Preference Diffusion for Recommendation

ICLR 2025

Shuo Liu[1], An Zhang[2*], Guoqing Hu[3], Hong Qian[1], Tat-Seng Chua[2]

[1]East China Normal University, China [2]National University of Singapore, Singapore

[3]University of Science and Technology of China, China

Contact: anzhang@u.nus.edu, shuoliu@stu.ecnu.edu.cn

Source Code Main Paper

華東師範大學 EAST CHINA NORMAL UNIVERSITY

NUS National University of Singapore

NExT++

中国科学技术大学 University of Science and Technology of China

Overview

Existing Methods:

1. DM-based recommenders inheriting generative objective functions lack a comprehensive understanding of user preference sequences (cf. Figure(a)).
2. DM-based recommenders simply employ standard recommendation training objectives, hindering their generative ability.

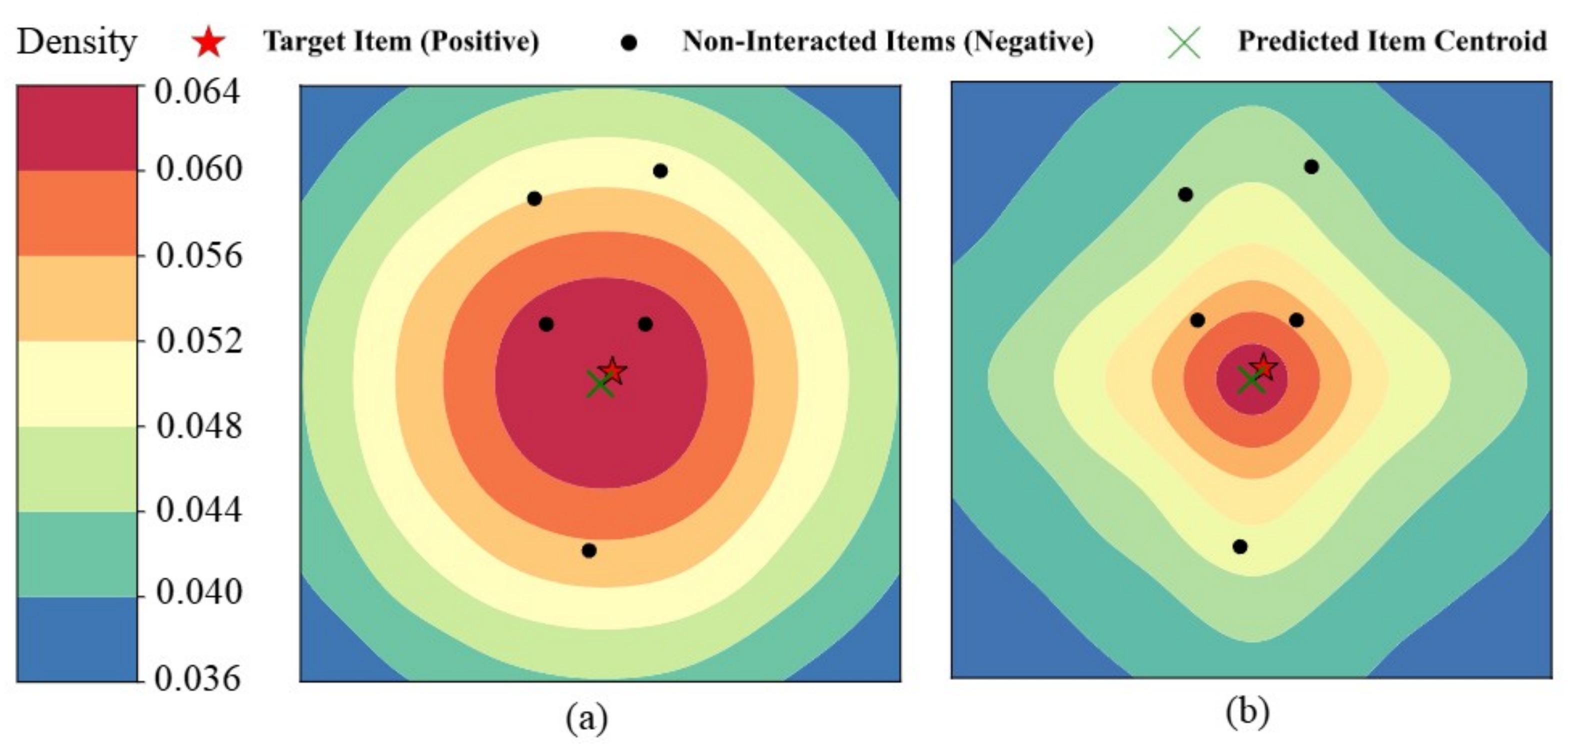

(a) (b)

Question: Can we **preserve the generative power** of diffusion-based recommenders while **fully capture user preference distribution**?

Goal: Propose an objective specifically designed for DM-based recommenders, which effectively integrates negative samples to better capture user preference distributions.

Contribution:

- Transform Bayesian personalized ranking from rating ranking into **log-likelihood ranking**
- Derive a variational upper bound using variational inference, which serves as a **surrogate optimization target** for direct optimization
- Incorporate multiple negative samples to enrich preference signals during training, while employing an efficient strategy to **avoid redundant denoising** from excessive negatives.

The Proposed PreferDiff

1. Rating ranking to log-likelihood ranking

$$\mathcal{L}_{\text{BPR}} = -\mathbb{E}_{(\mathbf{e}_0^+,\mathbf{e}_0^-,\mathbf{c})}\left[\log\sigma\left(f_\theta(\mathbf{e}_0^+ \mid \mathbf{c}) - f_\theta(\mathbf{e}_0^- \mid \mathbf{c})\right)\right]$$

Normalization

$$p_\theta(\mathbf{e}_0 \mid \mathbf{c}) = \frac{\exp(f_\theta(\mathbf{e}_0 \mid \mathbf{c}))}{Z_\theta}$$

$$\mathcal{L}_{\text{BPR-Diff}}(\theta) = -\mathbb{E}_{(\mathbf{e}_0^+,\mathbf{e}_0^-,\mathbf{c})}\left[\log\sigma\left(\log p_\theta(\mathbf{e}_0^+ \mid \mathbf{c}) - \log p_\theta(\mathbf{e}_0^- \mid \mathbf{c})\right)\right].$$

2. Deriving Surrogate Optimization Target

$$\mathcal{L}_{\text{BPR-Diff}}(\theta) = -\mathbb{E}_{(\mathbf{e}_0^+,\mathbf{e}_0^-,\mathbf{c})}\left[\log\sigma\left(\log\mathbb{E}_{q(\mathbf{e}_{1:T}\mid\mathbf{e}_0^+)}\frac{p_\theta(\mathbf{e}_{0:T}^+ \mid \mathbf{c})}{q(\mathbf{e}_{1:T}^+ \mid \mathbf{e}_0^+)} - \log\mathbb{E}_{q(\mathbf{e}_{1:T}\mid\mathbf{e}_0^-)}\frac{p_\theta(\mathbf{e}_{0:T}^- \mid \mathbf{c})}{q(\mathbf{e}_{1:T}^- \mid \mathbf{e}_0^-)}\right)\right]$$

Variational Inference Cosine Error

$$\mathcal{L}_{\text{Upper}}(\theta) = -\mathbb{E}_{(\mathbf{e}_0^+,\mathbf{e}_0^-,\mathbf{c}),\, t\sim\mathcal{U}(1,T)}\left[\log\sigma\left(-\left(S(\hat{\mathbf{e}}_0^+,\mathbf{e}_0^+) - S(\hat{\mathbf{e}}_0^-,\mathbf{e}_0^-)\right)\right)\right].$$

3. Extend to multiple negatives

$$\mathcal{L}_{\text{BPR-Diff-V}} = -\log\sigma\left(-|\mathcal{H}|\cdot\left(S(\hat{\mathbf{e}}_0^+,\mathbf{e}_0^+) - \frac{1}{|\mathcal{H}|}\sum_{v=1}^{|\mathcal{H}|} S(\hat{\mathbf{e}}_0^{-v},\mathbf{e}_0^{-v})\right)\right).$$

Jensen's inequality

$$\mathcal{L}_{\text{BPR-Diff-C}} = -\log\sigma\left(-|\mathcal{H}|\cdot\left[S(\hat{\mathbf{e}}_0^+,\mathbf{e}_0^+) - S\left(\mathcal{F}_\theta(\bar{\mathbf{e}}_t^-,t,\mathcal{M}(\mathbf{c})),\bar{\mathbf{e}}_0^-\right)\right]\right).$$

Main Takeaways

1. Standard Normal Initiation Matters

Embedding Initialization	Sports	Beauty	Toys
Uniform	0.0039/0.0026	0.0013/0.0037	0.0015/0.0011
Kaiming_Uniform	0.0025/0.0019	0.0040/0.0027	0.0051/0.0028
Kaiming_Normal	0.0023/0.0021	0.0049/0.0028	0.0041/0.0029
Xavier_Uniform	0.0011/0.0007	0.0036/0.0021	0.0051/0.0029
Xavier_Normal	0.0014/0.0007	0.0067/0.0037	0.0042/0.0023
Standard Normal	**0.0185/0.0147**	**0.0429/0.0323**	**0.0473/0.0367**

$$\mathbf{E}_0^t = \sqrt{\alpha_t}\,\mathbf{E}_0 + \sqrt{1-\alpha_t}\,\boldsymbol{\epsilon}$$

Variance

$$\mathrm{Var}(\mathbf{E}_0^t) = \alpha_t\,\mathrm{Var}(\mathbf{E}_0) + (1-\alpha_t)\,\mathbf{I}$$

Initialize item embeddings with **standard normal distribution**

2. Large Embedding Size Matters

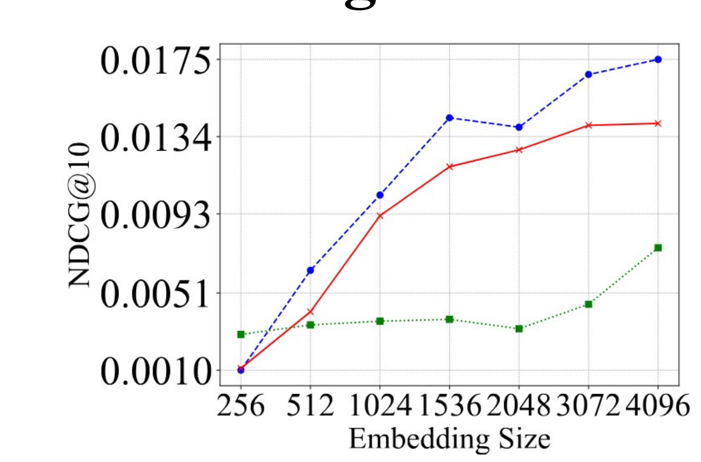

Sports

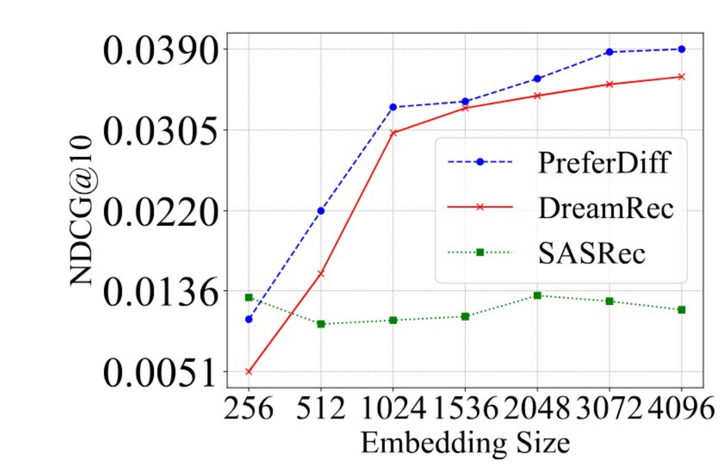

Toys

The identity-like covariance structure suggests that diffusion-based recommenders distribute variance evenly across embedding dimensions, **requiring more dimensions** to capture the complexity and diversity of the item space effectively.

Experimental Results

Model	Sports and Outdoors				Beauty				Toys and Games			
	R@5	N@5	R@10	N@10	R@5	N@5	R@10	N@10	R@5	N@5	R@10	N@10
GRU4Rec	0.0022	0.0020	0.0030	0.0023	0.0093	0.0078	0.0102	0.0081	0.0097	0.0087	0.0100	0.0090
SASRec	0.0047	0.0036	0.0067	0.0042	0.0138	0.0090	0.0219	0.0116	0.0133	0.0097	0.0170	0.0109
BERT4Rec	0.0101	0.0060	0.0157	0.0078	0.0174	0.0112	0.0286	0.0148	0.0226	0.0139	0.0304	0.0163
CL4SRec	0.0105	0.0070	0.0159	0.0085	0.0221	0.0123	0.0345	0.0178	0.0224	0.0142	0.0321	0.0169
TIGER	0.0093	0.0073	0.0166	0.0089	0.0236	0.0151	0.0366	0.0193	0.0185	0.0135	0.0252	0.0156
DiffRec	0.0125	0.0068	0.0200	0.0101	0.0195	0.0121	0.0409	0.0188	0.0268	0.0142	0.0426	0.0193
DreamRec	0.0155	0.0130	0.0211	0.0140	0.0406	0.0299	0.0483	0.0326	0.0440	0.0323	0.0490	0.0353
DiffuRec	0.0093	0.0078	0.0121	0.0087	0.0286	0.0215	0.0335	0.0230	0.0330	0.0262	0.0355	0.0271
MoRec	0.0056	0.0045	0.0076	0.0051	0.0259	0.0189	0.0353	0.0219	0.0154	0.0115	0.0191	0.0127
LLM2BERT4Rec	0.0118	0.0076	0.0183	0.0097	0.0379	0.0262	0.0474	0.0265	0.0339	0.0246	0.0443	0.0263
PreferDiff	**0.0185**	**0.0147**	**0.0247**	**0.0167**	**0.0429**	0.0323	0.0514	0.0350	**0.0473**	**0.0367**	**0.0535**	**0.0387**
PreferDiff-T	0.0182	0.0145	0.0222	0.0158	0.0429	**0.0327**	**0.0532**	**0.0360**	0.0460	0.0351	0.0525	0.0380
Improve	**19.35%**	**16.94%**	**17.06%**	**19.28%**	**5.66%**	**9.36%**	**10.43%**	**7.36%**	**7.50%**	**13.62%**	**9.18%**	**9.63%**

Model	Sports and Outdoors				Beauty				Toys and Games			
	R@5	N@5	R@10	N@10	R@5	N@5	R@10	N@10	R@5	N@5	R@10	N@10
PreferDiff	**0.0185**	**0.0147**	**0.0247**	**0.0167**	**0.0429**	**0.0323**	**0.0514**	**0.0350**	**0.0473**	**0.0367**	**0.0535**	**0.0387**
w/o-N	0.0165	0.0139	0.0214	0.0149	0.0415	0.0304	0.0492	0.0333	0.0445	0.0349	0.0495	0.0367
w/o-C	0.0180	0.0139	0.0230	0.0159	0.0393	0.0282	0.0496	0.0322	0.0458	0.0356	0.0521	0.0374
w/o-C&N	0.0155	0.0130	0.0211	0.0140	0.0406	0.0299	0.0483	0.0326	0.0440	0.0323	0.0490	0.0353

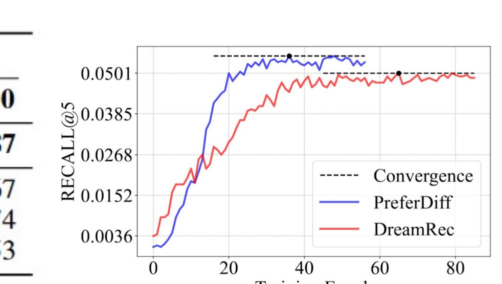

Results:

- PreferDiff significantly outperforms other DM-based recommenders across all metrics on three public benchmarks.
- PreferDiff can benefit from advanced text-embeddings.
- PreferDiff converges faster than other DM-based sequential recommenders.