

# TokenRec: Learning to Tokenize ID for LLM-based Generative Recommendations

Haohao Qu

The Hong Kong Polytechnic University  
haohao.qu@connect.polyu.hk

Zihuai Zhao

The Hong Kong Polytechnic University  
zihuai.zhao@connect.polyu.hk

Wenqi Fan\*

The Hong Kong Polytechnic University  
wenqifan03@gmail.com

Qing Li

The Hong Kong Polytechnic University  
qing-prof.li@polyu.edu.hk

## ABSTRACT

There is a growing interest in utilizing large-scale language models (LLMs) to advance next-generation Recommender Systems (RecSys), driven by their outstanding language understanding and in-context learning capabilities. In this scenario, tokenizing (i.e., indexing) users and items becomes essential for ensuring a seamless alignment of LLMs with recommendations. While several studies have made progress in representing users and items through textual contents or latent representations, challenges remain in efficiently capturing high-order collaborative knowledge into discrete tokens that are compatible with LLMs. Additionally, the majority of existing tokenization approaches often face difficulties in generalizing effectively to new/unseen users or items that were not in the training corpus. To address these challenges, we propose a novel framework called **TokenRec**, which introduces not only an effective ID tokenization strategy but also an efficient retrieval paradigm for LLM-based recommendations. Specifically, our tokenization strategy, Masked Vector-Quantized (MQ) Tokenizer, involves quantizing the masked user/item representations learned from collaborative filtering into discrete tokens, thus achieving a smooth incorporation of high-order collaborative knowledge and a generalizable tokenization of users and items for LLM-based RecSys. Meanwhile, our generative retrieval paradigm is designed to efficiently recommend top- $K$  items for users to eliminate the need for the time-consuming auto-regressive decoding and beam search processes used by LLMs, thus significantly reducing inference time. Comprehensive experiments validate the effectiveness of the proposed methods, demonstrating that TokenRec outperforms competitive benchmarks, including both traditional recommender systems and emerging LLM-based recommender systems.

## CCS CONCEPTS

• **Information systems** → **Recommender systems.**

\* Corresponding author: Wenqi Fan, Department of Computing, and Department of Management and Marketing, The Hong Kong Polytechnic University.

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 ACM 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](mailto:permissions@acm.org).

Conference'17, July 2017, Washington, DC, USA

© 2024 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

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

## KEYWORDS

Recommender Systems, Large Language Models (LLMs), ID Tokenization, Vector Quantization, Collaborative Filtering (CF).

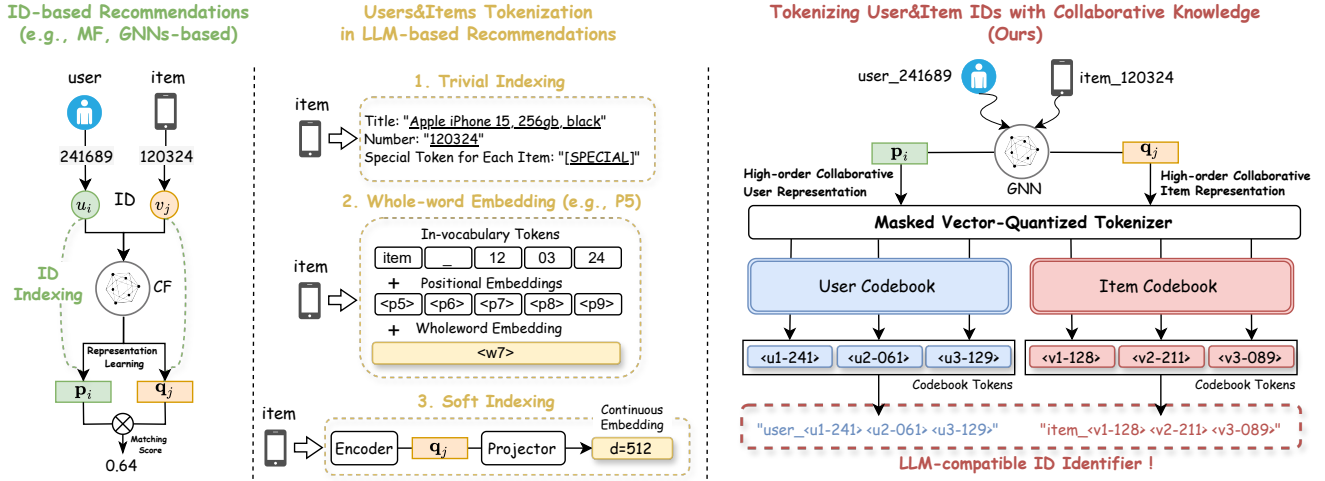
### ACM Reference Format:

Haohao Qu, Wenqi Fan\*, Zihuai Zhao, and Qing Li. 2024. TokenRec: Learning to Tokenize ID for LLM-based Generative Recommendations. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

As a prominent branch in the data mining field, recommender systems (**RecSys**) serve as an indispensable and effective technique in addressing information overload problems and enriching user experience across diverse applications [6, 48], such as e-commerce, job search, and social media platforms. To provide personalized recommendations that accord with user preferences, one of the most representative techniques is collaborative filtering (CF), which aims to capture collaborative knowledge by modeling the history of user-item interactions [15]. For example, as the most classic CF method, Matrix Factorization (MF) [26] decomposes the use-item matrix into two low-rank matrices (i.e., user&item representations) to represent each user and item with a unique ID (namely ID indexing/tokenizing), and calculates matching scores to predict user behaviors and generate recommendations via the inner product between their representations. Due to the superior ability in representation learning on graph-structured data, Graph Neural Networks (GNNs) (e.g., LightGCN [21] and GTN [14]) have been recently employed to significantly enhance CF by capturing higher-order collaborative knowledge on user-item interaction graphs for recommendations in IDs manners. The main idea of the existing methods is to obtain learnable representations (i.e., token embeddings) for discrete user&item IDs (i.e., tokens).

More recently, the rapid development of Large Language Models (LLMs) techniques (e.g., ChatGPT and LLaMA) has showcased notable milestones for revolutionizing natural language processing techniques [4, 47]. Technically, the unprecedented power of LLMs can be attributed to the scaling up of model/parameter size alongside a tremendous amount of training corpus. In particular, LLMs equipped with billion-scale parameters have exhibited unprecedented language understanding and generation abilities, along with remarkable generalization capability and in-context learning skills that facilitate LLMs to better generalize to unseen tasks and domains [47]. Given the emerging trends and aforementioned advancements of LLMs, LLM-empowered recommender systems have



**Figure 1: Comparison of ID tokenization methods in LLM-based recommendations. Unlike the existing methods, our approach can tokenize users and items with LLM-compatible tokens by leveraging high-order collaborative knowledge.**

drawn increasing attention from recent studies and demonstrated distinctive abilities for advancing recommender systems [48]. For example, ChatRec [19] integrates ChatGPT into conversational recommendations, which enables users to deliver their explicit preferences in natural language. P5 [20] introduces an LLM-based recommendation framework that unifies diverse recommendation tasks by multi-task prompt-based pre-training.

Unlike the ID-based recommendation methods (e.g., MF and GNN-based), users&items tokenization is one of the most critical steps to take advantage of LLMs within recommendations, as shown in Figure 1. To be more specific, the naive approach known as Independent Indexing (IID) assigns special tokens (i.e., ID) to tokenize each user and item within language models directly. However, this approach becomes infeasible and unrealistic when dealing with large-scale real-world recommender systems, in which the sizes of users and items typically reach billions, dramatically expanding the token vocabulary in LLMs. Meanwhile, as a natural solution, textual title indexing is proposed to utilize LLMs’ in-vocabulary tokens to represent items based on their titles and descriptions, such as the example "Apple iPhone 15, 256 GB, black", thus avoiding vocabulary size explosion in LLMs [1, 19]. To achieve a closer alignment between recommendations and natural language, P5 employs *whole-word embedding* [39] to indicate whether consecutive sub-word tokens originate from the same entities (i.e., user/item). Some studies use continuous embedding (i.e., soft indexing) learned from encoders to represent users&items in LLMs-based recommendations [29, 46].

Despite the success mentioned above, the majority of existing methods for users&items tokenization for LLM-based recommender systems still have several limitations. For example, the use of whole-word embedding cannot effectively capture high-order collaborative knowledge and generalize well to unseen users/items for recommendation. In addition, due to the nature of discrete tokens in language models, using continuous indexing makes it challenging to align LLMs in recommender systems closely.

To address such challenges, we propose a novel LLM-based framework for recommender systems (**TokenRec**), in which a novel tokenization strategy is proposed to tokenize numerical ID (i.e., identifiers) of users and items by seamlessly integrating high-order collaborative knowledge into LLMs. Meanwhile, a generative retrieval paradigm is developed to generate item representations and retrieve appropriate items for collaborative recommendations. Our major contributions are summarized as follows:

- We introduce a principle strategy named **Masked Vector-Quantized Tokenizer** to tokenize users and items tailored to LLMs, which contributes to incorporating high-order collaborative knowledge in LLM-based recommendations. More specifically, two novel mechanisms (i.e., masking and  $K$ -way encoder) are designed to enhance the generalization capability of the proposed tokenization method in LLM-based recommendations.
- We propose a novel framework (**TokenRec**) for recommender systems in the era of LLMs, where a generative retrieval paradigm is designed to effectively and efficiently recommend top- $K$  items for users rather than directly generating tokens in natural language.
- We conduct extensive experiments on four widely used real-world datasets to empirically demonstrate the effectiveness of our proposed TokenRec, including the superior recommendation performance and its generalization ability in predicting new and unseen users’ preferences.

The remainder of this paper is structured as follows. Section 2 introduces the proposed approach, which is evaluated in Section 3. Then, Section 4 summarizes the recent development of collaborative filtering and LLM-based RecSys. Finally, conclusions are drawn in Section 5.

## 2 THE PROPOSED METHOD

This section will begin by reformulating the collaborative recommendation as a language-processing task. Then, we provide an overview of the proposed TokenRec, followed by a detailed

explanation of each model component. Finally, we will discuss the model training and inference of TokenRec.

## 2.1 Notations and Definitions

Let  $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$  and  $\mathcal{V} = \{v_1, v_2, \dots, v_m\}$  be the sets of users and items, respectively, where  $n$  is the number of users, and  $m$  is the number of items. Moreover, we use  $\mathcal{N}_{(u_i)}$  to denote the item set that user  $u_i$  has interacted in the history. As the traditional collaborative filtering methods, user  $u_i$  and item  $v_j$  can be embedded into low-dimensional latent vectors (i.e., collaborative representations of users and items)  $\mathbf{p}_i \in \mathbb{R}^d$  and  $\mathbf{q}_j \in \mathbb{R}^d$  respectively, where  $d$  is the length of the vector.

In general, the goal of a recommender system is to understand users’ preferences by modeling interactions (e.g., clicks and bought) between users  $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$  and items  $\mathcal{V} = \{v_1, v_2, \dots, v_m\}$ . As a widely used solution, collaborative filtering (CF) techniques are developed to learn user and item representations from historical user-item interactions. Thus, we reformulate the CF recommendation in the language model paradigm. Assume that we only have token IDs  $\mathcal{T}_i$  and  $\mathcal{T}_j$  for each user  $u_i$  and item  $v_j$ . By integrating these IDs into textual prompts  $\mathcal{P}$ , LLM generates the representation  $\mathbf{z}_i \in \mathbb{R}^d$  of items that a user  $u_i$  may like, expressed as:

$$\mathbf{z}_i = \text{LLM}(\mathcal{P}, \mathcal{T}_i, \{\mathcal{T}_j | v_j \in \mathcal{N}_{(u_i)}\}). \quad (1)$$

Notably, the interacted items  $\mathcal{N}_{(u_i)}$  will be placed into the language model in a *non-sequential* way to accommodate the setting of collaborative filtering.

## 2.2 An Overview of the Proposed Framework

To better align natural language with recommendation tasks, we propose a novel LLM-based generative recommendation framework (TokenRec). As shown in Figure 2, the proposed framework consists of two key modules, namely **Masked Vector-Quantized (MQ) Tokenizer for Users and Items** and **Generative Retrieval for Recommendations**. The first module aims to address the fundamental task of ID tokenization in LLM-based recommendations so as to seamlessly integrate users&items (i.e., numeric ID) into natural language form. However, tokenizing users&items faces tremendous challenges due to the huge number of users and items in recommender systems. To address the emerging challenges, we introduce a Masked Vector-Quantized Tokenizer (MQ-Tokenizer) to learn specific codebooks and represent users&items with a list of special tokens through encoder and decoder networks. The goal of the second module is to perform user modeling via LLM for personalized recommendations. To achieve this, a generative retrieval paradigm is introduced to retrieve the  $K$ -nearest items from the whole item set for generating a personalized top- $K$  recommendation list in an effective and efficient manner. The details of the proposed method, TokenRec, will be described in the following sections.

## 2.3 Masked Vector-Quantized Tokenizers for Users and Items

Instead of assigning each user and item with a specific token (causing a significant increase in vocabulary size), we propose a novel strategy to tokenize users and items to align with natural

language, where quantization techniques are developed to represent each user and item with certain discrete indices (i.e., tokens). More specifically, as one of the most representative quantization techniques, Vector Quantization (VQ) aims to convert inputs into a set of discrete tokens (i.e., codes) and learn a discrete latent representation space (i.e., codebook) by reconstructing the original inputs [10, 36, 41]. Moreover, in order to capture *high-order collaborative knowledge* from users’ interaction history, we propose to conduct the vector quantization on the well-trained representations learned from advanced GNN-based recommendations. However, naively applying vector quantization to index/tokenize users and items is fraught with critical challenges due to poor generalization capability. For example, certain noise and propagation errors in the cascade process of user and item tokenizations in the LLMs greatly hinder the expressiveness of the vanilla VQ.

To tackle these challenges, we propose a Masked Vector-Quantized Tokenizer (**MQ-Tokenizer**) for tokenizing users and item IDs, where two novel strategies, namely masking and  $K$ -way encoding, are designed to enhance the generalization capability of our proposed tokenization in LLM-based recommendations. More specifically, the proposed MQ-Tokenizer consists of a masking operation on the inputs (i.e., users&items representations), a  $K$ -way encoder for multi-head feature extraction with a corresponding  $K$ -way codebook for latent feature quantization, and a  $K$ -to-1 decoder that reconstructs the input representations with the quantized features. It is worth mentioning that the proposed tokenization method will design the specific MQ-Tokenizers for the user and item respectively (i.e., *Item MQ-Tokenizer* and *User MQ-Tokenizer*), in which they share the same architecture. Thus, for simplicity, we will detail the item MQ-Tokenizer for item ID tokenization while omitting the user MQ-Tokenizer, as shown in Figure 2.

**2.3.1 Collaborative Knowledge.** The primary objective of our proposed MQ-Tokenizer lies in capturing high-order collaborative knowledge into latent representations through vector quantization. Collaborative knowledge has been widely proven informative for predicting user preferences in recommender systems [15, 44], as it reveals the in-depth behavioral similarity among users (or items). However, it is challenging for LLMs to map textual descriptions of the high-order collaborative relationships between users and items explicitly with natural language [42]. Meanwhile, as the most advanced collaborative filtering techniques, graph neural networks (GNNs) are proposed to learn user and item representations by capturing high-order collaborative signals among users and items in recommendations [14, 21]. Despite great potential, it is challenging to fully harness the high-order signals from user-item interactions in natural language for LLM-based recommendations. Thus, we propose to perform vector quantization on the collaborative representations learned by advanced GNN-based methods for indexing/tokenizing users and items. Specifically, our MQ-Tokenizer assigns each user and item a small number of discrete tokens generated by quantizing the GNN-based representations. In other words, it is suggested that users and items in close proximity to each other in the latent space of collaborative knowledge (i.e., user/item representations) are more likely to share similar tokens/indices, which naturally aligns LLMs with recommendations by tokenizing users and items in natural language.

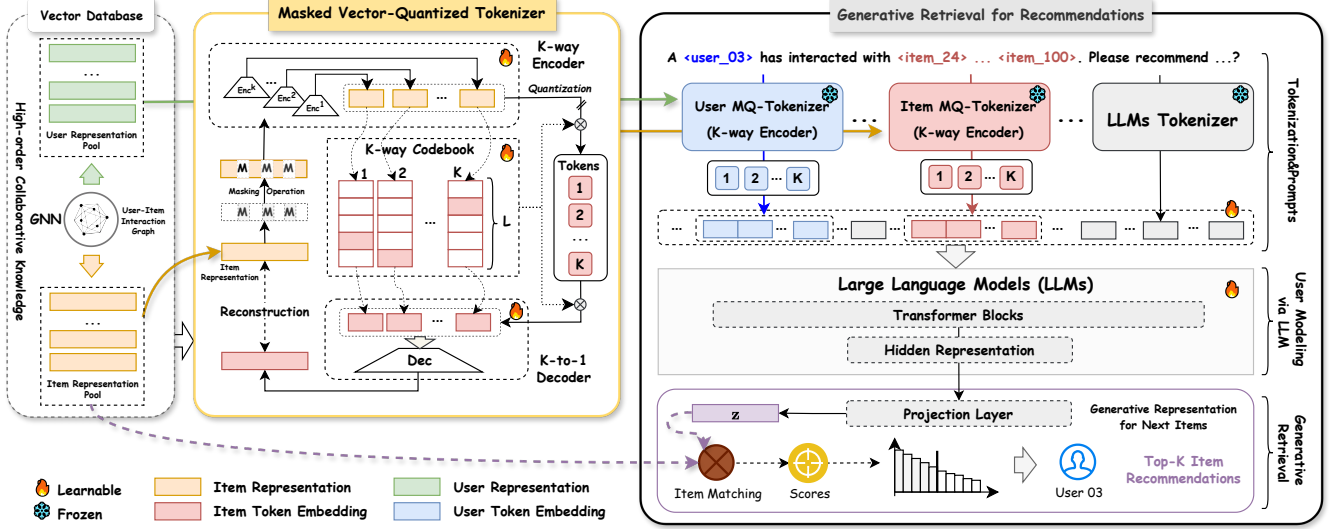


Figure 2: The overall framework of the proposed TokenRec, which consists of the masked vector-quantized tokenizer with a  $K$ -way encoder for item ID tokenization, and generative retrieval for recommendation. Note that we detail the item MQ-Tokenizer while omitting the user MQ-Tokenizer for simplicity.

**2.3.2 Masking Operation.** On the basis of introducing the collaborative representations learned from GNN-based collaborative filtering methods, which are stored in a vector database, we propose to randomly mask these representations to create a challenging task that enables the tokenizer to build a comprehensive understanding and generalize well. Specifically, we introduce an element-wise masking strategy  $\mathcal{E}$  following Bernoulli distribution as follows:

$$\mathcal{E} \sim \text{Bernoulli}(\rho), \quad (2)$$

where  $\rho$  is the masking ratio. The Bernoulli distribution is the discrete probability distribution of a random variable, which takes the value 1 with probability  $\rho$  and the value 0 with probability  $1 - \rho$ . Given the collaborative representations  $\mathbf{p}_i$  and  $\mathbf{q}_j$ , the masking process can be presented by

$$\mathbf{p}'_i = \text{Mask}(\mathbf{p}_i, \mathcal{E}), \quad \mathbf{q}'_j = \text{Mask}(\mathbf{q}_j, \mathcal{E}), \quad (3)$$

where  $\mathbf{p}'_i$  and  $\mathbf{q}'_j$  are the masked representations of user  $u_i$  and item  $v_j$ , respectively. It should be noted that  $\mathcal{E}$  randomly generates masks at each training epoch to create multifarious samples so as to enhance the generalization capability in the proposed tokenizers.

**2.3.3  $K$ -way Encoder and Codebook.** In light of the masked collaborative representations of items, we propose a novel vector quantization framework to tokenize each item as a handful of discrete tokens. To be more specific, a  $K$ -way encoder  $\text{Enc}(\cdot)$  is developed to learn a corresponding  $K$ -way codebook  $\mathbf{C} = \{\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^k, \dots, \mathbf{c}^K\}$  for items, where  $\mathbf{c}^k \in \mathbb{R}^{L \times d_c}$  is a latent space (i.e., the  $k$ -th sub-codebook) with  $L$  pairs of codeword (i.e., token<sup>1</sup>) and  $w^k$  and its  $d_c$ -dimensional codeword embedding  $\mathbf{c}_{w^k}^k$ . The idea behind this is to represent the users/items with entries of a learned

<sup>1</sup>In this work, the terms 'codeword' and 'token' are used interchangeably. 'Codeword' is a common term in the context of codebooks, while 'token' is more widely used in the language modeling literature.

codebook (indexed by discrete tokens), providing great potential to tokenize the discrete users and items based on their collaborative representations in an LLM-compatible manner.

In general, the proposed quantization method for user/item tokenization involves two main steps. Taking item  $v_j$  as an example, the  $K$ -way encoder first employs  $K$  different encoders  $\text{Enc}^k(\cdot)$  to encode the masked item representation  $\mathbf{q}'_j$  and generate  $K$  corresponding latent vectors  $\{\mathbf{a}_j^k\}_{k=1}^K$  as follows:

$$\mathbf{a}_j^k = \text{Enc}^k(\mathbf{q}'_j) = \text{MLP}^k(\mathbf{q}'_j), \quad (4)$$

where  $\mathbf{a}_j^k \in \mathbb{R}^{d_c}$ , and each encoder  $\text{Enc}^k(\cdot)$  can be implemented as a multilayer perceptron network (MLP) with three hidden layers. Powered by different encoders, the proposed  $K$ -way Encoder enables multiple attentions and uncovers different patterns on the inputs to enhance the generalization capability towards users&items representations.

The next step is to quantize the encoded vectors  $\{\mathbf{a}_j^k\}_{k=1}^K$  into discrete tokens (i.e., indices) by looking up nearest neighbors in the learnable  $K$ -way codebook  $\mathbf{C} = \{\mathbf{c}^1, \mathbf{c}^2, \dots, \mathbf{c}^k, \dots, \mathbf{c}^K\}$ . More specifically, given the item  $v_j$  and a sub-codebook  $\mathbf{c}^k \in \mathbb{R}^{L \times d_c}$  from the  $k$ -th encoder  $\text{Enc}^k$ , a practical approach is leveraging the Euclidean distance to calculate similarity scores between the encoded vector  $\mathbf{a}_j^k$  and all the codeword embeddings  $\{\mathbf{c}_l^k\}_{l=1}^L$ , so as to find the nearest embedding (i.e., token embedding) that can be used to effectively represent the encoded vector  $\mathbf{a}_j^k$ :

$$w_j^k = \arg \min_l \|\mathbf{a}_j^k - \mathbf{c}_l^k\|^2, \quad (5)$$

$$\text{Quantize}(\mathbf{a}_j^k) = \mathbf{c}_{w_j^k}^k, \quad (6)$$

where  $w_j^k$  denotes the codeword (i.e., ID token) of the nearest neighbor at sub-codebook  $c^k$  for the item  $v_j$ , and  $\|\cdot\|^2$  denotes the  $l_2$  norm of the variables.

In other words, the proposed Item MQ-Tokenizer can tokenize the discrete ID of item  $v_j$  to  $K$  discrete codebook tokens along with their corresponding codeword embeddings in LLM-based recommender systems as follows:

$$\text{item } v_j \rightarrow \text{tokens: } \{w_j^1, w_j^2, \dots, w_j^K\} \quad (7)$$

$$\rightarrow \text{tokens' embeddings: } [c_{w_j^1}^1, c_{w_j^2}^2, \dots, c_{w_j^K}^K]. \quad (8)$$

It is worth noting that we can also apply a similar quantification process to tokenize users in LLM-based recommendations.

**2.3.4  $K$ -to-1 Decoder.** After the  $K$ -way encoder, a  $K$ -to-1 decoder  $\text{Dec}(\cdot)$  is introduced to conduct input reconstruction. The basic idea is that  $K$  different embeddings indexed by  $K$  discrete tokens in the  $K$ -way codebook  $C$  are fed into the  $K$ -to-1 decoder to reconstruct the input representations of user  $\mathbf{p}_i$  or item  $\mathbf{q}_j$ . Mathematically, given item  $v_j$  and its quantized tokens  $\{w_j^1, w_j^2, \dots, w_j^K\}$ , the decoder first performs average pooling and then generates the reconstructed input representation  $\mathbf{r}_j$  via a three-layer MLP as follows:

$$\mathbf{r}_j = \text{Dec}(\{w_j^1, w_j^2, \dots, w_j^K\}) = \text{MLP}\left(\frac{1}{K} \sum_{k=1}^K c_{w_j^k}^k\right). \quad (9)$$

**2.3.5 Learning Objective.** In order to effectively learn the  $K$ -way encoder, codebook, and  $K$ -to-1 Decoder for user and item MQ-Tokenizers, a *reconstruction loss* is designed to reconstruct the prototype inputs from the set of discrete vectors. To be more specific, given the item  $v_j$ , the reconstruction loss encourages the reconstructed representation  $\mathbf{r}_j$  from the  $K$ -to-1 decoder to approximate original item representation  $\mathbf{q}_j$  learned from GNNs, which can be defined as:

$$\mathcal{L}_{recon}^{Item} = \|\mathbf{q}_j - \mathbf{r}_j\|^2. \quad (10)$$

However, the *arg min* operation in Eq. (5) is non-differentiable [3, 32], leading to an intractable computation for the gradient back-propagation in the reconstruction loss  $\mathcal{L}_{recon}^{Item}$  during optimization. To this end, a straight-through gradient estimator [2, 41] is introduced to directly assign the gradients of the decoder inputs (i.e., the selected tokens' embeddings) to the encoder outputs (i.e., the encoded representations), so as to optimize the encoders and decoder.

Meanwhile, optimizing the reconstruction loss  $\mathcal{L}_{recon}^{Item}$  cannot provide any gradients to update items'  $K$ -way codebook  $C$ . To address this, a *codebook loss*  $\mathcal{L}_{cb}^{Item}$  for items is further designed to bring the selected token's embedding  $c_{w_j^k}^k$  close to the outputs of the  $K$ -way encoder  $\text{Enc}^k(\cdot)$  using the  $l_2$  error for updating  $K$ -way codebook  $C$  as follows:

$$\mathcal{L}_{cb}^{Item} = \sum_{k=1}^K \|\text{sg}[\text{Enc}^k(\mathbf{q}_j)] - c_{w_j^k}^k\|^2, \quad (11)$$

where  $\text{sg}[\cdot]$  denotes the a stop-gradient operator. To be specific, the gradient of the variable in  $\text{sg}[\cdot]$  takes the value 0 when performing back-propagated gradient calculation.

Furthermore, to encourage a smooth gradient passing for the *arg min* operation from Eq. (5), a *commitment loss*  $\mathcal{L}_{cm}^{Item}$  is introduced to prevent the encoded features from fluctuating too frequently from one codeword to another. Unlike the codebook loss, the commitment loss only applies to the encoder weights, which can be calculated by:

$$\mathcal{L}_{cm}^{Item} = \sum_{k=1}^K \|\text{Enc}^k(\mathbf{q}_j) - \text{sg}[c_{w_j^k}^k]\|^2. \quad (12)$$

Finally, the overall optimization objective for item MQ-Tokenizer can be formulated by incorporating the aforementioned losses to jointly update the  $K$ -way encoder, the  $K$ -way codebook, and the  $K$ -to-1 decoder:

$$\mathcal{L}_{MQ}^{Item} = \mathcal{L}_{recon}^{Item} + \mathcal{L}_{cb}^{Item} + \beta^{Item} \times \mathcal{L}_{cm}^{Item}, \quad (13)$$

where  $\beta^{Item}$  is a hyper-parameter that aims to balance the importance of the commitment loss in the overall objective. Additionally, the aforementioned optimization objective for optimizing the user MQ-Tokenizer can be designed as follows:

$$\mathcal{L}_{MQ}^{User} = \mathcal{L}_{recon}^{User} + \mathcal{L}_{cb}^{User} + \beta^{User} \times \mathcal{L}_{cm}^{User}. \quad (14)$$

## 2.4 Generative Retrieval for Recommendations

In this subsection, we introduce a novel framework that takes advantage of LLMs for recommendations. Particularly, a generative retrieval paradigm is designed to generate item representations and retrieve appropriate items for collaborative recommendations. To align large language models with collaborative recommendations, the proposed framework involves three key components to reformulate collaborative filtering: tokenization&prompting, user modeling via LLM, and generative retrieval.

**2.4.1 Tokenization&Prompts.** Tokenization is the most crucial process of splitting the textual input and output into smaller units that can be processed by language models. In general, LLMs have thousands of tokens in their vocabularies (i.e., *in-vocabulary words*). For instance, a representative LLM, LLaMA [40], has a vocabulary size of 32,000. Nonetheless, the number of tokens in LLMs remains relatively small compared to the enormous number of users and items in real-world recommendation scenarios, often reaching millions or billions. Therefore, we introduce *out-of-vocabulary tokens* learned from the proposed *user and item MQ-Tokenizers* to facilitate the tokenization of user and item IDs. Using our proposed MQ-Tokenizers, only a limited number of  $K \times L$  out-of-vocabulary (OOV) tokens are required to tokenize millions or billions of users or items effectively. For example, in our experiments (see Section 3), we can use only 1,536 (i.e.,  $3 \times 512$ ) OOV tokens to tokenize a total of 39,387 items in the Amazon-Clothing dataset. The vocabulary expansion for tokenizing users and items by MQ-Tokenizers is much more efficient and affordable for LLM-based recommendations. In other words, textual input contents are tokenized by the LLM tokenizer (e.g., SentencePiece), while user and item IDs are tokenized into  $K$  discrete tokens (i.e., out-of-vocabulary tokens in LLMs) via the designed user and item MQ-Tokenizers, respectively.

To further enhance the capabilities of LLMs, prompting has been developed to provide explicit guidance to LLMs, enabling

them to comprehend textual inputs better and make predictions for downstream tasks [47, 48]. Hence, to generalize to unseen prompts, we design several prompts instructing the LLM backbone to understand users’ preferences for making recommendations. In these prompts, the ID tokens provided by the well-established MQ-Tokenizers are used to represent users and items in the language space of LLM-based recommender systems. For example, given the sub-codebook number  $K = 3$ , two representative prompts in our method and their user&item tokenizations are defined as follows:

Prompt 1 (without user’s historical interactions):  
I wonder what the **user\_03** will like. Can you help me decide?

---

⇒ I wonder what the **user\_⟨u1-128⟩⟨u2-21⟩⟨u3-35⟩** will like. Can you help me decide?

Prompt 2 (with user’s historical interactions):  
According to what items the **user\_03** has interacted with:  
**item\_08**, **item\_24**, **item\_63**. Can you describe the user’s preferences?

---

⇒ According to what items the **user\_⟨u1-128⟩⟨u2-21⟩⟨u3-35⟩** has interacted with:  
**item\_⟨v1-42⟩⟨v2-12⟩⟨v3-98⟩**,  
**item\_⟨v1-42⟩⟨v2-12⟩⟨v3-87⟩**,  
**item\_⟨v1-42⟩⟨v2-53⟩⟨v3-128⟩**.  
Can you describe the user’s preferences?

Technically, Prompt 1 and Prompt 2 showcase the input prompts without and with item interactions as supporting information, respectively, in which  $\langle uk-\cdot \rangle$  denotes the out-of-vocabulary (OOV) tokens in the  $k_{th}$  sub-codebook. For instance,  $\langle u2-21 \rangle$  represents the 21<sub>st</sub> token in the second sub-codebook for the **user\_03**. This also applies to tokenizing items in LLMs.

**2.4.2 User Modeling via LLM.** The goal of user modeling component is to capture users’ preferences for generating the representations of items that user  $u_i$  may like. A typical input  $\mathcal{X}_i$  of our LLM backbone can be formed by selecting a prompt template  $\mathcal{P}$  and the corresponding ID tokens for user  $u_i$  and his/her interacted items  $\mathcal{N}_{(u_i)}$  in the history as follows:

$$\mathcal{X}_i \rightarrow (\mathcal{P}, \mathcal{T}_{u_i}^c) \text{ or } (\mathcal{P}, \mathcal{T}_{u_i}^c, \{\mathcal{T}_{v_j}^c | v_j \in \mathcal{N}_{(u_i)}\}), \quad (15)$$

where  $\mathcal{T}_{u_i}^c$  represents the ID tokens generated by the user MQ-Tokenizer for user  $u_i$ .  $\{\mathcal{T}_{v_j}^c | v_j \in \mathcal{N}_{(u_i)}\}$  denotes the ID tokens generated by our item MQ-Tokenizer for the items that user  $u_i$  has interacted with. It is worth noting that we can randomly shuffle their interactions in  $\mathcal{N}_{(u_i)}$  for neglecting the consideration of user  $u_i$ ’s sequential signals towards items.

Under the conventional text-to-text generation paradigm (e.g., P5), the user modeling process involves receiving a text input  $\mathcal{X}_i$  for user  $u_i$  and generating descriptive texts for potential items in

an **auto-regressive manner**, which can be formalized as:

$$\mathcal{T}_t = \text{LLM}(\mathcal{X}_i, \mathcal{T}_{1:t-1}), \quad (16)$$

where  $\mathcal{T}_t$  represents the token being generated at the  $t_{th}$  position, while  $\mathcal{T}_{1:t-1}$  denotes the previously generated tokens from the LLM. In contrast, our user modeling process differs from the previous approach. To be more specific, we take the input  $\mathcal{X}_i$  and pass it through the LLM backbone, denoted as  $\text{LLM4Rec}(\cdot)$  to generate a hidden representation  $\mathbf{h}_i$  that reflects the model’s comprehension of user  $u_i$ ’s preferences for next-items recommendations. Mathematically, we can express this process as follows:

$$\mathbf{h}_i = \text{LLM4Rec}(\mathcal{X}_i). \quad (17)$$

In other words, the final encoded representation  $\mathbf{h}_i$  can be considered as user  $u_i$ ’s generative preferences of the next items for making personalized recommendations.

**2.4.3 Generative Retrieval.** In general, LLM-based recommender systems employ auto-regressive generation to decode recommendations in natural language [20, 49], such as producing textual strings like "item\_1234" or "the user love electronics ...". However, the beam search decoding in LLMs can be very time-consuming during auto-regressive generation [43], which is impractical for various real-time recommendation scenarios. Moreover, due to the hallucination issue, it is highly challenging to generate accurate item titles and descriptions when making personalized recommendations. For example, items’ title "iPhone SE, 256 GB, starlight" and "iPhone 15, 256 GB, starlight" share most tokens but are significantly different products - with "iPhone 15, 256 GB, starlight" being a non-factual product. The hallucination issue is likely more severe in e-commerce platforms, in which billions of products are sold, leading to invalid item identifiers for recommendations. Furthermore, title generation in LLM-based recommendations cannot generate unseen items in the fine-tuning stage, which is often infeasible in practice. To this end, we propose a generative retrieval paradigm for LLM-based recommendations, where a simple but effective and efficient strategy is designed to project generative users’ preferences for retrieving potential items from the whole item pool.

More specifically, the hidden state  $\mathbf{h}_i$  from  $\text{LLM4Rec}(\cdot)$  will be projected to a latent representation  $\mathbf{z}_i \in \mathbb{R}^d$  through a projection layer  $\text{Proj}(\cdot)$ , so as to make the alignment between the LLM-generated representation and item representations learned from GNNs as follows:

$$\mathbf{z}_i = \text{Proj}(\mathbf{h}_i), \quad (18)$$

where  $\text{Proj}(\cdot)$  can be modeled by a three-layer MLPs. Note that  $\mathbf{z}_i$  can be considered the generative representation of the next recommended items for user  $u_i$ . After that, we propose to retrieve the  $K$ -nearest items from the whole item set  $\mathcal{V}$  for generating the personalized top- $K$  recommendation list. This can be achieved by measuring the similarity scores between the target user’s generative preference and high-order items’ representations from the well-trained GNNs-based recommendation method. For example, the predicted similarity score  $y_{ij}$  of user  $u_i$  towards item  $v_j$  can be calculated by a matching function (e.g., cosine similarity) between user  $u_i$ ’s generative item representation  $\mathbf{z}_i$  and GNN-based item’s

$v_j$  representation  $\mathbf{q}_j$  which is stored in a vector database:

$$y_{ij} = \frac{\mathbf{z}_i \mathbf{q}_j}{\|\mathbf{z}_i\| \|\mathbf{q}_j\|}. \quad (19)$$

Given the scores of all items in the item base, the proposed TokenRec can easily retrieve top- $K$  items for users rather than directly generating tokens in natural language, thus offering advantages in terms of efficiency in inference and avoiding hallucinations. Furthermore, unseen items in the fine-tuning stage can be retrieved by only updating the item representations pool  $\mathcal{V}$  instead of retraining the entire LLM model.

## 2.5 TokenRec’s Training and Inference

**2.5.1 Training.** Technically, the proposed MQ-Tokenizer is responsible for incorporating high-order collaborative knowledge into ID tokenization, while the LLM4Rec backbone is used to capture user preferences and generate the list of items for recommendations. One straightforward approach is to update these two components jointly. However, the large gap between quantization and language processing makes it difficult to achieve updates synchronously. Here, we first train MQ-Tokenizers for users/items and then freeze the well-trained MQ-Tokenizers components to guide the tuning process of our LLM backbone. More specifically, our training process is as follows:

- **Step 1. Training Users&Items MQ-Tokenizers.** In order to learn users&items ID tokenization, our initial focus lies in training MQ-Tokenizers to quantize the collaborative representations for users and items independently. Specifically, we use the combined losses as given in Eq. (13) and Eq. (14) to train the MQ-Tokenizers for items  $\mathcal{V}$  and users  $\mathcal{U}$ , separately.
- **Step 2. Tuning the LLM4Rec for Generative Retrieval.** In this step, we tune the LLM backbone (e.g., T5), LLM token embeddings, and the projection layer for generative retrieval recommendations while keeping the MQ-Tokenizers frozen. More specifically, the objective of generative retrieval is to identify potential items from the whole item pool, such that top- $K$  ranked items are relevant to the specific query (i.e., generative item representations of users). To achieve such an objective, a general solution is to perform metric learning by predicting the relative similarity or distance based on the dense representations between inputs [5, 31, 34]. In other words, the proposed TokenRec aims to perform the nearest neighbor retrieval, achieved by calculating ranking scores between the user’s generative item representation and collaborative item representations learned from GNNs, so as to retrieve top- $K$  next items to the target user for personalized recommendations. Mathematically, we utilize a pairwise ranking loss for our tuning process, which encourages the query closer to the positive set (i.e., when  $\lambda = 1$ ) than to the negative set by a fixed margin  $\gamma$  (i.e., when  $\lambda = -1$ ) as follows:

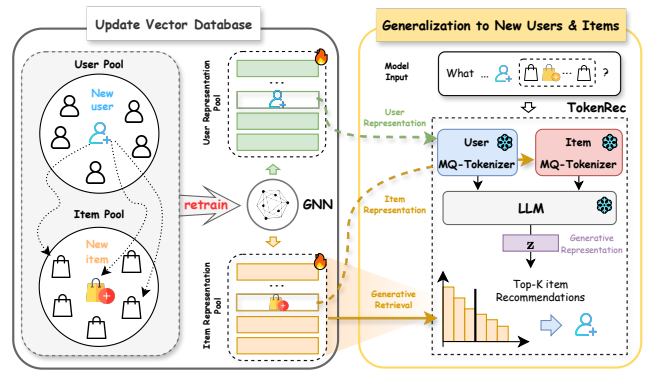
$$\mathcal{L}_{\text{LLM4Rec}} = \begin{cases} 1 - \text{sim}(\mathbf{z}_i, \mathbf{q}_j), & \text{if } \lambda = 1 \\ \max(0, \text{sim}(\mathbf{z}_i, \mathbf{q}_j) - \gamma), & \text{if } \lambda = -1 \end{cases} \quad (20)$$

where  $\mathbf{z}_i$  and  $\mathbf{q}_j$  denote the generative item representation of user  $u_i$  and the collaborative representation of item  $v_j$ , respectively.  $\text{sim}(\cdot, \cdot)$  is a metric function to measure the similarity between dense representations, such as cosine similarity.  $\lambda$  indicates whether user  $u_i$  has interacted with item  $v_j$ . Moreover,  $\gamma$  is

the margin value for negative pairs. It ensures that when the representations of a negative pair are already adequately distant, there is no need to expend additional effort in increasing the distance between them. This mechanism allows more focused training on pairs that are more challenging for recommendations. More specifically, we use cosine similarity as the metric function to maximize the similarity between user  $u_i$ ’s generative representation  $\mathbf{z}_i$  of next item and the *positive item*  $v_j$  representation  $\mathbf{q}_j$  (i.e., item  $v_j$  in user  $u_i$ ’s historical interactions) when  $\lambda = 1$ , while  $\lambda = -1$  otherwise.

**2.5.2 Inference.** The typical inference of LLMs refers to the process of generating target tokens. It suffers from laborious generation and insufficient modeling for unseen users and items. By utilizing the generative retrieval framework, our inference process overcomes these challenges in an efficient way. Specifically, the advantages of our proposed method can be summarized as follows:

- **Efficient Recommendations.** TokenRec proposes a novel LLM-empowered collaborative recommendation framework in generative retrieval paradigms, bypassing the time-consuming decoding process. In particular, the proposed LLM-based recommendation framework aims to output generative item representation and retrieve appropriate items for collaborative recommendations instead of a sequence of discrete tokens (e.g., item titles). This contributes to reducing the considerable computing cost of online RecSys. The efficiency evaluation can be found in Section 3.4 and Table 5.
- **Generalizability to New Users and Items.** By employing masking and  $K$ -way encoder mechanisms in vector quantization, the proposed architecture can provide robust ID tokenization for unseen users and items in fine-tuning the LLM4Rec component. For instance, as shown in Figure 3, when new users and items



**Figure 3: The TokenRec’s efficiency and generalization capability for new users and items during the inference stage. Rather than retraining the entire model, which can be computationally expensive and time-consuming, only the GNN needs to be updated for learning representations for new users and items while keeping the MQ-Tokenizers and LLM backbone frozen. This significantly reduces the computational burden and allows faster adaptation to new users and items.**

are added to the system, retraining is only required for the collaborative filtering component (e.g., MF and LightGCN) to capture collaborative knowledge in learning representations of new users and items (i.e., updating vector database). The MQ-Tokenizers and LLM backbone can remain frozen, saving computational resources and eliminating the need for retraining these components. Compared to finetuning LLMs, the training of GNNs is far more efficient [45]. The capability in adapting to new users and items is demonstrated in Section 3.3 and Table 4.

- **Concise Prompts.** TokenRec provides an inference alternative that relies solely on user ID tokens, e.g., Prompt 1 in Section 2.4.1, for LLM-based recommendation generation. This is made possible by incorporating the collaborative knowledge of users into user ID tokens through our MQ-Tokenizer. By doing so, TokenRec eliminates the necessity of including interacted items in inputs, thus reducing significant computing resources during inference. Additionally, this mechanism proves advantageous when dealing with users who have interacted with a large number of items, effectively avoiding the prevalent issue of context length limitation of many LLMs [24, 30], e.g., 512 tokens in T5 [35] and 2048 tokens in ChatGPT. The recommendation performance using user ID tokens only for model input can be observed in Table 2 and Table 3.

### 3 EXPERIMENT

#### 3.1 Experimental Settings

**3.1.1 Datasets.** To demonstrate the effectiveness of the proposed method, we conduct comprehensive experiments over four benchmark datasets: Amazon-Beauty (Beauty for short), Amazon-Clothing (Clothing for short), LastFM, and MovieLens 1M (ML1M for short). The first two datasets<sup>2</sup> are obtained from the *amazon.com* e-commerce platform, encompassing a broad spectrum of user interactions with Beauty and Clothing products. The LastFM<sup>3</sup> dataset provides music artist listening records from users at the Last.fm online music system. The ML1M<sup>4</sup> dataset offers a collection of movie ratings made by MovieLens users. Table 1 shows the statistics of these four datasets, where the maximum item sequence length is set to 100 to accommodate the input length of the LLM backbone T5 (512 tokens). Moreover, our training-validation-testing combination follows the leave-one-out policy [20], i.e., using all but the last observation in users’ interaction history as the training set. Last but not least, we randomly shuffle users’ interaction history to provide a non-chronological item list to align with the settings of collaborative filtering methods.

Table 1: Basic statistics of benchmark datasets.

| Datasets | User-Item Interaction |        |               |             |
|----------|-----------------------|--------|---------------|-------------|
|          | #Users                | #Items | #Interactions | Density (%) |
| LastFM   | 1,090                 | 3,646  | 37,080        | 0.9330      |
| ML1M     | 3,416                 | 6,040  | 447,294       | 2.1679      |
| Beauty   | 22,363                | 12,101 | 197,861       | 0.0731      |
| Clothing | 23,033                | 39,387 | 278,641       | 0.0307      |

<sup>2</sup><https://nijianmo.github.io/amazon/>

<sup>3</sup><https://grouplens.org/datasets/hetrec-2011/>

<sup>4</sup><https://grouplens.org/datasets/movielens/1m/>

**3.1.2 Baselines.** Here, we compare our approach with four representative collaborative filtering methods (i.e., MF, NeuCF, LightGCN, GTN, and LTGNN), three widely-used sequential recommendation methods (i.e., SASRec, BERT4Rec, and S<sup>3</sup>Rec), and four state-of-the-art LLM-based recommendation methods (i.e., P5, CID, POD, and CoLLM).

- **Collaborative Filtering: MF** [37] is the most classic CF method, while **NeuCF** [22] is the very first DNNs based collaborative filtering. **LightGCN** [21] and **GTN** [14] are representative CF based on GNNs techniques. **LTGNN** [45] is the most advanced GNNs-based collaborative filtering method.
- **Sequential Recommendations: SASRec** [25] is a self-attention based sequential recommendation model. **BERT4Rec** [38] is a bidirectional Transformer-based recommender trained with the BERT-style cloze task. **S<sup>3</sup>Rec** [50] is a representative sequential recommendation model trained by self-supervised learning.
- **P5** [20] is a pioneering work on LLM-based RecSys, which describes recommendation tasks in a text-to-text format and employs LLMs to capture deeper semantics for personalization and recommendation. In our experiments, we deploy two indexing methods, i.e., *random indexing (RID)* and *sequential indexing (SID)*, on the P5 model. Among these, P5-SID is chosen as a baseline for our efficiency evaluation, generalizability evaluation, and ablation study in our experiments.
- **CID** [23] is a non-trivial indexing approach that considers the co-occurrence matrix of items to design numeric IDs so that items co-occur in user-item interactions will have similar numeric IDs. To be consistent, we employ the P5 model as its LLM backbone.
- **POD** [28] encodes discrete prompts into continuous embeddings to reduce the excessive input length of LLMs based on P5 architecture.
- **CoLLM** [46] employs GNNs to provide continuous embeddings representing items and users for LLM-based recommendations. Notably, in our experiments, CoLLM’s binary classification output (i.e., whether or not a user likes an item) is reformatted to generate item IDs consistent with the output setting of top-K recommendations.

**3.1.3 Evaluation Metrics.** In order to evaluate the quality of recommendation results, we adopt two widely used evaluation metrics: top-K Hit Ratio (HR@K) and top-K Normalized Discounted Cumulative Gain (NDCG@K) [44, 45], in which higher values indicate better performance for recommendations. And the average metrics for all users in the test set are reported. In addition, we set the values of K as 10, 20, and 30, among which 20 is the default value for ablation experiments.

**3.1.4 Hyper-parameter Settings.** Our proposed model is implemented based on Hugging Face and PyTorch. The codebook number K (i.e., the subencoder), the token number L at each subcodebook, and the ratio  $\rho$  of our masking operation are search in the ranges of {1, 2, 3, 4, 5}, {128, 256, 512, 1024}, and {0 to 1} in 0.1 increments, respectively. Moreover, during the fine-tuning process of our proposal LLM-based recommendation framework, the ratio of negative sampling  $\lambda$  presented in Eq.(20) is fixed at 1:1, in which we randomly select an un-interacted item from the whole item base as the negative sample for each positive sample. The margin  $\gamma$  in



Eq.(20) is set to 0 to 0.2. We optimize the MQ-Tokenizers and the LLM backbone with AdamW [33] in a mini-batch manner. Note that the high-order collaborative representations for users and items are obtained from a representative collaborative filtering method, namely LightGCN [21]. For prompting, we design 11 templates for TokenRec: 10 of them are set to be *seen* prompts, and the remaining one is evaluated as an *unseen* prompt. For a fair comparison, we employ a widely-used lightweight LLM, i.e., *T5-small* [35], for our LLM4Rec model and all LLM-based baselines. The other default hyper-parameters for baseline methods are set as suggested by the corresponding papers.

### 3.2 Performance Comparison of Recommender Systems

We first compare the recommendation performance between TokenRec and all baselines over four benchmark datasets. Table 2 and Table 3 present the overall performance comparison on the four datasets, where \* denotes the proposed TokenRec, and the best and second best results are marked by **Bold** and underlined. Notably, the proposed model using unseen prompts will be indicated by the suffix (Unseen Prompt). Besides, the suffix of (User ID Only) denotes the case of using **user ID tokens** only for LLM-empowered recommendation without considering users’ interaction history towards items. We make the following observations:

- Our proposed TokenRec achieves the best performance and consistently outperforms all the baselines across all datasets regarding metrics with either unseen personalized prompts. On average, TokenRec significantly exceeds the strongest baselines by 19.08% on HR@20 and 9.09% on NCDG@20 in the LastFM dataset. Such improvement demonstrates the effectiveness of our proposed method and the great potential of exploring collaborative indexing (i.e., tokenization) in LLM-based RecSys.
- Even when using only user ID tokens, TokenRec surpasses most baselines in terms of accuracy, indicating its superior performance in collaborative recommendations. Such a result implies that TokenRec is capable of effectively modeling user even in the absence of their interaction history. Rather than spending a large number of tokens describing users’ interacted items, this allows TokenRec to generate recommendations using a concise input, thus circumventing the input length restrictions imposed by LLMs and saving considerable computation resources.
- As an earlier non-trivial indexing method to capture hidden knowledge in co-occurrence frequency, Collaborative Indexing (CID) outperforms Random Indexing (RID) and Sequential Indexing (SID) given the same setting of P5. These observations suggest the potential of integrating collaborative knowledge for item&user tokenization/indexing. However, the P5 variants and POD are inferior to the existing advanced GNN-based collaborative filtering (i.e., LightGCN, GTN, and LTGNN), implying their inability to capture collaborative information using LLMs.
- The GNN-based collaborative filtering methods perform relatively better than the traditional CF methods (i.e., MF and NCF) and representative sequential recommendation methods (i.e., BERT4Rec, SASRec, and S<sup>3</sup>Rec). The results demonstrate the effectiveness of GNNs in capturing collaborative signals via high-order connectivity.

### 3.3 Generalizability Evaluation

In most e-commerce and social media platforms, a significant number of new users and items are added daily to recommender systems. As a result, well-established systems are required to frequently conduct updates and enhancements to accommodate and generalize to the preferences of new users and the characteristics of new items, thereby providing personalized recommendations that cater to the evolving dynamics of their user base. Since newly added users and items lack their interactions in fine-tuning LLM4Rec, most existing LLM-based RecSys fail to retrieve suitable items as potential candidates, thus requiring extensive retraining. In contrast, our proposed TokenRec can easily generalize effectively even in cases where the user is not present in the training and fine-tuning corpus.

For this analysis, we consider the Beauty and LastFM datasets and exclude the 5% of users (ref as *unseen users*) with the least interaction history from the training data split to simulate newly added users. In this evaluation, the performance of TokenRec is compared to other LLM-based recommendation models, i.e., P5, POD, CID, and CoLLM. To ensure the absence of data leakage concerning unseen users, only the training split is utilized for establishing these LLM4Rec models, while it is allowed to update the vector database to provide collaborative representations of unseen users and items (as shown in Figure 3), which consumes far less computational resources than updating the LLM backbone of LLM-based RecSys. From the recommendation results presented in Table 4, several key observations can be made:

- Existing LLM-based recommendation methods face challenges in generalizability, as evidenced by a significant drop of over 40% in HR@20 and NDCG@20 for P5 and POD when recommending items to unseen users.
- The inclusion of collaborative knowledge in CID and CoLLM leads to a relatively improved performance in model generalization, as indicated by a reduced performance degradation in P5 and POD. However, these methods still experience more than a 20% drop, indicating that they overlook the importance of stable ID tokenization for LLM-based recommendations.
- In comparison, TokenRec outperforms all the compared methods for not only the training users but also the unseen users in both datasets. For instance, in the Amazon-Beauty dataset, the performance of TokenRec decreases by only 7% on average, demonstrating the strong generalization capability of TokenRec for newly added users. Such a superiority can be attributed to our MQ-Tokenizer for robust ID tokenization and the generative retrieval paradigm for flexible recommendation generation.

### 3.4 Efficiency Evaluation

In this subsection, we analyze the inference efficiency of TokenRec compared with the general LLM-based recommendation methods (i.e., P5, CID, and POD), which generate the tokens (e.g., IDs, titles, and descriptions) for Top-K items by using auto-regressive decoding and beam search. Our proposed method discards this text-decoding generation solution and introduces a generative retrieval paradigm to perform collaborative recommendations. The results, as presented in Table 5, demonstrate that TokenRec can achieve superior inference efficiency with a significant improvement of

**Table 2: Performance comparison of recommendation algorithms on the LastFM and ML1M datasets.**

| Model              | LastFM        |               |               |               |               |               | ML1M          |               |               |               |               |               |
|--------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
|                    | HR@10         | HR@20         | HR@30         | NG@10         | NG@20         | NG@30         | HR@10         | HR@20         | HR@30         | NG@10         | NG@20         | NG@30         |
| BERT4Rec           | 0.0319        | 0.0461        | 0.0640        | 0.0128        | 0.0234        | 0.0244        | 0.0779        | 0.1255        | 0.1736        | 0.0353        | 0.0486        | 0.0595        |
| SASRec             | 0.0345        | 0.0484        | 0.0658        | 0.0142        | 0.0236        | 0.0248        | 0.0785        | 0.1293        | 0.1739        | 0.0367        | 0.052         | 0.0622        |
| S <sup>3</sup> Rec | 0.0385        | 0.0490        | 0.0689        | 0.0177        | 0.0266        | 0.0266        | 0.0867        | 0.1270        | 0.1811        | 0.0361        | 0.0501        | 0.0601        |
| MF                 | 0.0239        | 0.0450        | 0.0569        | 0.0114        | 0.0166        | 0.0192        | 0.078         | 0.1272        | 0.1733        | 0.0357        | 0.0503        | 0.0591        |
| NCF                | 0.0321        | 0.0462        | 0.0643        | 0.0141        | 0.0252        | 0.0254        | 0.0786        | 0.1273        | 0.1738        | 0.0363        | 0.0504        | 0.0601        |
| LightGCN           | 0.0385        | 0.0661        | 0.0982        | 0.0199        | 0.0269        | 0.0336        | 0.0877        | 0.1288        | 0.1813        | 0.0374        | 0.0509        | 0.0604        |
| GTN                | 0.0394        | 0.0688        | 0.0963        | 0.0199        | 0.0273        | 0.0331        | 0.0883        | 0.1307        | 0.1826        | 0.0378        | 0.0512        | 0.0677        |
| LTGNN              | 0.0471        | 0.076         | 0.0925        | 0.0234        | 0.0318        | 0.0354        | 0.0915        | 0.1387        | 0.1817        | 0.0419        | 0.0570        | 0.0659        |
| P5-RID             | 0.0312        | 0.0523        | 0.0706        | 0.0144        | 0.0199        | 0.0238        | 0.0867        | 0.1248        | 0.1811        | 0.0381        | 0.0486        | 0.0662        |
| P5-SID             | 0.0375        | 0.0536        | 0.0851        | 0.0224        | 0.0255        | 0.0261        | 0.0892        | 0.1380        | 0.1784        | 0.0422        | 0.0550        | 0.0641        |
| CID                | 0.0381        | 0.0552        | 0.0870        | 0.0229        | 0.0260        | 0.0277        | 0.0901        | 0.1294        | 0.1863        | 0.0379        | 0.0525        | 0.0706        |
| POD                | 0.0367        | 0.0572        | 0.0747        | 0.0184        | 0.0220        | 0.0273        | 0.0886        | 0.1277        | 0.1846        | 0.0373        | 0.0487        | 0.0668        |
| CoLLM              | 0.0483        | 0.0786        | 0.1017        | 0.0234        | 0.0319        | 0.0366        | 0.0923        | 0.1499        | 0.1998        | 0.0456        | 0.0620        | 0.0719        |
| * (User ID Only)   | 0.0505        | 0.0881        | <u>0.1128</u> | <u>0.0251</u> | <u>0.0345</u> | 0.0397        | 0.0964        | 0.1546        | 0.2043        | 0.0493        | 0.0640        | 0.0745        |
| * (Unseen Prompt)  | <u>0.0514</u> | <u>0.0917</u> | <b>0.1294</b> | <b>0.0252</b> | 0.0343        | <b>0.0422</b> | <b>0.1012</b> | <u>0.1672</u> | <u>0.2144</u> | <b>0.0532</b> | <b>0.0698</b> | <b>0.0798</b> |
| TokenRec           | <b>0.0532</b> | <b>0.0936</b> | 0.1248        | 0.0247        | <b>0.0348</b> | <u>0.0415</u> | <u>0.1008</u> | <b>0.1677</b> | <b>0.2149</b> | <u>0.0528</u> | <u>0.0697</u> | <u>0.0797</u> |

\* are the variants of **TokenRec**, namely the cases of using user ID tokens only for model inputs without considering item interaction history and using the unseen prompt during evaluation.

**Table 3: Performance comparison of recommendation algorithms on the Beauty and Clothing datasets.**

| Model              | Beauty        |               |               |               |               |               | Clothing      |               |               |               |               |               |
|--------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
|                    | HR@10         | HR@20         | HR@30         | NG@10         | NG@20         | NG@30         | HR@10         | HR@20         | HR@30         | NG@10         | NG@20         | NG@30         |
| BERT4Rec           | 0.0329        | 0.0464        | 0.0637        | 0.0162        | 0.0205        | 0.0255        | 0.0135        | 0.0217        | 0.0248        | 0.0061        | 0.0074        | 0.0079        |
| SASRec             | 0.0338        | 0.0472        | 0.0637        | 0.0170        | 0.0213        | 0.0260        | 0.0136        | 0.0221        | 0.0256        | 0.0063        | 0.0076        | 0.0081        |
| S <sup>3</sup> Rec | 0.0351        | 0.0471        | 0.0664        | 0.0169        | 0.0237        | 0.0278        | 0.0140        | 0.0213        | 0.0256        | 0.0069        | 0.0081        | 0.0086        |
| MF                 | 0.0127        | 0.0195        | 0.0245        | 0.0063        | 0.0081        | 0.0091        | 0.0116        | 0.0175        | 0.0234        | 0.0074        | 0.0088        | 0.0101        |
| NCF                | 0.0315        | 0.0462        | 0.0623        | 0.0160        | 0.0196        | 0.0237        | 0.0119        | 0.0178        | 0.024         | 0.0072        | 0.0090        | 0.0103        |
| LightGCN           | 0.0344        | 0.0498        | 0.0630        | 0.0194        | 0.0233        | 0.0261        | 0.0157        | 0.0226        | 0.0279        | 0.0085        | 0.0103        | 0.0114        |
| GTN                | 0.0345        | 0.0502        | 0.0635        | 0.0198        | 0.0241        | 0.0268        | 0.0158        | 0.0226        | 0.0282        | 0.0084        | 0.0103        | 0.0111        |
| LTGNN              | 0.0385        | 0.0564        | 0.0719        | 0.0207        | 0.0252        | 0.0285        | 0.0155        | 0.0218        | 0.0272        | 0.0082        | 0.0110        | 0.0116        |
| P5-RID             | 0.0330        | 0.0511        | 0.0651        | 0.0146        | 0.0200        | 0.0144        | 0.0148        | 0.0225        | 0.0263        | 0.0071        | 0.0086        | 0.0095        |
| P5-SID             | 0.0340        | 0.0516        | 0.0672        | 0.0154        | 0.0231        | 0.0176        | 0.0143        | 0.0222        | 0.0258        | 0.0070        | 0.0086        | 0.0091        |
| CID                | 0.0341        | 0.0516        | 0.0673        | 0.0165        | 0.0236        | 0.0177        | 0.0146        | 0.0226        | 0.0276        | 0.0070        | 0.0087        | 0.0092        |
| POD                | 0.0339        | 0.0498        | 0.0639        | 0.0185        | 0.0222        | 0.0221        | 0.0147        | 0.0225        | 0.0261        | 0.0074        | 0.0087        | 0.0091        |
| CoLLM              | 0.0391        | 0.0606        | 0.0772        | 0.0200        | 0.0259        | 0.0303        | 0.0150        | 0.0218        | 0.0274        | 0.0079        | 0.0091        | 0.0117        |
| * (User ID Only)   | 0.0396        | 0.0599        | 0.0763        | 0.0214        | 0.0265        | 0.0300        | 0.0160        | 0.0228        | 0.0282        | 0.0092        | 0.0109        | 0.0119        |
| * (Unseen Prompt)  | <u>0.0402</u> | <b>0.0622</b> | <b>0.0791</b> | <u>0.0215</u> | <u>0.0270</u> | <b>0.0306</b> | <u>0.0164</u> | <u>0.0233</u> | <u>0.0286</u> | <u>0.0096</u> | <u>0.0111</u> | <u>0.0124</u> |
| TokenRec           | <b>0.0407</b> | <u>0.0615</u> | <u>0.0782</u> | <b>0.0222</b> | <b>0.0276</b> | <u>0.0303</u> | <b>0.0171</b> | <b>0.0240</b> | <b>0.0291</b> | <b>0.0108</b> | <b>0.0112</b> | <b>0.0130</b> |

\* are the variants of **TokenRec**, namely the cases of using user ID tokens only for model inputs without considering item interaction history and using the unseen prompt during evaluation.

approximately 1306.5% compared to the LLM-based recommendation baselines. This can be attributed to our generative retrieval paradigm, bypassing the most time-consuming auto-regressive decoding and beam search processes of LLMs [43].

### 3.5 Ablation Study

In order to assess the effectiveness of the proposed key components, we conducted ablation experiments on the LastFM and Amazon-Beauty datasets, where the influence of each component was eliminated separately as follows:

- w/o High-Order Collaborative Knowledge (HOCK): Use Matrix Factorization (MF) to learn the collaborative representations of users and items for ID tokenization in MQ-Tokenizers. Compared to advanced GNN-based methods (e.g., LightGCN), it is challenging to capture high-order collaborative signals among user-item interactions explicitly.
- w/o  $K$ -way: Replace the  $K$ -way encoder and codebook with 1-way ones while keeping the total number of codebook tokens (i.e., codewords) consistent.

**Table 4: Performance comparison on seen and unseen users for generalizability evaluation.**

| Dataset | Model    | Seen          |               | Unseen        |               |
|---------|----------|---------------|---------------|---------------|---------------|
|         |          | HR@20         | NG@20         | HR@20         | NG@20         |
| LastFM  | P5       | 0.0704        | 0.0320        | 0.0399        | 0.0137        |
|         | POD      | 0.0709        | 0.0323        | 0.0401        | 0.0138        |
|         | CID      | 0.0697        | 0.0314        | 0.0452        | 0.0196        |
|         | CoLLM    | <u>0.0812</u> | <u>0.0336</u> | <u>0.0574</u> | <u>0.0235</u> |
|         | TokenRec | <b>0.0973</b> | <b>0.0353</b> | <b>0.0773</b> | <b>0.0268</b> |
| Beauty  | P5       | 0.0511        | 0.0236        | 0.0274        | 0.0130        |
|         | POD      | 0.0507        | 0.0225        | 0.0269        | 0.0123        |
|         | CID      | 0.5234        | 0.0240        | 0.3336        | 0.0146        |
|         | CoLLM    | <u>0.0612</u> | <u>0.0261</u> | <u>0.0477</u> | <u>0.0195</u> |
|         | TokenRec | <b>0.0629</b> | <b>0.0289</b> | <b>0.0591</b> | <b>0.0266</b> |

**Table 5: Average inference time (milliseconds) per user.**

| Inference Time | LastFM | ML1M   | Beauty | Clothing |
|----------------|--------|--------|--------|----------|
| P5             | 96.04  | 99.75  | 86.39  | 93.38    |
| POD            | 96.30  | 101.42 | 87.69  | 94.48    |
| CID            | 94.96  | 99.42  | 84.87  | 92.02    |
| TokenRec       | 6.92   | 8.43   | 5.76   | 6.00     |
| Acceleration*  | 1284%  | 1089%  | 1398%  | 1455%    |

\* The average improvement in inference time compared to the baselines.

- w/o Masking: Deactivate the masking operation in users&items MQ-Tokenizers.

From the ablation results in Table 6, we can have the following observations. First, each component in our approach contributes to the overall performance since eliminating any of them would result in performance degradation. Second, the masking operation and the  $K$ -way framework introduced in our approach not only enhance the generalizability, as demonstrated in Table 4, but also result in moderate performance improvements during the ablation experiment. Finally, the removal of the high-order collaborative knowledge learned by advanced GNNs resulted in a significant drop in performance. This indicates the importance of incorporating such knowledge for the alignment of LLMs and personalized recommendations.

**Table 6: Results of Ablation Studies**

| Module       | LastFM |        | Beauty |        |
|--------------|--------|--------|--------|--------|
|              | HR@20  | NG@20  | HR@20  | NG@20  |
| w/o HOCK     | 0.0549 | 0.0172 | 0.0407 | 0.0149 |
| w/o $K$ -way | 0.0820 | 0.0309 | 0.0592 | 0.0250 |
| w/o Masking  | 0.0848 | 0.0332 | 0.0573 | 0.0253 |
| Full*        | 0.0936 | 0.0348 | 0.0615 | 0.0276 |

\* "Full" denotes the complete version of TokenRec.

### 3.6 Hyper-parameter Analysis

In TokenRec, we introduce three critical hyper-parameters, namely the masking ratio  $\rho$  in vector quantization, the number of sub-encoders/sub-codebooks  $K$ , and the number of tokens  $L$  in each sub-codebook. Their value sensitivities are evaluated in this section to facilitate the future application of our proposed model.

**3.6.1 Effect of Masking Ratio  $\rho$ .** We first investigate the impact of the hyper-parameter  $\rho$  in MQ-Tokenizer, which controls the masking ratio of our vector quantization for users& items ID tokenization. Figure 4 shows the performance change of TokenRec w.r.t. HR@20 and NDCG@20. We can find that introducing a small ratio of masking brings performance improvements. In most cases, the recommendation performance of our proposed method improves when  $\rho < 0.5$ , among which 0.2 can achieve the best improvement in our experiments. The experimental results also reveal that the recommendation performance degrades when the masking ratio  $\rho \geq 0.5$ , suggesting excessive masking should be avoided.

**3.6.2 Effect of Codebook Settings  $K$  and  $L$ .** To study whether our proposed method TokenRec can benefit from stacking more sub-codebooks and introducing more codebook tokens, we vary the numbers of the hyper-parameters  $K$  and  $L$  in the range of  $\{1, 2, 3, 4, 5\}$  and  $\{128, 256, 512, 1024\}$ , respectively, and report the performance on all datasets in Figure 5, with the optimal cases highlighted by black box. From the figure, we make the following observations.

- With the increase of codebook depth, a progressive performance improvement of our model for recommendations is witnessed in all datasets. Nevertheless, when  $K > 3$ , the improvement of recommendation performance becomes relatively marginal. As a result,  $K = 3$  is suggested by considering the trade-off between effectiveness and efficiency.
- The optimal value of  $L$  varies in relation to the differing sizes of users and items. Specifically, the best-balanced performance in terms of effectiveness and efficiency can be observed at a  $L$  of 256 and 512 on the LastFM/ML1M (smaller sizes) and the Amazon-Beauty/Clothing (larger sizes) datasets, respectively. This implies that slightly more codebook tokens should be used in each sub-codebook for datasets with more users/items.
- Additionally, in the case of a single codebook, merely increasing the number of codebook tokens cannot effectively deliver performance gains in recommendations. This demonstrates the effectiveness of the proposed  $K$ -way mechanism in our MQ-Tokenizer for recommendations.

## 4 RELATED WORK

In this section, we briefly review related work, including collaborative filtering and Large Language Models (LLMs).

### 4.1 Collaborative Filtering

In order to provide personalized recommendations that accord with user preferences, collaborative filtering (CF) serves as a representative technique for modeling collaborative information in recommender systems, such as user-item interactions, to capture similar user patterns and predict future interactions [11, 12]. As a representative example, matrix factorization (MF) [13, 26]




Figure 4: The effect of masking ratio  $\rho$  under HR@20 and NDCG@20 metrics.




Figure 5: The effect of the number of sub-codebooks  $K$  and the number of tokens in each sub-codebook  $J$  under HR@20 and NDCG@20 metrics.

vectorizes users and items into dense representations and models their interactions (i.e., user-item ID matrix) by calculating the inner products between the vectorized representations. Subsequently, NeuCF [22] incorporates neural networks with MF to decompose the use-item interactions into two low-rank matrices representing user&item embeddings. Later on, DSCF [17] takes advantage of deep language models to enhance user representations for collaborative recommendations by capturing auxiliary information from neighbors in social networks. Due to the superior representation learning capability in graphs, Graph Neural Networks (GNNs), such as LightGCN [21], GTN [14], and LTGNN [45], are proposed to

capture high-order collaborative knowledge on user-item interaction graphs for enhancing the performance of recommender systems. More specifically, GNNs take advantage of the graph-structured nature of user-item interactions and model similar user behavior patterns toward items through information propagation on user-item interaction graphs. For example, GraphRec [15, 16] introduces a graph attention network-based framework to encode user-item interactions and user-user social relations for social recommendations. LightGCN [21] is introduced to largely simplify the GNN-based recommendation methods by removing feature

transformation and nonlinear activation, achieving state-of-the-art prediction performance for recommendations. In addition, GTN [16] and LTGNN [45] provide an improvement by capturing the adaptive reliability and high-order linear-time patterns of interactions, respectively.

## 4.2 LLM-based Recommender Systems

With the rapid development of Large Language Models (LLMs), such as ChatGPT and GPT-4, notable milestones have been showcased for revolutionizing natural language processing techniques [8, 9]. Technically, the unprecedented power of LLMs can be attributed to the scaling up of model/parameter size alongside a tremendous amount of training corpus [18, 27]. In particular, LLMs equipped with billion-scale parameters have exhibited unprecedented language understanding and generation ability, along with remarkable generalization capability and reasoning skills that facilitate LLMs to better generalize to unseen tasks and domains [7]. Given the emerging trends and aforementioned advancements of LLMs, LLM-empowered recommender systems have drawn increasing attention from recent studies and demonstrated distinctive abilities for advancing recommender systems [43, 48]. Notably, to harness the distinctive capabilities of LLMs for advancing recommender systems, existing studies have actively investigated various paradigms [47], including pre-training, fine-tuning, and prompting (e.g., in-context learning, chain-of-thought, and instruction tuning), for adapting LLMs to recommendation tasks. For example, P5 [20] introduces an LLM-based recommendation model that unifies diverse recommendation tasks by multi-task prompt-based pre-training, which achieves impressive zero-shot generalization capability to unseen recommendation tasks with the help of personalized prompts.

Despite effectiveness, most LLM-based recommendation methods still have an intrinsic limitation on indexing users and items' IDs in language models. The naive approach known as Independent Indexing (IID) directly assigns special tokens to represent users and items within language models. Moreover, as a natural solution, textual title indexing is proposed to utilize textual contents (e.g., titles and description) to tokenize items using LLMs' in-vocabulary tokens [1, 19], such as the example "*Apple iPhone 15, 256 GB, black*". P5 [20] and POD [28] apply positional and whole-word embeddings to highlight the tokens representing items and users. Most recently, several studies borrow the conception of soft prompt and utilize exogenous tokens with continuous embeddings to represent users and items in LLM-based recommendations [29, 46]. Although the use of title indexing and whole-word embeddings can mitigate the issue of vocabulary explosion, it fails to effectively capture high-order collaborative knowledge and lacks generalizability for recommending to unseen users or items, as seen in P5 [20] and POD [28]. Furthermore, the discrete nature of tokens in language models presents a challenge in achieving tight alignment of LLMs in recommendations when utilizing continuous representations. Additionally, most LLM-based RecSys suffer from time-consuming inference due to auto-regressive decoding and beam search. However, the majority of existing methods for users&items tokenization fail to effectively capture high-order collaborative knowledge in personalized recommendations. Thus, this paper proposes a novel framework for LLM-based recommendations, TokenRec, which not

only introduces a generalizable ID tokenization strategy to capture high-order collaborative knowledge but also proposes a generative retrieval paradigm to generate top- $K$  items efficiently.

## 5 CONCLUSION

While existing large language model based recommendation methods achieve promising prediction performance, they fail to capture high-order collaborative knowledge and suffer from inferior generalization capability for tokenizing users&items. Additionally, the time-consuming inference remains as the emerging challenge in LLM-based recommender systems. To tackle these shortcomings, we propose a novel approach, named TokenRec, which not only introduces a generalizable user&item ID tokenization strategy to capture high-order collaborative knowledge but also presents a generative retrieval paradigm for the efficient generation of top- $K$  items. Particularly, Masked Vector-Quantized Tokenizer (MQ-Tokenizer) is developed to tokenize users and items in LLM-based recommendations by incorporating high-order collaborative knowledge. Through comprehensive experiments on four distinct datasets, we demonstrated that our model can achieve state-of-the-art recommendation performance while also exhibiting the capacity to generalize to unseen users.

## REFERENCES

- [1] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*.
- [2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [3] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. 2020. Learning with differentiable perturbed optimizers. *Advances in neural information processing systems* 33 (2020), 9508–9519.
- [4] 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.
- [5] Fatih Cakir, Kun He, Xide Xia, Brian Kulis, and Stan Sclaroff. 2019. Deep metric learning to rank. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1861–1870.
- [6] Xiao Chen, Wenqi Fan, Jingfan Chen, Haochen Liu, Zitao Liu, Zhaoxiang Zhang, and Qing Li. 2023. Fairly adaptive negative sampling for recommendations. In *Proceedings of the ACM Web Conference 2023*. 3723–3733.
- [7] Zhikai Chen, Haitao Mao, Hang Li, Wei Jin, Hongzhi Wen, Xiaochi Wei, Shuaiqiang Wang, Dawei Yin, Wenqi Fan, Hui Liu, et al. 2024. Exploring the potential of large language models (llms) in learning on graphs. *ACM SIGKDD Explorations Newsletter* 25, 2 (2024), 42–61.
- [8] Yujuan Ding, Wenqi Fan, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. A Survey on RAG Meets LLMs: Towards Retrieval-Augmented Large Language Models. *arXiv preprint arXiv:2405.06211* (2024).
- [9] Yujuan Ding, Yunshan Ma, Wenqi Fan, Yige Yao, Tat-Seng Chua, and Qing Li. 2024. FashionReGen: LLM-Empowered Fashion Report Generation. *arXiv preprint arXiv:2403.06660* (2024).
- [10] Patrick Esser, Robin Rombach, and Bjorn Ommer. 2021. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 12873–12883.
- [11] Wenqi Fan, Tyler Derr, Yao Ma, Jianping Wang, Jiliang Tang, and Qing Li. 2019. Deep Adversarial Social Recommendation. In *28th International Joint Conference on Artificial Intelligence (IJCAI-19)*. International Joint Conferences on Artificial Intelligence, 1351–1357.
- [12] Wenqi Fan, Tyler Derr, Xiangyu Zhao, Yao Ma, Hui Liu, Jianping Wang, Jiliang Tang, and Qing Li. 2021. Attacking black-box recommendations via copying cross-domain user profiles. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 1583–1594.
- [13] Wenqi Fan, Qing Li, and Min Cheng. 2018. Deep modeling of social relations for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*,

- [14] Wenqi Fan, Xiaorui Liu, Wei Jin, Xiangyu Zhao, Jiliang Tang, and Qing Li. 2022. Graph trend filtering networks for recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 112–121.
- [15] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *WWW*. 417–426.
- [16] Wenqi Fan, Yao Ma, Qing Li, Jianping Wang, Guoyong Cai, Jiliang Tang, and Dawei Yin. 2020. A graph neural network framework for social recommendations. *IEEE Transactions on Knowledge and Data Engineering* 34, 5 (2020), 2033–2047.
- [17] Wenqi Fan, Yao Ma, Dawei Yin, Jianping Wang, Jiliang Tang, and Qing Li. 2019. Deep social collaborative filtering. In *Proceedings of the 13th ACM conference on recommender systems*. 305–313.
- [18] Wenqi Fan, Shijie Wang, Jiani Huang, Zhikai Chen, Yu Song, Wenzhuo Tang, Haitao Mao, Hui Liu, Xiaorui Liu, Dawei Yin, et al. 2024. Graph Machine Learning in the Era of Large Language Models (LLMs). *arXiv preprint arXiv:2404.14928* (2024).
- [19] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. 2023. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524* (2023).
- [20] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems*. 299–315.
- [21] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 639–648.
- [22] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [23] Wenyue Hua, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2023. How to index item ids for recommendation foundation models. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*. 195–204.
- [24] Masahiro Kaneko and Naoaki Okazaki. 2023. Reducing Sequence Length by Predicting Edit Spans with Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 10017–10029.
- [25] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [26] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.
- [27] Jiatong Li, Yunqing Liu, Wenqi Fan, Xiao-Yong Wei, Hui Liu, Jiliang Tang, and Qing Li. 2024. Empowering molecule discovery for molecule-caption translation with large language models: A chatgpt perspective. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [28] Lei Li, Yongfeng Zhang, and Li Chen. 2023. Prompt distillation for efficient LLM-based recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 1348–1357.
- [29] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2023. LLaRA: Aligning Large Language Models with Sequential Recommenders. *arXiv preprint arXiv:2312.02445* (2023).
- [30] Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics* 12 (2024), 157–173.
- [31] Tie-Yan Liu et al. 2009. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval* 3, 3 (2009), 225–331.
- [32] Guy Lorberbom, Andreea Gane, Tommi Jaakkola, and Tamir Hazan. 2019. Direct Optimization through argmax for Discrete Variational Auto-Encoder. *Advances in neural information processing systems* 32 (2019).
- [33] Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Proceeding of the 7th International Conference on Learning Representations*.
- [34] Brian McFee and Gert Lanckriet. 2010. Metric learning to rank. (2010).
- [35] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research* 21, 140 (2020), 1–67.
- [36] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. Zero-shot text-to-image generation. In *International Conference on Machine Learning*. PMLR, 8821–8831.
- [37] Steffen Rendle, C. Freudenthaler, Zeno Gantner, and L. Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *UAI*.
- [38] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [39] Sho Takase and Sosuke Kobayashi. 2020. All word embeddings from one embedding. *Advances in Neural Information Processing Systems* 33 (2020), 3775–3785.
- [40] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [41] Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems* 30 (2017).
- [42] Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. 2024. Can language models solve graph problems in natural language? *Advances in Neural Information Processing Systems* 36 (2024).
- [43] Hanbing Wang, Xiaorui Liu, Wenqi Fan, Xiangyu Zhao, Venkataramana Kini, Devendra Yadav, Fei Wang, Zhen Wen, Jiliang Tang, and Hui Liu. 2024. Rethinking Large Language Model Architectures for Sequential Recommendations. *arXiv preprint arXiv:2402.09543* (2024).
- [44] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*. 165–174.
- [45] Jiahao Zhang, Rui Xue, Wenqi Fan, Xin Xu, Qing Li, Jian Pei, and Xiaorui Liu. 2024. Linear-Time Graph Neural Networks for Scalable Recommendations. In *Proceedings of the ACM on Web Conference 2024*. 3533–3544.
- [46] Yang Zhang, Fuli Feng, Jizhi Zhang, Keqin Bao, Qifan Wang, and Xiangnan He. 2023. Collm: Integrating collaborative embeddings into large language models for recommendation. *arXiv preprint arXiv:2310.19488* (2023).
- [47] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).
- [48] Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, et al. 2024. Recommender systems in the era of large language models (llms). *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [49] Bowen Zheng, Yupeng Hou, Hongyu Lu, Yu Chen, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Adapting large language models by integrating collaborative semantics for recommendation. *arXiv preprint arXiv:2311.09049* (2023).
- [50] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 1893–1902.