Rethinking Large Language Model Architectures for Sequential Recommendations

Hanbing Wang[1] , Xiaorui Liu[3] , Wenqi Fan[4] , Xiangyu Zhao[5] , Venkataramana Kini[2]
Devendra Yadav[2] , Fei Wang[2] , Zhen Wen[2] , Jiliang Tang[1] , Hui Liu[1]
[1]Michigan State University, [2]Amazon, [3]North Carolina State University, USA
[4]The Hong Kong Polytechnic University, [5]City University of Hong Kong, China
{wangh137,liuhui7,tangjili}@msu.edu,{venkini,feiww,yaddevn,zhenwen}@amazon.com
liu96@ncsu.edu,wenqifan03@gmail.com,xianzhao@cityu.edu.hk

ABSTRACT

Recently, sequential recommendation has been adapted to the LLM paradigm to enjoy the power of LLMs. LLM-based methods usually formulate recommendation information into natural language and the model is trained to predict the next item in an auto-regressive manner. Despite their notable success, the substantial computational overhead of inference poses a significant obstacle to their real-world applicability. In this work, we endeavor to streamline existing LLM-based recommendation models and propose a simple yet highly effective model Lite-LLM4Rec. The primary goal of Lite-LLM4Rec is to achieve efficient inference for the sequential recommendation task. Lite-LLM4Rec circumvents the beam search decoding by using a straight item projection head for ranking scores generation. This design stems from our empirical observation that beam search decoding is ultimately unnecessary for sequential recommendations. Additionally, Lite-LLM4Rec introduces a hierarchical LLM structure tailored to efficiently handle the extensive contextual information associated with items, thereby reducing computational overhead while enjoying the capabilities of LLMs. Experiments on three publicly available datasets corroborate the effectiveness of Lite-LLM4Rec in both performance and inference efficiency (notably 46.8% performance improvement and 97.28% efficiency improvement on ML-1m) over existing LLM-based methods. Our implementations will be open sourced.

CCS CONCEPTS

• **Information systems → Recommender systems**.

KEYWORDS

Recommender Systems, Large Language Models, Sequential Recommendation

1 INTRODUCTION

Sequential recommendation is to predict next item a user will interact with based on his/her interaction history. Because user interests are dynamic and evolving over time, it is important to capture the sequential pattern, leading to accurate recommendations. Traditional methods model the item transition patterns based on Markov Chain [10, 11, 34]. With the development of deep learning, a variety of deep neural networks, such as Transformer [18, 35], RNN [12, 40] and CNN [37], have been proposed to advance the task, achieving remarkable performance. Furthermore, side information (e.g., attributes, titles) has been incorporated [14, 17, 47], which helps achieve remarkable improvement, demonstrating its importance and potential.

Recently, the widespread success of large language models (LLMs), such as GPT [3], T5 [31], and Llama [39], has demonstrated their exceptional ability of contextual understanding and offers a promising direction to improve recommendation systems with heightened personalization and adaptability. Existing LLM-based recommendation algorithms [7, 21, 46] mainly adapt recommendation tasks to the LLM paradigm by formulating relevant information, i.e., interaction information, meta data, or candidate items through various indexing strategies into natural language. As shown in Figure 1, such information will be wrapped in a prompt template, and fed into LLMs as input. Subsequently, the input will be transformed into an informative latent representation and the model will auto-regressively generate the recommendations in natural language through decoding with beam search [42].

Although existing LLM-based methods have achieved remarkable success, the exorbitant cost of inference hinders their real-world applications and poses a formidable obstacle to seamless, real-time user experiences [27, 32, 45]. To gain deeper insights into this efficiency challenge, we perform a preliminary study in Section 3.2 on the impact of the key components of LLM-based methods. First, we observe that the beam search decoding is the most time-consuming component. This component is employed to generate k recommendations for a user, leading to k times greater model computation complexity. Second, we find that inference time diminishes with reduced input length. Longer input often leads to more model computations [21]. Furthermore, we reveal that existing item indexing will lead to serious redundant computation problem because 1) the item indexing will be tokenized into several

tokens before being fed into the LLM, which will result in longer input; and 2) some items may appear frequently in the data and the model will compute on the item token string every time when it appears. Our further study of the impact of beam search and item indexing on model performance in Section 3.3 suggests that the beam search decoding and item indexing are burdensome and unnecessary for sequential recommendation.

Grounded on our preliminary studies, we aim to streamline the model architecture of existing LLM-based recommendations, introducing Lite-LLM4Rec as an efficient solution for sequential recommendation. We find that the beam search decoding is unnecessary and resource-intensive for this specific task. Therefore, Lite-LLM4Rec circumvents the beam search decoding process. Though representing items with their context information can achieve superior performance, it makes the input very long and correspondingly increases computational costs. To tackle this challenge, we propose a hierarchical LLM structure with two key components - Item LLM and Recommendation LLM - to process extensive context information efficiently. We demonstrate that formulating relevant information such as interaction information, meta information into natural language is superfluous because LLMs can actually understand the semantic information encoded in the latent representation. This design can reduce the input length and alleviate redundant computation while enjoying the power of LLMs. Experimental results indicate that Lite-LLM4Rec can significantly improve not only the inference efficiency but also the overall performance. The main contributions of this paper are summarized as follows:

- We simplify existing LLM-based recommendation models and propose a simple but effective sequential recommendation model Lite-LLM4Rec. Through eliminating the decoder, Lite-LLM4Rec can achieve better performance as well as significant inference efficiency.
- We design a novel hierarchical LLM structure to efficiently process long context information in LLM-based recommendations, which can reduce the computational demands while enjoying the power of LLMs.
- We demonstrate the effectiveness of Lite-LLM4Rec on various benchmark datasets.

2 RELATED WORK

In this section, we summarize relevant works on sequential recommendation and LLM-based solutions.

2.1 Sequential Recommendation

Sequential recommendations [18, 35] leverage user historical interactions to infer the next item the user will interact with. Since user interests are dynamic and evolving over time, it is important to capture this sequential pattern and provide appropriate recommendations. Early studies mainly depend on Markov Chain to model item transition patterns [11, 34]. Recently, deep-learning based methods have dominated the area. For example, GRU4rec [13] proposes to use RNN in session-based recommendation. Some approaches introduce CNN into sequential recommendation [37, 41]. SASrec and BERT4rec [18, 35] employ the self-attention mechanism into sequential recommendation and achieve excellent performance. However, these methods neglect rich contextual information about
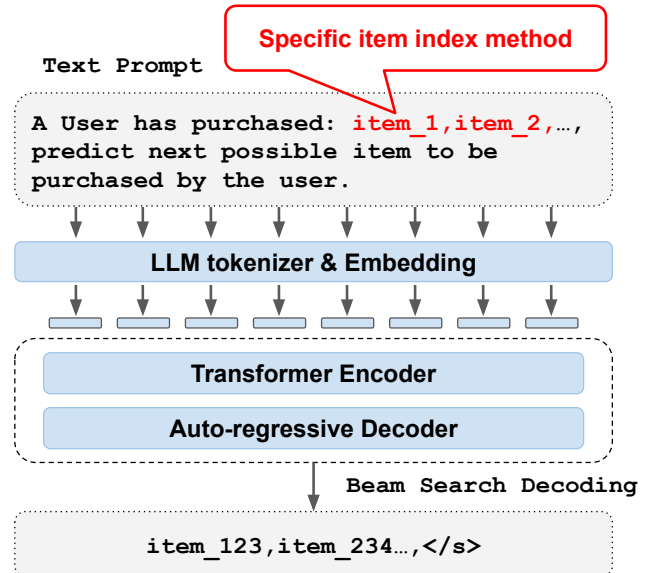

Figure 1: An Illustration of LLM-based sequential recommendations.

items, which is important for item modeling. To tackle this problem, several algorithms are proposed [25, 26, 43, 47, 48]. For example, FDSA [47] proposes a self-attention block to leverage item attribute information. S^3rec [48] maximizes the mutual information of context information in different forms to improve sequential recommendation. Despite the remarkable improvements made by these methods, they can still be further improved via the excellent world knowledge and context understanding ability of LLMs.

2.2 LLM-based Recommendation

The prevalence of LLMs has introduced a paradigm shift into recommender systems [6]. Early approaches investigate the practicality of textual representations generated by a language model for recommendations [9, 22, 30, 44]. By pre-training and then fine-tuning on downstream datasets, enhanced representations can be obtained [5, 15, 20]. The emergence of generative LLMs shifts recommendation system towards generative paradigm [32]. Early attempts explore the potential of LLMs on recommendation via prompt or in-context learning [24, 36, 38]. TallRec [1] trains the LLM to predict whether a user will like a new item given users' interaction history, underscoring the importance of instruction tuning. P5 [7] reformulates several recommendation tasks into a natural language generation task via personalized prompts. Refer to Figure 1 for an overall understanding of its architecture.

In order to leverage the power of LLMs, items are usually represented in natural language form via specific item indexing methods, which can be roughly divided into two categories. The first category proposes to reflect item relations via share tokens, such as semantic IDs [16, 27]. Given the continued significance of ID information, recent studies keep the IDs and represent items as 'item_1234', which will be tokenized into a token sequence 'item','_','12','34' before being input to LLMs [7, 21, 29]. The second category represents items via context information. Tallrec [1] uses item title to represent

Table 1: Inference time (ms) and input length

Method	Avg length	Time/Batch
Bert4rec	21.0	2.37
P5	112.2	2280
POD	85.0	2170

items. RECFORMER [20] flattens the key-value attribute pairs of the item as a sequence to accommodate more textual information, and employs a Longformer [2] to process the long context.

Although these methods have achieved remarkable improvements, they still face slow inference problem. Recent studies have delved into this issue. POD [21] improves P5 [7] by distilling discrete prompt into continuous prompt to reduce the input length, thus reducing inference time. llamaRec [45] proposes a two stage framework. Specifically, it retrieves candidate items via traditional ID-based methods, and designs a verbalizer approach for re-ranking. E4SRec [23] proposes to integrate ID embeddings extracted from a pre-trained SASRec with instruction prompt but neglects the rich context information of items. LightLM [27] proposes a tailored transformer-based architecture to achieve effective recommendations. It improves the inference efficiency by reducing the number of neurons thus reducing the computation demands. However, the modification of architecture may destroy the pre-trained knowledge of LLMs, leading to sub-optimal performance. Despite the strides made in improving inference efficiency, these methods still fall short in addressing the most time-consuming component and the issue of redundant computation.

3 PRELIMINARIES

In this section, we empirically study the inference efficiency problem of existing LLM-based recommendations. First, we compare the inference time of LLM-based methods with traditional sequential recommendation algorithm BERT4rec [35] and analyse the impact of each component in LLM-based recommendations in terms of inference efficiency. Then, we investigate the impact of different components in terms of performance effectiveness. Before detailing the aforementioned preliminary studies, we first introduce the problem statement and some key notations we will use in this work.

3.1 Problem Statement and Notations

We denote $u \in \mathcal{U}$ and $i \in \mathcal{I}$ for a user and an item, where $\mathcal{U}$ and $\mathcal{I}$ indicate the user set and the item set, respectively. The interaction history of a user u can be organized as a sequence $\mathcal{I}_u = (i_1, i_2, ..., i_t)$ in a chronological order, where t is the length of $\mathcal{I}_u$ and each item i is associated with textual information $\mathcal{T}_i$ (e.g., title and genre). Given the interaction history $\mathcal{I}_u$ of user u, sequential recommendation algorithms aim to predict next item i_{t+1} the user is most likely to interact with from $\mathcal{I} \setminus \mathcal{I}_u$, which represents the item set formed by excluding the items already interacted with the user from the complete set of items $\mathcal{I}$.

3.2 Efficiency Analysis

In this subsection, we will explore the following questions: 1) In terms of inference efficiency, how do LLM-based recommendation
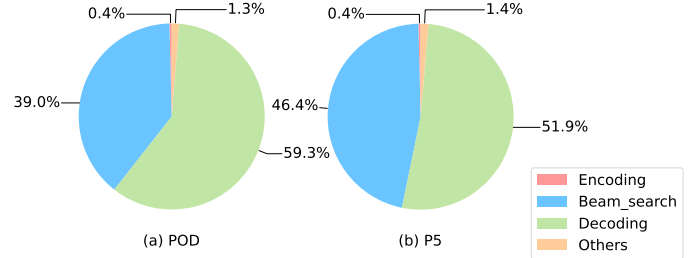
Figure 2: Inference time of different components for a batch of 32 users. 'Beam_search' refers to the beam search decoding process, 'Decoding' transfers the latent representation to token IDs, and 'Others' includes data preparation, metrics calculations etc.

algorithms perform compared with traditional recommendations; and 2) Why do they face such serious inference efficiency problem. To answer these questions, we first demonstrate the huge inference efficiency gap between LLM-based and traditional sequential recommendation systems and then investigate the inference time cost of each component, and the relation between input length and inference time. In this study, we focus on two typical LLM-based recommendation algorithms, i.e., P5 and POD [7, 21], and one representative traditional sequential recommendation system, i.e., BERT4rec [35].

For all the following experiments, we use the implementations released by POD [21][1] and we keep the same experimental settings. Our analysis is based on ML-1m and Movies, two popular publicly-available datasets. More details of the datasets can be found in Section 5.1.1. We calculate the inference time by measuring the total time of all the test data going through each component during inference, and then we can obtain the average inference time for a batch of 32 users. We denote the length of the input as the number of tokens of the tokenized context (e.g., 'item_1234' will be tokenized as 'item' '_' '12' '34' by the tokenizer of T5, and the input length will be 4). The results are shown in Table 1. We can make the following observations: 1) Compared with Bert4rec, the efficiency of P5 and POD has been significantly compromised, experiencing a slowdown of nearly a thousand times; and 2) POD is more efficient than P5 via reducing the input length which can lead to the decrease of the computational costs. To further explore why LLM-based methods face such formidable inference problem, we conduct time analysis on the inference time cost of each component in P5 and POD [7, 21]. Figure 2 demonstrates the average inference time cost of each component. From the figure, we note that the beam search decoding process is most time-consuming which takes approximately 98.2% on P5 [7] and 98.3% on POD [21] of the inference time.

3.3 Effectiveness Analysis

In the previous subsection, we have demonstrated that the inefficiency of LLM-based recommendations comes from the beam search decoding process and input length. In this subsection, we investigate how the beam search decoding and item indexing affect

---

[1]https://github.com/lileipisces/POD

Table 2: Performance of P5 and POD and their two variants

Datasets	ML-1m			Movies		
Methods	R@20	N@20	Time	R@20	N@20	Time
P5	0.2985	0.1442	2,280	0.1080	0.0761	1,770
w/o_d	0.3109	0.1459	40.06	0.1217	0.0794	37.01
w/o_d_TID	0.3354	0.1586	95.40	0.1401	0.0905	99.02
POD	0.2992	0.1403	2,170	0.1089	0.0761	1,400
w/o_d	0.3022	0.1416	35.21	0.1330	0.0866	32.13
w/o_d_TID	0.3339	0.1567	90.26	0.1406	0.0908	80.38

the inference efficiency and the performance. We implement two variants of P5 and POD as follows:

(1) w/o_d. It eliminates the decoder and uses an item projection head to perform the recommendation task (Details in Section 4.3).

(2) w/o_d_TID. On the basis of eliminating the decoder, it represents items with their titles instead of random numbers.

We report the performance and inference efficiency in Table 2. We can make the following observations: 1) eliminating the beam search decoding can significantly improve the inference efficiency; 2) although representing items with their titles can improve the performance due to the incorporation of contextual information, it impairs the inference efficiency because the length of input becomes longer (increases to 297.4 on ML-1m), resulting in more computational costs.

3.4 Discussion

In this subsection, we summarize key findings from the preliminary studies as follows: (1) Compared with Bert4Rec, LLM-based recommendations are much more time-consuming. (2) Reducing the input length will improve the inference efficiency. (3) Beam search decoding will have negative impacts on the efficiency of sequential recommendation. (4) Compared with random number, item title can better represent items and achieve better performance due to the incorporation of contextual information. (5) Despite the advantages of item title, they are usually very long and will increase computational costs. These findings provide the groundwork for us to simplify existing LLM-based recommendations and propose a simple but effective framework Lite-LLM4Rec for sequential recommendations.

4 THE PROPOSED FRAMEWORK

Motivated by our findings, in this section, we aim to simplify the architecture to obtain a better sequential recommendation system, which is easier to train and can achieve low-latency inference and better performance. In this section, we introduce the proposed framework Lite-LLM4Rec. We first give an overview of Lite-LLM4Rec. Then we detail its key components and finally give its training details.

4.1 An Overview

Figure 3 demonstrates the whole architecture of Lite-LLM4Rec. To mitigate redundant computation and enhance inference efficiency,
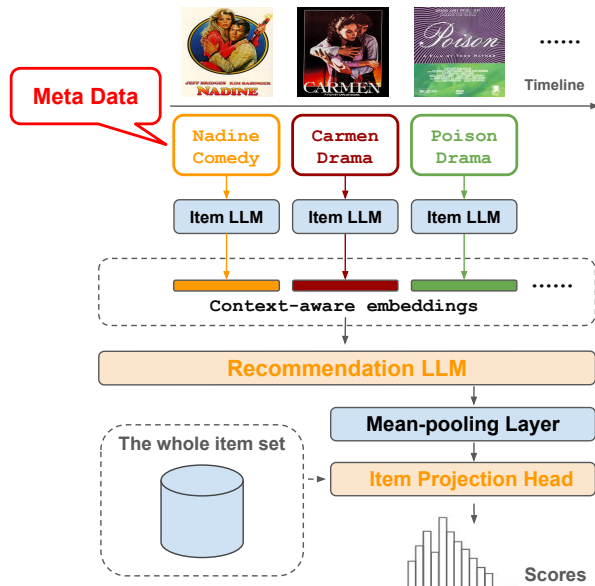
Figure 3: An overview of the architecture of Lite-LLM4Rec.

Lite-LLM4Rec proposes a hierarchical LLM structure which contains two LLM components: Item LLM and Recommendation LLM. We first encode the extensive context information of items via the Item LLM into context-aware vectors and then the Recommendation LLM takes these context-aware vectors as input instead of the original lengthy contexts. Then, Lite-LLM4Rec proposes to replace the beam search decoding process with an item projection head. As a result, it can generate recommendations more effectively. In the following subsections, we will detail our model.

4.2 Hierarchical LLM Structure

In order to leverage the power of LLMs, existing LLM-based recommendations usually formulate items into natural language through various indexing strategies, which can be fed into LLMs. These item indexing methods can be roughly divided into two categories. The first category encodes item relations into the item indexing such as semantic IDs [16, 21, 27]. These methods can capture item relations via shared tokens but the indexing by itself doesn't contain any semantic information, which may not be able to fully explore the potential of LLMs. The second category denotes an item by its meta data such as title or genre [1, 20]. These methods take item contextual information into consideration but the inputs are usually very long, which needs more time to process or even worse they need truncation or special architectures to process when they exceed the length limitation of LLMs.

In addition, the above two kinds of item indexings will be tokenized before being fed into the LLM. This tokenization process leads to increased computational complexity, as each item will be represented by multiple tokens. Moreover, redundant computations occur when the same item appears multiple times in the input. For example, in the ML-1m dataset, the movie 'Star Wars: Episode I - The Phantom Menace (1999)' will be tokenized into 11 tokens by

T5 [31], and appears 539 times in the input during inference, leading to severe redundant computation problem and correspondingly hurting the inference efficiency. .

Lite-LLM4Rec aims to simplify existing item indexing methods while also being able to leverage the power of LLMs. Particularly, we propose a hierarchical LLM structure comprising two distinct LLM components: Item LLM and Recommendation LLM. The Item LLM is to encode extensive context information of an item into a compact, context-aware vector representation. Leveraging its capabilities, the Item LLM can effectively capture the contextual nuances and dependencies within the input sequence, facilitating the creation of a context-aware vector for each item. Then, the Recommendation LLM processes the sequence of context-aware vectors rather than the original lengthy context sequences. Thus the input length of LLM can be significantly reduced. In the following, we will give the details about acquiring context-aware vectors for items and sequence representations.

We represent the context information of item i (e.g., title, genre) after tokenization by $\mathcal{T}_i^e = (w_1, w_2, ..., w_L)$, where w denotes tokens of the context information and L is the length of the context. Then, we input $\mathcal{T}_i^e$ into the item LLM which can be denoted as:

$$Input(\mathcal{T}_i^e) = (w_1, w_2, ..., w_L). \tag{1}$$

Through the model, we can obtain the output representation for each token:

$$Output(\mathcal{T}_i^e) = (h_{w_1}, h_{w_2}, ..., h_{w_L}). \tag{2}$$

Then we can obtain the context-aware embedding for the item:

$$h_i = Mean\_pooling(h_{w_1}, h_{w_2}, ..., h_{w_L}), \tag{3}$$

where $h_{w_i} \in \mathbb{R}^{1 \times d}$ is the representation for the corresponding token w_i. Notice that deriving high quality item representation is not our focus, so we just apply a simple mean-pooling over all tokens' representation to get the final context-aware embedding for each item. More complicated representation methods will be explored as one future work.

After obtaining context-aware embeddings for items. The Recommendation LLM will take a sequence of context-aware embeddings as input instead of the lengthy natural language input. Similarly, we also apply a simple mean-pooling function to the output of the Recommendation LLM to obtain the representation of sequence:

$$h_u = Mean\_pooling(Rec\_LLM(h_1, h_2, ..., h_t)), \tag{4}$$

where h represents the context-aware embedding for an item, Rec_LLM indicates the Recommendation LLM with per-trained weights.

This kind of architecture has several advantages: First, we encode the long context information of an item via the Item LLM into a single embedding, which greatly reduces the length of the input sequence while preserving the semantic information. Second, it mitigates the redundant computation problem since the item representation can be directly obtained from the Item LLM instead of computing repetitively every time when the item appears. Thus, Lite-LLM4Rec can improve inference efficiency via reducing computational costs.

Algorithm 1: Training Process of Lite-LLM4Rec

Input: The sequential data $\mathcal{D}$, Context information of items $\mathcal{T}$, hyper-parameter settings;

Output: Model parameters θ;

1 **while** *not coverage* **do**
2   Draw a batch of data $\mathcal{B}$ from $\mathcal{D}$;
3   **for** $(u, \mathcal{I}_u)$ *in* $\mathcal{B}$ **do**
4   Sample an item segment I_u^{seg} from $\mathcal{I}_u^{|s_u|-2}$;
5   **for** i *in* I_u^{seg} **do**
6   Acquire context information $\mathcal{T}_i$;
7   Calculate context-aware embeddings h_i through Eq. 1 - Eq. 3;
8   **end**
9   Acquire the user representation h_u through Eq. 4;
10   **end**
11   Compute logits through Eq. 5 ;
12   Compute L_{CE} through Eq. 6;
13   Update the parameters θ of the model through L_{CE} ;
14 **end**

4.3 Item Projection Head

Our findings in preliminary studies suggest that the beam search decoding during the auto-regressive generation is the most time-consuming component. This operation is originally designed for natural language generation, which generates the tokens step by step. When it is applied to sequential recommendations, it will generate the recommendation items in natural languages form (e.g., generating a string 'item_1123').

Our preliminary study suggests that this operation is unnecessary for sequential recommendations due to the following disadvantages. First, the head layer of LLM will output the probability distribution over the LLM vocabulary, which is unnecessary since item tokens only account for a small proportion of the whole vocabulary. Second, The model will generate non-existent and repetitive items. Third, generating k recommendations for a user will result in k times more computational costs which significantly contributes to inference delays, thereby exacerbating the efficiency problem.

Lite-LLM4Rec proposes to circumvent the beam search decoding process and introduces a simple but effective item projection head, which will directly output the probability over the whole item set. Lite-LLM4Rec treats the probability as the ranking scores and takes items with the highest scores as the final recommendations. In particular, Lite-LLM4Rec adopts a simple one-layer Multi Layer Perceptron (MLP) without bias [19] as the projection head, which can be formulated as follows:

$$logits = W_{proj} h_u, \tag{5}$$

where h_u stands for sequence representation for user u obtained by Eq. 4, W_{proj} is the projection matrix of the MLP, and *logits* represents the output scores over the whole item set.

4.4 Model Training

Since we use an item projection head to acquire recommendation scores, there is no need to do negative-sampling because it implicitly

Table 3: Statistics of the datasets after pre-processing.

Dataset	Users	Items	Interact	Sparsity(%)
ML-1M	6,040	3,706	994,169	0.0446
Movies	79,276	29,946	1,775,011	0.0007
Toys	11,803	8,569	206,103	0.0020

considers all the un-interacted items as negative samples. We train the hierarchical model with a cross-entropy loss as shown below:

$$L_{CE} = -\sum_{i=1}^{N} y_i log r_i, \tag{6}$$

where N is the number of items, y_i represents the ground-truth for item i, which is 1 if item i is the ground-truth item, otherwise 0; and r_i is the predicted score of item i. We provide the training procedure as well as the implementation details in Algorithm 1. We adopt two training strategies which are denoted as 'Lite-LLM4Rec _sampling' and 'Lite-LLM4Rec _all' in Table 4 and their details are given in 5.1.4.

5 EXPERIMENT

In this section, we conduct comprehensive experiments to verify the effectiveness and efficiency of the proposed Lite-LLM4Rec. In particular, we try to answer the following questions:

- Can the proposed Lite-LLM4Rec achieve better overall performance? (Section 5.2)
- Can the proposed Lite-LLM4Rec improve inference efficiency? (Section 5.3)
- How does Lite-LLM4Rec perform on Top-N recommendation task? (Section 5.4)
- How do different components of Lite-LLM4Rec affect the recommendation performance? (Section 5.5)

5.1 Experimental Settings

5.1.1 Datasets. To evaluate the effectiveness of Lite-LLM4Rec, we conduct a series of experiments on three real-world benchmark datasets, including ML-1m[2] [8], Amazon-Movies and Amazon-Goys&Games[3] [28]. We partition them into training, validation and test sets with the commonly used leave-one-out strategy. It takes the second-to-last item as the validation item, the last item as the test item and all other items as training items in each user's interaction history. Table 3 shows the statistics of these datasets and some details are as follows:

- The ML-1m dataset is an open dataset for movie recommendations. There are approximately 100k interactions. We adopt 5-core filtering strategy where we filter out users and items with less than 5 interactions.
- We consider two categories of Amazon dataset corpora: Movies and Toys&Games (denoted as Toys for clarity). These datasets are collected from the e-commerce platform Amazon[4] with item meta data, user reviews and ratings. We adopt

10-core filtering strategy to filter the users and items with less than 10 interactions to ensure data quality.

5.1.2 Evaluation protocols. We adopt two widely used metrics Recall@k and NDCG@k, where $k = 10, 20$. Recall@k represents the coverage of ground-truth items that appear in the final recommendation list. NDCG@k measures the ranking quality of the final recommendation items. For both metrics, a larger value indicates better performance. For our method and the baselines, we evaluate the performance on the whole item set, and the reported results are the average values over all users.

5.1.3 Baselines. We choose representative methods from three groups as baselines, i.e., traditional ID-based sequential models, Context-aware ID-based models, LLM-based recommendation models. We consider the following tradition ID-based sequential models: (i) **SASRec** [18] proposes an unidirectional attention-based sequential model which can capture long-term semantics to predict the next item. (ii) **BERT4Rec** [35] introduces a bidirectional attention-based transformer to model user behavior sequences. It introduces the Cloze objective into sequential recommendations.

The context-aware ID-based sequential models include: (i) **S3Rec** [48] applies self-supervised learning to the sequential recommendation task. It proposes four self-supervised optimization objectives to maximize the mutual information of context information to learn the correlation between items. (ii) **FDSA** [47] proposes to model feature transitions through different self-attention blocks. It integrates with item-level transitions for modeling user's sequential intents.

We choose the following LLM-based sequential models: (i) **P5** [7] transforms various recommendation tasks into the conditional natural language generation task via personalized prompts and integrates them into a unified framework. (ii) **POD** [21] proposes to distill knowledge in the discrete prompt into continuous prompt vectors, which is more flexible and expressive and can reduce the inference time. (iii) **LightLM** [27] proposes a tailored Transformer-based recommender, which is effective and efficient for generative recommendations. Since **llamaRec [45]** is a two-stage framework while Lite-LLM4Rec is single-stage, we do not choose it as one baseline.

5.1.4 Implementation details. We use T5-small from Huggingface[5] as our backbone in the main experiments. The encoder and decoder in this model both have 6 layers, each of which is an 8-headed attention layer. We find that after training, further fine-tuning the context-aware embeddings and the Recommendation LLM will result in better performance. In the ablation study 5.5.2, we also investigate the impact of different backbones for Item LLM. Following [21], we randomly sample a segment of no more than 21 items from a user's interaction history for each iteration which is denoted as 'Lite-LLM4Rec _sampling' in Table 4. We also implement another training strategy where we traverse all the training data without sampling which is represented as 'Lite-LLM4Rec _all'. We set the batch_size for three datasets to 256 and the learning rate to 0.0005. The embedding dimension is set to 512. The dropout rate is 0.8 and the weight_decay is 0.1 for Movies and Toys. The dropout rate is 0.7 for ML-1m. The warm_up rate and the adam_eps are set to 0.1 and 1e-6 for Movies and Toys, respectively. All methods

---

[2]https://grouplens.org/datasets/movielens/1m/
[3]https://cseweb.ucsd.edu/~jmcauley/datasets.html#amazon_reviews
[4]https://www.amazon.com/

[5]https://huggingface.co/t5-small

Table 4: Overall performance comparison of on ML-1m, Movies and Toys with conventional sequential baselines and LLM-based sequential methods. R denotes Recall, and N denotes NDCG. Boldface represents the best results. Underscore indicates the second best results.

Datasets	ML-1m				Movies				Toys			
Method	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20
BERT4Rec	0.1981	0.2892	0.1093	0.1323	0.0761	0.1048	0.0493	0.0565	0.0352	0.0543	0.0179	0.0227
SASRec	0.1834	0.2785	0.0801	0.1040	0.1082	0.1449	0.0714	0.0807	0.0561	0.0776	0.0312	0.0366
S^3-Rec	0.1776	0.2733	0.0882	0.1123	0.0764	0.1074	0.0481	0.0560	0.0538	0.0811	0.0276	0.0345
FDSA	0.1962	0.2854	0.1058	0.1283	0.1151	0.1500	0.0804	0.0891	0.0568	0.0860	0.0344	0.0417
LightLM	0.1705	0.2531	0.0928	0.1135	0.0751	0.0958	0.0521	0.0574	0.0509	0.0721	0.0302	0.0355
P5	0.2149	0.2985	0.1232	0.1442	0.0905	0.1080	0.0718	0.0761	0.0543	0.0661	0.0356	0.0416
POD	0.2185	0.2992	0.1201	0.1403	0.0904	0.1089	0.0715	0.0761	0.0565	0.0690	0.0421	0.0452
Lite-LLM4Rec _sampling	0.2733	0.3770	0.1518	0.1780	0.1156	0.1447	0.0856	0.0929	**0.0682**	**0.0927**	**0.0426**	**0.0488**
Lite-LLM4Rec _all	**0.3209**	**0.4255**	**0.1866**	**0.2129**	**0.1253**	**0.1596**	**0.0906**	**0.0992**	0.0627	0.0894	0.0382	0.0449

are implemented using Pytorch with an AdamW optimizer. For the hyper-parameters of baselines, we use the values suggested by the original papers with carefully fine-tuning on the three datasets. We check the validation performance every epoch and adopt early-stop when the validation performance of R@10 doesn't improve for 20 consecutive times.

5.2 Effectiveness Comparison

The performance comparison is shown in Table 4. From the results, we can make the following observations: 1) Lite-LLM4Rec exhibits significantly better performance than LLM-based recommendation baselines. Notably, the average Recall@10 improvements over the best results of LLM-based recommendation baselines are 46.8% on ML-1m, 38.4% on Movies, and 20.7% on Toys, highlighting the effectiveness of our design. The potential reasons for the performance improvement are two-fold. First, the beam search decoding process will score the whole LLM vocabulary where item tokens are a small proportion while the item projection head in Lite-LLM4Rec only scores items directly. Second, we propose a hierarchical LLM structure, where item information is encoded into context-aware vectors instead of the original natural language form. 2) Lite-LLM4Rec consistently outperforms the best performance over baselines. The improvement on NDCG@20 is approximately 51.4% on ML-1m, 11.3% on Movies, and 7.9% on Toys, which can be attributed to the power of LLMs. Notice that traditional ID-based and context-aware recommendation algorithms are still competitive. We will perform ablation study in Section 5.5.1 to further study how Lite-LLM4Rec works.

5.3 Efficiency Comparison

In this subsection, we analyze the inference efficiency of Lite-LLM4Rec with LLM-based recommendation baselines. Since operations like obtaining context-aware embeddings and metrics computing and so on can be done off-line, we just measure the time between inputting the data to the model and obtaining the final recommendations as the inference time. The results are shown in Table 5. Apparently, Lite-LLM4Rec can achieve superior inference efficiency. The improvement over POD is approximately 99.71% on ML-1m, 99.63% on Movies, and 99.66% on Toys. We contribute the

Table 5: Comparison of inference time (ms) for a batch of 32 users.

Datasets	ML-1m	Movies	Toys
P5	2,280	1,770	1,620
POD	2,170	1,400	1,490
Lite-LLM4Rec	**6.13**	**5.14**	**5.03**
Improvement	99.71%	99.63%	99.66%

Table 6: Comparison of average input length for a batch of 32 users.

Datasets	ML-1m	Movies	Toys
P5	112.2	112.1	107.1
POD	85.0	89.1	83.5
Lite-LLM4Rec	**21**	**20.7**	**20.8**
Improvement	75.2%	76.7%	75.0%

efficiency improvement to the following reasons. First, we remove the most time-consuming part - beam search decoding. Second, the hierarchical LLM structure we propose to process long context can mitigate the redundant computation problem and reduce the computational costs. We also report the comparison of input length in Table 6. As can be seen, the advantage of Lite-LLM4Rec is evident, as the improvement of input length is 75.2%, 76.7%, and 75.0% over POD on three datasets, respectively.

5.4 Top-N Recommendation

Besides sequential recommendation, we also apply Lite-LLM4Rec to Top-N recommendation task. We follow the setting of POD [21]. We input one ground-truth item along with 99 negative items to the LLM and fine-tune the model to predict the ground-truth item. Finally, We test the model over the 99 negative examples. The results are reported in Figure 5.

From the figure, we can find that Lite-LLM4Rec can also perform well on this task. Our method can achieve much better performance on NDCG than P5 [7] and POD [21]. This indicates that our method
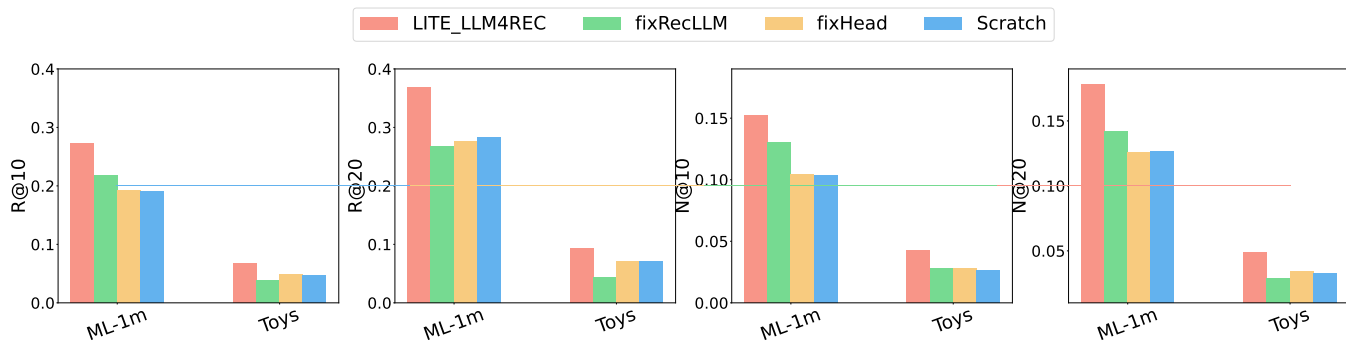
Figure 4: Ablation study results of R@10, R@20, N@10 and N@20 on ML-1m and Toys dataset.
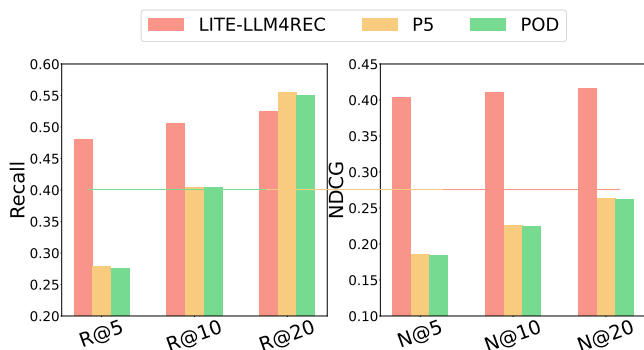

Figure 5: Results of Top-N recommendation on ML-1m.

Table 7: Impact of Item LLMs

ML-1m	R@10	R@20	N@10	N@20
Lite-LLM4Rec	0.2733	0.3770	0.1518	0.1780
T5-base	0.2728	0.3742	0.1512	0.1767
Bert	0.2706	0.3749	0.1529	0.1793
Sbert	0.2668	0.3707	0.1486	0.1748

can greatly improve the ranking quality especially when candidates are given. For Recall, our method can perform better when k is small. This is because the leave-one-out strategy we adopt where we only have one ground-truth item and our method can already rank the ground truth item in high position.

5.5 Ablation Study

5.5.1 Effectiveness of each component. In this section, we aim to analyse how different components influence the overall performance. We conduct experiments on ML-1m and Toys datasets with Lite-LLM4Rec _sampling strategy to assess each component. We design the following variants of our model:

(1) fixT5: Recommendation LLM (the encoder of T5) is fixed.
(2) fixHead: item projection head is fixed.
(3) Scratch: the parameters of the Recommendation LLM is randomly initialized instead of loading pre-trained weights.

We demonstrate the results in Figure 4. From the results, we can have the following observations. First, each component in our framework contributes to the overall performance since fixing any one of the components results in the performance drop. Second, if we train a T5 from scratch instead of using the pre-trained weights for the recommendation LLM, the performance will drop either. This demonstrates the knowledge stored in LLM is of help to our recommendation task.

5.5.2 Impact of Item LLMs. In order to explore whether the performance gains come from the alignment between the Item LLM and

the Recommendation LLM (both of them are T5), we conduct experiments with other Item LLM backbones like Bert [4], Sbert [33] and T5-base [31]. For Bert, we use 'bert-base-uncased' version and take the output of the pooler layer as the sequence representation. For Sbert, we use 'all-MiniLM-L6-v2' version. Since the dimension of hidden state is not matched, one-layer MLP is adopted to transform the dimension. The results are reported in Table 7. We can find that other Item LLMs can also achieve satisfying performance, which indicates that the influence of the backbones of Item LLMs is limited.

6 CONCLUSION

In this work, we propose Lite-LLM4Rec, a simplified but effective LLM-based sequential recommendation model, which can achieve low-latency inference as well as better performance. Through experimental explorations, we find that the beam search decoding process is burdensome and unnecessary for sequential recommendation task. We propose to circumvent it and use an item projection head to acquire recommendations. Additionally, we find that existing item indexing methods will lead to high computational costs and redundant computations. To address the issue, we propose a novel hierarchical LLM structure to process the long context information of items efficiently while also enjoying the power of LLMs. These two designs effectively tackle the inference problem in existing LLM-based recommendation algorithms. Experiments conducted on three real-world datasets can demonstrate that Lite-LLM4Rec can achieve superior inference efficiency while also improving the overall performance. In the future, we are interested in empowering our models with inductive learning ability. Also, we are interested in how the item indexing impacts the training of the model. Meanwhile, the size of the backbone of recommendation LLMs may affect the performance and we would investigate it as one future work.

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. *arXiv preprint arXiv:2305.00447* (2023).

[2] Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150* (2020).

[3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.

[4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[5] Hao Ding, Yifei Ma, Anoop Deoras, Yuyang Wang, and Hao Wang. 2021. Zero-shot recommender systems. *arXiv preprint arXiv:2105.08318* (2021).

[6] Wenqi Fan, Zihuai Zhao, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Jiliang Tang, and Qing Li. 2023. Recommender systems in the era of large language models (llms). *arXiv preprint arXiv:2307.02046* (2023).

[7] 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.

[8] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)* 5, 4 (2015), 1–19.

[9] 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*. 1096–1102.

[10] Ruining He, Wang-Cheng Kang, Julian J McAuley, et al. 2018. Translation-based Recommendation: A Scalable Method for Modeling Sequential Behavior.. In *IJCAI*. 5264–5268.

[11] Ruining He and Julian McAuley. 2016. Fusing similarity models with markov chains for sparse sequential recommendation. In *2016 IEEE 16th international conference on data mining (ICDM)*. IEEE, 191–200.

[12] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent neural networks with top-k gains for session-based recommendations. In *Proceedings of the 27th ACM international conference on information and knowledge management*. 843–852.

[13] Balázs Hidasi, Alexandros Karatzoglou, Linas Baltrunas, and Domonkos Tikk. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939* (2015).

[14] Balázs Hidasi, Massimo Quadrana, Alexandros Karatzoglou, and Domonkos Tikk. 2016. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 241–248.

[15] Yupeng Hou, Shanlei Mu, Wayne Xin Zhao, Yaliang Li, Bolin Ding, and Ji-Rong Wen. 2022. Towards universal sequence representation learning for recommender systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 585–593.

[16] Wenyue Hua, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2023. How to Index Item IDs for Recommendation Foundation Models. *arXiv preprint arXiv:2305.06569* (2023).

[17] Jin Huang, Zhaochun Ren, Wayne Xin Zhao, Gaole He, Ji-Rong Wen, and Daxiang Dong. 2019. Taxonomy-aware multi-hop reasoning networks for sequential recommendation. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 573–581.

[18] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.

[19] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. 2023. Do LLMs Understand User Preferences? Evaluating LLMs On User Rating Prediction. *arXiv preprint arXiv:2305.06474* (2023).

[20] Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. 2023. Text Is All You Need: Learning Language Representations for Sequential Recommendation. *arXiv preprint arXiv:2305.13731* (2023).

[21] 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.

[22] Ruyu Li, Wenhao Deng, Yu Cheng, Zheng Yuan, Jiaqi Zhang, and Fajie Yuan. 2023. Exploring the Upper Limits of Text-Based Collaborative Filtering Using Large Language Models: Discoveries and Insights. *arXiv preprint arXiv:2305.11700* (2023).

[23] Xinhang Li, Chong Chen, Xiangyu Zhao, Yong Zhang, and Chunxiao Xing. 2023. E4SRec: An Elegant Effective Efficient Extensible Solution of Large Language Models for Sequential Recommendation. *arXiv preprint arXiv:2312.02443* (2023).

[24] Junling Liu, Chao Liu, Renjie Lv, Kang Zhou, and Yan Zhang. 2023. Is chatgpt a good recommender? a preliminary study. *arXiv preprint arXiv:2304.10149* (2023).

[25] Qiang Liu, Shu Wu, Diyi Wang, Zhaokang Li, and Liang Wang. 2016. Context-aware sequential recommendation. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 1053–1058.

[26] Jarana Manotumruksa, Craig Macdonald, and Iadh Ounis. 2018. A contextual attention recurrent architecture for context-aware venue recommendation. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 555–564.

[27] Kai Mei and Yongfeng Zhang. 2023. LightLM: A Lightweight Deep and Narrow Language Model for Generative Recommendation. *arXiv preprint arXiv:2310.17488* (2023).

[28] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*. 188–197.

[29] Junyan Qiu, Haitao Wang, Zhaolin Hong, Yiping Yang, Qiang Liu, and Xingxing Wang. 2023. ControlRec: Bridging the Semantic Gap between Language Model and Personalized Recommendation. *arXiv preprint arXiv:2311.16441* (2023).

[30] Zhaopeng Qiu, Xian Wu, Jingyue Gao, and Wei Fan. 2021. U-BERT: Pre-training user representations for improved recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 4320–4327.

[31] 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 (2020), 5485–5551.

[32] Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan H Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Q Tran, Jonah Samost, et al. 2023. Recommender Systems with Generative Retrieval. *arXiv preprint arXiv:2305.05065* (2023).

[33] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084* (2019).

[34] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing personalized markov chains for next-basket recommendation. In *Proceedings of the 19th international conference on World wide web*. 811–820.

[35] 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.

[36] Weiwei Sun, Lingyong Yan, Xinyu Ma, Pengjie Ren, Dawei Yin, and Zhaochun Ren. 2023. Is ChatGPT Good at Search? Investigating Large Language Models as Re-Ranking Agent. *arXiv preprint arXiv:2304.09542* (2023).

[37] 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*. 565–573.

[38] Paul Thomas, Seth Spielman, Nick Craswell, and Bhaskar Mitra. 2023. Large language models can accurately predict searcher preferences. *arXiv preprint arXiv:2309.10621* (2023).

[39] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).

[40] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-based recommendation with graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 346–353.

[41] An Yan, Shuo Cheng, Wang-Cheng Kang, Mengting Wan, and Julian McAuley. 2019. CosRec: 2D convolutional neural networks for sequential recommendation. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 2173–2176.

[42] 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* (2023).

[43] Weihua Yuan, Hong Wang, Xiaomei Yu, Nan Liu, and Zhenghao Li. 2020. Attention-based context-aware sequential recommendation model. *Information Sciences* 510 (2020), 122–134.

[44] Zheng Yuan, Fajie Yuan, Yu Song, Youhua Li, Junchen Fu, Fei Yang, Yunzhu Pan, and Yongxin Ni. 2023. Where to go next for recommender systems? id-vs. modality-based recommender models revisited. *arXiv preprint arXiv:2303.13835* (2023).

[45] 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* (2023).

[46] Junjie Zhang, Ruobing Xie, Yupeng Hou, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2023. Recommendation as instruction following: A large language model empowered recommendation approach. *arXiv preprint arXiv:2305.07001* (2023).

[47] 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*. 4320–4326.

[48] 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.