R}}_{u,i} \leftarrow \mathcal{F}_{\Theta}(u,i), for all (u,i,\mathbf{R}_{u,i}) \in \mathcal{B}
  4:
                            \begin{array}{l} \mathcal{L}_{\mathrm{MSE}}^{w} \leftarrow \mathcal{L}_{\mathrm{MSE}}^{w}(\tilde{\mathbf{R}}_{u,i},\mathbf{R}_{u,i},\mathbf{w};\mathcal{B}) \text{ with (10)} \\ \mathbf{\Theta}' \leftarrow \mathbf{\Theta} - \boldsymbol{\eta} \cdot \nabla_{\mathbf{\Theta}} \mathcal{L}_{\mathrm{MSE}}^{w}, \text{update meta model.} \end{array}
  5:
  6:
                            \mathbf{R}'_{u,i} \leftarrow \mathcal{F}_{\mathbf{\Theta}'}(u,i), for all (u,i,\mathbf{R}_{u,i}) \in \mathcal{B}
  7:
                            \mathcal{L}_{\text{JOINT}} \leftarrow \beta \cdot \mathcal{L}_{\text{MSE}}^{w} - (1 - \beta) \cdot \mathcal{L}_{\text{DDRO}}\tilde{\mathbf{w}} \leftarrow \max(-\nabla_{\mathbf{w}} \mathcal{L}_{\text{JOINT}}, 0);
  8:
  9:
                           \mathbf{w} \leftarrow \frac{\tilde{\mathbf{w}}}{\sum_{j} \tilde{w} + \epsilon}; normalization according to Ren et al. (2018)
10:
                            \Theta \leftarrow \Theta - \eta \cdot \nabla_{\Theta} \mathcal{L}_{\mathrm{MSE}}^{w}(\tilde{\mathbf{R}}_{u,i}, \mathbf{R}_{u,i}, \mathbf{w}; \mathcal{B})
11:
12: return \Theta
```

user–music recommendations. These datasets contain user ratings on a 5-point scale [1, 5], along with genre or category annotations. For Netflix (respectively, Yahoo-R2), we randomly sample 3 000 items (respectively, 5 000) and retain users with at least 20 ratings (respectively, 100+ ratings). In the *product recommendation* setting, we use the Coat (Schnabel et al., 2016), which captures user-coat interactions in e-commerce. It contains [1, 5] ratings and item 'meta' attributes. Finally, for the *social recommendation* scenario, we consider the KuaiRec (Gao et al., 2022), which is collected from a mobile video-sharing platform, which includes play duration, video length, and 'watch ratios' from 0 (never watched) to 2 (twice watched), which we linearly interpolate to 5-star ratings for consistency.

Baselines. We compare our algorithm against a broad set of state-of-the-art recommender-systems methods, as well as diversification techniques covering greedy, probabilistic, and graph-based strategies. To study the impact of diversification, we employ classical *Non-negative MF* (NMF) (Lee & Seung, 1999) as a relevance-only baseline that does not use any diversity mechanisms. We also include two baselines from the post-processing family: *Maximal Marginal Relevance* (MMR) (Carbonell & Goldstein, 1998) and *Diversity-weighted Utility Maximization* (DUM) (Ashkan et al., 2015), both greedy diversification techniques applied on top of NMF as the underlying model.

MMR greedily selects top-k items that maximizes a weighted combination of relevance and dissimilarity with previously selected items. DUM, on the other hand, uses a submodular combination of relevance and category-based diversity reward. Determinantal Point Processes (DPP) (Chen et al., 2018b) estimates the likelihood of item sets to be diverse and relevant as the determinant of an item-item similarity kernel matrix, from which we select the top-k using a greedy selection. Finally, we include a recent embedding-based method, Diversified GNN Recommender (DGRec) (Yang et al., 2023), which introduces a diversity-aware aggregation mechanism into graph neural networks by selecting neighbors that maximize coverage over item categories.

Evaluation criteria. We evaluate model accuracy using metrics suitable for top-k recommendations, hit rate, precision, and recall metrics. For each user u, let \mathcal{R}_u denote the top-k recommended items

Table 5: Diversity gain achieved by Direct Diversity Tuning (DDT) on top-1 $\sim k$ and $k+1 \sim 2k$ recommendations across five datasets Coat, Yahoo-R2, Netflix, MovieLens, and KuaiRec. We vary the diversity reward parameter $k \in \{5, 10, 20, 30, 40\}$ and apply MDR to pre-trained NMF and NCF, reporting diversity gain. We report the mean and standard deviation across 10 runs.

		Diversity gain (%)									
	Dataset			$1 \sim k$	$k+1 \sim 2k$						
		k=5	k=10	k=20	k=30	k=40	k=5	k=10	k=20	k=30	k=40
NMF	KuaiRec Coat MovieLens Netflix Yahoo-R2	8.5 (1.6) 5.2 (3.4) 36.4 (11.1) 14.2 (2.3) 49.8 (19.1)	5.2 (2.0) 2.7 (4.0) 22.0 (9.9) 4.8 (3.2) 36.0 (10.8)	6.5 (0.6) 2.3 (1.1) 32.3 (6.9) 13.0 (1.3) 59.0 (12.4)	9.9 (1.3) -1.2 (3.7) 12.1 (4.4) 2.1 (2.1) 34.2 (7.7)	8.7 (0.8) 1.1 (1.3) 25.4 (4.1) 10.3 (0.9) 65.1 (5.8)	5.8 (1.3) 0.4 (1.2) 7.3 (2.9) 1.6 (1.7) 33.1 (2.4)	10.1 (0.7) 1.4 (0.9) 21.7 (3.1) 9.2 (0.6) 66.5 (3.2)	2.9 (0.9) -0.4 (2.0) 6.8 (1.7) 1.4 (2.0) 31.4 (1.3)	10.3 (0.8) 1.8 (0.7) 19.7 (2.6) 8.7 (0.5) 66.5 (2.3)	0.9 (0.8) -1.2 (1.1) 6.4 (1.4) 1.4 (1.5) 29.6 (1.0)
NCF	KuaiRec Coat MovieLens Netflix Yahoo-R2	1.3 (1.8) 0.8 (2.3) 33.2 (11.8) 4.6 (2.4) 27.0 (15.3)	-0.0 (1.2) 0.9 (1.2) 6.1 (4.3) 1.4 (1.4) -0.2 (8.8)	0.3 (0.4) 0.8 (0.9) 17.8 (5.9) 2.8 (1.3) 20.8 (10.2)	1.2 (1.8) 0.3 (0.7) 3.6 (2.3) -0.1 (1.1) 2.3 (8.3)	1.1 (0.9) 0.7 (0.9) 9.5 (3.8) 1.1 (0.9) 19.5 (8.5)	0.9 (1.2) -0.3 (0.8) 1.9 (2.0) -0.5 (0.8) 7.3 (7.1)	1.3 (0.9) 0.3 (0.6) 4.9 (2.8) 0.4 (0.4) 16.7 (8.1)	0.2 (0.9) -0.1 (0.3) 1.1 (1.3) -0.2 (0.3) 2.0 (5.4)	1.4 (0.6) 0.2 (0.4) 3.4 (1.7) 0.2 (0.3) 16.0 (6.9)	0.0 (0.8) 0.1 (0.6) 0.6 (1.0) -0.0 (0.3) 3.2 (3.2)

and \mathcal{T}_u the set of ground-truth relevant items (i.e., rated above 4 in the test set). *Hit rate* measures whether at least one relevant item appears in \mathcal{R}_u ; *precision* is the fraction of items in \mathcal{R}_u that are in \mathcal{T}_u ; and *recall* is the fraction of relevant items in \mathcal{T}_u that are retrieved in \mathcal{R}_u .

However, the above requires known ground-truth. To evaluate the relevance of unseen items, we compute the potential user satisfaction as the *relevance score*

$$Relevance(u) = \frac{1}{k} \sum_{i \in \mathcal{R}_u} \mathbb{I}[p(\tilde{\mathbf{R}}_{u,i}) > \tau], \tag{11}$$

using preference likelihoods $p(\tilde{\mathbf{R}}_{u,i}) \in [0,1]$.

Experiment environment. We conduct all our experiments on 2 AMD Epyc 7742 CPUs, 1 TB of RAM and 1 NVIDIA DGX-A100 GPU. Our code is written in Python v3.11.7. All results are averaged over 10 independent runs. We adopt the Adam optimizer during the optimization.

Detailed Parameter setting. In all experiment, we set $\epsilon = 10^{-4}$ for a exact rank approximation. During the optimization with adaptive β , we choose a learning rate l=0.01 for the matrix factorization model, and l=0.001 for the neural network model. We keep the Adam weight decay parameter as the default.

D ADDITIONAL EXPERIMENT

We report the additional experiment results that are omitted in the main content. Figure 3 shows the adaptive β based optimization of two datasets (MovieLens and Netflix) and two models (NMF and NCF). The patterns are consistently aligned with the analysis. Figure 4 shows the 4 accuracy metric results: recall, mean squared error, hit rate, and precision, along with the diversity reword objective score (DRO). Similar to Table 3, Table 5 presents the *in-objective* and *out-objective* diversity gain of NMF and NCF with MDR.

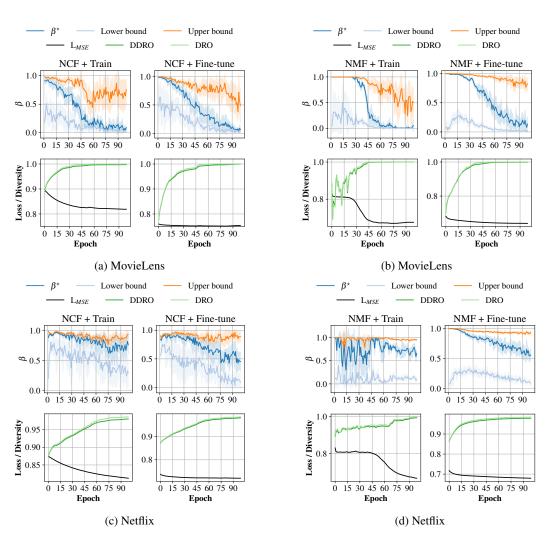


Figure 3: Optimization of relevance-diversity joint objective with adaptive β .

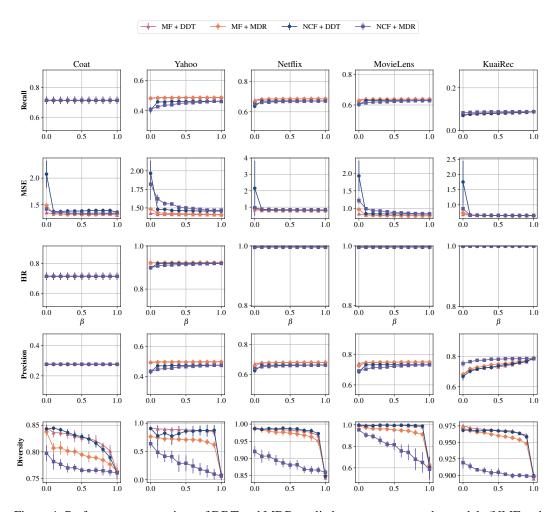


Figure 4: Performance comparison of DDT and MDR applied to two recommender models (NMF and NCF) across five datasets: Coat, Yahoo-R2, Netflix, MovieLens, and KuaiRec. We vary the parameter $\beta \in [0,1]$ to control the trade-off between relevance and diversity, where $\beta=1$ corresponds to optimizing only the relevance loss $\mathcal{L}_{\rm MSE}$ and $\beta=0$ corresponds to optimizing only the diversity loss $\mathcal{L}_{\rm DDRO}$. The x-axis indicates the value of β . The y-axis shows recall (top row), mean squared error(second row), hit rate (third row), precision (fourth row), and diversity reword objective (DRO) score with k=10 (bottom row). In each setting, we initialize from a pre-trained model (using $\mathcal{L}_{\rm MSE}$ only), then fine-tune with the joint loss $\mathcal{L}_{\rm JOINT}$ for 10 epochs, selecting the best result by diversity score. All experiments are repeated 10 times, and we report the mean and standard deviation.