Gradient-based Diversity Optimization with Differentiable Top-K Objective

Tianyi Zhou[†], Sebastian Dalleiger[†], Ece Calikus[‡], and Aristides Gionis[†]

†KTH Royal Institute of Technology

‡Uppsala University

Recommendation Homogeneity

Prediction models for user preference

Recommender systems trained to **maximize relevance** tend to suggest similar, popular items repeatedly.

Consequence:

- Top-k recommendations become homogeneous—dominated by a few popular categories.
- Long-tail items (niche content, minority creators) are systematically under-exposed.
- Users are trapped in filter bubbles, reducing serendipity and exploration.

Accuracy-optimal ≠ user-optimal: diversity matters for satisfaction.

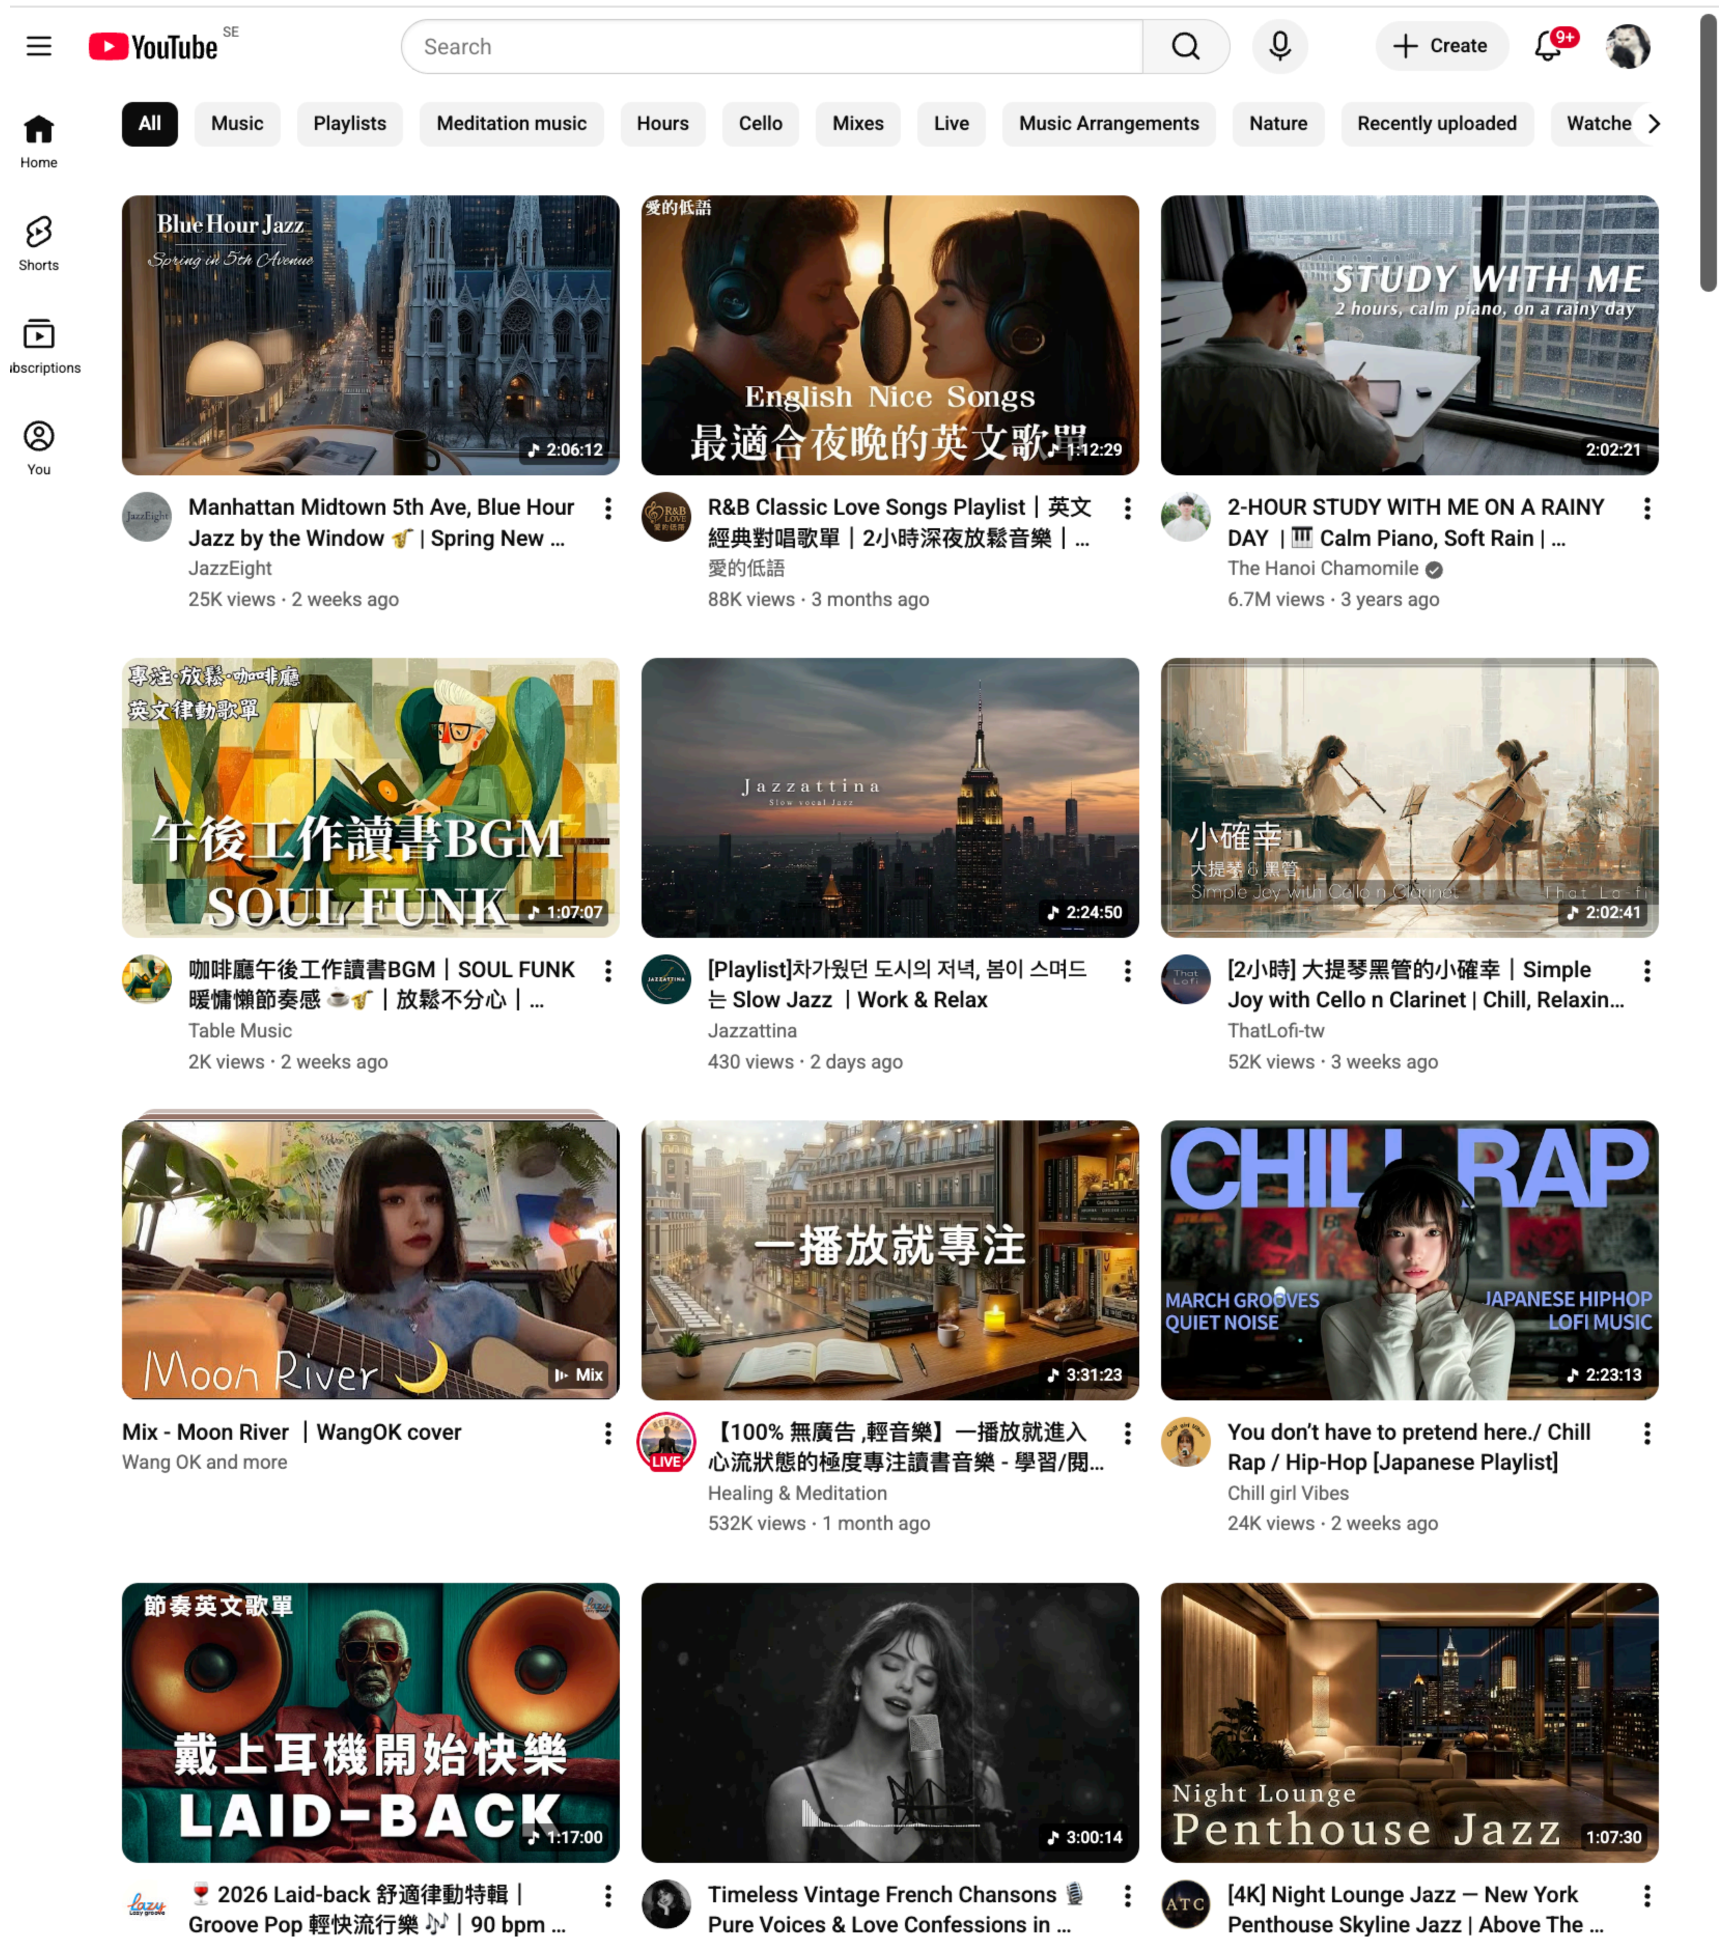

100% music video content with similar genre! I am quite disappointed with this recommendation. 😫

Example: Relevance-only Recommendation

Relevance-only optimization leads to:

- Homogeneous outputs — models repeatedly select similar items
- Filter bubbles — users trapped in narrow content loop
- Popularity bias — long-tail items are systematically over looked

Why diversity matters:

- User studies confirm diverse outputs improve satisfaction without major relevance loss
- Critical in recommender systems (entertainment, news, shopping, etc.)

Case study:

- A non-negative matrix factorization (NMF) model as a recommender system
- Popular mean-squared error (MSE) as training objective
- 60% recommendation are Drama

Relevance-only objective

Rank	Genre
1	Drama
2	Drama
3	Drama
4	Drama, Romance
5	Comedy
6	Documentary
7	Documentary
8	Comedy, Romance
9	Drama
10	Drama, War

6/10 items are Drama!

Problem Setup

Setting: Users $U = \{u_1, \ldots, u_n\}$, items $I = \{i_1, \ldots, i_m\}$, partial observations $\Omega = \{(u, i, \mathbf{R}_{u,i})\}$.

Relevance model F_Θ: predicts $\tilde{\mathbf{R}}_{u,i} = F_\Theta(u, i)$ --- e.g., matrix factorization or MLP.

Standard training minimizes MSE:

$$\Theta^\star = \arg\min_{\Theta} \sum_{(u,i)\in\Omega_T} (\mathbf{R}_{u,i} - \tilde{\mathbf{R}}_{u,i})^2 + \lambda\|\Theta\|_2^2$$

Learn latent embedding for recommendation

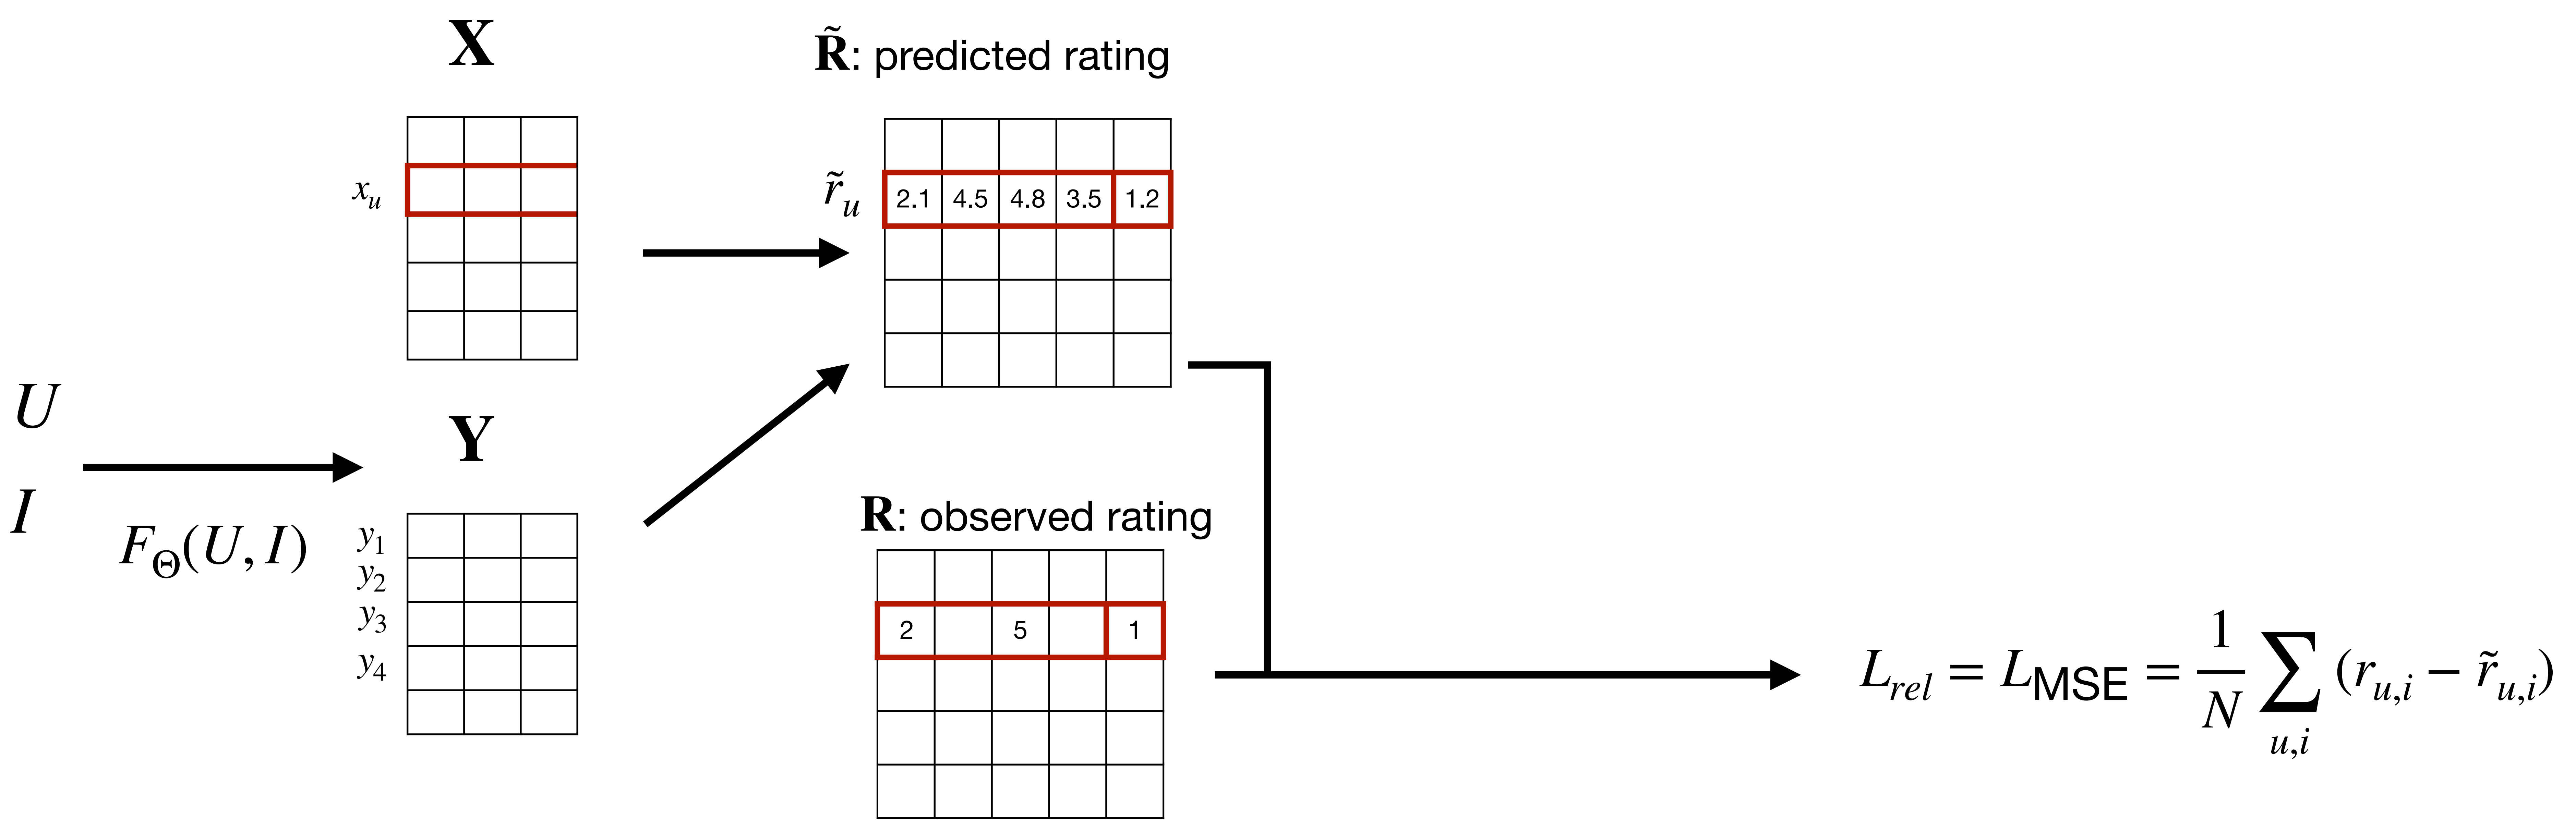

Problem Setup

Given pairwise item affinity $\mathbf{S} \in \mathbb{R}^{m \times m}$ and top-k indicator $\mathbf{l}_u = \mathrm{top}_k(\tilde{\mathbf{r}}_u)$:

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

Diversity reword objective (DRO): $L_{\mathrm{DRO}}(k) = \frac{1}{n} \sum_{u=1}^{n} D_{\mathbf{S}}(Z_u(k))$

Learn latent embedding for recommendation

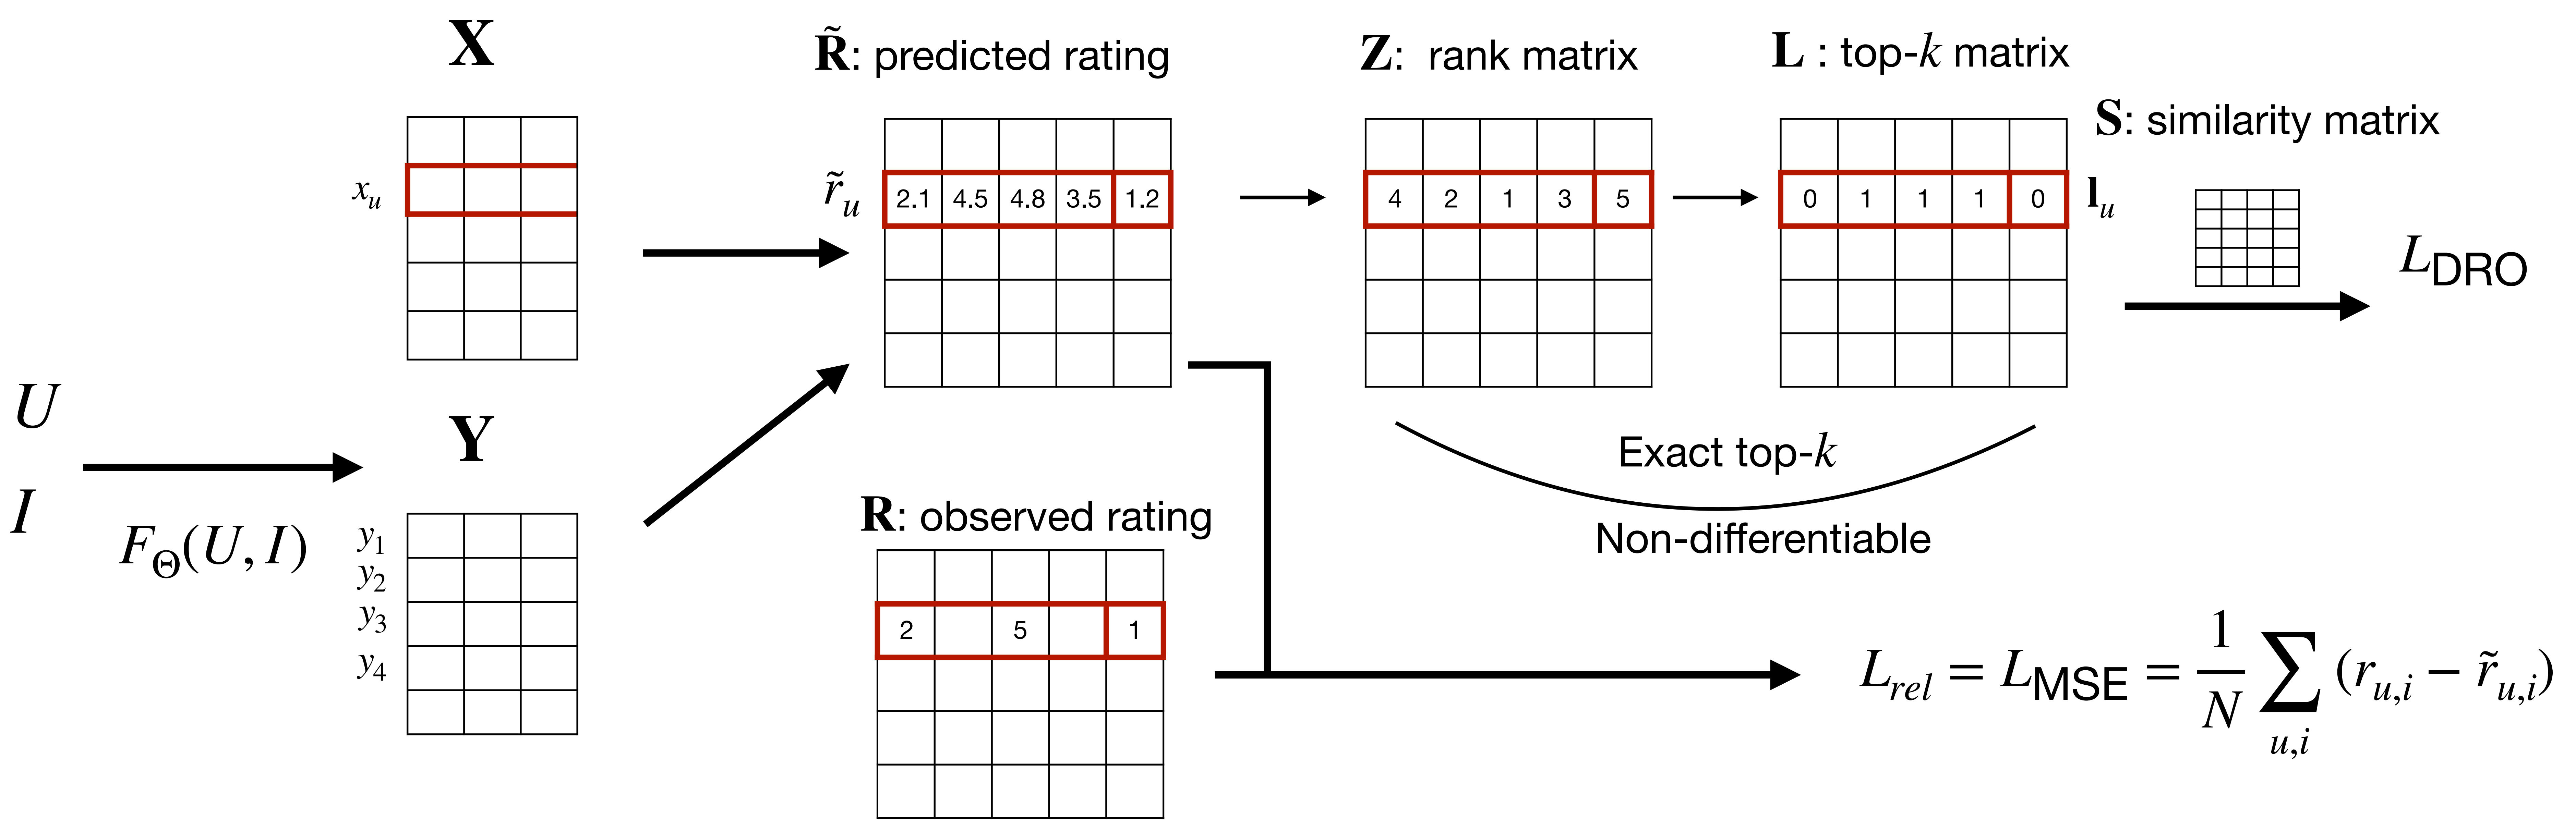

Making Diversity Differentiable

Key technique: Replace discrete ranking with **differentiable soft ranking** (Blondel et al., 2020)

Soft ranking via permutahedron

$$\tilde{\mathbf{z}}_u^{(\epsilon)} = \mathrm{softrank}(\tilde{\mathbf{r}}_u) := \arg\min_{\mathbf{r}\in\mathbf{P}_m} \left\{ -\frac{1}{\epsilon} < \tilde{\mathbf{r}}_u, \mathbf{r} > + H(\mathbf{r}) \right\}$$

- $\mathbf{P}_m$ — the convex hull of all permutations $(1,2,\ldots,m)$ embedded in an m dimensional space
- ϵ — smoothness parameter, as $\epsilon \to 0$: soft ranks $\to$ true ranks
- Complexity — $O(n \log n)$ time, $O(n)$ space
- $\epsilon = 1/n$ yield approximation error almost 0 in DDRO

Soft top-k indicator: $\tilde{\mathbf{I}}_u(i) = \sigma_\tau(k - \tilde{z}_u(i))$, where σ_τ is sigmoid function.

Differentiable Diversity Reward Objective

$$L_{\mathrm{DDRO}} = \frac{1}{n \cdot N} \sum_u \sum_i \sum_j \tilde{\mathbf{I}}_u(i)\, \tilde{\mathbf{I}}_u(j)\, (1 - \mathbf{S}_{i,j})$$

Making Diversity Differentiable

Joint Optimization Problem

Find Θ minimizing:

$$L_{\text{JOINT}}(\beta, \Theta) = \beta L_{\text{rel}} + (1-\beta) L_{\text{div}}$$

- L_{rel}— relevance objective, L_{MSE} or rank-based relevance objective
- $L_{\text{div}} = -L_{\text{DDRO}}$
- $\beta \in [0,1]$: relevance--diversity trade-off

Making Diversity Differentiable

Joint Optimization Problem

Find Θ minimizing:

$$L_{\text{JOINT}}(\beta, \Theta) = \beta L_{\text{rel}} + (1 - \beta) L_{\text{div}}$$

- L_{rel}— relevance objective, L_{MSE} or rank-based relevance objective
- $L_{\text{div}} = -L_{\text{DDRO}}$
- $\beta \in [0,1]$: relevance--diversity trade-off

Contribution

A unified, differentiable, model-agnostic framework that:

- Optimizes relevance & diversity **jointly** during training
- Works with **any** gradient-based model
- Requires **no** architecture changes or post-processing

Two Strategies to Apply Diversity Objective

Direct Diversity-Guided Tuning (DDT)

Idea: Directly optimize L_{JOINT} via gradient descent with the trade-off between diversity and relevance.

Intuition: Diversity objective naturally constrains the learned model to prefer both diversity and relevance.

Challenge: Relevance and diversity gradients can conflict ($\mathbf{g}_{\text{rel}}$ vs. $\mathbf{g}_{\text{div}}$).

Solution: We give the feasible region of that optimize both objective.

- Adaptive — choose optimal β at each step.
- Static — choose static β to balance the trade-off.

Two Strategies to Apply Diversity Objective

Direct Diversity-Guided Tuning (DDT)

Idea: Directly optimize L_{JOINT} via gradient descent with the trade-off between diversity and relevance.

Intuition: Diversity objective naturally constrains the learned model to prefer both diversity and relevance.

Challenge: Relevance and diversity gradients can conflict ($\mathbf{g}_{\text{rel}}$ vs. $\mathbf{g}_{\text{div}}$).

Solution: We give the feasible region of that optimize both objective.

- Adaptive — choose optimal β at each step.
- Static — choose static β to balance the trade-off.

Meta-Diversity Reweighting (MDR)

Idea: Keep standard relevance training, but reweight training samples using diversity as a meta-objective.

Intuition: Redistribute the data distribution that prefer diversity.

Challenge: New parameter $w_{u,i}$ is introduces for each data example during optimization.

Solution: We adopt L_{JOINT} as a meta-objective.

- The outer-objective adjusts weights using meta-objective.
- The inner-objective is relevance-only.

Direct Diversity-Guided Tuning (DDT)

Direct Diversity-Guided Tuning (DDT)

Idea: Directly optimize L_{JOINT} via gradient descent with the trade-off between diversity and relevance.

Intuition: Diversity objective naturally constrains the learned model to prefer both diversity and relevance.

Challenge: Relevance and diversity gradients can conflict ($\mathbf{g}_{\text{rel}}$ vs. $\mathbf{g}_{\text{div}}$).

Solution: We give the feasible region of that optimize both objective.

- Adaptive — choose optimal β at each step.
- Static — choose static β to balance the trade-off.

Feasible Region for Common Descent

Let $a = ||\mathbf{g}_{\text{rel}}||$, $b = ||\mathbf{g}_{\text{div}}||$, $\rho = \dfrac{<\mathbf{g}_{\text{rel}}, \mathbf{g}_{\text{div}}>}{(ab)}$,

for a step along $-\mathbf{g}_\beta$ to decrease both losses:

- Aligned ($\rho > 0$): all $\beta \in [0,1]$ are feasible
- Orthogonal $\rho = 0$: $\beta \in (0,1)$
- Opposing ($\rho < 0$): $\beta \in \left(\dfrac{b|\rho|}{a + b|\rho|}, \dfrac{b}{b + a|\rho|}\right)$

Direct Diversity-Guided Tuning (DDT)

Direct Diversity-Guided Tuning (DDT)

Idea: Directly optimize L_{JOINT} via gradient descent with the trade-off between diversity and relevance.

Intuition: Diversity objective naturally constrains the learned model to prefer both diversity and relevance.

Challenge: Relevance and diversity gradients can conflict ($\mathbf{g}_{\text{rel}}$ vs. $\mathbf{g}_{\text{div}}$).

Solution: We give the feasible region of that optimize both objective.

- Adaptive — choose optimal β at each step.
- Static — choose static β to balance the trade-off.

Feasible Region for Common Descent

Let $a = ||\mathbf{g}_{\text{rel}}||$, $b = ||\mathbf{g}_{\text{div}}||$, $\rho = \dfrac{< \mathbf{g}_{\text{rel}}, \mathbf{g}_{\text{div}} >}{(ab)}$,

for a step along $-\mathbf{g}_\beta$ to decrease both losses:

- Aligned ($\rho > 0$): all $\beta \in [0,1]$ are feasible
- Orthogonal $\rho = 0$: $\beta \in (0,1)$
- Opposing ($\rho < 0$): $\beta \in \left(\dfrac{b|\rho|}{a + b|\rho|}, \dfrac{b}{b + a|\rho|}\right)$

Adaptive $\beta^\star$ via multiple gradient descent algorithm

Maximize the minimum per-step decrease:

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

It guarantees convergence to Pareto stationary points.

Experiments of DDT

Experiment setting

- Neural collaborative filtering (NCF)
- Train from scratch
- Fin-tune relevance-only trained model

Main result:

- $\beta^{\star}$ stays within feasible bounds; shifts from relevance →diversity
- Fine-tuning preserves relevance while boosting diversity
- Train with L_{JOINT} converge to worse stationary point

Adaptive β for optimization

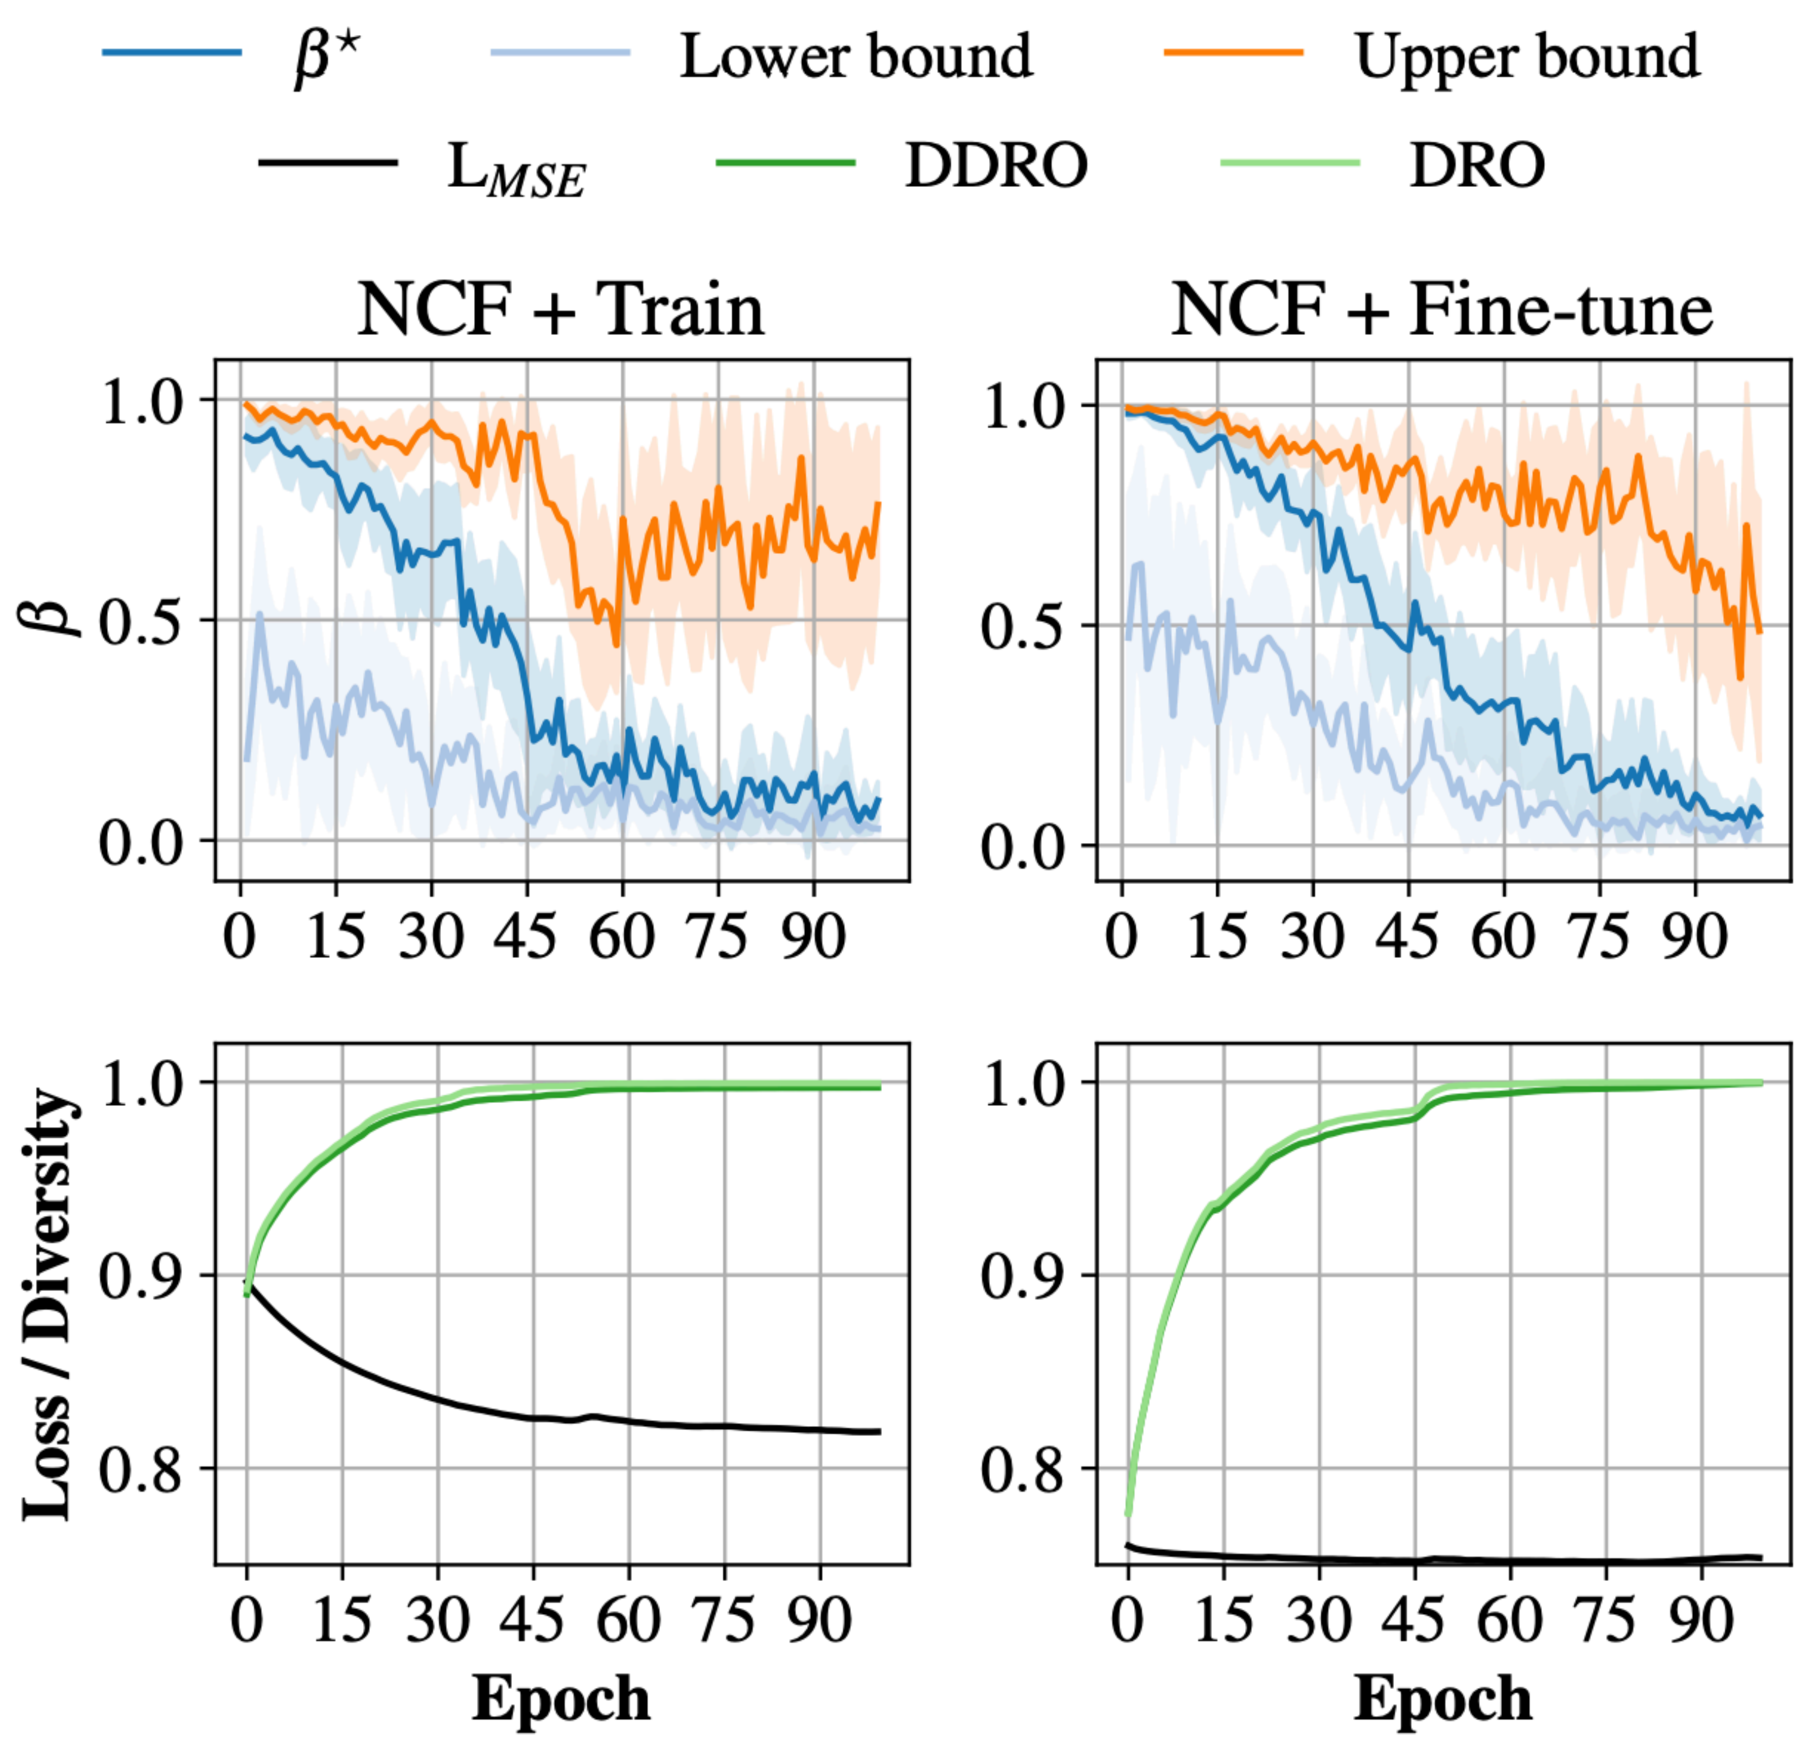

Meta-Diversity Reweight (MDR)

Meta-Diversity Reweighting (MDR)

Idea: Keep standard relevance training, but reweight training samples using diversity as a meta-objective.

Intuition: Redistribute the data distribution that prefer diversity.

Challenge: New parameter $w_{u,i}$ is introduces for each data example during optimization.

Solution: We adopt L_{JOINT} as a meta-objective.

- The outer-objective adjusts weights using meta-objective.
- The inner-objective is relevance-only.

Meta-learning loop

1. Initialize $w_{u,i} = 0$ for all $(u, i) \in B$
2. One-step inner update → temporal model Θ'
3. Compute L_{JOINT} on Θ'
4. Weight gradient: $\tilde{\mathbf{w}} = \max(-\nabla_{\mathbf{w}} L_{\text{JOINT}}, 0)$
5. Normalize: $w_j = \dfrac{\tilde{w}_j}{\sum_k \tilde{w}_k + \epsilon}$
6. Update Θ with reweighted MSE.

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

Example of Training Data Reweighting

Setup: 1 user, 5 items, all ratings = 5.0

Pairwise distance matrix

$$D = \begin{pmatrix} 0 & 1 & 0.5 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 \\ 0.5 & 1 & 0 & 1 & 0.5 \\ 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0.5 & 1 & 0 \end{pmatrix}$$

Cluster structure:

- Cluster A: Items 1 & 5 (dist = 0)
- Cluster B: Items 2 & 4 (dist = 0)
- Bridge: Item 3 (dist 0.5 to A, 1 to B)

MDR weight Computation

Step 1: Initialize $\mathbf{w} = \{0,0,0,0,0\}$

Step 2: One-step meta-update $\Theta \to \Theta'$

Step 3: Gradient $\nabla_{\mathbf{w}} L_{\text{JOINT}}$:

$$[-22449, +1675, -60342, -39423, -35152]$$

Step 4: Rectify: $\tilde{w}_j = \max(-\nabla_{w_j}, 0)$:

$$[22449, 0, 60342, 39423, 35152]$$

Step 5: Normalization

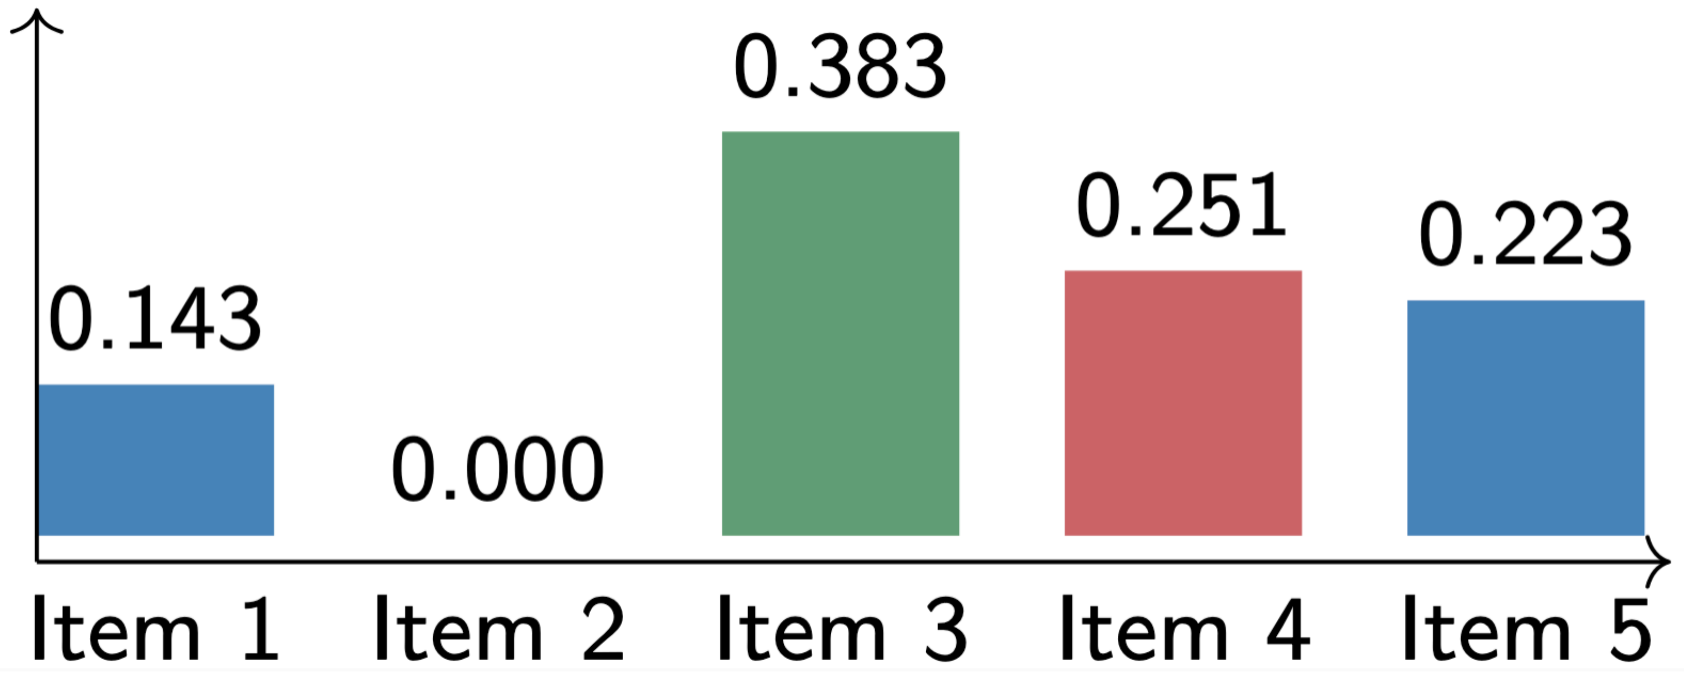

Experiments

Experiment setting:

- Matrix factorization (MF)
- Neural collaborative filtering (NCF)
- Dataset: MovieLens and Netflix
- Static $\beta \in [0,1]$

Main result:

- Recall/precision stable over wide β
- Diversity increases as $\beta \downarrow$
- DDT show strongest performance
- MDR is competitive esp. with simple model

Optimization with varying static β

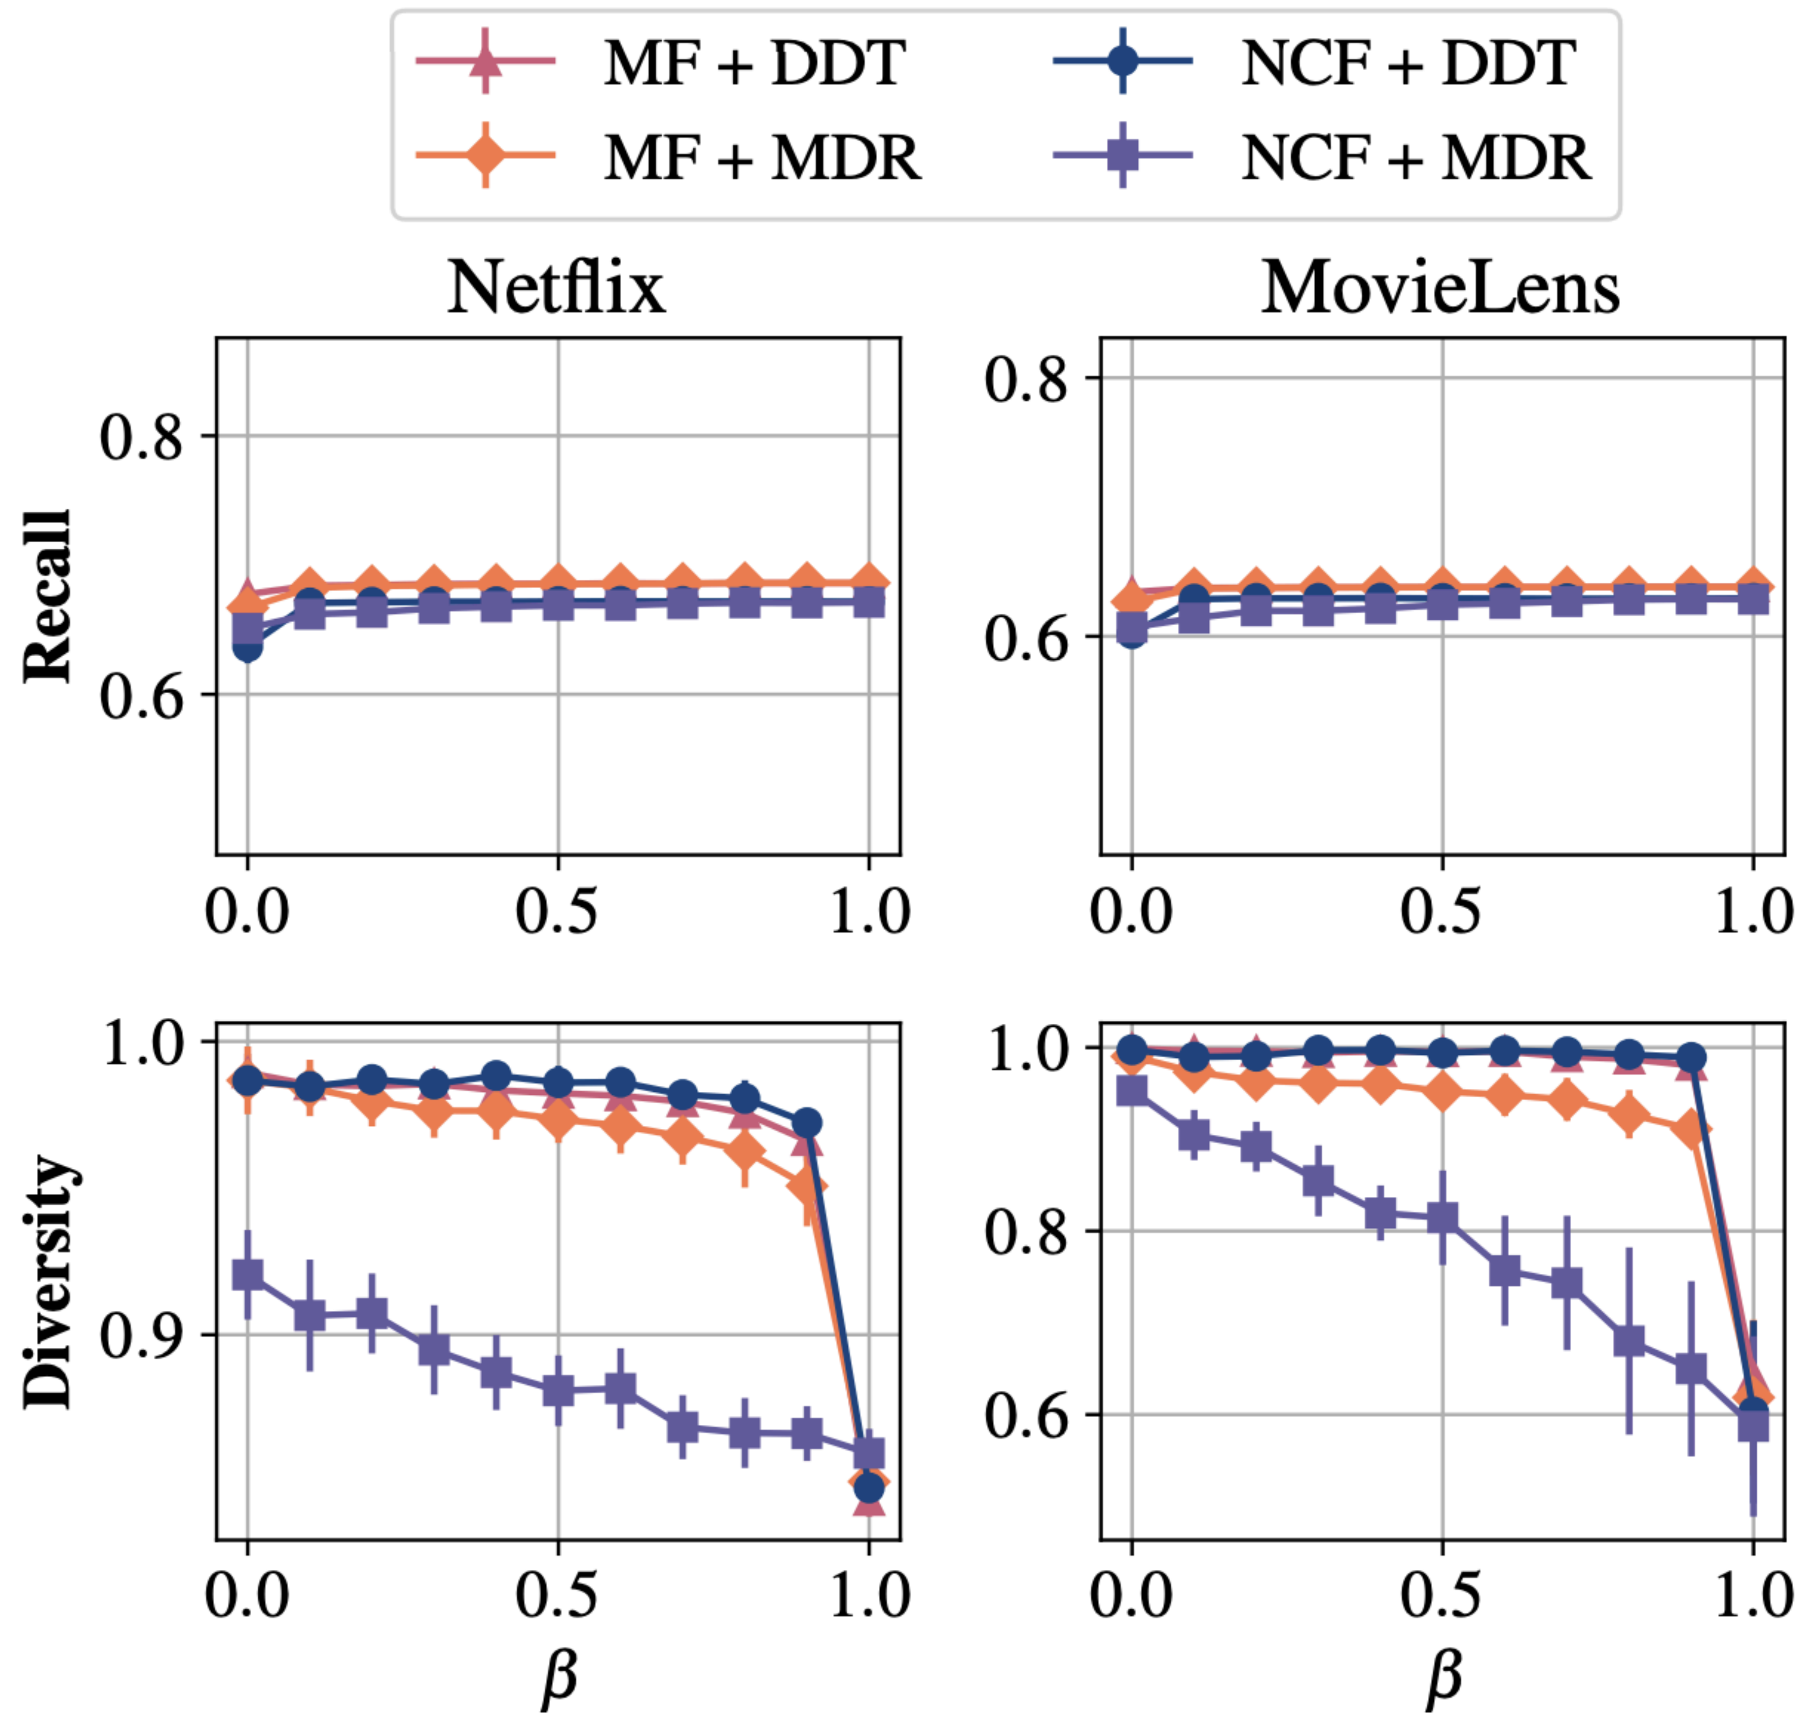

Diversity Gain beyond Objective Optimization

Generality beyond the optimized range

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

Thank you!