

CESRec: Constructing Pseudo Interactions for Sequential Recommendation via Conversational Feedback

Yifan Wang¹, Shen Gao^{1,*}, Jiabao Fang², Rui Yan³,
Billy Chiu⁴, Shuo Shang^{1,*}

¹School of Computer Science, University of Electronic Science and Technology of China

²School of Computer Science and Technology, Shandong University

³School of Artificial Intelligence, Wuhan University

⁴Department of Computing and Decision Sciences, Lingnan University

Abstract

Sequential Recommendation Systems (SRS) have become essential in many real-world applications. However, existing SRS methods often rely on collaborative filtering signals and fail to capture real-time user preferences, while Conversational Recommendation Systems (CRS) excel at eliciting immediate interests through natural language interactions but neglect historical behavior. To bridge this gap, we propose CESRec, a novel framework that integrates the long-term preference modeling of SRS with the real-time preference elicitation of CRS. We introduce semantic-based pseudo interaction construction, which dynamically updates users' historical interaction sequences by analyzing conversational feedback, generating a pseudo-interaction sequence that seamlessly combines long-term and real-time preferences. Additionally, we reduce the impact of outliers in historical items that deviate from users' core preferences by proposing dual alignment outlier items masking, which identifies and masks such items using semantic-collaborative aligned representations. Extensive experiments demonstrate that CESRec achieves state-of-the-art performance by boosting strong SRS models, validating its effectiveness in integrating conversational feedback into SRS¹.

1 Introduction

Sequential Recommendation Systems (SRS) are pivotal in various applications, such as e-commerce (Zhou et al., 2018) and streaming platforms (Pan et al., 2023), by providing personalized item recommendations based on users' historical interaction sequences (Fang et al., 2020). Recently, large language models (LLMs) have demonstrated remarkable reasoning capabilities (Mann et al., 2020; Zhang et al., 2022; Gao et al., 2024), making




Figure 1: Comparison of sequential recommendation, conversational recommendation, and our CESRec, which leverage advantage of conversational recommendation to enhance sequential recommendation.

them promising method for enhancing recommendation tasks. Several studies (Liao et al., 2024; Bao et al., 2023; Gao et al., 2019) have demonstrated the superiority of directly applying LLMs to sequential recommendation tasks. In contrast, Conversational Recommendation Systems (CRS) employ natural language interactions to inquire about user preferences and predict personalized item recommendations (Friedman et al., 2023; Mysore et al., 2023; Shang et al., 2014). However, existing SRS methods usually rely on collaborative filtering signals while neglecting the rich semantic information associated with items. A significant limitation of these approaches is their inability to capture users' real-time interests, as immediate preferences are not dynamically reflected in the behavior sequence. Conversely, while CRS methods excel at capturing immediate interests through natural language conversations, they typically fail to incorporate historical interaction sequences into their frameworks (Zhou et al., 2020; Lei et al., 2020; He et al., 2023; Feng et al., 2023). Consequently, the first challenge lies in dynamically integrating

*Corresponding authors

¹Code is available at <https://github.com/NNNNyifan/CESRec>

the *long-term preference modeling* of SRS with the *real-time interests modeling* facilitated by natural language feedback in CRS.

In this paper, we propose **Conversation Enhanced Sequential Recommendation (CESRec)**. To address the first challenge, we introduce **semantic-based pseudo interaction construction**, a novel method that directly refines the historical interaction sequence based on users' natural language feedback. Specifically, this approach analyzes users' natural language feedback to model their current preferences and refines their historical interaction sequence, generating a *pseudo-interaction sequence* that seamlessly integrates both long-term and real-time preferences. Next, we use the pseudo-interaction sequence as input to SRS, which effectively combines the collaborative filtering signals of SRS with the semantic signals derived from natural language feedback. This enables accurate recommendations based on natural language interactions without requiring extensive modifications to existing SRS-based systems, ensuring seamless integration and enhanced user experience.

Since historical interaction sequences often contain items that deviate substantially from users' main preferences, such as mistakenly clicked items or transient interests, as observed in many recent studies (Lin et al., 2023; Wang et al., 2021), these outliers can adversely affect the modeling of user behavior. The inclusion of these items can negatively influence the LLM's modeling of user behavior, potentially misleading the construction of the pseudo-interaction sequence. For example, if a user's primary preference is horror films, the inclusion of a comedy movie in the interaction sequence might lead the LLMs to utilize "horror-comedy" films to construct the pseudo-interaction sequence, rather than a pure horror film. In this work, we refer to such items as **outlier items**. Therefore, the second challenge is how to accurately identify these outlier items and mask them in the interaction sequence to minimize their impact on the generation of the pseudo-interaction sequence.

To address this, we propose dual alignment outlier items masking, a method that accurately identifies outlier items from the user's historical interaction sequence based on semantic-collaborative aligned representations and subsequently masks these items. Specifically, we leverage LLMs to obtain semantic embeddings of items and extract collaborative representations from the SRS model. We then introduce a dual alignment mechanism to

derive hybrid item representations, which simultaneously capture co-occurrence relationships and semantic information among items. Based on these hybrid representations, we identify items that substantially deviate from the user's core preferences, ensuring precise masking while preserving the integrity of the user's historical behavior sequence. The experimental results demonstrate that our CESRec can boost the performance of several state-of-the-art SRS models in terms of HR and NDCG, which verifies that our CESRec can effectively integrate long-term preferences with real-time interests through natural language feedback. The main contributions of this work are as follows:

- We propose CESRec, which combines the advantage of real-time natural language feedback with the efficiency of learning user preferences from historical behavior.
- We introduce semantic-based pseudo interaction construction method to refine user historical interaction sequences by leveraging user natural language feedback.
- We propose dual alignment outlier items masking method to optimize item selection during the sequence refinement process.
- Extensive experiments demonstrate that our proposed CESRec achieves state-of-the-art performance by boosting the performance of several strong SRS models.

2 Related Work

Sequential Recommendation Sequential recommendation aims to predict the next item that aligns with a user's preferences based on their historical interaction sequence (Fang et al., 2020; Li et al., 2023a,c; Shang et al., 2015; Chen and Shang, 2019; Chen et al., 2013). Traditional sequential recommendation models capture user preferences by leveraging item co-occurrence relationships. To model complex sequential patterns, CNN-based (Tang and Wang, 2018) and GNN-based (He et al., 2020; Han et al., 2019) methods have been introduced. Additionally, transformer-based approaches, such as SASRec (Kang and McAuley, 2018) and BERT4Rec (Sun et al., 2019), have been developed to capture long-term dependencies between arbitrary items. However, most of these methods primarily model user preferences based on long-term interaction histories, making it challenging to effectively capture dynamic shifts in user interests. As a result, they struggle to reflect

users’ real-time preferences within interaction sequences, leading to recommendations that may not accurately align with users’ immediate interests.

LLMs for Recommendation Large Language Models (LLMs) have demonstrated remarkable capabilities across various domains. Recent research explores hybrid approaches to enhance sequential recommendation by integrating traditional Recsys with LLMs (Dai et al., 2023; Geng et al., 2022; Hou et al., 2024; Bao et al., 2023). (Liao et al., 2024) combines ID-based embeddings and textual features via hybrid prompting, while (Rajput et al., 2023) introduces generative retrieval using semantic ID decoding. (Liu et al., 2024) leverages LLMs to generate item embeddings, and (Hu et al., 2024) aligns ID embeddings with text via a projector module. However, they do not fully exploit the rich semantic information contained in users’ conversational feedback, limiting their ability to dynamically adapt recommendation strategies based on real-time user preferences.

3 Problem Definition

In this paper, we follow the problem definition commonly used in sequential recommendation tasks (Hu et al., 2024). Given a user $u \in \mathcal{U}$, where $\mathcal{U}$ represents the set of all users, and a historical interaction sequence $\mathcal{I}(u) = \{v_1^{(u)}, v_2^{(u)}, \dots, v_{N_u}^{(u)}\}$, the model aims to predict the next item the user is likely to interact with based on $\mathcal{I}(u)$. Here, $v_i^{(u)}$ denotes the i -th item interacted by user u , and all items belong to the item set $\mathcal{V}$. The sequence length of $\mathcal{I}(u)$ is denoted by N_u .

4 CESRec

4.1 Overview

In this section, we show the details of the Conversation Enhanced Sequential Recommendation (CESRec), which is illustrated in Figure 2. The proposed model consists of two main components: Semantic Pseudo Sequence Construction and Dual Alignment Outlier Items Masking. The **Semantic Pseudo Sequence Construction** module is designed to construct the pseudo-interaction sequence by refining the historical interaction sequence via users’ conversational feedback. Subsequently, the **Dual Alignment Outlier Items Masking** module further enhances the refinement process by

identifying and masking items that deviate from the user’s core preferences.

4.2 Dual Alignment Outlier Items Masking

In the process of constructing a semantic-based pseudo-interaction sequence, the model leverages the users historical sequences to capture their core preferences and selects appropriate replacement items based on natural language feedback. However, during the modification of the original interaction sequence, items in the historical sequence that deviate from the user’s core preferences can interfere with the LLM’s modeling of user behavior. This misalignment can introduce bias, potentially leading to the inappropriate replacement of items. In this work, we refer to such items as outlier items. To address this issue, we propose a dual-alignment outlier items masking method to ensure that such deviating items are appropriately masked.

According to a recent study (Sheng et al., 2024), LLMs can implicitly encode user preference information, and items sharing similar content tend to exhibit similar semantic embeddings. Based on this observation, we extract item embeddings from LLMs, which are rich in semantic information. Given an item $v_i^{(u)}$ with content information c_i such as title, we employ an LLM to obtain the semantic embeddings e_i^{LLM} :

$$e_i^{LLM} = \text{Extractor}(c_i), \quad (1)$$

where $\text{Extractor}(\cdot)$ refers to the LLM tokenizer and model layers, and we utilize the output of the last hidden layer e_i^{LLM} as the semantic embedding.

Relying solely on semantic embeddings to identify outlier items may compromise the integrity of the user’s historical behavior sequence, thereby limiting the effectiveness of SRS in accurately modeling user preferences. We introduce a trainable adapter to align the semantic embeddings derived from LLMs with the collaborative signals typically used in SRS. This adapter is specifically trained to fuse the positional influence and co-occurrence information while utilizing semantic embeddings for masking:

$$e_i^{\text{hybrid}} = \text{Adapter}(\theta_{\text{collab}}; e_i^{LLM}), \quad (2)$$

where $e_i^{\text{hybrid}} \in \mathbb{R}^{d_r}$ denote the hybrid embedding that combines both semantic and collaborative information through an adaptation module. The adaptation is performed by $\text{Adapter}_\theta : \mathbb{R}^{d_l} \rightarrow \mathbb{R}^{d_r}$ with




Figure 2: Overview of CESRec. In our proposed framework, we first employ the conventional sequential recommendation method (*a.k.a.*, Recsys) to predict an item based on the user’s historical interaction sequence. Next, our CESRec refines the interaction sequence by constructing the pseudo-interaction sequence and masking the outlier items. Finally, we employ Recsys to give a new recommendation by using the refined sequence.

trainable parameters θ_{collab} , which projects embeddings from the LLM’s latent space (d_l -dimensional) to the Recsys’s embedding space (d_r -dimensional). The detailed architecture and training procedure of the adaptor are described in Appendix 9.1

Finally, to identify outlier items in interactions, we rank items based on the similarity between user representation and each item. We first obtain all the hybrid embeddings of all the user interacted items in $\mathcal{I}(u)$, and fuse all the item representation as the user embedding u^{hybrid} :

$$u^{\text{hybrid}} = \text{Fuse}(\{e_1^{\text{hybrid}}, e_2^{\text{hybrid}}, \dots, e_{N_u}^{\text{hybrid}}\}), \quad (3)$$

where the $\text{Fuse}(\cdot)$ denotes the mean-pooling operator. Then, we calculate the similarity between each item representation e_i^{hybrid} and user representation u^{hybrid} :

$$s_i = \text{Similarity}(e_i^{\text{hybrid}}, u^{\text{hybrid}}), \quad (4)$$

where $s_i \in [0, 1]$ denotes the similarity score, and we employ the cosine similarity as the $\text{Similarity}(\cdot)$ function to measure the semantic gap between e_i^{hybrid} and u^{hybrid} . To identify outlier items in interaction sequence, we rank items based on their similarity scores s_i . The top k items with the lowest similarity scores are considered as the outlier

items and will be subsequently masked from the user interaction sequence.

The input and output format of the final dual alignment outlier items masking is as follows:

$$I(u)' = \text{Dual-Alignment}(I(u)), \quad (5)$$

where $I(u)' = \{v_1^{(u)}, \dots, v_{N_u-k}^{(u)}, \hat{v}_1^{(u)}, \dots, \hat{v}_k^{(u)}\}$ represents interaction sequence after masking, $\hat{v}_i^{(u)}$ represents the top k items with the lowest similarity scores. Using these hybrid representations, we identify and mask the outlier items that deviate from the user’s core preferences while preserving the integrity of their historical behavior sequence. This optimization enables the CESRec to better concentrate on core preferences when constructing a semantic-based pseudo interaction sequence.

4.3 Semantic Pseudo Sequence Construction

To address the challenge of dynamically integrating long-term preference modeling of SRS with real-time interest modeling driven by natural language interactions in CRS, we propose a semantic-based pseudo sequence construction approach. This method leverages natural language feedback from users to directly capture their current preferences, and generates semantic-based

pseudo sequence by incorporating current preferences to historical interaction sequence. Specifically, we introduce a *constructor* that constructs semantic-based pseudo interaction sequences based on user-provided feedback. Following the previous works (Fang et al., 2024), we ask the user for preference about the target item attributes.

$$\text{feedback} = \text{User-Interaction}(v_{rec}^{(u)}, \text{Attr}_{\text{target}}) \quad (6)$$

where $v_{rec}^{(u)}$ represents the recommended item generated by an SRS with input $I(u)$, $\text{Attr}_{\text{target}}$ refers to attributes of the target item, and feedback denotes a natural language feedback derived from the user that describes the user preference of the item attributes. For instance, if the SRS recommends <Avatar> to the user, but the user prefers films directed by Christopher Nolan, the user may respond with feedback such as: “I don’t like film directed by James Cameron; I prefer Christopher Nolan.”

Next, the Constructor integrates user feedback to iterative refine the historical interaction sequence $I'(u)$ and generate the pseudo-interaction sequence $I_{\text{pseudo}}(u)$:

$$I_{\text{pseudo}}(u) = \text{Constructor}(I'(u), \text{feedback}), \quad (7)$$

where $I_{\text{pseudo}}(u)$ represents the pseudo-interaction sequence generated by the Constructor.

Training of Constructor To achieve a synergistic integration of user long-term interests and real-time preferences, we fine-tune LLM as the Constructor. The training objective for the Constructor is formulated as a sequence prediction task:

$$\mathcal{L}_{\text{seq}} = - \sum_{t=1}^{|I_{\text{pseudo}}^*(u)|} \log P_{\Psi}(v_t^* | v_{<t}^*, I(u), \text{feedback}) \quad (8)$$

I_{pseudo}^* denotes the pseudo-interaction sequence, which optimally combines the user’s long-term interests and real-time preference reflected in user feedback.

To construct the training data for the constructor module, we generate the pseudo-interaction sequence by replacing items that no longer align with the user’s current preferences. During the training data construction process, we randomly sample items from the interaction sequence as “Outlier Items”. The target item, which reflects the user’s current preference, serves as the ground truth for model training. The feedback derived from the transition between the Outlier Items and the target

item is utilized as the input feature for the model. The training instruction is as follows:

Instruction of Constructor

Instruction: Based on the preferences mentioned in the user feedback and the information about <items> contained in the historical interaction sequence, replace the <items> the user dislikes with <items> user may currently prefer.
Input: historical interaction sequence: <sequence>; user feedback: <feedback>.
Output: pseudo-interaction sequence: <pseudo-sequence>

Finally, after refining the interaction sequence of the user by the Constructor, we use the semantic pseudo interaction sequence $I_{\text{pseudo}}(u)$ as the input to the SRS to regenerate recommended items.

$$v_{N_u+1}^{(u)} = \text{SRS}(I_{\text{pseudo}}(u)), \quad (9)$$

where SRS represents sequential recommendation models, $v_{N_u+1}^{(u)}$ represents the regenerated recommended item based on the semantic pseudo interaction sequence. Since our proposed CESRec is model-agnostic, it can be seamlessly integrated with existing sequential recommendation models.

5 Experimental Setup

5.1 Dataset and Evaluation Metric

Dataset	#User	#Item	#Review	#Density
Video Games	55,223	17,408	496,315	0.051628%
Toys	208,180	78,772	1,826,430	0.011138%
MovieLens	6,040	3,883	1,000,209	4.264680%

Table 1: Statistics of three datasets.

We conduct experiments on two commonly used recommendation datasets, Video Games and Toys, constructed from the Amazon review datasets (Ni et al., 2019). We also employ the MovieLens datasets (Harper and Konstan, 2015) which is a widely adopted dataset for sequential recommendation tasks, which contains user interactions with movies. Statistics are shown in Table 1.

We adopt two widely used metrics to evaluate the performance: Normalized Discounted Cumulative Gain (NDCG@K) and Hit Rate (HR@K) with K=5,10. We select 100 non-interacted items to construct the candidate set, ensuring the inclusion of the correct subsequent item.

Dataset	Traditional Model	HR@5	NDCG@5	HR@10	NDCG@10	LLM-based Model	HR@5	NDCG@5	HR@10	NDCG@10
Video Games	SASRec	0.590	0.4629	0.717	0.5042	LLaRA	0.270	0.2277	0.360	0.2558
	+CESRec-LLaMA2	0.633	0.4847	0.725	0.5144	+CESRec-LLaMA2	0.380	0.3097	0.450	0.3316
	+CESRec-LLaMA3	0.646	0.4923	0.745	0.5242	+CESRec-LLaMA3	0.380	0.3254	0.440	0.3445
Movielens	SASRec	0.757	0.5688	0.866	0.6045	LLaRA	0.170	0.1416	0.210	0.1542
	+CESRec-LLaMA2	0.824	0.6076	0.882	0.6264	+CESRec-LLaMA2	0.260	0.2192	0.310	0.2347
	+CESRec-LLaMA3	0.810	0.5996	0.886	0.6244	+CESRec-LLaMA3	0.280	0.2348	0.330	0.2508
Toys	SASRec	0.431	0.3173	0.537	0.3509	LLaRA	0.420	0.3957	0.430	0.3986
	+CESRec-LLaMA2	0.472	0.3376	0.557	0.3647	+CESRec-LLaMA2	0.500	0.4671	0.590	0.4955
	+CESRec-LLaMA3	0.478	0.3408	0.557	0.3659	+CESRec-LLaMA3	0.500	0.4671	0.600	0.4993

Table 2: Performance on three datasets. We apply our proposed CESRec on two strong SRS: SASRec and LLaRA, and we implement CESRec based on two LLM backbones: LLaMA2 and LLaMA3.

Dataset	Method	HR@5	NDCG@5	HR@10	NDCG@10
Video Games	+CESRec-LLaMA3	0.646	0.4923	0.745	0.5242
	+CESRec w/o d.a.	0.634	0.4849	0.723	0.5136
	+CESRec w/o c.	0.610	0.4711	0.723	0.5077
	SASRec	0.590	0.4629	0.717	0.5042
Movielens	+CESRec-LLaMA3	0.810	0.5996	0.886	0.6244
	+CESRec w/o d.a.	0.805	0.5940	0.880	0.6186
	+CESRec w/o c.	0.774	0.5766	0.866	0.6061
	SASRec	0.757	0.5688	0.866	0.6045
Toys	+CESRec-LLaMA3	0.478	0.3408	0.557	0.3659
	+CESRec w/o d.a.	0.468	0.3354	0.557	0.3638
	+CESRec w/o c.	0.443	0.3222	0.530	0.3501
	SASRec	0.431	0.3173	0.537	0.3509

Table 3: Performance of ablation models. We conduct ablation study on SASRec+CESRec.

5.2 Implementation Detail

To evaluate our method, we employ gpt-4o-mini as a user simulator to provide natural language feedback (details in Appendix 9.2).

We employ LoRA technique to fine-tune the LLM as the constructor module.

For the sequential recommendation method, SASRec (Kang and McAuley, 2018), we train the model on all three datasets using the Adam optimizer (Kingma, 2014) for 200 epochs, with a learning rate of 0.001 and a batch size of 256. For the LLM-based recommendation method, LLaRA (Liao et al., 2024), the original configuration selects the top-ranked item from the candidate set as the recommendation result.

To ensure consistency with our experimental setup, we adopt the ranking method from (Wang et al., 2024), which ranks the candidate items based on the cosine similarity between item embeddings and the output embeddings of LLaRA. In our CESRec, we mask one item in three datasets. We implement our CESRec using two LLMs as the backbone: LLaMA-2-7b (Touvron et al., 2023) and LLaMA-3-8b (Dubey et al., 2024). And we use the same user simulator as the previous conversational recommendation studies Fang et al. (2024) when training and evaluating the models.

5.3 Baselines

We conducted experiments using two strong SRS backbones: (1) **SASRec** (Kang and McAuley, 2018) is a widely used sequential recommendation model that employs a self-attention mechanism to effectively capture relationships between items within a user’s interaction sequence. (2) **LLaRA** (Liao et al., 2024) is an LLM-based recommendation model that utilizes a hybrid prompting approach, combining ID-based and text-based representations of items as input. This model aims to enhance recommendation accuracy by integrating both structured and unstructured data sources.

6 Experimental Results

6.1 Main Results

We evaluate the performance of our proposed CESRec and baseline methods on three datasets using four evaluation metrics. As shown in Table 2, SASRec+CESRec and LLaRA+CESRec consistently outperform their corresponding base SRS model (*a.k.a.*, SASRec and LLaRA) across all datasets and metrics. This demonstrates that the semantic-based pseudo interaction sequences, which incorporate users’ current feedback, enable recommendation models to more effectively capture users’ real-time preferences.

Secondly, CESRec demonstrates improved performance when leveraging larger LLMs as the backbone, suggesting that more powerful LLMs possess the stronger capability to accurately model user preferences and select relevant replacement items.

6.2 Ablation Study

To validate the effectiveness of each module, we compare the performance of the following variants of CESRec-LLaMA3 on the SASRec backbone: (1) **CESRec w/o d.a.**: we solely employ user conversational feedback to construct pseudo interaction

LLM Backbones	HR@5	NDCG@5	HR@10	NDCG@10
SASRec	0.590	0.4629	0.717	0.5042
+CESRec-Qwen2.5-3B-Instruct	0.649	0.4929	0.717	0.5146
+CESRec-LLaMA2-7B-Instruct	0.633	0.4847	0.725	0.5144
+CESRec-Mistral-7B-Instruct-v0.3	0.646	0.4906	0.734	0.5188
+CESRec-Qwen2.5-7B-Instruct	0.644	0.4911	0.732	0.5195
+CESRec-LLaMA3-8B-Instruct	<u>0.646</u>	<u>0.4923</u>	0.745	0.5242

Table 4: Performance comparison of various LLM backbones on the Video Game dataset.

sequences and remove the dual alignment from CESRec. (2) **CESRec w/o c.**: we only leverage dual alignment method to mask outlier items and do not construct pseudo sequence. The results, as shown in Table 3, demonstrate that all modules in the model contribute to enhancing sequential recommendation. The superior performance of CESRec-LLaMA3 over CESRec w/o d.a. indicates that the dual alignment outlier items masking method enables CESRec to concentrate on user’s main preference, and construct semantic pseudo sequences that better align with user preferences.

6.3 The Influence of LLM Backbones

To assess the influence of varying LLM backbones on recommendation accuracy, we evaluated diverse LLM architectures and scales on Video Games dataset. Employing the proposed CESRec framework, the small language model (SLM) Qwen2.5-3B-Instruct shows improved recommendation performance in resource-constrained settings, highlighting CESRec’s effectiveness under limited resources. Its fewer hidden layers compared to LLMs reduce semantic loss when mapping from the LLM’s latent space to the Recsys’s embedding space, enabling more accurate masking of outlier items. For models of comparable size, Qwen2.5-7B-Instruct and Mistral-7B-Instruct-v0.3 consistently surpass Llama-2-7B-Instruct, suggesting their generated pseudo-interaction sequences more accurately reflect real-time user interests.

6.4 The Impact of Historical Interaction Sequence Length




Figure 3: Performance of using different lengths of the historical interaction sequence.

To investigate the impact of historical interaction sequence length, we evaluate model performance using different sequence lengths in terms of HR@5 and NDCG@5 on the Toys and Video Games datasets. As shown in Figure 3, the results demonstrate that our proposed CESRec consistently outperforms the baseline SASRec across all three sequence length ranges. This demonstrates the robustness of our model in effectively handling historical interaction sequences of varying lengths, further confirming its adaptability in diverse recommendation scenarios.

6.5 Analysis of Interaction Numbers




Figure 4: The influence of different numbers of user natural language feedback.

We further investigate the impact of the number of natural language feedback of CESRec-LLaMA3, based on SASRec. As illustrated in Figure 4, as the number of interactions between users and the CESRec-LLaMA3 increases, the performance of the Recsys consistently improves. The HR@K and NDCG@K metrics (with K=5, 10) demonstrate a steady upward trend across all three real-world datasets. This indicates that as users provide more feedback, the Recsys becomes increasingly effective at capturing users’ real-time interests. By constructing semantic-based pseudo interaction sequences that reflect these interests, the system generates recommendations that better align with users’ current preferences.

The improvement in both HR and NDCG metrics demonstrate the model’s enhanced ability to both identify and effectively rank relevant items, yielding more user-centric recommendations.

6.6 Analysis of Masking Outlier Items

We further investigated the impact of the number of masked outlier items on the performance of CESRec. The results show that for the MovieLens and




Figure 5: The impact of masking different numbers of outlier items.

Video Games datasets, the model achieves optimal performance when the number of masked items is set to 2. Beyond this threshold, performance begins to decline as the number of masked items increases. This decline can be attributed to the fact that excessive masking reduces the length of the user’s historical sequence, leading to a loss of valuable information regarding user preferences. Consequently, the model struggles to accurately capture user behavior and predict items that align with these preferences. In contrast, for the Toys dataset, the model’s performance improves as the number of masked items increases. This trend can be attributed to the higher sparsity of the Toys dataset compared to other two datasets, as shown in Table 1. With greater sparsity, the items in the constructed sequences exhibit more variability, and as the model adjusts these sequences based on user feedback expressed in natural language, the impact on the recommendation outcomes becomes more notable. Therefore, by masking items that deviate from the user’s preferences, the model can concentrate on the most relevant interactions, resulting in improved performance.

6.7 Case Study

To intuitively validate the effectiveness of our proposed CESRec, we randomly select an example from MovieLens dataset, as shown in Figure 6. The detail user’s historical interactions with movies are shown in Appendix 9.7. Given this sequence as input, SASRec generates “Jack Frost” as a recommended item by capturing the co-occurrence relationships between movies. However, “Jack Frost” is a comedy film, which does not align with the user’s current preference for horror films. To encourage the model’s focus on the user’s core interests, we employ the dual alignment outlier items



Figure 6: A case study of CESRec.

masking method. This method masks the “Super Mario Bros.”, which belongs to the action/animation genre and deviates from user’s core preference for horror films. Thus, the model can better align with user’s primary interests and improve recommendation accuracy. This masking process enables the CESRec to better concentrate on the user’s core preferences. Since “Jack Frost” is inconsistent with the user’s preference, CESRec constructs a semantic-based pseudo-interaction sequence incorporating the user’s conversational feedback: “I don’t like comedy; I prefer horror.”. During this process, CESRec replaces “Cops and Robbersons (comedy)” with “Carnosaur 2 (horror)” to reinforce the user’s stated preference. Ultimately, based on this refined interaction sequence, CESRec predicts “Halloween: H20” as the recommended item.

7 Conclusion

In this paper, we proposed Conversation Enhanced Sequential Recommendation (CESRec), a novel framework that seamlessly integrates the long-term preference modeling of SRS with the real-time preference elicitation of CRS. By leveraging users’ natural language feedback, CESRec dynamically refines historical interaction sequences to generate pseudo-interaction sequences that capture both long-term preferences and real-time interests. Additionally, the dual alignment outlier items masking method addresses the challenge of outlier items in historical sequences by accurately identifying and masking items that deviate from users’ core prefer-

ences. Extensive experiments on three real-world datasets demonstrate that CESRec enhances the performance of SOTA SRS models, achieving superior results in terms of HR and NDCG metrics.

Limitations

Our method relies on user conversational feedback to dynamically refine the historical interaction sequence, aiming to better align with the user's real-time preferences. However, if the user's feedback is expressed in a vague, ambiguous, or unclear manner, the model may fail to capture the user's real-time preferences accurately, leading to the generation of an imprecise pseudo-interaction sequence, which in turn affects the recommendation performance. In future work, we will investigate more sophisticated dialogue mechanisms that can effectively guide users to articulate their latent preferences.

Ethical Considerations

The research conducted in this paper centers on investigating the effectiveness of leveraging LLMs to bridge the gap between conversational recommendation and sequential recommendation. Our work systematically benchmarks LLMs under various real-world scenarios and evaluates their performance. In the process of conducting this research, we have adhered to ethical standards to ensure the integrity and validity of our work. To minimize potential bias and ensure fairness, we employ the same prompts and experimental setups as those used in existing publicly accessible and freely available studies. We have made every effort to ensure that our research does not harm individuals or groups and does not involve any form of deception or misuse of information.

8 Acknowledgments

This work was supported by the National Natural Science Foundation of China (T2293773, 62432002, 62406061), and the Natural Science Foundation of Shandong Province (ZR2023QF159), sponsored by the CCF-DiDi GAIA Collaborative Research Funds (CCF-DiDi GAIA 202504), the CCF-Huawei Populus Grove Fund (CCF-HuaweiDB202509).

References

- Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1007–1014.
- Lisi Chen, Gao Cong, Christian S Jensen, and Dingming Wu. 2013. Spatial keyword query processing: An experimental evaluation. *Proceedings of the VLDB Endowment*, 6(3):217–228.
- Lisi Chen and Shuo Shang. 2019. Region-based message exploration over spatio-temporal data streams. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 873–880.
- 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. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1126–1132.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Hui Fang, Danning Zhang, Yiheng Shu, and Guibing Guo. 2020. Deep learning for sequential recommendation: Algorithms, influential factors, and evaluations. *ACM Transactions on Information Systems (TOIS)*, 39(1):1–42.
- Jiabao Fang, Shen Gao, Pengjie Ren, Xiuying Chen, Suzan Verberne, and Zhaochun Ren. 2024. A multi-agent conversational recommender system. *arXiv preprint arXiv:2402.01135*.
- Yue Feng, Shuchang Liu, Zhenghai Xue, Qingpeng Cai, Lantao Hu, Peng Jiang, Kun Gai, and Fei Sun. 2023. A large language model enhanced conversational recommender system. *arXiv preprint arXiv:2308.06212*.
- Luke Friedman, Sameer Ahuja, David Allen, Zhenning Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, et al. 2023. Leveraging large language models in conversational recommender systems. *arXiv preprint arXiv:2305.07961*.
- Shen Gao, Xiuying Chen, Piji Li, Zhaochun Ren, Li-dong Bing, Dongyan Zhao, and Rui Yan. 2019. Abstractive text summarization by incorporating reader comments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6399–6406.
- Shen Gao, Zhengliang Shi, Minghang Zhu, Bowen Fang, Xin Xin, Pengjie Ren, Zhumin Chen, Jun Ma, and

- Zhaochun Ren. 2024. Confucius: Iterative tool learning from introspection feedback by easy-to-difficult curriculum. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 18030–18038.
- Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems*, pages 299–315.
- Peng Han, Peng Yang, Peilin Zhao, Shuo Shang, Yong Liu, Jiayu Zhou, Xin Gao, and Panos Kalnis. 2019. Gcn-mf: disease-gene association identification by graph convolutional networks and matrix factorization. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 705–713.
- F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19.
- Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 639–648.
- Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian McAuley. 2023. Large language models as zero-shot conversational recommenders. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 720–730.
- Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*, pages 364–381. Springer.
- Jun Hu, Wenwen Xia, Xiaolu Zhang, Chilin Fu, Weichang Wu, Zhaoxin Huan, Ang Li, Zuoli Tang, and Jun Zhou. 2024. Enhancing sequential recommendation via llm-based semantic embedding learning. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 103–111.
- Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*, pages 197–206. IEEE.
- Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Wenqiang Lei, Gangyi Zhang, Xiangnan He, Yisong Miao, Xiang Wang, Liang Chen, and Tat-Seng Chua. 2020. Interactive path reasoning on graph for conversational recommendation. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 2073–2083.
- Chengxi Li, Yejing Wang, Qidong Liu, Xiangyu Zhao, Wanyu Wang, Yiqi Wang, Lixin Zou, Wenqi Fan, and Qing Li. 2023a. Strec: Sparse transformer for sequential recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 101–111.
- Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. 2023b. [Text is all you need: Learning language representations for sequential recommendation](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '23, page 1258–1267, New York, NY, USA. Association for Computing Machinery.
- Muyang Li, Zijian Zhang, Xiangyu Zhao, Wanyu Wang, Minghao Zhao, Runze Wu, and Ruocheng Guo. 2023c. Automlp: Automated mlp for sequential recommendations. In *Proceedings of the ACM Web Conference 2023*, pages 1190–1198.
- Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. Llra: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1785–1795.
- Yujie Lin, Chenyang Wang, Zhumin Chen, Zhaochun Ren, Xin Xin, Qiang Yan, Maarten de Rijke, Xiuzhen Cheng, and Pengjie Ren. 2023. A self-correcting sequential recommender. In *Proceedings of the ACM Web Conference 2023*, pages 1283–1293.
- Qidong Liu, Xian Wu, Wanyu Wang, Yejing Wang, Yuanshao Zhu, Xiangyu Zhao, Feng Tian, and Yefeng Zheng. 2024. Large language model empowered embedding generator for sequential recommendation. *arXiv preprint arXiv:2409.19925*.
- Ben Mann, N Ryder, M Subbiah, J Kaplan, P Dhariwal, A Neelakantan, P Shyam, G Sastry, A Askell, S Agarwal, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 1.
- Sheshera Mysore, Andrew McCallum, and Hamed Zamani. 2023. Large language model augmented narrative driven recommendations. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 777–783.
- Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 188–197.
- Yunzhu Pan, Chen Gao, Jianxin Chang, Yanan Niu, Yang Song, Kun Gai, Depeng Jin, and Yong Li. 2023.

- Understanding and modeling passive-negative feedback for short-video sequential recommendation. In *Proceedings of the 17th ACM conference on recommender systems*, pages 540–550.
- Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36:10299–10315.
- Shuo Shang, Lisi Chen, Zhewei Wei, Christian S Jensen, Ji-Rong Wen, and Panos Kalnis. 2015. Collective travel planning in spatial networks. *IEEE Transactions on Knowledge and Data Engineering*, 28(5):1132–1146.
- Shuo Shang, Ruogu Ding, Kai Zheng, Christian S Jensen, Panos Kalnis, and Xiaofang Zhou. 2014. Personalized trajectory matching in spatial networks. *The VLDB Journal*, 23(3):449–468.
- Leheng Sheng, An Zhang, Yi Zhang, Yuxin Chen, Xiang Wang, and Tat-Seng Chua. 2024. Language models encode collaborative signals in recommendation.
- Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 1441–1450.
- Jiaxi Tang and Ke Wang. 2018. Personalized top-n sequential recommendation via convolutional sequence embedding. In *Proceedings of the eleventh ACM international conference on web search and data mining*, pages 565–573.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Bohao Wang, Feng Liu, Jiawei Chen, Yudi Wu, Xingyu Lou, Jun Wang, Yan Feng, Chun Chen, and Can Wang. 2024. Llm4dsr: Leveraging large language model for denoising sequential recommendation. *arXiv preprint arXiv:2408.08208*.
- Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. 2021. Denoising implicit feedback for recommendation. In *Proceedings of the 14th ACM international conference on web search and data mining*, pages 373–381.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.
- Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1059–1068.
- Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards topic-guided conversational recommender system. *arXiv preprint arXiv:2010.04125*.

9 Appendix

9.1 Theoretical Foundation for Dual Alignment Strategies

Our dual alignment strategy maps LLM-derived semantic embeddings $\mathbf{e}_i^{\text{LLM}}$ to the recommendation space via an MLP-based **Adapter**, minimizing the MSE loss between hybrid embeddings $\mathbf{e}_i^{\text{hybrid}}$ and collaborative embeddings $\mathbf{e}_i^{\text{collab}}$.

$$\begin{aligned}\mathbf{e}_i^{\text{hybrid}} &= \text{Adapter}(\mathbf{e}_i^{\text{LLM}}) \\ &= W_2 \cdot \text{GELU}(W_1 \cdot \mathbf{e}_i^{\text{LLM}} + b_1) + b_2\end{aligned}\quad (10)$$

$$L_{\text{align}} = \|\mathbf{e}_i^{\text{hybrid}} - \mathbf{e}_i^{\text{collab}}\|_2^2 \quad (11)$$

This ensures $\mathbf{e}_i^{\text{hybrid}}$ fuses semantic and collaborative signals.

9.2 User Simulation

Inspired by the principle that "while users may not always articulate their preferences, they can reliably identify what they dislike," our model initially generates recommendations based on the user's historical interactions. If the user expresses dissatisfaction with a recommended item, they provide natural language feedback highlighting the specific features they find undesirable.

Following the work Fang et al. (2024), we adopt a similar user simulation mechanism. To mitigate potential data leakage, the user simulator is not directly exposed to the target item itself; instead, it receives only descriptive information about the target item. The user simulator is prompted to provide feedback as follows: "You are a user interacting with a recommender system. Based on the information about your <target item> and the <recommended item> provided by the recommender, give feedback to the recommender."

Model	NDCG@10	Recall@10	NDCG@50	Recall@50	MRR
Recformer	0.0323	0.0625	0.04489	0.1125	0.02906
+ CESRec + Mistral v0.3-7B	0.0323	0.0625	0.05089	0.1375	0.03053
+ CESRec + LLaMA3-8B	0.0368	0.0750	0.05859	0.1625	0.03321
+ CESRec + Qwen2.5-7B	0.0411	0.0875	0.06561	0.1875	0.03518

Table 5: The Recformer experimental results on the Video Game dataset.

9.3 Influence of Feedback Types

The polarity of user feedback (positive or negative) influences the interaction sequence reconstruction process. The table 9.7 illustrates CESRec’s adaptive capability in responding to varying feedback polarities.

9.4 Incorporate with Language-based Baseline

We integrate the CESRec into the Recformer (Li et al., 2023b) and evaluated this approach on the Video Game dataset. The experimental results are shown in Table 5, demonstrating that after inserting our CESRec, Recformer shows improvements in NDCG@10, Recall@10, NDCG@50, Recall@50, and MRR metrics. Additionally, we selected different LLMs as baselines. Due to differences in semantic understanding capabilities, the enhancement effects vary across LLM baselines. In the experiments enhancing Recformer, the best performance was achieved using Qwen2.5-7b as the baseline.

9.5 Alternative Similarity Functions

We adopted cosine similarity for semantic alignment, as prior work (Sheng et al., 2024) demonstrated its effectiveness in evaluating semantic proximity. To assess robustness, we replaced cosine similarity with Euclidean distance on the Video Game dataset. As shown in table 6, the results confirmed that cosine similarity outperformed alternatives in identifying outlier items and improving the precision of the recommendations.

Model	HT@5	NDCG@5	HT@10	NDCG@10
SASRec	0.590	0.4629	0.717	0.5042
+CESRec-llama3-L2	0.644	0.4941	0.722	0.5191
+CESRec-llama3-cosine	0.646	0.4923	0.745	0.5242

Table 6: Performance of Different Similarity Function

9.6 Inference Latency

To evaluate practical deployment feasibility, we measure the average latency of LLaMA3-8B-CESRec and Qwen2.5-7B-CESRec when perform-

ing two key operations: (1) outlier item masking and (2) pseudo-sequence construction, across three benchmark datasets:

Model	Component	Video Game	Toys	Movielens
CESRec-LLaMA3-8B	Outlier Items Masking	0.1407s	0.1597s	0.1757s
	Pseudo Sequence Construction	10.4124s	7.6398s	9.1576s
	Total Time	10.5532s	7.7996s	9.3334s
CESRec-Qwen2.5-7B	Outlier Items Masking	0.1520s	0.1806s	0.1722s
	Pseudo Sequence Construction	6.8974s	7.9245s	4.0569s
	Total Time	7.0494s	8.1051s	4.2292s

Table 7: Inference Latency Comparison (Average Time)

We calculated the average inference time across 10 samples. The real-time inference latency for outlier item masking ranges between 0.14-0.17s, while pseudo sequence construction constitutes the major portion of the inference delay. Based on real-time user feedback, the total inference time for LLaMA3-8B to construct pseudo sequences and generate new recommendations ranges between 7-10s, while Qwen2.5-7B ranges between 4-8s. Since different models have varying inference speeds, faster models can be selected based on strictly demanding production requirements.

9.7 Detail Information of Case study

User Interaction Sequence of Case Study

- I Still Know What You Did Last Summer
- Jungle 2 Jungle
- Two if by Sea
- M. Butterfly
- Super Mario Bros
- Blank Check
- Repossessed
- The Evening Star
- The Beautician and the Beast
- Mr. Wrong
- A Night at the Roxbury
- Halloween: The Curse of Michael Myers
- Stop! Or My Mom Will Shoot
- Cops and Robbersons

Examples of bias migration analysis of CESRec

User Interaction Sequence:

['Dumb and Dumber', 'The Hangover',
'Bridesmaids', 'Anchorman', 'The Exorcist', 'Hereditary', 'The Conjuring', 'A Nightmare on Elm Street', 'The Babadook', 'It', 'Superbad', 'Step Brothers']
.....

Positive Feedback ("I like comedies")

Replaces <The Conjuring (horror)> to
<Dogma (comedy)> to amplify comedy recommendations.

Sequence reconstructed by CESRec:

["Dumb and Dumber", "The Hangover",
"Bridesmaids", "Anchorman", "The Exorcist", "Hereditary", "Dogma", "A Nightmare on Elm Street", "The Babadook", "It", "Superbad", "Step Brothers"]
.....

Negative Feedback ("I dislike comedies")

Replaces <Bridesmaids (comedy)> to
<Sleepy Hollow (horror)> to align with implicit horror preferences.

Sequence reconstructed by CESRec:

["Dumb and Dumber", "The Hangover",
"Sleepy Hollow", "Anchorman", "The Exorcist", "Hereditary", "The Conjuring", "A Nightmare on Elm Street", "The Babadook", "It", "Superbad", "Step Brothers"]