



# Prompt-to-Slate: Diffusion Models for Prompt-Conditioned Slate Generation

Federico Tomasi  
Spotify  
London, United Kingdom  
federicot@spotify.com

Francesco Fabbri  
Spotify  
Barcelona, Spain  
francescof@spotify.com

Justin Carter  
Spotify  
New York, USA  
jcarter@spotify.com

Elias Kalomiris  
Spotify  
New York, USA  
eliask@spotify.com

Mounia Lalmas  
Spotify  
London, United Kingdom  
mounia@acm.org

Zhenwen Dai  
Spotify  
London, United Kingdom  
zhenwend@spotify.com

## Abstract

*Slate generation* is a common task in streaming and e-commerce platforms, where multiple items are presented together as a list or “slate”. Traditional systems focus mostly on item-level ranking and often fail to capture the coherence of the slate as a whole. A key challenge lies in the combinatorial nature of selecting multiple items jointly. To manage this, conventional approaches often assume users interact with only one item at a time, assumption that breaks down when items are meant to be consumed together.

In this paper, we introduce DMSG, a generative framework based on diffusion models for prompt-conditioned slate generation. DMSG learns high-dimensional structural patterns and generates coherent, diverse slates directly from natural language prompts. Unlike retrieval-based or autoregressive models, DMSG models the joint distribution over slates, enabling greater flexibility and diversity.

We evaluate DMSG in two key domains: music playlist generation and e-commerce bundle creation. In both cases, DMSG produces high-quality slates from textual prompts without explicit personalization signals. Offline and online results show that DMSG outperforms strong baselines in both relevance and diversity, offering a scalable, low-latency solution for prompt-driven recommendation. A live A/B test on a production playlist system further demonstrates increased user engagement and content diversity.

## Keywords

diffusion model, slate generation, music playlist generation, bundle generation

## ACM Reference Format:

Federico Tomasi, Francesco Fabbri, Justin Carter, Elias Kalomiris, Mounia Lalmas, and Zhenwen Dai. 2025. Prompt-to-Slate: Diffusion Models for Prompt-Conditioned Slate Generation. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys '25)*, September 22–26, 2025, Prague, Czech Republic.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

RecSys '25, Prague, Czech Republic

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1364-4/25/09

<https://doi.org/10.1145/3705328.3748072>

Prague, Czech Republic. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3705328.3748072>

## 1 Introduction

Slate generation is an increasingly important research area, designed for scenarios where multiple items are recommended together, such as in streaming platforms like YouTube and e-commerce sites like Amazon. Unlike traditional ranking-based approaches that evaluate each candidate individually with respect to a query, slate generation focuses on selecting a group of items that collectively maximize user satisfaction [12, 35]. This involves optimizing the entire slate of items shown to a user accounting for factors like co-consumption patterns and diversity across items.

Increasingly, real-world systems are being tasked with generating slates in response to prompts, themes, or editorial intents, and often without access to user interaction data. This shift is driven in part by the rise of language-based interfaces, which require models that can flexibly respond to natural language input. However, most existing methods, especially retrieval- and ranking-based approaches, lack the generative flexibility needed in these scenarios. As a result, there is growing interest in models capable of generating entire slates from scratch, while preserving the structural coherence that makes slates meaningful.

A major challenge in slate generation lies in its combinatorial nature: jointly optimizing all items in a slate causes the number of possible item combinations to grow exponentially. Traditional methods often address this by assuming that users interact with only a single item from the slate, treating subsequent interactions as separate but conditionally dependent events [12]. While this assumption simplifies modeling and reduces computational complexity, it overlooks scenarios where users engage with multiple items simultaneously. This limitation is particularly evident in applications like music playlist creation [36] or bundle recommendations [33, 34], where items are naturally consumed together.

We propose a *prompt-based slate generation model*, where the objective is to generate a slate of items conditioned solely on a natural language prompt. This framework addresses settings where historical user data is unavailable, insufficient, or used indirectly, for example by incorporating user preference in a textual format.

We formulate slate generation as a generative modeling problem and employ *diffusion models (DMs)* to sample coherent and diverse slates from a learned distribution, conditioned on a prompt. DMs

are well-suited to this problem: they excel at modeling complex structures in high-dimensional spaces and generating diverse, high-quality outputs [11, 29]. Our novel application of DMs to prompt-based slate generation eliminates the need for explicit combinatorial optimization. Instead, the model learns the implicit structure of desirable slates from the data, enabling the direct generation of item sets directly from prompts.

A key strength of this approach lies in its stochasticity-driven diversity: the model can generate multiple high-quality, yet distinct slates for the same prompt. This capability can enhance content freshness and promote the discovery of less popular items—an especially valuable trait in large catalogs where many relevant combinations may exist. Our evaluations show that this generation process improves both relevance and diversity, ultimately enhancing the user discovery experience [10].

Our framework, *Diffusion Model for Slate Generation (DMSG)*, generates a distribution of slates conditioned on a natural language prompt. To support dynamic, large-scale catalogs with millions of items, we represent items as embeddings in a continuous space and train the diffusion model on concatenated item embeddings. At inference time, the generated latent representations are decoded into catalog items using a decoding module. This design ensures that the core model remains stable as the catalog evolves—only the item mapping function needs updating, making it adaptable to real-world deployments. Moreover, this approach enables efficient training through the use of closed-form lower bounds [15, 25]. To ensure the system is suitable for low-latency, real-time recommendation, we incorporate DDIM [30], a noise scheduler that significantly accelerates the generation process while maintaining output quality.

We evaluate DMSG across two key tasks: music playlist generation and e-commerce bundle creation. Our model consistently outperforms state-of-the-art retrieval-based baselines, achieving improvements of up to +12.9% in average precision and +17% in NDCG over the next-best model. DMSG also produces high-quality recommendations, with an average BertScore of 0.8, while delivering strong content freshness. To validate the approach in a real-world environment, we ran a two-week live A/B test on a production music playlist generation system. Results demonstrated tangible benefits: a −13.4% reduction in duplicated recommended content and a +6.8% increase in active user interactions, highlighting both the quality and impact on users of prompt-based slate generation.

The remaining of this paper is structured as follows. Section 2 reviews related work and the state of the art in slate generation and diffusion models. Section 3 formalizes the slate generation problem. Sections 4 and 5 introduce our methodology based on continuous diffusion models. Section 6 presents our offline experiments, while Section 7 reports on our online evaluation. Finally, Section 8 discusses key contributions and outlines directions for future work.

## 2 Related Work

We review relevant literature, highlighting key methodologies and findings, as well as the gaps our work aims to address.

*Slate Generation.* Recommendation systems are typically trained to present users with the most relevant items, often as a ranked

list. In such lists, top-ranked items receive disproportionate attention due to positional bias [9]. However, in applications like music playlist creation [6, 28] or e-commerce bundles [33], users are expected to engage with most or all recommended content. In these cases, the combination of items, known as a slate, becomes essential to optimize. This is known as the *slate generation* problem [12, 35, 36].

Several approaches have been proposed to address slate generation. Reinforcement learning (RL) is a common strategy [12], though it often depends on simplifying assumptions, such as modeling users as selecting only the “best” item from a slate, to reduce the combinatorial complexity. Some methods employ user preference models to guide slate construction and train RL agents [36]. However, RL-based methods typically require extensive interaction data and strong modeling assumptions, limiting their practicality. Recent work explores transformer-based models that bypass RL altogether [7]. While prior work largely focuses on personalization using historical interactions, our approach diverges by generating slates from scratch using only natural language prompts, enabling applicability even in settings where personalization is infeasible.

*Generative Modeling.* Generative models, particularly large language models (LLMs), have shown remarkable generalization and output quality across language-centric tasks. Transformer-based, decoder-only architectures [3] have become standard and are increasingly adapted for recommendation [17]. These models offer a compelling alternative to RL, as they do not require predefined reward functions or suffer from slow convergence. They also incorporate textual context, such as user queries or profile descriptions.

While LLMs are adaptable for recommendation, they often fall short in large-scale retrieval tasks [4] due to mismatches between their general-purpose knowledge and the domain-specific nature of catalogs. In particular, LLMs require prior exposure to catalog items or external mappings between textual and catalog representations [4]. This makes them less practical in dynamic real-world settings where catalogs frequently change and models must scale to unseen items [41]. In contrast, our approach conditions generation on text prompts but employs diffusion models rather than LLM-style decoding. Diffusion models are better suited for learning the structure of high-cardinality, dynamic item spaces, making them more robust and scalable for real-world deployment.

*Diffusion Models for Recommendations.* Diffusion models (DMs) [25] are a powerful class of generative models capable of producing structured, high-dimensional outputs, such as images or audio, from text prompts. Their ability to generalize to out-of-distribution prompts, coupled with their inherent stochasticity, makes them an attractive option for recommendation tasks. Recent studies show that DMs can outperform collaborative filtering and traditional retrieval methods by generating diverse, coherent user-item interaction pairs [39]. Some efforts have extended DMs to sequential recommendation [18, 40], where models generate the next item in a user sequence based on historical data. However, these methods still rely heavily on personalized signals, limiting their applicability in data-sparse settings.

Our work is the first to apply diffusion models to prompt-based slate generation, where full item slates are generated from scratch

using only natural language input. This enables our model to operate effectively in cold-start and editorial settings, where personalization is not feasible or desired. By modeling the joint distribution over entire slates, rather than individual items, our approach captures higher-order dependencies and enables flexible, diverse slate generation from a single prompt.

This expands the role of diffusion models beyond personalized sequence generation to broader applications like content curation, editorial tooling, and theme-based recommendation—especially in dynamic environments with limited behavioral data [41].

*Diversity in Recommender Systems.* Diversity has become a key concept in recommendation systems, often considered as important as relevance [37]. Prior work shows that exposure to diverse content correlates with long-term user retention [1], while low diversity can lead to filter bubbles and echo chambers [24, 31]. Recent studies argue for treating diversity not merely as a list-level metric but as a mechanism for user exploration and knowledge discovery. For example, Coppolillo et al. [8] model recommendation as a process of knowledge acquisition, where diverse recommendations maximize informational gain. Parallel efforts have explored using LLMs for diversity-aware re-ranking. Carraro and Bridge [5] show that LLM-based re-rankers can improve diversity over standard ranking pipelines, though they often underperform compared to specialized models. These findings suggest that generation-based approaches may play an increasingly important role in diversity-aware recommendation.

Our method contributes to this line of work by using diffusion models to stochastically generate slates rather than re-ranking fixed outputs. This generative process naturally yields slates that are both relevant and diverse, as multiple valid item combinations can be sampled from the model’s learned distribution. As shown in our experiments, this leads to high-quality recommendations that significantly improve content freshness and variety, without requiring explicit diversity constraints [13].

### 3 Problem Formulation

We address the problem of slate generation, defined as selecting a list (or slate) of items to present to a user with the goal of maximizing user satisfaction. In our offline evaluation, we quantify user satisfaction through two key metrics: relevance and diversity. Slate generation plays a central role in many practical recommender systems, including bundle recommendation [2, 33], video recommendation [12], and music playlist generation [28, 36]. The objective in these settings is to construct the most effective combination of items under the assumption that users will engage with a substantial portion (or the entirety) of the recommended slate.

We focus on two example domains: music playlist generation and e-commerce bundle recommendations. We refer to “catalog items” as the basic units of recommendation, which could be music tracks in a playlist or products on an e-commerce platform. Catalogs often span thousands to millions of items, with new ones being continuously added or removed, especially in dynamic environments like streaming services and online retail platforms.

We assume access to data in the form of item sequences, each paired with a textual description (*labels*). Each item is associated with metadata, such as title, artist, and release date for music, or

category and product information in e-commerce. The task is to generate a slate of items that aligns with a text query, optimizing user satisfaction without relying on historical user interaction data.

We formulate the problem using continuous diffusion processes (DMs) [25], inspired by their demonstrated ability to generate high-quality, structured content such as images and audio. This motivates our exploration of their potential in the context of slate generation.

### 4 Preliminaries

Diffusion models [11] are probabilistic generative models based on Markov chains, trained to learn a data distribution  $q(\mathbf{x}_0)$  where  $\mathbf{x}_0 \in \mathbb{R}^d$ . These models generate samples through a finite sequence of steps, ultimately producing an output distribution  $p_\theta(\mathbf{x}_0) = \int p_\theta(\mathbf{x}_{0:T}) d\mathbf{x}_{1:T}$ , where  $\mathbf{x}_1, \dots, \mathbf{x}_T$  are latent variables with the same dimensionality as the data  $\mathbf{x}_0 \sim q(\mathbf{x}_0)$ . DMs consist of two key components: the forward and the reverse process. The forward process, also known as the diffusion process, is a fixed, non-learned Markov chain  $q(\mathbf{x}_{1:T}|\mathbf{x}_0)$  that progressively adds Gaussian noise to the input data over time. This process is governed by a predefined variance schedule  $\beta_1, \dots, \beta_T$  and is defined as:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) := \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}),$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$$

The forward process incrementally adds a fixed amount of noise at each step, progressively diminishing the signal from the original data. Importantly, sampling  $\mathbf{x}_t$  at step  $t$  can be done in closed form:

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I})$$

where  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=1}^t \alpha_s$ . Using the reparametrization trick, we can sample  $\mathbf{x}_t$  in closed form as:  $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t$ , where  $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

The reverse process models the joint distribution  $p_\theta(\mathbf{x}_{0:T})$  as a Markov chain with Gaussian transitions, starting from  $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$ . In practical terms, this means that DMs gradually denoise a normally distributed sample  $p(\mathbf{x}_T)$  over  $T$  steps to recover a sample from the data distribution. The full reverse process is defined as:

$$p_\theta(\mathbf{x}_{0:T}) := p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t),$$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t)),$$

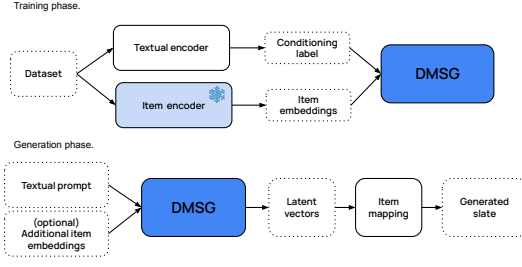
where  $\boldsymbol{\mu}_\theta$  and  $\Sigma_\theta$  are functions parameterized by neural networks. Following from [29], we train only  $\boldsymbol{\mu}_\theta$  and set  $\Sigma_\theta = \sigma_t^2 \mathbf{I}$  using a small fixed variance.

Training is performed by optimizing the variational bound on the negative log likelihood,

which can be further simplified by ignoring the variance terms and edge effects, as shown in [11]:

$$L_{\text{simple}}(\theta) := \mathbb{E}_{t, \mathbf{x}_0, \epsilon} \left[ \left\| \epsilon - \epsilon_\theta \left( \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t \right) \right\|^2 \right] \quad (1)$$

where  $t \sim \mathcal{U}(1, T)$  is uniformly sampled between 1 and  $T$ . This approach stabilizes training, making it more efficient and leading



**Figure 1: Overview of DMSG during training (top) and inference (bottom) phases.**

to higher-quality generation. The network  $\epsilon_\theta$  is trained to approximate the true reverse distribution  $p_\theta$ . Different parametrizations of  $p_\theta$  are possible, each corresponding to different loss functions. In our model, we adopt the *v-prediction* [27]. Specifically, we predict the velocity value  $v_t = \alpha_t \epsilon_t - \sigma_t x_t$ , with a slight modification of the original loss function:

$$L_{\text{DMSG}}(\theta) := \mathbb{E}_{t, x_0, v} \left[ \left\| v - v_\theta \left( \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon, t, c \right) \right\|^2 \right], \quad (2)$$

where  $c$  represents the conditional information. This formulation correspond to “SNR+1” weighting [27]. Empirically, we found that this loss function leads to better convergence and overall training performance compared to directly predicting  $\epsilon$  or  $x$ .

Once the model  $p_\theta$  is trained, the diffusion process can be used to generate new samples starting from  $x_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , and then iteratively sampling  $x_{t-1} \sim p_\theta(x_{t-1}|x_t)$  for  $t \in [0, T]$ .

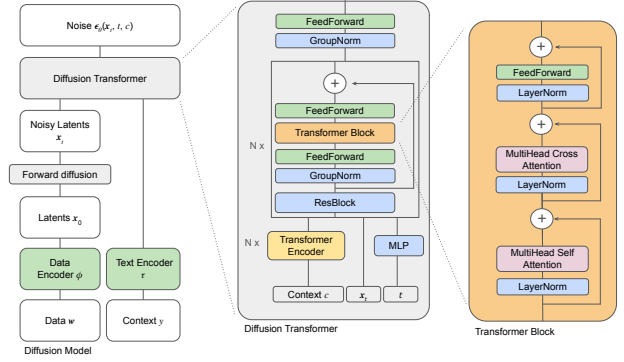
## 5 Methodology

Slate generation refers to the task of sampling an ordered list of items from a potentially large and dynamic catalog. Each slate consists of items grouped by shared characteristics and may vary in length. For simplicity, we standardize the input by padding sequences to a fixed length and discarding items from slates that exceed this maximum size. Each slate is also paired with a textual description that serves as the conditioning context.

Our model, which we call the Diffusion Model for Slate Generation (DMSG), is trained to reconstruct slates based on this conditioning input. An overview of DMSG is shown in Figure 1. DMSG consists of three main components: (i) an *encoding module*, which maps discrete catalog items into a continuous latent space; (ii) a *conditioning module*, which encodes the textual context into a fixed-length vector; and (iii) a *diffusion process module*, a probabilistic generative model that iteratively transforms a latent vector into Gaussian noise and back, enabling the generation of coherent item sequences.

**Encoding Module.** We assume access to a discrete sequence  $\mathbf{w}$  of length  $n$ , where each  $w_i \in \mathcal{V}$  is an item selected from a catalog  $\mathcal{V}$  (also referred to as *vocabulary*), similarly to text generation setting. Importantly, we do not assume a fixed vocabulary; instead, the catalog can dynamically change during generation.

To use diffusion models, we convert the discrete sequence  $\mathbf{w}$  into a continuous representation  $x$  in a low-dimensional latent space [15]. This is achieved through an embedding function  $\phi$ ,



**Figure 2: Diffusion model and transformer details. The core module is a diffusion transformer to predict the noise artificially added to latent variables. At generation time, this is used to denoise a latent vector.**

which maps each item  $w_i$  to a continuous vector in  $x_i \in \mathbb{R}^d$ . The full sequence embedding is thus  $\phi(\mathbf{w}) = \{\phi(w_1), \dots, \phi(w_n)\}$ , where  $\phi(\mathbf{w}) \in \mathbb{R}^{n \times d_\phi}$ . This process is flexible and can be implemented in various ways to support different types of input data [16, 21].

In alignment with the continuous diffusion model formulation described in Section 4, we define the latent encoding of each item  $w_i$  as  $x_i = \phi(w_i)$ . This enables the direct use of diffusion models in the continuous domain, where  $x_i$  serves as the representation of catalog item  $w_i$ .

We explored two types of encoding modules: 1) a trainable encoder, jointly optimized with the diffusion model, and 2) a fixed encoder, trained separately and held constant during the DM training. Training the encoder jointly with the diffusion model introduces additional parameters and requires modifying the optimization objective (see Equation (1)). However, preliminary experiments [15, 25, 32] suggest that using a fixed encoder offers practical benefits, including improved training stability and comparable performance. Therefore, in this work, we adopt a fixed encoding module, leaving a comprehensive comparison with trainable alternatives for future work.

**Conditioning Module.** To guide the diffusion process toward generating contextually appropriate slates, we model conditional distributions of the form  $p(x|c)$ , where  $c$  represents the *context* variable [25]. The context can take various forms, such as a textual description (also referred to as *prompt*) or additional related catalog items. By conditioning on such information, the model can generate outputs that align with specific semantic intents [23].

We preprocess the conditional text using a domain-agnostic encoder  $\tau$ , which maps an input variable  $y$  (e.g., a textual prompt) into a vector representation  $\tau(y) \in \mathbb{R}^{M \times d_\tau}$ . For simplicity, we refer to this encoded vector as context  $c$ . To implement  $\tau$ , we use a stack of transformer modules, well suited for capturing the structure and semantics of textual inputs [38]. The resulting context vector  $c$  is used to condition the diffusion process via a cross-attention mechanism within the transformer-based architecture. Figure 2 shows a detailed illustration of DMSG and its components.

*Diffusion Process.* At the core of our model is a conditional diffusion process—a probabilistic generative model that iteratively transforms a random sample drawn from a Gaussian distribution into a structured sample from a target distribution, represented in a compact latent space [29]. As described in Section 4, the diffusion process consists of two components: a forward, fixed components that gradually adds noise to the latent variables, and a learnable reverse process that removes this noise to recover the original signal. To implement the reverse process, we use a transformer-based architecture designed to capture sequential dependencies within the slate and enable scalable training. This design follows recent successful applications of transformers in image generation tasks [22]. We refer to this core component as the *diffusion transformer*.

The diffusion transformer incorporates a cross-attention mechanism within its core blocks [38], similar to latent DMs that condition generation on external information such as class labels and textual prompts [25]. In our setup, the transformer processes the context vector  $c$ , the noisy latent input  $x_t$  and the diffusion step  $t$ , and is trained to predict the noise component originally added to the clean sample  $x_0$ . This module, denoted  $v_\theta(x_t, t, c)$ , predicts the *velocity* (described in Section 4), and is ultimately used to reconstruct the original latent sample  $x_0$ .

At generation time, the model takes the conditioning information, denoted as  $y$ , typically in the form of natural language. The input is processed by the text encoder  $\tau$ , producing a latent representation  $c = \tau(y)$ . To initiate the generation process, we sample an initial latent vector  $x_T \sim \mathcal{N}(0, I)$  from a standard Gaussian distribution and iteratively apply the reverse diffusion process (or a subset of its steps) to generate a new sample. Each step is conditioned on  $c$ , ensuring that the final output reflects the original prompt  $y$ . Formally, the reverse process at each timestep computes:

$$x_{t-1} = \frac{1}{1 - \bar{\alpha}_t} \left( \sqrt{\bar{\alpha}_{t-1}} \beta_t \hat{x}_0 + \sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1}) x_t \right)$$

where the predicted clean sample  $\hat{x}_0$  is given by:

$$\hat{x}_0 = \sqrt{\bar{\alpha}_t} x_t - \sqrt{1 - \bar{\alpha}_t} v_\theta(x_t, t, c).$$

After the final step, the continuous representation  $x_0$  is converted into a discrete slate  $\mathbf{w}$  using a fixed (non-learned) *rounding* step, parametrized by  $p_{\mathbf{w}}(\mathbf{w}|x_0)$  [15].

This decoding step enables flexibility in handling dynamic catalogs at inference time, as new items can be added or removed without retraining the model. By simply updating the mapping function  $p_{\mathbf{w}}(\mathbf{w}|x_0)$ , the system can adjust to changes in the item space. For example, to avoid recommending duplicate items within a single session, small variations can be introduced to  $x_0$ . If a nearest neighbor mapping is used, it may yield the same items repeatedly; in such cases, duplicate recommendations can be masked, and the next nearest item in the embedding space can be selected instead.

*Noise Scheduler.* A well-known limitation of DMs is the need to sequentially iterate through all  $T$  diffusion steps. While using a large  $T$  can improve training stability and sample quality, it also requires the reverse diffusion process to traverse each step  $t \in \{1, \dots, T\}$ , often numbering in the thousands. This makes generation computationally expensive, posing challenges for low-latency, user-facing applications.

To address this, we leverage fast sampling techniques, specifically denoising diffusion implicit models (DDIMs) [30]. DDIMs generalize the Markovian diffusion process of DDPM to a non-Markovian formulation that retains the same training objective. This enables us to skip intermediate steps during inference, bypassing the full step-by-step trajectory used during training. Using DDIMs, we achieve fast generation times (on the order of milliseconds per sample) while maintaining comparable sample quality. In practice, we reduce the number of inference steps to 50, allowing us to meet the strict latency requirements of real-time recommendation systems.

## 6 Offline Experiments

We evaluate DMSG on two distinct slate generation tasks: music playlist generation and e-commerce bundle recommendation. Our experiments are conducted using a subset of (prompt, slate) pairs, and assume that all information about the relevance of the slate is encapsulated within the prompt itself.

### 6.1 Experimental setup

*Baselines.* We compared our diffusion-based approach with several unpersonalized baselines that also generate recommended slates from textual prompts: (1) **Popularity**: A heuristic that recommends the most popular items in the dataset, independent of the input prompt. Popularity is computed using the frequency of items within the dataset. (2) **Prompt2Vec**: An encoder-only model that maps a sentence embedding of the input query through two fully connected layers to produce a dense vector. Similar to DMSG, we retrieve the closest catalog items using a nearest neighbor mapping in the embedding space. The model is trained at the item level using associated metadata as input. (3) **BM25**: A traditional retrieval method that scores items based on their textual relevance to the prompt, using the BM25 ranking function [24]. (4) **S2S**: A sequence-to-sequence model built on a transformer architecture [38], consisting of an encoder-decoder pair. The encoder processes the textual prompt, while the decoder generates a probability distribution over the item vocabulary, from which we select the top candidates for each prompt to generate the slate.

*Datasets.* We evaluate our approach on the following datasets: (1) **Spotify Million Playlist Dataset (MPD)**: This dataset contains 1 million playlists created by Spotify users between January 2010 and October 2017, and include playlist title, description, and ordered lists of tracks with corresponding artist and track names.<sup>1</sup> For training and evaluation efficiency, we randomly sample a subset of 100K playlists for our experiments. (2) **Curated playlists (Curated)**: A proprietary dataset from a popular music streaming platform, comprising 30,000 high-quality, editorially curated playlists. Each playlist includes the same metadata as MPD. (3) **Bundle Recommendation Dataset (Bundle)**: Based on Amazon datasets, this dataset includes 11,000 unique catalog items that appear in 5,444 bundles [33, 34].<sup>2</sup>

Each dataset is split into training and test sets using a 80-20 ratio. The split is slate-wise, where 80% of the slates are used for training,

<sup>1</sup> [aicrowd.com/challenges/spotify-million-playlist-dataset-challenge](https://aicrowd.com/challenges/spotify-million-playlist-dataset-challenge)

<sup>2</sup> [github.com/BundleRec/bundle\\_recommendation](https://github.com/BundleRec/bundle_recommendation)

and 20% for the test set which is used exclusively for evaluation. All reported metrics are computed on the test set.

**Model training.** We trained DMSG for 1,000 epochs using the Adam optimizer [14] with a learning rate of  $10^{-4}$ . The velocity prediction model  $v_\theta$  consists of three transformer layers (Figure 2), each preceded by a 320-dimensional residual block with *swish* activation and group normalization (20 groups). Each transformer module includes a group normalization layer (20 groups), 10 heads with 32 dimensions each, a single transformer block between two dense layers, and a final *swish* activation and group normalization layer at the end of the sequence. We encode the context with 3 encoder-only transformer layers, each with 8 attention heads, a hidden dimension of 128, and a 10% dropout rate with a ReLU activation.

To generate training labels for each slate, we used item metadata—including title, descriptions, and curated keywords—processed through a pretrained text encoder.<sup>3</sup> This encoder supports input sequences of up to 77 tokens [25]. This setup effectively captures the nuances of textual prompts in our target applications. However, other domains may require different encoders to accommodate longer inputs. Notably, our framework supports plug-and-play integration of arbitrary text encoders with minimal or no changes to the model architecture.

For catalog item embeddings, we set the encoding dimension to 80, with a sequence length of 32 for the Curated and MPD datasets, and 5 for the Bundle dataset. This is the maximum slate size we consider, padding smaller sequences and trimming longer sequences to such sequence length. For Curated and MPD datasets we use a Word2Vec-based encoder that encodes tracks based on co-occurrence in playlists [20], while for Bundle dataset we use a pre-trained LLM<sup>4</sup> to encode item textual features in a zero-shot setting. For comparison, we also trained the S2S and Prompt2Vec baselines for 1,000 epochs using the Adam optimizer with a learning rate of  $10^{-4}$ . Prompt2Vec uses a single hidden layer dimension of 4096, while S2S uses an output dimensionality of 2048 for both its decoder and encoder components.

## 6.2 Results

We present results evaluating the effectiveness of DMSG across our two slate generation tasks: music playlist generation and e-commerce bundle recommendation.

**6.2.1 Is DMSG effective at generating relevant slates?** We start by evaluating the relevance of items generated by DMSG in response to a given prompt. We assume that all information needed to assess the relevance of a slate is contained within the prompt itself. This approach makes DMSG directly comparable to retrieval-based baselines such as BM25, which also rely solely on the input prompt.

To evaluate relevance, we compare DMSG against baseline models using reference oracle slates and standard ranking metrics. For each model, we generate candidate slates conditioned on the reference prompt and assess how closely the generated items match the reference slate. Catalog items are considered relevant if they appear in the reference slate, aligning with common recommender system

Curated	NDCGSim@30	MAPSim@30	BertScore	CategorySim@30
Popularity	0.0800 (+/- 0.11)	0.1010 (+/- 0.16)	0.2952 (+/- 0.23)	0.0057 (+/- 0.04)
Prompt2Vec	0.0674 (+/- 0.10)	0.0762 (+/- 0.14)	0.2670 (+/- 0.22)	0.0048 (+/- 0.02)
BM25	0.6054 (+/- 0.23)	0.7955 (+/- 0.23)	0.8928 (+/- 0.14)	0.4677 (+/- 0.27)
S2S	0.3761 (+/- 0.28)	0.4763 (+/- 0.34)	0.6531 (+/- 0.31)	0.2153 (+/- 0.24)
DMSG (ours)	<b>0.7087 (+/- 0.18)</b>	<b>0.8978 (+/- 0.11)</b>	<b>0.9314 (+/- 0.06)</b>	<b>0.4783 (+/- 0.24)</b>

MPD	NDCGSim@30	MAPSim@30	BertScore	CategorySim@30
Popularity	0.0064 (+/- 0.02)	0.0602 (+/- 0.14)	0.0657 (+/- 0.16)	0.0066 (+/- 0.03)
Prompt2Vec	0.0054 (+/- 0.02)	0.0289 (+/- 0.06)	0.0684 (+/- 0.17)	0.0094 (+/- 0.02)
BM25	<b>0.5614 (+/- 0.15)</b>	0.7373 (+/- 0.13)	0.8444 (+/- 0.10)	0.3674 (+/- 0.20)
S2S	0.3506 (+/- 0.16)	0.5247 (+/- 0.23)	0.7762 (+/- 0.17)	0.2337 (+/- 0.21)
DMSG (ours)	0.5210 (+/- 0.14)	<b>0.8119 (+/- 0.09)</b>	<b>0.8796 (+/- 0.06)</b>	<b>0.4792 (+/- 0.18)</b>

Bundle	NDCGSim@5	MAPSim@5	BertScore	CategorySim@5
Popularity	0.0870 (+/- 0.09)	0.1410 (+/- 0.13)	0.1659 (+/- 0.14)	0.0822 (+/- 0.15)
Prompt2Vec	0.1862 (+/- 0.10)	0.3292 (+/- 0.18)	0.3452 (+/- 0.15)	0.0678 (+/- 0.11)
BM25	0.0955 (+/- 0.04)	0.4998 (+/- 0.13)	0.5338 (+/- 0.17)	0.1940 (+/- 0.11)
S2S	0.1406 (+/- 0.10)	0.2112 (+/- 0.15)	0.2930 (+/- 0.18)	0.1600 (+/- 0.21)
DMSG (ours)	<b>0.3157 (+/- 0.13)</b>	<b>0.5690 (+/- 0.11)</b>	<b>0.5784 (+/- 0.05)</b>	<b>0.4069 (+/- 0.21)</b>

**Table 1: Comparison of DMSG with baselines, including average score and standard deviation across all generated slate.**

evaluation practices [26]. As a result the reported scores should be interpreted as lower bound estimates, since models may generate equally relevant items that are not present in the reference slate.

To better account for semantically similar but non-identical items, we introduce the cosine similarity between embeddings as part of the ground truth. Specifically, for a generated item  $item_i$  in position  $i$ , if it does not appear in the reference slate, it is assigned a relevance score based on cosine similarity to the corresponding reference item  $ref_i$ , i.e.,  $s_i = \cos\_sim(item_i, ref_i)$ . This approach reduces the penalty for recommending items not present in the ground truth, but still highly similar to those and, at the same time, mitigates mismatches via approximate nearest-neighbor lookups on the item embedding space. We introduce two similarity-based metrics: NDCGSim and MAPSim. The first one assesses the quality of the slate as whole, while the latter focuses on the relevance of the retrieved items. We also report BertScore [42], which measures semantic similarity between the generated and reference slates using item vectors as token embeddings, and CategorySim, a category-level similarity score that weights metadata overlaps by their frequency of appearance.

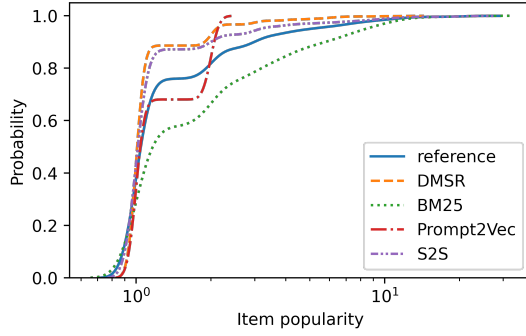
Table 1 presents the average and standard deviation of each metric across all test prompts, providing insight into the model stability. To account for the stochasticity nature of DMSG, we sample five slates per prompt and compute average scores, and then aggregate across prompts to report overall performance.

DMSG outperforms the baselines across nearly all metrics. On the Curated dataset, DMSG shows notable improvements in NDCGSim and MAPSim (+17% and +12.9% over the second-best BM25), indicating both high-quality item selection and effective slate structuring. The smallest gain is seen in CategorySim (+2.3%), which is expected since retrieval models like BM25 already perform well in category alignment. We observe a similar trend on the MPD dataset. While BM25 slightly outperforms DMSG on NDCGSim, DMSG achieves higher MAPSim and BERTScore, suggesting that it generally recommends more semantically relevant items even if not ranked identically to the reference.

<sup>3</sup>huggingface.co/ianstenbit/keras-sd2.1

<sup>4</sup>https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2





**Figure 3: Cumulative distribution of exposure generated by the recommendation algorithms when compared to initial level of popularity in the Bundle dataset.**


Results on the Bundle dataset echoes the findings from the music datasets. Due to the shorter length of bundles, we compute metrics using only the first five items. DMSG continues to outperform all baselines, confirming its ability to handle both long and short slates effectively. This demonstrates the broad applicability of DMSG across diverse slate generation scenarios.

**6.2.2 Can DMSG provide a diverse selection of items?** Diversity is a critical aspect of recommender and retrieval systems. In this work, we focus on two complementary dimensions of diversity: 1) recommending less popular items to avoid over-concentration on frequent content and 2) providing fresh and varied results across multiple generations for the same prompt, while maintaining relevance. The first aspect encourages exposure to a broader range of catalog items by reducing reliance on popularity signals. The second, tied to the notion of freshness, ensures that repeated prompts return different—but still relevant—slates over time. This ability to produce varied recommendations is known to help maintain user interest and supports discovery.

Figure 3 compares the popularity distribution of items recommended by different models on the Bundle dataset, relative to the reference data. Item popularity (x-axis) is defined by the number of user interactions logged for each item, while the y-axis shows the proportion of recommended items at each popularity level. DMSG, along with S2S, tends to recommend less popular items overall, falling below the popularity level of the reference set, whereas BM25 consistently favors more popular items, often exceeding reference popularity. These results, alongside the relevance metrics in Table 1, show that DMSG successfully promotes lower-popularity items while maintaining high-quality slates, ultimately enabling better exploration of the catalog.

Another strength of DMs is their ability to naturally generate fresh content due to their stochastic sampling process in continuous latent space. We evaluate this by measuring freshness—the number of new, non-repeated items—across multiple slates generated for the same prompt. To ensure these new items are still relevant, we also measure BERTScore for each slate, assessing content quality across generations.

Figure 4 summarizes results on the MPD test set, showing freshness and quality across multiple runs per prompt. Each run produces



**Figure 4: Repeated playlist generation with freshness and quality score. Multiple playlist from the same prompt have similar qualities and low overlap.**

	NDCGSim@30	MAPSim@30	BertScore	CategorySim@30
DMSG (random sort)	0.5134 (+/- 0.15)	0.7330 (+/- 0.10)	0.8744 (+/- 0.08)	0.4788 (+/- 0.18)
DMSG (no post-process)	0.5138 (+/- 0.15)	0.7343 (+/- 0.10)	0.8744 (+/- 0.08)	0.4789 (+/- 0.18)
DMSG (centroid sort)	<b>0.5210 (+/- 0.15)</b>	<b>0.8119 (+/- 0.09)</b>	<b>0.8796 (+/- 0.08)</b>	<b>0.4792 (+/- 0.18)</b>

**Table 2: Post-processing methods on MPD data.**

a slate of items; we plot the first 16 positions for clarity. The right-most column shows the average BERTScore per run, which remains consistently around 0.8, confirming that freshness does not come at the cost of quality. Notably, the model surfaces new items throughout the entire slate, not just in later positions. This is evident from the spread of blue squares across all positions in the figure, indicating DMSG’s ability to inject freshness uniformly across the slate. This property is essential for delivering dynamic and engaging user experiences, as it ensures that each recommendation instance feels new, even for identical prompts.

**6.2.3 What is the effect of post-processing?** Starting from a conditioning input, DMSG generates a slate of items in a structured order to optimize coherence and relevance based on the learned latent space. However, in real-world recommender systems, it is common to apply a post-processing step to reorder items, either to incorporate additional relevance signals or to personalize the slate for a specific user.

To assess DMSG flexibility in such settings, we evaluate its performance both as a standalone recommender and as a candidate generator, followed by three post-processing strategies (i) no post-processing, preserving the model’s original item order; (ii) sorting by item centroid, where items are ranked by proximity to the centroid of the latent slate embedding; and (iii) random shuffling, serving as a control. We focus on the MPD dataset, where the effect of post-processing is more pronounced due to the longer length of generated slates.

Table 2 presents the results. Overall, the differences between variants are modest. However, sorting by item centroid slightly improves performance, suggesting that DMSG is able to position the most relevant tracks near the center of the latent distribution.

By sorting these toward the top, we can ensure that the first items in the slate are most representative of the input prompt, a useful property for tasks like playlist generation, where early items often receive more attention.

## 7 Online Experiments

We evaluated our slate generation approach to assess its effectiveness in powering real-world recommender systems. Although trained offline on a limited subset of generic music playlists, we expected DMSG to deliver satisfactory user listening experiences. Notably, we do not update or fine-tune the model during the testing period—we leave such enhancements for future work.

*Experiment Settings.* To evaluate the effectiveness of our proposed approach, we conducted a live A/B test on our production platform. We compared our model, DMSG, against the platform’s default playlist generation model over a two-week period. Users were randomly assigned to one of two non-overlapping groups: the control group received playlists generated by the production model, while the treatment group received playlists generated by DMSG.

The experiment collected a large sample of interaction data, spanning 1 million users, 2.8 million unique tracks, and 4.8 million distinct sessions. Each session was divided into multiple slates, with each slate consisting of five consecutive tracks. For users in the treatment group, these tracks were generated using DMSG; for those in the control group, they were generated by the existing production model.


The production model (control) is a personalized retrieval-based system that, given textual keywords representing the playlist’s content hypothesis and user familiarity with the content hypothesis, retrieves a broader pool of relevant tracks. These are then filtered using additional user-specific features (e.g., content familiarity). Content familiarity is computed using past music tracks that the users interacted with. To ensure a fair comparison, we retained the same post-processing and sequencing pipeline across both groups, replacing only the retrieval component with DMSG, which generates the initial track set based on the same content hypothesis.

DMSG was deployed on a single NVIDIA L4 GPU system for inference. During testing, the model utilized 50 diffusion steps and consistently achieved a P99 latency of 150ms. In comparison, the baseline control system exhibited a P99 latency of 500ms. This reduction provides potential for future iterations or an overall improvement in user-perceived latency.

*Online Evaluation.* We evaluate two key performance metrics commonly used in the streaming domain: *stream curations*, which capture active interaction with a track (e.g., adding it to a playlist or liking it), and *listening behavior*, capturing user session listening (e.g., number of tracks skipped or minutes played).

On average, the results show a +6.8% uplift in stream curations, indicating that users engaged more actively with tracks recommended by DMSG compared to the control group. When focusing specifically on curations involving the addition of tracks to a user’s “Liked Songs”, the uplift increases to +10.5%.

To better understand this effect, we break down stream curation performance by user type: 1) new users: no prior interaction on



**Figure 5: Differences in stream curations between treatment (DMSG) and control group divided by user type, with error bars indicating the 95% confidence interval.**

the platform before the experiment; 2) existing users: some previous but infrequent engagement; and 3) habitual users: regular, high-frequency engagement. As shown in Figure 5, the uplift in stream curations is most pronounced among existing and habitual users, highlighting DMSG’s particular value for users familiar with the platform. A possible explanation is that these users, having developed some baseline familiarity or intent, may be more open to discovering new content.

Adding a track to user’s Liked Songs playlist is an indication that the recommendation was both relevant and novel, as tracks already liked cannot be added again. While this reflects a significant improvement in active engagement, it is accompanied by a small decline in listening duration: −5.6% decrease in minutes played and +3% increase in skip rate (where users skipped tracks before completion). This trade-off is expected, as DMSG does not rely on explicit personalization signals and instead prioritizes content freshness, which, as discussed earlier, better supports discovery.

Importantly, active curations can serve as a proxy for how users respond to novel or previously unseen content, which is often related to content diversity. While novelty reflects a user’s exposure to new items, diversity captures the broader range of content types or categories presented. The two are closely linked, as more diverse recommendations increase the likelihood of surfacing novel content. To directly assess content freshness, we measured the proportion of repeated tracks served to the same user during the first week post-exposure. DMSG showed a −13.4% reduction in repetition compared to the control group (lower is better), indicating greater diversity in recommendations. This aligns with both our underlying assumptions and offline evaluations, highlighting how DMSG’s inherent stochastic process supports the generation of diverse, high-quality recommendations.

## 8 Discussion and Future Work

We introduced DMSG, a novel method for slate generation based on diffusion models. By leveraging the generative capabilities of diffusion models, DMSG conditions on contextual information, such as textual prompts, to produce structured and coherent item slates. Its



ability to generalize to unseen catalog items, combined with the stochastic nature of iterative sampling, enables DMSG to dynamically balance content freshness and relevance, making it well-suited for integration into larger recommender system pipelines.

As demonstrated in our experiments, DMSG can generate relevant item combinations from simple text prompts, opening up a range of applications. In music playlist generation, it can reconstruct reference playlists from minimal input or generate new ones based on user-provided or system-generated descriptions, such as preferred music genres. While our study focuses on text-based conditioning, our framework supports alternative conditioning signals, including catalog items or user interaction patterns. In e-commerce bundle recommendations, DMSG can be paired with an external module that detects user interest in categories (e.g., “TV accessories”) based on browsing behavior or past interactions. Users can also explicitly explore these categories, receiving fresh, dynamically updated bundles. DMSG’s stochasticity enables diverse yet relevant recommendations, supporting discovery and ultimately enhancing user satisfaction.

To further enhance DMSG’s capabilities, future work could explore integrating richer metadata into the generation process. Unlike pixels in images or samples in audio, catalog items often contain detailed metadata that could improve alignment between prompts and generated slates. Additionally, incorporating prompt reformulation techniques, as explored in recent studies [19], may offer finer control over slate generation and improve output quality.

## References

- [1] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of the web conference 2020*. 2155–2165.
- [2] Tzof Avny Brosh, Amit Livne, Oren Sar Shalom, Bracha Shapira, and Mark Last. 2022. Bruce: Bundle recommendation using contextualized item embeddings. In *Proceedings of the 16th ACM Conference on Recommender Systems*. 237–245.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [4] Yuwei Cao, Nikhil Mehta, Xinyang Yi, Raghunandan Keshavan, Lukasz Heldt, Lichan Hong, Ed H Chi, and Maheswaran Sathiamoorthy. 2024. Aligning Large Language Models with Recommendation Knowledge. *arXiv preprint arXiv:2404.00245* (2024).
- [5] Diego Carraro and Derek Bridge. 2024. Enhancing recommendation diversity by re-ranking with large language models. *ACM Transactions on Recommender Systems* (2024).
- [6] Jia-Wei Chang, Ching-Yi Chiou, Jia-Yi Liao, Ying-Kai Hung, Chien-Che Huang, Kuan-Cheng Lin, and Ying-Hung Pu. 2021. Music recommender using deep embedding-based features and behavior-based reinforcement learning. *Multimedia Tools and Applications* 80, 26 (2021), 34037–34064.
- [7] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems* 34 (2021), 15084–15097.
- [8] Erica Coppolino, Giuseppe Manco, and Aristides Gionis. 2024. Relevance meets diversity: A user-centric framework for knowledge exploration through recommendations. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 490–501.
- [9] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*. 87–94.
- [10] Fernando Diaz, Bhaskar Mitra, Michael D Ekstrand, Asia J Biega, and Ben Carterette. 2020. Evaluating stochastic rankings with expected exposure. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 275–284.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems* 33 (2020), 6840–6851.
- [12] Eugene Ie, Vihan Jain, Jing Wang, Sanmit Narvekar, Ritesh Agarwal, Rui Wu, Heng-Tze Cheng, Morgane Lustman, Vince Gatto, Paul Covington, Jim McFadden, Tushar Chandra, and Craig Boutilier. 2019. Reinforcement Learning for Slate-based Recommender Systems: A Tractable Decomposition and Practical Methodology. (May 2019). *arXiv:1905.12767* [cs.LG]
- [13] Marius Kaminskis and Derek Bridge. 2016. Diversity, serendipity, novelty, and coverage: a survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM Transactions on Interactive Intelligent Systems (TiIS)* 7, 1 (2016), 1–42.
- [14] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [15] Xiang Li, John Thickstun, Ishaan Gulrajani, Percy S Liang, and Tatsunori B Hashimoto. 2022. Diffusion-lm improves controllable text generation. *Advances in Neural Information Processing Systems* 35 (2022), 4328–4343.
- [16] Paul Pu Liang, Manzil Zaheer, Yuan Wang, and Amr Ahmed. 2020. Anchor & transform: Learning sparse embeddings for large vocabularies. *arXiv preprint arXiv:2003.08197* (2020).
- [17] Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. 2022. A survey of transformers. *AI open* 3 (2022), 111–132.
- [18] Xiao Lin, Xiaokai Chen, Chenyang Wang, Hantao Shu, Linfeng Song, Biao Li, et al. 2023. Discrete Conditional Diffusion for Reranking in Recommendation. *arXiv preprint arXiv:2308.06982* (2023).
- [19] Bingqi Ma, Zhuofan Zong, Guanglu Song, Hongsheng Li, and Yu Liu. 2024. Exploring the Role of Large Language Models in Prompt Encoding for Diffusion Models. *arXiv preprint arXiv:2406.11831* (2024).
- [20] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th acm international conference on information and knowledge management*. 2243–2251.
- [21] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [22] William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4195–4205.
- [23] Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative adversarial text to image synthesis. In *International conference on machine learning*. PMLR, 1060–1069.
- [24] Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations and Trends® in Information Retrieval* 3, 4 (2009), 333–389.
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2021. High-resolution image synthesis with latent diffusion models. 2022 IEEE. In *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10674–10685.
- [26] Alan Said and Alejandro Bellogin. 2014. Comparative recommender system evaluation: benchmarking recommendation frameworks. In *Proceedings of the 8th ACM Conference on Recommender systems*. 129–136.
- [27] Tim Salimans and Jonathan Ho. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512* (2022).
- [28] Jagendra Singh and Vijay Kumar Bohat. 2021. Neural Network Model for Recommending Music Based on Music Genres. In *2021 International Conference on Computer Communication and Informatics (ICCCI)*. IEEE, 1–6.
- [29] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*. PMLR, 2256–2265.
- [30] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020).
- [31] Ivan Srba, Robert Moro, Matus Tomlein, Branislav Pecher, Jakub Simko, Elena Stefancova, Michal Kompan, Andrea Hrkova, Juraj Podrouzek, Adrian Gavornik, et al. 2023. Auditing YouTube’s recommendation algorithm for misinformation filter bubbles. *ACM Transactions on Recommender Systems* 1, 1 (2023), 1–33.
- [32] Robin Strudel, Corentin Tallec, Florent Althé, Yilun Du, Yaroslav Ganin, Arthur Mensch, Will Grathwohl, Nikolay Savinov, Sander Dieleman, Laurent Sifre, et al. 2022. Self-conditioned embedding diffusion for text generation. *arXiv preprint arXiv:2211.04236* (2022).
- [33] Zhu Sun, Kaidong Feng, Jie Yang, Hui Fang, Xinghua Qu, Yew-Soon Ong, and Wenyuan Liu. 2024. Revisiting Bundle Recommendation for Intent-aware Product Bundling. *ACM Transactions on Recommender Systems* (2024).
- [34] Zhu Sun, Jie Yang, Kaidong Feng, Hui Fang, Xinghua Qu, and Yew Soon Ong. 2022. Revisiting Bundle Recommendation: Datasets, Tasks, Challenges and Opportunities for Intent-Aware Product Bundling. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- [35] Peter Sunehag, Richard Evans, Gabriel Dulac-Arnold, Yori Zwols, Daniel Visentin, and Ben Coppin. 2015. Deep Reinforcement Learning with Attention for Slate Markov Decision Processes with High-Dimensional States and Actions. (Dec.

- 2015). [arXiv:1512.01124 \[cs.AI\]](#)
- [36] Federico Tomasi, Joseph Cauteruccio, Surya Kanoria, Kamil Ciosek, Matteo Rinaldi, and Zhenwen Dai. 2023. Automatic Music Playlist Generation via Simulation-based Reinforcement Learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4948–4957.
  - [37] Saúl Vargas and Pablo Castells. 2011. Rank and relevance in novelty and diversity metrics for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*. 109–116.
  - [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
  - [39] Wenjie Wang, Yiyang Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. 2023. Diffusion recommender model. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 832–841.
  - [40] Zhengyi Yang, Jiancan Wu, Zhicai Wang, Xiang Wang, Yancheng Yuan, and Xiangnan He. 2024. Generate What You Prefer: Reshaping Sequential Recommendation via Guided Diffusion. *Advances in Neural Information Processing Systems* 36 (2024).
  - [41] Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Michael He, et al. 2024. Actions speak louder than words: Trillion-parameter sequential transducers for generative recommendations. *arXiv preprint arXiv:2402.17152* (2024).
  - [42] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675* (2019).