

Collaborative Cross-modal Fusion with Large Language Model for Recommendation

Zhongzhou Liu
Singapore Management University
Singapore
zzliu.2020@phdcs.smu.edu.sg

Kuicai Dong
Nanyang Technological University
Singapore
kuicai001@e.ntu.edu.sg

Hao Zhang
Nanyang Technological University
Singapore
hzhzhang26@outlook.com

Yuan Fang
Singapore Management University
Singapore
yfang@smu.edu.sg

Abstract

Despite the success of conventional collaborative filtering (CF) approaches for recommendation systems, they exhibit limitations in leveraging semantic knowledge within the textual attributes of users and items. Recent focus on the application of large language models for recommendation (LLM4Rec) has highlighted their capability for effective semantic knowledge capture. However, these methods often overlook the collaborative signals in user behaviors. Some simply instruct-tune a language model, while others directly inject the embeddings of a CF-based model, lacking a synergistic fusion of different modalities. To address these issues, we propose a framework of Collaborative Cross-modal Fusion with Large Language Models, termed CCF-LLM, for recommendation. In this framework, we translate the user-item interactions into a hybrid prompt to encode both semantic knowledge and collaborative signals, and then employ an attentive cross-modal fusion strategy to effectively fuse latent embeddings of both modalities. Extensive experiments demonstrate that CCF-LLM outperforms existing methods by effectively utilizing semantic and collaborative signals in the LLM4Rec context.

CCS Concepts

• Information systems → Users and interactive retrieval; Collaborative filtering; • Applied computing;

Keywords

Large Language Models; Recommendation Systems; Cross-modal; Collaborative Filtering

ACM Reference Format:

Zhongzhou Liu, Hao Zhang, Kuicai Dong, and Yuan Fang. 2024. Collaborative Cross-modal Fusion with Large Language Model for Recommendation. 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, 10 pages. <https://doi.org/10.1145/3627673.3679596>

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

© 2024 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA, <https://doi.org/10.1145/3627673.3679596>.

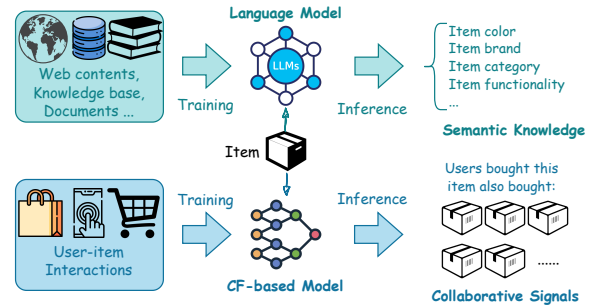


Figure 1: An illustration of heterogeneous characteristics between the semantic knowledge from LLMs and the collaborative signals from conventional recommendation systems.

1 Introduction

Collaborative filtering (CF)-based recommendation systems, which aim to learn users' preferences from historical user-item interactions, have demonstrated significant success across multiple domains [3, 10, 15, 16, 24, 36]. Their success is attributed to effectively modeling the collaborative signal that encapsulates similarities among user-user, item-item, and user-item co-occurrences [5, 40]. However, traditional CF models still struggle to process the rich semantic knowledge in users' and items' textual features [34], urging the development of more advanced models with semantic awareness. Recently, large language models (LLMs) [1, 41, 48] have demonstrated their remarkable capabilities in many tasks, given their strong capacity for assimilating human knowledge about society and the physical world. Motivated by the power of LLMs, numerous researchers are exploring their potential in recommendation systems, referred to as LLM4Rec [21, 31, 32].

Initial efforts [17, 19, 22, 33, 43, 49] try to transform user-item interactions into natural language sequences and then instruct LLMs to conduct zero-shot recommendations. However, these approaches usually underperform conventional CF-based systems, e.g., standard matrix factorization (MF) models [26]. Although LLMs are proficient in grasping the semantic knowledge of users' or items' textual attributes, solely relying on semantic relatedness is inadequate for modeling user preferences. This limitation stems from the focus of LLMs on semantic similarities, overlooking the collaborative signals.

Taking the well-known “*Beer and Diapers*” story as an example¹, beer and diapers indeed are semantically unrelated products but often co-occur in a single transaction [18], where such co-occurrence relationship can be easily modeled via collaborative signals. Thus, we posit that the collaborative signal is essential for LLM4Rec. In its absence, LLMs cannot effectively leverage and exploit historical user-item interactions solely from textual descriptions, leading to inaccurate user preference modeling.

Figure 1 illustrates the heterogeneous characteristics that diverge between semantic knowledge and collaborative signals. LLMs are typically trained to capture the semantic knowledge among users and items based on their textual attributes, whereas CF-based systems are trained to model co-occurring correlations based on the patterns of user behaviors. Due to their different training data and objectives, collaborative signals can be exploited to complement semantic knowledge. As a result, recent studies [2, 6, 12, 29, 50] have explored the incorporation of collaborative signals into LLMs through two primary approaches: (i) transforming user-item interactions into natural language descriptions, or (ii) directly integrating collaborative embeddings into LLMs.

(i) *Collaborative Signals in Natural Language Descriptions*. Some recent work has attempted to integrate the collaborative signals into LLMs by converting them into natural language descriptions. These approaches [2, 6, 12] train or use instruct-tuning on LLMs with the language descriptions of collaborative signals, which implicitly capture such signals during the training process. Specifically, each instance of user-item interaction is processed as a sentence, e.g., “Will user u like item i ?”, followed by a “Yes” or “No” response as the optimization objective. However, the improvement of these methods is marginal compared to CF-based methods. The primary reason is that the user-item interactions, as described in unstructured natural language, cannot effectively instruct the LLM to understand the non-linear, high-order correlation concealed in the huge and sparse co-occurrence logs. In contrast, CF-based systems explicitly model the correlation in a structured manner.

(ii) *Collaborative Signals in Embeddings*. Some approaches attempt to directly insert latent representations of the collaborative signals into the prompts [29, 50]. For instance, both LLaRA [29] and CoLLM [50] represent each item as $[\text{text}_i][\text{emb}_i]$, where $[\text{text}_i]$ denotes the textual attributes of item i (e.g., “Titanic”), and $[\text{emb}_i]$ is to be replaced by the item embedding extracted from a CF-based model such as LightGCN [15]. That is, these approaches directly inject the embeddings of collaborative signals into the input sequence of LLMs along with the semantic token embeddings. However, LLMs may not effectively understand the collaborative signals in this form, as the semantic and CF embeddings are from two disparate modalities and reside in two distinct spaces. Consequently, such a naïve injection leads to only a marginal improvement.

In this paper, we argue that current LLM4Rec methods do not fully leverage the collaborative signals for recommendation due to two major challenges: (1) *How to design an input template to assist LLMs in effectively assimilating the collaborative signals*; and

(2) *how to synergistically align and fuse two different modalities capturing semantic knowledge and collaborative signals*? To tackle these challenges, we propose a framework of Collaborative Cross-modal Fusion with Large Language Models for recommendation (CCF-LLM) that can adaptively fuse semantic knowledge with collaborative signals. To mitigate the first challenge, we propose a hybrid prompt translation step to translate user-item interactions into a prompt sequence that encodes collaborative signals and semantic knowledge. Though existing approaches [29, 50] also encode the two modalities within a prompt, they use two separate tokens for each item to capture the semantic knowledge and collaborative signals, respectively, decoupling the two modalities. On the contrary, we employ a *single* token to capture both modalities simultaneously, paving the way for the subsequent multi-modal fusion. To address the second challenge, we propose an attentive cross-modal fusion strategy to fuse information from the two modalities. Inspired by the cross-gate mechanism [39], we propose a GATE network to fuse the two modalities effectively. With the GATE network, the fusion process can be optimized adaptively in a finer dimension-wise granularity with more flexibility. Compared to previous LLM4Rec systems, CCF-LLM makes better use of collaborative signals to achieve optimal recommendation results.

In summary, we highlight our contributions as stated below. (1) We underscore the significance of integrating both semantic and collaborative signals into LLM4Rec. In particular, in Section 4.2 we empirically illustrate the limitations of capturing collaborative signals using solely natural language descriptions. (2) We propose CCF-LLM, a novel framework that enhances the current LLM4Rec paradigm through translating user-item interactions into a hybrid prompt sequence, and further mapping the prompt into the LLM and CF embedding spaces for cross-modal fusion. (3) We conduct extensive experiments on two public datasets. The experimental results demonstrate the importance of integrating collaborative signals and the effectiveness of CCF-LLM.

2 Related Work

In this section, we provide a literature review pertaining to CF-based recommendation systems and LLM4Rec approaches. Our work is inspired by them in combining CF and LLM embeddings for cross-modal fusion.

CF-based Recommendation System. The conventional scheme of collaborative filtering (CF)-based recommendation systems involves extracting the user-user and item-item similarities from users’ historical interactions and predicting relatedness scores between user-item pairs [3, 10, 15, 16, 24, 36]. The fundamental assumption of collaborative filtering is that if two users have similar ratings or behaviors on the same items, they will also have similar ratings or behaviors on other items [40]. Though CF has achieved great successes in various domains, the conventional CF-based models still face certain limitations. One of the key issues is that they ignore the rich semantic knowledge within users’ and items’ textual attributes, which significantly limits the generalizability and recommendation capacity of these models [38]. To overcome this limitation, many works incorporate auxiliary information into their modeling process, such as social network [25, 46], comments and reviews [35], multi-behavior interactions [7] and knowledge graphs [44]. While

¹We use this classical example only for illustrating the limitations of LLMs. We will present a more detailed case study on real-world data in Section 4.6.

external knowledge can provide some help, the capacity of these models is restricted by the availability of training data, lacking open-world knowledge and reasoning capability. Moreover, not all types of external knowledge can be generalized to different domains.

LLM4Rec Approaches. Recently there has been a growing trend in utilizing large language models (LLMs) for recommendation. Having been pre-trained on a large number of corpora with billions of parameters, LLMs can potentially answer queries for recommendation directly. It has been demonstrated that LLMs can learn to solve unseen tasks through just a few carefully designed instructions, using their inherent reasoning strength and open-world knowledge [4]. They usually take queries and target users' historical records in natural languages as input and output the recommendation results [17, 19, 22, 33, 43, 49]. Note that due to the gap between the pre-training objective and downstream recommendation task, the recommendation performance is often unsatisfactory under zero-shot settings [30]. To solve this problem, a popular trend is to integrate collaborative signals into LLMs to enhance their performance. One way is to conduct instruct-tuning on a pre-trained language model on recommendation tasks [2, 6, 9, 12, 47]. In their tuning paradigm, each prompt consists of a prompt input and a prompt output. They fed the prompt input to an LLM and optimized it by minimizing the loss between the generated output and prompt output. The prompt input usually contains a fixed task description and personalized input data for a user. The input data can be the user's historical behavior [2, 6, 6, 9] or knowledge related to the user's behavior [47]. Other than fine-tuning, P5 [12] trains a BERT model [11] from scratch to let it adapt to multiple recommendation tasks, which performs well in certain conditions.

Though the above methods have demonstrated the capabilities of LLMs for recommendation, they failed to explicitly leverage the collaborative signals, which are usually represented by latent user/item embeddings in a CF-based model. Only very few methods have considered this issue [28, 29, 50], and tried to feed both CF and LLM embeddings as input to an LLM. What they did was simply inject the CF embeddings at the LLM's input embedding sequence. To realize the injection, they first align the CF embeddings to the identical dimension as LLM embeddings and then map the predefined placeholders to the aligned CF embeddings. However, the naïve injection would suffer from an inconsistency between the semantic and collaborative signals, resulting in negative impact on the recommendation performance. In contrast, our work proposes to fuse the two modalities adaptively to reconcile the inconsistency, which helps LLMs fully leverage and integrate collaborative signals for recommendation.

3 Methodology

The overall framework of CCF-LLM is illustrated in Figure 2. Firstly, it translates the user-item interactions into hybrid prompt sequences (as shown in Figure 2a), which encodes heterogeneous modalities from collaborative signals and semantic knowledge. Subsequently, it utilizes an attentive cross-modal fusion strategy (as shown in Figure 2b) to fuse these heterogeneous modalities. This includes a mapping phase to map the tokens in hybrid prompt to LLM or CF embeddings accordingly (as shown in Figure 2b-1), and a fusion phase to adaptively align and fuse the heterogeneous embeddings

via a GATE network (as shown in Figure 2b-2). The fully fused latent embeddings are integrated with an LLM for recommendation.

In Section 3.1, we elaborate the preliminaries for the recommendation task. In Section 3.2, we explain our hybrid prompt translation step. Following this, we explain the attentive cross-modal strategy in Section 3.3, and summarize the novelty of CCF-LLM in terms of the motivation and techniques, compared to existing works. Finally, we present the training steps in Section 3.4.

3.1 Preliminaries

Task Formulation. We address the click-through rate (CTR) prediction task. CTR prediction is the task of predicting the likelihood of a person clicking on an advertisement or item. The CTR prediction is typically formulated as a supervised binary classification task. This usually involves a set of users \mathcal{U} and items \mathcal{I} . Each instance of user-item interaction is represented as a triplet (u, i, y) , where $u \in \mathcal{U}$ represents a user, $i \in \mathcal{I}$ denotes an item, and $y \in \{0, 1\}$ indicates the groundtruth interaction status (i.e., clicked or not clicked). Given the feature x_u of the user u and feature x_i of the item i , the goal of a CTR model is to accurately fit and predict the posterior probability $P(y|x_u, x_i)$.

Conventional CF-based recommendation. A conventional CF-based recommendation system aims to predict the interaction status given the l -dimensional embeddings representing the collaborative signals $x_u^{CF} \in \mathbb{R}^l$ for user u and $x_i^{CF} \in \mathbb{R}^l$ for item i . In most conventional CF-based recommendation models like NCF [16] and LightGCN [15], the embeddings x_u^{CF} and x_i^{CF} are derived from the encoding network ENC with the trainable parameters $\Theta_E = \{\Theta_E^U, \Theta_E^I\}$, as follows.

$$x_u^{CF} = \text{ENC}(u; \Theta_E^U), \quad x_i^{CF} = \text{ENC}(i; \Theta_E^I). \quad (1)$$

Then the encoded user and item embeddings are fed into an interaction network INT to make predictions with $\hat{y} = \text{INT}(x_u^{CF}, x_i^{CF}; \Theta_T)$, where Θ_T is the set of trainable parameters of the interaction network. In some models, INT can be a non-parametric function, i.e., $\Theta_T = \emptyset$. The model is optimized by minimizing the cross-entropy loss function that measures the prediction error between prediction \hat{y} and ground-truth label y on the training data. When conducting recommendation with an LLM, however, we need to convert the interaction history between u and i into a hybrid prompt $p_{u,i}$ via the hybrid prompt translation step (demonstrated in Section 3.2), and make predictions with $\hat{y} = M(p_{u,i}; \Theta)$, where M is the LLM4Rec model parameterized by Θ .

3.2 Hybrid Prompt Translation

To work with LLMs, we need to translate the user-item interaction triplets $\{(u, i, y)\}$ into a prompt sequence. The purpose is twofold: we aim to encode the task description and user-item interactions into natural language, and we also seek to encode the collaborative signals in the form of latent embeddings.

Our prompt template is designed to be similar to the templates used in previous studies [2, 50]. Specifically, we formulate two types of item feature $x_i = \{x_i^{CF}, x_i^{SM}\}$ via a hybrid prompt, where $x_i^{SM} = \{x_i^{SM}[1], x_i^{SM}[2], \dots, x_i^{SM}[T_i]\}$ is the set of embedding representing the semantic knowledge of item's textual attributes. T_i is

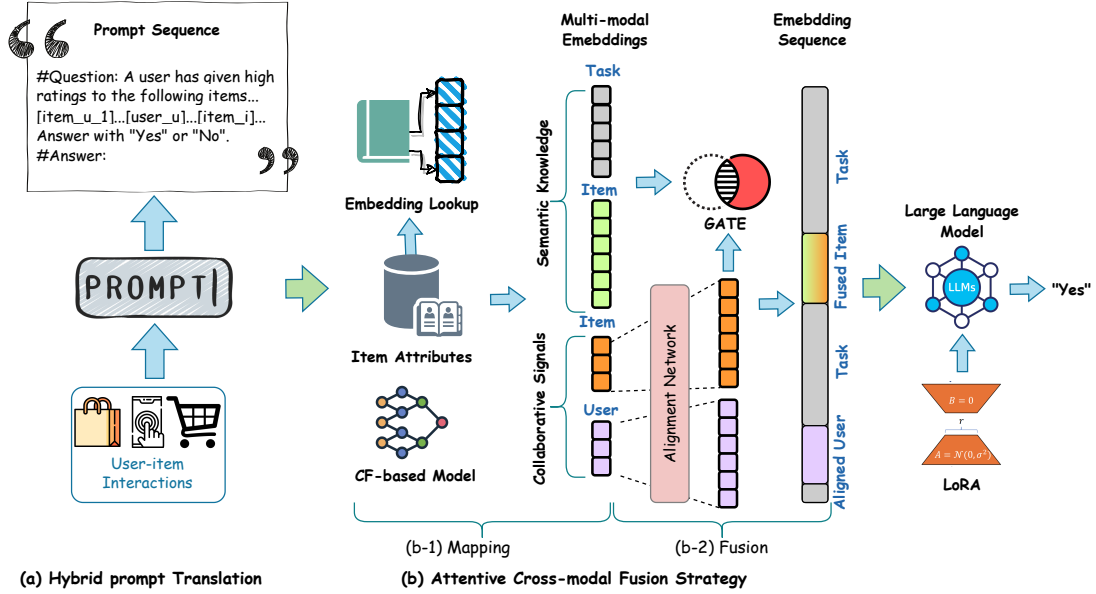


Figure 2: The overall framework of the proposed Collaborative Cross-modal Fusion with Large Language Model (CCF-LLM).

the total number of tokens in item i 's attribute. Each token embedding $x_i^{SM}[t]$ ($1 \leq t \leq T_i$) is a d -dimensional vector. Similarly, the user feature is defined as $x_u = \{x_u^{CF}\}$. Note that we do not include the user's semantic knowledge in the framework of CCF-LLM due to two reasons: (1) the availability of the user's textual attributes is highly inconsistent across different datasets, and (2) evaluating the impact of such inconsistent features is challenging. Hence, we leave it for future work to incorporate user semantic knowledge. Finally, we present our hybrid prompt template as follows.

```
#Question: A user has given high ratings to the following items:
{[Item_u]}. Additionally, we have information about the user's
preferences encoded in the feature [User_u]. Using all available
information, make a prediction about whether the user would enjoy
the item [Item_i]. Answer with "Yes" or "No". #Answer:
```

We use several special tokens as placeholders to represent user and item features within the prompt template. For the item side, we use `[Item_i]` to represent the item entity with ID i , and `{[Item_u]}` is a list of special tokens representing the items that have been interacted by user u . For the user side, we use `[User_u]` as a placeholder to represent the user entity with ID u . Hence, the prompt template consists of three parts: task descriptions in natural languages, user entities, and item entities. Note that, distinct from previous works [29, 50] which separates the embeddings for two modalities x_i^{CF} and x_i^{SM} in the prompt, we utilize a single token to represent the two heterogeneous modalities for each item to facilitate the cross-modal fusion.

3.3 Attentive Cross-modal Fusion Strategy

Our attentive cross-modal fusion strategy consists of two phases, mapping and fusion, that help to fuse information from heterogeneous modalities. Unlike previous works that utilize solely natural

languages [2] or inject CF embeddings [29, 50], our framework adaptively fuse the two modalities into an input embedding sequence for the LLM, as shown in Figure 2b. We explain the two phases of mapping and fusion in the following.

Mapping Phase. The mapping phase aims to convert the prompt sequences to latent embeddings. It has two primary goals: first, for task descriptions and user/item attributes (e.g., title), it maps them to latent embeddings representing semantic knowledge for subsequent processing by an LLM. Second, for the user and item's collaborative signals, it maps them to the latent embeddings from a pre-trained CF-based recommendation system.

Formally, let the hybrid prompt for user u and item i be $p_{u,i}$. For each token tk in the tokenized $p_{u,i}$, we obtain its latent embedding with the following rules. (1) If tk is a natural language token, we obtain its corresponding latent embedding directly via the LLM's input embedding lookup. (2) If tk is a placeholder for users, i.e., `[User_u]`, we obtain the latent embedding x_u^{CF} from a pre-trained CF-based model. (3) If tk is a placeholder for items, i.e., `[Item_i]` or `{[Item_u]}`, we obtain the embeddings representing the two modalities for semantic knowledge and collaborative signals. Specifically, we first obtain the item's CF embedding, x_i^{CF} , from a CF-based model similarly as the user side. Then, we extract the item's textual attributes from a database and obtain its LLM embedding, x_i^{SM} , via the LLM's input embedding lookup. The obtained x_i^{CF} and x_i^{SM} will be fused adaptively in the next part.

In this way, each token in the prompt sequence $p_{u,i}$ is mapped to its corresponding latent embeddings, which will be used to prepare a fused embedding sequence as the input to the LLM.

Fusion Phase. The CF embeddings cannot be directly fused with the LLM embeddings due to their inconsistent latent spaces. Similar to prior work [29, 50], we utilize an alignment network to map the

CF embeddings to the LLM embedding space. Given the alignment network $\text{ALG} : \mathbb{R}^l \rightarrow \mathbb{R}^d$, this process can be presented as follows.

$$\tilde{x}_u^{CF} = \text{ALG}(x_u^{CF}; \Theta_A), \quad \tilde{x}_i^{CF} = \text{ALG}(x_i^{CF}; \Theta_A), \quad (2)$$

where Θ_A is the set of trainable parameters. $\tilde{x}_u^{CF} \in \mathbb{R}^d$ and $\tilde{x}_i^{CF} \in \mathbb{R}^d$ are aligned CF embeddings which are ready to be fused. In practice, the alignment network is usually implemented as a multi-layer perceptron (MLP).

However, only aligning the dimension of latent embeddings between CF embeddings and the LLM embeddings is not enough to fully leverage the collaborative signal for the recommendation. Because it is still hard for LLM to understand the latent embeddings of collaborative signals due to their distinct nature. In response to this challenge, we propose an attentive cross-modal fusion strategy shown in Figure 2b. Inspired by the cross-gate mechanism [39], we train a learnable network GATE to fuse the two embeddings in different modalities in a dimension-wise manner adaptively. Specifically, for each token of the item i 's attribute, GATE learns a fusion weight vector $\alpha \in \mathbb{R}^d$, which is defined as

$$\begin{aligned} \alpha &= \text{GATE}(\tilde{x}_i^{CF}, x_i^{SM}[t]; \Theta_G) \\ &= \text{MLP}(\tilde{x}_i^{CF}; \Theta_{G_1}) + \text{MLP}(x_i^{SM}[t]; \Theta_{G_2}), \end{aligned} \quad (3)$$

where $\Theta_G = \{\Theta_{G_1}, \Theta_{G_2}\}$ is the set of trainable parameters. Finally, we utilize the weight vector α , which is adapted to the collaborative signal and semantic knowledge of each item, to fuse the LLM embedding x_i^{SM} and the aligned CF embedding \tilde{x}_i^{CF} , as follows.

$$\tilde{x}_i[t] = x_i^{SM}[t] + \alpha \odot \tilde{x}_i^{CF}, \quad (4)$$

where the operator \odot represents element-wise multiplication.

Finally, we use the mapped and fused embeddings to assemble the embedding sequence, which is fed into the LLM for prediction. For each token in the prompt $p_{u,i}$, we replace it with the corresponding embedding following the same rules as in the mapping step. (1) If it is a natural language token, we replace it with its LLM embedding, which is denoted by the grey parts in the embedding sequence in Figure 2(b-2). (2) If it is a placeholder for a user u , we replace it with the aligned CF embedding \tilde{x}_u^{CF} , which is denoted by the purple part. (3) If it is a placeholder for an item i , we replace it with $\tilde{x}_i = \{\tilde{x}_i[1], \tilde{x}_i[2], \dots, \tilde{x}_i[T_i]\}$, which is denoted by the green-orange mixture. The embeddings are concatenated in the same sequence as the tokens in the prompt sequence $p_{u,i}$.

3.4 Training

In this section, we describe how we optimize CCF-LLM, including the learning objectives as well as our two-stage training strategy.

Learning Objectives. Though a pre-trained LLM can perform many tasks in a zero-shot setting, it has not been specifically trained on collaborative signals. Thus, it is suboptimal to directly utilize the fused embedding sequence for recommendation without fine-tuning. Following the prevalent training paradigm in existing LLM4Rec works, we adopt the LoRA module [20] to tune the LLM. LoRA trains pairs of low-rank weights as an adaptor with significantly fewer parameters in comparison to the LLM, instead of updating all its trainable parameters. Hence, the full set of trainable parameters in our CCF-LLM is $\Theta = \{\Theta_E, \Theta_A, \Theta_G, \Theta_L\}$, where Θ_L is

the set of LoRA weights. Note that the tuning of Θ_E is optional as it is already pre-trained, and we only use the ENC network from the CF-based model. Θ_A and Θ_G are the training parameters of the ALG and GATE networks, respectively.

Conventional CTR models generate a Bernoulli distribution indicating the likelihood of interaction. In contrast, language models produce a probability distribution across the entire vocabulary, thereby differing from the output distribution of CTR models. To align \hat{y} to the output of a conventional CTR model, we define the probability $p_{yes} = P(\hat{y} = \text{“Yes”} | p_{u,i})$ as the indicator of a positive prediction, and $p_{no} = P(\hat{y} = \text{“No”} | p_{u,i})$ as a negative prediction. Here, “Yes” and “No” are two tokens in the LLM vocabulary. We optimize the LLM from two perspectives. First, we aim to align the predictions with the ground-truth label y . Second, we aim to further constrain the relationship between p_{yes} and p_{no} , so that $p_{yes} > p_{no}$ if $y = 1$ and vice versa. Therefore, we optimize the LLM by minimizing the following loss.

$$\min_{\Theta} \mathcal{L} = \mathcal{L}_1(p_{yes}, y) + \mathcal{L}_1(p_{no}, 1 - y) + k \times \mathcal{L}_2(p_{yes}, p_{no}, y), \quad (5)$$

where \mathcal{L}_1 denotes a classification loss that corresponds to the first goal and \mathcal{L}_2 is a ranking loss that corresponds to the second goal. We implement \mathcal{L}_1 as a binary cross-entropy (BCE) loss and \mathcal{L}_2 as a Bayesian personalized ranking (BPR) loss [37]. k is a scalar to balance the weight between the two objectives.

Two-stage Training. A straightforward training strategy is to train all parameters together in an end-to-end manner. Nevertheless, as empirically suggested [50], this approach could potentially diminish LLM4Rec performance, particularly in cold-start situations. In particular, before the convergence of the end-to-end tuning, the LLM may misguide the training of attentive cross-modal fusion. Consequently, we employ a two-stage training setting, a technique also used earlier [50]. Specifically, we divide Θ into Θ_1 and Θ_2 , two non-overlapping subsets such that $\Theta_1 = \{\Theta_L\}$, and $\Theta_2 = \{\Theta_E, \Theta_A, \Theta_G\}$. We first minimize Equation (5), updating only Θ_1 until it converges, followed by minimizing Equation (5) through the sole updating of Θ_2 until it also converges.

The intuition of the two-stage training is as follows: after the first stage, the LLM has been adapted to the recommendation task. However, at this point, the collaborative signal is still plagued with noises because both ALG and GATE networks are not trained well, thereby curtailing the recommendation capability. As a result, in the second stage, we strive to augment the representation potential of the collaborative signals by focusing on the attentive cross-modal fusion. Given that the LLM has already been tuned, for efficiency we do not further fine-tune it during the second stage.

3.5 Comparison to Existing Work

For a fair comparison with existing LLM4Rec, we focus on the different approaches to incorporating collaborative signals. Thus, we intentionally adopt a similar input prompt format as existing work such as CoLLM [50] and LLaRA [29]. However, our proposed CCF-LLM is different from them in the following three aspects.

Motivation. At a high level, existing work [29, 50] aims to incorporate the collaborative signals by directly injecting into LLMs the latent embeddings pre-trained by a disparate CF model. However,

our CCF-LLM aims to solve the problem that LLM cannot effectively understand the CF embeddings by unifying both modalities in a single space rather than just introducing a new modality, as further explained below.

Hybrid Prompt Translation. In existing work [29, 50], the hybrid prompt encodes texts and collaborative signals separately in different fields (i.e., tokens) without considering their fusion. In CCF-LLM, we utilize a single token to unify the representation of the two modalities for each item, and further consider the following fusion steps.

Mapping and Fusion of Two Modalities. In existing work, there is no fusion of the two modalities. Both latent embeddings for semantic knowledge (compatible with a pre-trained LLM) and collaborative signals (extracted from a pre-trained CF model) are fed into an LLM as is. However, in CCF-LLM, we fuse the two modalities with a trainable cross-gate mechanism.

To sum up, though some existing work has realized the importance of introducing collaborative signals, they merely inject the pre-trained collaborative signals directly into LLMs, overlooking the disparity between the two modalities. Thus, LLMs may not be able to understand the pre-trained collaborative signals. In contrast, CCF-LLM utilizes the attentive cross-modal fusion strategy to fully fuse the latent embeddings for these two heterogeneous modalities in a single space for the first time in the field of LLM4Rec. Compared to these previous methods, we offer a new insight into fusing the two modalities in the LLM4Rec paradigm.

4 Experiment

In this section, we first elaborate on our experimental settings. Then, we evaluate our CCF-LLM on the real-world CTR task and address the following research questions (RQs). **RQ1:** Can LLMs utilize collaborative signals solely through natural language descriptions without CF embeddings? **RQ2:** Can our proposed CCF-LLM fully leverage collaborative signals to enhance the recommendation? **RQ3:** How can we effectively fuse the modalities from LLM and CF embeddings? **RQ4:** How does our proposed attentive cross-modal fusion strategy benefit recommendation?

4.1 Experiment Settings

Datasets. We conduct experiments on the MovieLens-1M [13] and Amazon-Book [14] datasets. For the MovieLens-1M dataset, the 20 most recent months of user-item interactions are used. The train/validation/test data are split by 10/5/5 months. For the Amazon-Book dataset, a total of 16 months of user-item interactions from 2017 onwards are used. The train/validation/test data are split by 11/2.5/2.5 months. In both datasets, users rate items on a 1 to 5 scale. To convert ratings to binary labels, a threshold of 3 is used for MovieLens-1M and 4 is used for Amazon-Book. The statistics of datasets used in this paper are shown in Table 1.

Evaluation Metric. Since our recommendation task is CTR prediction, following the literature [2, 27, 45, 50, 52] in both conventional CF-based and LLM4Rec models, we adopt the standard Area Under the Curve (AUC) as the metric to evaluate the CTR performance. To compute the AUC of an LLM4Rec approach, we use the normalized p_{yes} as the prediction score. To evaluate the relative improvement,

Table 1: Statistics of Datasets.

Dataset	User	Item	Train	Validation	Test
MovieLens-1M	839	3,256	33,891	10,401	7,331
Amazon-Books	22,967	34,154	727,468	25,747	25,747

we adopt the RelImpr metric [52], which is defined as follows.

$$RelImpr = \left(\frac{AUC \text{ of measured model} - 0.5}{AUC \text{ of base model} - 0.5} - 1 \right) \times 100\% \quad (6)$$

Baselines. We compare our proposed CCF-LLM to baselines in various categories to answer the aforementioned research questions. For RQ1, we investigate two well-known LLMs: (1) a widely used **closed-source LLM** that can be accessed via APIs, and (2) **Vicuna-7B** [8], a powerful LLM fine-tuned from LLaMA [41]. For RQ2, we compare CCF-LLM with several conventional CF-based models and the state-of-the-art LLM4Rec models. The CF-based models include: (1) **MF** [26], the classical matrix factorization for recommendation; (2) **LightGCN** [15], a state-of-the-art graph-based recommendation model; (3) **SASRec** [23], a sequence recommendation model with self-attention. The LLM4Rec models include: (1) **Softprompt** [51], a soft prompt-tuning method; (2) **TallRec** [2], an instruct-tuning-based method; (3) **CoLLM** [50], a method that directly injects aligned collaborative signals into the embedding sequences of an LLM. For RQ3 and RQ4, we compare several variants of CCF-LLM, which will be introduced in Sections 4.3 and 4.5.

Impelmentation details. We tune hyper-parameters and configure experiment settings based on the validation set and guidance from the literature. We align the backend LLM for all LLM4Rec approaches to Vicuna-7B, and implement CCF-LLM with three backend CF-based models used as our baselines, denoting them as CCF-LLM (MF), CCF-LLM (LightGCN) and CCF-LLM (SASRec). We set the dimension d to 128 for all CF embeddings x_u^{CF} and x_i^{CF} . The dimension l for the backend Vicuna-7B LLM is 4096. The hyper-parameter k in Equation (5) is set to 2.0. When pre-training CF-based models, we set the batch size to 1024 and use the Adam optimizer with a learning rate of either 1e-3 or 1e-2 until converged. When training CCF-LLM with the two-stage training strategy, we have referenced the official training script² for Vicuna. In particular, we use the AdamW optimizer with a learning rate of 1e-4 and batch size of 3. We also limit the maximum length of the input sequence to 700 tokens to prevent out-of-memory issues, such that inputs exceeding this length are truncated. For other parameters, we adhere to the configurations outlined in their original papers or the default values in their software. All experiments were conducted on a Linux server with 4×V100 GPUs and 128GB RAM.

4.2 Pilot Study (RQ1)

In this section, we aim to explore a compelling hypothesis: Given the significance of collaborative signals in recommendation systems and the exceptional performance of LLMs on diverse natural language tasks, can LLMs directly leverage collaborative signals encoded using natural language descriptions?

²<https://github.com/lm-sys/FastChat>

Table 2: Pilot study on the recommendation capability of LLMs, encoding collaborative signals in natural language.

Method		Closed-source LLM	Vicuna-7B
		AUC	AUC
Zero-shot	w/ CF	0.5902	0.6202
	w/o CF	0.5783	0.6184
Instruct-tuning	w/ CF	0.6235	0.7108
	w/o CF	0.5998	0.7084
MF		0.6482 [†]	
CoLLM (MF)		0.7295 [†]	
CCF-LLM (MF)		0.7315	

Results are reported as the average of 5 runs.

[†]Results are reproduced from Zhang et al. [50].

Experiment Setup. As mentioned in Section 4.1, we resort to two well-known LLMs to evaluate this research question: a widely used closed-source LLM (gpt-3.5-turbo-0613) and Vicuna-7B (the same LLM used in CCF-LLM). The key step is to encode collaborative signals in natural languages. Building on the fundamental assumption of collaborative filtering that items interacted by users of similar interests should be highlighted, we identify items that are potentially interesting to the target user through a pre-trained CF-based model. These potential items are then communicated directly to the LLM as collaborative signals by explicitly mentioning their titles. Specifically, we replace the second sentence of the prompt template described in Section 3.2 with the following: “Additionally, we have information that users like these items may also enjoy: {[sim_items]}.”, where {[sim_items]} is a list of top-10 items out of all items in test set recommended by MF. We assess each LLM through a zero-shot setting and an instruct-tuning setting. In the instruct-tuning setting, we set the prompt output to “Yes” or “No” according to the ground-truth labels. Note that we are unable to access the parameters of the closed-source model. Thus we only evaluate the closed-source model by treating its generated output as a binary value when computing AUC.

Results and Discussion. We report the results of only using natural language descriptions for LLMs in Table 2. “w/ CF” indicates that we encode the collaborative signals into prompts, and “w/o CF” indicates that no collaborative signal is encoded. Besides, we also report the results of MF, CoLLM, and the proposed CCF-LLM for reference, where CoLLM and CCF-LLM use MF as the backend for a fair comparison. It can be observed that collaborative signals improve recommendations by an average of 0.0178 in both zero-shot and instruct-tuning settings for the closed-source LLM and 0.0021 for Vicuna-7B. On the one hand, the improvement confirms that LLM can benefit from collaborative signals in the form of natural language descriptions. Our results also align with previous studies [2, 6, 12] that tuning LLMs with user-item interactions can adapt the LLMs to recommendation tasks, thus achieving favorable performance over the zero-shot setting. On the other hand, the improvement is marginal. For LLM4Rec approaches that integrate CF embeddings, CoLLM surpasses the best performance achieved through using solely natural language descriptions by 0.0187, while

Table 3: Comparison between the proposed CCF-LLM and the baselines. The best result in each group is bolded.

Method	MovieLens-1M		Amazon-Book	
	AUC	RelaImpr	AUC	RelaImpr
MF	0.6482 [†]	-	0.7134 [†]	-
CoLLM (MF)	0.7295 [†]	54.86%	0.8109 [†]	45.69%
CCF-LLM (MF)	0.7315	56.21%	0.8150	47.61%
LightGCN	0.5959 [†]	-	0.7103 [†]	-
CoLLM (LightGCN)	0.7100 [†]	118.98%	0.7978 [†]	41.61%
CCF-LLM (LightGCN)	0.7427	153.08%	0.8049	44.98%
SASRec	0.7078 [†]	-	0.6887 [†]	-
CoLLM (SASRec)	0.7235 [†]	7.56%	0.7746 [†]	45.52%
CCF-LLM (SASRec)	0.7526	21.56%	0.7792	47.96%
Softprompt	0.7071 [†]	-	0.7224 [†]	-
TallRec	0.7097 [†]	1.25%	0.7375 [†]	6.79%
CoLLM (Best)	0.7295 [†]	10.82%	0.8109 [†]	39.79%
CCF-LLM (Best)	0.7526	21.97%	0.8150	41.64%

Results are reported as the average of 5 runs.

[†]Results are reproduced from Zhang et al. [50].

CCF-LLM outperforms by 0.0207. We hypothesize that the significant gap arises as natural language descriptions are inadequate to fully convey the collaborative signals, thus making it difficult to instruct LLMs to enhance recommendations. It is more effective to utilize CF embeddings as the carrier for collaborative signals. Therefore, we choose to integrate both modalities in CF and LLM embeddings under the context of LLM4Rec.

4.3 Main Results of CCF-LLM (RQ2)

In this section, we investigate the performance of the proposed CCF-LLM on the MovieLens-1M and Amazon-book datasets.

Experiment Setup. We implement our proposed CCF-LLM with three backend CF-based models and compare them individually in three subgroups. Besides, we also include CoLLM [50] in each subgroup with the corresponding backend CF-based model. Finally, we compare with other state-of-the-art LLM4Rec approaches that do not rely on CF embeddings in the fourth subgroup, including Softprompt [51] and TallRec [2]. We denote CCF-LLM (Best) and CoLLM (Best) as the implementation with the backend CF-based models that have achieved the best performance.

Results and Discussion. Table 3 displays a comparison of the overall performance. To compute RelaImpr, the base models are the corresponding conventional CF-based models in the first three subgroups. For the fourth subgroup, the base model is Softprompt. Based on the results, we can make the following observations: (1) CCF-LLM outperforms all other baselines, including conventional CF-based models and state-of-the-art LLM4Rec methods, demonstrating the effectiveness of our proposed framework. (2) Compared to conventional CF-based models, CCF-LLM has a relative improvement of up to 153.08% in MovieLens-1M and 47.96% in Amazon-Book. This reveals that the absence of semantic knowledge significantly limits the performance of conventional CF-based models. Thus, it is necessary to explore the potential of LLMs in




Figure 3: Ablation study on cross-modal fusion strategies.

recommendation. (3) CCF-LLM has a relative improvement of up to 41.64% compared to LLM4Rec methods that do not rely on CF embeddings. This observation suggests utilizing latent embeddings for collaborative signals is useful in improving recommendation performance. (4) Compared to CoLLM which also utilizes both CF and LLM embeddings, CCF-LLM has better performance. This is due to the proposed attentive cross-modal fusion strategy that contributes to a more synergistic integration between the heterogeneous modalities. (5) Our CCF-LLM is similar to Softprompt by tuning the embeddings for special tokens. However, the trainable embeddings in Softprompt are not constrained with a CF-based model as CCF-LLM does. Note that the performance of Softprompt is even worse than TallRec which only relies on natural language descriptions. This observation suggests that improperly tuning the embeddings for an LLM may result in a negative impact.

4.4 Ablation Study (RQ3)

In this section, we investigate the influence of the proposed attentive cross-modal fusion strategy. Our default fusion strategy, as illustrated in Equations (3) and (4), relies on the weight vector α to perform an adaptive dimensional-wise fusion between the CF and LLM embeddings. Though our fusion strategy has proven to be effective, we are still curious if there are other alternatives.

Experiment Setup. We conduct the ablation studies from two angles, including *how to fuse* and *what to fuse*. For the first angle, we examine two alternatives by modifying Equations (3) and (4). Firstly, we fix α to a constant value 1, so that the fused item embedding will be computed as $\tilde{x}_i[t] = x_i^{SM}[t] + \tilde{x}_i^{CF}$. We denote this variant as “ $\alpha = 1$ ”. Then we modify the output dimension of GATE network to 1, i.e., to conduct cross-modal fusion in a token-wise manner. We denote this variant as “ $\alpha \in R$ ”. For a fair comparison, we utilize LightGCN as the backend model for all the above variants.

Besides, we are also interested in *what to fuse*, particularly in fusing multiple CF embeddings from different CF-based models at a time. We denote the implementations of CCF-LLM with different backend CF-based models in Section 4.3 as “L” (for LightGCN), “M” (for MF), and “S” (for SASRec). Here, we further examine the combination of collaborative signals from LightGCN and MF (denoted as “L+M”), LightGCN and SASRec (denoted as “L+S”) as well as MF and SASRec (denoted as “M+S”). In this way, the fusion phase defined in Equation (4) can be presented as follows.

$$\tilde{x}_i[t] = x_i^{SM}[t] + \alpha \odot \tilde{x}_i^{CF_1} + \beta \odot \tilde{x}_i^{CF_2}, \quad (7)$$

Table 4: Impact of the training strategies.

Training Strategy	AUC	RelaImpr
Default	0.7427	-
Stage 1 Only	0.7025	-16.56%
End-to-end	0.6390	-42.73%

where $\tilde{x}_i^{CF_1}$ and $\tilde{x}_i^{CF_2}$ are CF embeddings generated from different CF-based models for the same item i . β is another fusion weight vector output from a second GATE network.

Results and Discussion. We present the findings of our ablation study as shown in Figure 3. Our analysis in Figure 3a focuses on *how to fuse* and reveals that adding CF embeddings directly onto LLM embeddings ($\alpha = 1$) results in the worst performance. This is due to the fact that such a fusion strategy completely disregards the correlation between the two modalities. In contrast, $\alpha \in R$ and CoLLM have quite a better performance than $\alpha = 1$. It is because, though inadequate, they both capture the correlation to some extent. In $\alpha \in R$ it is conducted with a GATE network adaptively in a token-wise manner, and in CoLLM it is conducted with the LLM which applies an attention mechanism on the input token sequence. Our proposed CCF-LLM, however, utilizes the cross-gate mechanism adaptively in a finer dimensional-wise granularity, leading to optimal performance.

In Figure 3b, we examine *what to fuse* and find that combining multiple CF embeddings does not improve the recommendation performance compared to using only a single backend CF-based model. One possible reason is that the collaborative signals from various models are redundant, offering no additional insight beyond that of a single model. For some items, there could also be conflicting signals that are incompatible for fusion.

4.5 Model Analyses (RQ4)


In this section, we aim to analyze several key factors in our proposed CCF-LLM that can benefit the recommendations, including the training strategies as well as the distribution of fused embeddings.

Training Steps. As demonstrated in Section 3.4, our training strategy involves two stages. We examine two other strategies in Table 4, including only training Θ_1 without the second stage (denoted as “Stage 1 Only”), and training Θ_1 and Θ_2 in an end-to-end manner within a single stage until convergence (denoted as “End-to-end”). Experiments are conducted on the MovieLens-1M dataset with LightGCN as the backend CF-based model.

We can observe that without training the attentive cross-modal fusion strategy in the second stage, the fusion between the two modalities are incomplete, resulting in inferior performance. On the other hand, training all parameters in an end-to-end manner will significantly decrease the performance. A possible reason could be that earlier in the training process, an insufficiently tuned LLM may misguide the training of the ALG and GATE networks, leading to undesirable fusion. In summary, our two-stage training strategy empowers the effective optimization of the cross-modal fusion strategy. It helps us integrate the two modalities in a unified latent space

Table 5: Case study of a user who has given high ratings to action, drama, and adventure movies.

Movies	Genres	Ground-truth Labels	Rank by LightGCN	Rank by Vicuna-7B	Rank by CCF-LLM
Deconstructing Harry	Comedy, Drama	1	5	2	1
Big Daddy	Comedy	0	4	1	4
But I'm a Cheerleader	Comedy	1	2	3	2
Desperately Seeking Susan	Comedy, Romance	1	1	4	3
Lake Placid	Horror, Thriller	0	3	5	5

**Figure 4: Visualization of the item embeddings in different modalities with t-SNE. Green: Aligned CF embeddings; Yellow: LLM embeddings; Purple: fused embeddings.**

for a synergistic fusion, thereby benefiting the recommendation outcomes.

Visualization of Cross-modal Fusion. To further examine how the embeddings are fused into a unified latent space, we randomly select 1,000 items from each dataset and visualize the item embeddings from the two modalities as well as their fused embeddings with t-SNE [42] in Figure 4. The CF embeddings for collaborative signals are extracted from SASRec.

From Figure 4 we can observe that the aligned CF embeddings (green) and LLM embeddings (yellow, which are pre-processed with mean pooling for each item before t-SNE) are clearly separated from each other with well-defined boundaries. This observation suggests that directly injecting collaborative signals into LLMs is infeasible, even after the CF embeddings are already aligned by ALG as in previous work [29, 50]. In CCF-LLM, we further fuse the two modalities with an attentive cross-modal fusion strategy. We observe that the fused embeddings (purple) become indistinguishable from the LLM embeddings that capture semantic knowledge, suggesting that the fused embeddings are more consistent with the LLM space. The visualization on the two datasets indicate that our CCF-LLM can better fuse the two types of modalities with the proposed attentive cross-modal fusion strategy.

4.6 Case Study

We select a typical case to investigate the impact of fusing collaborative signals and semantic knowledge in LLM4Rec. In the test set of MovieLens-1M, a user with ID 4 has given high ratings to “Star Wars: Episode IV - A New Hope”, “Blade Runner”, “Raiders of the Lost Ark”, “The Godfather”, “Die Hard”, “True Romance”, “Rocky”, “The Untouchables”, and “Diva”. From the list, we observe that those movies are primarily actions, dramas, or adventures. In the case study, we select five movies in the test set to examine the

predictions made by LightGCN, the original Vicuna-7B, and our proposed CCF-LLM using LightGCN as the backend model.

Table 5 presents the results of the case study. On the one hand, we can observe that Vicuna-7B utilizes semantic knowledge to predict high scores for the movies “Deconstructing Harry” and “Big Daddy”, which have similar genres or themes with the user’s previously favored movies. On the other hand, LightGCN makes predictions that benefit from collaborative signals, and it predicts high scores for “But I’m a Cheerleader” and “Desperately Seeking Susan”. Both of these two movies are positive in the test set but are not similar to the user’s previously favored movies in semantics, which shows the significance of collaborative signals for recommendation.

Finally, by fusing collaborative signals with semantic knowledge, CCF-LLM can leverage both modalities for recommendation. For instance, CCF-LLM shows an increase in ranks for “But I’m a Cheerleader” and “Desperately Seeking Susan” when compared to predictions made solely using natural language input by Vicuna-7B. Moreover, guided by the collaborative signal, CCF-LLM decreases the rank for “Big Daddy” from 1 to 4. The changes in rank illustrate that collaborative signals can help LLMs to improve the predictions. It is worth noting that CCF-LLM raises the rank for “Deconstructing Harry” even though LightGCN does not rate it highly. This observation suggests that CCF-LLM is robust to noisy collaborative signals and can correct the inaccurate information from CF embeddings. As a result, our proposed CCF-LLM outperforms both CF-based models and LLMs in recommendation.

5 Conclusion

In this paper, we proposed a novel framework for collaborative cross-modal fusion with large language models for recommendation (CCF-LLM). Through the hybrid prompt translation and the attentive cross-modal fusion strategy, CCF-LLM enhanced LLM4Rec by adaptively fusing the CF and LLM embeddings, achieving synergy across heterogeneous modalities. Our work shed light on how to effectively leverage the collaborative signals in the context of LLM4Rec. In future studies, we aim to explore the full potential of LLMs in recommendations with more types of modalities and more advanced fusion strategies.

Acknowledgments

This research / project is supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (Proposal ID: T2EP20122-0041). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. arXiv:2303.08774
- [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. In *RecSys*. 1007–1014.
- [3] Matthew Brand. 2003. Fast online SVD revisions for lightweight recommender systems. In *SDM*. 37–46.
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. In *NeurIPS*. 1877–1901.
- [5] Rui Chen, Qingyi Hua, Yan-Shuo Chang, Bo Wang, Lei Zhang, and Xiangjie Kong. 2018. A survey of collaborative filtering-based recommender systems: From traditional methods to hybrid methods based on social networks. *IEEE Access* 6 (2018), 64301–64320.
- [6] Zheng Chen. 2023. PALR: Personalization Aware LLMs for Recommendation. arXiv:2305.07622
- [7] Zhiyong Cheng, Ying Ding, Xiangnan He, Lei Zhu, Xuemeng Song, and Mohan S Kankanhalli. 2018. A³NCF: An Adaptive Aspect Attention Model for Rating Prediction. In *IJCAL*. 3748–3754.
- [8] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. <https://vicuna.lmsys.org>
- [9] Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. M6-rec: Generative pretrained language models are open-ended recommender systems. (2022). arXiv:2205.08084
- [10] Mukund Deshpande and George Karypis. 2004. Item-based top-n recommendation algorithms. *ACM Transactions on Information Systems* 22, 1 (2004), 143–177.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *ACL*. 4171–4186.
- [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 *RecSys*. 299–315.
- [13] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm Transactions on Interactive Intelligent Systems* 5, 4 (2015), 1–19.
- [14] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *WWW*. 507–517.
- [15] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [16] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *WWW*. 173–182.
- [17] Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian McAuley. 2023. Large language models as zero-shot conversational recommenders. In *CIKM*. 720–730.
- [18] John D Holt and Soon M Chung. 1999. Efficient mining of association rules in text databases. In *CIKM*. 234–242.
- [19] Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *ECIR*. 364–381.
- [20] 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*.
- [21] Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *RecSys*. 1281–1283.
- [22] Jianchao Ji, Zelong Li, Shuyuan Xu, Wenyue Hua, Yingqiang Ge, Juntao Tan, and Yongfeng Zhang. 2024. GenRec: Large language model for generative recommendation. In *ECIR*. 494–502.
- [23] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *ICDM*. 197–206.
- [24] James M Keller, Michael R Gray, and James A Givens. 1985. A fuzzy k-nearest neighbor algorithm. *IEEE transactions on systems, man, and cybernetics* 4 (1985), 580–585.
- [25] Ioannis Konstas, Vassilios Stathopoulos, and Joemon M Jose. 2009. On social networks and collaborative recommendation. In *SIGIR*. 195–202.
- [26] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.
- [27] Xiangyang Li, Bo Chen, Lu Hou, and Ruiming Tang. 2023. CTRL: Connect Tabular and Language Model for CTR Prediction. arXiv:2306.02841
- [28] 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:2312.02443
- [29] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. Llara: Large language-recommendation assistant. In *SIGIR*. 1785–1795.
- [30] Jianghao Lin, Xinyi Dai, Yunjia Xi, Weiwen Liu, Bo Chen, Xiangyang Li, Chenxu Zhu, Hui Feng Guo, Yong Yu, Ruiming Tang, et al. 2023. How Can Recommender Systems Benefit from Large Language Models: A Survey. (2023). arXiv:2306.05817
- [31] Huanshuo Liu, Bo Chen, Menghui Zhu, Jianghao Lin, Jiarui Qin, Yang Yang, Hao Zhang, and Ruiming Tang. 2024. Retrieval-Oriented Knowledge for Click-Through Rate Prediction. arXiv:2404.18304
- [32] Peng Liu, Lemei Zhang, and Jon Atle Gulla. 2023. Pre-train, Prompt, and Recommendation: A Comprehensive Survey of Language Modeling Paradigm Adaptations in Recommender Systems. *Transactions of the Association for Computational Linguistics* 11 (2023), 1553–1571.
- [33] Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. 2024. LLM-Rec: Personalized Recommendation via Prompting Large Language Models. In *NAACL*. 583–612.
- [34] Yashar Moshfeghi, Benjamin Piwowarski, and Joemon M Jose. 2011. Handling data sparsity in collaborative filtering using emotion and semantic based features. In *SIGIR*. 625–634.
- [35] Zhaopeng Qiu, Xian Wu, Jingyue Gao, and Wei Fan. 2021. U-BERT: Pre-training user representations for improved recommendation. In *AAAI*. 4320–4327.
- [36] Steffen Rendle. 2010. Factorization machines. In *ICDM*. 995–1000.
- [37] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*. 452–461.
- [38] Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *Comput. Surveys* 47, 1 (2014), 1–45.
- [39] Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. arXiv:1505.00387
- [40] Xiaoyuan Su and Taghi M Khoshgoftaar. 2009. A survey of collaborative filtering techniques. *Advances in artificial intelligence* (2009).
- [41] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. (2023). arXiv:2302.13971
- [42] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research* 9, 86 (2008), 2579–2605.
- [43] Lei Wang and Ee-Peng Lim. 2023. Zero-Shot Next-Item Recommendation using Large Pretrained Language Models. arXiv:2304.03153
- [44] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *SIGIR*. 950–958.
- [45] 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:2306.10933
- [46] Xi Wang, Harald Steck, Yang Guo, and Yong Liu. 2012. On top-k recommendation using social networks. In *RecSys*. 67–74.
- [47] Bin Yin, Junjie Xie, Yu Qin, Zixiang Ding, Zhichao Feng, Xiang Li, and Wei Lin. 2023. Heterogeneous Knowledge Fusion: A Novel Approach for Personalized Recommendation via LLM. In *RecSys*. 599–601.
- [48] Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2023. Glm-130b: An open bilingual pre-trained model. In *ICLR*.
- [49] Yuhui Zhang, Hao Ding, Zeren Shui, Yifei Ma, James Zou, Anoop Deoras, and Hao Wang. 2021. Language models as recommender systems: Evaluations and limitations. In *NeurIPS Workshop ICBINB*.
- [50] Yang Zhang, Fuli Feng, Jizhi Zhang, Keqin Bao, Qifan Wang, and Xiangnan He. 2023. Collm: Integrating collaborative embeddings into large language models for recommendation. arXiv:2310.19488
- [51] Zizhuo Zhang and Bang Wang. 2023. Prompt Learning for News Recommendation. In *SIGIR*. 227–237.
- [52] 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 *KDD*. 1059–1068.