Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quality-Constrained Preference Fading Model for Discrete Diffusion Recommendation

Weisong Zhang
Yunnan Normal University
Tianwei Xu
xutianwei@ynnu.edu.cn

Yunnan Normal University
Juxiang Zhou
Yunnan Normal University
Bingkun Wang
Yunnan Normal University

Article

Keywords: Discrete Diffusion Models, Recommendation Systems, Preference Fading, Negative Sampling, Top-K Recommendation

Posted Date: July 9th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10217705/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Quality-Constrained Preference Fading Model for Discrete Diffusion Recommendation

Weisong Zhang[1,2], Tianwei Xu[1,2*], Juxiang Zhou[1,2], Bingkun Wang[1,2]

[1*]Key Laboratory of Education Informatization for Nationalities, Ministry of Education, Yunnan Normal University, Juxian Street 768, Chenggong District, Kunming, 650500, Yunnan, China.
[2]Yunnan Key Laboratory of Smart Education, Yunnan Normal University, Juxian Street 768, Chenggong District, Kunming, 650500, Yunnan, China.

*Corresponding author(s). E-mail(s): xutianwei@ynnu.edu.cn;
Contributing authors: zhangweisong@ynnu.edu.cn;
zhoujuxiang@ynnu.edu.cn; wangbingkun@ynnu.edu.cn;

Abstract

In recent years, discrete diffusion models have demonstrated substantial potential in recommendation systems by modeling the recommendation process as a mechanism of "preference fading and growing" within a discrete item space. However, existing methods typically rely on a uniform random sampling strategy for item replacement during the forward preference fading stage. In massive and sparse recommendation spaces, this mechanism tends to generate corrupted preference states that lack informational value. Consequently, the reverse denoising process suffers from inadequate discriminative supervision, which ultimately restricts the model's performance in Top-K recommendation tasks. To address this issue, we propose a Quality-Constrained Preference Fading (QCPF) model. QCPF reconstructs the replacement distribution during the forward fading phase into a quality-aware mixture distribution. This distribution dynamically integrates random negative samples to maintain exploration coverage, in-batch hard candidates to provide discriminative signals, and popularity-based proxy negative samples to reflect exposure characteristics. Concurrently, QCPF introduces a lightweight false-negative-aware filtering mechanism that eliminates invalid replacements through identity and historical interaction checks, thereby enhancing the reliability of the preference degradation trajectory. Serving purely as a forward enhancement module during the training phase, this mechanism

fully preserves the original diffusion transition backbone during inference, introducing no additional parameters or inference latency. Extensive experiments conducted on five real-world public benchmark datasets demonstrate that QCPF significantly improves the Top-K Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG). These results validate the effectiveness of constructing high-quality forward noise in enhancing the performance of discrete diffusion recommendations.

Keywords: Discrete Diffusion Models, Recommendation Systems, Preference Fading, Negative Sampling, Top-K Recommendation

1 Introduction

Recommender systems play a pivotal role in personalized services and online decision-making, and reinforcement learning has also been explored as a way to optimize recommendation policies Lin et al. (2023); Chen et al. (2025); Wu et al. (2022). However, real-world recommendation scenarios typically operate within extremely large discrete item spaces, where systems can only acquire highly sparse implicit feedback data Liang et al. (2018); Chen et al. (2023). Learning highly discriminative ranking signals from a candidate set inundated with unobserved interactions remains a core challenge for Top-K recommendation algorithms Wu et al. (2021); Wang et al. (2021).

In recent years, the rapid advancement of generative artificial intelligence has propelled a paradigm shift in recommendation methodologies, transitioning from traditional retrieval-based paradigms to generative ones Deldjoo et al. (2024); Wu et al. (2024). Among these, diffusion models have demonstrated remarkable potential in handling sparsity and capturing complex preference structures through forward noise injection and reverse denoising Ho et al. (2020); Wang et al. (2023); Li et al. (2023). As systematically summarized in a recent comprehensive survey on diffusion models for recommendation Wei and Fang (2026), diffusion models are profoundly reshaping the landscape of core tasks like collaborative filtering and sequential recommendation. This is attributed to their advantages in fitting complex distributions and providing highly flexible denoising backbones. Beyond the recommendation setting itself, diffusion-based generative learning has also been explored in GANs, latent diffusion, denoising diffusion implicit models, score-based generative modeling, score matching, continuous-time denoising, and discrete-data score matching formulations Goodfellow et al. (2014); Rombach et al. (2022); Song et al. (2022); Lim et al. (2023); Song and Ermon (2019); Campbell et al. (2022); Meng et al. (2022). Traditional diffusion recommendation models mostly follow the continuous diffusion approach, injecting continuous Gaussian noise into dense item embeddings or preference vectors Yang et al. (2023); Wang et al. (2023); Li et al. (2023). However, recommendation data is inherently discrete, and continuous Gaussian perturbations often exhibit an intrinsic mismatch with the item-level discrete ranking objectives Austin et al. (2021); Benigni et al. (2026). To this end, recent discrete diffusion recommendation frameworks (e.g., PreferGrow Hu et al. (2025)) break this paradigm by formulating the recommendation

process directly as a "fading and growing" mechanism over discrete items. In the forward diffusion process, a user's target positive items gradually fade, being replaced by other candidates or transitioning into a non-preference state. In the reverse denoising process, the model learns to recover the original preferences from this corrupted state by estimating relative preference ratios. In parallel, diffusion has also been extended to other recommendation scenarios such as cross-domain recommendation, multimedia recommendation, generative outfit recommendation, and multi-interest generative recommendation Xuan (2024); Deldjoo et al. (2020); Xu et al. (2024); Li et al. (2025). In collaborative filtering, diffusion has also been adapted to model higher-order connectivity and binomial corruption processes for recommendation data, showing that discrete or structured diffusion can better align with the combinatorial nature of user-item interactions Hou et al. (2024); B'en'edict et al. (2023).

Although discrete diffusion-based recommendation methods have achieved initial success, existing research has largely overlooked the "degradation quality" during the forward preference fading stage Wei and Fang (2026); Lin et al. (2024). Current discrete diffusion models typically employ a simple uniform random sampling strategy for item replacement Hu et al. (2025); Lee et al. (2024). In extremely sparse and massive candidate spaces, randomly sampled replacement items often deviate significantly from the target user's true interests, yielding limited informational value Lee et al. (2024); Mansoury et al. (2024). Such an untargeted preference degradation trajectory implies that the reverse generative model only needs to rely on weak discriminative signals to accomplish the distinction task, thereby failing to provide sufficient supervisory evidence for fine-grained Top-K ranking Mansoury et al. (2024); Choi et al. (2026). Exposure-aware recommendation research suggests that using high-exposure but unclicked items can help reduce false-negative risk, while contrastive and hard-negative sampling strategies can strengthen discriminative learning Mansoury et al. (2024); Lee et al. (2024); Choi et al. (2026); He et al. (2025). Meanwhile, recent diffusion and discrete-diffusion studies show that the quality of the forward corruption trajectory can materially affect downstream performance, especially when the denoising process is guided by more informative corrupted states Mao et al. (2025); Lou et al. (2023); Campbell et al. (2022). In addition, recent studies on diffusion-based sequential recommendation and diffusion augmentation further indicate that semantically guided corruption or augmentation can preserve more useful contextual signals than random perturbation Xie et al. (2024); Liu et al. (2023); Choi et al. (2026).

To address the technical deficiency of uninformative corrupted preference states caused by uniform random replacement in discrete diffusion models, this paper proposes a Quality-Constrained Preference Fading (QCPF) method. QCPF overcomes the limitations of a single uniform distribution by reconstructing the replacement distribution during the forward fading phase into a quality-aware mixed distribution Lee et al. (2024); Choi et al. (2026). Specifically, this mixed distribution hierarchically integrates three types of negative evidence: simple random negative samples to maintain spatial exploration coverage, batch-level hard negatives to provide strong discriminative ranking supervision, and popularity-based exposure proxy metrics to reflect implicit feedback characteristics He et al. (2025); Lee et al. (2024); Chen et al. (2023). Furthermore, to prevent the "false-negative" misleading effect that occurs when a high-quality

replacement item coincidentally hits the user's true preference, this paper introduces a lightweight false-negative-aware filtering mechanism Wu et al. (2021); Chen et al. (2023); Mansoury et al. (2024). This design is consistent with prior work on exposure bias and other safety-oriented recommendation strategies that distinguish true negatives from merely unobserved interactions. By dynamically eliminating invalid replacements through identity and historical interaction filtering and re-normalizing over the valid candidate set, it prevents the risk of model collapse under extreme noise distributions. Related ideas also appear in contrastive augmentation and semantic-guided diffusion, where better sample construction improves representation learning stability Choi et al. (2026); Liu et al. (2023); Xie et al. (2024).

It is worth emphasizing that the QCPF mechanism serves purely as a data augmentation and forward fading optimization module during the training phase. It fully preserves the Markov state transition matrix and the reverse preference score generation flow of the discrete diffusion backbone (e.g., PreferGrow). Consequently, without introducing any additional inference parameters or computational overhead, this method generates corrupted preference states characterized by both high informational value and high reliability.This separation between training-time corruption design and inference-time generation also echoes broader diffusion and score-based generative modeling practice, where the forward process can be customized without changing the core sampling backbone Ho et al. (2020); Wang et al. (2023). Similar training-only perturbation ideas have also been explored in DDIM-style and other discrete-state diffusion frameworks Song et al. (2022); Austin et al. (2021).

The main contributions of this paper are summarized as follows:

- **Proposed the QCPF quality-constrained mechanism:** We design a hybrid fading strategy that integrates multi-source data—namely, simple negative samples, in-batch hard negative samples, and popularity-based proxy negative samples. This strategy significantly enhances the informational value of the preference degradation trajectories.
- **Introduced a false-negative-aware constraint:** While ensuring the reverse generative model architecture remains unaltered and imposing zero inference burden, we implement a strict filtering mechanism that effectively mitigates the risks of noisy supervision caused by invalid replacements.
- **Extensive empirical validation:** We conduct comprehensive Top-K main experiments on five public recommendation benchmark datasets, alongside detailed ablation studies and hyperparameter sensitivity tests specifically on the MovieLens (ML1M) dataset. The results demonstrate that QCPF achieves consistent and significant performance improvements in terms of Hit Ratio (HR) and Normalized Discounted Cumulative Gain (NDCG) under Top-K ($K = 5, 10, 20$) scenarios.

2 Related Work

This section reviews existing literature from three dimensions: sequential and implicit feedback recommendation, diffusion-based recommendation algorithms, and negative sampling strategies in recommender systems.

2.1 Traditional and Sequential Recommendation

Sequential recommendation aims to capture a user's dynamic preferences and predict future interactions by analyzing their chronologically ordered historical interaction sequences Chen et al. (2025); Wang et al. (2021). Early research predominantly relied on Markov Chains or Recurrent Neural Networks (RNNs, such as GRU4Rec Hidasi et al. (2015)) to model short-term transition patterns in user behavior sequences Tang and Wang (2018). With the widespread adoption of the self-attention mechanism, Transformer-based architectures (e.g., SASRec Kang and McAuley (2018) and BERT4Rec Sun et al. (2019)) have gradually dominated the sequential recommendation domain, as they can more effectively capture long-range dependencies and sequential features Hu et al. (2025); Wang et al. (2021); Li et al. (2023). Subsequent studies further incorporated techniques such as contrastive learning and graph-augmented representations to boost performance Wu et al. (2021, 2022); Choi et al. (2026). However, despite the significant success of these traditional methods in feature extraction, they typically only learn deterministic preference representations from highly sparse implicit feedback, making it difficult to effectively characterize the uncertainty of user preferences in massive candidate spaces Li et al. (2023); Wang et al. (2023). Variational autoencoders have also been used for collaborative filtering, but they still focus on deterministic or latent probabilistic preference representation rather than diffusion-style iterative corruption and recovery Liang et al. (2018).

2.2 Diffusion-based Recommendation Algorithms

Recently, due to their powerful capability in generating complex data distributions and their training stability, diffusion models have been widely introduced into recommender systems, propelling the transition from retrieval paradigms to generative paradigms Lin et al. (2024); Mao et al. (2025); Deldjoo et al. (2024). According to the comprehensive survey by Wei and Fang Wei and Fang (2026), existing works can be deconstructed along three orthogonal axes: "core recommendation tasks" (e.g., collaborative filtering, sequential recommendation), "multi-modal/multi-domain environments," and "trustworthy constraints" (e.g., fairness, robustness). Diffusion models primarily serve two practical paradigms in these tasks: "data augmentation" or "direct generative recommendation" Wei and Fang (2026), successfully capturing structured transitions and combinatorial interactions even in traditional collaborative filtering Hou et al. (2024); B'en'edict et al. (2023). Recent survey work further highlights that diffusion-based recommendation research has expanded beyond single-domain settings into a broader landscape of sequential, multimodal, and trustworthy recommendation scenarios Wu et al. (2024).

Early diffusion recommendation models mostly belonged to the Continuous Diffusion school. These methods typically encode target items or user sequences into dense embeddings and inject continuous Gaussian noise into them (e.g., DreamRec Yang et al. (2023) and DiffuRec Li et al. (2023)). Other approaches apply Gaussian noise over global user-item preference scores, i.e., the probability simplex (e.g., DiffRec Wang et al. (2023)). However, inputs and outputs in recommender systems are inherently discrete Austin et al. (2021). The intrinsic mismatch between continuous perturbation

mechanisms and discrete ranking objectives often causes the generated item representations to deviate from the true semantic space Benigni et al. (2026); Austin et al. (2021).

To tackle this, Discrete Diffusion models, which perturb and recover directly within the discrete space, have become a research hotspot Austin et al. (2021); Xie et al. (2024). For instance, the state-of-the-art framework PreferGrow proposed a matrix-form discrete diffusion framework that formulates recommendation as a "fading and growing" process. Positive items are retained, replaced, or faded to a non-preference state over diffusion timesteps according to transition probabilities, thereby bypassing the rigid assumptions of continuous models on Gaussian noise priors. However, current discrete diffusion recommendation algorithms predominantly focus on designing reverse denoising networks and optimizing multi-interest/quantization modules Wei and Fang (2026); Mao et al. (2025), while lacking an in-depth investigation into the quality of the "replacement distribution" during the forward preference fading stage. When facing massive and sparse item pools, uniform random replacement is highly prone to producing degradation trajectories devoid of informational value Chen et al. (2025); Lee et al. (2024).

2.3 Negative Sampling in Recommender Systems

In implicit feedback recommendation, negative sampling is critical for model learning Liang et al. (2018); Lee et al. (2024). Since unobserved user-item pairs cannot simply be equated to true negative samples that the user dislikes, constructing high-quality negative evidence profoundly impacts model performance Chen et al. (2023); Lee et al. (2024). Classical pairwise learning methods, such as Bayesian Personalized Ranking (BPR), rely on contrasting a positive item with a randomly sampled negative item Wu et al. (2021). Subsequent research has demonstrated that more sophisticated negative sampling strategies (e.g., popularity-biased sampling, hard negative mining, dynamic sampling) can significantly enhance a model's discriminative ability Lee et al. (2024); He et al. (2025). Exposure-aware studies also point out that unobserved items are not missing at random; reasonably utilizing items with high exposure but no clicks can mitigate the risk of false negatives Mansoury et al. (2024).

Although the proposed QCPF mechanism is closely related to the ideology of negative sampling, it fundamentally differs in its operational mechanism. Traditional negative sampling is typically employed merely as an auxiliary training technique to construct point-wise or pair-wise ranking loss functions Wu et al. (2021). In contrast, QCPF directly integrates simple random negative samples, batch-level hard candidates, and popularity-based exposure proxies as multi-level negative evidence into the forward state transition process of discrete diffusion Hu et al. (2025); Lee et al. (2024). This means that the corrupted state constructed by quality-constrained sampling inherently serves as the input to the reverse preference growing model Ho et al. (2020). By elevating the quality of the corrupted state, QCPF enables the contaminated preference state to provide more discriminative and informative supervision targets for reverse denoising, compelling the model to more accurately learn the relative preference ratios among candidate items Lou et al. (2023); Hu et al. (2025).

3 Methodology

3.1 Overall Framework Overview

The Quality-Constrained Preference Fading (QCPF) model for discrete diffusion recommendation is built upon a bi-directional Markov chain framework of "preference fading and growing" within a discrete space. Its core idea is to reconstruct the perturbation trajectory during the forward diffusion phase to enhance the discriminability and informational value of noisy items, without altering the architecture of the reverse denoising generative model, as illustrated in Figure 1.

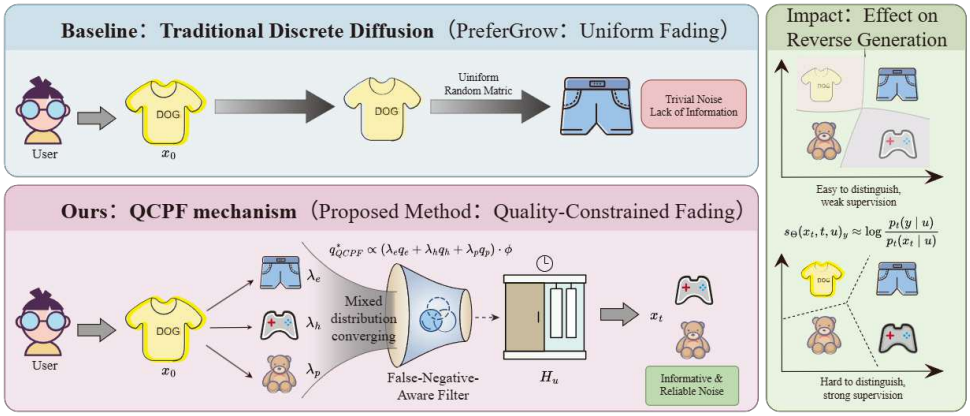


Fig. 1 Comparison of preference fading strategies in discrete diffusion recommenders. The top panel illustrates traditional uniform fading (e.g., PreferGrow), where random sampling tends to generate uninformative noise. The bottom panel presents the proposed QCPF mechanism, which provides a high-quality degradation trajectory with strong discriminative supervision for reverse generation by integrating multi-source high-quality negative samples and applying false-negative filtering.

The overall pipeline of QCPF is divided into two phases: forward quality-constrained preference fading and reverse preference ratio growing. In the forward phase, given a user u's historical interaction sequence $\mathcal{H}_u$ and the target positive item x_0, the system constructs a high-quality corrupted preference state x_t by dynamically integrating multi-level negative evidence and applying a false-negative-aware filtering mechanism. In the reverse phase, the denoising network takes the corrupted state x_t, the user context, and the diffusion timestep t as inputs. It utilizes the Score Entropy Loss to learn the relative preference ratios in the discrete space, ultimately generating a Top-K recommendation list through preference growing over reverse timesteps. Because QCPF exclusively optimizes the replacement distribution during the forward training phase, it introduces zero additional parameters or computational overhead during the inference phase.

3.2 Construction of Quality-Constrained Replacement Distribution

To overcome the limitation of traditional discrete diffusion models, which employ a single uniform distribution during the forward phase, leading to an overly simplistic noise space, QCPF reconstructs the forward item replacement distribution into a quality-aware distribution. This distribution is formed by a weighted combination of multi-level negative evidence, as illustrated in Figure 2.

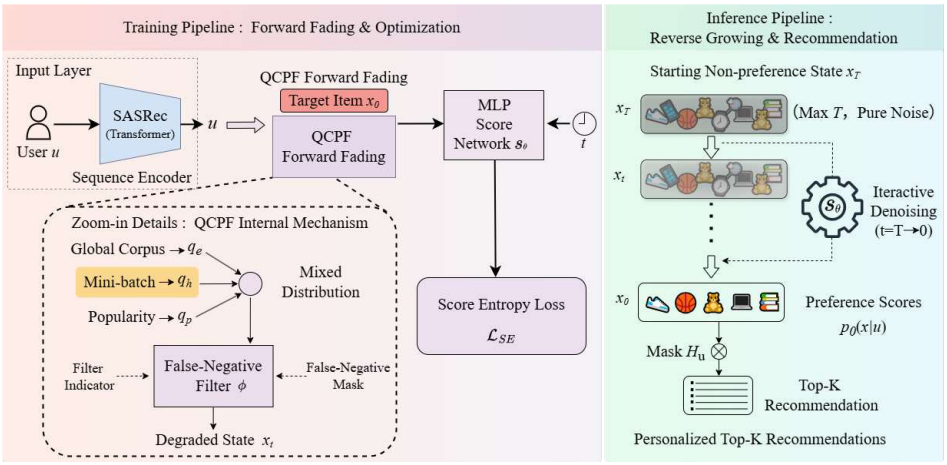


Fig. 2 Overall architecture and bi-directional diffusion process of QCPF. In the forward diffusion phase, the model utilizes a dynamic mixed distribution and parallel mask filtering to generate a high-quality corrupted state x_t. In the reverse phase, the denoising network accurately recovers the target positive preference x_0 based on the context containing highly deceptive noise.

3.2.1 Fusion of Multi-Level Negative Evidence

Let the candidate replacement item source set be defined as:

$$\mathcal{M} = \{easy, hard, pop\} \tag{1}$$

where each component captures negative supervision signals from a distinct dimension. The mathematical expression of the initial mixed replacement distribution $q_{QCPF}(i \mid u, x_0)$ is formulated as follows:

$$q_{QCPF}(i \mid u, x_0) = \sum_{m \in \mathcal{M}} \lambda_m q_m(i \mid u, x_0) = \lambda_e q_e(i) + \lambda_h q_h(i \mid \mathcal{B}) + \lambda_p q_p(i \mid \mathcal{B}) \tag{2}$$

where $i \in \mathcal{I}$ denotes a candidate replacement item from the item corpus $\mathcal{I}$, and λ_e, λ_h, λ_p are the corresponding weighted mixing coefficients for each negative sample source, satisfying $\lambda_m \geq 0$ and $\sum_{m \in \mathcal{M}} \lambda_m = 1$.

The specific construction rules for the three candidate distributions are as follows:

1. **Simple Random Negative Samples (q_e):** Uniformly sampled from the global item corpus $\mathcal{I}$. The mathematical expression is:

$$q_e(i) = \frac{1}{|\mathcal{I}|} \tag{3}$$

   This component is designed to maintain the overall exploration coverage of the diffusion process, preventing the degradation trajectory from collapsing into local regions.
2. **Batch-level Hard Candidates (q_h):** Drawn from the deduplicated set of items $\mathcal{B}_{unique}$ within the current training mini-batch $\mathcal{B}$. The expression is:

$$q_h(i \mid \mathcal{B}) = \frac{1}{|\mathcal{B}_{unique}|} \sum_{x \in \mathcal{B}_{unique}} \delta(i = x) \tag{4}$$

   where $\delta(\cdot)$ is the indicator function that equals 1 if and only if the condition is met, and 0 otherwise. These items reflect the genuine interests of other concurrent users, posing a strong deception to the current user and providing highly discriminative ranking supervision signals.
3. **Popularity-based Exposure Proxy Negative Samples (q_p):** Utilizing the occurrence frequency of items within the current mini-batch $\mathcal{B}$ as a lightweight exposure proxy metric. It is defined as:

$$q_p(i \mid \mathcal{B}) = \frac{freq_{\mathcal{B}}(i)^{\alpha}}{\sum_{j \in \mathcal{B}} freq_{\mathcal{B}}(j)^{\alpha}} \tag{5}$$

   where $freq_{\mathcal{B}}(i)$ denotes the frequency of item i in the current batch, and α is the frequency smoothing coefficient. This component prioritizes the sampling of high-exposure but unclicked items during forward fading, narrowing the gap between the generated noise and the true high-exposure distribution.

3.2.2 False-Negative-Aware Filtering and Normalization

More discriminative hard replacement candidates (e.g., popular items or in-batch hard negatives) provide rich information, but they are also highly susceptible to mistakenly sampling potential positive items that the user genuinely likes but have not yet observed (i.e., generating false-negative errors). To address this, QCPF introduces a lightweight validity indicator function $\phi(i, u, x_0)$ for active constraint filtering:

$$\phi(i, u, x_0) = \mathbb{I}(i \neq x_0) \cdot \mathbb{I}(i \notin \mathcal{H}_u) \tag{6}$$

where $\mathbb{I}(\cdot)$ is the indicator function. This mechanism simultaneously satisfies two filtering conditions:

- **Identity Filtering ($\mathbb{I}(i \neq x_0)$):** Removes items that are exactly identical to the current target positive item x_0, completely eliminating meaningless identity state transitions.

- **Historical Interaction Filtering ($\mathbb{I}(i \notin \mathcal{H}_u)$):** Eliminates all pre-existing items in the user's historical interaction sequence $\mathcal{H}_u$, avoiding the misclassification of known preferences as invalid noise.

After the false-negative-aware filtering, the system re-normalizes the mixed replacement distribution over the remaining valid candidate item set, yielding the final high-quality replacement distribution $q^*_{QCPF}(i \mid u, x_0)$:

$$q^*_{QCPF}(i \mid u, x_0) = \frac{q_{QCPF}(i \mid u, x_0)\phi(i, u, x_0)}{\sum_{j \in \mathcal{I}} q_{QCPF}(j \mid u, x_0)\phi(j, u, x_0)} \tag{7}$$

To ensure efficient convergence during the training phase, the aforementioned false-negative-aware filtering mechanism employs fully parallelized vectorized tensor operations at the code implementation level, bypassing traditional iterative retrieval. Specifically, the historical interaction records $\mathcal{H}_u$ are converted into a static sparse lookup tensor, allowing the validity indicator $\phi(i, u, x_0)$ to be computed via highly efficient sparse-dense tensor operations and batch-level tensor broadcasting on GPUs. Furthermore, the summation over the massive item corpus $\mathcal{I}$ in the denominator of Eq. (7) is accelerated through optimized parallel reduction primitives, which avoids explicit sequential traversals of the item space. This engineering design achieves an effective time complexity of $\mathcal{O}(1)$ per batch element, completely eliminating the need for redundant set operations or CPU-GPU data transfers. Experimental data indicates that this filtering process introduces minimal computational overhead and barely alters the training throughput of the original diffusion model. Consequently, it perfectly achieves the engineering goal of zero inference overhead while safeguarding the quality of supervision signals.

3.3 Forward Preference Fading Perturbation Process

After constructing the high-quality replacement distribution, QCPF integrates it into the Markov state transition framework of discrete diffusion. At any diffusion timestep $t \in [1, T]$, the forward preference fading process can be abstractly represented as:

$$q_{QCPF}(x_t \mid x_0, u) = (1-\rho_t)\delta(x_t = x_0) + \rho_t \left[\gamma q^*_{QCPF}(x_t \mid u, x_0) + (1-\gamma)\delta(x_t = [NP])\right] \tag{8}$$

where:

- $\rho_t \in [0, 1]$ denotes the total state transition probability controlled by the diffusion Noise Schedule. It gradually increases as timestep t progresses from $0 \to T$, causing the original preference signal to progressively degrade towards a non-preference state.
- $\gamma \in [0, 1]$ is a hyperparameter balancing coefficient used to precisely regulate the relative probability of the positive item being "replaced by an item sampled from the replacement distribution q^*_{QCPF}" versus "transitioning to a Non-preference Absorbing State $[NP]$".
- $\delta(\cdot)$ is the state indicator function, which equals 1 when the condition in the parentheses is satisfied, and 0 otherwise.

Through this transition framework, QCPF does not alter the original structural backbone of the diffusion model. Instead, by introducing multi-level negative evidence and false-negative constraints, it ensures that the corrupted state x_t carries more challenging information during the training phase. This consequently constructs a high-quality preference degradation trajectory for the reverse network.

3.4 Reverse Preference Recovery and Model Optimization

The objective of the reverse preference growing generation is to train a neural network $s_\Theta(x_t, t, u)$ capable of accurately recovering the user's original preference x_0 from a corrupted state x_t that contains high-quality noise.

Based on the relative ratio modeling principle of discrete diffusion, the network does not directly predict the absolute click probability of isolated items. Instead, it learns to fit the marginal log preference ratios between any candidate item y and the current corrupted state x_t:

$$s_\Theta(x_t, t, u)_y \approx \log \frac{p_t(y \mid u)}{p_t(x_t \mid u)}, \quad \forall y \in \mathcal{I} \tag{9}$$

In the specific implementation, we employ a self-attention model (e.g., SASRec) as the sequence encoder to extract the context vector representation $\mathbf{u} = \mathrm{SeqRec}(\mathcal{H}_u)$ from the historical sequence $\mathcal{H}_u$. Subsequently, a Multi-Layer Perceptron (MLP) is utilized to map and interact the item embeddings, timestep vector, and context vector:

$$s_\Theta(x_t, t, u)_y = \mathbf{e}_y^\top \mathrm{MLP}(\mathrm{concat}(\mathbf{e}_{x_t}, \mathbf{t}, \mathbf{u})) \tag{10}$$

where $\mathbf{e}_y$ and $\mathbf{e}_{x_t}$ represent the low-dimensional dense vector embeddings for item y and the corrupted state x_t, respectively.

The training and optimization of the model parameters Θ adopt the Score Entropy Loss under the discrete diffusion framework:

$$\mathcal{L}_{SE} = \mathbb{E}_{x_0}\mathbb{E}_t\mathbb{E}_{x_t} \sum_{y \in \mathcal{I}} Q_t(x_t, y) \cdot l_{SE}(x_0, x_t, y \mid u) \tag{11}$$

$$l_{SE}(x_0, x_t, y \mid u) = e^{s_\Theta(x_t,t,u)_y} - e^{r_t(x_0,x_t,y)} s_\Theta(x_t, t, u)_y + e^{r_t(x_0,x_t,y)}[r_t(x_0, x_t, y) - 1] \tag{12}$$

where Q_t is the derivative rate matrix corresponding to the transition probability matrix; and $r_t(x_0, x_t, y) = \log \frac{p_{t|0}(y|x_0,u)}{p_{t|0}(x_t|x_0,u)}$ represents the true forward reference log ratio under the current degraded state.

Because the forward reference trajectory constructed by QCPF integrates in-batch hard samples and high-exposure popular items, it forces the denoising network s_Θ to learn how to overcome these deeply deceptive noises during reverse parameter optimization. This compels the model to meticulously distinguish the relative ranking ratios between the target item and various high-quality candidate items. Consequently, this significantly enhances the Top-K fine-grained ranking precision of the ultimately learned preference score vectors in the massive and sparse space.

4 Experiments

To comprehensively evaluate the effectiveness of the proposed Quality-Constrained Preference Fading (QCPF) model, we conduct extensive experiments on five real-world public recommendation datasets, aiming to answer the following core research questions (RQs):

- **RQ1:** Compared with existing sequential recommendation algorithms and diffusion-based recommendation models, can QCPF achieve superior Top-K ranking performance?
- **RQ2:** How much do the individual data source modules (e.g., batch-level hard candidates, popularity proxy) of the mixed replacement distribution and the false-negative-aware filtering mechanism contribute to the overall performance of QCPF?
- **RQ3:** How do variations in hyperparameters, such as the mixing coefficients and the smoothing parameter, affect the recommendation performance of the model?
- **RQ4:** Does the QCPF mechanism maintain the high efficiency of the base model during the inference phase while improving recommendation accuracy?

4.1 Experimental Setup

Datasets: We evaluate the proposed QCPF model on five real-world benchmark datasets:

- **MovieLens:** A commonly used movie recommendation dataset that contains user ratings, movie titles, and movie genres.
- **Steam:** Encompasses user reviews for video games on the Steam Store.
- **Beauty:** Contains metadata and user reviews for beauty products from Jun 1996 to Sep 2023.
- **Toys:** Includes user reviews and metadata for toys and games from Jun 1996 to Jul 2014.
- **Sports:** Comprises reviews and metadata for sports and outdoor products from 1996 to 2014.

Following prior works, we adopt the user-splitting strategy, which has been shown to effectively prevent information leakage in test sets. Specifically, we sort all sequences chronologically for each dataset and then split the data into training, validation, and test sets with an 8:1:1 ratio, while preserving the last 10 interactions as the historical sequence. The statistical characteristics of the processed datasets are shown in Table 1. As observed from the table, the recommendation datasets face a significant challenge of severe data sparsity.

Evaluation Metrics: We select Hit Ratio (HR@K) and Normalized Discounted Cumulative Gain (NDCG@K) as the core evaluation metrics.

Baselines: We compare our proposed model with both traditional discriminative recommenders using negative sampling and diffusion-based generative recommenders. The baselines include:

Table 1 Statistics of datasets after preprocessing.

Dataset	# users	# items	# Interactions	sparsity
Movies	6040	3883	1001456	04.27%
Steam	39795	9265	2949605	00.80%
Beauty	22,363	12,101	198,502	00.07%
Toys	19,412	11,924	138,444	00.06%
Sports	35,598	18,357	256,598	00.04%

- **SASRec** Kang and McAuley (2018): Leverages the self-attention mechanism in Transformer to model user preference scores from interaction histories, addressing data sparsity through negative sampling.
- **Caser** Tang and Wang (2018): Utilizes horizontal and vertical convolutional filters to capture sequential patterns at the point-level and union-level, allowing for skip behaviors, and models user preference scores while addressing data sparsity through negative sampling.
- **GRU4Rec** Hidasi et al. (2015): Adopts RNNs to model user preference scores from interaction histories, mitigating data sparsity through negative sampling.
- **DreamRec** Yang et al. (2023): Reshapes sequential recommendation as oracle item generation, addressing data sparsity by adding Gaussian noise to dense item embeddings.
- **PreferDiff** Liu et al. (2025): Introduces an optimization objective specifically designed for item-level DM-based recommenders, which can integrate multiple negative samples, addressing data sparsity by adding noise to dense item embeddings and using negative sampling.
- **DiffRec** Wang et al. (2023): A preference score-level diffusion-based generative recommender assuming a Gaussian prior, addressing data sparsity by adding Gaussian noise to preference scores, without considering the constraints of the probability simplex.
- **DDSR** Xie et al. (2024): A preference score-level diffusion-based generative recommender assuming a categorical prior, addressing data sparsity by adding discrete noise to preference scores while respecting the constraints of the probability simplex.
- **PreferGrow** Hu et al. (2025): A state-of-the-art discrete diffusion framework that formulates recommendation as a "fading and growing" process over discrete items, serving as the direct structural baseline for our proposed QCPF mechanism.

4.2 Overall Performance Comparison (RQ1)

Table 2 details the overall performance of all models across the five public benchmark datasets. The experimental data explicitly exposes the "task mismatch" bottleneck of the continuous diffusion paradigm in recommendation scenarios. Methods represented by DreamRec Yang et al. (2023) and PreferDiff Liu et al. (2025) fail to outperform the strong traditional baseline SASRec on certain datasets, corroborating that forcibly mapping discrete items into a continuous space severely distorts feature semantics.

Table 2 Overall performance comparison of different models across five benchmark datasets.

Method	K	MovieLens		Steam		Beauty		Toys		Sports	
		HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG	HR	NDCG
SASRec	5	0.0905	0.0502	0.0321	0.0192	0.0326	0.0221	0.0340	0.0258	0.0154	0.0101
	10	0.1709	0.0760	0.0572	0.0272	0.0438	0.0257	0.0438	0.0289	0.0230	0.0126
	20	0.2899	0.1059	0.0957	0.0368	0.0595	0.0297	0.0499	0.0305	0.0298	0.0143
Caser	5	0.0925	0.0576	0.0258	0.0163	0.0170	0.0106	0.0093	0.0072	0.0051	0.0030
	10	0.1605	0.0794	0.0446	0.0223	0.0219	0.0122	0.0129	0.0083	0.0093	0.0043
	20	0.2592	0.1042	0.0736	0.0295	0.0295	0.0141	0.0185	0.0098	0.0149	0.0057
GRU4Rec	5	0.0892	0.0551	0.0255	0.0156	0.0130	0.0082	0.0124	0.0078	0.0070	0.0048
	10	0.1534	0.0757	0.0441	0.0216	0.0188	0.0101	0.0180	0.0092	0.0107	0.0059
	20	0.2501	0.1000	0.0769	0.0298	0.0313	0.0132	0.0304	0.0117	0.0171	0.0076
DreamRec	5	0.0676	0.0437	0.0109	0.0073	0.0300	0.0247	0.0381	0.0304	0.0095	0.0084
	10	0.1083	0.0568	0.0157	0.0088	0.0353	0.0265	0.0412	0.0314	0.0112	0.0089
	20	0.1610	0.0701	0.0218	0.0104	0.0402	0.0277	0.0463	0.0327	0.0149	0.0098
PreferDiff	5	0.0538	0.0349	0.0167	0.0105	0.0335	0.0272	0.0386	0.0308	0.0168	0.0130
	10	0.0852	0.0450	0.0297	0.0145	0.0434	0.0304	0.0494	0.0343	0.0211	0.0144
	20	0.1270	0.0555	0.0526	0.0203	0.0577	0.0340	0.0644	0.0380	0.0256	0.0155
DiffRec	5	0.0266	0.0121	0.0268	0.0143	0.0095	0.0055	0.0053	0.0028	0.0063	0.0035
	10	0.0899	0.0320	0.0657	0.0268	0.0218	0.0094	0.0153	0.0059	0.0117	0.0052
	20	0.1768	0.0541	0.1159	0.0394	0.0355	0.0128	0.0226	0.0078	0.0194	0.0072
DDSR	5	0.0871	0.0533	0.0252	0.0157	0.0291	0.0217	0.0386	0.0302	0.0146	0.0109
	10	0.1523	0.0742	0.0437	0.0216	0.0434	0.0262	0.0479	0.0332	0.0213	0.0130
	20	0.2450	0.0975	0.0736	0.0291	0.0608	0.0306	0.0618	0.0367	0.0298	0.0151
PreferGrow	5	0.1413	0.0912	0.0615	0.0406	0.0420	0.0322	0.0419	0.0311	0.0207	0.0142
	10	0.2240	0.1177	0.0935	0.0508	0.0532	0.0358	0.0513	0.0337	0.0267	0.0162
	20	0.3362	0.1460	0.1413	0.0628	0.0708	0.0402	0.0642	0.0370	0.0343	0.0181
QCPF (Ours)	5	**0.1476**	**0.0951**	**0.0624**	**0.0433**	**0.0434**	**0.0330**	**0.0430**	**0.0341**	**0.0211**	**0.0144**
	10	**0.2327**	**0.1224**	**0.0940**	**0.0527**	**0.0554**	**0.0370**	**0.0528**	**0.0354**	**0.0271**	**0.0165**
	20	**0.3463**	**0.1510**	**0.1433**	**0.0638**	**0.0721**	**0.0415**	**0.0653**	**0.0382**	**0.0348**	**0.0185**

Building upon the established superiority of discrete modeling, QCPF further breaks through the performance bottlenecks of existing state-of-the-art discrete baselines. Compared to PreferGrow, QCPF achieves significant improvements across all evaluation metrics. Given that both models are perfectly aligned in their reverse generative architectures, this substantial gain can be explicitly attributed to QCPF. By reconstructing the forward fading phase via the quality-aware mixed distribution, QCPF compels the model to overcome more challenging corrupted states, thereby capturing finer-grained ranking signals.

Furthermore, QCPF exhibits exceptional generalization capability across varying degrees of data sparsity. Whether on the relatively interaction-rich MovieLens (ML1M) dataset (where HR@5 reaches 0.1476) or the extremely sparse Sports and Beauty datasets, QCPF maintains comprehensive dominance over all competing baselines. This consistent performance gain demonstrates that the multi-source negative evidence integrated within the model (e.g., popularity proxies and in-batch hard samples) can adaptively address diverse data distribution challenges. This not only establishes a robust supervisory foundation for the reverse denoising process but also profoundly validates, from an empirical perspective, that optimizing the forward degradation trajectory is the critical determinant of the upper bound in discrete diffusion recommendation.

4.3 Ablation Study (RQ2)

To comprehensively evaluate the independent contributions of the core components in QCPF, we conducted an in-depth ablation study on the ML1M dataset. We constructed the following four model variants for comparative analysis: (1) **w/o Batch Hard**, removing the batch-level hard candidates ($\lambda_h = 0$); (2) **w/o Easy Proxy**, removing the simple random negative samples ($\lambda_e = 0$); (3) **w/o Pop Proxy**, removing the popularity-based exposure proxy ($\lambda_p = 0$); and (4) **w/o Filter**, disabling the false-negative-aware filtering mechanism. The experimental results are presented in Table 3.

Table 3 Ablation study of QCPF on the MovieLens (ML1M) dataset.

Model Variant	HR@5	HR@10	HR@20	NDCG@5	NDCG@10	NDCG@20
QCPF (Full)	**0.1476**	**0.2327**	**0.3463**	**0.0951**	**0.1224**	**0.1510**
Uniform Fading Baseline	0.1413	0.2240	0.3362	0.0912	0.1177	0.1460
w/o Batch Hard ($\lambda_h = 0$)	0.1466	0.2312	0.3412	0.0948	0.1220	0.1497
w/o Easy Proxy ($\lambda_e = 0$)	0.1470	0.2301	0.3412	0.0949	0.1216	0.1495
w/o Pop Proxy ($\lambda_p = 0$)	0.1469	0.2296	0.3414	0.0947	0.1213	0.1495
w/o Filter	0.1474	0.2321	0.3454	0.0949	0.1218	0.1504

The results indicate that removing any of the distribution enhancement modules (w/o Batch Hard, w/o Easy Proxy, w/o Pop Proxy) causes the model's performance to degrade towards the uniform fading baseline. Notably, the removal of batch-level

hard samples (w/o Batch Hard) results in the most significant drop across all metrics. This is because this module utilizes the genuine interactions of concurrent users as the forward degradation trajectory, providing the most discriminative signals for the reverse denoising network. Its absence directly weakens the model's capability to capture high-order collaborative information. The complementary synergy among the multi-source distributions effectively enhances the overall representational capacity of the recommendation model.

Furthermore, disabling the false-negative filtering mechanism (w/o Filter) results in only an extremely marginal decline in model performance. However, the core function of this mechanism is to provide a theoretical robust safety net. In complex mixed sampling scenarios, without the dual constraints of identity and historical interaction checks, the model is highly susceptible to misclassifying true positive preferences as forward noise. Benefiting from the mask-based parallelized tensor operations, this filtering mechanism effectively circumvents the risk of training collapse under extreme data distributions. It accomplishes this without incurring any noticeable computational overhead, thereby safeguarding the stability and the lower bound of the model's training process.

4.4 Hyperparameter Sensitivity Analysis (RQ3)

We further investigate the impact of key hyperparameters on the performance of QCPF, including the mixing coefficients $(\lambda_h, \lambda_e, \lambda_p)$ and the popularity smoothing coefficient (α), as illustrated in Figure 3.

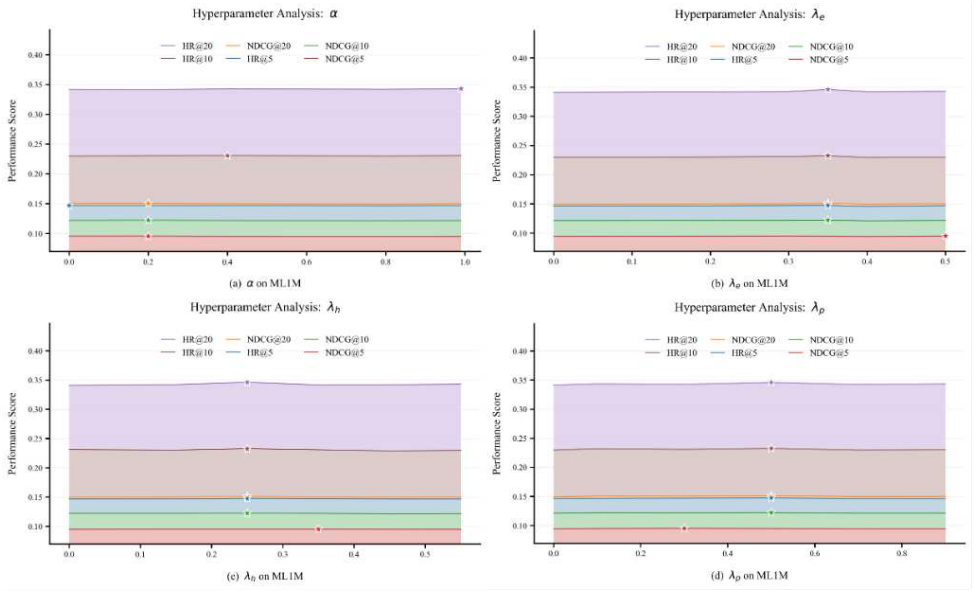


Fig. 3 Hyperparameter sensitivity analysis of QCPF on the ML1M dataset. The subfigures demonstrate the performance trends across different values of the popularity smoothing coefficient α and the mixing coefficients $(\lambda_e, \lambda_h, \lambda_p)$.

Regarding the allocation of mixing weights, the results indicate that the recommendation performance peaks when the model maintains a stable exploration baseline (λ_e) while striking a balance between batch-level hard samples (λ_h) and popularity proxies (λ_p). This phenomenon essentially reflects the dynamic requirements of the forward diffusion process: it necessitates sufficient random noise to ensure the ergodicity of spatial exploration, while simultaneously requiring the introduction of moderate, high-difficulty interference signals to enhance the discriminative power of the reverse gradients.

As for the popularity smoothing coefficient α, experiments reveal that an excessively large value causes the forward sampling to lean disproportionately towards globally popular items. Such extreme skewness can overshadow users' fine-grained personalized preference features, subsequently leading to a degradation in recommendation performance. This demonstrates that a reasonable smoothing constraint is crucial for preventing the model from falling into popularity bias.

Comprehensively observing the performance across different parameter values, QCPF exhibits a broad, stable plateau within the reasonable ranges of these hyperparameters. This relatively low parameter sensitivity substantiates the robustness of the model, substantially reducing hyperparameter tuning costs and resource consumption for large-scale industrial deployment.

4.5 Model Complexity and Inference Efficiency Analysis (RQ4)

By virtue of its architectural design, QCPF perfectly achieves the engineering deployment objective of imposing "zero extra inference overhead" while significantly enhancing Top-K ranking performance.

This high efficiency fundamentally stems from the non-invasive modular design of QCPF. The internal complex quality-aware mixed sampling and mask-based false-negative filtering strategies of QCPF are strictly confined to and completely decoupled within the forward noise-injection training phase. During the reverse denoising and online generation phase (i.e., the actual online recommendation service period), QCPF seamlessly retains the Markov state transition backbone and the reverse network of the base discrete diffusion model.

From the perspective of model complexity, since QCPF introduces no additional learnable parameters or extra attention layers to the reverse generative network, its space complexity remains strictly identical to that of the baseline model. Regarding time complexity, the forward enhancement mechanism remains completely dormant during the inference phase. Consequently, when generating recommendation lists in test or online environments, the time complexity strictly aligns with the most fundamental discrete diffusion generation process. This design elegantly circumvents the online computational latency typically incurred by incorporating complex negative evidence. As a result, QCPF can easily meet the stringent online service requirements of large-scale recommendation platforms for high concurrency and low latency, demonstrating immense potential for real-world industrial deployment.

5 Conclusion

Addressing the core pain point in discrete diffusion recommendation models where uniform random replacement during the forward fading phase leads to uninformative corrupted states, thereby restricting the fine-grained ranking capability of the reverse model this paper proposes QCPF: an innovative, training-only enhancement mechanism based on Quality-Constrained Preference Fading.

Unlike existing diffusion recommendation paradigms, QCPF breaks the limitation of a single uniform distribution by reconstructing the forward preference replacement into a multi-level quality-aware mixed distribution. This distribution dynamically integrates simple random negative samples, batch-level hard candidates, and popularity-based exposure proxies. This design not only ensures foundational exploration coverage within the massive and sparse space during the diffusion process but also provides highly discriminative and challenging ranking supervision signals for the reverse preference growing network. Concurrently, to eliminate the potential noise risks introduced by high-quality replacements, QCPF incorporates a lightweight false-negative-aware filtering mechanism. Through strict identity filtering and historical interaction constraints, it ensures the high reliability of the preference degradation trajectory.

The prominent advantage of this study is that the QCPF mechanism seamlessly embeds into existing generative frameworks without altering the Markov state transition backbone of discrete diffusion or the structure of the reverse denoising generative network. Consequently, while significantly enhancing feature learning capabilities during the training phase, the model perfectly achieves its anticipated goal of "zero extra parameters and no additional inference latency" during the inference phase, substantially boosting its feasibility for industrial deployment. Extensive experiments on five public recommendation benchmark datasets thoroughly validate the exceptional performance of QCPF in improving Top-K recommendation accuracy (especially for smaller values of K). This profoundly reveals an important scholarly proposition: a high-quality forward preference degradation trajectory is the critical determinant of the upper bound in discrete diffusion recommendations.

6 Limitations and Future Work

Although QCPF has achieved significant success in enhancing discrete diffusion recommendation performance, this study still possesses certain limitations that warrant further exploration, pointing the way for future work:

- **Lack of Deep Semantic-Level Constraints:** In the current implementation of QCPF, the mixed replacement distribution primarily relies on statistical features such as random sampling, batch co-occurrence, and global popularity. It has yet to deeply integrate the semantic features of items. In future work, we will consider incorporating semantic similarity calculations based on Large Language Models (LLMs) or multi-modal representations to construct a fine-grained "semantic replacement distribution," thereby unearthing high-quality degradation trajectories even among extremely niche, long-tail items.

- **Static Allocation of Mixing Weights:** Currently, the mixing coefficients $(\lambda_e, \lambda_h, \lambda_p)$ for the various negative evidence data sources are pre-defined hyperparameters. Across different training stages, the reverse model's capacity to adapt to noise difficulty changes dynamically. We plan to explore an adaptive Curriculum Learning mechanism, allowing the model to lean towards simple random negative samples in the early training stages to ensure stable convergence, while gradually increasing the weight of batch-level hard samples in the later stages to approach the performance limit.
- **Generalization Capability of the False-Negative Filter:** The current filtering mechanism relies exclusively on explicit historical interaction records (Hard Interactions). In highly sparse scenarios, a substantial number of unobserved yet potentially preferred items might still bypass the filtering defenses. Subsequent research could integrate Knowledge Graphs or Causal Inference to develop more flexible and intelligent false-negative evaluation metrics.

7 Supplementary information

No supplementary information is available for this article.

Author Contributions Tianwei Xu contributed to the conceptualization, methodology, and writing of the original draft. Weisong Zhang, Juxiang Zhou, and Bingkun Wang provided supervision, and were responsible for review and editing. All authors have read and approved the final manuscript.

Funding This work was supported by Major Science and Technology Project of Yunnan Province (No.202402AD080002).

Data Availability The benchmark datasets analysed during the current study (including MovieLens 1M, Steam, Beauty, Toys, and Sports) are publicly available: https://drive.google.com/file/d/1fsQUa92UV9_MqcGKqDhK9Sh4JrlbblLZ/view?usp=drive_link.

Declarations

Conflict of interest The authors declare no conflict of interest.

Ethics approval and consent to participate Not applicable.

Consent for publication Not applicable.

Competing interests The authors declare no competing interests.

Code availability The code used in this study is available from the corresponding author upon reasonable request.

References

Austin J, Johnson DD, Ho J, et al (2021) Structured denoising diffusion models in discrete state-spaces. ArXiv abs/2107.03006. URL https://api.semanticscholar.org/CorpusID:235755106

B'en'edict G, Jeunen O, Papa S, et al (2023) Recfusion: A binomial diffusion process for 1d data for recommendation. ArXiv abs/2306.08947. URL https://api.semanticscholar.org/CorpusID:259164453

Benigni M, Ferrari Dacrema M, Jannach D (2026) Diffusion recommender models and the illusion of progress: A concerning study of reproducibility and a conceptual mismatch. ACM Trans Recomm Syst 4(3). https://doi.org/10.1145/3795792, URL https://doi.org/10.1145/3795792

Campbell A, Benton J, De Bortoli V, et al (2022) A continuous time framework for discrete denoising models. In: Koyejo S, Mohamed S, Agarwal A, et al (eds) Advances in Neural Information Processing Systems, vol 35. Curran Associates, Inc., pp 28266–28279, URL https://proceedings.neurips.cc/paper_files/paper/2022/file/b5b528767aa35f5b1a60fe0aaeca0563-Paper-Conference.pdf

Chen J, Dong H, Wang X, et al (2023) Bias and debias in recommender system: A survey and future directions. ACM Transactions on Information Systems 41(3):1–39

Chen Z, Gan W, Wu J, et al (2025) Data scarcity in recommendation systems: A survey. ACM Trans Recomm Syst 3(3). https://doi.org/10.1145/3639063, URL https://doi.org/10.1145/3639063

Choi J, Noh Y, Park D (2026) Simdiffrec: Semantic similarity-guided diffusion for contrastive sequential recommendation. In: Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining. Association for Computing Machinery, New York, NY, USA, WSDM '26, p 111–121, https://doi.org/10.1145/3773966.3777978, URL https://doi.org/10.1145/3773966.3777978

Deldjoo Y, Schedl M, Cremonesi P, et al (2020) Recommender systems leveraging multimedia content. ACM Comput Surv 53(5). https://doi.org/10.1145/3407190, URL https://doi.org/10.1145/3407190

Deldjoo Y, He Z, McAuley J, et al (2024) A review of modern recommender systems using generative models (gen-recsys). arxiv. arXiv preprint arXiv:240400579

Goodfellow IJ, Pouget-Abadie J, Mirza M, et al (2014) Generative adversarial nets. In: Ghahramani Z, Welling M, Cortes C, et al (eds) Advances in Neural Information Processing Systems, vol 27. Curran Associates, Inc., URL https://proceedings.neurips.cc/paper_files/paper/2014/file/f033ed80deb0234979a61f95710dbe25-Paper.pdf

He X, Fan W, Wang R, et al (2025) Balancing user preferences by social networks: A condition-guided social recommendation model for mitigating popularity bias. Neural Networks 187:107317. https://doi.org/https://doi.org/10.1016/j.neunet.2025.107317, URL https://www.sciencedirect.com/science/article/pii/S0893608025001960

Hidasi B, Karatzoglou A, Baltrunas L, et al (2015) Session-based recommendations with recurrent neural networks. CoRR abs/1511.06939. URL https://api.semanticscholar.org/CorpusID:260446846

Ho J, Jain A, Abbeel P (2020) Denoising diffusion probabilistic models. In: Larochelle H, Ranzato M, Hadsell R, et al (eds) Advances in Neural Information Processing Systems, vol 33. Curran Associates, Inc., pp 6840–6851, URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf

Hou Y, Park JD, Shin WY (2024) Collaborative filtering based on diffusion models: Unveiling the potential of high-order connectivity. Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval URL https://api.semanticscholar.org/CorpusID:269293824

Hu G, Zhang A, Liu S, et al (2025) Fading to grow: Growing preference ratios via preference fading discrete diffusion for recommendation. ArXiv abs/2509.26063. URL https://api.semanticscholar.org/CorpusID:281682044

Kang WC, McAuley J (2018) Self-attentive sequential recommendation. In: 2018 IEEE International Conference on Data Mining (ICDM), pp 197–206, https://doi.org/10.1109/ICDM.2018.00035

Lee C, Choi J, Wi H, et al (2024) Scone: A novel stochastic sampling to generate contrastive views and hard negative samples for recommendation. URL https://arxiv.org/abs/2405.00287, arXiv:2405.00287

Li W, Huang R, Zhao H, et al (2025) Dimerec: A unified framework for enhanced sequential recommendation via generative diffusion models. In: Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining. Association for Computing Machinery, New York, NY, USA, WSDM '25, p 726–734, https://doi.org/10.1145/3701551.3703555, URL https://doi.org/10.1145/3701551.3703555

Li Z, Sun A, Li C (2023) Diffurec: A diffusion model for sequential recommendation. ACM Trans Inf Syst 42(3). https://doi.org/10.1145/3631116, URL https://doi.org/10.1145/3631116

Liang D, Krishnan RG, Hoffman MD, et al (2018) Variational autoencoders for collaborative filtering. In: Proceedings of the 2018 World Wide Web Conference. International World Wide Web Conferences Steering Committee, Republic and

Canton of Geneva, CHE, WWW '18, p 689–698, https://doi.org/10.1145/3178876.3186150, URL https://doi.org/10.1145/3178876.3186150

Lim S, YOON EB, Byun T, et al (2023) Score-based generative modeling through stochastic evolution equations in hilbert spaces. In: Oh A, Naumann T, Globerson A, et al (eds) Advances in Neural Information Processing Systems, vol 36. Curran Associates, Inc., pp 37799–37812, URL https://proceedings.neurips.cc/paper_files/paper/2023/file/76c6f9f2475b275b92d03a83ea270af4-Paper-Conference.pdf

Lin J, Liu J, Zhu J, et al (2024) A survey on diffusion models for recommender systems. ArXiv abs/2409.05033. URL https://api.semanticscholar.org/CorpusID:272524935

Lin Y, Liu Y, Lin F, et al (2023) A survey on reinforcement learning for recommender systems. IEEE Transactions on Neural Networks and Learning Systems 35(10):13164–13184

Liu Q, Yan F, Zhao X, et al (2023) Diffusion augmentation for sequential recommendation. URL https://arxiv.org/abs/2309.12858, arXiv:2309.12858

Liu S, Zhang A, Hu G, et al (2025) Preference diffusion for recommendation. URL https://arxiv.org/abs/2410.13117, arXiv:2410.13117

Lou A, Meng C, Ermon S (2023) Discrete diffusion modeling by estimating the ratios of the data distribution. In: International Conference on Machine Learning, URL https://api.semanticscholar.org/CorpusID:264451832

Mansoury M, Mobasher B, van Hoof H (2024) Mitigating exposure bias in online learning to rank recommendation: A novel reward model for cascading bandits. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. Association for Computing Machinery, New York, NY, USA, CIKM '24, p 1638–1648, https://doi.org/10.1145/3627673.3679763, URL https://doi.org/10.1145/3627673.3679763

Mao W, Liu S, Liu H, et al (2025) Distinguished quantized guidance for diffusion-based sequence recommendation. In: Proceedings of the ACM on Web Conference 2025. Association for Computing Machinery, New York, NY, USA, WWW '25, p 425–435, https://doi.org/10.1145/3696410.3714955, URL https://doi.org/10.1145/3696410.3714955

Meng C, Choi K, Song J, et al (2022) Concrete score matching: Generalized score matching for discrete data. In: Koyejo S, Mohamed S, Agarwal A, et al (eds) Advances in Neural Information Processing Systems, vol 35. Curran Associates, Inc., pp 34532–34545, URL https://proceedings.neurips.cc/paper_files/paper/2022/file/df04a35d907e894d59d4eab1f92bc87b-Paper-Conference.pdf

Rombach R, Blattmann A, Lorenz D, et al (2022) High-resolution image synthesis with latent diffusion models. In: 2022 IEEE/CVF Conference on Computer

Vision and Pattern Recognition (CVPR), pp 10674–10685, https://doi.org/10.1109/CVPR52688.2022.01042

Song J, Meng C, Ermon S (2022) Denoising diffusion implicit models. URL https://arxiv.org/abs/2010.02502, arXiv:2010.02502

Song Y, Ermon S (2019) Generative modeling by estimating gradients of the data distribution. In: Wallach H, Larochelle H, Beygelzimer A, et al (eds) Advances in Neural Information Processing Systems, vol 32. Curran Associates, Inc., URL https://proceedings.neurips.cc/paper_files/paper/2019/file/3001ef257407d5a371a96dcd947c7d93-Paper.pdf

Sun F, Liu J, Wu J, et al (2019) Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. Association for Computing Machinery, New York, NY, USA, CIKM '19, p 1441–1450, https://doi.org/10.1145/3357384.3357895, URL https://doi.org/10.1145/3357384.3357895

Tang J, Wang K (2018) Personalized top-n sequential recommendation via convolutional sequence embedding. Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining URL https://api.semanticscholar.org/CorpusID:39847715

Wang S, Cao L, Wang Y, et al (2021) A survey on session-based recommender systems. ACM Comput Surv 54(7). https://doi.org/10.1145/3465401, URL https://doi.org/10.1145/3465401

Wang W, Xu Y, Feng F, et al (2023) Diffusion recommender model. In: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. Association for Computing Machinery, New York, NY, USA, SIGIR '23, p 832–841, https://doi.org/10.1145/3539618.3591663, URL https://doi.org/10.1145/3539618.3591663

Wei TR, Fang Y (2026) Diffusion models in recommendation systems: A survey. URL https://arxiv.org/abs/2501.10548, arXiv:2501.10548

Wu J, Wang X, Feng F, et al (2021) Self-supervised graph learning for recommendation. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval. Association for Computing Machinery, New York, NY, USA, SIGIR '21, p 726–735, https://doi.org/10.1145/3404835.3462862, URL https://doi.org/10.1145/3404835.3462862

Wu L, Zheng Z, Qiu Z, et al (2024) A survey on large language models for recommendation. World Wide Web 27(5). https://doi.org/10.1007/s11280-024-01291-2, URL https://doi.org/10.1007/s11280-024-01291-2

Wu S, Sun F, Zhang W, et al (2022) Graph neural networks in recommender systems: A survey. ACM Comput Surv 55(5). https://doi.org/10.1145/3535101, URL https://doi.org/10.1145/3535101

Xie W, Wang H, Zhang L, et al (2024) Breaking determinism: Fuzzy modeling of sequential recommendation using discrete state space diffusion model. ArXiv abs/2410.23994. URL https://api.semanticscholar.org/CorpusID:273707451

Xu Y, Wang W, Feng F, et al (2024) Diffusion models for generative outfit recommendation. In: Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval. Association for Computing Machinery, New York, NY, USA, SIGIR '24, p 1350–1359, https://doi.org/10.1145/3626772.3657719, URL https://doi.org/10.1145/3626772.3657719

Xuan Y (2024) Diffusion cross-domain recommendation. URL https://arxiv.org/abs/2402.02182, arXiv:2402.02182

Yang Z, Wu J, Wang Z, et al (2023) Generate what you prefer: Reshaping sequential recommendation via guided diffusion. In: Oh A, Naumann T, Globerson A, et al (eds) Advances in Neural Information Processing Systems, vol 36. Curran Associates, Inc., pp 24247–24261, URL https://proceedings.neurips.cc/paper_files/paper/2023/file/4c5e2bcbf21bdf40d75fddad0bd43dc9-Paper-Conference.pdf