



Large Language Models Enhanced Collaborative Filtering

Zhongxiang Sun*
Zihua Si*
Gaoling School of Artificial Intelligence
Renmin University of China
Beijing, China
{sunzhongxiang,zihua_si}@ruc.edu.cn

Xiaoxue Zang
Kai Zheng
Yang Song
Kuaishou Technology Co., Ltd.
Beijing, China
{zangxiaoxue,zhengkai}@kuaishou.com
ys@sonyis.me

Xiao Zhang
Jun Xu†
Gaoling School of Artificial Intelligence
Renmin University of China
Beijing, China
{zhangx89,junxu}@ruc.edu.cn

ABSTRACT

Recent advancements in Large Language Models (LLMs) have attracted considerable interest among researchers to leverage these models to *enhance* Recommender Systems (RSs). Existing work predominantly utilizes LLMs to generate knowledge-rich texts or utilizes LLM-derived embeddings as features to improve RSs. Although the extensive world knowledge embedded in LLMs generally benefits RSs, the application can only take limited number of users and items as inputs, without adequately exploiting collaborative filtering information. Considering its crucial role in RSs, one key challenge in enhancing RSs with LLMs lies in providing better collaborative filtering information through LLMs. In this paper, drawing inspiration from the in-context learning and chain of thought reasoning in LLMs, we propose the Large Language Models enhanced Collaborative Filtering (**LLM-CF**) framework, which distills the world knowledge and reasoning capabilities of LLMs into collaborative filtering. We also explored a concise and efficient instruction-tuning method, which improves the recommendation capabilities of LLMs while preserving their general functionalities (e.g., not decreasing on the LLM benchmark). Comprehensive experiments on three real-world datasets demonstrate that LLM-CF significantly enhances several backbone recommendation models and consistently outperforms competitive baselines, showcasing its effectiveness in distilling the world knowledge and reasoning capabilities of LLM into collaborative filtering.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Large Language Models, Collaborative Filtering

*Equal Contribution. Work done during their internship at Kuaishou.

†Corresponding author.

Work partially done at Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education.

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

CIKM '24, October 21–25, 2024, Boise, ID, USA.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0436-9/24/10

<https://doi.org/10.1145/3627673.3679558>

ACM Reference Format:

Zhongxiang Sun, Zihua Si, Xiaoxue Zang, Kai Zheng, Yang Song, Xiao Zhang, and Jun Xu. 2024. Large Language Models Enhanced Collaborative Filtering. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3627673.3679558>

1 INTRODUCTION

Large Language Models (LLMs) [30, 40] have made rapid advancements, showcasing remarkable capabilities [43] in context comprehension, reasoning, generalization, and modeling world knowledge, and so on. With the advancement of Large Language Models (LLMs), many researchers are focusing on how to utilize LLMs in recommendation systems (RSs). Many studies have already applied LLMs to various aspects of RSs, including ranking [51], Click-Through Rate (CTR) prediction [2, 9, 49], sequential recommendation [14], rating prediction [22], and data augmentation [29, 44]. Considering specific methods to utilize LLMs for RSs, current applications can be classified into two categories. (1) **LLMs as RSs**: LLMs can be directly prompted or be fine-tuned to function as specialized RSs [2, 9, 37, 49, 51]. (2) **LLM-enhanced RSs**: Based on world knowledge and reasoning abilities, LLM-derived embedding vectors and LLM-generated texts can enhance RSs [7, 14, 47, 49, 52].

Despite their effectiveness, there are still several challenges to be addressed. **LLMs as RSs** suffers from low efficiency due to the resource-intensive nature of LLMs, making their practical application challenging. **LLM-enhanced RSs** inadequately exploit collaborative filtering information because the LLM can only take a limited number of users and items as inputs. How to better leverage LLMs to provide enhanced collaborative filtering information to existing RSs becomes key in LLM-enhanced RSs.

Considering the challenges in deploying LLMs as RSs due to their inherently extensive parameterization, we focus on **LLM-enhanced RSs**, which are more applicable and flexible for existing RSs. In order to better guide collaborative filtering to enhance existing RSs with LLMs. Inspired by Chain-of-Thought (CoT) and In-Context Learning [4, 11] in LLMs, we propose a novel **Large Language Models enhanced Collaborative Filtering (LLM-CF) Framework**, which distills the world knowledge and reasoning capabilities of LLM into collaborative filtering in an in-context, chain of thought methodology. As shown in Figure 1, LLM-CF can be decoupled into two parts: (1) **offline service part (§ 4)**: Fine-tune LLM to enhance its recommendation capabilities, generate CoT reasoning with collaborative filtering information, and construct in-context CoT dataset. (2) **online service part (§ 5)**: Retrieve the

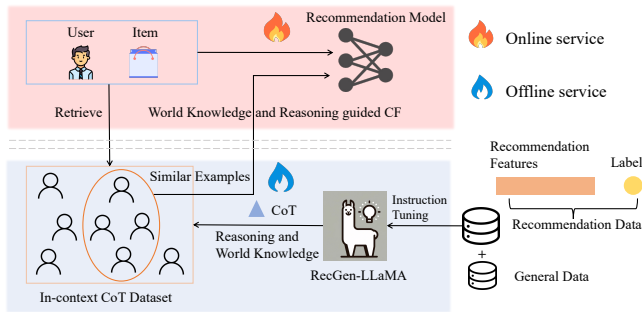


Figure 1: LLM-CF integrates LLM-based world knowledge and reasoning with collaborative filtering to improve recommendation performance, using LLMs with recommender capability and decoupled latency-free offline generation.

in-context CoT examples, learn the world knowledge and reasoning guided Collaborative Filtering (CF) feature, and use this feature to enhance existing RSs.

In the offline service, we perform instruction tuning on LLM to obtain CF information about users and items in the recommendation data. However, our initial findings indicate that full parameter tuning LLMs could result in substantial forgetting of their general capabilities, as discussed in § 4.2. We leveraged a simple but effective data mixing method to finetune LLaMA2 [40], and successfully trained a model **RecGen-LLaMA**, which achieves an optimal balance between general and recommendation capabilities. Then, we use RecGen-LLaMA to generate CoT reasoning for a subset of examples in training data, forming the in-context CoT dataset.

In the online service, the retrieval module uses a query composed of the textual features of the current recommendation features to perform embedding-based retrieval on the in-context CoT dataset, forming in-context CoT examples. These retrieved examples contain similar recommendation features, as well as CoT reasoning generated by RecGen-LLaMA. The in-context CoT examples are concatenated with the current recommendation features and then fed into the **In-context Chain of Thought (ICT)** module of LLM-CF to learn world-knowledge and reasoning guided CF feature. Finally, the enhanced CF feature is fed into the backbone recommendation model for making the final prediction.

Advantage: LLM-CF not only leverages LLMs to provide enhanced collaborative filtering information to existing RSs but also achieves exceptional deployment efficiency. (1) We teach collaborative filtering knowledge from recommendation data to LLMs, ensuring that the generated texts include a comprehensive understanding of user behaviors and item features. (2) When integrating the world knowledge and reasoning capabilities from LLMs into RSs, We design ICT modules that embody both these capabilities and collaborative filtering information. This information comes from the explicit collaborative filtering information contained in the retrieved similar user histories and the RecGen-LLaMA generated CoT reasoning. (3) The time cost of LLM-CF is manageable, as we only require the offline serving of LLM without the need for online inference alongside the RSs.

We summarize the major contributions of this paper as follows:

(1) We represent pioneering work in using the LLMs-based world knowledge and reasoning to guide collaborative filtering features, thereby enhancing conventional recommendation models.

(2) The proposed LLM-CF is inspired by in-context learning and chain of thought reasoning in LLMs, which effectively distills the world knowledge and reasoning capabilities of LLMs into conventional recommendation models in an in-context chain of thought manner. Moreover, the LLM-CF is more efficient compared to existing LLM-enhanced RSs by decoupling LLM generation from the recommendation system’s online services.

(3) We conducted extensive experiments on three public datasets. The experimental results demonstrate that the LLM-CF could significantly improve the recommendation performance of conventional recommendation models in both ranking and retrieval tasks, verifying the effectiveness of the LLM-CF.

2 RELATED WORK

2.1 LLM as RSs

LLMs as RSs involve directly prompting LLMs to make recommendations using natural language-based queries or adapting LLMs to serve as RSs after fine-tuning them with recommendation data. P5 [12] transforms user interaction data into text prompts using item indices for training language models. In contrast, TALLRec [2] utilizes instructional designs to outline recommendation tasks and adapts LLMs through fine-tuning to follow these guidelines, thereby producing recommendations. Further, ReLLa [25] uses retrieved user history to fine-tune LLMs, addressing the issue of LLMs’ weak capability in processing long user sequences. LLAMARec [51] initially applies small-scale recommenders to select candidates from user interaction history. Then, this history and the chosen items are fed into the LLM as text using a specially crafted prompt template. However, directly using LLMs as inference models for recommendation tasks presents challenges, such as high computational costs and slow inference times, causing challenges in meeting the requirements for online services and deployment.

2.2 LLM-enhanced RSs

LLM-enhanced RSs leverage the world knowledge and reasoning abilities of LLMs by utilizing them to generate knowledge-rich texts or employ LLM-derived embeddings as features to enhance RSs. GIRL [52] applies an LLM, fine-tuned with job datasets, to create job descriptions from CVs, boosting traditional job RSs. KAR [47] uses LLMs for generating user Preference Reasoning and item Factual Knowledge, enhancing RSs through hybrid-expert adaptors. ONCE [26] explores both open and closed-source LLMs for RS enhancement; open-source LLMs as feature encoders and closed-source via prompt learning. HKF [49] employs LLMs to merge diverse user behavior data, improving RSs by integrating these semantic features. However, these models primarily focus on limited user-item information, neglecting collaborative filtering information, and suffer from efficiency issues due to real-time LLM processing for new interactions or items. The proposed LLM-CF addresses these by distilling LLMs’ knowledge and reasoning with collaborative filtering into existing RSs, separating LLM generation from online services to achieve efficient LLM-enhanced RSs.

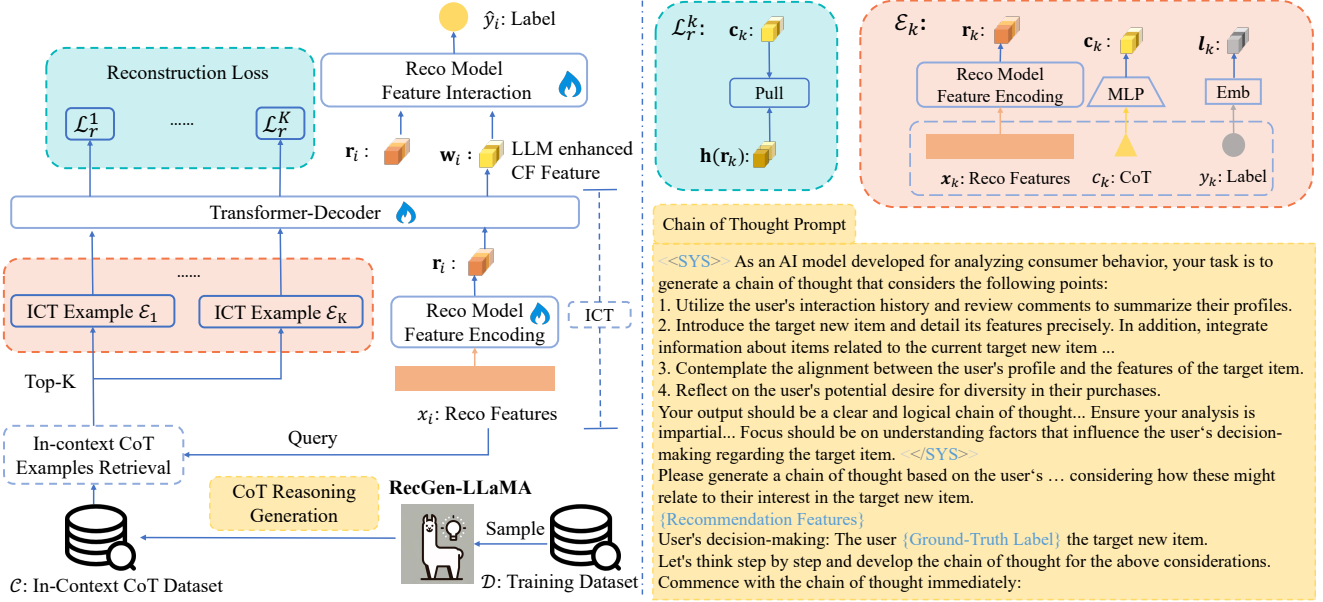


Figure 2: The overall architecture of the proposed model LLM-CF.

3 PRELIMINARIES

In this section, we define the recommendation task as a binary classification problem using multi-field categorical data. The dataset $\mathcal{D} = \{(x_n, y_n)\}_{n=1}^N$ comprises pairs of recommendation features x_i and binary labels y_i . Recommendation features include user features (e.g., User ID, gender, age) and item features (e.g., item ID, brand). The label y_i indicates a click (1) or no click (0). The goal is to learn a function $f(\cdot)$ with parameters θ to predict click probabilities, $\hat{y}_i = f(x_i; \theta)$, for each x_i .

To meet the requirements of large language models, we follow the instruction prompt used in [2, 25], which involves extracting textual format recommendation features x_i^t from recommendation features x_i and organizing it into a "Task Instruction". This instruction guides the LLM to determine whether a user is likely to be interested in the target item based on the user's historical interactions and the other user features. The LLM generates a binary response of "Yes" or "No", with "Yes" indicating a positive interaction (click) and "No" indicating a negative interaction (no click).

To study the world knowledge and reasoning guided collaborative filtering feature based on LLMs, we propose LLM-CF (Large Language Models enhanced Collaborative Filtering). The overview of LLM-CF is illustrated in Figure 2. The proposed LLM-CF can be decoupled into the **offline service part** (§ 4) and the **online service part** (§ 5).

4 OFFLINE SERVICE OF LLM-CF

In this section, we introduce the offline service of LLM-CF in detail.

4.1 Overview

The offline service part of LLM-CF includes the following process:

Training of RecGen-LLaMA: LLMs equipped with recommendation capability can better comprehend the collaborative filtering information within recommendation data, thereby generating improved textual to enhance conventional recommendation models. However, our experiments reveal that directly following the previous work [2] by using recommendation data to fine-tune LLMs leads to catastrophic forgetting of general capabilities. This results in a significant decline in the LLM benchmark to the extent that cannot generate meaningful text. To solve this challenge, we conducted extensive experiments and found a concise and efficient full-parameter instruction-tuning method. This method, which integrates recommendation data with general instruction-tuning data, optimizes the balance between the model's general and recommendation abilities. By applying this method to the widely-used LLaMA2 [40], we successfully trained **RecGen-LLaMA**, achieving an optimal balance between general and recommendation capabilities.

CoT Reasoning Generation: For the recommendation data $(x_i, y_i) \in \mathcal{D}$, we designed a zero-shot CoT prompt that decomposes user-item interaction and then reconstructs them, thereafter inducing RecGen-LLaMA to generate CoT reasoning c based on the textual representation of recommendation features x^t . These $\{c_1, \dots, c_m, \dots, c_M\}$ along with the original recommendation data form the In-Context CoT Dataset $\mathcal{C} = \{(x_m, c_m, y_m)\}_{m=1}^M$.

4.2 RecGen-LLaMA

In this section, we explore how to fine-tune an LLM that balances recommendation capability with general capabilities by conducting extensive experimental analysis.

Analysis Setup: For the empirical study, we selected the Amazon [15] Beauty and Sports datasets to evaluate the recommendation capability, with detailed statistical data provided in § 6.1.1. For evaluating the general capability of LLMs, we employed the MMLU

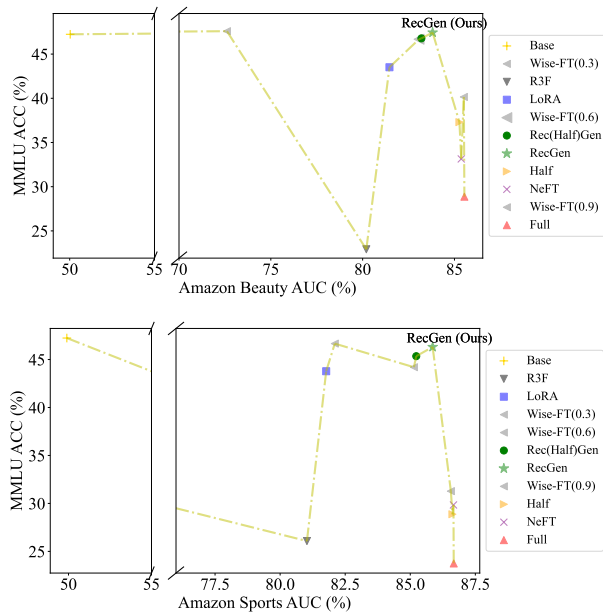


Figure 3: The recommendation and general capabilities of LLaMA2 after fine-tuning on different recommendation data using various methods.

benchmark [16], which is widely utilized in evaluating LLM general capabilities [30, 39, 40]. The evaluation metrics include the AUC (Area Under the Curve) for recommendation capability and ACC (Accuracy) for general capability. Regarding the LLM, we chose the currently widely used LLaMA2-7B-chat [40] as our base model.

We systematically explore diverse methodologies to optimize LLaMA’s recommendation performance while retaining its general capabilities. Our investigation encompasses Continual Learning, Robust Representation Fine-tuning, and Parameter-efficient Fine-tuning approaches: (1) **Base**: Utilize LLaMA2 directly without any fine-tuning. (2) **Half**: Fine-tune LLaMA2 directly using half the amount of the recommendation data. (3) **Full**: Fine-tune LLaMA2 directly using all the recommendation data. (4) **LoRA**: Given that parameter-efficient fine-tuning methods can preserve most of the model parameters and thus retain the model’s general capabilities, we adopted Low-Rank Adaptation [19] on the Query and Value matrices of LLaMA2, adding parameters with a rank of 8. (5) **NeFT**: Adding noise during fine-tuning can prevent overfitting or representation collapse of the pre-trained model [20, 50]. We incorporated noise in the embedding layer of the input when fine-tuning LLaMA2 with recommendation data. (6) **R3F**: Aghajanyan et al. [1] proposed a Fine-tuning method rooted in an approximation to the trust region, which ensures that the pre-trained models do not forget the original pre-trained representations when they are fine-tuned for new tasks. (7) **Wise-FT**: Wortsman et al. [45] leveraged the weighted ensemble of the pre-trained model and the fine-tuned model to enhance various capabilities of the visual models, with the ensemble weight defined as $W_{\text{ensemble}} = \alpha \cdot W_{\text{fine-tuned}} + (1 - \alpha) \cdot W_{\text{pre-trained}}$. We ensemble the original LLaMA2 and LLaMA2 fine-tuned with recommendation data, setting α to 0.3, 0.6, and 0.9 to obtain **Wise-FT(0.3)**, **Wise-FT(0.6)**, and **Wise-FT(0.9)** respectively. (8) **RecGen**: Dong

et al. [10] found that the coding and mathematical abilities of LLMs, along with their general capabilities, can be better balanced through a mixed-data fine-tuning approach. We adopt to utilize a mix of recommendation data and high-quality general data (i.e. LIMA [53] and alpaca-gpt4 [31]) for fine-tuning LLaMA2.

Results and Findings: We presented the results in Figure 3. As shown in the figure, using the off-the-shelf LLaMA2 model (**Base**) for recommendation tasks has poor performance. However, fine-tuning LLaMA2 with recommendation data (**Full**) leads to optimal recommendation performance on both Amazon Sports and Amazon Beauty, but it diminishes LLaMA2’s general capabilities. We also found that **RecGen** and **Wise-FT(0.6)** achieve a good balance between general capabilities and recommendation performance, with **RecGen** showing superior results. Based on **RecGen**, we developed **RecGen-LLaMA** to enhance conventional recommendation models.

4.3 CoT Reasoning Generation

To enable RecGen-LLaMA to generate CoT reasoning with world knowledge and collaborative filtering information, we design the Chain of Thought prompt \mathcal{P} . As shown in Figure 2, this prompt can decompose user-item interactions and then reconstruct them to analyse the relation between them, aiming to analyze user data by simulating the step-by-step reasoning process of humans. Taking product recommendation as an example: Firstly, we had RecGen-LLaMA systematically analyze the user’s interaction history and feedback comments to establish a detailed user profile. Next, RecGen-LLaMA introduced the target new product and its related features in detail to better understand the target product. Finally, RecGen-LLaMA further analyses the alignment between the user profile and the target product features, reflecting on the user’s potential needs for shopping diversity. For the training recommendation data $(x_i, y_i) \in \mathcal{D}$, this generation process can be represented as:

$$c_i = \text{RecGen-LLaMA}(x_i, y_i, \mathcal{P}). \quad (1)$$

Considering the resource constraints in real recommendation scenarios, we can sample M training examples from \mathcal{D} for CoT reasoning generation. The sampling process can be random sampling or select representative user interaction histories. In this paper, we adopted the uniform random sampling, resulting in $\{c_1, \dots, c_m, \dots, c_M\}$ combined with the original recommendation examples to form the In-context CoT dataset $\mathcal{C} = \{(x_m, c_m, y_m)\}_{m=1}^M$.

4.4 Efficiency Analysis of Offline Service

When deployed, LLM-CF shows high efficiency compared to previous LLM-enhanced RSs work. It avoids real-time generation by LLMs, requiring only periodic updates on the history dataset and successfully decoupling LLM generation from the recommendation system’s online services. Next, we will analyze the time efficiency from the Training and Generation of RecGen-LLaMA.

Training: Based on the experiment results in Figure 3, we discovered that training with half the amount of recommendation data (i.e., **Half**) as compared to using the full dataset (i.e., **Full**) achieved better general capabilities and nearly the same recommendation performance. This phenomenon was also observed in the case of RecGen-LLaMA trained with half the recommendation

Table 1: Comparing LLM-CF with KAR in terms of the efficiency of using LLMs to generate data for dynamic scenarios. ✓ indicates that it does not require the real-time generation of new data; ✗ indicates the opposite.

Scenario	KAR	LLM-CF
New users	✗	✓
New items	✗	✓
New interaction (no new items/users)	✓	✓

data (i.e., **Rec(Half)Gen**). This further proves that in real-world Recommender Systems scenarios, we can use a small amount of recommendation data to fine-tune LLMs to enhance recommendation capability, thereby significantly reducing the training overhead of RecGen-LLaMA in real-world Recommender Systems.

Generation: In the generation phase, we compared LLM-CF with the currently most efficient LLM-enhanced RSs approach KAR [47], which achieves a certain speed-up with prestore generation when user features and item features are relatively fixed. In our efficiency analysis, we have more thoroughly considered three scenarios (i.e., new interactions, new users, new items). Our model’s decoupling of generation and online recommendation effectiveness is evident from Table 1. In all three scenarios, our model does not require real-time generation to meet online recommendation needs. This is particularly significant for scenarios like short-video recommendations, where many new items appear daily, further reducing the time delay in system online services.

5 ONLINE SERVICE OF LLM-CF

In this section, we introduce the online service of LLM-CF in detail.

5.1 Overview

The online service part of LLM-CF includes the following components:

In-context CoT Examples Retrieval: The Retrieval process involves finding the top- K historical recommendation examples similar to the current recommendation data in order to provide explicit collaborative filtering information. This involves identifying $\mathcal{I}_i = \{(x_k, c_k, y_k)\}_{k=1}^K$ from C , which includes recommendation data similar to the current recommendation data, as well as CoT reasoning c_k containing world-knowledge and collaborative filtering information.

In-context Chain of Thought (ICT) Module: Inspired by the success of in-context learning and chain of thought in LLMs, we use \mathcal{I}_i as in-context CoT examples and x_i as the query. By employing a transformer decoder layer for In-context Chain of Thought learning, we learnt the world-knowledge and reasoning guided collaborative filtering feature w_i , which can be used to enhance underlying recommendation models in a model-agnostic manner.

Training: During the training phase, we designed a reconstruction loss for the CoT reasoning in the in-context examples, to further strengthen the world-knowledge and reasoning capabilities contained in the collaborative filtering features generated by ICT module.

5.2 In-context CoT Examples Retrieval

The Retrieval module is responsible for retrieving similar In-context CoT examples for the current recommendation data (x_i, y_i) . Recent studies in retrieval-augmented recommendation have shown the potential of using current recommendation features (i.e., user features and target item features) for history retrieval to enhance collaborative filtering capabilities [3, 32, 33]. Our approach extends this concept by not only leveraging the collaborative filtering information but also incorporating the world knowledge and reasoning abilities of RecGen-LLaMA.

We use the text format features x_i^t of x_i as the query to retrieve similar In-context CoT examples from the In-context CoT dataset C , where the key for the examples in C is composed of the text format features: $\mathcal{K} = [x_1^t, \dots, x_m^t, \dots, x_M^t]$.

We employed embedding-based retrieval, including the encoding and ranking processes. To implement this process efficiently, the retrieval process is based on the Approximate Nearest Neighbor search [28]. In the Encoding phase, we used the BGE embedding [48] as the text encoder to convert the query x_i^t and the candidate keys of the In-context CoT dataset \mathcal{K} into embedding formats:

$$e(x_i^t), e(\mathcal{K}) = \text{encoder}(x_i^t), \text{encoder}(\mathcal{K}),$$

where $e(\mathcal{K}) = [e(x_1^t), \dots, e(x_m^t), \dots, e(x_M^t)]$.

In the ranking phase, the cosine similarity between the query embedding $e(x_i^t)$ and each candidate key’s embedding in $e(\mathcal{K})$ is computed. Specifically, for each embedding $e(x_m^t)$ in $e(\mathcal{K})$, the cosine similarity with $e(x_i^t)$ is calculated as follows:

$$\text{sim}(e(x_i^t), e(x_m^t)) = \frac{e(x_i^t) \cdot e(x_m^t)}{\|e(x_i^t)\| \|e(x_m^t)\|}$$

where \cdot denotes the dot product and $\|\cdot\|$ denotes the norm of a vector. Subsequently, the indices of the top- K most similar candidate keys are identified based on their cosine similarity scores. These indices correspond to the examples in the In-context CoT dataset C . The In-context CoT examples associated with these indices are extracted from C , forming:

$$\mathcal{I}_i = \{\mathcal{E}_1, \dots, \mathcal{E}_k, \dots, \mathcal{E}_K\},$$

where $\mathcal{E}_k = (x_k, c_k, y_k)$. The \mathcal{I}_i contains the world-knowledge and reasoning enhanced explicitly collaborative filtering information of x_i .

Remark. (1) To prevent data leakage, we ensure that \mathcal{I}_i will definitely not contain interaction information from future time steps of x_i . (2) To ensure that the proportion of positive and negative labels in \mathcal{I}_i does not affect the prediction results of downstream recommendation tasks, we constrain the ratio of positive and negative labels in \mathcal{I}_i to remain equal, thereby blocking shortcuts and highlighting the effectiveness of our model.

5.3 In-context Chain of Thought Module

Inspired by the success of in-context learning and chain of thought in LLMs, the In-context Chain of Thought (ICT) Module learns world knowledge and reasoning-guided Collaborative Filtering features through an in-context chain of thought methodology.

The ICT module utilizes I_i as in-context examples and \mathbf{x}_i as the query, forming the ICT tokens:

$$\mathbf{T} = [\mathbf{x}_1, c_1, y_1, \dots, \mathbf{x}_K, c_K, y_K, \mathbf{x}_i]. \quad (2)$$

The ICT module first encodes the recommendation features (i.e., user features and target item features), CoT reasoning, and label of \mathbf{T} into corresponding tokens embedding \mathbf{E} :

$$\mathbf{E} = [\mathbf{r}_1, \mathbf{c}_1, \mathbf{l}_1, \dots, \mathbf{r}_K, \mathbf{c}_K, \mathbf{l}_K, \mathbf{r}_i]. \quad (3)$$

The high-dimensional sparse one-hot recommendation features in \mathbf{x} are mapped into low-dimensional dense space via an ID Embedding layer with an embedding lookup operation, and the text features of \mathbf{x} are encoded by the text encoder (the same as the In-context CoT examples Retrieval), then the underlying recommendation model’s feature encoding module encodes them into the token \mathbf{r} . The CoT reasoning c is also mapped into low-dimensional dense space via the text encoder and MLP projection, forming the token \mathbf{c} . The label l is the binary one-hot vector, which can also be mapped into a low-dimensional dense space via the ID embedding layer to obtain \mathbf{l} .

Considering existing research indicating a correlation between the in-context learning capability and the structural attributes of transformer decoders [8, 36], we use transformer decoder layers to encode \mathbf{E} . The ICT token embedding \mathbf{E} is then fed through L Transformer Decoder blocks, generating hidden representations \mathbf{H} of ICT tokens:

$$\mathbf{H} = [\mathbf{h}(\mathbf{r}_1), \mathbf{h}(\mathbf{c}_1), \mathbf{h}(\mathbf{l}_1), \dots, \mathbf{h}(\mathbf{r}_i))] = \text{Decoder}(\mathbf{E}). \quad (4)$$

The world knowledge and reasoning guided collaborative filtering feature \mathbf{w} is the next token of \mathbf{T} (i.e., the last hidden representations of \mathbf{H}):

$$\mathbf{w} = \mathbf{h}(\mathbf{r}_i).$$

Please note that the length of decoder input is not long in our experiments. This ensures that the ICT module does not impose much time overhead on the recommendation models. We leave details in Section 5.6.

5.4 Model-Agnostic Application

The primary goal of recommendation models is to learn a function $f(\cdot)$ characterized by parameters θ , which can skillfully predict the click probability $P(y_i = 1|\mathbf{x}_i)$ for each sample \mathbf{x}_i , formalized as $\hat{y}_i = f(\mathbf{x}_i; \theta)$. The world-knowledge and reasoning guided collaborative filtering feature \mathbf{w}_i can be directly utilized in enhanced underlying recommendation models:

$$\hat{y}_i = f([\mathbf{x}_i, \mathbf{w}_i]; \theta). \quad (5)$$

For ranking models, there is a greater need for interactions between recommendation features; here, $[\cdot, \cdot]$ can be a concatenation operation, followed by the use of a feature interaction module to learn deep interactions. In contrast, retrieval models emphasize efficiency more, and $[\cdot, \cdot]$ can be an addition operation.

5.5 Model Training

In the model training phase, for each data $(\mathbf{x}_i, y_i) \in \mathcal{D}$, in addition to the original loss of the underlying recommendation model \mathcal{L}_o^i , we also designed a reconstruction loss for the CoT reasoning in the In-context CoT Examples to further strengthen the world-knowledge

and reasoning capabilities contained in the collaborative filtering features \mathbf{w}_i . The reconstruction loss is to minimize the distance between the predicted embedding and the ground-truth embedding of CoT reasoning:

$$\mathcal{L}_r^i = \frac{1}{K} \sum_{i=1}^K \left(1 - \frac{\mathbf{c}_i \cdot \mathbf{h}(\mathbf{r}_i)}{\|\mathbf{c}_i\| \|\mathbf{h}(\mathbf{r}_i)\|}\right),$$


where K is the number of In-context CoT Examples.

Finally, the total loss \mathcal{L} is computed as:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (\alpha \mathcal{L}_r^i + \mathcal{L}_o^i), \quad (6)$$

where α are hyperparameters that control the importance of the two parts of the loss, and N is the number of recommendation data.

5.6 Efficiency Analysis of Online Service

In the online service of LLM-CF, as shown in Figure 2, only the blue fire  module requires online computation. The ICT module’s inputs are short sequences that need minimal computational overhead. Experiments (§ 6.5) show that the number of ICT examples only needs $K = 4$, meaning the sequence length of ICT tokens \mathbf{T} is only a dozen or so. Furthermore, it can utilize other acceleration techniques for transformer decoders, which are currently being widely studied [23]. The computation time of the recommendation model is essentially the same as that of underlying models, and the inputs ICT token embedding for the ICT module \mathbf{r} and the final underlying model are only needed to compute once, requiring no additional computational overhead. Given the reduced necessity for maintaining an extensive In-Context CoT Dataset, the retrieval process for In-Context CoT Examples can be efficiently accelerated through the utilization of precomputed embedding vectors in conjunction with Approximate Nearest Neighbor search algorithms [28, 35].

The analysis of online service efficiency for LLM-CF, combined with the offline efficiency analysis in § 4.4, demonstrates the feasibility of our framework for deployment.

6 EXPERIMENT

In this section, we empirically verify the effectiveness of LLM-CF by extensive experiments. The experimental details and source codes can be found at (<https://github.com/Jeryi-Sun/LLM-CF>).

6.1 Experimental Setup

6.1.1 Datasets. We conducted experiments on three widely used datasets with varying domains, following existing work [5, 12, 27, 55, 56]. The Amazon¹ review datasets [15] are one of the most widely used benchmarks for recommendation. We adopt three subsets of them: *Sports & Outdoors*, *Beauty*, and *Toys & Games*. Following the common practices [12, 27, 55, 56], we use the reviews between January 1, 2019, to December 31, 2019 and treat all the recorded reviews as positive samples.

¹<http://jmcauley.ucsd.edu/data/amazon/>

6.1.2 Baselines. To validate the versatility of LLM-CF on the ranking phase, we selected representative CTR models as backbone models. CTR prediction is commonly applied in the ranking phase of recommendation systems, focusing on modeling the probability of users clicking on items through feature interactions. The selected six **backbone** models are as follows: **DeepFM** [13], **xDeepFM** [24], **AutoInt** [38], **DCNv1** [41], **DCNv2** [42], and **DIN** [54].

To validate the effectiveness of LLM-CF, we chose representative methods that leverage LLMs to enhance recommendation models as the competitive framework. Considering that our framework initially fine-tunes LLaMA2, and LLaMA2 performs well on CTR tasks, a natural idea is to employ knowledge distillation using fine-tuned LLaMA2 to enhance CTR models. Following the previous methods [18, 57], we implemented the **KD** framework, which utilizes the logits of fine-tuned LLaMA2 as soft labels for CTR models during training. We also compared our framework with a recent work, **KAR** [47], which utilizes the open-world knowledge of LLMs to enhance recommendations. For a fair comparison, we used the same LLM (i.e., LLaMA2-7B-chat) as other models in **KAR**.

6.1.3 Evaluation Metrics & Protocol. Following the common practices [12, 27, 55, 56], we use the *leave-one-out* strategy to partition the training, validation, and test datasets. We consider all recorded interactions in the datasets as positive samples and randomly sample a negative sample for each positive sample. For the evaluation metrics, we employ widely used *AUC* (Area under the ROC curve), *LogLoss* (binary cross-entropy loss), and *RelImpr* (the relative improvement of *AUC*) [54] for the CTR prediction task.

6.1.4 Implementation Details. As for the LLM, we utilized the widely used LLaMA2-7B-chat [40]. To finetune LLaMA2, we employed the previously mentioned strategy of mixing recommendation data and general data (i.e., **RecGen** in § 4.2). We use 8 Nvidia A800 GPUs to perform full-parameter SFT on the LLM. We adopt the DeepSpeed zero-2 strategy [34], where each GPU has a batch size of 16, resulting in a total batch size of 128 without gradient accumulation. We used the Adam optimizer and set the initial learning rate as $1e-5$ with the cosine learning rate schedule.

For the application of LLM-CF to backbone models, we trained them using the Adam optimizer with a learning rate of 0.001 and a batch size of 128. Additionally, we adopted early-stopped training to avoid over-fitting. The dimension of all feature embeddings is set to 32 for all methods. The number of layers and heads in the ICT module’s transformer decoder is set to 2. The number of In-context CoT Dataset M is $\frac{1}{10}$ of the Training Dataset. The number of In-context CoT examples K is set to 4. The α is tuned among $(0, 1]$ in the step of 0.1. Given the fact that LLM-CF and competitive framework are applied over backbone models, for a fair comparison, we maintain consistency of the backbone models across all frameworks. For instance, the backbone parts of LLM-CF (xDeepFM), KD(xDeepFM), and KAR(xDeepFM) share the same model architectures and basic hyper-parameters as xDeepFM.

6.2 Experimental Results

Table 2 reports the overall performance of LLM-CF over 6 backbone models on three real-world datasets. Based on the results presented in Table 2, we found that LLM-CF outperformed the corresponding

backbone models in terms of all evaluation metrics on real-world datasets, with statistical significance. The results verified the effectiveness of the LLM-CF, which can effectively integrate world-knowledge and reasoning-guided collaborative filtering features into the underlying recommendation models using the in-context chain of thought manner.

We also observed that LLM-CF achieved better results than other frameworks that leverage LLMs to enhance recommendation models (i.e., KAR and KD). Compared with LLM-CF, KAR overlooks the importance of collaborative filtering information, only considering current user and item features to enhance recommendation models, and does not use an LLM fine-tuned with recommendation data, leading to a significant performance drop in real-world in-domain datasets used in our experiments. The suboptimal performance of KD is mainly because the training objective of LLMs is Cross-Entropy for all tokens, while the underlying recommendation training mainly involves Binary Cross-Entropy loss, resulting in a misalignment of optimization goals. Moreover, LLM-CF is orthogonal to KD and could potentially be used in conjunction to improve recommendation performance further. The results verify that LLM-CF can better leverage LLMs to provide enhanced collaborative filtering information to underlying recommendation models.

6.3 Ablation Study

LLM-CF consists of several key operations, and to understand the effects of each operation, we conducted several ablation experiments for LLM-CF. They are: **(a) w/o text**: Remove the text features in recommendation features. **(b) w/o CoT**: Remove the CoT reasoning along with reconstruction loss in in-context CoT examples. **(c) w/o RecGen-LLaMA**: Replace the RecGen-LLaMA with the original LLaMA2-7B-chat to generate CoT reasoning. **(d) w/o TF decoder**: Replace the transformer decoder layers with mean pooling in the ICT module.

Table 4 presents the performance of various model variants over 6 backbone models on three real-world datasets. The results reveal several insights: (1) After removing CoT reasoning, there is a significant decline in recommendation performance, demonstrating the effectiveness of CoT reasoning in providing world knowledge and reasoning for collaborative filtering features. (2) Replacing RecGen-LLaMA with the original LLaMA2-7B-chat to generate CoT reasoning shows a downward trend in recommendation performance, highlighting the necessity of enhancing the recommendation capabilities of LLM when used to augment conventional recommendation models. (3) After replacing the transformer decoder layers with mean pooling in the ICT module, there is a noticeable decline in recommendation performance, underscoring the importance of using transformer decoder layers for learning in-context CoT examples. (4) Removing text features leads to a slight decline in performance, but the change is not significant, suggesting that text features have a limited role in the framework.

6.4 Application of Retrieval Tasks

LLM-CF is versatile, applicable not only to models in the ranking stage but also to models in the retrieval stage. In the preceding experiments, we validated the efficacy of LLM-CF in the CTR prediction

Table 2: Overall performance comparisons. The best and the second-best performance methods are denoted in bold and underlined fonts, respectively. * means improvements over the second-best methods are significant (p -value < 0.05). The RelImpr is calculated over backbone models, i.e., ‘None’ framework. KD, KAR, and LLM-CF are all based on the LLaMA2-7B.

Backbone	Framework	Sports			Beauty			Toys		
		AUC↑	Logloss↓	RelImpr↑	AUC↑	Logloss↓	RelImpr↑	AUC↑	Logloss↓	RelImpr↑
DeepFM	None	0.7990	0.5471	0.000%	0.7853	0.5545	0.000%	0.7681	0.5770	0.000%
	KD	<u>0.8043</u>	<u>0.5404</u>	1.773%	<u>0.7959</u>	<u>0.5442</u>	3.715%	<u>0.7713</u>	<u>0.5716</u>	1.194%
	KAR	0.7991	0.5469	0.033%	0.7870	0.5546	0.596%	0.7698	0.5718	0.634%
	LLM-CF	0.8137*	0.5306*	4.916%	0.8044*	0.5366*	6.695%	0.7881*	0.5581*	7.460%
xDeepFM	None	0.8158	0.5318	0.000%	0.8065	0.5359	0.000%	0.7836	0.5589	0.000%
	KD	<u>0.8169</u>	0.5298	0.348%	<u>0.8104</u>	0.5345	1.272%	0.7865	0.5553	1.023%
	KAR	0.8161	<u>0.5279</u>	0.094%	0.8101	<u>0.5315</u>	1.175%	<u>0.7898</u>	<u>0.5529</u>	2.186%
	LLM-CF	0.8196*	0.5248*	1.203%	0.8113	0.5311	1.566%	0.7947*	0.5473*	3.914%
AutoInt	None	0.8003	0.5444	0.000%	0.7949	0.5469	0.000%	0.7630	0.5770	0.000%
	KD	0.8012	0.5439	0.300%	<u>0.7961</u>	<u>0.5444</u>	0.407%	0.7635	0.5770	0.190%
	KAR	<u>0.8039</u>	0.5390	1.199%	0.7939	0.5476	-0.339%	<u>0.7683</u>	<u>0.5741</u>	2.015%
	LLM-CF	0.8088*	<u>0.5391</u>	2.831%	0.8090*	0.5321*	4.781%	0.7754*	0.5685*	4.714%
DCNv1	None	0.8023	0.5442	0.000%	0.8146	0.5255	0.000%	0.7621	0.5831	0.000%
	KD	<u>0.8040</u>	<u>0.5441</u>	0.562%	0.8147	0.5286	0.031%	<u>0.7652</u>	0.5847	1.183%
	KAR	0.8024	0.5469	0.033%	<u>0.8165</u>	<u>0.5229</u>	0.604%	0.7651	<u>0.5821</u>	1.144%
	LLM-CF	0.8092*	0.5368*	2.282%	0.8182*	0.5216*	1.144%	0.7702*	0.5745*	3.090%
DCNv2	None	0.8110	0.5331	0.000%	0.8028	0.5378	0.000%	0.7774	0.5650	0.000%
	KD	<u>0.8112</u>	<u>0.5320</u>	0.064%	0.8057	0.5343	0.958%	0.7827	0.5609	1.911%
	KAR	0.8087	0.5363	-0.739%	0.8003	0.5404	-0.825%	0.7759	0.5662	-0.541%
	LLM-CF	0.8131*	0.5307*	0.675%	<u>0.8033</u>	<u>0.5372</u>	0.165%	<u>0.7812</u>	<u>0.5619</u>	1.370%
DIN	None	0.7986	0.5519	0.000%	0.7861	0.5613	0.000%	0.7586	0.5885	0.000%
	KD	<u>0.8023</u>	<u>0.5422</u>	1.239%	<u>0.7934</u>	<u>0.5518</u>	2.551%	<u>0.7652</u>	<u>0.5847</u>	2.552%
	KAR	0.7971	0.5525	-0.502%	0.7861	0.5604	0.000%	0.7620	0.5874	1.315%
	LLM-CF	0.8089*	0.5374*	3.449%	0.7967*	0.5492*	3.705%	0.7783*	0.5699*	7.618%

Table 3: Overall performance of LLM-CF on retrieval tasks. The best performance methods are denoted in bold.

Models	Beauty				Sports				Toys			
	HIT@5	HIT@10	NDCG@5	NDCG@10	HIT@5	HIT@10	NDCG@5	NDCG@10	HIT@5	HIT@10	NDCG@5	NDCG@10
SASREC	0.0379	0.0648	0.0206	0.0293	0.0161	0.0203	0.0088	0.0101	0.0329	0.0646	0.0143	0.0251
+LLM-CF	0.0395	0.0652	0.0210	0.0293	0.0205	0.0349	0.0107	0.0154	0.0466	0.0723	0.0238	0.0321
YoutubeDNN	0.0150	0.0272	0.0089	0.0128	0.0109	0.0177	0.0069	0.0091	0.0127	0.0229	0.0077	0.0110
+LLM-CF	0.0160	0.0291	0.0099	0.0139	0.0112	0.0191	0.0069	0.0094	0.0165	0.0285	0.0099	0.0138
GRU4REC	0.0279	0.0470	0.0180	0.0241	0.0165	0.0268	0.0102	0.0135	0.0257	0.0411	0.0152	0.0202
+LLM-CF	0.0303	0.0504	0.0186	0.0250	0.0203	0.0317	0.0127	0.0164	0.0220	0.0380	0.0139	0.0192
SRGNN	0.0199	0.0341	0.0123	0.0169	0.0085	0.0146	0.0054	0.0073	0.0158	0.0284	0.0096	0.0137
+LLM-CF	0.0214	0.0359	0.0128	0.0174	0.0118	0.0210	0.0072	0.0101	0.0177	0.0302	0.0104	0.0144

task. In this section, we employed LLM-CF over retrieval models (SASREC [21], YoutubeDNN [6], GRU4REC [17] and SRGNN [46]) and experimentally validated its effectiveness on four ranking metrics. We trained the models using the sampled softmax loss, with 128 randomly sampled negatives for each interaction. Since these models use inner product to estimate probabilities, when applying LLM-CF, we fused the vectors obtained from LLM-CF with the user vector in the backbone model by element-wise addition. To prevent data leakage, for the retrieval task, we mask the information of the target item in the current interaction in the in-context chain of thought learning. Table 3 reports the overall performance of LLM-CF over 4 backbone models on three real-world datasets. LLM-CF

enhances the performance of backbone retrieval models, further validating the versatility of the collaborative filtering features enhanced by world knowledge and reasoning.

6.5 Effects of In-context CoT Examples

Considering the length of In-context CoT (ICT) examples and the ratio of positive to negative examples has an important impact on the final prediction results. We examine the influence of the length of ICT examples and the effect of relaxing positive/negative (P/N) example ratio constraints on the performance of the LLM-CF across three real-world datasets and six backbone models.

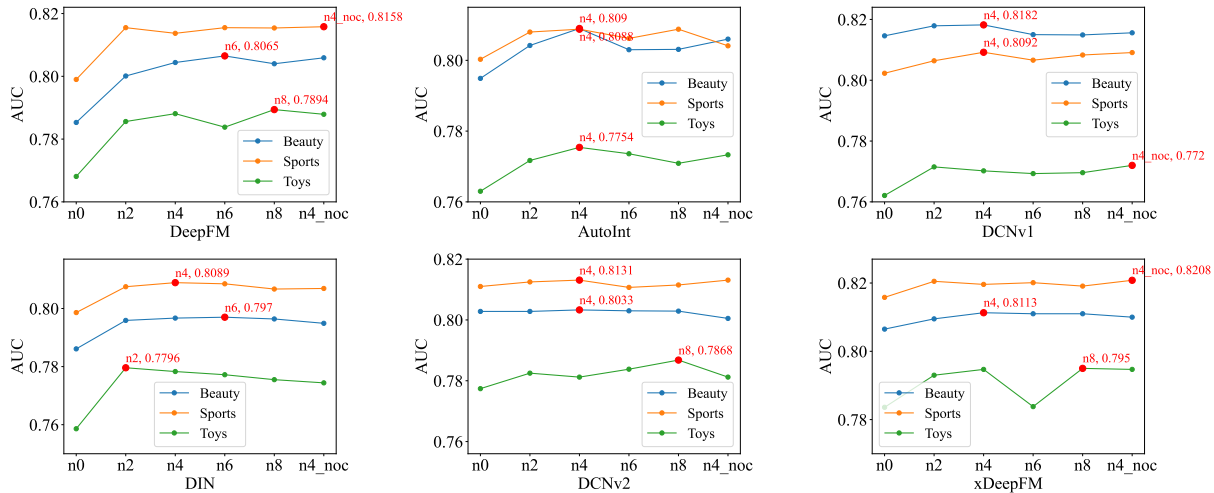


Figure 4: The recommendation performance of LLM-CF on 6 backbone models across three real-world datasets as it varies with the change in In-context CoT examples length and positive/negative example constraints. The red dots indicate the best-performing setting for the current backbone model and dataset.

Table 4: Ablation Study of LLM-CF. Boldface indicates best.

Backbone	Framework	Sports		Beauty		Toys	
		AUC \uparrow	Logloss \downarrow	AUC \uparrow	Logloss \downarrow	AUC \uparrow	Logloss \downarrow
DeepFM	None	0.7990	0.5471	0.7853	0.5545	0.7681	0.5770
	w/o text	0.8137	0.5308	0.8030	0.5382	0.7880	0.5583
	w/o CoT	0.8009	0.5448	0.7980	0.5423	0.7766	0.5662
	w/o RecGen-LLaMA	0.8132	0.5301	0.7973	0.5452	0.7806	0.5640
	w/o TF decoder	0.8071	0.5378	0.7986	0.5411	0.7739	0.5682
	LLM-CF	0.8137	0.5306	0.8044	0.5366	0.7881	0.5581
xDeepFM	None	0.8158	0.5318	0.8065	0.5356	0.7836	0.5589
	w/o text	0.8200	0.5248	0.8088	0.5308	0.7898	0.5547
	w/o CoT	0.8177	0.5290	0.8081	0.5349	0.7915	0.5561
	w/o RecGen-LLaMA	0.8191	0.5246	0.8111	0.5311	0.7936	0.5490
	w/o TF decoder	0.8192	0.5246	0.8092	0.5349	0.7912	0.5571
	LLM-CF	0.8196	0.5248	0.8113	0.5311	0.7947	0.5473
AutoInt	None	0.8003	0.5444	0.7949	0.5468	0.7630	0.5770
	w/o text	0.8071	0.5380	0.8043	0.5385	0.7710	0.5725
	w/o CoT	0.7992	0.5468	0.7940	0.5490	0.7581	0.5812
	w/o RecGen-LLaMA	0.8041	0.5398	0.7993	0.5412	0.7706	0.5702
	w/o TF decoder	0.8017	0.5420	0.7942	0.5474	0.7592	0.5816
	LLM-CF	0.8088	0.5391	0.8090	0.5321	0.7754	0.5685
DcNv1	None	0.8023	0.5442	0.8146	0.5255	0.7621	0.5831
	w/o text	0.8099	0.5408	0.8155	0.5338	0.7692	0.5751
	w/o CoT	0.8017	0.5441	0.7622	0.5802	0.7622	0.5802
	w/o RecGen-LLaMA	0.8036	0.5432	0.7891	0.5617	0.7720	0.5733
	w/o TF decoder	0.8006	0.5465	0.8118	0.5294	0.7598	0.5858
	LLM-CF	0.8092	0.5368	0.8182	0.5216	0.7702	0.5745
DcNv2	None	0.8110	0.5331	0.8028	0.5378	0.7774	0.5650
	w/o text	0.8130	0.5297	0.7971	0.5447	0.7790	0.5628
	w/o CoT	0.8081	0.5365	0.7968	0.5457	0.7753	0.5681
	w/o RecGen-LLaMA	0.8117	0.5313	0.8023	0.5381	0.7826	0.5602
	w/o TF decoder	0.8113	0.5314	0.7978	0.5437	0.7810	0.5612
	LLM-CF	0.8131	0.5307	0.8033	0.5372	0.7812	0.5619
DIN	None	0.7986	0.5519	0.7861	0.5613	0.7586	0.5885
	w/o text	0.8053	0.5415	0.7938	0.5509	0.7742	0.5731
	w/o CoT	0.8025	0.5445	0.7907	0.5557	0.7718	0.5735
	w/o RecGen-LLaMA	0.8057	0.5410	0.7953	0.5571	0.7783	0.5737
	w/o TF decoder	0.8004	0.5484	0.7850	0.5602	0.7666	0.5794
	LLM-CF	0.8089	0.5374	0.7967	0.5492	0.7783	0.5699

As shown in Figure 4, ICT lengths were varied across 0 (n0), 2 (n2), 4 (n4), 6 (n6), and 8 (n8) with a consistent P/N sample ratio. Our empirical results reveal that n4 consistently outperforms n2, while lengths of n6 and n8 show variable performance, occasionally surpassing but generally not exceeding n4. A notable performance decline for most datasets and models at n8 suggests a diminishing

return with longer ICT, posing n4 as an effective balance between performance gains and computational efficiency.

Furthermore, the removal of P/N ratio constraints (denoted as n4_noc) shown in Figure 4 did not significantly deviate from the performance of the constrained scenarios. While some datasets showed marginal improvements without P/N constraints, others exhibited a decrease. Maintaining an equal P/N sample ratio effectively validated the superiority of LLM-CF, reducing potential biases arising from disproportionate label distributions.

7 CONCLUSION

In this paper, we proposed the Large Language Model enhanced Collaborative Filtering (LLM-CF) Framework, which integrates the world knowledge and reasoning capabilities of Large Language Models into Recommender Systems. LLM-CF leverages LLMs' world knowledge and reasoning, particularly through the In-context Chain of Thought module, to enhance collaborative filtering in RSs. The key contribution of LLM-CF is its effective distillation of LLM capabilities into RSs, balancing recommendation accuracy with operational efficiency. Our experiments across various datasets demonstrate that LLM-CF significantly outperforms conventional recommendation models in both ranking and retrieval tasks, confirming the benefits of integrating LLMs into RSs using LLM-CF.

ACKNOWLEDGMENTS

This work was funded by the National Key R&D Program of China (2023YFA1008704), the National Natural Science Foundation of China (No. 62377044), Beijing Key Laboratory of Big Data Management and Analysis Methods, Major Innovation & Planning Interdisciplinary Platform for the "Double-First Class" Initiative, funds for building world-class universities (disciplines) of Renmin University of China, and PCC@RUC. Supported by Kuaishou Technology. Supported by the Outstanding Innovative Talents Cultivation Funded Programs 2024 of Renmin University of China.

REFERENCES

- [1] Armen Aghajanyan, Akshat Shrivastava, Anchit Gupta, Naman Goyal, Luke Zettlemoyer, and S. Gupta. 2020. Better Fine-Tuning by Reducing Representational Collapse. *ArXiv abs/2008.03156* (2020).
- [2] 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. *Proceedings of the 17th ACM Conference on Recommender Systems* (2023).
- [3] Jianxin Chang, Chenbin Zhang, Zhiyi Fu, Xiaoxue Zang, Lin Guan, Jing Lu, Yiqun Hui, Dewei Leng, Yanan Niu, Yang Song, and Kun Gai. 2023. TWIN: Two-stage Interest Network for Lifelong User Behavior Modeling in CTR Prediction at Kuaishou. *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023).
- [4] Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2023. A survey of chain of thought reasoning: Advances, frontiers and future. *arXiv preprint arXiv:2309.15402* (2023).
- [5] Zhixuan Chu, Hongyan Hao, Xin Ouyang, Simeng Wang, Yan Wang, Yue Shen, Jinjie Gu, Qing Cui, Longfei Li, Siqiao Xue, James Y Zhang, and Sheng Li. 2023. Leveraging Large Language Models for Pre-trained Recommender Systems. *arXiv:2308.10837* [cs.LG].
- [6] Paul Covington, Jay K. Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. *Proceedings of the 10th ACM Conference on Recommender Systems* (2016).
- [7] Yu Cui, Feng Liu, Pengbo Wang, Bohao Wang, Heng Tang, Yi Wan, Jun Wang, and Jiawei Chen. 2024. Distillation Matters: Empowering Sequential Recommenders to Match the Performance of Large Language Model. <https://api.semanticscholar.org/CorpusID:269484474>
- [8] Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui, and Furu Wei. 2023. Why Can GPT Learn In-Context? Language Models Secretly Perform Gradient Descent as Meta-Optimizers. In *Findings of the Association for Computational Linguistics: ACL 2023*. Association for Computational Linguistics, Toronto, Canada, 4005–4019.
- [9] Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. 2023. Uncovering ChatGPT’s Capabilities in Recommender Systems. *Proceedings of the 17th ACM Conference on Recommender Systems* (2023). <https://api.semanticscholar.org/CorpusID:258461170>
- [10] Guanting Dong, Hongyi Yuan, Keming Lu, Chengpeng Li, Mingfeng Xue, Dayiheng Liu, Wei Wang, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2023. How Abilities in Large Language Models are Affected by Supervised Fine-tuning Data Composition. *ArXiv abs/2310.05492* (2023).
- [11] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234* (2022).
- [12] 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* (Seattle, WA, USA) (*RecSys '22*). Association for Computing Machinery, New York, NY, USA, 299–315. <https://doi.org/10.1145/3523227.3546767>
- [13] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine Based Neural Network for CTR Prediction. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence* (Melbourne, Australia) (*IJCAI'17*). AAAI Press, 1725–1731.
- [14] Jesse Harte, Wouter Zorgdrager, Panos Louridas, Asterios Katsifodimos, D. Janach, and Marios Fragkoulis. 2023. Leveraging Large Language Models for Sequential Recommendation. *Proceedings of the 17th ACM Conference on Recommender Systems* (2023). <https://api.semanticscholar.org/CorpusID:261823711>
- [15] Ruining He and Julian McAuley. 2016. Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering. In *Proceedings of the 25th International Conference on World Wide Web* (Montréal, Québec, Canada) (*WWW '16*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 507–517. <https://doi.org/10.1145/2872427.2883037>
- [16] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300* (2020).
- [17] Balázs Hidasi and Alexandros Karatzoglou. 2018. Recurrent Neural Networks with Top-k Gains for Session-based Recommendations. In *CIKM*. ACM, 843–852.
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. *arXiv:1503.02531* [stat.ML].
- [19] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*. OpenReview.net.
- [20] Neel Jain, Ping-yeh Chiang, Yuxin Wen, John Kirchenbauer, Hong-Min Chu, Gowthami Somepalli, Brian R Bartoldson, Bhavya Kailkhura, Avi Schwarzschild, Aniruddha Saha, et al. 2023. NEFTune: Noisy Embeddings Improve Instruction Finetuning. *arXiv preprint arXiv:2310.05914* (2023).
- [21] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [22] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed H. Chi, and Derek Zhiyuan Cheng. 2023. Do LLMs Understand User Preferences? Evaluating LLMs On User Rating Prediction. *ArXiv abs/2305.06474* (2023). <https://api.semanticscholar.org/CorpusID:258615591>
- [23] Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2019. Reformer: The Efficient Transformer. In *International Conference on Learning Representations*.
- [24] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. XDeepFM: Combining Explicit and Implicit Feature Interactions for Recommender Systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (London, United Kingdom) (*KDD '18*). Association for Computing Machinery, New York, NY, USA, 1754–1763. <https://doi.org/10.1145/3219819.3220023>
- [25] Jianghao Lin, Rong Shan, Chenxu Zhu, Koujianhua Du, Bo Chen, Shigang Quan, Ruiming Tang, Yong Yu, and Weinan Zhang. 2023. ReLLa: Retrieval-enhanced Large Language Models for Lifelong Sequential Behavior Comprehension in Recommendation. *arXiv preprint arXiv:2308.11131* (2023).
- [26] Qijiong Liu, Nuo Chen, Tetsuya Sakai, and Xiao-Ming Wu. 2023. A First Look at LLM-Powered Generative News Recommendation. *arXiv preprint arXiv:2305.06566* (2023).
- [27] Kai Mei and Yongfeng Zhang. 2023. LightLM: A Lightweight Deep and Narrow Language Model for Generative Recommendation. *arXiv:2310.17488* (2023).
- [28] Marius Muja and David G. Lowe. 2009. Fast Approximate Nearest Neighbors with Automatic Algorithm Configuration. In *International Conference on Computer Vision Theory and Applications*.
- [29] Sheshera Mysore, Andrew McCallum, and Hamed Zamani. 2023. Large Language Model Augmented Narrative Driven Recommendations. *Proceedings of the 17th ACM Conference on Recommender Systems* (2023).
- [30] OpenAI. 2023. GPT-4 Technical Report. *ArXiv abs/2303.08774* (2023). <https://api.semanticscholar.org/CorpusID:257532815>
- [31] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277* (2023).
- [32] Qi Pi, Xiaoqiang Zhu, Guorui Zhou, Yujing Zhang, Zhe Wang, Lejian Ren, Ying Fan, and Kun Gai. 2020. Search-based User Interest Modeling with Lifelong Sequential Behavior Data for Click-Through Rate Prediction. *Proceedings of the 29th ACM International Conference on Information & Knowledge Management* (2020).
- [33] Jiarui Qin, Weinan Zhang, Rong Su, Zhirong Liu, Weiwen Liu, Ruiming Tang, Xiuqiang He, and Yong Yu. 2021. Retrieval & Interaction Machine for Tabular Data Prediction. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (2021).
- [34] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–16.
- [35] Jie Ren, Minjia Zhang, and Dong Li. 2020. HM-ANN: Efficient Billion-Point Nearest Neighbor Search on Heterogeneous Memory. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) (*NIPS'20*). Curran Associates Inc., Red Hook, NY, USA, Article 895, 13 pages.
- [36] Ruifeng Ren and Yong Liu. 2023. In-context Learning with Transformer Is Really Equivalent to a Contrastive Learning Pattern. *arXiv preprint arXiv:2310.13220* (2023).
- [37] 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*. 890–896.
- [38] Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. 2019. AutoInt: Automatic Feature Interaction Learning via Self-Attentive Neural Networks. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (Beijing, China) (*CIKM '19*). Association for Computing Machinery, New York, NY, USA, 1161–1170. <https://doi.org/10.1145/3357384.3357925>
- [39] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).
- [40] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutli Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [41] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. 2017. Deep & Cross Network for Ad Click Predictions. In *Proceedings of the ADKDD '17* (Halifax, NS, Canada) (*ADKDD '17*). Association for Computing Machinery, New York, NY, USA, Article 12, 7 pages. <https://doi.org/10.1145/3124749.3124754>

- [42] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. 2021. DCN V2: Improved Deep & Cross Network and Practical Lessons for Web-Scale Learning to Rank Systems. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (*WWW '21*). Association for Computing Machinery, New York, NY, USA, 1785–1797. <https://doi.org/10.1145/3442381.3450078>
- [43] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* (2022).
- [44] Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2023. LLMRec: Large Language Models with Graph Augmentation for Recommendation. *ArXiv* (2023).
- [45] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. 2022. Robust Fine-Tuning of Zero-Shot Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7959–7971.
- [46] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-Based Recommendation with Graph Neural Networks. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 01 (July 2019), 346–353. <https://doi.org/10.1609/aaai.v33i01.3301346>
- [47] Yunjia Xi, Weiwen Liu, Jianghao Lin, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. 2023. Towards Open-World Recommendation with Knowledge Augmentation from Large Language Models. *arXiv preprint arXiv:2306.10933* (2023).
- [48] Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighoff. 2023. C-Pack: Packaged Resources To Advance General Chinese Embedding. *arXiv:2309.07597* [cs.CL]
- [49] Bin Yin, Jun Xie, Yu-Xia Qin, Zixiang Ding, Zhichao Feng, Xiang Li, and Wei Lin. 2023. Heterogeneous Knowledge Fusion: A Novel Approach for Personalized Recommendation via LLM. *Proceedings of the 17th ACM Conference on Recommender Systems* (2023). <https://api.semanticscholar.org/CorpusID:260682982>
- [50] Hongyi Yuan, Zheng Yuan, Chuanqi Tan, Fei Huang, and Songfang Huang. 2022. HyPe: Better Pre-trained Language Model Fine-tuning with Hidden Representation Perturbation. In *Annual Meeting of the Association for Computational Linguistics*.
- [51] Zhenrui Yue, Sara Rabhi, Gabriel De Souza Pereira Moreira, Dong Wang, and Even Oldridge. 2023. LamaRec: Two-Stage Recommendation using Large Language Models for Ranking. *ArXiv abs/2311.02089* (2023). <https://api.semanticscholar.org/CorpusID:265033634>
- [52] Zhi Zheng, Zhaopeng Qiu, Xiao Hu, Likang Wu, Hengshu Zhu, and Hui Xiong. 2023. Generative job recommendations with large language model. *arXiv preprint arXiv:2307.02157* (2023).
- [53] Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206* (2023).
- [54] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep Interest Network for Click-Through Rate Prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (London, United Kingdom) (KDD '18)*. Association for Computing Machinery, New York, NY, USA, 1059–1068. <https://doi.org/10.1145/3219819.3219823>
- [55] 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 *CIKM '20: The 29th ACM International Conference on Information and Knowledge Management, Virtual Event, Ireland, October 19-23, 2020*. ACM, 1893–1902.
- [56] Kun Zhou, Hui Yu, Wayne Xin Zhao, and Ji-Rong Wen. 2022. Filter-Enhanced MLP is All You Need for Sequential Recommendation. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) (*WWW '22*). Association for Computing Machinery, New York, NY, USA, 2388–2399. <https://doi.org/10.1145/3485447.3512111>
- [57] Jieming Zhu, Jinyang Liu, Weiqi Li, Jincai Lai, Xiuqiang He, Liang Chen, and Zhibin Zheng. 2020. Ensembled CTR Prediction via Knowledge Distillation. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (Virtual Event, Ireland) (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 2941–2958. <https://doi.org/10.1145/3340531.3412704>