

LLMRec: Large Language Models with Graph Augmentation for Recommendation

Wei Wei
University of Hong Kong
weiweics@connect.hku.hk

Xubin Ren
University of Hong Kong
xubinren@connect.hku.hk

Jiabin Tang
University of Hong Kong
jiabintang77@gmail.com

Qinyong Wang
Baidu Inc.
wangqinyong@baidu.com

Lixin Su
Baidu Inc.
sulixini@baidu.com

Suqi Cheng
Baidu Inc.
chengsuqi@gmail.com

Junfeng Wang
Baidu Inc.
wangjunfeng@baidu.com

Dawei Yin
Baidu Inc.
yindawei@acm.org

Chao Huang*
University of Hong Kong
chaohuang75@gmail.com

ABSTRACT

The problem of data sparsity has long been a challenge in recommendation systems, and previous studies have attempted to address this issue by incorporating side information. However, this approach often introduces side effects such as noise, availability issues, and low data quality, which in turn hinder the accurate modeling of user preferences and adversely impact recommendation performance. In light of the recent advancements in large language models (LLMs), which possess extensive knowledge bases and strong reasoning capabilities, we propose a novel framework called LLMRec that enhances recommender systems by employing three simple yet effective LLM-based graph augmentation strategies. Our approach leverages the rich content available within online platforms (e.g., Netflix, MovieLens) to augment the interaction graph in three ways: (i) reinforcing user-item interaction edge, (ii) enhancing the understanding of item node attributes, and (iii) conducting user node profiling, intuitively from the natural language perspective. By employing these strategies, we address the challenges posed by sparse implicit feedback and low-quality side information in recommenders. Besides, to ensure the quality of the augmentation, we develop a denoised data robustification mechanism that includes techniques of noisy implicit feedback pruning and MAE-based feature enhancement that help refine the augmented data and improve its reliability. Furthermore, we provide theoretical analysis to support the effectiveness of LLMRec and clarify the benefits of our method in facilitating model optimization. Experimental results on benchmark datasets demonstrate the superiority of our LLM-based augmentation approach over state-of-the-art techniques. To

ensure reproducibility, we have made our code and augmented data publicly available at: <https://github.com/HKUDS/LLMRec.git>.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Large Language Models, Graph Learning, Data Augmentation, Content-based Recommendation, Multi-modal Recommendation, Collaborative Filtering, Data Sparsity, Bias in Recommender System

ACM Reference Format:

Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. LLMRec: Large Language Models with Graph Augmentation for Recommendation. In *Proceedings of The 17th ACM International Conference on Web Search and Data Mining (WSDM '2024)*. ACM, Merida, Mexico, 10 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Recommender systems play a crucial role in mitigating information overload by providing online users with relevant content [4, 30, 35]. To achieve this, an effective recommender needs to have a precise understanding of user preferences, which is not limited to analyzing historical interaction patterns but also extends to incorporating rich side information associated with users and items [52].

In modern recommender systems, such as Netflix, the side information available exhibits heterogeneity, including item attributes [44], user-generated content [8], and multi-modal features [43] encompassing both textual and visual aspects. This diverse content offer distinct ways to characterize user preferences. By leveraging such side information, models can obtain informative representations to personalize recommendations. However, despite significant progress, these methods often face challenges related to data scarcity and issues associated with handling side information.

Sparse Implicit Feedback Signals. Data sparsity and the cold-start problem hinder collaborative preference capturing [39]. While many efforts (e.g., NGCF [32], LightCGN [12]) tried powerful graph neural networks (GNNs) in collaborative filtering (CF), they face limits due to insufficient supervised signals. Some studies [27] used contrastive learning to add self-supervised signals (e.g., SGL [42],

*Chao Huang is the corresponding author.

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

WSDM '2024, March 4th-8th, 2024, Merida, Yucatan, Mexico

© 2024 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

SimGCL [45]). However, considering that real-world online platforms (e.g., Netflix, MovieLens) derive benefits from modal content, recent approaches, unlike general CF, are dedicated to incorporating side information as auxiliary for recommenders. For example, MMGCN [41] and GRCN [40] incorporate item-end content into GNNs to discover high-order content-aware relationships. LATTICE [50] leverages auxiliary content to conduct data augmentation by establishing i-i relationships. Recent efforts (e.g., MSSL [36], MICRO [49]) address sparsity by introducing self-supervised tasks that maximize the mutual information between multiple content-augmented views. However, strategies for addressing data sparsity in recommender systems, especially in multi-modal content, can sometimes be limited. This is because the complexity and lack of side information relevance to CF can introduce distortions in the underlying patterns [40]. Therefore, it becomes crucial to ensure the accurate capture of realistic user preferences when incorporating side information in CF, in order to avoid suboptimal results.

Data Quality Issues of Side Information. Recommender systems that incorporate side information often encounter significant issues that can negatively impact their performance. i) **Data Noise** is an important limitation faced by recommender systems utilizing side information is the issue of data noise, where attributes or features may lack direct relevance to user preferences. For instance, in a micro video recommender, the inclusion of irrelevant textual titles that fail to capture the key aspects of the video’s content introduces noise, adversely affecting representation learning. The inclusion of such invalid information confuse the model and lead to biased or inaccurate recommendations. ii) **Data heterogeneity**[5] arises from the integration of different types of side information, each with its own unique characteristics, structures, and representations. Ignoring this heterogeneity leads to skewed tensor space distributions [44]. Bridging heterogeneous gap is crucial for successfully incorporating side information uniformly. iii) **Data incompleteness** [16] occurs when side information lacks certain attributes or features. For instance, privacy concerns[47] may make it difficult to collect sufficient user profiles to learn their interests. Additionally, items may have incomplete textual descriptions or missing key attributes. This incompleteness impairs the model’s ability to fully capture the unique characteristics of users and items, thereby affecting the accuracy of recommendations.

Having gained insight into data sparsity and low-quality encountered by modern recommenders with auxiliary content, this work endeavors to overcome these challenges through explicit augment potential user-item interactive edges as well as enhances user/item node side information (e.g., language, genre). Inspired by the impressive natural language understanding ability of large language models (LLMs), we utilize LLMs to augment the interaction graph. *Firstly*, LLMRec embraces the shift from an ID-based recommendation framework to a modality-based paradigm [18, 46]. It leverages large language models (LLMs) to predict user-item interactions from a natural language perspective. Unlike previous approaches that rely solely on IDs, LLMRec recognizes that valuable item-related details are often overlooked in datasets [19]. Natural language representations provide a more intuitive reflection of user preferences compared to indirect ID embeddings. By incorporating LLMs, LLMRec captures the richness and context of natural

language, enhancing the accuracy and effectiveness of recommendations. *Secondly*, to elaborate further, the low-quality and incomplete side information is enhanced by leveraging the extensive knowledge of LLMs, which brings two advantages: i) LLMs are trained on vast real-world knowledge, allowing them to understand user preferences and provide valuable completion information, even for privacy-constrained user profiles. ii) The comprehensive word library of LLMs unifies embeddings in a single vector space, bridging the gap between heterogeneous features and facilitating encoder computations. This integration prevents the dispersion of features across separate vector spaces and provide more accurate results.

Enabling LLMs as effective data augmentors for recommenders poses several technical challenges that need to be addressed:

- **C1:** How to enable LLMs to reason over user-item interaction patterns by explicitly augmenting implicit feedback signals?
- **C2:** How to ensure the reliability of the LLM-augmented content to avoid introducing noise that could compromise the results? The potential of LLM-based augmentation to enhance recommenders by addressing sparsity and improving incomplete side information is undeniable. However, effectively implementing this approach requires addressing the aforementioned challenges. Hence, we have designed a novel framework **LLMRec** to tackle these challenges.

Solution. Our objective is to address the issue of sparse implicit feedback signals derived from user-item interactions while simultaneously improving the quality of side information. Our proposed LLMRec incorporates three LLM-based strategies for augmenting the interaction graph: *i)* Reinforcing user-item interaction edges, *ii)* Enhancing item attribute modeling, and *iii)* Conducting user profiling. To tackle **C1** for ‘*i)*’, we devise an LLM-based Bayesian Personalized Ranking (BPR)[28] sampling algorithm. This algorithm uncover items that users may like or dislike based on textual content from from natural language perspective. These items are then used as positive and negative samples in the BPR training process. It is important to note that LLMs are unable to perform all-item ranking, so the selected items are chosen from a candidate item pool provided by the base recommender for each user. During the node attribute generation process (corresponding to ‘*ii)*’ and ‘*iii)*’), we create additional attributes for each user/item using existing text and interaction history. However, it is important to acknowledge that both the augmented edges and node features can contain noise. To address **C2**, our denoised data robustification mechanism comes into play by integrating noisy edge pruning and feature MAE to ensure the quality of the augmented data.

In summary, our contributions can be outlined as follows:

- The LLMRec is the pioneering work that using LLMs for graph augmentation in recommender by augmenting: user-item interaction edges, ii) item node attributes, iii) user node profiles.
- The proposed LLMRec addresses the scarcity of implicit feedback signals by enabling LLMs to reason explicitly about user-item interaction patterns. Additionally, it resolves the low-quality side information issue through user/item attribute generation and a denoised augmentation robustification mechanism with the noisy feedback pruning and MAE-based feature enhancement.
- Our method has been extensively evaluated on real-world datasets, demonstrating its superiority over state-of-the-art baseline methods. The results highlight the effectiveness of our approach in

improving recommendation accuracy and addressing sparsity issues. Furthermore, in-depth analysis and ablation studies provide valuable insights into the impact of our LLM-enhanced data augmentation strategies, further solidifying the model efficacy.

2 PRELIMINARY

Recommendation with Graph Embedding. Collaborative filtering (CF) learns from sparse implicit feedback \mathcal{E}^+ , with the aim of learning collaborative ID-corresponding embeddings $\mathbf{E}_u, \mathbf{E}_i$ for recommender prediction, given user $u \in \mathcal{U}$ and item $i \in \mathcal{I}$. Recent advanced recommenders employ GNNs to model complex high-order[32] u-i relation by taking \mathcal{E}^+ as edges of sparse interactive graph. Therefore, the CF process can be separated into two stages, bipartite graph embedding, and u-i prediction. Optimizing collaborative graph embeddings $\mathbf{E} = \{\mathbf{E}_u, \mathbf{E}_i\}$ aims to maximize the posterior estimator with \mathcal{E}^+ , which is formally presented below:

$$\mathbf{E}^* = \arg \max_{\mathbf{E}} p(\mathbf{E}|\mathcal{E}^+) \quad (1)$$

Here, $p(\mathbf{E}|\mathcal{E}^+)$ is to encode as much u-i relation from \mathcal{E}^+ into $\mathbf{E}_u, \mathbf{E}_i$ as possible for accurate u-i prediction $\hat{y}_{u,i} = \mathbf{e}_u \cdot \mathbf{e}_i$.

Recommendation with Side Information. However, sparse interactions in \mathcal{E}^+ pose a challenge for optimizing the embeddings. To handle data sparsity, many efforts introduced side information in form of node features \mathbf{F} , by taking recommender encoder f_{Θ} as feature graph. The learning process of the f_{Θ} (including $\mathbf{E}_u, \mathbf{E}_i$ and feature encoder) with side information \mathbf{F} is formulated as maximizing the posterior estimator $p(\Theta|\mathbf{F}, \mathcal{E}^+)$:

$$\Theta^* = \arg \max_{\Theta} p(\Theta|\mathbf{F}, \mathcal{E}^+) \quad (2)$$

f_{Θ} will output the final representation \mathbf{h} contain both collaborative signals from \mathbf{E} and side information from \mathbf{F} , i.e., $\mathbf{h} = f_{\Theta}(\mathbf{f}, \mathcal{E}^+)$.

Recommendation with Data Augmentation. Despite significant progress in incorporating side information into recommender, introducing low-quality side information may even undermine the effectiveness of sparse interactions \mathcal{E}^+ . To address this, our LLMRec focuses on user-item interaction feature graph augmentation, which involves LLM-augmented u-i interactive edges $\mathcal{E}_{\mathcal{A}}$, and LLM-generated node features $\mathbf{F}_{\mathcal{A}}$. The optimization target with augmented interaction feature graph is as:

$$\Theta^* = \arg \max_{\Theta} p(\Theta|\{\mathbf{F}, \mathbf{F}_{\mathcal{A}}\}, \{\mathcal{E}^+, \mathcal{E}_{\mathcal{A}}\}) \quad (3)$$

The recommender f_{Θ} input union of original and augmented data, which consist of edges $\{\mathcal{E}^+, \mathcal{E}_{\mathcal{A}}\}$ and node features $\{\mathbf{F}, \mathbf{F}_{\mathcal{A}}\}$, and output quality representation \mathbf{h} to predicted preference scores $\hat{y}_{u,i}$ by ranking the likelihood of user u will interact with item i .

3 METHODOLOGY

To conduct LLM-based augmentation, in this section, we address these questions: **Q1:** How to enable LLMs to predict u-i interactive edges? **Q2:** How to enable LLMs to generate valuable content? **Q3:** How to incorporate augmented contents into original graph contents? **Q4:** How to make model robust to the augmented data?

3.1 LLMs as Implicit Feedback Augmentor (Q1)

To directly confront the scarcity of implicit feedback, we employ LLM as a knowledge-aware sampler to sample pair-wise [28] u-i training data from a natural language perspective. This increases

potential effective supervision signals and helps gain a better understanding of user preferences by integrating contextual knowledge into the u-i interactions. Specifically, we feed each user’s historical interacted items with side information (e.g., year, genre) and an item candidates pool $C_u = \{i_{u,1}, i_{u,2}, \dots, i_{u,|C_u|}\}$ into LLM. LLM then is expected to select items that user u might be likely (i_u^+) or unlikely (i_u^-) to interact with from C_u . Here, we introduce C_u because LLMs can’t rank all items. Selecting items from the limited candidate set recommended by the base recommender (e.g., MMSL [36], MICRO [49]), is a practical solution. These candidates C_u are hard samples with high prediction score \hat{y}_{ui} to provide potential, valuable positive samples and hard negative samples. It is worth noting that we represent each item using textual format instead of ID-corresponding indexes [19]. This kind of representation offers several advantages: (1) It enables recommender to fully leverage the content in datasets, and (2) It intuitively reflects user preferences. The process of augmenting user-item interactive edges and incorporating it into the training data can be formalized as:

$$i_u^+, i_u^- = LLM(\mathbb{P}_u^{UI}); \quad \mathcal{E}_{BPR} = \mathcal{E} \cup \mathcal{E}_{\mathcal{A}} \quad (4)$$

where i_u^+, i_u^- are positive and negative samples for BPR selected by LLMs from candidates C_u for user u based on input prompt \mathbb{P}_u^{UI} . The augmented dataset $\mathcal{E}_{\mathcal{A}}$ comprises pairwise training triplets (u, i_u^+, i_u^-) , i.e., $\mathcal{E}_{\mathcal{A}} = \{(u, i_u^+, i_u^-) | (u, i_u^+) \in \mathcal{E}_{\mathcal{A}}^+, (u, i_u^-) \in \mathcal{E}_{\mathcal{A}}^-\}$. The textual u-i augmentation prompt \mathbb{P}_u^{UI} encompasses different components: i) task description, ii) historical interactions, iii) candidates, and iv) output format description, as illustrated in Fig. 2 (a).

The utilization of LLMs-based sampler in this study to some extent alleviate *noise* (i.e., false positive) and *non-interacted items issue* (i.e., false negative) [2, 17] exist in raw implicit feedback. In this context, (i) **false positive** are unreliable u-i interactions, which encompass items that were not genuinely intended by the user, such as accidental clicks or instances influenced by popularity bias [31]; (ii) **false negative** represented by non-interacted items, which may not necessarily indicate user dispreference but are conventionally treated as negative samples [3]. By taking LLMs as implicit feedback augmentor, LLMRec enables the acquisition of more meaningful and informative samples by leveraging the remarkable reasoning ability of LLMs with the support of LLMs’ knowledge. The specific analysis is supported by theoretical discussion in Sec. 3.4.1.

3.2 LLM-based Side Information Augmentation

3.2.1 User Profiling & Item Attribute Enhancing (Q2). Leveraging knowledge base and reasoning abilities of LLMs, we propose to summarize user profiles by utilizing users’ historical interactions and item information to overcome limitation of privacy. Additionally, the LLM-based item attributes generation aims to produce space-unified, and informative item attributes. Our LLM-based side information augmentation paradigm consists of two steps:

- i) **User/Item Information Refinement.** Using prompts derived from the dataset’s interactions and side information, we enable LLM to generate user and item attributes that were not originally part of the dataset. Specific examples are shown in Fig. 2(b)(c).
- ii) **LLM-enhanced Semantic Embedding.** The augmented user and item information will be encoded as features and used as input for the recommender. Using LLM as an encoder offers efficient and state-of-the-art language understanding, enabling profiling user interaction preferences and debiasing item attributes.

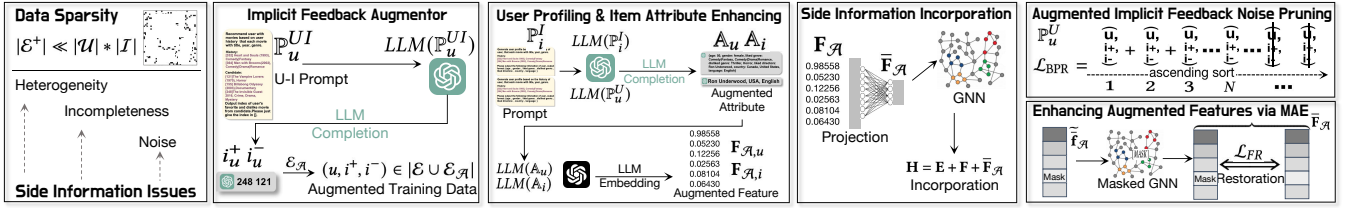


Figure 1: The LLMRec framework: (1) Three types of data augmentation strategies: i) augmenting user-item interactions; ii) enhancing item attributes, and iii) user profiling. (2) Augmented training with and denoised data robustification mechanism.

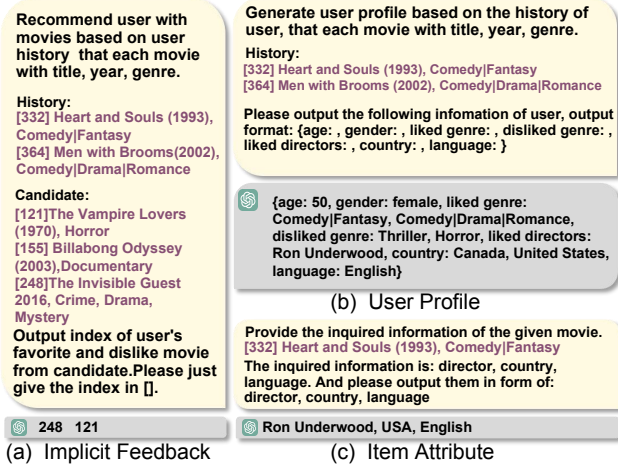


Figure 2: Constructed prompt $\mathbb{P}_u^{UI}, \mathbb{P}_u^U, \mathbb{P}_i^I$ for LLMs' completion including i) task description, ii) historical interactions, iii) candidates, and iv) output format description.

Formally, the LLM-based side information augmentation is as:

$$\begin{cases} \text{user} : \mathbb{A}_u = LLM(\mathbb{P}_u^U) \longrightarrow \mathbf{f}_{\mathcal{A},u} = LLM(\mathbb{A}_u) \\ \text{item} : \mathbb{A}_i = LLM(\mathbb{P}_i^I) \longrightarrow \mathbf{f}_{\mathcal{A},i} = LLM(\mathbb{A}_i) \end{cases} \quad (5)$$

where $\mathbf{f}_{\mathcal{A},u}, \mathbf{f}_{\mathcal{A},i} \in \mathbb{R}^{d_{LLM}}$ are LLM-augmented user/item features with LLM's hidden dimension d_{LLM} . The textual prompts \mathbb{P}_u^U and \mathbb{P}_i^I are used for attribute refinement for user u and item i , respectively. \mathbb{A}_u and \mathbb{A}_i represent generated textual attributes that to be encoded as features $\mathbf{F}_{\mathcal{A},u}, \mathbf{F}_{\mathcal{A},i}$ using the embedding capability of $LLM(\cdot)$.

3.2.2 Side Information Incorporation (Q3). After obtaining the augmented side information for user/item, an effective incorporation method is necessary. LLMRec includes a standard procedure: (1) Augmented Semantic Projection, (2) Collaborative Context Injection, and (3) Feature Incorporation. Let's delve into each:

- **Augmented Semantic Projection.** Linear layers with dropout are employed to not only reduce the dimensionality of LLM-enhanced semantic features but also map such augmented features into their own space [37]. This process can be represented as $\bar{\mathbf{F}}_{\mathcal{A}} = \text{Linear}(\mathbf{F}_{\mathcal{A}})$, where $\mathbf{f}_{\mathcal{A}} \in \mathbb{R}^{1 \times d_{LLM}}$ is the input feature and $\bar{\mathbf{f}}_{\mathcal{A}} \in \mathbb{R}^{1 \times d}$ is the output feature after projection.
- **Collaborative Context Injection.** To inject high-order [32] collaborative connectivity into augmented features $\bar{\mathbf{f}}_{\mathcal{A},u}$ and $\bar{\mathbf{f}}_{\mathcal{A},i}$, LLMRec employs light weight GNNs [12] as the encoder.
- **Semantic Feature Incorporation.** Instead of taking augmented features $\bar{\mathbf{F}}_{\mathcal{A}}$ as initialization of learnable vectors of recommender

$\bar{\mathbf{f}}_{\mathcal{A}}$, we opt to treat $\bar{\mathbf{F}}_{\mathcal{A}}$ as additional compositions added to the ID-corresponding embeddings $(\mathbf{e}_u, \mathbf{e}_i)$. This allows flexibly adjust the influence of LLM-augmented features using scale factors and normalization. Formally, the $\bar{\mathbf{F}}_{\mathcal{A}}$'s incorporation is presented as:

$$\mathbf{h}_u = \mathbf{e}_u + \omega_1 \cdot \sum_{k \in \mathcal{M} \cup \mathbb{A}_u} \frac{\bar{\mathbf{f}}_u^k}{\|\bar{\mathbf{f}}_u^k\|_2}; \quad \mathbf{h}_i = \mathbf{e}_i + \omega_1 \cdot \sum_{k \in \mathcal{M} \cup \mathbb{A}_i} \frac{\bar{\mathbf{f}}_i^k}{\|\bar{\mathbf{f}}_i^k\|_2}$$

The final prediction representations \mathbf{h}_u and \mathbf{h}_i , are in $\mathbb{R}^{1 \times d}$. User profiles are \mathbb{A}_u , debiased item attributes are \mathbb{A}_i , and original multimodal side information is \mathcal{M} . The specific type of feature is $\bar{\mathbf{f}}^k$. We adjust feature vectors using the aggregation weight ω_1 and L_2 normalization to mitigate distribution gaps [9], ensuring the effectiveness of additional features within the recommender encoder.

3.3 Training with Denoised Robustification (Q4)

In this section, we outline how LLMRec integrate augmented data into the optimization. We also introduce two quality constraint mechanisms for augmented edges and node features: i) Noisy user-item interaction pruning, and ii) MAE-based feature enhancement.

3.3.1 Augmented Optimization with Noise Pruning. We train our recommender using the union set $\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}$, which includes the original training set \mathcal{E} and the LLM-augmented set $\mathcal{E}_{\mathcal{A}}$. The objective is to optimize the BPR \mathcal{L}_{BPR} loss with increased supervisory signals $\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}$, aiming to enhance the recommender's performance by leveraging the incorporated LLM-enhanced user preference:

$$\mathcal{L}_{BPR} = \sum_{(u, i^+, i^-)}^{|\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}|} -\log(\sigma(\hat{y}_{u, i^+} - \hat{y}_{u, i^-})) + \omega_2 \cdot \|\Theta\|^2 \quad (6)$$

$$\mathcal{E}_{\mathcal{A}} \subseteq \{LLM(\mathbb{P}_u) | u \in \mathcal{U}\}, \quad |\mathcal{E}_{\mathcal{A}}| = \omega_3 * B$$

The training triplet (u, i^+, i^-) is selected from the union training set $\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}$. The predicted scores of positive-negative sample pairs are obtained through inner products of final representation \mathbf{h} , i.e., $\hat{y}_{u, i^+} = \mathbf{h}_u \cdot \mathbf{h}_{i^+}$, $\hat{y}_{u, i^-} = \mathbf{h}_u \cdot \mathbf{h}_{i^-}$. The augmented dataset $\mathcal{E}_{\mathcal{A}}$ is a subset of the overall LLM-generated data $\{LLM(\mathbb{P}_u) | u \in \mathcal{U}\}$, obtained by sampling. This is because excessive inclusion of pseudo label may lead to a degradation in result accuracy. The number of samples $|\mathcal{E}_{\mathcal{A}}|$ is controlled by the batch size B and a rate ω_3 . Weight-decay regularization $\|\Theta\|^2$ weighted by ω_2 , mitigates overfitting. $\sigma(\cdot)$ is activation function sigmoid to introduce non-linearity.

Noise Pruning. To enhance the effectiveness of augmented data, we prune out unreliable u-i interaction noise. Technically, the largest values before minus are discarded after sorting each iteration. This helps prioritize and emphasize relevant supervisory

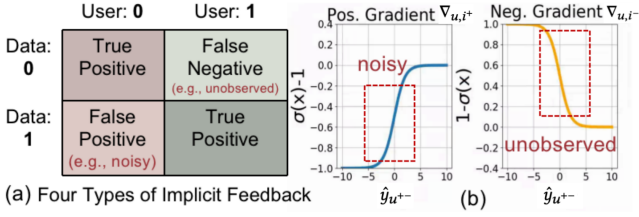


Figure 3: (a) Implicit feedback encompasses both false positive and false negative samples. (b) The gradient ∇ of the BPR loss for positive \hat{y}_{u,i^+} and negative \hat{y}_{u,i^-} scores, despite having a large magnitude, can have an incorrect direction that notably impacts the robustness and effectiveness of training.

signals while mitigating the influence of noise. Formally, the objective \mathcal{L}_{BPR} in Eq. 6 with noise pruning can be rewritten as follows:

$$(1-\omega_4)^{|\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}|} \sum_{(u,i^+,i^-)} -\text{SortAscend}(\log(\sigma(\hat{y}_{u,i^+} - \hat{y}_{u,i^-}))) [0 : N] + \omega_2 \cdot \|\Theta\|^2 \quad (7)$$

The function $\text{SortAscend}(\cdot) [0 : N]$ sorts values and selects the top- N . The retained number N is calculated by $N = (1 - \omega_4) \cdot |\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}|$, where ω_4 is a rate. This approach allows for controlled pruning of loss samples, emphasizing relevant signals while reducing noise. This can avoid the impact of unreliable gradient backpropagation, thus making optimization more stable and effective.

3.3.2 Enhancing Augmented Semantic Features via MAE.

To mitigate the impact of noisy augmented features, we employ the Masked Autoencoders (MAE) for feature enhancement [10]. Specifically, the masking technique is to reduce the model’s sensitivity to features, and subsequently, the feature encoders are strengthened through reconstruction objectives. Formally, we select a subset of nodes $\tilde{\mathcal{V}} \subset \mathcal{V}$ and mask their features using a mask token [MASK], denoted as $\mathbf{f}_{[MASK]}$ (e.g., a learnable vector or mean pooling). The mask operation can be formulated as follows:

$$\tilde{\mathbf{f}}_{\mathcal{A}} = \begin{cases} \mathbf{f}_{[MASK]} & v \in \tilde{\mathcal{V}} \\ \mathbf{f}_{\mathcal{A}} & v \notin \tilde{\mathcal{V}} \end{cases} \quad (8)$$

The augmented feature after the mask operation is denoted as $\tilde{\mathbf{f}}_{\mathcal{A}}$. It is substituted as mask token $\mathbf{f}_{[MASK]}$ if the node is selected ($\tilde{\mathcal{V}} \subset \mathcal{V}$), otherwise, it corresponds to the original augmented feature $\mathbf{f}_{\mathcal{A}}$. To strengthen the feature encoder, we introduce the feature restoration loss \mathcal{L}_{FR} by comparing the masked attribute matrix $\tilde{\mathbf{f}}_{\mathcal{A},i}$ with the original augmented feature matrix $\mathbf{f}_{\mathcal{A}}$, with a scaling factor γ . The restoration loss function \mathcal{L}_{FR} is as follows:

$$\mathcal{L}_{FR} = \frac{1}{|\tilde{\mathcal{V}}|} \sum_{v \in \tilde{\mathcal{V}}} \left(1 - \frac{\tilde{\mathbf{f}}_{\mathcal{A}} \cdot \mathbf{f}_{\mathcal{A}}}{\|\tilde{\mathbf{f}}_{\mathcal{A}}\| \cdot \|\mathbf{f}_{\mathcal{A}}\|}\right) \gamma \quad (9)$$

The final optimization objective is the weighted sum of the noise-pruned BPR loss \mathcal{L}_{BPR} and the feature restoration (FR) loss \mathcal{L}_{FR} .

3.4 In-Depth Analysis of our LLMRec

3.4.1 LLM-based Augmentation Facilitates Optimization.

This section highlights challenges addressed by LLM-based augmentation in recommender systems. False negatives (non-interacted interactions) and false positives (noise) as in Fig. 3 (a) can affect data

Table 1: Statistics of the Original and Augmented Datasets

Dataset		Netflix			MovieLens		
Graph	Ori.	# U	# I	# E	# U	# I	# E
	Aug.	# E:	26374		# E:	24990	
Ori. Sparsity		99.970%			99.915%		
Att.	Ori.	U: None	I: year, title		U: None	I: title, year, genre	
	Aug.	U[1536]:		age, gender, liked genre, disliked genre,			
		I[1536]:		liked directors, country, and language			
Modality	Textual[768], Visual [512]			Textual [768], Visual [512]			

* Att. represents attribute, Ori. represents original, and Aug. represents augmentation. Number in [] represents the feature dimensionality.

quality and result accuracy [3, 31]. Non-interacted items do not necessarily imply dislike [3], and interacted one may fail to reflect real user preferences due to accidental clicks or misleading titles, etc. Mixing of unreliable data with true user preference poses a challenge in build accurate recommender. Identifying and utilizing reliable examples is key to optimizing the recommender [2].

In theory, non-interacted and noisy interactions are used for negative \hat{y}_{u,i^-} and positive \hat{y}_{u,i^+} scores, respectively. However, their optimization directions oppose the true direction with large magnitudes, *i.e.*, the model optimizes significantly in the wrong directions (as in Fig. 3 (b)), resulting in sensitive suboptimal results.

Details. By computing the derivatives of the \mathcal{L}_{BPR} (Eq. 10), we obtain positive gradients $\nabla_{u,i^+} = 1 - \sigma(\hat{y}_{u,i^+})$ and negative gradients $\nabla_{u,i^-} = \sigma(\hat{y}_{u,i^-}) - 1$, where $\hat{y}_{u,i^+} = \hat{y}_{u,i^+} - \hat{y}_{u,i^-}$. Fig. 3 (b) illustrates these gradients and unveils some observations. Noisy interactions, although treated as positives, often have small values \hat{y}_{u,i^+} as false positives, resulting in large gradients ∇_{u,i^+} . Conversely, unobserved items, treated as negatives, tend to have relatively large values \hat{y}_{u,i^-} as false negatives, leading to small \hat{y}_{u,i^+} and large gradients ∇_{u,i^-} .

$$\begin{cases} \nabla_{u,i^+} = \frac{\partial \mathcal{L}_{BPR}}{\partial \hat{y}_{u,i^+}} = -\frac{\partial \log \sigma(\hat{y}_{u,i^+})}{\partial \sigma(\hat{y}_{u,i^+})} \frac{\partial \sigma(\hat{y}_{u,i^+})}{\partial \hat{y}_{u,i^+}} \\ = -\frac{1}{\sigma(\hat{y}_{u,i^+})} \cdot \sigma(\hat{y}_{u,i^+}) \cdot (1 - \sigma(\hat{y}_{u,i^+})) \cdot 1 = \sigma(\hat{y}_{u,i^+} - \hat{y}_{u,i^-}) - 1 \quad \text{noisy} \\ \nabla_{u,i^-} = \frac{\partial \mathcal{L}_{BPR}}{\partial \hat{y}_{u,i^-}} = -\frac{\partial \log \sigma(\hat{y}_{u,i^+})}{\partial \sigma(\hat{y}_{u,i^+})} \frac{\partial \sigma(\hat{y}_{u,i^+})}{\partial \hat{y}_{u,i^-}} \\ = \frac{1}{\sigma(\hat{y}_{u,i^+})} \cdot \sigma(\hat{y}_{u,i^+}) \cdot (1 - \sigma(\hat{y}_{u,i^+})) \cdot 1 = 1 - \sigma(\hat{y}_{u,i^+} - \hat{y}_{u,i^-}) \quad \text{unobserved} \end{cases} \quad (10)$$

Conclusion. Wrong samples possess incorrect directions but are influential. LLM-based augmentation uses the natural language space to assist the ID vector space to provide a comprehensive reflection of user preferences. With real-world knowledge, LLMRec gets quality samples, reducing the impact of noisy and unobserved implicit feedback, improving accuracy, and speeding up convergence.

3.4.2 Time Complexity. We analyze the time complexity. The projection of augmented semantic features has a time complexity of $\mathcal{O}(|\mathcal{U} \cup \mathcal{I}| \times d_{LLM} \times d)$. The GNN encoder for graph-based collaborative context learning takes $\mathcal{O}(L \times |\mathcal{E}^+| \times d)$ time. The BPR loss function computation has a time complexity of $\mathcal{O}(d \times |\mathcal{E} \cup \mathcal{E}_{\mathcal{A}}|)$, while the feature reconstruction loss has a time complexity of $\mathcal{O}(d \times |\tilde{\mathcal{V}}|)$, where $|\tilde{\mathcal{V}}|$ represents the count of masked nodes.

4 EVALUATION

To evaluate the performance of LLMRec, we conduct experiments, aiming to address the following research questions:

- **RQ1:** How does our LLM-enhanced recommender perform compared to the current state-of-the-art baselines?
- **RQ2:** What is the impact of key components on the performance?
- **RQ3:** How sensitive is the model to different parameters?
- **RQ3:** Are the data augmentation strategies in our LLMRec applicable across different recommendation models?
- **RQ5:** What is the computational cost associated with our devised LLM-based data augmentation schemes?

4.1 Experimental Settings

4.1.1 Datasets. We perform experiments on publicly available datasets, *i.e.*, Netflix and MovieLens, which include multi-modal side information. Tab. 1 presents statistical details for both the original and augmented datasets for both user and item domains. **MovieLens.** We utilize the MovieLens dataset derived from ML-10M¹. Side information includes movie title, year, and genre in textual format. Visual content consists of movie posters obtained through web crawling by ourselves. **Netflix.** We collected its multi-modal side information through web crawling. The implicit feedback and basic attribute are sourced from the Netflix Prize Data² on Kaggle. For both datasets, CLIP-ViT[25] is utilized to encode visual features. **LLM-based Data Augmentation.** The study employs the OpenAI package, accessed through LLMs’ APIs, for augmentation. The OpenAI Platform documentation provides details³. Augmented implicit feedback is generated using the "gpt-3.5-turbo-0613" chat completion model. Item attributes such as directors, country, and language are gathered using the same model. User profiling, based on the "gpt-3.5-turbo-16k" model, includes age, gender, preferred genre, disliked genre, preferred directors, country, and language. Embedding is performed using the "text-embedding-ada-002" model. The approximate cost of augmentation strategies on two datasets is 15.65 USD, 20.40 USD, and 3.12 USD, respectively.

4.1.2 Implementation Details. The experiments are conducted on a 24 GB Nvidia RTX 3090 GPU using PyTorch[23] for code implementation. The AdamW optimizer[22] is used for training, with different learning rate ranges of $[5e^{-5}, 1e^{-3}]$ and $[2.5e^{-4}, 9.5e^{-4}]$ for Netflix and MovieLens, respectively. Regarding the parameters of the LLMs, we choose the *temperature* from larger values $\{0.0, 0.6, 0.8, 1\}$ to control the randomness of the generated text. The value of *top-p* is selected from smaller values $\{0.0, 0.1, 0.4, 1\}$ to encourage probable choices. The *stream* is set to false to ensure the completeness of responses. For more details on the parameter analysis, please refer to Section 4.4. To maintain fairness, both our method and the baselines employ a unified embedding size of 64.

4.1.3 Evaluation Protocols. We evaluate our approach in the top-K item recommendation task using three common metrics: Recall (R@k), Normalized Discounted Cumulative Gain (N@k), and Precision (P@k). To avoid potential biases from test sampling, we employ the all-ranking strategy[38, 40]. We report averaged results from five independent runs, setting K to 10, 20, and 50 (reasonable

for all-ranking). Statistical significance analysis is conducted by calculating *p*-values against the best-performing baseline.

4.1.4 Baseline Description. Four distinct groups of baseline methods for thorough comparison. i) General CF Methods: **MF-BPR** [28], **NGCF** [32] and **LightGCN** [12]. ii) Methods with Side Information: **VBPR** [11], **MMGCN** [41] and **GRCN** [40]. iii) Data Augmentation Methods: **LATTICE** [50]. iv) Self-supervised Methods: **CLCRec** [39], **MMSSL** [36] and **MICRO** [49].

4.2 Performance Comparison (RQ1)

Tab. 2 compares our proposed LLMRec method with baselines.

- **Overall Model Superior Performance.** Our *LLMRec* outperforms the baselines by explicitly augmenting u-i interactive edges and enhancing the quality of side information. It is worth mentioning that our model based on LATTICE’s [50] encoder, consisting of a ID-corresponding encoder and a feature encoder. This improvement underscores the effectiveness of our framework.
- **Effectiveness of Side Information Incorporation.** The integration of side information significantly empowers recommenders. Methods like MMSSL [36] and MICRO [49] stand out for their effective utilization of multiple modalities of side information and GNNs. In contrast, approaches rely on limited content, such as VBPR [11] using only visual features, or CF-based architectures like NGCF [32], without side information, yield significantly diminished results. This highlights the importance of valuable content, as relying solely on ID-corresponding records fails to capture the complete u-i relationships.
- **Inaccurate Augmentation yields Limited Benefits.** Existing methods, such as LATTICE[50], MICRO[49] that also utilize side information for data augmentation have shown limited improvements compared to our LLMRec. This can be attributed to two main factors: (1) The augmentation of side information with homogeneous relationships (e.g., i-i or u-u) may introduce noise, which can compromise the precise of user preferences. (2) These methods often not direct augmentation of u-i interaction data.
- **Advantage over SSL Approaches.** Self-supervised models like, MMSSL[36], MICRO[49], have shown promising results in addressing sparsity through SSL signals. However, they do not surpass the performance of LLMRec, possibly because their augmented self-supervision signals may not align well with the target task of modeling u-i interactions. In contrast, we explicitly tackle the scarcity of training data by directly establishing BPR triplets.

4.3 Ablation and Effectiveness Analyses (RQ2)

We conduct an ablation study of our proposed LLMRec approach to validate its key components, and present the results in Table 3.

4.3.1 Effectiveness of Data Augmentation Strategies.

- (1). *w/o*-u-i: Disabling the LLM-augmented implicit feedback $\mathcal{E}_{\mathcal{A}}$ results in a significant decrease. This indicates that LLMRec increases the potential supervision signals by including contextual knowledge, leading to a better grasp of user preferences.
- (2). *w/o*-u: Removing our augmentor for user profiling result in a decrease in performance, indicating that our LLM-enhanced user side information can effectively summarize useful user preference profile using historical interactions and item-end knowledge.

¹<https://files.grouplens.org/datasets/movieLens/ml-10m-README.html>

²<https://www.kaggle.com/datasets/netflix-inc/netflix-prize-data>

³<https://platform.openai.com/docs/api-reference>

Table 2: Performance comparison on different datasets in terms of Recall@10/20/50, and NDCG@10/20/50, and Precision@20.

Baseline	Netflix							MovieLens						
	R@10	N@10	R@20	N@20	R@50	N@50	P@20	R@10	N@10	R@20	N@20	R@50	N@50	P@20
General Collaborative Filtering Methods														
MF-BPR	0.0282	0.0140	0.0542	0.0205	0.0932	0.0281	0.0027	0.1890	0.0815	0.2564	0.0985	0.3442	0.1161	0.0128
NGCF	0.0347	0.0161	0.0699	0.0235	0.1092	0.0336	0.0032	0.2084	0.0886	0.2926	0.1100	0.4262	0.1362	0.0146
LightGCN	0.0352	0.0160	0.0701	0.0238	0.1125	0.0339	0.0032	0.1994	0.0837	0.2660	0.1005	0.3692	0.1209	0.0133
Recommenders with Side Information														
VBPR	0.0325	0.0142	0.0553	0.0199	0.1024	0.0291	0.0028	0.2144	0.0929	0.2980	0.1142	0.4076	0.1361	0.0149
MMGCN	0.0363	0.0174	0.0699	0.0249	0.1164	0.0342	0.0033	0.2314	0.1097	0.2856	0.1233	0.4282	0.1514	0.0147
GRCN	0.0379	0.0192	0.0706	0.0257	0.1148	0.0358	0.0035	0.2384	0.1040	0.3130	0.1236	0.4532	0.1516	0.0150
Data Augmentation Methods														
LATTICE	0.0433	0.0181	0.0737	0.0259	0.1301	0.0370	0.0036	0.2116	0.0955	0.3454	0.1268	0.4667	0.1479	0.0167
MICRO	<u>0.0466</u>	0.0196	<u>0.0764</u>	0.0271	<u>0.1306</u>	0.0378	<u>0.0038</u>	0.2150	<u>0.1131</u>	<u>0.3461</u>	<u>0.1468</u>	<u>0.4898</u>	<u>0.1743</u>	<u>0.0175</u>
Self-supervised Methods														
CLCRec	0.0428	0.0217	0.0607	0.0262	0.0981	0.0335	0.0030	0.2266	0.0971	0.3164	0.1198	0.4488	0.1459	0.0158
MMSSL	0.0455	<u>0.0224</u>	0.0743	<u>0.0287</u>	0.1257	<u>0.0383</u>	0.0037	<u>0.2482</u>	0.1113	0.3354	0.1310	0.4814	0.1616	0.0170
LLMRec	0.0531	0.0272	0.0829	0.0347	0.1382	0.0456	0.0041	0.2603	0.1250	0.3643	0.1628	0.5281	0.1901	0.0186
<i>p</i> -value	$2.9e^{-4}$	$3.0e^{-3}$	$9.4e^{-5}$	$1.5e^{-3}$	$2.8e^{-5}$	$2.2e^{-3}$	$3.4e^{-5}$	$2.8e^{-5}$	$1.6e^{-2}$	$3.1e^{-3}$	$4.1e^{-4}$	$1.9e^{-3}$	$1.3e^{-2}$	$1.8e^{-3}$
Improv.	13.95%	21.43%	8.51%	20.91%	5.82%	19.06%	7.89%	4.88%	10.52%	5.26%	10.90%	7.82%	9.06%	6.29%

Table 3: Ablation study on key components (i.e., data augmentation strategies, denoised data robustification mechanisms)

	Metrics	R@10	N@10	R@20	N@20	R@50	N@50	P@20
Aug.	w/o-u-i	0.0477	0.0239	0.0791	0.0317	0.1376	0.0432	0.0037
	w/o-u	0.0423	0.0196	0.0656	0.0255	0.1192	0.0360	0.0033
	w/o-u&i	0.0309	0.0127	0.0602	0.0202	0.1051	0.0289	0.0030
Q.C.	w/o-prune	0.0504	0.0258	0.0786	0.0328	0.1363	0.0447	0.0039
	w/o-QC	0.0488	0.0244	0.0786	0.0318	0.1279	0.0416	0.0038
	LLMRec	0.0531	0.0272	0.0829	0.0347	0.1382	0.0456	0.0041

* "Aug": data augmentation operations; Q.C.: denoised data robustification.

Table 4: Parameter analysis of temperature τ and top-p ρ .

Para.	Temperature τ				Top-p ρ			
	$\tau=0$	$\tau=0.6$	$\tau=0.8$	$\tau=1$	$\rho=0$	$\rho=0.1$	$\rho=0.4$	$\rho=1$
R@10	0.0558 \uparrow	0.0531	0.0553 \uparrow	0.0531=	0.0537 \uparrow	0.0531	0.0520 \downarrow	0.0531=
R@20	0.0808 \downarrow	0.0829	0.0813 \downarrow	0.0775 \downarrow	0.0802 \downarrow	0.0829	0.0796 \downarrow	0.0770 \downarrow
R@50	0.1344 \downarrow	0.1382	0.1360 \downarrow	0.1312 \downarrow	0.1360 \downarrow	0.1382	0.1344 \downarrow	0.1333 \downarrow

- (3). *w/o-u&i*: when we remove the augmented side information for both users and items ($\mathcal{F}_{\mathcal{A},u}, \mathcal{F}_{\mathcal{A},i,1}$), lower recommendation accuracy is observed. This finding indicates that the LLM-based augmented side information provides valuable augmented data to the recommender system, assisting in obtaining quality and informative representations.

4.3.2 Impact of the Denoised Data Robustification.

- w/o-prune*: The removal of noise pruning results in worse performance. This suggests that the process of removing noisy implicit feedback signals helps prevent incorrect gradient descent.
- w/o-QC*: The performance suffer when both the limits on implicit feedback and semantic feature quality are simultaneously removed (i.e., *w/o-prune* + *w/o-MAE*). This indicates the benefits of our denoised data robustification mechanism by integrating noise pruning and semantic feature enhancement.

4.4 Hyperparameter Analysis (RQ3)

4.4.1 Parameters Affecting Augmented Data Quality.

Table 5: Analysis of key parameter (i.e., # candidate $|C|$) for LLM w.r.t implicit feedback augmentation $\mathcal{E}_{\mathcal{A}}$.

Data	Netflix			MovieLens		
	Metrics	$ C =3$	$ C =10$	$ C =30$	$ C =3$	$ C =10$
R@20	0.0786 \downarrow	0.0829	0.0808 \downarrow	0.3567 \downarrow	0.3643	0.3695 \uparrow
N@20	0.0314 \downarrow	0.0347	0.0330 \downarrow	0.1603 \downarrow	0.1628	0.1614 \downarrow
P@20	0.0039 \downarrow	0.0041	0.0040 \downarrow	0.0179 \downarrow	0.0186	0.0182 \downarrow

- Temperature τ of LLM*: The temperature parameter τ affects text randomness. Higher values (>1.0) increase diversity and creativity, while lower values (<0.1) result in more focus. We use τ from $\{0, 0.6, 0.8, 1\}$. As shown in Table 4, increasing τ initially improves most metrics, followed by a decrease.
- Top-p p of LLM*: Top-p Sampling[13] selects tokens based on a threshold determined by the top-p parameter p . Lower p values prioritize likely tokens, while higher values encourage diversity. We use p from $\{0, 0.1, 0.4, 1\}$ and smaller p values tend to yield better results, likely due to avoiding unlisted candidate selection. Higher p values cause wasted tokens due to repeated LLM inference.
- # of Candidate C* : We use C to limit item candidates for LLM-based recommendation. $\{3, 10, 30\}$ are explored due to cost limitations, and Table 5 shows that $C = 10$ yields the best results. Small values limit selection, and large values increase recommendation difficulty.
- Prune Rate ω_4* : LLMRec uses ω_4 to control noise in augmented training data to be pruned. We set ω_4 to $\{0.0, 0.2, 0.4, 0.6, 0.8\}$ on both datasets. As shown in Fig. 4 (a), $\omega_4 = 0$ yields the worst result, highlighting the need to constrain noise in implicit feedback.

4.4.2 Sensitivity of Recommenders to the Augmented Data.

- # of Augmented Samples per Batch $|\mathcal{E}_{\mathcal{A}}|$* : LLMRec uses ω_3 and batch size B to control the number of augmented BPR training data samples per batch. ω_3 is set to $\{0.0, 0.1, 0.2, 0.3, 0.4\}$ on Netflix and $\{0.0, 0.2, 0.4, 0.6, 0.8\}$ on MovieLens. Suboptimal results occur when ω_3 is zero or excessively large. Increasing diversity and randomness can lead to a more robust gradient descent.
- Scale ω_2 for Incorporating Augmented Features*: LLMRec uses ω_2 to control feature magnitude, with values set to $\{0.0, 0.8, 1.6, 2.4,$

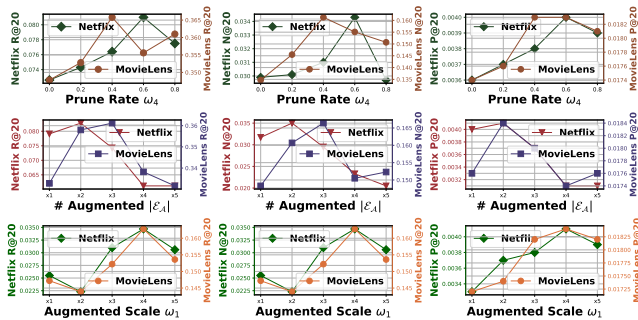


Figure 4: Impact of hyperparameters (i.e., prune rate ω_4 , # augmented BPR training data $|\mathcal{E}_A|$, and augmented feature incorporate scale ω_1).

Table 6: Model-agnostic experiment to evaluate the effectiveness of LLM-based data augmentation on different recommender in terms of R@20, N@20, and P@20.

	Method	LATTICE	MICRO	MMSSL
Aug.	R@20	0.0821 \uparrow 11.40%	0.0835 \uparrow 9.29%	0.0833 \uparrow 11.11%
	N@20	0.0287 \uparrow 10.81%	0.0301 \uparrow 11.07%	0.0313 \uparrow 9.06%
	P@20	0.0039 \uparrow 8.33%	0.0041 \uparrow 7.89%	0.0041 \uparrow 10.81%

3.2} on Netflix and {0.0, 0.1, 0.2, 0.3, 0.4} on MovieLens. Optimal results depend on the data, with suboptimal outcomes occurring when ω_2 is too small or too large, as shown in Fig. 4 (c).

4.5 Model-agnostic Property (RQ4)

We conducted model-agnostic experiments on Netflix to validate the applicability of our data augmentation. Specifically, we incorporated the augmented implicit feedback \mathcal{E}_A and features $F_{A,u}$, $F_{A,i}$ into baselines MICRO, MMSSL, and LATTICE. As shown in Tab. 6, our LLM-based data improved the performance of all models, demonstrating their effectiveness and reusability. Some results didn't surpass our model, maybe due to: i) the lack of a quality constraint mechanism to regulate the stability and quality of the augmented data, and ii) the absence of modeling collaborative signals in the same vector space, as mentioned in Sec. 3.2.2.

4.6 Cost/Improvement Conversion Rate (RQ5)

To evaluate the cost-effectiveness of our augmentation strategies, we compute the CIR as presented in Tab. 7. The CIR is compared with the ablation of three data augmentation strategies and the best baseline from Tab. 3 and Tab. 2. The cost of the implicit feedback augmentor refers to the price of GPT-3.5 turbo 4K. The cost of side information augmentation includes completion (using GPT-3.5 turbo 4K or 16K) and embedding (using text-embedding-ada-002). We utilize the HuggingFace API tool for tokenizer and counting. The results in Tab. 7 show that 'U' (LLM-based user profiling) is the most cost-effective strategy, and the overall investment is worthwhile.

5 RELATED WORK

Content-based Recommendation. Existing recommenders have explored the use of auxiliary multi-modal side knowledge, with methods like VBPR [11] combine traditional collaborative filtering with visual features, while MMGCN [41], GRCN [40] leverage GNNs to capture modality-aware higher-order collaborative signals.

Table 7: Comparison of the cost and improvement rate(CIR) of data augmentation strategies and LLMRec. 'Cost': expenditure of utilizing LLM, 'Imp': the average improvement rate in R@10/N@10. 'CIR': the ratio of improvement to cost.

	Cost(USD)	R@10		N@10	
		Imp.(%)	CIR(%)	Imp.(%)	CIR(%)
U	10.92	25.53	233.79	38.78	355.13
I	1.96	2.31	117.86	1.12	57.14
U-I	8.26	11.32	137.05	13.81	167.19
LLMAug	21.14	13.95	65.99	21.43	101.37

Recent approaches like MMSSL [36] and MICRO [49] align modal signals with collaborative signals through contrastive SSL, revealing the informative aspects of modal signals that benefit recommendations. However, the data noise, heterogeneity, and incompleteness may can lead to bias. To overcome this, LLMRec explores LLM-based augmentation to improve the quality of the data.

Large Language Models (LLMs) for Recommendation. LLMs have garnered attention in recommendation systems, with multiple attempts to utilize them for modeling user behavior [15, 33]. LLMs have been employed as an inference model in diverse recommendation tasks, including rating prediction, sequential recommendation, and direct recommendation [1, 6, 48]. LLMs have also been used to simulate point-wise, pair-wise, and list-wise scenarios in ranking tasks [7]. For example, RLMRec[26] enhance the recommendation performance by leveraging LLMs and aligning their semantic space with collaborative relation modeling through contrastive or generative modeling techniques. Besides, efforts such as GraphGPT [29] also tried to utilize LLMs to model structure relations. However, existing LLM-based recommendation methods lack comparative analysis against state-of-the-art recommenders. We enhance our model with LLM-based data augmentation and collaborative relation learning, validating its effectiveness across various advanced recommendation systems with side information.

Data Augmentation for Recommendation. Extensive research has explored data augmentation in recommendation systems [14, 17]. Various operations, such as permutation, deletion, swap, insertion, and duplication, have been proposed for sequential recommendation [21, 24]. Commonly used techniques include counterfactual reasoning and contrastive learning [20, 34, 42, 51]. Our LLMRec use LLMs as an inference model to augment edge and enhance node features by leveraging consensus knowledge from the large model.

6 CONCLUSION

This study focuses on the design of LLM-enhanced models to address the challenges of sparse implicit feedback signals and low-quality side information by profiling user interaction preferences and debiasing item attributes. To ensure the quality of augmented data, a denoised augmentation robustification mechanism is introduced. The effectiveness of LLMRec is supported by theoretical analysis and experimental results, demonstrating its superiority over state-of-the-art recommendation techniques on benchmark datasets. Future directions for investigation include integrating causal inference into side information debiasing and exploring counterfactual factors for context-aware user preference.

REFERENCES

- [1] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. TALLRec: An Effective and Efficient Tuning Framework to Align Large Language Model with Recommendation. *arXiv preprint arXiv:2305.00447* (2023).
- [2] Chong Chen, Weizhi Ma, Min Zhang, et al. 2023. Revisiting negative sampling vs. non-sampling in implicit recommendation. *TOIS* 41, 1 (2023), 1–25.
- [3] Chong Chen, Min Zhang, Yongfeng Zhang, et al. 2020. Efficient neural matrix factorization without sampling for recommendation. *TOIS* 38, 2 (2020), 1–28.
- [4] Hanxiong Chen, Yunqi Li, Shaoyun Shi, Shuchang Liu, He Zhu, and Yongfeng Zhang. 2022. Graph collaborative reasoning. In *WSDM*. 75–84.
- [5] Mengru Chen, Chao Huang, Lianghao Xia, Wei Wei, et al. 2023. Heterogeneous graph contrastive learning for recommendation. In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*. 544–552.
- [6] Zheng Chen. 2023. PALR: Personalization Aware LLMs for Recommendation. *arXiv preprint arXiv:2305.07622* (2023).
- [7] Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. 2023. Uncovering ChatGPT’s Capabilities in Recommender Systems. *arXiv preprint arXiv:2305.02182* (2023).
- [8] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *WWW*. 417–426.
- [9] Xinyu Fu, Jiani Zhang, et al. 2020. Maggn: Metapath aggregated graph neural network for heterogeneous graph embedding. In *WWW*. 2331–2341.
- [10] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *CVPR*. 16000–16009.
- [11] Ruining He and Julian McAuley. 2016. VBPR: visual bayesian personalized ranking from implicit feedback. In *AAAI*, Vol. 30.
- [12] 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.
- [13] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751* (2019).
- [14] Tinglin Huang, Yuxiao Dong, Ming Ding, Zhen Yang, Wenzheng Feng, Xinyu Wang, and Jie Tang. 2021. MixGCF: An Improved Training Method for Graph Neural Network-based Recommender Systems. In *KDD*.
- [15] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, et al. 2023. Do LLMs Understand User Preferences? Evaluating LLMs On User Rating Prediction. *arXiv preprint arXiv:2305.06474* (2023).
- [16] Hyeyoung Ko, Suyeon Lee, Yoonseo Park, and Anna Choi. 2022. A survey of recommendation systems: recommendation models, techniques, and application fields. *Electronics* 11, 1 (2022), 141.
- [17] Dongha Lee, SeongKu Kang, Hyunjun Ju, et al. 2021. Bootstrapping user and item representations for one-class collaborative filtering. In *SIGIR*. 317–326.
- [18] Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. 2023. Text Is All You Need: Learning Language Representations for Sequential Recommendation. In *KDD*.
- [19] Jinming Li, Wentao Zhang, Tian Wang, Guanglei Xiong, et al. 2023. GPT4Rec: A Generative Framework for Personalized Recommendation and User Interests Interpretation. *arXiv preprint arXiv:2304.03879* (2023).
- [20] Zhiwei Liu, Yongjun Chen, Jia Li, Philip S Yu, Julian McAuley, and Caiming Xiong. 2021. Contrastive self-supervised sequential recommendation with robust augmentation. *arXiv preprint arXiv:2108.06479* (2021).
- [21] Zhiwei Liu, Ziwei Fan, et al. 2021. Augmenting sequential recommendation with pseudo-prior items via reversely pre-training transformer. In *SIGIR*. 1608–1612.
- [22] Ilya Loshchilov et al. 2017. Decoupled weight decay regularization. In *ICLR*.
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS* 32 (2019).
- [24] Aleksandr Petrov and Craig Macdonald. 2022. Effective and Efficient Training for Sequential Recommendation using Recency Sampling. In *Recsys*. 81–91.
- [25] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*. PMLR, 8748–8763.
- [26] Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2023. Representation Learning with Large Language Models for Recommendation. *arXiv preprint arXiv:2310.15950* (2023).
- [27] Xubin Ren, Lianghao Xia, Yuhao Yang, Wei Wei, Tianle Wang, Xuheng Cai, and Chao Huang. 2023. SSLRec: A Self-Supervised Learning Library for Recommendation. *arXiv preprint arXiv:2308.05697* (2023).
- [28] Steffen Rendle, Christoph Freudenthaler, et al. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [29] Jiabin Tang, Yuhao Yang, Wei Wei, Lei Shi, Lixin Su, Suqi Cheng, Dawei Yin, and Chao Huang. 2023. GraphGPT: Graph Instruction Tuning for Large Language Models. *arXiv preprint arXiv:2310.13023* (2023).
- [30] Changxin Tian, Yuexiang Xie, Yaliang Li, et al. 2022. Learning to denoise unreliable interactions for graph collaborative filtering. In *SIGIR*. 122–132.
- [31] Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. 2021. Denoising implicit feedback for recommendation. In *WSDM*. 373–381.
- [32] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*. 165–174.
- [33] Xiaolei Wang, Xinyu Tang, Wayne Xin Zhao, Jingyuan Wang, and Ji-Rong Wen. 2023. Rethinking the Evaluation for Conversational Recommendation in the Era of Large Language Models. *arXiv preprint arXiv:2305.13112* (2023).
- [34] Zhenlei Wang, Jingsen Zhang, Hongteng Xu, Xu Chen, Yongfeng Zhang, Wayne Xin Zhao, and Ji-Rong Wen. 2021. Counterfactual data-augmented sequential recommendation. In *SIGIR*. 347–356.
- [35] Wei Wei, Chao Huang, Lianghao Xia, Yong Xu, Jiashu Zhao, and Dawei Yin. 2022. Contrastive meta learning with behavior multiplicity for recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*. 1120–1128.
- [36] Wei Wei, Chao Huang, Lianghao Xia, and Chuxu Zhang. 2023. Multi-Modal Self-Supervised Learning for Recommendation. In *WWW*. 790–800.
- [37] Wei Wei, Lianghao Xia, and Chao Huang. 2023. Multi-Relational Contrastive Learning for Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 338–349.
- [38] Yinwei Wei, Xiang Wang, et al. 2021. Hierarchical user intent graph network for multimedia recommendation. *Transactions on Multimedia (TMM)* (2021).
- [39] Yinwei Wei, Xiang Wang, Qi Li, Liqiang Nie, Yan Li, et al. 2021. Contrastive learning for cold-start recommendation. In *ACM MM*. 5382–5390.
- [40] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, and Tat-Seng Chua. 2020. Graph-refined convolutional network for multimedia recommendation with implicit feedback. In *MM*. 3541–3549.
- [41] Yinwei Wei, Xiang Wang, Liqiang Nie, Xiangnan He, Richang Hong, and Tat-Seng Chua. 2019. MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video. In *MM*. 1437–1445.
- [42] Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, et al. 2021. Self-supervised graph learning for recommendation. In *SIGIR*. 726–735.
- [43] Zixuan Yi, Xi Wang, Iadh Ounis, and Craig Macdonald. 2022. Multi-modal Graph Contrastive Learning for Micro-video Recommendation. In *SIGIR*. 1807–1811.
- [44] Yuxin Ying, Fuzhen Zhuang, Yongchun Zhu, Deqing Wang, and Hongwei Zheng. 2023. CAMUS: Attribute-Aware Counterfactual Augmentation for Minority Users in Recommendation. In *WWW*. 1396–1404.
- [45] Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. 2022. Are graph augmentations necessary? Simple graph contrastive learning for recommendation. In *SIGIR*. 1294–1303.
- [46] Zheng Yuan, Fajie Yuan, Yu Song, Youhua Li, Junchen Fu, Fei Yang, Yunzhu Pan, and Yongxin Ni. 2023. Where to go next for recommender systems? id-vs. modality-based recommender models revisited. In *SIGIR*.
- [47] Honglei Zhang, Fangyuan Luo, Jun Wu, Xiangnan He, and Yidong Li. 2023. LightFR: Lightweight federated recommendation with privacy-preserving matrix factorization. *ACM Transactions on Information Systems* 41, 4 (2023), 1–28.
- [48] Junjie Zhang, Ruobing Xie, Yupeng Hou, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2023. Recommendation as instruction following: A large language model empowered recommendation approach. *arXiv preprint arXiv:2305.07001* (2023).
- [49] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, et al. 2022. Latent structure mining with contrastive modality fusion for multimedia recommendation. *TKDE* (2022).
- [50] Jinghao Zhang, Yanqiao Zhu, Qiang Liu, Shu Wu, et al. 2021. Mining Latent Structures for Multimedia Recommendation. In *MM*. 3872–3880.
- [51] Shengyu Zhang, Dong Yao, Zhou Zhao, et al. 2021. Causerec: Counterfactual user sequence synthesis for sequential recommendation. In *SIGIR*. 367–377.
- [52] Ding Zou, Wei Wei, Xian-Ling Mao, Ziyang Wang, Minghui Qiu, Feida Zhu, and Xin Cao. 2022. Multi-level cross-view contrastive learning for knowledge-aware recommender system. In *SIGIR*. 1358–1368.