

# Enhancing High-order Interaction Awareness in LLM-based Recommender Model

Xinfeng Wang<sup>†</sup>, Jin Cui<sup>†</sup>, Fumiyo Fukumoto<sup>‡</sup>, and Yoshimi Suzuki<sup>‡</sup>

<sup>†</sup>Graduate School of Engineering

<sup>‡</sup>Interdisciplinary Graduate School

University of Yamanashi, Kofu, Japan

{g22dtsa7, g22dtsa5, fukumoto, ysuzuki}@yamanashi.ac.jp

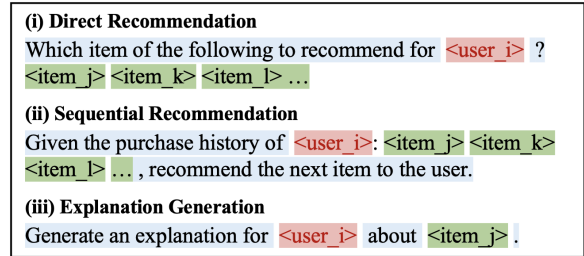
## Abstract

Large language models (LLMs) have demonstrated prominent reasoning capabilities in recommendation tasks by transforming them into text-generation tasks. However, existing approaches either disregard or ineffectively model the user–item high-order interactions. To this end, this paper presents an enhanced LLM-based recommender (ELMRec). We enhance whole-word embeddings to substantially enhance LLMs’ interpretation of graph-constructed interactions for recommendations, without requiring graph pre-training. This finding may inspire endeavors to incorporate rich knowledge graphs into LLM-based recommenders via whole-word embedding. We also found that LLMs often recommend items based on users’ earlier interactions rather than recent ones, and present a reranking solution. Our ELMRec outperforms state-of-the-art (SOTA) methods in both direct and sequential recommendations. Our code is available online<sup>1</sup>.


## 1 Introduction

Large language models (LLMs) with powerful reasoning capabilities have attracted considerable attention in recommendations (Li et al., 2022; Huang et al., 2023; Yang et al., 2023; Wei et al., 2024). They typically exploit prompt learning (Liu et al., 2023) to integrate user and item IDs into LLMs, including discrete prompts, which leverage alternative words (e.g., a movie title) to represent IDs (Li et al., 2023a, 2024), continuous prompts to directly fine-tune ID vectors (Yu et al., 2024), and hybrid prompts (Liao et al., 2023; Wang et al., 2024a). Furthermore, several attempts (Geng et al., 2022; Li et al., 2023b; Wang et al., 2024b) convert multiple recommendation tasks into generation tasks via textual prompts for training LLMs. As illustrated in Fig. 1 (a), these tasks are (i) direct recommendation, (ii) sequential recommendation, and (iii)

<sup>1</sup><https://github.com/WangXFng/ELMRec>



(a) Large language model (LLM)-based recommendation



(b) Graph neural network (GNN)-based recommendation

Figure 1: Illustration of our motivation. In (a), LLM-based recommenders bridge users (pink) and items (green) via text prompts (blue), failing to capture high-order interactive signals. Conversely, GNNs can capture these signals, e.g., 3-hop neighbors (red arrows) in (b).

explanation generation based on user reviews. This enables LLMs to understand diverse perspectives in behavioral semantics for recommendations. However, there still remain two major drawbacks for these LLM-based recommenders. One is their inability to learn the high-order collaborative signal within the user–item interaction graph, even though this signal could greatly assist models which identify the target from millions of candidates. As depicted in Fig. 1 (b), graph neural networks (GNNs), by propagating node embeddings via user–item interaction edges, have been widely utilized to capture these signals for recommendations (He et al., 2020; Zhao et al., 2022; Yu et al., 2022, 2023; Wang et al., 2023). More recently, researchers have exploited LLMs to enhance graph learning for recommendations (Wei et al., 2024; Du et al., 2024; Wang et al., 2024e,f). Nevertheless, integrating the learning process of GNNs into LLMs for recom-

mentations still poses an unresolved challenge.

Another issue with LLM-based recommenders is their struggle with token decomposition, where an ID token (e.g., “user\_1234”) is often split into multiple tokens by LLMs (e.g., “user”, “\_”, “12”, and “34”). To address the issue, [Geng et al. \(2022\)](#) and [Li et al. \(2023b\)](#) propose representing each user or item (e.g., “user\_1234”) with both its ID tokens (e.g., “user”, “\_”, “12”, and “34”) and whole-word embeddings. Each ID’s tokens are covered by one whole-word embedding vector, while all other non-ID tokens share an identical embedding vector. This can emphasize the self-attention correlations among tokens representing the same user or item. However, such whole-word embeddings cannot independently represent users and items, failing to address spurious relatedness among their IDs, e.g., “1234” and “8912” are independent, while they share the same token “12”.

Motivated by the issues, we propose an enhanced LLM-based recommender model (ELMRec). We present novel whole-word embeddings using random feature propagation ([Eliasof et al., 2023](#)), which remarkably enhance high-order interaction awareness of LLMs for recommendations. This is inspired by the fact that GNNs propagate node embeddings to their neighbors, leading them to convey similar features ([Wang et al., 2024c](#)), which can reflect their relationships in the interactive graph. This also mitigates the spurious relatedness between IDs, as each ID is represented by both its tokens and their corresponding whole-word embeddings. Our approach allows LLMs to seamlessly absorb high-order interactive signals for recommendations without requiring graph pre-training.

Furthermore, LLM-based recommenders may prioritize predicting the next items based on earlier interactions rather than recent ones. We attribute this to two factors: (1) during training, they repeatedly sample random segments from interaction sequences to predict the last item in each segment, potentially causing them to overly emphasize past interactions; and (2) the Transformer ([Vaswani et al., 2017](#)) in LLMs excels at capturing long-term dependencies among historical interactions, which potentially leads LLMs to pay less attention to recent ones. Hence, we introduce a simple, effective, and training-free reranking approach.

The main contributions of our work are summarized as follows: (1) we present a novel whole-word embedding approach, which significantly boosts

the ability of LLMs to capture the relative positions of users and items within the interaction graph. This finding may inspire endeavors to incorporate knowledge graphs into LLMs via whole-word embedding for recommendations; (2) we demonstrate that LLMs might recommend items based on users’ past interactions and propose a plug-and-play reranking solution; and (3) extensive experiments demonstrate that EMLRec achieves state-of-the-art (SOTA) performance in both direct and sequential recommendations.

## 2 Related Work

### 2.1 LLM-based Recommendation

LLMs have exhibited remarkable capabilities in item recommendations ([Li et al., 2022](#); [Wei et al., 2024](#); [Huang et al., 2023](#); [Li et al., 2024](#); [Wang et al., 2024a](#)) and explainable recommendations ([Yang et al., 2023](#); [Cheng et al., 2023](#)). Many efforts have leveraged the LLMs’ knowledge to enhance recommendations, addressing zero-/few-shot and cold-start scenarios ([He et al., 2023](#); [Sanner et al., 2023](#)), reranking candidates ([Yue et al., 2023](#); [Carraro and Bridge, 2024](#)), and refining representations ([Harte et al., 2023](#); [Ren et al., 2023](#); [Lin et al., 2023](#); [Lei et al., 2023](#); [Viswanathan et al., 2023](#)).

Recently, [Geng et al. \(2022\)](#) convert multi-modal data, e.g., reviews, user–item interactions, and sequential behaviors, into text-to-text prompts for LLM-based recommenders. [Li et al. \(2023b\)](#) further propose a novel prompt distillation, reducing inference time without compromising performance.

### 2.2 GNN-based Recommendation

Many studies ([He et al., 2020](#); [Wang et al., 2022](#); [Zhao et al., 2022](#); [Lin et al., 2022](#); [Yu et al., 2023](#); [Wang et al., 2024d](#)) have been conducted to capture high-order features from the interaction graph. A surge of research has recently leveraged LLMs to enhance graph learning for recommendation tasks. This includes LLM-based graph structure augmentation ([Wei et al., 2024](#)), enhancing the text attributes of nodes and their neighbors ([Chen et al., 2024](#); [Du et al., 2024](#); [Damianou et al., 2024](#)), aligning graph-structured interactive patterns with textual inference ([Guo et al., 2024](#); [Huang et al., 2024](#)), and leveraging massive textual prompts to generate the interaction edges for LLM-based recommenders ([Wang et al., 2024e](#)). However, they pay no attention to incorporating the learning process of GNNs into LLMs for recommendations.

### 3 Preliminary

#### 3.1 Model Basics

Following prior research (Li et al., 2023b), ELM-Rec utilizes three recommendation tasks and distills their long textual prompts into prompt words (i.e., “P1”, “P2”, and “P3” in Fig. 2). These tasks are (i) direct recommendations, suggesting an item from a given candidate list to a user, (ii) sequential recommendations, forecasting the next item through user sequential interactions, and (iii) explanation generation for user interactions. This work aims to enhance (i) and (ii), while the task (iii) follows the previous work by Li et al. (2023b).

We define the input and output words of each task as  $X = [x_1, \dots, x_{|X|}]$  and  $Y = [y_1, \dots, y_{|Y|}]$ . We concatenate the input tokens with prompt vectors via  $\mathbf{X}_p = [\mathbf{x}_1, \dots, \mathbf{x}_{|X|}, \mathbf{p}_1, \dots, \mathbf{p}_{|P|}]$ . Given the whole-word embedding  $\mathbf{X}_\omega = [\omega_1, \dots, \omega_{|X|+|P|}]$ , we feed  $\hat{\mathbf{X}} = \mathbf{X}_p + \alpha\mathbf{X}_\omega$  into LLMs to obtain a probability distribution  $p(y|Y_{<t}, X)$  over a vocabulary at step  $t$ , where  $Y_{<t}$  denotes the tokens generated before step  $t$ . We adopt a log-likelihood loss function to optimize the model parameters  $\Theta$ :

$$\mathcal{L}_\Theta = \frac{1}{|\mathcal{D}|} \sum_{(X,Y) \in \mathcal{D}} \frac{1}{|Y|} \sum_{t=1}^{|Y|} -\log p(y|Y_{<t}, X), \quad (1)$$

where  $\mathcal{D}$  indicates a set of input-output pairs for three tasks.  $|\mathcal{D}|$  and  $|Y|$  denote the number of pairs, and words in each output, respectively.

#### 3.2 Whole-word Embedding

To comprehend the functioning of whole-word embedding (Geng et al., 2022; Li et al., 2023b), we revisit the self-attention value  $a_{i,j}$  in LLMs with the whole-word embeddings (i.e.,  $\omega_i$  and  $\omega_j$ ):

$$\begin{cases} \mathbf{q}_i = (\mathbf{x}_i + \omega_i)\mathbf{W}_Q, \\ \mathbf{k}_j = (\mathbf{x}_j + \omega_j)\mathbf{W}_K, \\ \mathbf{v}_i = (\mathbf{x}_i + \omega_i)\mathbf{W}_V, \\ a_{i,j} = \text{softmax}(\mathbf{q}_i\mathbf{k}_j^\top). \end{cases} \quad (2)$$

We further unfold  $\mathbf{q}_i\mathbf{k}_j^\top$  as follows:

$$\begin{aligned} \mathbf{q}_i\mathbf{k}_j^\top &= (\mathbf{x}_i + \omega_i)\mathbf{W}_Q\mathbf{W}_K^\top(\mathbf{x}_j + \omega_j)^\top, \\ &= \underbrace{\mathbf{x}_i\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{x}_j^\top}_{\text{Self-attention Correlations}} + \mathbf{x}_i\mathbf{W}_Q\mathbf{W}_K^\top\omega_j^\top \\ &\quad + \omega_i\mathbf{W}_Q\mathbf{W}_K^\top\mathbf{x}_j^\top + \underbrace{\omega_i\mathbf{W}_Q\mathbf{W}_K^\top\omega_j^\top}_{\text{Whole-word Correlations}}. \end{aligned} \quad (3)$$

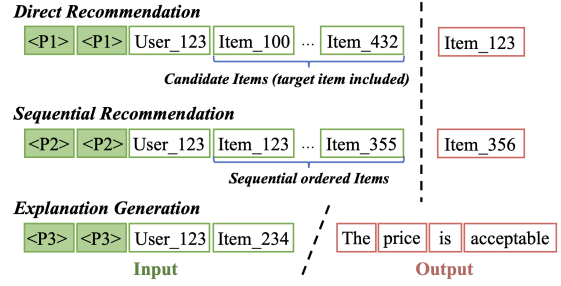


Figure 2: Illustration of input and output words.

The whole-word correlations can improve the self-attention weights among tokens of the same user or item. This is attributed to the fact that  $\omega_i$  and  $\omega_j$  are identical if they are the same ID’s tokens. Therefore, we propose enhancing graph awareness in LLMs by representing user–item relationships with  $\omega_i$  and  $\omega_j$ .

### 4 Approach

We present our ELMRec from both direct and sequential recommendation perspectives.

#### 4.1 Enhancing Direct Recommendation

Motivated by Wang et al. (2024c)’s work, we utilize the random feature propagation (Eliasof et al., 2023) based on LightGCN (He et al., 2020) to generate interaction graph-aware whole-word embeddings. The whole-word embeddings can enhance LLMs’ ability to understand the graph-constructed user–item interactions.

**Random Feature Propagation.** As illustrated in Fig. 3, the random feature propagation via LightGCN makes close nodes comprise similar information, thereby sharing semantic similarity. Specifically, we initiate user and item embeddings by the normal distribution with mean 0 and standard deviation  $\sigma$  as follows:

$$\psi(v) \sim \mathcal{N}(0, \sigma^2), \phi(u) \sim \mathcal{N}(0, \sigma^2), \quad (4)$$

where  $\psi(v)$  and  $\phi(u)$  indicate the embeddings of item  $v$  and user  $u$ , respectively. We then utilize LightGCN to perform the graph convolutions due to its simplification and effectiveness, which is formulated layer by layer as follows:

$$\begin{aligned} \psi(v)^{(l+1)} &= \sum_{u \in \mathcal{N}_v} \frac{1}{\sqrt{|\mathcal{N}_v|}\sqrt{|\mathcal{N}_u|}} \phi(u)^{(l)}, \\ \phi(u)^{(l+1)} &= \sum_{v \in \mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u|}\sqrt{|\mathcal{N}_v|}} \psi(v)^{(l)}, \end{aligned} \quad (5)$$

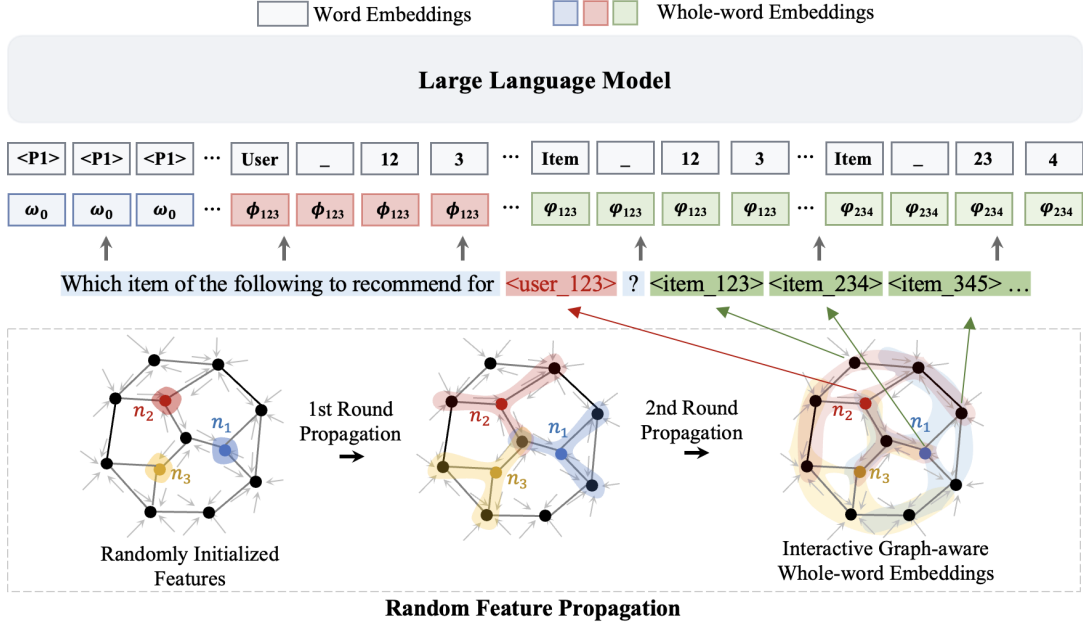


Figure 3: Illustration of integrating interaction graph awareness into LLMs. We first leverage random feature propagation based on LightGCN to obtain whole-word embeddings, which can reflect user–item positions in the interaction graph by their semantic similarity conveyed by red, blue, and yellow edges. We then merge whole-word and word embeddings to enhance LLMs with interaction graph awareness.

where  $l$  refers to the layer of LightGCN, and  $\mathcal{N}_v$  and  $\mathcal{N}_u$  denote the neighbor nodes of item  $v$  and user  $u$ , respectively. By repeatedly applying Eq. (5), we can obtain the coefficient  $c_{u_j \rightarrow u_i}$  representing the influence of  $u_i$ 's 2-order neighbor  $u_j$  on  $u_i$  as follows:

$$c_{u_j \rightarrow u_i} = \frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_v|}} \sum_{k \in \mathcal{N}_{u_i} \cap \mathcal{N}_{u_j}} \frac{1}{|\mathcal{N}_k|}. \quad (6)$$

We can observe that the influence is determined by: (1) the number of their co-interactions, the more the stronger; (2) the popularity of the shared items, the less popular the stronger; and (3) the activity level of user  $u_i$ , the less active the stronger. The stronger influence indicates that more information is propagated to its high-order neighbors, reflecting the high-order interactive correlations between users and items in the graph.

**Interaction Graph-aware Embeddings.** We define the final embeddings after performing random feature propagation as follows:

$$\hat{\psi}(v) = \frac{\sum_{l=0}^L \psi(v)^{(l)}}{L+1}, \hat{\phi}(u) = \frac{\sum_{l=0}^L \phi(u)^{(l)}}{L+1}, \quad (7)$$

where  $L$  indicates the number of LightGCN layers.

We substitute the whole-word embeddings by:

$$\omega_i = \begin{cases} \hat{\psi}(v) & \omega_i \text{ represents item } v, \\ \hat{\phi}(u) & \omega_i \text{ represents user } u, \\ \omega_0 & \text{otherwise,} \end{cases} \quad (8)$$

where  $\omega_0$  refers to the general embedding vector for all non-ID tokens. As such, the enhanced whole-word embedding not only mitigates the spurious relatedness between user and item IDs but also captures high-order collaborative signals among users and items. This is because each ID is represented by both its tokens and whole-word embeddings.

**Remarks.** The majority of tokenizers separate large digital numbers into subwords using Byte-Pair Encoding (Sennrich et al., 2016) or Unigram (Kudo and Richardson, 2018). As a result, only a small number of subwords in the vocabulary are leveraged to represent all digital tokens (e.g., 529<sup>2</sup> out of 32,100 subwords in T5-11b<sup>3</sup>). This leads to a low rank in embedding matrices due to the limited number of digital subwords compared to users and items. In contrast, incorporating the whole-word embedding into the matrices allows a higher rank, further enhancing LLMs' ID representation power (Chen et al., 2021).

<sup>2</sup>There are 326 digital subtokens such as “12” and “34” and 203 continuation subwords such as “\_12” and “\_34”.

<sup>3</sup><https://huggingface.co/google-t5/t5-11b>






Figure 4: Influence of relative positions within the interaction graph on direct and sequential recommendations.

## 4.2 Enhancing Sequential Recommendation

To prevent LLMs from overly prioritizing earlier interactions, we introduce incremental whole-word embeddings and a reranking approach.

**Incremental whole-word embeddings.** Indeed, negative samples for direct recommendations are randomly drawn from the sparse interaction graph, often making only the ground-truth item a close node for the target user. As shown in Fig. 4, by assessing their positions, ELMRec can effectively distinguish the close node from remote ones (i.e., negative samples). An input example of a direct recommendation is given as follows:

$$\langle P1 \rangle \underbrace{user\_123}_{\text{target user}} \quad item\_567 \quad \underbrace{item\_334}_{\text{target item}} \dots item\_1238. \quad (9)$$

The self-attention values between token embeddings of *user\_123* and *item\_334* are enhanced since they are close in the graph and have strong graph-based correlations. In contrast, an input example of a sequential recommendation would be given by:

$$\langle P2 \rangle \underbrace{user\_123}_{\text{target user}} \quad \underbrace{item\_100 \quad item\_104 \dots item\_136}_{\text{tokens whose self-attention is highlighted}}. \quad (10)$$

All these items of  $\{item\_100 \dots item\_136\}$  are the user’s interacted items, which are close to the user in the graph as shown in Fig. 4. As a result, the self-attention values among all these token embeddings are highlighted, potentially exacerbating the model’s tendency to prioritize earlier interactions overly and weakening the sequential patterns needed for LLMs to generate the next item.

We thus employ incremental whole-word embeddings (Geng et al., 2022; Li et al., 2023b), in which we assign increasing indices to users and items based on their appearance order in the input sequence as follows:

$$\langle P2 \rangle \underbrace{user\_123}_{\#0} \quad \underbrace{item\_100}_{\#1} \quad \underbrace{item\_104}_{\#2} \dots \underbrace{item\_136}_{\#N}. \quad (11)$$

Dataset	#User	#Item	#Review	Avg.	Density (%)
Sports	48,993	34,298	296,337	8.3	0.0453
Beauty	22,363	12,101	198,502	8.9	0.0734
Toys	19,804	22,086	167,597	8.6	0.0724

Table 1: Statistics of dataset. “#User”, “#Item”, “#Review”, and “#Avg” denote the number of users, items, reviews, and the average user interactions, respectively.

Subsequently, we utilize the indices to extract the corresponding embeddings from the global whole-word embedding matrix  $\mathbf{X}_{\hat{\omega}}$  ( $\mathbf{X}_{\hat{\omega}} \in \mathbb{R}^{I \times d_n}$ , where  $I$  refers to the maximum incremental index). Therefore, each ID is represented by its literal representation and the appearance order, which not only identifies tokens belonging to the same ID but also indicates their appearance order. The approach potentially emphasizes recent interactions with whole-word embeddings of larger index numbers.

**Reranking approach.** To capture the sequential dependencies among items, recent research (Geng et al., 2022; Li et al., 2023b) leaves the last item in each user’s historical interactions for testing, and then iteratively and randomly selects interaction subsequences to predict the last item within each subsequence for training LLMs. We posit that this training strategy would lead LLMs to overemphasize early subsequences. For instance, when presented with the sequential interactions ( $i_1 \rightarrow i_2 \rightarrow i_3 \rightarrow i_4 \rightarrow i_5$ ) of user  $u_1$ , LLMs are highly likely to recommend items  $i_3$  and  $i_4$ . This is because the subsequences ( $i_1 \rightarrow i_2 \rightarrow i_3$ ) and ( $i_2 \rightarrow i_3 \rightarrow i_4$ ) might be included in the training data caused by this strategy.

To alleviate the issue, we propose a training-free, straightforward, yet effective reranking approach. Specifically, for each prediction, we prompt LLMs to provide  $N$  more items with their ordered probabilities. We then rerank the first  $(k+N)$  candidates to filter out the items that the user has interacted with. We refine the probability distribution  $p(y|Y_{<t}, X)$  at step  $t$  as follows:

$$p_{(:,k)}(y|Y_{<t}, X) \leftarrow \text{rerank}(p_{(:,k+N)}(y|Y_{<t}, X)). \quad (12)$$

## 5 Experiments

### 5.1 Experimental Setup

**Dataset and Metrics.** We performed experiments on three benchmark datasets, i.e., Sports & Outdoors, Beauty, and Toys & Games<sup>4</sup>. Following

<sup>4</sup><https://www.amazon.com/>

Models	Sports				Beauty				Toys			
	H@5	N@5	H@10	N@10	H@5	N@5	H@10	N@10	H@5	N@5	H@10	N@10
Traditional Approach												
SampleX	0.2362	0.1505	0.3290	0.1800	0.2247	0.1441	0.3090	0.1711	0.1958	0.1244	0.2662	0.1469
Large Language Model-based Approach												
P5	0.1955	0.1355	0.2802	0.1627	0.1564	0.1096	0.2300	0.1332	0.1322	0.0889	0.2023	0.1114
RSL	0.2092	0.1502	0.3001	0.1703	0.1564	0.1096	0.2300	0.1332	0.1423	0.0825	0.1926	0.1028
POD	0.2105	0.1539	0.2889	0.1782	0.1931	0.1404	0.2677	0.1639	0.1461	0.1029	0.2119	0.1244
<b>ELMRec</b>	<b>0.5782</b>	<b>0.4792</b>	<b>0.6479</b>	<b>0.4852</b>	<b>0.6052</b>	<b>0.4852</b>	<b>0.6794</b>	<b>0.4973</b>	<b>0.5178</b>	<b>0.4051</b>	<b>0.6045</b>	<b>0.4141</b>
Impv.	174.7%*	211.4%*	124.3%*	172.3%*	213.4%*	245.6%*	153.8%*	203.4%*	254.4%*	293.7%*	185.3%*	232.9%*
Graph Neural Network-based Approach												
LightGCN	0.4150	0.3002	0.5436	0.3418	0.4205	0.3067	0.5383	0.3451	0.3879	0.2874	0.5106	0.3272
NCL	0.4292	0.3131	0.5592	0.3551	0.4378	0.3228	0.5542	0.3607	0.3975	0.2925	0.5120	0.3325
XSimGCL	0.3547	0.2689	0.4486	0.2992	0.3530	0.2734	0.4392	0.3012	0.3351	0.2614	0.4186	0.2885
<b>ELMRec</b>	<b>0.5782</b>	<b>0.4792</b>	<b>0.6479</b>	<b>0.4852</b>	<b>0.6052</b>	<b>0.4852</b>	<b>0.6794</b>	<b>0.4973</b>	<b>0.5178</b>	<b>0.4051</b>	<b>0.6045</b>	<b>0.4141</b>
Impv.	34.7%*	53.1%*	15.9%*	36.6%*	38.2%*	50.3%*	22.6%*	37.9%*	30.3%*	38.5%*	16.0%*	24.5%*

Table 2: Performance comparison on direct recommendation. **Bold**: Best, underline: Second best. “\*” indicates that the improvement is statistically significant ( $p$ -value  $< 0.05$ ) in the 10-trial T-test.

Models	Sports				Beauty				Toys			
	H@5	N@5	H@10	N@10	H@5	N@5	H@10	N@10	H@5	N@5	H@10	N@10
Traditional Approach												
Caser	0.0116	0.0072	0.0194	0.0097	0.0205	0.0131	0.0347	0.0176	0.0166	0.0107	0.0270	0.0141
GRU4Rec	0.0129	0.0086	0.0204	0.0110	0.0164	0.0099	0.0283	0.0137	0.0097	0.0059	0.0176	0.0084
HGN	0.0189	0.0120	0.0313	0.0159	0.0325	0.0206	0.0512	0.0266	0.0321	0.0221	0.0497	0.0277
SASRec	0.0233	0.0154	0.0350	0.0192	0.0387	0.0249	0.0605	0.0318	0.0463	0.0306	0.0675	0.0374
Transformer-based Approach												
BERT4Rec	0.0115	0.0075	0.0191	0.0099	0.0203	0.0124	0.0347	0.0170	0.0116	0.0071	0.0203	0.0099
FDSA	0.0182	0.0122	0.0288	0.0156	0.0267	0.0163	0.0407	0.0208	0.0228	0.0140	0.0381	0.0189
S <sup>3</sup> -Rec	0.0251	0.0161	0.0385	0.0204	0.0387	0.0244	0.0647	0.0327	0.0443	0.0294	0.0700	0.0376
Large Language Model-based Approach												
P5	0.0387	0.0312	0.0460	0.0336	0.0508	0.0379	0.0664	0.0429	0.0648	0.0567	0.0709	0.0587
RSL	0.0392	0.0330	0.0512	0.0375	0.0508	0.0381	0.0667	0.0446	0.0676	0.0583	0.0712	0.0596
POD	0.0497	0.0399	0.0585	0.0422	0.0559	0.0420	0.0696	0.0471	0.0692	0.0589	0.0744	0.0601
<b>ELMRec</b>	<b>0.0538</b>	<b>0.0453</b>	<b>0.0616</b>	<b>0.0471</b>	<b>0.0609</b>	<b>0.0486</b>	<b>0.0750</b>	<b>0.0529</b>	<b>0.0713</b>	<b>0.0608</b>	<b>0.0764</b>	<b>0.0618</b>
Impv.	8.2%*	13.5%*	5.3%*	11.6%*	8.9%*	13.5%*	7.8%*	12.3%*	3.0%	3.2%	2.7%	2.8%

Table 3: Performance comparison between ELMRec and baselines in the sequential recommendation task.

Li et al. (2023b)’s work, for direct and sequential recommendation tasks, we designate the last interaction as the test label, the second-to-last as the validation label, and the remaining interactions as training data. Likewise, for the explanation generation task, we split each dataset into training, validation, and test sets in an 8:1:1 ratio. Table 1 provides the statistics of datasets. We utilized hit rate (HR) $@k$  (H $@k$ ) and normalized discounted cumulative gain (NDCG) $@k$  (N $@k$ ) with  $k \in \{5, 10\}$  as evaluation metrics.

**Baselines.** We compared ELMRec with seven methods for direct recommendations: SimpleX (Mao et al., 2021), P5 (Geng et al., 2022), RLS (Chu et al., 2023), POD (Li et al., 2023b), LightGCN (He et al., 2020), NCL (Lin et al., 2022) and XSimGCL (Yu et al., 2023). We also compared ELMRec with ten methods for sequential recommendations: CASER (Tang and Wang, 2018), HGN (Ma et al., 2019), GRU4Rec (Hidasi et al., 2015), BERT4Rec (Sun et al., 2019), FDSA (Zhang

et al., 2019), SASRec (Kang and McAuley, 2018), S<sup>3</sup>-Rec (Zhou et al., 2020), P5 (Geng et al., 2022), RLS (Chu et al., 2023) and POD (Li et al., 2023b).

Consistent with the LLM-based competitors P5 and POD, ELMRec adopts T5-small (Raffel et al., 2020) as the LLM. Due to space limits, further experimental details are provided in Appendix A.1.

## 5.2 Main Results

Tables 2 and 3 show the comparative results between ELMRec and baselines for direct and sequential recommendations, respectively. We can see that ELMRec is consistently superior to all competitors on all three datasets. Specifically, the improvements compared with the runner-up NCL were 15.9%  $\sim$  53.1% for direct recommendations, and the improvements compared with the runner-up POD were 2.7%  $\sim$  13.5% for sequential recommendations across three datasets. We also have the following observations:

(1) LLM-based competitors notably fall short of GNN-based ones in direct recommendations, high-

Models	Sports		Beauty		Toys	
	H@10	N@10	H@10	N@10	H@10	N@10
Direct recommendation						
w/o Text Prompt	0.1270	0.1006	0.0381	0.0296	0.0190	0.0180
w/o Graph-aware	0.2890	0.1783	0.2687	0.1650	0.2141	0.1243
ELMRec	<b>0.6479</b>	<b>0.4852</b>	<b>0.6794</b>	<b>0.4973</b>	<b>0.6045</b>	<b>0.4141</b>
Impv.	124.2%	172.1%	152.8%	201.4%	182.3%	233.1%
Sequential recommendation						
w/o Reranking	0.0599	0.0435	0.0703	0.0471	0.0737	0.0572
ELMRec	<b>0.0616</b>	<b>0.0471</b>	<b>0.0750</b>	<b>0.0529</b>	<b>0.0764</b>	<b>0.0618</b>
Impv.	2.7%	8.3%	6.7%	7.7%	3.7%	8.0%

Table 4: Ablation study. “w/o Graph-aware” and “w/o Reranking” denote the ELMRec without the interaction graph awareness and reranking approach, respectively. “w/o Text Prompt” indicates that only whole-word embeddings are fed into the LLM for recommendations.

lighting the significance of high-order interactive signals. Our ELMRec, by integrating graph awareness into LLMs, surpasses the leading LLM-based competitor, POD, by 124.3% to 293.7%.

(2) GNN-based approaches, LightGCN, NCL, and XSimGCL, that learn high-order signals are competitive among baselines for direct recommendations. In contrast, ELMRec outperforms them by leveraging richer semantic information through learning text-to-text prompts of sequential behaviors and reviews.

(3) Traditional and Transformer-based methods without pre-trained knowledge are less effective than LLM-based methods in the sequential recommendation task. In contrast, by leveraging LLMs’ knowledge to extract user preferences and item characteristics from multimodal textual prompts, P5, RSL, POD, and ELMRec effectively infer users’ behavioral patterns, enhancing the potential of LLMs for sequential recommendations.

(4) Overall, the sequential recommendation task is more challenging than the direct recommendation task, as evidenced by the lower HR@ $k$  and N@ $k$  values across all methods. This is mainly because the number of candidates in sequential recommendations is much larger than those in direct recommendations. Exploring effective prompts to reduce the candidate range for the former task is an interesting direction for future work.

### 5.3 Ablation Study

To examine the effectiveness of each component in ELMRec, we conducted an ablation study in Tables 4 and 5. The result provides four findings as follows:

(1) The substantial performance drop of ELMRec without interaction graph awareness supports our assumption that the novel whole-word embed-

Models	Sports		Beauty		Toys	
	H@10	N@10	H@10	N@10	H@10	N@10
Direct recommendation						
w/o Graph-aware	0.2890	0.1783	0.2687	0.1650	0.2141	0.1243
Prepending	0.3456	0.2174	0.2730	0.1675	0.2362	0.1309
Addition	<b>0.6479</b>	<b>0.4852</b>	<b>0.6794</b>	<b>0.4973</b>	<b>0.6045</b>	<b>0.4141</b>

Table 5: Results by various methods of incorporating graph-aware whole-word embeddings. “Prepending” refers to prepending directly whole word embedding before the input sequence. “Addition” indicates adding graph-aware embeddings to ID tokens.

Whole-word Embeddings	Sports		Beauty		Toys	
	H@10	N@10	H@10	N@10	H@10	N@10
Direct recommendation						
- Random	0.2758	0.1736	0.2850	0.1752	0.2127	0.1240
- Incremental	0.2989	0.1820	0.2881	0.1753	0.2207	0.1334
- Graph-aware	<b>0.6479</b>	<b>0.4852</b>	<b>0.6794</b>	<b>0.4973</b>	<b>0.6045</b>	<b>0.4141</b>
Sequential recommendation						
- Random	0.0579	0.0413	0.0713	0.0473	0.0741	0.0572
- Incremental	<b>0.0616</b>	<b>0.0471</b>	<b>0.0750</b>	<b>0.0529</b>	<b>0.0764</b>	<b>0.0618</b>
- Graph-aware	0.0556	0.0407	0.0708	0.0481	0.0733	0.0601

Table 6: Effect of various whole-word embeddings. “Graph-aware” represents the interaction graph-aware whole-word embeddings. “Incremental” indicates that the graph-aware whole-word embeddings are replaced with incremental embeddings. The result of “Random” is obtained by disordering the indices for incremental whole-word embeddings.

dings can enhance the LLM’s comprehension of graph-based user-item relationships. This finding suggests that integrating comprehensive knowledge graphs into LLMs using whole-word embeddings could be a promising approach to further improve recommendation performance.

(2) The reranking approach boosts ELMRec’s performance across all datasets, especially in yielding noticeable improvement in NDCG values. This demonstrates that the approach can mitigate the issue of overly prioritizing users’ earlier interactions by optimizing the order of candidates.

(3) Without text prompts, the LLM struggles to effectively understand recommendation tasks. This shows that text prompts are essential for LLM-based recommenders, while graph awareness plays a supplementary role in improving its recommendation performance.

(4) We can see that prepending graph-aware embeddings can slightly improve performance but not as well as adding them with ID tokens. This supports our assumption that the graph-based correlations can enhance the self-attention values of tokens for direct recommendations.




Figure 5: Effect of  $\alpha$ . The x-axis and the y-axis indicate the values of  $\alpha$  and NDCG@10 (%), respectively.

Hyperparameter	Sports		Beauty		Toys	
	H@10	N@10	H@10	N@10	H@10	N@10
$N = 0$	0.0599	0.0435	0.0703	0.0471	0.0737	0.0572
$N = 5$	0.0614	0.0469	0.0744	0.0528	0.0763	0.0617
$N = 10$	<b>0.0616</b>	<b>0.0471</b>	0.0748	0.0528	<b>0.0764</b>	<b>0.0618</b>
$N = 15$	<b>0.0616</b>	<b>0.0471</b>	<b>0.0750</b>	<b>0.0529</b>	<b>0.0764</b>	<b>0.0618</b>

Table 7: Effect of  $N$  for reranking candidates.

## 5.4 In-depth Analyses of ELMRec

We conducted in-depth analyses to better understand ELMRec, including the effect of whole-word embedding, parameter sensitivity, computational complexity analysis, and visualization.

### 5.4.1 Effect of Whole-word Embedding

To show how to choose better whole-word embedding for different tasks, we conducted experiments and reported the result by ELMRec equipped with various whole-word embeddings in Table 6. We have the following insights:

(1) For the direct recommendation, the random indexed and incremental whole-word embeddings make similar contributions. The interaction graph-aware whole-word embeddings significantly improve LLMs’ performance by empowering them to understand the relative positions of users and items in the graph-constructed interactions.

(2) For the sequential recommendation, graph-aware whole-word embeddings slightly harm the performance, as the way it emphasizes all interacted items weakens the sequential patterns. The incremental whole-word embeddings work better than the random indexed ones. This is because the increasing indices can better represent the appearance order of users’ historical interactions.

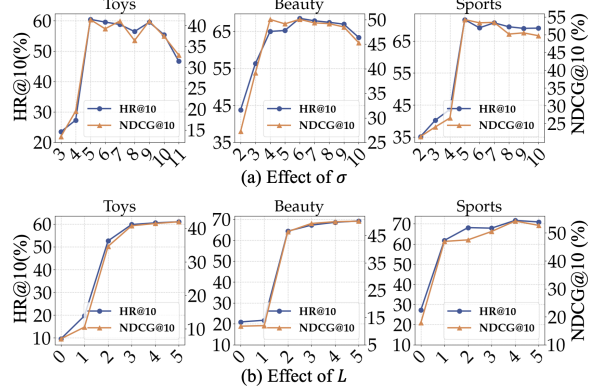


Figure 6: Effect of  $\sigma$  and  $L$  in direct recommendations.

### 5.4.2 Parameter Sensitivity

We investigated the impact of the four hyperparameters, i.e.,  $\alpha$ ,  $N$ ,  $\sigma$ , and  $L$  on ELMRec.

**Effect of  $\alpha$ .** The parameter  $\alpha$  was assigned to vary the influence level of whole-word embeddings. In Fig. 5, the optimal performance of sequential recommendation is distributed at smaller values of  $\alpha$ , whereas the direct recommendation task requires a larger value of  $\alpha$ .

**Effect of  $N$ .**  $N$  was to control the number of additional candidates in reranking approaching. The first five additional candidates, as shown in Table 7, contribute the most to the performance improvement, while increasing  $N$  to 10 or 15 results in minimal or negligible improvements.

**Effect of  $\sigma$ .**  $\sigma$  represents the standard deviation to initiate user and item embeddings. Fig. 6 (a) illustrates that performance improves with increasing  $\sigma$  until it reaches an optimum, after which it declines across all datasets, emphasizing the importance of an appropriate  $\sigma$ .

**Effect of  $L$ .**  $L$  indicates the number of LightGCN layers. In Fig. 6 (b), when the number of layers  $L$  exceeds 4, high-order features have limited or negative effects on performance, while resulting in higher computational costs.

### 5.4.3 Computational Complexity Analysis

The computational complexity of ELMRec hinges on Transformer ( $m^2$ ) and LightGCN ( $n^2$ ), where  $m$  signifies the number of average input tokens and  $n$  denotes the total count of users and items. Given  $m \ll n$ , the complexity of ELMRec in the fixed setting is ( $n^2$ ), underscoring that it is essential to consider the number of users and items in practical deployment. That also prompts us to explore GNNs of less computational complexity for ELMRec, such as SIGN (Frasca et al., 2020), Nodeformer (Wu






Figure 7: Visualization of whole-word embeddings of users and items after 1st and 3rd rounds of propagation. The dots in the same color denote users (or items) who have interacted with the same items (or users). The closer the dots, the greater their similarity. Further visualization results are provided in the Appendix. A.1.3.

Datasets	Stages		
	Pre-training	DirRec	SeqRec
Sports	04h06m58s	17m07s	07m38s
Beauty	02h08m48s	10m32s	04m49s
Toys	03h56m09s	12m44s	04m04s

Table 8: Running time in various stages on three datasets. “DirRec” and “SeqRec” denote the cumulative time cost in direct and sequential recommendations, respectively. “h”, “m”, and “s” indicate “hours” and “minutes”, and “seconds” respectively.

et al., 2022) and HGCN (Wang et al., 2023).

In Table 8, we present the running time by ELMRec in pretraining, direct recommendation, and sequential recommendation, respectively. Despite the substantial cost of the initial pretraining stage, ELMRec shows efficient performance in both direct and sequential recommendation inference.

#### 5.4.4 Visualization

We visualized whole-word embeddings via t-SNE (Van der Maaten and Hinton, 2008). The result in Fig. 7 prompts the observations as stated below:

(1) The whole-word embeddings, learned via random feature propagation within the interaction graph, effectively represent similarity and diversity in collaborative signals among users and items. These embeddings guarantee that similar users and items are closely located in semantic space, so as to enhance the self-attention mechanism and integrate graph awareness into LLMs for recommendations.

(2) Both high- and low-order collaborative sig-

nals reveal user and item similarities. Compared to low-order signals, high-order signals offer three advantages: (i) they better disperse embeddings of active users at the center of distributions on Toys and Beauty, (ii) they reinforce cohesion within clusters of items in the same color on Beauty, and (iii) they facilitate the detection of latent high-order relations among nodes (e.g., items in purple on Toys).

## 6 Conclusion

This paper presents ELMRec, an enhanced LLM-based recommender model. We introduce novel whole-word embeddings to improve the interaction graph awareness of LLMs and a reranking approach to prevent LLMs from prioritizing earlier interactions. Extensive experiments show the effectiveness of ELMRec in both direct and sequential recommendations. Future work involves (i) exploring computationally efficient GNNs (Frasca et al., 2020; Wu et al., 2022) and (ii) investigating more effective and learnable reranking approaches.

## Acknowledgements

We would like to thank anonymous reviewers for their thorough comments and suggestions. This work is supported by the international exchange grant from Tateisi Science and Technology Foundation (No.2242112), the research grant from JKA, and the China Scholarship Council (No.202208330093).

## Limitations

In the fine-tuning stage, the ELMRec model adopts time-consuming modules LightGCN ( $O(n^2)$ ) where  $n$  refers to the number of nodes in the interaction graph, therefore its computational cost heavily relies on the number of users and items in the recommender system. Another limitation is the same as LightGCN in its inability to handle cold start scenarios. It requires retraining the entire model from scratch if a new user or item is introduced, which can be costly.

## Ethics Statement

This paper does not involve the presentation of a new dataset, an NLP application, and the utilization of demographic or identity characteristics information.

## References

- Sheldon Axler. 2015. *Linear algebra done right*. Springer.
- Diego Carraro and Derek Bridge. 2024. Enhancing recommendation diversity by re-ranking with large language models. *arXiv preprint arXiv:2401.11506*.
- Pu-Chin Chen, Henry Tsai, Srinadh Bhojanapalli, Hyung Won Chung, Yin-Wen Chang, and Chun-Sung Ferng. 2021. A simple and effective positional encoding for transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2974–2988.
- 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):42–61.
- Hao Cheng, Shuo Wang, Wensheng Lu, Wei Zhang, Mingyang Zhou, Kezhong Lu, and Hao Liao. 2023. Explainable recommendation with personalized review retrieval and aspect learning. In *the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 51–64.
- Zhixuan Chu, Hongyan Hao, Xin Ouyang, Simeng Wang, Yan Wang, Yue Shen, Jinjie Gu, Qing Cui, Longfei Li, Siqiao Xue, et al. 2023. Leveraging large language models for pre-trained recommender systems. *arXiv preprint arXiv:2308.10837*.
- Andreas Damianou, Francesco Fabbri, Paul Giglioli, Marco De Nadai, Alice Wang, Enrico Palumbo, and Mounia Lalmas. 2024. Towards graph foundation models for personalization. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1798–1802.
- Yingpeng Du, Ziyang Wang, Zhu Sun, Haoyan Chua, Hongzhi Liu, Zhonghai Wu, Yining Ma, Jie Zhang, and Youchen Sun. 2024. Large language model with graph convolution for recommendation. *arXiv preprint arXiv:2402.08859*.
- Moshe Eliasof, Fabrizio Frasca, Beatrice Bevilacqua, Eran Treister, Gal Chechik, and Haggai Maron. 2023. Graph positional encoding via random feature propagation. In *International Conference on Machine Learning*, pages 9202–9223. PMLR.
- Fabrizio Frasca, Emanuele Rossi, Davide Eynard, Ben Chamberlain, Michael Bronstein, and Federico Monti. 2020. Sign: Scalable inception graph neural networks. *arXiv preprint arXiv:2004.11198*.
- 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*, pages 299–315.
- Naicheng Guo, Hongwei Cheng, Qianqiao Liang, Linxun Chen, and Bing Han. 2024. Integrating large language models with graphical session-based recommendation. *arXiv preprint arXiv:2402.16539*.
- Jesse Harte, Wouter Zorgdrager, Panos Louridas, Asterios Katsifodimos, Dietmar Jannach, and Marios Fragkoulis. 2023. Leveraging large language models for sequential recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1096–1102.
- 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*, pages 639–648.
- Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian McAuley. 2023. Large language models as zero-shot conversational recommenders. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 720–730.
- Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.
- Xu Huang, Jianxun Lian, Yuxuan Lei, Jing Yao, Defu Lian, and Xing Xie. 2023. Recommender ai agent: Integrating large language models for interactive recommendations. *arXiv preprint arXiv:2308.16505*.
- Xuanwen Huang, Kaiqiao Han, Yang Yang, Dezheng Bao, Quanjin Tao, Ziwei Chai, and Qi Zhu. 2024. Can gnn be good adapter for llms? *arXiv preprint arXiv:2402.12984*.

- Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, pages 197–206. IEEE.
- Taku Kudo and John Richardson. 2018. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. arXiv 2018. *arXiv preprint arXiv:1808.06226*.
- Yuxuan Lei, Jianxun Lian, Jing Yao, Xu Huang, Defu Lian, and Xing Xie. 2023. Recexplainer: Aligning large language models for recommendation model interpretability. *arXiv preprint arXiv:2311.10947*.
- Jian Li, Jieming Zhu, Qiwei Bi, Guohao Cai, Lifeng Shang, Zhenhua Dong, Xin Jiang, and Qun Liu. 2022. Miner: multi-interest matching network for news recommendation. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 343–352.
- Lei Li, Yongfeng Zhang, and Li Chen. 2023a. Personalized prompt learning for explainable recommendation. *ACM Transactions on Information Systems*, 41(4):1–26.
- Lei Li, Yongfeng Zhang, and Li Chen. 2023b. Prompt distillation for efficient llm-based recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1348–1357.
- Yaoyiran Li, Xiang Zhai, Moustafa Alzantot, Keyi Yu, Ivan Vulić, Anna Korhonen, and Mohamed Hammad. 2024. Calrec: Contrastive alignment of generative llms for sequential recommendation. *arXiv preprint arXiv:2405.02429*.
- 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*.
- Xinyu Lin, Wenjie Wang, Yongqi Li, Fuli Feng, See-Kiong Ng, and Tat-Seng Chua. 2023. A multi-facet paradigm to bridge large language model and recommendation. *arXiv preprint arXiv:2310.06491*.
- Zihan Lin, Changxin Tian, Yupeng Hou, and Wayne Xin Zhao. 2022. Improving graph collaborative filtering with neighborhood-enriched contrastive learning. In *Proceedings of the ACM Web Conference 2022*, pages 2320–2329.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Chen Ma, Peng Kang, and Xue Liu. 2019. Hierarchical gating networks for sequential recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 825–833.
- Kelong Mao, Jieming Zhu, Jinpeng Wang, Quanyu Dai, Zhenhua Dong, Xi Xiao, and Xiuqiang He. 2021. Simplex: A simple and strong baseline for collaborative filtering. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1243–1252.
- 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. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2023. Representation learning with large language models for recommendation. *arXiv preprint arXiv:2310.15950*.
- Scott Sanner, Krisztian Balog, Filip Radlinski, Ben Wedin, and Lucas Dixon. 2023. Large language models are competitive near cold-start recommenders for language-and item-based preferences. In *Proceedings of the 17th ACM conference on recommender systems*, pages 890–896.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.
- Gilbert Strang. 2022. *Introduction to linear algebra*. SIAM.
- 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*, pages 1441–1450.
- Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 565–573.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- 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.

- Vijay Viswanathan, Luyu Gao, Tongshuang Wu, Pengfei Liu, and Graham Neubig. 2023. Datafinder: Scientific dataset recommendation from natural language descriptions. In *the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10288–10303.
- Chenyang Wang, Yuanqing Yu, Weizhi Ma, Min Zhang, Chong Chen, Yiqun Liu, and Shaoping Ma. 2022. Towards representation alignment and uniformity in collaborative filtering. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1816–1825.
- Tianhao Wang, Sheng Wu, Fen Yi, Lidan Kuang, You Wang, and Jin Zhang. 2024a. Hybrid prompt recommendation explanation generation combined with graph encoder. *Neural Processing Letters*, 56(1):32.
- Xinfeng Wang, Jin Cui, Yoshimi Suzuki, and Fumiyo Fukumoto. 2024b. Rdrec: Rationale distillation for llm-based recommendation. *arXiv preprint arXiv:2405.10587*.
- Xinfeng Wang, Fumiyo Fukumoto, Jin Cui, Yoshimi Suzuki, Jiyi Li, and Dongjin Yu. 2023. Eedn: Enhanced encoder-decoder network with local and global context learning for poi recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 383–392.
- Xinfeng Wang, Fumiyo Fukumoto, Jin Cui, Yoshimi Suzuki, Jiyi Li, and Dongjin Yu. 2024c. Cadrec: Contextualized and debiased recommender model. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 405–415.
- Xinfeng Wang, Fumiyo Fukumoto, Jin Cui, Yoshimi Suzuki, and Dongjin Yu. 2024d. Nfarec: A negative feedback-aware recommender model. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 935–945.
- Xinyuan Wang, Liang Wu, Liangjie Hong, Hao Liu, and Yanjie Fu. 2024e. Llm-enhanced user-item interactions: Leveraging edge information for optimized recommendations. *arXiv preprint arXiv:2402.09617*.
- Yan Wang, Zhixuan Chu, Xin Ouyang, Simeng Wang, Hongyan Hao, Yue Shen, Jinjie Gu, Siqiao Xue, James Zhang, Qing Cui, et al. 2024f. Llmrg: Improving recommendations through large language model reasoning graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19189–19196.
- Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. Llmrec: Large language models with graph augmentation for recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 806–815.
- Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. 2022. Nodeformer: A scalable graph structure learning transformer for node classification. *Advances in Neural Information Processing Systems*, 35:27387–27401.
- Zhengyi Yang, Jiancan Wu, Yanchen Luo, Jizhi Zhang, Yancheng Yuan, An Zhang, Xiang Wang, and Xiangnan He. 2023. Large language model can interpret latent space of sequential recommender. *arXiv preprint arXiv:2310.20487*.
- Junliang Yu, Xin Xia, Tong Chen, Lizhen Cui, Nguyen Quoc Viet Hung, and Hongzhi Yin. 2023. Xsimgl: Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on Knowledge and Data Engineering*.
- Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are graph augmentations necessary? simple graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*, pages 1294–1303.
- Xiaohan Yu, Li Zhang, Xin Zhao, Yue Wang, and Zhongrui Ma. 2024. Ra-rec: An efficient id representation alignment framework for llm-based recommendation. *arXiv preprint arXiv:2402.04527*.
- Zhenrui Yue, Sara Rabhi, Gabriel de Souza Pereira Moreira, Dong Wang, and Even Oldridge. 2023. Llamarec: Two-stage recommendation using large language models for ranking. *arXiv preprint arXiv:2311.02089*.
- Tingting Zhang, Pengpeng Zhao, Yanchi Liu, Victor S Sheng, Jiajie Xu, Deqing Wang, Guanfeng Liu, Xiaofang Zhou, et al. 2019. Feature-level deeper self-attention network for sequential recommendation. In *IJCAI*, pages 4320–4326.
- Sen Zhao, Wei Wei, Ding Zou, and Xianling Mao. 2022. Multi-view intent disentangle graph networks for bundle recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 4379–4387.
- 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*, pages 1893–1902.



## A Appendix

### A.1 Experimental Details

This section provides further experimental details, including implementation setup, visualization results, computational complexity analysis, running time, and model optimization and inference.

#### A.1.1 Implementation Details

For a fair comparison, the hyperparameters of ELMRec were set to the same as POD. Both the encoder and decoder consist of 6 layers with 8-headed attention in each layer. The number of negative items for direct recommendation is set to 99 for both training and evaluation. We set the number of prompt words to 3 for all tasks, and the batch size for training all three tasks to 64. Consistent with the LLM-based competitors P5 and POD, ELMRec adopts T5-small (Raffel et al., 2020) as the LLM for comparison.

For  $\alpha$ ,  $\sigma$ ,  $N$ , and  $L$ , which are not included in POD, we tuned them to attain the best performance. The search range for  $\alpha$ ,  $\sigma$ ,  $N$ , and  $L$  were  $1 \sim 10$ ,  $1 \sim 10$ ,  $\{5, 10, 15\}$  and  $\{1, 2, 3, 4, 5\}$ , respectively. Table 9 displayed the result. Specifically, in the direct recommendation task,  $\alpha$  and  $\sigma$  were 5 and 5 for Sports, 9 and 6 for Beauty, and 9 and 5 for Toys, while in the sequential recommendation task, they were 11 and 5 for Sports, 6 and 6 for Beauty, and 1 and 5 for Toys, respectively.  $N$  and  $L$  were 15 and 4 for Beauty, and 10 and 4 for others. The embedding size of word tokens and whole-word tokens was 512. We employed a learning rate of 0.01 for all datasets.

We iteratively and randomly sample a segment from a user’s item sequence for training the sequential recommendation task. We exploit the discrete prompt templates for different tasks from (Geng et al., 2022). We utilized the AdamW optimizer (Loshchilov and Hutter, 2017). We implemented and evaluated our ELMRec using PyTorch on an Nvidia GeForce RTX 4090 (24GB memories).

#### A.1.2 Baselines

To evaluate the performance of sequential and direct recommendations, we compared our ELMRec with the following fourteen baselines which are divided into four groups:

##### - Traditional Approach

- **GRU4Rec** (Hidasi et al., 2015) regards the entire item sequence of each user as the user’s session and employs GRUs to recommend.

Dataset	DirRec				SeqRec			
	$\alpha$	$\sigma$	$N$	$L$	$\alpha$	$\sigma$	$N$	$L$
Sports	5	5	10	4	11	5	10	4
Beauty	9	6	15	4	6	6	15	4
Toys	9	5	10	4	1	5	10	4

Table 9: Best values of hyperparameter for the three datasets. “DirRec” and “SeqRec” denote direct and sequential recommendations, respectively.

- **CASER** (Tang and Wang, 2018) treats user interactions as images and employs 2-dimensional convolutions to capture sequential patterns.
- **SASRec** (Kang and McAuley, 2018) exploits Markov Chains to excavate short-term semantics in users’ sequential patterns.
- **HGN** (Ma et al., 2019) exploits a novel gating strategy to model users’ long- and short-term interests in candidate items.
- **SimpleX** (Chen et al., 2021) incorporates the cosine contrastive loss into simple unified collaborative filtering to deliver recommendations.
- *Transformer-based Approach*
  - **BERT4Rec** (Sun et al., 2019) proposes to leverage the BERT-style cloze task for the sequential recommender algorithm.
  - **FDSA** (Zhang et al., 2019) incorporates item features with item sequences of users to perform recommendations.
  - **S<sup>3</sup>-Rec** (Zhou et al., 2020) learns users’ latent behavioral features via employing a self-supervised learning paradigm.
- *Graph Neural Network (GNN)-based Approach*
  - **LightGCN** (He et al., 2020) simplifies the design of GCN to make it more suitable for recommendation purposes.
  - **NCL** (Lin et al., 2022) regards users (or items) and their structural neighbors as contrastive pairs to augment the user and item embeddings.
  - **XSimGCL** (Yu et al., 2023) creates contrastive views by adding uniform noise to enhance the robustness of user and item representations.
- *Large Language Model (LLM)-based Approach*
  - **P5** (Geng et al., 2022) converts three different recommendation tasks into textual generation tasks using LLMs for recommendations.
  - **RSL** (Chu et al., 2023) adopts novel training and inference strategies to deliver LLM-based recommendations.
  - **POD** (Li et al., 2023b) refines P5 through prompt distillation to make efficient and precise recommendations.

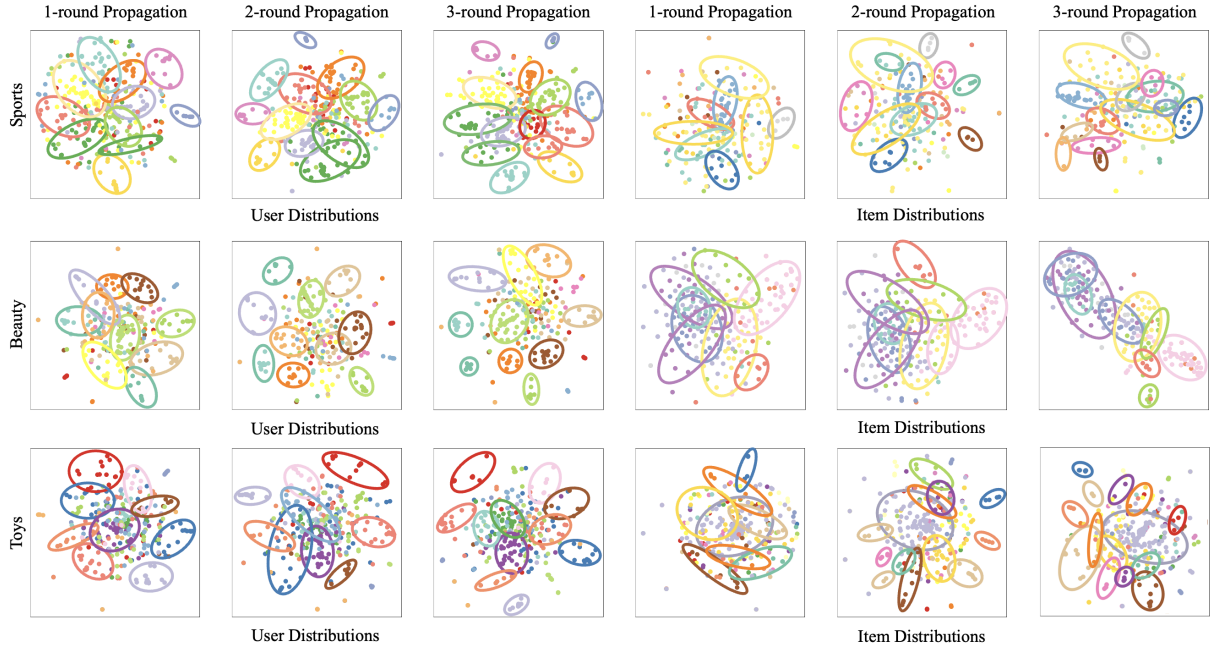


Figure 8: Visualization of user and item whole-word embeddings at each round of random feature propagation. The dots in the same color denote users who have interacted with the same items or the items that are interacted with by the same user. Similar users and items are close to each other in the distribution.

### A.1.3 Visualization Details

Fig. 8 shows the visualization result of user and item whole-word embeddings at each round of random feature propagation. We found that the random feature propagation based on LightGCN makes the embeddings of similar users and items closer in the semantic space for all three datasets. This explains why the whole-word embeddings can enhance the self-attention mechanism and integrate graph awareness into LLMs for recommendations. We can also observe that both high- and low-order collaborative signals can reveal which users and items are similar. Also, high-order signals can better disperse the embeddings of active nodes (e.g., user distributions in the Toys and Beauty datasets), reinforce cohesion within clusters (e.g., item distribution in the Beauty dataset), and detect latent high-order relations (e.g., some dots of the same color begin to cluster in the high-order signals).

It is noteworthy that we only feed the whole-word embeddings and their colors into the t-SNE tool to generate visualizations. The t-SNE tool maps high-dimensional data into two or three dimensions and keeps similar data points close.

### A.1.4 Model Optimization and Inference

The input for the explanation task comprises solely two IDs (i.e., user and item), whereas for direct and sequential recommendation tasks, it could contain

a hundred or more IDs due to the negative samples and the user’s historical item sequence. The diverse input and output lengths across tasks hinder the efficiency of training the model with mixed samples from various tasks (Geng et al., 2022). Therefore, to avoid unnecessary padding and extra computing costs, we train our ELMRec model alternately, using a batch of samples from one task and then switching to another task.

During training, we shuffle the input-output pairs of the three tasks and repeatedly select samples from each task in a specific batch size. We save a checkpoint if the total validation loss of the model in all tasks is the lowest for the current epoch. If this doesn’t occur 5 times, we terminate the training process and load the best checkpoint for evaluation. During inference, we employ a beam search algorithm to generate results by selecting the candidate sequences with the highest likelihood over the vocabulary (Li et al., 2023b). We set the number of beams at 20 for sequential and direct recommendations. For generation tasks, we apply group beam search with the number of beams and beam groups set to 21 and 3, respectively.

### A.2 Reranking Approach

To capture sequential dependencies among items, Geng et al. (2022) and Li et al. (2023b) use the last item in each user’s interaction history for testing,

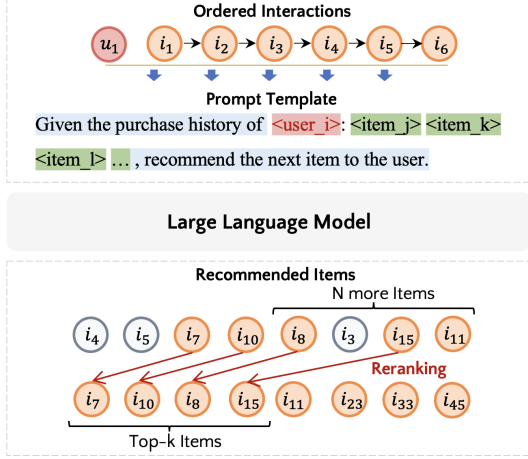


Figure 9: Illustration of reranking approach. The gray nodes such as  $i_4$  and  $i_5$  indicate the user’s interacted items. The red arrows refer to the reranking processes.

and randomly select subsequences from users’ histories to predict the last item in each subsequence for training LLMs to make generative recommendations. However, this strategy often causes LLMs to emphasize fragmentary early subsequences overly. For example, in Fig. 9, given user  $u_1$ ’s interactions ( $i_1 \rightarrow i_2 \rightarrow i_3 \rightarrow i_4 \rightarrow i_5 \rightarrow i_6$ ), the LLM is likely to recommend items  $i_4$  and  $i_5$  because subsequences ( $i_1 \rightarrow i_2 \rightarrow i_3 \rightarrow i_4$ ) and ( $i_1 \rightarrow i_2 \rightarrow i_3 \rightarrow i_4 \rightarrow i_5$ ) are included in the training data.

To address this issue, we propose a straightforward, training-free reranking approach. As shown in Fig. 9, for each prediction, we prompt the LLM to provide  $N$  additional items with their ordered probabilities (where  $N$  is a hyperparameter). We then filter out items the user has already interacted with and rerank the top  $(k+N)$  candidates for the sequential recommendation.

### A.3 Matrix Form of Random Feature Propagation

We present the matrix form of random feature propagation, which enables straightforward implementation. Specifically, we define the user–item interaction matrix as  $\mathbf{R} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{V}|}$ , where  $|\mathcal{U}|$  and  $|\mathcal{V}|$  denote the number of users and items, respectively. Each entry  $\mathbf{R}_{u,v}$  is 1 if user  $u$  has interacted with item  $v$  otherwise 0. We then obtain the adjacency matrix of the user–item graph  $\mathbf{A} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{V}|) \times (|\mathcal{U}|+|\mathcal{V}|)}$  as follows:

$$\mathbf{A} = \begin{bmatrix} 0 & \mathbf{R} \\ \mathbf{R}^\top & 0 \end{bmatrix}. \quad (13)$$

Consider the initial user–item embedding matrix to be the 0-th layer embedding matrix,  $\mathbf{E}^{(0)} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{V}|) \times d_n}$ , where user and item embeddings are randomly generated via the normal distribution. Subsequently, the matrix equivalent form of information propagation is given by:

$$\mathbf{E}^{(l+1)} = (\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}) \mathbf{E}^{(l)}, \quad (14)$$

where  $\mathbf{D} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{V}|) \times (|\mathcal{U}|+|\mathcal{V}|)}$  is a diagonal matrix in which each entry  $d_{ij}$  denotes the number of nonzero entries in the  $i$ -th row vector of matrix  $\mathbf{A}$ . The final graph-aware whole-embeddings matrix  $\Omega \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{V}|) \times d_n}$  is obtained by:

$$\Omega = \frac{1}{L+1} \sum_{l=0}^L \tilde{\mathbf{A}} \mathbf{E}^{(l)}, \quad (15)$$

where  $\tilde{\mathbf{A}} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  is the symmetrically normalized adjacency matrix.

### A.4 Mitigation of Low Rank Issue

Since tokenizers separate large digital numbers into subwords (Sennrich et al., 2016; Kudo and Richardson, 2018), only a small number of subwords in the vocabulary are leveraged to represent all digital tokens. We argue that when considering all samples, the whole-word embedding improves the ID representation power in terms of the rank of attention matrices for T5-base and T5-11B.

**Theorem 1.** *Let  $d_x$  be the number of digital subwords in vocabulary, and  $d_p$  be the number of users and items. Consider input ID embeddings  $X$  and their whole-word embeddings  $P \in \mathbb{R}^{n \times d_n}$ , where  $n$  is the number of samples. Let  $\mathbf{W}_Q$  and  $\mathbf{W}_K \in \mathbb{R}^{d_n \times d_h}$  be weight matrices for queries and keys, respectively, where  $(n > d_p > d_h = d_n > d_x)$ <sup>5</sup>. Define  $\mathbf{A}_x = X \mathbf{W}_Q \mathbf{W}_K^\top X^\top$  and  $\mathbf{A}_{x+p} = (X + P) \mathbf{W}_Q \mathbf{W}_K^\top (X + P)^\top$  as attention matrices with and without whole-word embeddings, respectively. Then, for any  $\mathbf{W}_Q$  and  $\mathbf{W}_K$ , we have:*

$$\begin{aligned} \text{rank}(\mathbf{A}_x) &= \min\{d_x, d_n\} = d_x, \\ \text{rank}(\mathbf{A}_{x+p}) &= \min\{d_x + d_n, d_n\} = d_n > d_x. \end{aligned} \quad (16)$$

Let  $d_x$  represent the number of digital subwords within the vocabulary, and let  $d_p$  denote the total count of users and items. Consider the input ID embeddings  $X \in \mathbb{R}^{n \times d_n}$  and their corresponding

<sup>5</sup>In T5 models,  $d_n = d_h = 768$  for T5-base and 1024 for T5-11B, and  $d_x = 529$ .

whole-word embeddings  $P \in \mathbb{R}^{n \times d_n}$ , where  $n$  denotes the total number of samples.

Notably,  $X$  denotes pre-trained embeddings extracted from the vocabulary, while  $P$  is randomly generated from a normal distribution. Due to their distinct origins, the matrices  $X$  and  $P$  are linearly independent (Axler, 2015; Strang, 2022), i.e.,  $X \cap P = \emptyset$ . We thus have:

$$\begin{aligned} & \text{rank}(X + P) \\ &= \text{rank}(X) + \text{rank}(P) - \text{rank}(X \cap P) \\ &= \min(d_x, d_n) + \min(d_p, d_n) - 0 \\ &= d_x + d_n. \end{aligned}$$

Define  $\mathbf{W}_Q \in \mathbb{R}^{d_n \times d_h}$  and  $\mathbf{W}_K \in \mathbb{R}^{d_n \times d_h}$  as weight matrices for queries and keys, respectively, where  $n > d_p > d_h = d_n > d_x$ . Define  $\mathbf{A}_x = X\mathbf{W}_Q\mathbf{W}_K^\top X^\top$  and  $\mathbf{A}_{x+p} = (X + P)\mathbf{W}_Q\mathbf{W}_K^\top (X + P)^\top$  as the attention matrices with and without whole-word embeddings, respectively. Since the rank of the product of two matrices is upper bounded by the minimum of their ranks, we derive:

$$\begin{aligned} & \text{rank}(\mathbf{A}_x) \\ &= \text{rank}(X\mathbf{W}_Q\mathbf{W}_K^\top X^\top) \\ &= \min(\text{rank}(X), \text{rank}(\mathbf{W}_Q), \text{rank}(\mathbf{W}_K)) \\ &= \min(d_x, d_n) \\ &= d_x. \end{aligned}$$

$$\begin{aligned} & \text{rank}(\mathbf{A}_{x+p}) \\ &= \text{rank}((X + P)\mathbf{W}_Q\mathbf{W}_K^\top (X + P)^\top) \\ &= \min(\text{rank}(X + P), \text{rank}(\mathbf{W}_Q), \text{rank}(\mathbf{W}_K)) \\ &= \min(d_x + d_n, d_n) \\ &= d_n > d_x. \end{aligned}$$

The above derivation shows that, when considering all samples, whole-word embeddings increase the lower bound of the rank of the attention matrix from  $d_x$  to  $d_n$ . This can enhance the model's ID representation power in terms of the rank of the attention matrix (Chen et al., 2021).