

Generative Diffusion Models for Sequential Recommendations

Sharare Zolghadr¹, Ole Winther² and Paul Jeha²

¹University of Padua, Padua, Italy

²Technical University of Denmark, Copenhagen, Denmark

Abstract

Generative models such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) have shown promise in sequential recommendation tasks. However, they face challenges, including posterior collapse and limited representation capacity. The work by Li et al., 2023 is a novel approach that leverages diffusion models to address these challenges by representing item embeddings as distributions rather than fixed vectors. This approach allows for a more adaptive reflection of users' diverse interests and various item aspects. During the diffusion phase, the model converts the target item embedding into a Gaussian distribution by adding noise, facilitating the representation of sequential item distributions and the injection of uncertainty. An Approximator then processes this noisy item representation to reconstruct the target item. In the reverse phase, the model utilizes users' past interactions to reverse the noise and finalize the item prediction through a rounding operation. This research introduces enhancements to the DiffuRec architecture, particularly by adding offset noise in the diffusion process to improve robustness and incorporating a cross-attention mechanism in the Approximator to better capture relevant user-item interactions. These contributions led to the development of a new model, DiffuRecSys, which improves performance. Extensive experiments conducted on three public benchmark datasets demonstrate that these modifications enhance item representation, effectively capture diverse user preferences, and outperform existing baselines in sequential recommendation research.

Keywords

Diffusion Models, Recommender Systems, Generative Models

1. Introduction

Recommender systems are algorithms that suggest items to users by analyzing various forms of input data. Their primary goal is to enhance the customer experience through personalized recommendations, often based on prior implicit feedback. These systems track user behaviors, such as purchase history, viewing habits, and browsing activity, to model user preferences. Sequential Recommendation, a specific type of recommendation, is particularly relevant for applications where user behavior is naturally sequential. It focuses on predicting the next item a user will interact with by considering the order of previous interactions.

Mainstream solutions to Sequential Recommendation (SR) [2] represent items with fixed vectors, which have a limited ability to capture the latent aspects of items and the diversity of user preferences. Generative models like Generative Adversarial Networks (GANs) [3] and Variational Auto-Encoders (VAEs) [4] have been widely applied in personalized recommendations, using adversarial training and encoder-decoder architectures, respectively, to model user behavior and preferences. However, Diffusion Models have shown significant advantages over GANs and VAEs, such as greater stability and higher generation quality in various tasks.

Diffusion Models (DMs) [5, 6, 7] have achieved state-of-the-art results in image synthesis tasks [7, 8, 9, 10, 11]. These models alleviate the trade-off between stability and quality by gradually corrupting images in a forward process and iteratively learning to reconstruct them. DMs progressively corrupt \mathbf{x}_0 with random noise and then recover \mathbf{x}_0 from the corrupted state \mathbf{x}_T step by step. This forward process creates a tractable posterior [6], enabling the iterative modeling of complex distributions through flexible neural networks in the reverse generation process.

The objectives of recommender models align well with DMs, as recommendation essentially involves inferring future interaction probabilities from noisy historical interactions, where noise represents false-positive and false-negative items [12, 13]. This makes DMs a promising approach for accurately modeling complex interaction patterns with strong representational ability. Despite their success in other domains, applying diffusion models to recommender systems remains underexplored.

We further explore diffusion models for sequential recommendation (SR) by extending the method introduced by Li et al. (2023). Our work proposes significant enhancements to the existing architecture, resulting in a new model, **DiffuRecSys**¹. Specifically, our contributions are as follows:

- **Enhancing the Diffusion Recommender Model:** We incorporate cross-attention mechanisms within the Approximator of the model architecture. The model isn't just learning temporal dependencies (the sequential order of items) but also more complex relationships between past interactions and the target item by focusing on the most relevant past events.
- **Incorporation of Offset Noise:** We introduce offset noise into the diffusion process to increase model robustness and effectively handle variability in user interactions.
- **Comprehensive Experimental Validation:** We conduct extensive experiments across three datasets under various settings, demonstrating improvements of DiffuRec with our extensions over standard baselines.

2. Related Works

2.1. Recommender Systems

In the era of information overload, recommender systems are essential for filtering relevant information and have been widely adopted in e-commerce, online news, and social media. A core aspect of these systems is modeling

ROEGEN@RecSys'24: The 1st Workshop on Risks, Opportunities, and Evaluation of Generative Models in Recommender Systems, Co-located with ACM RecSys in Bari, Italy, October 2024

✉ sharare.zolghadr@studenti.unipd.it (S. Zolghadr); olwi@dtu.dk (O. Winther); pauje@dtu.dk (P. Jeha)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://youtu.be/bEpDfaAGL2I>

user preferences based on past interactions. Traditional recommender systems are generally categorized into three main types: content-based filtering [14], collaborative filtering [15, 16], and hybrid methods [17]. However, while conventional approaches focus on static user preferences, sequential recommender systems account for the temporal dynamics of user behavior. These systems recognize that user preferences evolve over time, and the order of interactions provides critical context for making accurate recommendations. Early sequential recommendation models utilized Markov chains to capture item transition probabilities. These models could make recommendations based on recent interactions by modeling the likelihood of a user moving from one item to another. However, Markov chains often struggle to capture long-term dependencies due to their limited memory capacity [18].

Recurrent Neural Networks (RNNs) have been widely adopted for sequential recommendation tasks due to their ability to model temporal sequences. Popular variants, such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRUs) effectively capture both short-term and long-term dependencies in user behavior. These models have significantly improved recommendation accuracy by leveraging sequential patterns in user interactions [19].

Attention mechanisms have further advanced sequential recommendation systems by enabling models to focus on the most relevant parts of the interaction history. Self-attention models, such as the Transformer, effectively capture complex dependencies within sequences, leading to state-of-the-art performance in various sequential recommendation tasks [2].

2.2. Generative Models

In the realm of generative modeling, the primary objective is to approximate an unknown distribution $p_{\text{data}}(x)$ using a model distribution $p_{\text{model}}(x)$. This is typically achieved by defining $p_{\text{model}}(x; \theta)$, where θ represents parameters that are adjusted to minimize the divergence between p_{data} and p_{model} . A prevalent method for achieving this is maximum likelihood estimation, which minimizes the Kullback-Leibler (KL) divergence between the true and model distributions.

Deep generative models have recently demonstrated high-quality sampling across various data modalities. Models such as Generative Adversarial Networks (GANs), autoregressive models, and Variational Autoencoders (VAEs) have shown impressive capabilities in generating complex data types, including images and audio. This progress has sparked significant interest in adapting these models for generative recommender systems. By modeling the underlying generative process, such systems infer user interaction probabilities for items that have not been interacted with. Typically, these models assume that latent factors, such as user preferences, drive users' interactions with items (e.g., clicks).

Generative recommender models [20] can be broadly categorized into two types. First, GAN-based models employ a generator to estimate user interaction probabilities and rely on adversarial training for optimization [21]. However, adversarial training can be unstable, which may lead to inconsistent performance. Second, VAE-based models use an encoder to approximate the posterior distribution over latent factors and aim to maximize the likelihood of observed interactions [22]. Although VAEs generally outperform GANs in recommendation tasks, they face a trade-off

between tractability and representation power [23]. Simple encoders may struggle to capture heterogeneous user preferences, while more complex models may result in intractable posterior distributions [24]. Ongoing research and innovation are crucial for advancing generative modeling and enhancing recommendation systems.

3. Problem Statement

In sequential recommendation, let U represent the set of users and I the set of items. For each user $u \in U$, we organize the items that the user has interacted with in chronological order, forming a sequence $S_u = [i_1^u, i_2^u, \dots, i_{N-1}^u, i_N^u]$, where $i_n^u \in I$ denotes the item that user u interacted with at the n -th timestamp, and N represents the maximum sequence length. The task of sequential recommendation is to predict the next item that user u is likely to interact with. During training, the model learns to maximize the probability of recommending the target item i_N based on the user's previously interacted items $[i_1^u, i_2^u, \dots, i_{N-1}^u]$. This probability is expressed as $p(i_N^u | i_1^u, i_2^u, \dots, i_{N-1}^u)$. During inference, the model predicts the probability of recommending the next item i_{N+1} based on the entire interaction sequence $[i_1^u, i_2^u, \dots, i_N^u]$, which can be written as $p(i_{N+1}^u | i_1^u, i_2^u, \dots, i_N^u)$. The ultimate goal of sequential recommendation is to generate a ranked list of items as candidates for the next item that user u is most likely to prefer.

4. Preliminaries

In this section, we provide a brief introduction to diffusion models and sequential recommendations. Subsequently, in Section 5 (Model Architecture), we will explain the methodology for adapting diffusion models to integrate with recommender systems.

4.1. Diffusion Model

Deep diffusion probabilistic models can be briefly described as hierarchical variational autoencoders (VAEs) with a bottom-up path defined by a diffusion process (e.g., Gaussian diffusion) and a top-down path parameterized by deep neural networks (DNNs), representing a reversed diffusion process. Diffusion models [6] are latent variable models characterized by the distribution:

$$p_{\theta}(x_0) := \int p_{\theta}(x_{0:T}) dx_{1:T}, \quad (1)$$

where x_1, \dots, x_T are latent variables with the same dimensionality as the data $x_0 \sim q(x_0)$. The joint distribution $p_{\theta}(x_{0:T})$ represents the reverse process and is defined as a Markov chain with learned Gaussian transitions starting from $p(x_T) = \mathcal{N}(x_T; 0, I)$:

$$p_{\theta}(x_{0:T}) := p(x_T) \prod_{t=1}^T p_{\theta}(x_{t-1} | x_t), \quad (2)$$

$$p_{\theta}(x_{t-1} | x_t) := \mathcal{N}(x_{t-1}; \mu_{\theta}(x_t, t), \Sigma_{\theta}(x_t, t)).$$

What distinguishes diffusion models from other latent variable models is that the approximate posterior $q(x_{1:T} | x_0)$, known as the forward process or diffusion process, is

fixed as a Markov chain that gradually adds Gaussian noise to the data according to a variance schedule β_1, \dots, β_T :

$$\begin{aligned} q(x_{1:T} | x_0) &:= \prod_{t=1}^T q(x_t | x_{t-1}), \\ q(x_t | x_{t-1}) &:= \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I). \end{aligned} \quad (3)$$

Here, x_t is sampled from this Gaussian distribution, where β_t controls the amount of noise added at the t -th diffusion step, and I is the identity matrix. The value of β_t is determined by a predefined noise schedule β , which specifies the noise amount for each step. Common noise schedules include square-root [25], cosine [26], and linear [5].

In the reverse phase, where the original representation x_0 is unknown, a deep neural network $f_\theta(\cdot)$ (e.g., Transformer [27] or U-Net [28]) is typically used to estimate x_0 . Given the original representation x_0 and the schedule β , training $f_\theta(\cdot)$ involves optimizing the variational lower bound (VLB) on the negative log-likelihood [6]:

$$\begin{aligned} \mathbb{E}[-\log p_\theta(x_0)] &\leq \mathbb{E}_q \left[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T} | x_0)} \right] = \\ \mathbb{E}_q \left[-\log p(x_T) - \sum_{t \geq 1} \log \frac{p_\theta(x_{t-1} | x_t)}{q(x_t | x_{t-1})} \right] &=: \mathcal{L}. \end{aligned} \quad (4)$$

The forward process variances, β_t , can either be adjusted via parameterization [4] or kept fixed as hyperparameters. The reverse process remains effective due to the use of Gaussian conditionals in $p_\theta(x_{t-1} | x_t)$, particularly when β_t values are small [6]. A notable property of the forward process is that it facilitates straightforward sampling of x_t at any time step t . Using the notation $\alpha_t := 1 - \beta_t$ and $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$, we express:

$$q(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I), \quad (5)$$

Efficient training can be achieved by optimizing the random terms of \mathcal{L} using stochastic gradient descent. To further enhance training stability, we can reduce variance by rewriting \mathcal{L} (4) as:

$$\begin{aligned} \mathbb{E}_q \left[\underbrace{D_{KL}(q(x_T | x_0) \| p(x_T))}_{\mathcal{L}_T} + \right. \\ \left. \sum_{t > 1} \underbrace{D_{KL}(q(x_{t-1} | x_t, x_0) \| p_\theta(x_{t-1} | x_t))}_{\mathcal{L}_{t-1}} \right. \\ \left. - \underbrace{\log p_\theta(x_0 | x_1)}_{\mathcal{L}_0} \right], \end{aligned} \quad (6)$$

Equation (6) leverages KL divergence to compare $p_\theta(x_{t-1} | x_t)$ against the forward process posteriors, which are tractable when conditioned on x_0 :

$$q(x_{t-1} | x_t, x_0) = \mathcal{N}(x_{t-1}; \tilde{\mu}_t(x_t, x_0), \tilde{\beta}_t I), \quad (7)$$

where

$$\begin{aligned} \tilde{\mu}_t(x_t, x_0) &:= \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}x_0 + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}x_t, \\ \tilde{\beta}_t &:= \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}\beta_t, \end{aligned} \quad (8)$$

Thus, all KL divergences in Equation (6) are Gaussian comparisons, which can be calculated exactly using closed-form expressions, rather than high-variance Monte Carlo estimates.

The loss function in Equation (6) is organized into three components: L_t , L_{t-1} , and L_0 . Component L_t aims to bring $q(x_t | x_0)$ close to a standard Gaussian distribution. Component L_{t-1} minimizes the KL divergence between the forward process posterior and $p_\theta(x_{t-1} | x_t)$, which is generated by the deep neural network for the reverse process. The last Component L_0 applies negative log-likelihood for the final prediction. This variational lower bound (VLB) loss could lead to unstable model training. To address this, each L_t can be further simplified to:

$$L_{t,\text{simple}} = \mathbb{E}_{t, x_0, \epsilon} \left[\left\| \epsilon - \epsilon_\theta \left(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t \right) \right\|^2 \right], \quad (9)$$

where $\epsilon \sim \mathcal{N}(0, I)$ is sampled from a standard Gaussian distribution for noise injection and diffusion, and $\epsilon_\theta(\cdot)$ serves as an approximator instantiated by a deep neural network, such as a Transformer or U-Net.

4.2. Sequential Recommendation

Let $R \in \mathbb{R}^{m \times n}$ represent the user-item interaction matrix, where m is the number of users and n is the number of items. Each entry r_{ui} indicates the interaction (e.g., rating or click) between user u and item i . Each user u has a sequence of interacted items $S_u = [i_1, i_2, \dots, i_{T=N_u}]$, where N_u is the length of the sequence. Items are embedded into a latent space using an embedding matrix $E \in \mathbb{R}^{n \times d}$, where d is the dimensionality of the embedding. The embedding of item i is denoted as $e_i = E[i]$.

RNNs are used to model the sequential nature of user interactions. For a sequence S_u , the hidden state h_t at time step t is updated as

$$h_t = \text{RNN}(h_{t-1}, e_{i_t}), \quad (10)$$

where e_{i_t} is the embedding of the item at time t , and h_{t-1} is the hidden state from the previous time step.

The self-attention mechanism can be used to weigh the importance of each item in the sequence. The attention score α_{tj} between item i_t and item i_j is computed as

$$\alpha_{tj} = \frac{\exp(e_{i_t}^\top e_{i_j})}{\sum_{k=1}^{T_u} \exp(e_{i_t}^\top e_{i_k})}, \quad (11)$$

where \exp denotes the exponential function.

The attended representation z_t is then

$$z_t = \sum_{j=1}^{T_u} \alpha_{tj} e_{i_j}. \quad (12)$$

The model predicts the next item \hat{i}_{T_u+1} based on the final hidden state h_{T_u} or the attended representation z_{T_u} as

$$\hat{y} = W h_{T_u} + b, \quad (13)$$

where W and b are learned parameters, and \hat{y} is the predicted score vector for the next item.

5. Model Architecture

The DiffuRec framework is shown in Figure 1. The input to the model is the embedding e_{n+1} of the target item i_{n+1} .

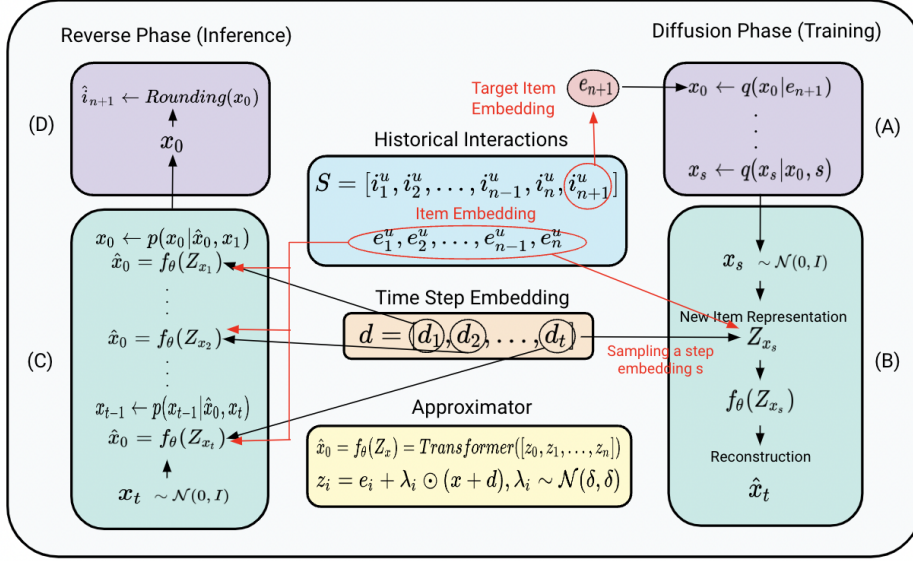


Figure 1: Overview of the diffusion process for sequential recommendation: (A) Injecting noise into the target item after s diffusion steps; (B) Generating new item representation based on user history and the last target item; (C) Reverse phase for target item reconstruction; (D) Rounding phase to map the continuous target representation to discrete item indices.

According to the diffusion model detailed in Section 4.1, we aim to reverse this process to retrieve the target item from the historical sequence S . To achieve this, we introduce noise into the target item embedding e_{n+1} via the diffusion process, which involves multiple samplings from Gaussian distributions (Equation 3). The noised target item representation x_s , which has undergone s diffusion steps, is considered as the distribution representation sampled from $q(\cdot)$. Consequently, x_s can still encapsulate multiple latent aspects of the item. Next, we use x_s to adjust the new representation of each historical item embedding in S , leveraging the target item’s guidance as additional semantic signals. The resulting representations $Z_{x_s} = [z_1, z_2, \dots, z_n]$ are then input into the approximator $f_\theta(\cdot)$. The model is trained to ensure the reconstructed \hat{x}_0 from the approximator closely matches the target item embedding e_{n+1} . In the reverse phase, we start by sampling the noised target item representation x_t from a standard Gaussian distribution, $x_t \sim \mathcal{N}(0, I)$. Similar to the diffusion process, the adjusted representations Z_{x_t} (obtained using x_t) are fed into the well-trained approximator $f_\theta(\cdot)$ for \hat{x}_0 estimation. Following Equation (2), the estimated \hat{x}_0 and x_t are utilized to reverse x_{t-1} through $p(\cdot)$. This iterative process continues until x_0 is reached. It is important to note that the reverse phase is stochastic, which models the uncertainty in user behavior. Finally, a rounding function maps the reversed continuous representation x_0 into discrete candidate item indices for predicting the target item \hat{i}_{n+1} . We present three versions of the Diffusion Recommender Model: 1) a baseline based on previous work [1], 2) a model with cross-attention in the approximator, and 3) a model with both offset noise and cross-attention. We compare the performance of all models to assess the effectiveness of each component.

5.1. Approximator

The transformer is employed as the backbone of the approximator $f_\theta(\cdot)$ to generate \hat{x}_0 , leveraging its effectiveness in modeling sequential dependencies during both the diffusion

and reverse phases:

$$\hat{x}_0 = f_\theta(Z_x) = \text{Transformer}([z_0, z_1, \dots, z_n]). \quad (14)$$

To enhance the representation of the user’s historical interactions with the last target item, we employ cross-attention. Instead of summing the last target item and timestep embedding with the historical interactions embedding, as shown in Equation (15), we utilize cross-attention, as depicted in Equation (16):

$$z_i = e_i + \lambda_i \odot (x + d), \quad (15)$$

$$z_i = \text{CrossAttention}(e_i, \lambda_i \odot (x + d)), \quad (16)$$

In Equation (16), \odot denotes element-wise multiplication, and d represents the step embeddings that incorporate specific information for each diffusion and reverse step. In the diffusion phase, x is the noised target item embedding x_s . In the reverse process, x refers to the reversed target item representation, i.e., $x = x_s$ for $s = t, t-1, \dots, 2, 1$.

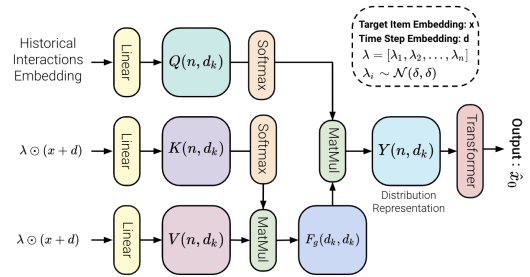


Figure 2: The architecture of approximation using cross-attention between the last target item and the user’s historical interactions.

Cross-attention allows the model to dynamically weigh the importance of historical interactions to the last target item, enabling a more nuanced aggregation of information.

This improves the model’s ability to capture relevant context and relationships, leading to better item representation, a multi-head attention Transformer is used to predict the next item in the sequence. According to Equation (16), λ_i is sampled from a Gaussian distribution, $\lambda_i \sim \mathcal{N}(\delta, \delta)$, where δ is a hyperparameter that defines both the mean and variance. During the diffusion phase, λ_i determines the level of noise injection.

In the diffusion phase, a step-index s is randomly sampled from a uniform distribution over the range $[1, t]$, where t represents the total number of steps. Conversely, in the reverse phase, s ranges from t to 1. Consequently, both the corresponding step embedding and the sequence item distribution are fed into the approximator for model training (see Equation (15)). Ultimately, the representation h_n of item i_n generated by the final layer is used as \hat{x}_0 .

5.2. Diffusion Phase

The offset noise technique was implemented to enhance the robustness and diversity of generated items in a diffusion model. The conventional approach in diffusion models involves sampling noise from a standard normal distribution, $\mathcal{N}(0, I)$, where the mean is zero and the variance is one. However, this approach can make the model sensitive to mean shifts in the input data, potentially limiting the diversity and quality of generated outputs.

To address this issue, an offset term was introduced in the noise distribution. Specifically, the noise was sampled from a distribution with a constant offset term, denoted as:

$$\epsilon_{\text{offset}} \sim \mathcal{N}(0.1\delta_c, I), \quad (17)$$

instead of the standard normal distribution. This offset is applied only during the training or forward process. During inference or the reverse process, noise is still sampled from a standard normal distribution. Here, δ_c represents a constant value, and the scaling factor 0.1 adjusts the mean of the noise distribution.

During training, we randomly sample a diffusion step s for each target item, specifically $s = \lfloor s' \rfloor$ with $s' \sim U(0, t)$. x_s is generated via:

$$q(x_s|x_0, s) = \mathcal{N}(x_s; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_s)I). \quad (18)$$

We derive x_0 through one-step diffusion from the target item embedding:

$$q(x_0|e_{n+1}) = \mathcal{N}(x_0; \sqrt{\bar{\alpha}_0}e_{n+1}, (1 - \alpha_0)I). \quad (19)$$

Utilizing the reparameterization trick with $\epsilon \sim \mathcal{N}(0.1\delta_c, I)$, x_s can be generated as follows:

$$x_s = \sqrt{\bar{\alpha}_s}x_0 + \sqrt{1 - \bar{\alpha}_s}\epsilon. \quad (20)$$

The diffusion process or model training is detailed in Algorithm 1.

5.3. Reverse Phase

In the reverse phase, our goal is to iteratively reconstruct the target item representation x_0 from a purely Gaussian noise x_t . However, obtaining x_0 directly at each reverse step is infeasible. Thus, we utilize a well-trained approximator to generate \hat{x}_0 as an estimate of x_0 , specifically $x_0 = \hat{x}_0$. The reverse phase is detailed in Algorithm 2. Subsequently,

Algorithm 1: Training (Diffusion Phase)

Input: User interactions \bar{X} , Initial parameters θ

Output: Optimized parameters θ

- 1 Sample a batch $X \subseteq \bar{X}$;
 - 2 **for** all $x_0 \in X$ **do**
 - 3 Sample $t \sim \mathcal{U}(1, T)$, $\epsilon \sim \mathcal{N}(0.1\delta_c, I)$;
 - 4 $x_t \leftarrow q(x_t|x_0)$, $\lambda_i \sim \mathcal{N}(\delta, \delta)$;
 - 5 $(z_1, \dots, z_n) \leftarrow$
 $(e_1 + \lambda_1(x_s + d_s), \dots, e_n + \lambda_n(x_s + d_s))$;
 - 6 $\hat{x}_0 \leftarrow f_\theta(z_1, \dots, z_n)$;
 - 7 Update θ via $L_{CE}(\hat{x}_0, i_{n+1})$;
 - 8 **While** not converged **do** Repeat above steps;
-

following Equation (2), the reverse process proceeds by applying the reparameterization trick:

$$\hat{x}_0 = f_\theta(Z_{x_t}), \quad x_{t-1} \leftarrow p(x_{t-1} | \hat{x}_0, x_t), \quad (21)$$

$$x_{t-1} = \tilde{\mu}_t(x_t, \hat{x}_0) + \tilde{\beta}_t \epsilon', \quad (22)$$

where

$$\begin{aligned} \tilde{\mu}_t(x_t, \hat{x}_0) &= \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t} \hat{x}_0 + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} x_t, \\ \tilde{\beta}_t &= \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t. \end{aligned} \quad (23)$$

After generating x_{t-1} using $\epsilon' \sim \mathcal{N}(0, I)$, repeat the reverse process until reaching x_0 .

Algorithm 2: Inference (Reverse Phase)

Input: Sequence (i_1, \dots, i_n) , Target $x_t \sim \mathcal{N}(0, I)$

Output: Predicted item i_{n+1}

- 1 **for** $s = T, \dots, 1$ **do**
 - 2 $\lambda_i \sim \mathcal{N}(\delta, \delta)$;
 - 3 $(z_1, \dots, z_n) \leftarrow$
 $(e_1 + \lambda_1(x_s + d_s), \dots, e_n + \lambda_n(x_s + d_s))$;
 - 4 $\hat{x}_0 \leftarrow f_\theta(z_1, \dots, z_n)$, $\epsilon' \sim \mathcal{N}(0, I)$;
 - 5 $x_{s-1} \leftarrow \tilde{\mu}_s(\hat{x}_0, x_s) + \tilde{\beta}_s \cdot \epsilon'$;
 - 6 $i_{n+1} \leftarrow \text{Rounding}(x_0)$
-

5.4. Loss Function and Rounding

Following previous works [29, 30, 31], we adopt the cross-entropy loss during the diffusion phase for model optimization as follows:

$$\hat{y}_i = \frac{\exp(\hat{x}_0 \cdot e_{n+1})}{\sum_{i \in I} \exp(\hat{x}_0 \cdot e_i)}, \quad (24)$$

$$L_{CE} = -\frac{1}{|U|} \sum_{u \in U} \log \hat{y}_u, \quad (25)$$

where \hat{x}_0 is reconstructed by the Transformer-based approximator, and \cdot denotes the inner product operation.

During the inference phase, the task is to map the reversed target item representation x_0 into the discrete item index space for the final recommendation. To achieve this, we compute the inner product between x_0 and all candidate item embeddings e_i and subsequently select the item index corresponding to the maximum value as the recommended result. This process can be formalized as follows:

$$i_{n+1} = \arg \max_{i \in I} \text{Rounding}(x_0) = x_0 \cdot e_i^T. \quad (26)$$

6. Experiment

In this section, we conduct experiments on three real-world datasets to address the following research question:

RQ1) Does modifying the baseline Diffusion Recommender Model (DiffuRec) enhance performance?

RQ2) How does the proposed DiffRecSys framework compare to competitive methods?

RQ3) What are the weaknesses of the DiffRecSys model, and what solutions can be proposed to address them?

6.1. Dataset

For the experiments conducted in this study, we utilized the RecBole framework [32], a comprehensive and unified tool for benchmarking recommender systems. RecBole provides an extensive collection of preprocessed datasets, which facilitated consistent and efficient data preparation. Specifically, we employed RecBole to download and preprocess the required datasets, available at <https://github.com/RUCAIBox/RecSysDatasets>, ensuring compatibility and comparability across various recommendation models. The standardized environment of RecBole, along with its support for various algorithms and evaluation metrics, enabled a fair and reliable assessment of model performance. We conducted experiments on three public benchmark datasets. Two sub-categories of the Amazon dataset (Beauty and Toys) were selected, both comprising user reviews of products. Additionally, we utilized the MovieLens 1M dataset (ML1M), a widely recognized benchmark dataset containing approximately 1 million user ratings on movies.

6.2. Experiment Setup

For the training, validation, and test splits, we adopt a sequential split approach. Specifically, for all datasets, given a sequence $S = \{i_1, i_2, \dots, i_n\}$, the most recent interaction (i_n) is used for testing, the penultimate interaction (i_{n-1}) is used for validation, and the earlier interactions ($\{i_1, i_2, \dots, i_{n-2}\}$) are used for training.

To assess the performance of sequential recommender systems, we evaluate all models using HR@K (Hit Rate) and NDCG@K (Normalized Discounted Cumulative Gain). HR@K measures the proportion of times the true item is among the top-K recommendations, while NDCG@K evaluates the ranking quality of the top-K items, taking into account the position of the correct items. The experimental results are reported for the top-K list with K set to 5, 10, and 20. We treat all reviews or ratings as implicit feedback (i.e., user-item interactions) and chronologically organize them by their timestamps. The maximum sequence length is set to 100 to balance capturing sufficient historical information with computational efficiency for the MovieLens-1M, Amazon Beauty, and Amazon Toys datasets. These datasets exhibit considerable diversity in sequence lengths and dataset sizes, covering a broad range of real-world scenarios. Additionally, we filter out unpopular items and inactive users who have fewer than five associated interactions.

For training, we employ the Adam optimizer with an initial learning rate of 0.001, a weight decay of 0.0001, and a learning rate decay step that reduces the learning rate starting at epoch 30, with the total number of epochs set to 41. Transformer parameters are initialized using Xavier initialization, and the number of blocks is set to 4.

We use the standard Transformer architecture and configuration as described by [27], which includes multi-head self-attention, a feed-forward network with ReLU activation, layer normalization, dropout, and residual connections. To enhance the model’s representation capacity, we stack multiple Transformer blocks.

Both the embedding dimension and hidden state size are fixed at 128, with a batch size of 1024. Dropout rates for the Transformer block and item embeddings are set to 0.1 and 0.3, respectively, across all datasets. For the hyperparameter λ , values are sampled from a Gaussian distribution with a mean of 0.001 and a standard deviation of 0.001. The total number of reverse steps t is set to 32. For the noise schedule β , a truncated linear schedule is used where the noise variance decreases linearly over the steps. The baseline model and two other variants are evaluated across multiple experimental runs, typically more than five, with average results reported. Statistical significance is assessed using a Student’s t-test.

6.3. Main Results (RQ1)

To evaluate the efficacy of our proposed modifications, we conducted a series of experiments on three distinct datasets. The experimental setup was designed to compare the performance of the baseline model (DiffuRec [1]) with two enhanced models: one incorporating cross-attention alone and another combining cross-attention with Offset Noise. Each experiment was repeated five times with different random seeds to ensure the robustness and reliability of the results. Our objective was to determine whether these modifications led to performance improvements.

Comprehensive details of the experimental configurations and the average outcomes for each model are provided in Tables 1, 2, and 3.

Table 1
Comparison of DiffuRec and its Variants

Model Version	Cross Attention	Offset Noise
DiffuRec (Baseline)	-	-
Modified Version 1	✓	-
Modified Version 2	✓	✓

Table 2
Convergence Training Time in Seconds

Model	MovieLens	Toys	Beauty
Baseline	65312.18	7458.10	8734.22
Modified V1	30612.66	4470.52	6776.10
Modified V2	27283.95	4623.85	6468.38

According to the results shown in Table 4, these subtle modifications led to a significant average accuracy improvement of 1.5% compared to the baseline model. For each dataset, the p-values from the paired t-tests were below the conventional significance threshold of 0.05, indicating that the performance differences between the extended models (Modified V1 and Modified V2) and the baseline (DiffuRec) are statistically significant across all datasets and metrics tested. Furthermore, the new variants not only outperformed the baseline but also reduced training time by approximately two-thirds, underscoring the effectiveness of our enhancements.

Table 3
Average Performance and Standard Deviation Over Multiple Experiments: Modified Versions vs. DiffuRec (Baseline)

Dataset	Metric	Modified V1	Modified V2	DiffuRec
Beauty	HR@5	0.0669 ± 0.0018	0.0667 ± 0.0014	0.0557
	HR@10	0.0974 ± 0.0021	0.0980 ± 0.0012	0.0790
	HR@20	0.1399 ± 0.0023	0.1400 ± 0.0019	0.1110
	NDCG@5	0.0458 ± 0.0012	0.0458 ± 0.0010	0.0400
	NDCG@10	0.0556 ± 0.0012	0.0559 ± 0.0009	0.0475
	NDCG@20	0.0663 ± 0.0011	0.0665 ± 0.0010	0.0556
Movielens-1M	HR@5	0.1940 ± 0.0029	0.1957 ± 0.0035	0.1797
	HR@10	0.2829 ± 0.0027	0.2843 ± 0.0043	0.2626
	HR@20	0.3947 ± 0.0044	0.3958 ± 0.0039	0.3679
	NDCG@5	0.1319 ± 0.0022	0.1319 ± 0.0017	0.1212
	NDCG@10	0.1607 ± 0.0012	0.1605 ± 0.0019	0.1479
	NDCG@20	0.1888 ± 0.0080	0.1886 ± 0.0013	0.1744
Toys	HR@5	0.0692 ± 0.0025	0.0684 ± 0.0018	0.0557
	HR@10	0.0986 ± 0.0035	0.0971 ± 0.0019	0.0746
	HR@20	0.1371 ± 0.0043	0.1350 ± 0.0010	0.0984
	NDCG@5	0.0458 ± 0.0010	0.0455 ± 0.0013	0.0417
	NDCG@10	0.0553 ± 0.0012	0.0548 ± 0.0013	0.0477
	NDCG@20	0.0650 ± 0.0010	0.0643 ± 0.0013	0.0537

Table 4
Paired t-test Results between Variants and Baseline

Dataset	Compared Model	t-statistic	p-value	Avg. Diff.
Beauty	Modified V2	3.9853	0.0105	1.3845%
	Modified V1	4.0182	0.0101	1.4029%
Movielens-1M	Modified V2	6.8432	0.0010	1.6575%
	Modified V1	6.5228	0.0013	1.7198%
Toys	Modified V2	3.1720	0.0248	1.6547%
	Modified V1	3.1530	0.0253	1.5553%

The comparison between the two extended models reveals that while the introduction of the cross-attention mechanism consistently improved performance across various scenarios, the impact of adding offset noise was more variable. Although offset noise occasionally enhanced performance compared to the cross-attention-only model, in some cases, the difference was minimal. Notably, extended training with offset noise did not lead to overfitting and instead promoted better generalization; however, this did not always translate into substantial performance gains.

6.4. Overall Comparison (RQ2)

Table 5 presents a comprehensive comparison of our diffusion-based recommender system model, DiffuRecSys, with various conventional methods [19, 33, 34, 35, 36, 2, 37] across three public benchmark datasets. The results demonstrate that the diffusion-based approach consistently outperforms all conventional models on both sparse datasets (e.g., Beauty, Toys) and dense dataset (e.g., ML-1M). This performance improvement is attributed to several key factors, including the use of Probabilistic Item Representations, which involves sampling from random Gaussian noise to effectively model the uncertainty of user behaviors in real-world scenarios. Additionally, by applying noise to the target item rather than the entire sequence, crucial information related to user preferences is preserved. Finally, the Diverse Inference Strategy introduces variability into the recommendation results, further enhancing overall performance.

Figure 3 shows the t-distributed Stochastic Neighbor Embedding (t-SNE) visualization of the last target item for each user across different datasets, including Beauty, Toys, and

Figure 3: The Reconstructed Target Item Distributions after applying clustering for different users across various datasets.

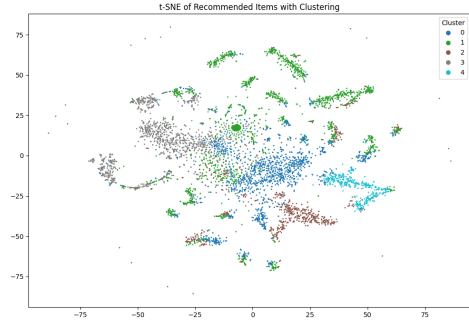


Figure 4: Amazon Beauty dataset




Figure 5: Amazon Toys dataset




Figure 6: Movielens1M dataset

ML-1M. This visualization provides insights into the distributions of the reconstructed target items across various user groups and datasets.

6.5. Case Study (RQ3)

In line with previous studies [38, 39, 40, 35], we classified the top 20% most frequent items as head items and designated the remainder as long-tail items for performance evaluation. The model exhibits reduced performance on long-tail items compared to head items, highlighting a persistent challenge in recommending less frequent items.

Additionally, we categorized the Amazon Beauty and

Table 5

Overall performance of different methods for the sequential recommendation. The best score and the second-best score in each row are bolded and underlined, respectively. The last column indicates improvements over the best baseline method.

Dataset	Metric	GRU4Rec	SASRec	BERT4Rec	STOSA	ACVAE	VSAN	MFGAN	CL4SRec	DuoRec	DiffuRecSys	Improv.
Beauty	HR@5	0.0206	0.0371	0.0370	0.0460	0.0438	0.0475	0.0382	0.0396	<u>0.0541</u>	0.0667	23.29%
	HR@10	0.0332	0.0592	0.0598	0.0659	0.0690	0.0759	0.0605	0.0630	<u>0.0825</u>	0.0980	18.79%
	HR@20	0.0526	0.0893	0.0935	0.0932	0.1059	0.1086	0.0916	0.0965	<u>0.1102</u>	0.1400	27.04%
	NDCG@5	0.0139	0.0233	0.0233	0.0318	0.0272	0.0298	0.0254	0.0232	<u>0.0362</u>	0.0458	26.52%
	NDCG@10	0.0175	0.0284	0.0306	0.0382	0.0354	0.0389	0.0310	0.0307	<u>0.0447</u>	0.0559	25.06%
	NDCG@20	0.0221	0.0361	0.0391	0.0451	0.0453	0.0471	0.0405	0.0392	<u>0.0531</u>	0.0665	25.24%
Toys	HR@5	0.0121	0.0429	0.0371	<u>0.0563</u>	0.0457	0.0481	0.0395	0.0503	0.0539	0.0684	21.49%
	HR@10	0.0184	0.0652	0.0524	<u>0.0769</u>	0.0663	0.0719	0.0641	0.0736	0.0744	0.0971	26.27%
	HR@20	0.0290	0.0957	0.0760	0.1006	0.0984	0.1029	0.0892	0.0990	<u>0.1056</u>	0.1350	27.84%
	NDCG@5	0.0077	0.0248	0.0259	<u>0.0393</u>	0.0291	0.0286	0.0257	0.0264	0.0340	0.0455	15.78%
	NDCG@10	0.0097	0.0320	0.0329	<u>0.0460</u>	0.0364	0.0363	0.0328	0.0339	0.0406	0.0548	19.13%
	NDCG@20	0.0123	0.0397	0.0368	<u>0.0519</u>	0.0432	0.0441	0.0381	0.0404	0.0472	0.0643	23.89%
MovieLens-1M	HR@5	0.0806	0.1078	0.1308	0.1230	0.1356	0.1220	0.1275	0.1142	<u>0.1679</u>	0.1957	16.56%
	HR@10	0.1344	0.1810	0.2219	0.1889	0.2033	0.2016	0.2086	0.1815	<u>0.2540</u>	0.2843	11.93%
	HR@20	0.2081	0.2745	0.3354	0.2724	0.3085	0.3015	0.3166	0.2818	<u>0.3478</u>	0.3958	13.80%
	NDCG@5	0.0475	0.0681	0.0804	0.0810	0.0837	0.0751	0.0778	0.0705	<u>0.1091</u>	0.1319	20.90%
	NDCG@10	0.0649	0.0918	0.1097	0.1040	0.1145	0.1007	0.1040	0.0920	<u>0.1370</u>	0.1605	17.15%
	NDCG@20	0.0834	0.1156	0.1384	0.1231	0.1392	0.1257	0.1309	0.1170	<u>0.1607</u>	0.1886	17.36%

MovieLens-1M datasets into five groups based on sequence length percentiles to assess model performance across different sequence lengths. In the Amazon Beauty dataset, the model generally performs better on longer sequences, as most sequences are relatively short. Conversely, in the MovieLens-1M dataset, performance gradually decreases as sequence length increases. This pattern suggests that shorter sequences may lack sufficient information to accurately predict user preferences, while excessively long sequences introduce challenges for model performance.

To further enhance recommendation performance, we propose averaging the predictions generated from different random seeds, akin to ensemble methods. This approach likely improves performance by accounting for various aspects of behavioral uncertainty through the aggregation of diverse recommendation outcomes. By leveraging this strategy, we can enhance the model’s performance on very long, short, or low-frequency items. Based on these findings, we propose an efficient inference procedure, detailed in Algorithm 3.

Algorithm 3: Efficient inference of i_{n+1}^u for user sequence s_u

Input: Inference sequence $s_u = [i_1^u, \dots, i_n^u]$, sequence length N , diffusion steps T , random seeds list **seeds**

Output: Average probability distribution

```

1 for  $i = 1$  to  $|\text{seeds}|$  do
2   Fix random seed as  $\text{seeds}[i]$ ;
3   Sample Gaussian noise  $\epsilon \sim \mathcal{N}(0, I)$ ;
4   Calculate  $\mathbf{x}_n^0$  via  $q_\phi(\mathbf{x}_n^0 | \mathbf{i}_n)$ ;
5   Calculate  $\mathbf{x}_n^T$  via  $q(\mathbf{x}_n^T | \mathbf{x}_n^0)$ ;
6   Predict  $\mathbf{x}_N^0$  via  $f_\theta(\mathbf{x}_n^T, T)$ ;
7   Obtain distribution  $\mathbf{p}_i \leftarrow \hat{p}(\mathbf{i}_{N+1} | \mathbf{x}_N^0)$ ;
8 return  $\frac{1}{|\text{seeds}|} \sum_{i=1}^{|\text{seeds}|} \mathbf{p}_i$ 

```

7. Conclusion

This work addresses the challenge of sequential recommendation systems, where existing single vector-based representations often fail to effectively capture the temporal dynamics and evolving user preferences. The introduction of diffusion models in sequential recommendation provides a more suitable framework for modeling the latent aspects of items and users’ multi-level interests.

In particular, our work enhanced the Diffusion Recommender model, DiffuRec, by incorporating cross-attention mechanisms between historical user interactions and target items within the Approximator component of the model architecture. This addition allows the model to dynamically weigh historical user interactions to target items, leading to more nuanced and context-aware representations. Additionally, we introduced offset noise into the diffusion process to improve model robustness, resulting in the new model, DiffuRecSys. Another area of interest is the exploration of guided diffusion models to further improve the accuracy and relevance of user preference modeling.

Our experiments demonstrate the effectiveness of DiffuRecSys. However, validating its performance across a diverse range of datasets and domains is essential for establishing its generalizability. As the field of recommendation systems continues to evolve, we believe our work will inspire further innovations and applications in user-centric recommendation technologies. In conclusion, DiffuRecSys represents a novel and effective advancement in sequential recommendation, setting a new standard for intelligent and personalized recommendation systems.

References

- [1] Z. Li, A. Sun, C. Li, Diffurec: A diffusion model for sequential recommendation, ACM Transactions on Information Systems 42 (2023) 1–28.
- [2] W.-C. Kang, J. McAuley, Self-attentive sequential recommendation, in: 2018 IEEE international conference on data mining (ICDM), IEEE, 2018, pp. 197–206.
- [3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial networks, Communications of the ACM 63 (2020) 139–144.
- [4] D. P. Kingma, M. Welling, Auto-encoding variational bayes, arXiv preprint arXiv:1312.6114 (2013).
- [5] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, Advances in neural information processing systems 33 (2020) 6840–6851.
- [6] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in: International conference on machine learning, PMLR, 2015, pp. 2256–2265.

- [7] R. Cai, G. Yang, H. Averbuch-Elor, Z. Hao, S. Belongie, N. Snavely, B. Hariharan, Learning gradient fields for shape generation, in: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, Springer, 2020, pp. 364–381.
- [8] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, T. Salimans, Cascaded diffusion models for high fidelity image generation, *Journal of Machine Learning Research* 23 (2022) 1–33.
- [9] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, D. J. Fleet, Video diffusion models, *Advances in Neural Information Processing Systems* 35 (2022) 8633–8646.
- [10] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, M. Chen, Hierarchical text-conditional image generation with clip latents, *arXiv preprint arXiv:2204.06125* 1 (2022) 3.
- [11] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022*, pp. 10684–10695.
- [12] M. Sato, S. Takemori, J. Singh, T. Ohkuma, Unbiased learning for the causal effect of recommendation, in: *Proceedings of the 14th ACM conference on recommender systems, 2020*, pp. 378–387.
- [13] W. Wang, F. Feng, X. He, L. Nie, T.-S. Chua, Denoising implicit feedback for recommendation, in: *Proceedings of the 14th ACM international conference on web search and data mining, 2021*, pp. 373–381.
- [14] P. Lops, M. De Gemmis, G. Semeraro, Content-based recommender systems: State of the art and trends, *Recommender systems handbook* (2011) 73–105.
- [15] G. Linden, B. Smith, J. York, Amazon.com recommendations: Item-to-item collaborative filtering, *IEEE Internet computing* 7 (2003) 76–80.
- [16] B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: *Proceedings of the 10th international conference on World Wide Web, 2001*, pp. 285–295.
- [17] R. Burke, Hybrid recommender systems: Survey and experiments, *User modeling and user-adapted interaction* 12 (2002) 331–370.
- [18] S. Rendle, C. Freudenthaler, L. Schmidt-Thieme, Factorizing personalized markov chains for next-basket recommendation, in: *Proceedings of the 19th international conference on World wide web, 2010*, pp. 811–820.
- [19] B. Hidasi, A. Karatzoglou, L. Baltrunas, D. Tikk, Session-based recommendations with recurrent neural networks, *arXiv preprint arXiv:1511.06939* (2015).
- [20] J. Wang, L. Yu, W. Zhang, Y. Gong, Y. Xu, B. Wang, P. Zhang, D. Zhang, Irgan: A minimax game for unifying generative and discriminative information retrieval models, in: *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval, 2017*, pp. 515–524.
- [21] B. Jin, D. Lian, Z. Liu, Q. Liu, J. Ma, X. Xie, E. Chen, Sampling-decomposable generative adversarial recommender, *Advances in Neural Information Processing Systems* 33 (2020) 22629–22639.
- [22] J. Ma, C. Zhou, P. Cui, H. Yang, W. Zhu, Learning disentangled representations for recommendation, *Advances in neural information processing systems* 32 (2019).
- [23] D. Liang, R. G. Krishnan, M. D. Hoffman, T. Jebara, Variational autoencoders for collaborative filtering, in: *Proceedings of the 2018 world wide web conference, 2018*, pp. 689–698.
- [24] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, M. Welling, Improved variational inference with inverse autoregressive flow, *Advances in neural information processing systems* 29 (2016).
- [25] A. Q. Nichol, P. Dhariwal, Improved denoising diffusion probabilistic models, *arXiv preprint arXiv:2102.09672* (2021).
- [26] P. Dhariwal, A. Q. Nichol, Diffusion models beat gans on image synthesis, in: *Advances in Neural Information Processing Systems, volume 34, 2021*, pp. 8780–8794.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [28] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, Springer, 2015, pp. 234–241.
- [29] R. Strudel, C. Tallec, F. Altché, Y. Du, Y. Ganin, A. Mensch, W. Grathwohl, N. Savinov, S. Dieleman, L. Sifre, et al., Self-conditioned embedding diffusion for text generation, *arXiv preprint arXiv:2211.04236* (2022).
- [30] R. K. Mahabadi, H. Ivison, J. Tae, J. Henderson, I. Beltagy, M. E. Peters, A. Cohan, Tess: Text-to-text self-conditioned simplex diffusion, *arXiv preprint arXiv:2305.08379* (2023).
- [31] X. Han, S. Kumar, Y. Tsvetkov, Ssd-lm: Semi-autoregressive simplex-based diffusion language model for text generation and modular control, *arXiv preprint arXiv:2210.17432* (2022).
- [32] W. X. Zhao, S. Mu, Y. Hou, Z. Lin, Y. Chen, X. Pan, K. Li, Y. Lu, H. Wang, C. Tian, et al., Recbole: Towards a unified, comprehensive and efficient framework for recommendation algorithms, in: *proceedings of the 30th acm international conference on information & knowledge management, 2021*, pp. 4653–4664.
- [33] X. Xie, F. Sun, Z. Liu, S. Wu, J. Gao, J. Zhang, B. Ding, B. Cui, Contrastive learning for sequential recommendation, in: *2022 IEEE 38th international conference on data engineering (ICDE), IEEE, 2022*, pp. 1259–1273.
- [34] R. Ren, Z. Liu, Y. Li, W. X. Zhao, H. Wang, B. Ding, J.-R. Wen, Sequential recommendation with self-attentive multi-adversarial network, in: *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval, 2020*, pp. 89–98.
- [35] Z. Fan, Z. Liu, Y. Wang, A. Wang, Z. Nazari, L. Zheng, H. Peng, P. S. Yu, Sequential recommendation via stochastic self-attention, in: *Proceedings of the ACM Web Conference 2022, 2022*, pp. 2036–2047.
- [36] Z. Xie, C. Liu, Y. Zhang, H. Lu, D. Wang, Y. Ding, Adversarial and contrastive variational autoencoder for sequential recommendation, in: *Proceedings of the Web Conference 2021, 2021*, pp. 449–459.
- [37] R. Qiu, Z. Huang, H. Yin, Z. Wang, Contrastive learning for representation degeneration problem in sequential recommendation, in: *Proceedings of the fifteenth ACM international conference on web search and data mining, 2022*, pp. 813–823.

- [38] Y. Zhang, D. Z. Cheng, T. Yao, X. Yi, L. Hong, E. H. Chi, A model of two tales: Dual transfer learning framework for improved long-tail item recommendation, in: Proceedings of the web conference 2021, 2021, pp. 2220–2231.
- [39] J. Zhao, P. Zhao, L. Zhao, Y. Liu, V. S. Sheng, X. Zhou, Variational self-attention network for sequential recommendation, in: 2021 IEEE 37th International Conference on Data Engineering (ICDE), IEEE, 2021, pp. 1559–1570.
- [40] N. Sachdeva, G. Manco, E. Ritacco, V. Pudi, Sequential variational autoencoders for collaborative filtering, in: Proceedings of the twelfth ACM international conference on web search and data mining, 2019, pp. 600–608.