LLM-Rec: Personalized Recommendation via Prompting Large Language Models

Hanjia Lyu[1] **Song Jiang**[2] **Hanqing Zeng**[3] **Qifan Wang**[3] **Si Zhang**[3]
Ren Chen[3] **Chris Leung**[3] **Jiajie Tang**[3] **Yinglong Xia**[3] **Jiebo Luo**[1]
[1]University of Rochester [2]UCLA [3]Meta AI
hlyu5@ur.rochester.edu yxia@meta.com jluo@cs.rochester.edu

Abstract

We investigate various prompting strategies for enhancing personalized recommendation performance with large language models (LLMs) through *input augmentation*. Our proposed approach, termed **LLM-Rec**, encompasses four distinct prompting strategies: (1) basic prompting, (2) recommendation-driven prompting, (3) engagement-guided prompting, and (4) recommendation-driven + engagement-guided prompting. Our empirical experiments show that incorporating the augmented input text generated by LLM leads to *improved* recommendation performance. Recommendation-driven and engagement-guided prompting strategies are found to elicit LLM's understanding of global and local item characteristics. This finding highlights the importance of leveraging diverse prompts and input augmentation techniques to enhance the recommendation capabilities with LLMs.
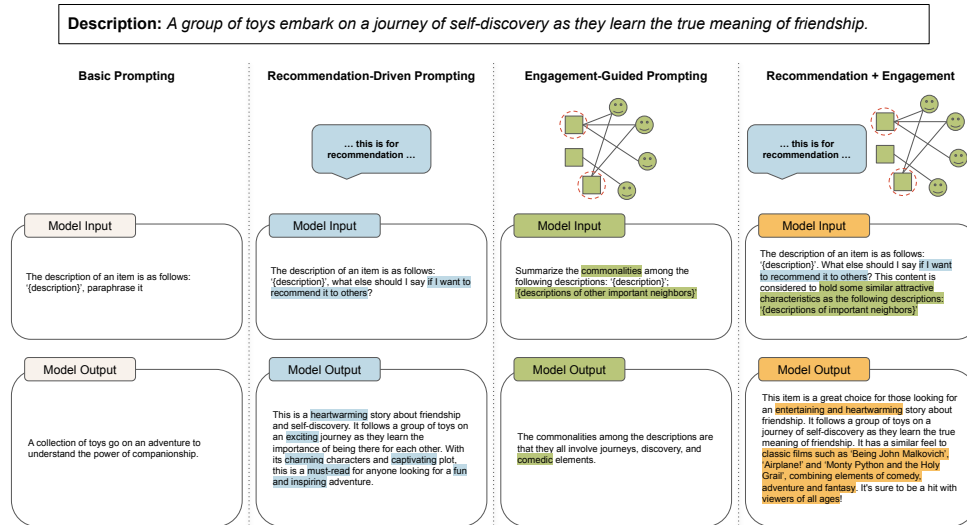
Figure 1: Recommendation-driven and engagement-guided prompting components play a crucial role in enabling large language models to focus on relevant context and align with user preferences. Recommendation-driven and engagement-guided prompting components, along with their corresponding augmented text are highlighted.

1   Introduction

The use of large language models in recommender systems has garnered significant attention in recent research. Numerous studies have explored the direct use of LLMs as recommender models. The underlying principle of these approaches involves constructing prompts that encompass the recommendation task, user profiles, item attributes, and user-item interactions. These task-specific prompts are then presented as input to the LLM, which is instructed to predict the likelihood of interaction between a given user and item [3, 4, 5, 8, 9, 21].

While these works demonstrate the potential of LLMs as powerful recommender models, the focus primarily revolves around utilizing the LLM directly for recommendation purposes. However, in this study, we approach the problem from a different perspective. Rather than using LLMs as recommender models, this study delves into the exploration of prompting strategies to *augment input text* with LLMs for personalized content recommendation. By leveraging LLMs, which have been fine-tuned on extensive language datasets, we seek to unlock their potential in generating high-quality and context-aware input text for enhanced recommendations.

Specifically, we propose **LLM-Rec** prompting, which encompasses various prompting strategies tailored for personalized content recommendation. These strategies include basic prompting, recommendation-driven prompting, engagement-guided prompting, and the combination of recommendation-driven and engagement-guided prompting. By leveraging these strategies, we aim to enhance the generation of input text by LLMs and improve the accuracy and relevance of content recommendations.

Through comprehensive empirical experiments, we evaluate the effectiveness of the **LLM-Rec** framework and compare it against baseline approaches. Our study provides insights into the impact of different prompting strategies on recommendation performance and sheds light on the potential of leveraging LLMs for personalized recommendation.

2   LLM-Rec Prompting

2.1   Basic prompting

We consider three basic prompting variants and refer to them as p_{para}, p_{tag}, and p_{infer}, respectively in the following experiments.

- p_{para}: This prompt instructs LLM to paraphrase the original content description, emphasizing the objective of maintaining the same information without introducing any additional details.
- p_{tag}: This prompt instructs LLM to summarize the content description by using tags, aiming to generate a more concise overview that captures key information.
- p_{infer}: This prompt instructs LLM to deduce the characteristics of the original content description and provide a categorical response that operates at a broader, less detailed level of granularity.

The exact prompts and corresponding responses by LLM are shown in Figure 2 (upper).

2.2   Recommendation-driven prompting

This prompting strategy is to add a recommendation-driven instruction into the basic prompting. We refer to the three recommendation-driven prompting as p_{para}^{rec}, p_{tag}^{rec}, and p_{infer}^{rec}, respectively in the following experiments, aligning with their counterparts in the basic prompts. The exact prompts and corresponding responses by LLM are shown in Figure 2 (lower).

The use of recommendation-driven prompting exhibits several compelling characteristics, making it an appealing approach for generating high-quality content descriptions:

1. **Enhanced Context:** By explicitly mentioning that the generated content description is intended for content recommendation, models gain a clearer understanding of the task at hand. This additional context helps models align their responses more closely with the purpose of generating content descriptions for recommendation purposes.
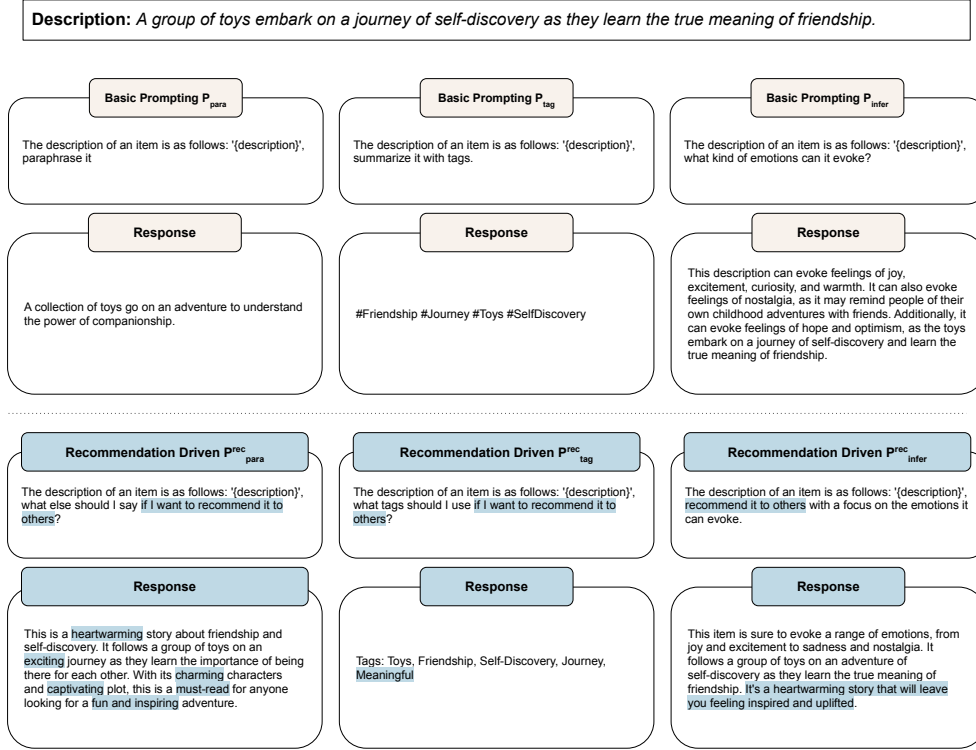
Description: *A group of toys embark on a journey of self-discovery as they learn the true meaning of friendship.*

Basic Prompting P_para

The description of an item is as follows: '{description}', paraphrase it

Basic Prompting P_tag

The description of an item is as follows: '{description}', summarize it with tags.

Basic Prompting P_infer

The description of an item is as follows: '{description}', what kind of emotions can it evoke?

Response

A collection of toys go on an adventure to understand the power of companionship.

Response

#Friendship #Journey #Toys #SelfDiscovery

Response

This description can evoke feelings of joy, excitement, curiosity, and warmth. It can also evoke feelings of nostalgia, as it may remind people of their own childhood adventures with friends. Additionally, it can evoke feelings of hope and optimism, as the toys embark on a journey of self-discovery and learn the true meaning of friendship.

Recommendation Driven P^rec_para

The description of an item is as follows: '{description}', what else should I say if I want to recommend it to others?

Recommendation Driven P^rec_tag

The description of an item is as follows: '{description}', what tags should I use if I want to recommend it to others?

Recommendation Driven P^rec_infer

The description of an item is as follows: '{description}', recommend it to others with a focus on the emotions it can evoke.

Response

This is a heartwarming story about friendship and self-discovery. It follows a group of toys on an exciting journey as they learn the importance of being there for each other. With its charming characters and captivating plot, this is a must-read for anyone looking for a fun and inspiring adventure.

Response

Tags: Toys, Friendship, Self-Discovery, Journey, Meaningful

Response

This item is sure to evoke a range of emotions, from joy and excitement to sadness and nostalgia. It follows a group of toys on an adventure of self-discovery as they learn the true meaning of friendship. It's a heartwarming story that will leave you feeling inspired and uplifted.

Figure 2: Examples of basic prompting (upper) - $\mathbf{p_{para}}$, $\mathbf{p_{tag}}$, $\mathbf{p_{infer}}$, recommendation-driven prompting (lower) - $\mathbf{p^{rec}_{para}}$, $\mathbf{p^{rec}_{tag}}$, $\mathbf{p^{rec}_{infer}}$, and their corresponding responses made by GPT-3 (text-davinci-003). Recommendation-driven instructions as well as the additional content provided in the response to the recommendation-driven prompts are highlighted.

2. **Guided Generation:** The specific instruction acts as a guiding cue for models, directing their attention towards generating content descriptions that are better suited for recommendation scenarios. The mention of "content recommendation" likely prompts LLM to focus on key features, relevant details, and aspects of the content that are more helpful in guiding users towards their preferred choices.

3. **Improved Relevance:** The instruction aids LLM in generating content descriptions that are tailored to the requirements of content recommendation. This alignment with the recommendation task leads to more relevant and informative descriptions, as LLM is primed to emphasize aspects that are important for users seeking recommendations.

2.3 Engagement-guided prompting

This prompting strategy is to leverage user behavior (*i.e.*, user-item engagement) to design prompts with the intention to guide LLM to better capture the characteristics inside the content description that align with user preferences. We aim to generate more meaningful description with this type of prompts for recommendation tasks. We refer to this variant as $\mathbf{p^{eng}}$.

To create the engagement-guided prompt, we combine the content description of the target item, denoted as d_{target}, with the content descriptions of T **important** neighbor items, represented as $d_1, d_2, \cdots, d_T$. The importance is measured based on user engagement. We will discuss more details in the *Experiment* section. This fusion of information forms the basis of the prompt, which is designed to leverage user engagement and preferences in generating more contextually relevant content descriptions:

Table 1: Statistics of the evaluation datasets.

Dataset	# Interaction	# Item	# User
MovieLens-1M	1,000,209	3,706	6,040
Recipe	132,246	4,125	2,589

"Summarize the commonalities among the following descriptions: 'd_{target}'; 'd_1; d_2; ... d_T'."

An engagement-guided prompt can assist the Language Model (LLM) in generating more useful content descriptions for content recommendation due to several reasons:

1. **Contextual Relevance:** By incorporating information from both the target item and its important neighbor items, the prompt provides LLM with a broader context and a more comprehensive understanding of the content. This contextual information helps LLM generate descriptions that are more relevant to the specific item and its related items, thereby increasing their usefulness in content recommendation scenarios.

2. **User Preference Alignment:** Including the content descriptions of important neighbor items, which are determined based on user engagement, enables LLM to align with user preferences. By considering items that are frequently engaged by users, the generated content descriptions are more likely to capture the content characteristics and features that are appealing to the target users. This alignment enhances the usefulness of the generated descriptions in effectively recommending items that align with user preferences.

3. **Enhanced Recommendation Quality:** The engagement-guided prompt leverages user engagement data to identify important neighbor items. By including information from these items in the prompt, LLM can potentially uncover meaningful connections, similarities, or relevant aspects between the target item and its neighbors. This can result in more accurate, informative, and high-quality content descriptions, thereby improving the overall performance of the content recommendation system.

2.4 Recommendation-driven + engagement-guided prompting

This type of prompt intends to incorporate both the recommendation-driven and engagement-guided instructions (Figure 1), which we denote as $p^{rec+eng}$. The prompt is designed as following:

"The description of an item is as follows: 'd_{target}'. What should I say if I want to recommend it to others? This content is considered to hold some similar attractive characteristics as the following descriptions: 'd_1; d_2; ... d_T'."

3 Experiment

3.1 Experiment Setup

We investigate the four prompting strategies for large language models on two widely-used recommendation benchmarks.

3.1.1 Benchmarks

Two datasets are used to evaluate the effect of prompting strategies on input augmentation. Their statistics are shown in Table 1.

1. **MovieLens-1M** [6] is a highly recognized benchmark dataset commonly used for evaluating item recommendation systems. It contains a vast collection of 1,000,209 ratings provided by 6,040 MovieLens users, covering 3,900 movies. Each user has at least 20 ratings. Following He *et al.* [7], we convert the rating data into implicit feedback. More specifically, each entry is marked as 0 or 1 indicating whether the user has rated the corresponding item. The original movie data only contain movie titles and genres. We employ GPT-3 (`text-davinci-003`) to generate the content description of each movie using the following prompt:

"Summarize the movie {title} with one sentence. The answer cannot include the movie title."

The response from GPT-3 is used as the content description. `Temperature` is set at 0 to generate more focused and deterministic responses.

2. **Recipe** [11] is another benchmark dataset we use to assess the recommendation performance. This dataset consists of recipe details and reviews sourced from `Food.com`. The metadata includes ratings, reviews, recipe names, descriptions, ingredients, directions, and so on. For instance, an example recipe description is *"all of the flavors of mac n' cheese in the form of a hot bowl of soup!"*. In our evaluation, we employ the recipe descriptions as item descriptions for the four prompting strategies. Similar to the **MovieLens-1M** dataset, we apply filtering criteria, excluding users with fewer than 20 ratings and items with fewer than 30 ratings.

3.1.2 Item module

- **Response generation:** In our evaluation, we focus on assessing the performance of **GPT-3** [1], particularly the variant known as `text-davinci-003`. This model is an advancement over the InstructGPT models [12], incorporating several improvements. We specifically select this variant due to its ability to consistently generate high-quality writing, effectively handle complex instructions, and demonstrate enhanced proficiency in generating longer form content [13].

- **Text encoder:** We use Sentence-BERT [14] to derive the textual embeddings from the original content description and augmented text. The embedding model is `all-MiniLM-L6-v2`.

- **Importance measurement for engagement-guided prompting:** In our study, we show an example of use Personalized PageRank (PPR) score as the importance measurement. In particular, we first construct the user-item bipartite graph $G = (V, E)$. In this notation, G represents the bipartite graph, E denotes the set of nodes, and E represents the set of edges. There are two types of nodes including users $V_{user} \subset V$ and items $V_{item} \subset V$ ($V_{user} \cup V_{item} = V, V_{user} \cap V_{item} = \varnothing$). An edge $e \in E$ between a user node $v_{user} \in V_{user}$ and an item node $v_{item} \in V_{item}$ is created if this user interacts with this item. Next, we proceed by calculating the Personalized PageRank (PPR) score for each item node, which quantifies their relative importance from an individual node's perspective. For every item node, we construct a set of significant neighboring items. By identifying the top T item nodes with the highest PPR scores that share the same genre as the target item node, we pinpoint essential neighbor items guided by user engagement. The rationale behind this approach lies in the observation that when users frequently engage with two items, there tends to be a greater similarity in terms of user preferences. By incorporating this information, we aim to capture user preferences more effectively, leading to enhanced performance in content recommendation. For both datasets, we set $T = 3$.

3.1.3 User module

We employ an embedding table to convert user ID into latent representations. For both **MovieLens-1M** and **Recipe**, the output dimension is set at 128.

3.1.4 Recommendation module

In our study, we explore four recommendation modules.

- **ItemPop**: This method makes recommendation based on item popularity.

- **MLP**: This recommendation module is a combination of Multi-Layer Perceptron (MLP) and dot product. For simplicity, we refer to it as MLP. The augmented text embeddings and the original content description embeddings are combined by concatenation and then passed through a two-layer MLP. The first MLP layer's output dimension, as well as the input/output dimensions of the second MLP layer, are all set to 128. A ReLU activation function and a dropout layer are applied to the first MLP layer. Next, the dot product of the latent embeddings of the user and the item is calculated, and the resulting value is then passed through a Sigmoid function. This Sigmoid function transforms the dot product into a final relevance score between the user and the item.

- **AutoInt** [16]: A multi-head self-attentive neural network with residual connections is proposed to explicitly model the feature interactions in the low-dimensional space.
- **DCN-V2** [18]: **DCN** [17] uses feature crossing explicitly at each layer. For our experiments, we employ the improved version of DCN, namely **DCN-V2** [18].

3.1.5 Model training.

To train the model, we employ the Binary Cross Entropy Loss. Each user-item interaction is considered as a positive sample. Each user-item interaction within the dataset is treated as a positive sample. In addition to positive samples, we randomly select negative samples by pairing users and items that do not have any recorded interactions. To prevent overfitting and optimize training efficiency, we employ an early stop mechanism. It is worth noting that we have also explored the possibility of using the Bayesian Personalized Ranking (BPR) Loss [15] within the framework. However, after experimentation, we find that the BPR Loss does not yield superior performance compared to the Binary Cross Entropy Loss. As a result, we choose to use the Binary Cross Entropy Loss as our primary loss function.

3.1.6 Evaluation protocols.

To assess the recommendation performance, we adopt the evaluation methodology employed by Wei *et al*. [19]. Initially, we randomly divide the dataset into training, validation, and testing sets using an 8:1:1 ratio. Negative training samples are created using random negative sampling, as mentioned earlier.

For the validation and testing sets, we pair each observed user-item interaction with 1,000 items that the user has not previously interacted with. It is important to note that there is *no* overlap between the negative samples in the training set and the unobserved user-item pairs in the validation and testing sets. This ensures the independence of the evaluation data.

To evaluate the performance of top-K recommendations, we employ widely-used metrics such as Precision@K, Recall@K, and NDCG@K. In our case, we set $K = 10$, indicating that we consider the top 10 recommendations. We report the average scores across five different splits of the testing sets, providing a comprehensive evaluation of the recommendation performance.

3.1.7 Hyper-parameter settings.

We initialize the model parameters randomly, following a Gaussian distribution. To optimize the framework, we employ the AdamW algorithm [10] with a weight decay value of 0.0005. The hyper-parameter grids for the learning rate and dropout rate are discussed in the Appendix. Settings that achieve the highest Recall@K on the validation set are chosen for the evaluation on the testing set.

3.1.8 Implementation details.

Our methods are implemented and experiments are conducted using PyTorch. The computation of PPR scores is facilitated by the use of the `torch-ppr` library. The experiments are conducted on a NVIDIA A100 GPU with 80 GB of memory. Each experiment is run on one GPU at a time.

4 Results

Table 2 summarizes the recommendation performance of the baselines. Remarkably, LLM-Rec boosts simple MLP models to achieve superior recommendation performance, surpassing other more complex feature-based recommendation methods.

To understand the effect of each prompting strategy, we design another experiment. Figure 3 shows the evaluation architecture. We keep the recommendation and user modules consistent across all experiments and only change the augmented text generated by our proposed prompting strategies. For each generated response, we first encode it and then concatenate the embeddings with the embeddings of the original content description. The responses generated from the basic prompting (*i.e.*, $\mathbf{p_{para}}$, $\mathbf{p_{tag}}$, $\mathbf{p_{infer}}$), recommendation-driven prompting (*i.e.*, $\mathbf{p_{para}^{rec}}$, $\mathbf{p_{tag}^{rec}}$, $\mathbf{p_{infer}^{rec}}$), engagement-guided prompting (*i.e.*, $\mathbf{p^{eng}}$), and recommendation-driven + engagemen guided prompting (*i.e.*, $\mathbf{p^{rec+eng}}$) are compared.

Table 2: **LLM-Rec empowers simple MLP models to achieve superior recommendation performance, surpassing other more complex feature-based recommendation methods.** The input feature for the MLP, AutoInt, and DCN-V2 models is the embeddings of the original content description. LLM-Rec in this table represents the MLP baseline whose input feature is the concatenation of the embeddings of the original content description and all responses generated by large language models via our proposed prompting strategies. Note that it is still just an MLP model.

	Movielens-1M			Recipe		
	Precision@10	Recall@10	NDCG@10	Precision@10	Recall@10	NDCG@10
ItemPop	0.0426	0.0428	0.0530	0.0116	0.0274	0.0201
MLP	0.2914	0.2440	0.3626	0.0325	0.0684	0.0580
AutoInt	0.2149	0.1706	0.2698	0.0351	0.0772	0.0658
DCN-V2	0.2961	0.2433	0.3689	0.0360	0.0786	0.0653
LLM-Rec	**0.3150**	**0.2766**	**0.3951**	**0.0394**	**0.0842**	**0.0706**
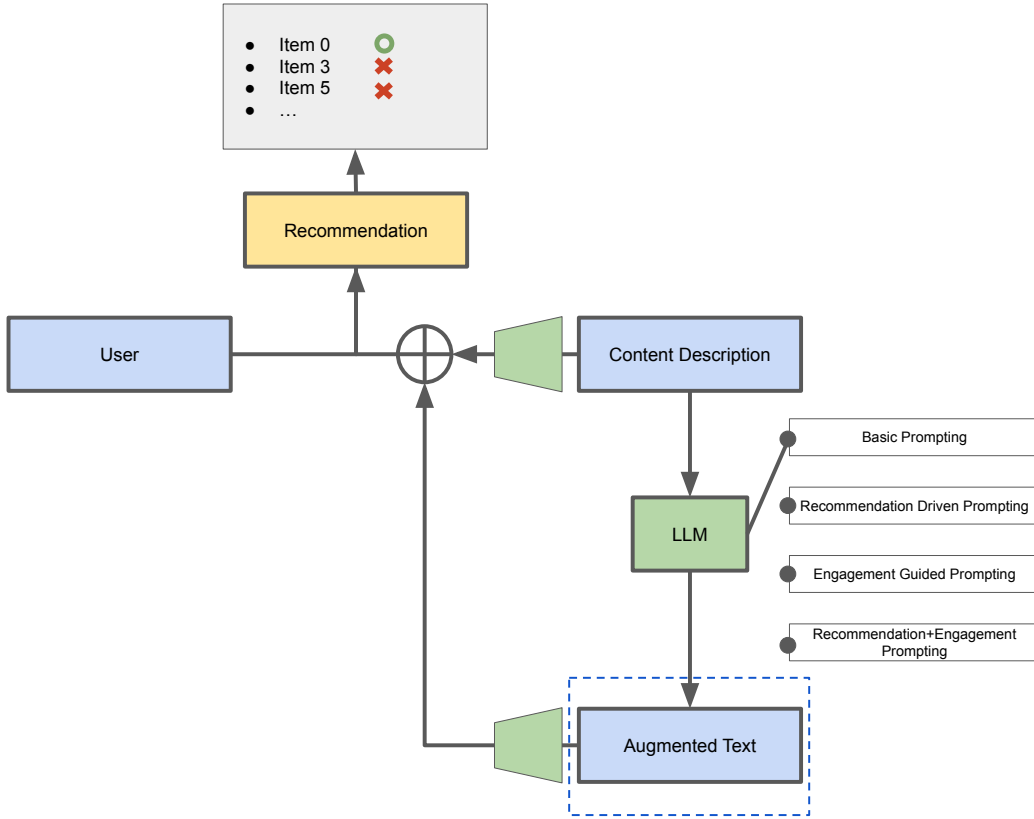
Figure 3: Evaluation architecture. $\oplus$ represents concatenation. Only prompts and corresponding augmented text which are highlighted with dashed lines are different across baselines. Other input and modules **remain consistent** throughout the evaluation process.

CD represents the original content description. No augmented text is introduced. The item input exclusively comprises the embeddings of the original content description.

Table 4 shows the recommendation performance of each prompting strategy. Two key takeaways can be observed from the figure. Firstly, the combination of augmented text and the original content description leads to an improvement in recommendation performance. This finding suggests that all three types of prompting, namely basic, recommendation-driven, and engagement-guided, provide additional and valuable information for the recommendation module to effectively model content recommendation.
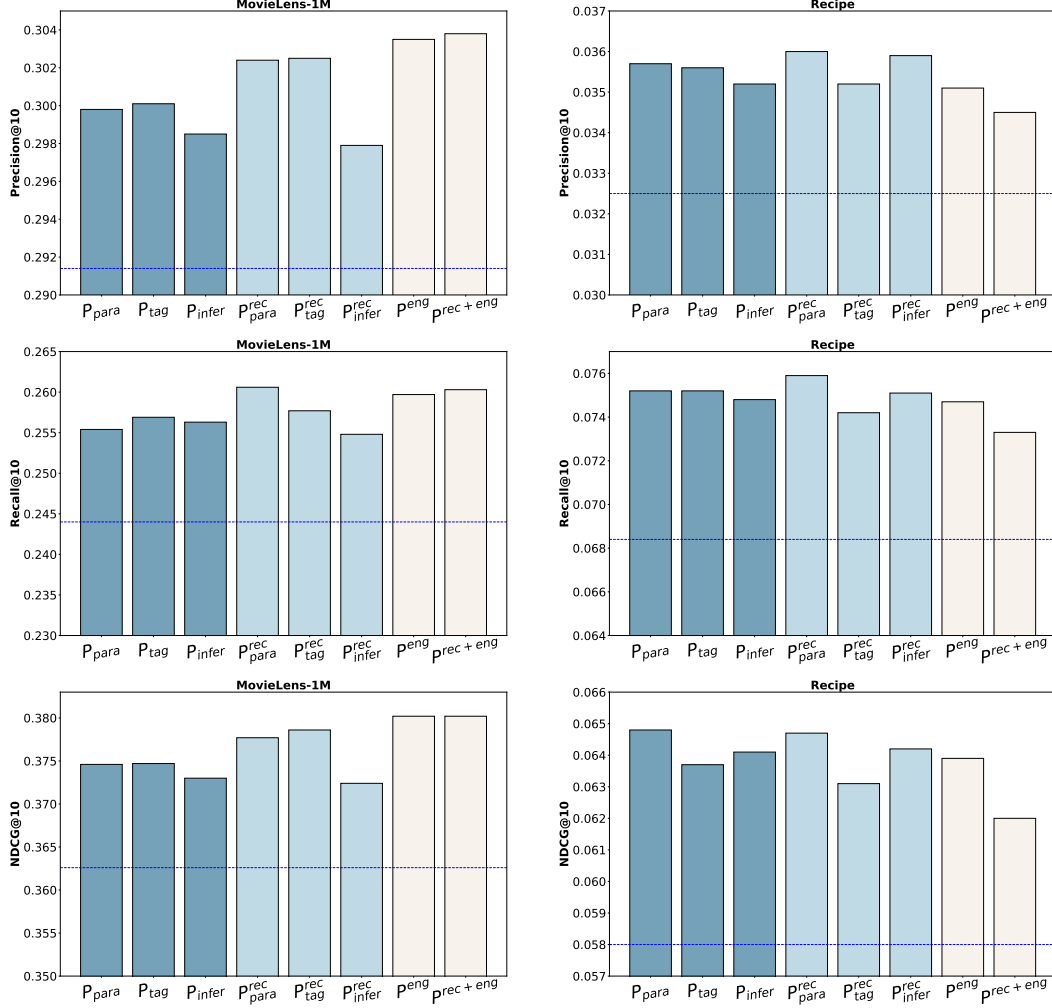
Figure 4: **The ablation study conducted on different prompting strategies shows that augmenting the input text with responses generated by large language models using our proposed prompting strategies enhances recommendation performance. However, the extent of this improvement may vary depending on the characteristics of the datasets used.** The *basic prompting* strategy includes three variants: P_{para} (paraphrase), P_{tag} (tag), and P_{infer} (infer). The *recommendation-driven* versions of these three variants are denoted by P_{para}^{rec}, P_{tag}^{rec}, and P_{infer}^{rec}. In addition, P^{eng} represents the *engagement-guided* prompts, and $P^{rec\text{-}eng}$ stands for the *recommendation + engagement* prompts. The blue line in each figure indicates the performance achieved by using only the original content description CD. Note that in this ablation study, all the baselines use a concatenation of the CD embeddings and prompt response embeddings as their model input. For instance, the bar corresponding to P_{para} represents the recommendation performance of the baseline that takes the concatenation of CD and P_{para} as input. For simplicity, we have omitted CD in this figure.

Secondly, the extent of this improvement may vary depending on the characteristics of the datasets used. To further investigate the reasons the recommendation performances vary across different prompting strategies, we conduct a case study comparing P_{para} with P_{para}^{rec}. For both datasets, we choose three example item descriptions and the corresponding generated responses (Figure 5). The three examples selected are the top three items that the recommendation is correct based on the response of P_{para}^{rec} while incorrect based on the response of P_{para}. We find that the most distinctive words in the response of P_{para}^{rec} are the words that are related with user preferences. These words include the words that can express users' preferences about items such as `exciting`, `thought-provoking`, `delicious`, and so on. We also discover words that are related to the well-
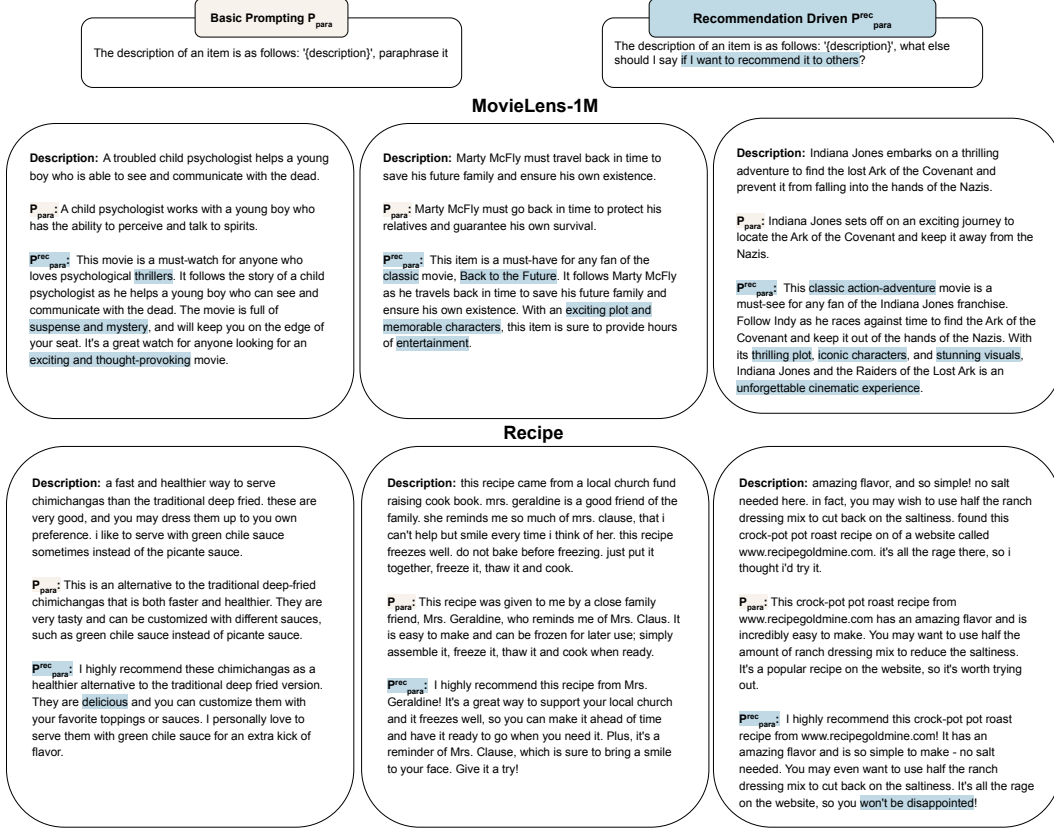
8

Figure 5: **Examples show that the distinctive words generated with the recommendation-driven prompting strategy are related with user preferences.** The augmented words in the response of P_{para}^{rec} are highlighted.

defined concept in terms of user preferences such as genres (*e.g.*, `classic`, `action`, `adventure`). We hypothesize that words generated with the recommendation-driven prompting strategy improve recommendation performance.

To validate this hypothesis, we design two variants of the response, namely P_{para}^{mask} and $P_{para}^{keyword}$. To construct P_{para}^{mask}, we mask the words that appear in the response of P_{para}^{rec} but are absent in the response of P_{para}. To construct $P_{para}^{keyword}$, we append the words that (1) appear in the response of P_{para}^{rec} and (2) are well-defined user-preference-related words such as genres to the end of the response of P_{para}. These two variants of the responses are then fed into MLP models to form baselines. The recommendation performances are shown in Figure 6. Upon a comparative analysis of the performance between P_{para}^{rec} and P_{para}^{mask}, a discernible decline in recommendation performance is observed when words unique to the response of P_{para}^{rec} are selectively masked. This outcome underscores the pivotal role played by the supplementary insights introduced through the augmented text. Furthermore, our investigation reveals that the incorporation of vital keywords, as opposed to the inclusion of all response words, can yield even superior recommendation performance. This phenomenon may be attributed to potential discrepancies or extraneous elements within the response of P_{para}^{rec}.

When comparing p_{infer} and p_{infer}^{rec}, which instruct LLM to deduce the characteristics of the original content description, a *decrease in performance* is observed. This suggests that the reduced recommendation performance may stem from the discrepancy between the inferred context and the original context. In other words, the desired response from p_{infer} requires inference *beyond* the information provided in the original context, making it less effective. Conversely, recommendation-driven prompting proves beneficial for p_{para} and p_{tag}, as these prompts do not rely on LLM inferring information beyond the original input.
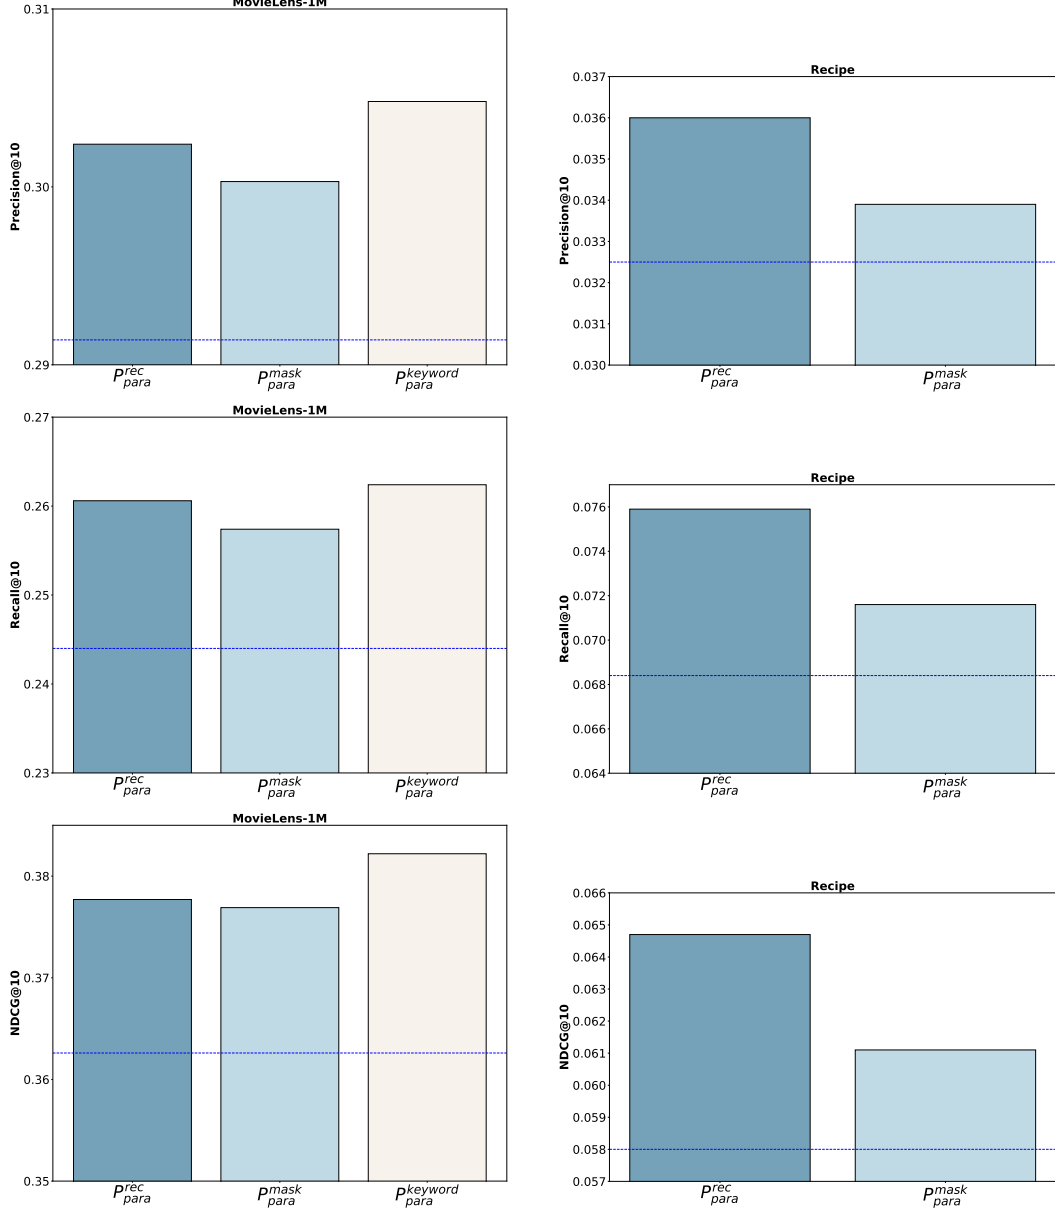
Figure 6: **The ablation study shows that the LLM-augmented words that align with user preferences such as genres boost recommendation performance.** P_{para}^{mask} masks all the words that appear in the response of P_{para}^{rec} but are absent in the response of P_{para}. $P_{para}^{keyword}$ append the user-preference-related words that appear in the response of P_{para}^{rec} onto the end of the response of P_{para}.

When comparing p^{eng} with CD, we observe improvements in recommendation performance. In a manner consistent with our preceding experiments, we curate exemplar responses from p^{eng} for closer scrutiny (Figure 7). Our analysis unveils a distinctive pattern where the response of p^{eng} encompasses descriptive vocabulary intricately tied to user preferences. Remarkably, upon contrasting these lexicons with the terms embedded in the response of p_{para}^{rec}, a discernible refinement is evident. Notably, the composite response stemming from $p^{rec+eng}$, achieved through the synergy of recommendation-driven and engagement-guided prompting strategies, not only integrates user-preference-related words of a more universal nature but also embraces well-defined concepts, such as genres.
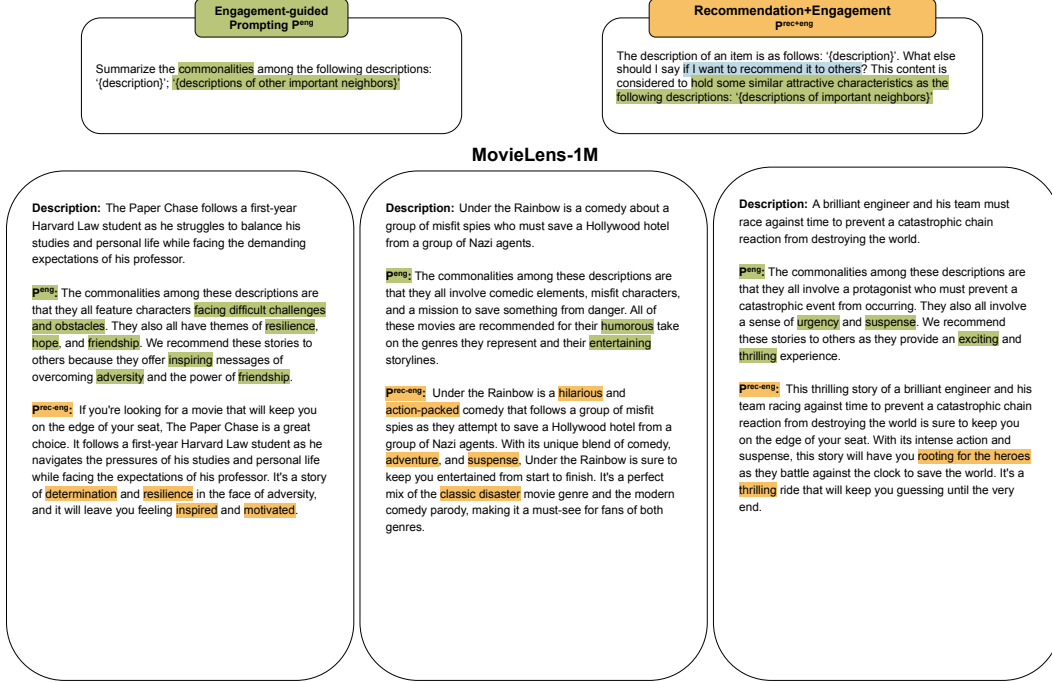
10

Figure 7: **Examples show that the distinctive words generated with the engagement-guided prompting strategy are fine-grained descriptive words related to user preferences.** The augmented words in the response of P^{eng} and $P^{rec+eng}$ are highlighted.

Table 3: **Concatenating embeddings of the responses augmented by LLM-Rec outperforms concatenating randomized embeddings and concatenating duplicate CD embeddings. It also achieves a superior performance than concatenating the raw text.**

	Movielens-1M			Recipe		
	Precision@10	Recall@10	NDCG@10	Precision@10	Recall@10	NDCG@10
CD	0.2914	0.2440	0.3626	0.0325	0.0684	0.0580
Randomizing embeddings	1e-4	8.85e-5	3e-4	0.0	0.0	0.0
Duplicating CD embeddings	0.2858	0.2417	0.3567	0.0327	0.0694	0.0590
Text concatenation	0.3038	0.2615	0.3796	0.0332	0.0714	0.0591
P^{all}	**0.3126**	**0.2731**	**0.3932**	**0.0394**	**0.0842**	**0.0706**

To conduct a more in-depth exploration into the caliber of the combined augmented text, we engage in a process of concatenating the embeddings derived from diverse prompting strategies' responses. This concatenation is performed in various permutations, interwoven with the embeddings of the original content description. We then conduct the same experiment again, searching for hyper-parameters in the similar manner as discussed previously. The results are shown in Figure 8. P^{all} concatenates the embeddings of all responses and the original content description. Overall, concatenating more information helps improve recommendation performance. This finding emphasizes the added value of incorporating augmented text over using the original content description alone.

Table 3 shows the recommendation performances of other concatenation variants:

- **Randomizing embeddings**: We randomize the embeddings of P^{all}.

- **Duplicating CD embeddings**: We concatenate multiple CD embeddings to match the dimension of the embeddings of P^{all}.

- **Text concatenation**: Instead of concatenating the embeddings of all response, we concatenate the responses first, and then convert it to embeddings.
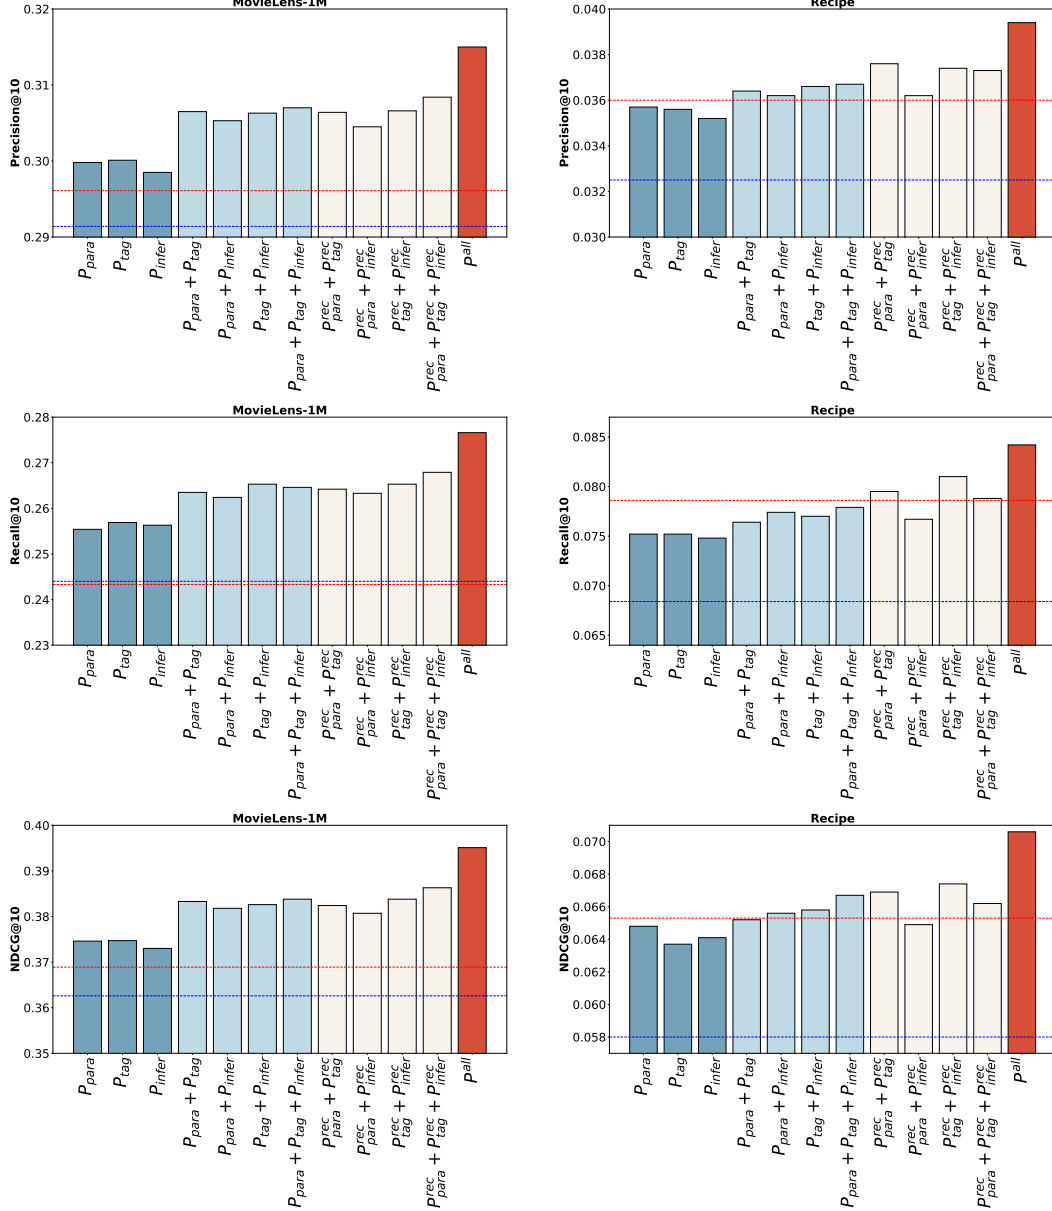
Figure 8: **The ablation study shows that overall, recommendation benefits from concatenating the embeddings of the input text augmented by large language models.** The blue line in each figure indicates the performance achieved by using only the original content description CD, while the red line represents the performance achieved by **DCN-V2** [18].

5 Discussion

In this study, we have investigated the effectiveness of **LLM-Rec** prompting as a straightforward yet impactful mechanism for improving personalized recommendation through large language models. Our findings reveal several key insights.

Firstly, we demonstrate that by combining augmented text with the original content description, we observe a significant enhancement in recommendation performance. It also empowers simple models such as MLPs to achieve superior recommendation performance than other more complex feature-based methods. This highlights the value of incorporating additional context to facilitate more accurate and relevant recommendations, coupled with an easier training process.

Furthermore, our experimental results on recommendation-driven and engagement-guided prompting strategies illustrate their ability to encourage the large language model to generate high-quality input text specifically tailored for recommendation purposes. These prompting strategies effectively leverage recommendation goals and user engagement signals to guide the model towards producing more desirable recommendations. More specifically, the recommendation-driven prompting strategy engenders a spectrum of broader user-preference-associated terms, including well-established concepts. This phenomenon signifies its adeptness at tapping into the global comprehension of the recommendation objective concerning the specific item to be suggested. On the other hand, the engagement-guided prompting strategy, integrating more immediate user co-engagement signals, encapsulates the capacity of LLMs to grasp nuanced, finely detailed, and localized insights about the item to be recommended.

Lastly, by combining all augmented text, we achieve the best overall recommendation performance. This suggests the complementary nature of these strategies and their collective impact in further improving recommendation quality.

Overall, our study showcases the effectiveness of **LLM-Rec** prompting in facilitating large language models to generate enhanced and relevant input text for personalized recommendation. These findings contribute to the advancement of recommendation systems, emphasizing the significance of thoughtful prompt design to enhance recommendation performance.

Throughout our experimental analysis, we also uncover a potential limitation when employing LLM for augmenting input text in recommendation systems. We notice a distinction between prompts that solely instruct LLM to modify the content description and those that prompt LLM to infer additional information. In the latter case, where inference beyond the original context is required, the recommendation-driven prompting strategy may not yield the expected benefits. Our hypothesis suggests that the quality, specifically in terms of recommendation relevance, of the inferred context might have an unknown impact on the overall recommendation performance.

This observation emphasizes the need for careful consideration and evaluation of the prompts employed, particularly when instructing LLM to infer information beyond the provided context. While recommendation-driven prompting strategies prove effective for prompts that do not necessitate inference, their effectiveness may be hindered when the prompts require LLM to extrapolate information. Further research is necessary to explore techniques for managing and improving the quality and impact of inferred context on recommendation outcomes.

In addition to its superior performance in personalized content recommendation, the incorporation of engagement signals in prompt designs may have broader associated benefits. The engagement-guided prompting strategy instructs the LLM to generate commonalities among different items, resembling the concept of neighborhood aggregation in Graph Neural Network (GNN) learning. In GNN, each target node is partially learned by aggregating information from its neighboring nodes. In this context, we highlight the potential of using engagement-guided prompts as a means to replace the learning process of GNN, thereby simplifying the overall model architecture.

Furthermore, leveraging the fine-tuned LLM opens up possibilities for zero-shot generation without incurring any additional learning cost. Since the LLM has already undergone training to capture linguistic patterns and semantic understanding, it can be harnessed to generate responses or recommendations in unseen scenarios without requiring further training. This zero-shot generation capability enables flexibility and scalability in recommendation systems, allowing for efficient adaptation to new domains or contexts.

The combination of engagement-guided prompting and the zero-shot generation potential of the fine-tuned LLM presents promising opportunities for streamlining model architectures, reducing computational complexity, and expanding the applicability of recommendation systems. Further exploration and investigation in this direction could unlock novel techniques for efficient and effective personalized recommendation.

6   Related Work

In addition to leveraging LLMs directly as recommenders, there have been efforts to use LLMs for augmenting the input side of personalized recommendation. For instance, Chen *et al.* [2] incorporated user history behaviors, such as clicks, purchases, and ratings, into LLMs to generate

user profiles. These profiles were then combined with the history interaction sequence and candidate items to construct the final recommendation prompt. LLMs were subsequently employed to predict the likelihood of user-item interaction based on this prompt. Xi *et al.* [20] introduced a method that leverages the reasoning knowledge of LLMs regarding user preferences and the factual knowledge of LLMs about items. However, our study takes a different approach, focusing specifically on input augmentation for items with LLMs' reasoning ability. By employing prompting strategies, we aim to generate augmented input text that better captures the characteristics and nuances of items, leading to improved personalized recommendations.

7   Conclusions

In this study, we introduce **LLM-Rec** prompting strategies, which leverage large language models (LLMs) for input augmentation, aiming to enhance personalized recommendation. Through rigorous experimentation across four variants of **LLM-Rec**, we observe that the combination of augmented input text and original content descriptions yields notable improvements in recommendation performance.

These findings emphasize the potential of using LLMs and strategic prompting techniques to enhance the accuracy and relevance of personalized recommendation with an easier training process. By incorporating additional context through augmented text, we enable the recommendation algorithms to capture more nuanced information and generate recommendations that better align with user preferences.

The experimental results of **LLM-Rec** highlights the importance of innovative approaches in leveraging LLMs for content recommendation and showcases the value of input augmentation in improving recommendation performance. As personalized recommendation continues to play a pivotal role in various domains, our study provides insights into effective strategies for leveraging LLMs to deliver enhanced recommendation experiences.

References

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

[2] Zheng Chen. Palr: Personalization aware llms for recommendation. *arXiv preprint arXiv:2305.07622*, 2023.

[3] Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. Uncovering chatgpt's capabilities in recommender systems. *arXiv preprint arXiv:2305.02182*, 2023.

[4] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524*, 2023.

[5] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems*, pages 299–315, 2022.

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

[7] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pages 173–182, 2017.

[8] Ruyu Li, Wenhao Deng, Yu Cheng, Zheng Yuan, Jiaqi Zhang, and Fajie Yuan. Exploring the upper limits of text-based collaborative filtering using large language models: Discoveries and insights. *arXiv preprint arXiv:2305.11700*, 2023.

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

[10] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

[11] Bodhisattwa Prasad Majumder, Shuyang Li, Jianmo Ni, and Julian McAuley. Generating personalized recipes from historical user preferences. *arXiv preprint arXiv:1909.00105*, 2019.

[12] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.

[13] Raf. How do text-davinci-002 and text-davinci-003 differ? *OpenAI*, 2023.

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

[15] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*, 2012.

[16] Weiping Song, Chence Shi, Zhiping Xiao, Zhijian Duan, Yewen Xu, Ming Zhang, and Jian Tang. Autoint: Automatic feature interaction learning via self-attentive neural networks. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1161–1170, 2019.

[17] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. Deep & cross network for ad click predictions. In *Proceedings of the ADKDD'17*, pages 1–7. 2017.

[18] Ruoxi Wang, Rakesh Shivanna, Derek Cheng, Sagar Jain, Dong Lin, Lichan Hong, and Ed Chi. Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the web conference 2021*, pages 1785–1797, 2021.

[19] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. Mmgcn: Multi-modal graph convolution network for personalized recommendation of micro-video. In *Proceedings of the 27th ACM international conference on multimedia*, pages 1437–1445, 2019.

[20] Yunjia Xi, Weiwen Liu, Jianghao Lin, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. Towards open-world recommendation with knowledge augmentation from large language models. *arXiv preprint arXiv:2306.10933*, 2023.

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

A Supplemental Material

A.1 Prompts

- p_1: The description of an item is as follows: '{item description}', paraphrase it.
- p_1^{rec}: The description of an item is as follows: '{item description}', what else should I say if I want to recommend it to others?
- p_2: The description of an item is as follows: '{item description}', summarize it with tags.
- p_2^{rec}: The description of an item is as follows: '{item description}', what tags should I use if I want to recommend it to others?
- p_3: The description of an item is as follows: '{item description}', what kind of emotions can it evoke?
- p_3^{rec}: The description of an item is as follows: '{item description}', recommend it to others with a focus on the emotions it can evoke.
- p^{eng}: Summarize the commonalities among the following descriptions: '{item description}'; '{descriptions of other important neighbors}'.
- $p^{rec+eng}$: The description of an item is as follows: '{item description}'. What else should I say if I want to recommend it to others? This content is considered to hold some similar attractive characteristics as the following descriptions: '{descriptions of other important neighbors}'.

A.2   Hyper-parameters

For the MLP model, the hyper-parameter grids for the learning rate and dropout rate are $\{0.0001, 0.0005, 0.001\}$ and $\{0.1, 0.3, 0.5\}$, respectively. For AutoInt, the hyper-parameter grids for the learning rate, dropout rate, hidden layer size, number of attention layers, and attention heads are $\{0.001, 0.005, 0.01\}$, $\{0.1, 0.3, 0.5\}$, $\{16, 32, 64, 128\}$, $\{1, 2\}$, and $\{1, 2\}$, respectively. For DCN-V2, the learning rate, dropout rate, hidden layer size, and number of cross layers are searched in $\{0.001, 0.005, 0.01\}$, $\{0.1, 0.3, 0.5\}$, $\{16, 32, 64, 128\}$, $\{1, 2\}$, and $\{1, 2\}$, respectively. The performance is evaluated every five epochs, and the early stop mechanism is configured to have a patience of 5. Additionally, we set the batch size to 4096 for all baselines except for AutoInt which is 1024 due to the memory limitation.