

# ThinkRec: Thinking-based recommendation via LLM

Anonymous ACL submission

## Abstract

Recent advances in large language models (LLMs) have enabled more semantic-aware recommendations through natural language generation. Existing LLM for recommendation (LLM4Rec) methods mostly operate in a System 1-like manner, relying on superficial features to match similar items based on click history, rather than reasoning through deeper behavioral logic. This often leads to superficial and erroneous recommendations. Motivated by this, we propose **ThinkRec**, a thinking-based framework that shifts LLM4Rec from System 1 to System 2 (rational system). Technically, ThinkRec introduces a thinking activation mechanism that augments item metadata with keyword summarization and injects synthetic reasoning traces, guiding the model to form interpretable reasoning chains that consist of analyzing interaction histories, identifying user preferences, and making decisions based on target items. On top of this, we propose an instance-wise expert fusion mechanism to reduce the reasoning difficulty. By dynamically assigning weights to expert models based on users’ latent features, ThinkRec adapts its reasoning path to individual users, thereby enhancing precision and personalization. Extensive experiments on real-world datasets demonstrate that ThinkRec significantly improves the accuracy and interpretability of recommendations<sup>1</sup>.

## 1 Introduction

Recommendation systems are indispensable in modern digital platforms, enabling users to navigate vast content efficiently (Glenski et al., 2021; Ma et al., 2023; Iana et al., 2024). Traditional sequential recommendation methods rely on implicit modeling of user interaction histories, cannot model context and incorporate broader knowledge, which restricts their reasoning and generalization

ability. Recent advances in LLMs offer strong contextual comprehension and extensive world knowledge to improve recommendation systems.

Prior solution paradigms of LLM4Rec can be separated into three categories: (1) *item scoring* (Liu et al., 2024; Zhang et al., 2025), where LLMs answer binary preference questions given user and item context; (2) *item generation* (Li et al., 2023; Bao et al., 2024), which maps natural language prompts to item IDs through aligned embedding spaces; and (3) *hybrid modeling* (Geng et al., 2023; Zhang et al., 2024b), where a single LLM performs multiple tasks such as pointwise prediction, pairwise ranking, or listwise generation. Although these LLM4Rec methods differ in output formats and representation learning strategies, they fundamentally resemble *System 1*—the intuitive system—in cognitive science (Kahneman, 2011). They tend to match similar items based on click history while overlooking the deeper behavioral logic. This limitation becomes evident in cases such as the one shown in Figure 1(b), where the user’s behavior over time is: dislikes “Dune”, likes “The Three-Body Problem”, and likes “Foundation” (all three are science fiction). Methods that rely on System 1-like intuition tend to infer that the user would also like “Hyperion” (science fiction) simply because it belongs to the same genre. In reality, the user might dislike philosophical or metaphysical themes in fiction, which are prominent in \*Hyperion\*, making it an unsuitable recommendation. Clearly, if we can leverage the vast knowledge encoded in LLMs and fully activate their reasoning capabilities for user behavior modeling, recommendation performance can be significantly improved.

This motivates our effort to push LLM4Rec from a *System 1* paradigm toward a more rational, *System 2-like reasoning framework*. We raise two key questions for the advancement: 1) How to effectively balance recommendation objectives with language modeling tasks to fully exploit the reason-

<sup>1</sup>Our implementations are available in anonymous Github: [https://anonymous.4open.science/r/ThinkRec\\_LLM](https://anonymous.4open.science/r/ThinkRec_LLM)

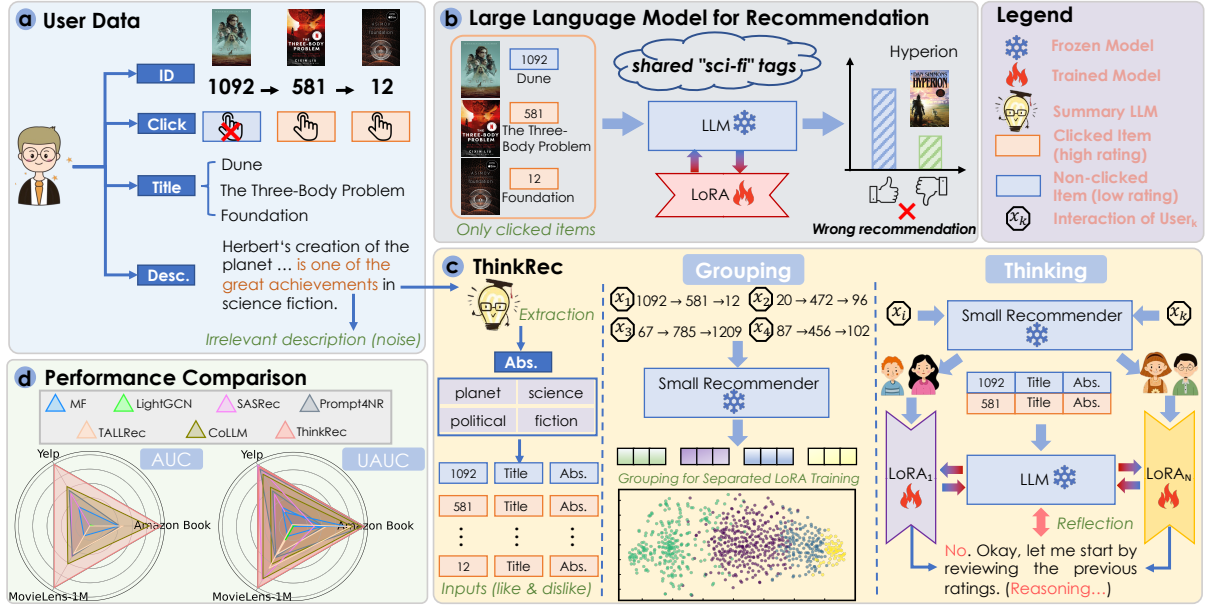


Figure 1: (a) shows the composition of user interaction data. (b) and (c) illustrate previous LLM-based recommendations and our ThinkRec, respectively. (d) compares ThinkRec with baselines in three real-world datasets.

ing capabilities of LLMs. Existing methods often prioritize direct recommendation metrics such as hit rate or ranking accuracy, overlooking LLMs' strengths in semantic reasoning. However, blind reinforcement thinking can lead to simple next token prediction, defeating the goal of recommendation. 2) How to think more effectively in the presence of diverse user behaviors and underlying preferences. As shown in Figure 1(c), user behaviors vary widely, and uniform modeling tends to obscure personality preferences while introducing noise that distracts the model from salient signals, impairing its reasoning capacity. Moreover, inferring user intent solely from high-rating items and generic world knowledge limits the informational basis for accurate preference reasoning.

To address these challenges, we propose **Thinking-based Recommendation** via LLM, abbreviated as **ThinkRec**. One of the main problems faced is that the data and optimization goals of recommendation tasks lack the ability to activate thinking in LLMs (Challenge 1). To overcome this challenge, we designed the item augmentation and thinking activation framework for finetuning. The fine-tuned model analyzes associations in historical item information, determines user preferences, and gives explicit reasons while deriving recommendations. We extracted the metadata keywords of the items with the help of an existing summarization model as the augmentation information of the items to support the reasoning. In addition, rea-

soning data is synthesized using a strong reasoning model, and the reasoning capability is distilled to the local model by mixed sampling of reasoning and recommendation data. Therefore, item augmentation and thinking activation become a bridge connecting the recommendation task and the language task, making recommendations traceable. To address the difficulty of thinking diversely with rich information (Challenge 2), we add the user's preferences (yes/no) of items to the prompts and generate personalized recommendation experts based on latent user features. In the technique, we design a dynamic Low-Rank Adaptation (LoRA) fusion method. Users are grouped by latent user features extracted from traditional recommendation models as shown in Figure 1(c). A set of base LoRAs can be represented and fine-tuned using the grouped data, and the engagement level of each LoRA can be determined through a gating mechanism. The difficulty of thinking is reduced through information classification and personalization mechanisms.

We conduct experiments on three real-world recommendation datasets, validating the rationality and effectiveness of ThinkRec. ThinkRec average outperforms state-of-the-art baselines by 7.96% in AUC and by 56.54% in METEOR. In summary, the main contributions of this work are threefold:

- We analyze the key challenges of LLM-based recommendations and almost the first emphasize the importance of activating thinking to make recommendations more interpretable.

- We propose the ThinkRec, including item augmentation and a thinking activation framework, as well as a latent user feature-based gating mechanism for instance-wise recommendation expert fusion.
- Extensive experiments on three real-world datasets validate the effectiveness and reasonableness of ThinkRec.

## 2 Related Work

### 2.1 LLM-based Recommendation

With the rapid progress of LLMs, there has been growing interest in adapting them to recommendation tasks using both textual and structured data. Early works (Bao et al., 2023; Ho et al., 2023; Lin et al., 2024a; Zhang et al., 2025) rely on prompt-based scoring by converting recommendations into binary question-answering, while others (Ji et al., 2023a; Li et al., 2023; Lin et al., 2024b) align language space with item embeddings to directly generate item IDs. More advanced approaches, such as P5 (Geng et al., 2022) and InstructRec (Zhang et al., 2024b) unify multiple sub-tasks—like rating prediction, pairwise comparison, or ranking, into a single language modeling framework. These methods demonstrate LLMs’ flexibility in expressing recommendation semantics, yet most focus on adapting output formats, treating LLMs as static scorers or selectors with limited interpretability.

Beyond static prompting, recent efforts shift toward viewing LLMs as interactive agents capable of multi-step reasoning and personalized decision-making. For example, RecMind (Wang et al., 2024a) and MACRec (Wang et al., 2024b) introduce modularized reasoning structures, enabling planning-based recommendation. Others explore integrating LLMs into various stages of the pipeline—for instance, generating user or item representations (Zhang et al., 2024c) or acting as high-level decision controllers (Park and Zhang, 2025). Meanwhile, instruction tuning (Zhang et al., 2024b), chain-of-thought prompting (Liu et al., 2025) have been proposed to enhance LLMs’ semantic understanding. Despite these developments, explicit reasoning supervision and fine-grained user modeling remain underexplored. ThinkRec addresses these gaps by integrating reasoning-augmented training and personalized expert fusion, enabling interpretable, user-specific recommendations grounded in structured reasoning.

### 2.2 Reasoning Model

Recent advances have transformed LLMs from passive token predictors into structured reasoning agents by incorporating process-level supervision. Techniques such as Chain-of-Thought prompting (Wei et al., 2022), ReAct (Yao et al., 2022), and Tree-of-Thoughts (Yao et al., 2023) enable models to generate intermediate reasoning steps, improving interpretability and multistep inference quality. To further reinforce reasoning capabilities, recent works have introduced process reward models (Zhang et al., 2024a) and self-improvement pipelines using techniques like Monte Carlo Tree Search (Luo et al., 2024) and reasoning distillation (Xu et al., 2023). These trends define the emerging paradigm of Large Reasoning Models, which prioritize explicit, verifiable thinking processes over direct output. ThinkRec draws inspiration from this direction by introducing reasoning-augmented training for recommendation, aligning to embed structured reasoning into decision making.

## 3 Method

### 3.1 Preliminary

**Problem Definition.** In the view of sequential recommendation, datasets can be formulated as  $\mathcal{S} = \{(x_{u,t}, y_{u,t})\}_{u=1,2,\dots,N, t=1,2,\dots,T_u}$ , where  $x_{u,t} = \{y_{u,1:(t-1)}\}$  denotes a user’s historical behaviors prior to the  $t$ th behavior  $y_{u,t}$  and arranged in a chronological order, and  $T_u$  denotes the number of behaviors for the user  $u$ . We will drop the sub-scripts occasionally and write  $(x, y)$  in place of  $(x_{u,t}, y_{u,t})$  for simplicity. Each behavior  $y$  is represented as  $(u, t, i_{id}, l)$ , where  $i_{id}$  is an item ID and  $l \in \{1, 0\}$  indicating the interaction label. Furthermore, additional textual information about items is provided, mainly the item title and description. Given  $x_{u,t}$ , get the textual information, and label to construct a historical text prompt to determine whether or not you would enjoy an item composed of the title and description of  $y_{u,t}$  and give a reason  $r_{u,t}$ . We aim to enable the LLM to provide the underlying thinking beyond merely making recommendations as in previous work, thus achieving more effective and interpretable recommendations.

**Integrating Collaborative Embeddings into LLMs.** The sequential recommenders, such as MF (Koren et al., 2009) and LightGCN (He et al., 2020), are used for encoding collaborative information, which can be integrated into LLM embed-

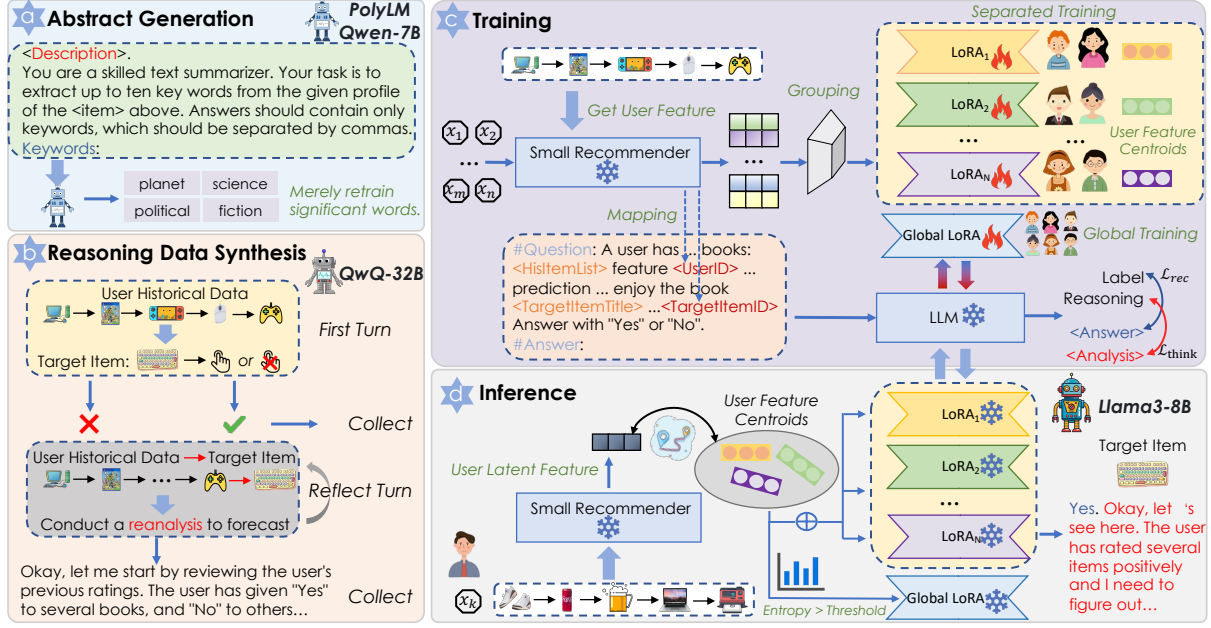


Figure 2: Overview of the proposed ThinkRec. Keywords are extracted from the description of items with the pretrained PolyLM-Qwen-7B. Reasoning data is synthesized with a reasoning model from a sample of training data. Global LoRA is trained with total data, and base LoRAs are tuned on it with data grouping by user feature. The used LoRA, fused or single, is determined by the similarity between the user feature and LoRA representations.

dings. Formally, for each sample  $(u, t, i_{id}, l)$ :

$$\mathbf{e}_s^u = f_\psi(u; \mathcal{S}); \mathbf{e}_s^i = f_\psi(i_{id}; \mathcal{S}), \quad (1)$$

where  $\mathbf{e}_s^u \in \mathbb{R}^{1 \times d_1}$  denotes the user's representation with dimension  $d_1$ ,  $f_\psi(u; \mathcal{S})$  denotes the process used to obtain this representation, similarly for  $i_{id}$ . Formally, for an item  $i$  with the text metadata  $i_{txt}$ , we obtain its textual tokens  $\mathbf{E}_{txt}$  as follows:

$$\mathbf{E}_{txt} = \text{WE}(\text{TKZ}(i_{txt})), \quad (2)$$

where  $\text{TKZ}(\cdot)$ ,  $\text{WE}(\cdot)$  present the LLM tokenizer and word embedding layer, respectively.  $\mathbf{E}_{txt} = \{\text{emb}_{txt}^{1:L}\}$ , where  $L$  is the number of tokens and  $\text{emb}_{txt}^1 \in \mathbb{R}^{1 \times d_2}$  denotes embeddings in the language space with dimension  $d_2$ . To align  $\mathbf{e}_s^i, \mathbf{e}_s^u$  to LLMs' language space,  $\text{proj}_\phi(\cdot)$  is used:

$$\text{emb}_s^i = \text{proj}_\phi(\mathbf{e}_s^i), \text{emb}_s^u = \text{proj}_\phi(\mathbf{e}_s^u). \quad (3)$$

### 3.2 Thinking Enhanced Recommendation

To align recommendation with language modeling tasks, we highlight the importance of thinking activation for LLMs and propose a framework that consists of data construction and jointly training.

#### 3.2.1 Reasoning data construction

To support LLMs for more effective reasoning, we need more textual information in addition to item titles. Given that the raw metadata often contains excessive or noisy content, directly using long-form

text may hinder the understanding and inference of LLMs. To address this, we leverage a pre-trained summarization model (Wei et al., 2023) to extract key concepts from the item metadata, and retain up to 10 representative keywords to enhance the semantic representation of each item. To model more comprehensive preference information, not just liking information, we add low-scoring interactions to the history sequence as well. Then we augment each item with title  $i_t$ , label  $i_l \in \{yes, no\}$ , feature, and keywords  $i_k$  as follows:

$$i_{txt} = i_t \text{ with feature } \text{emb}_s^i \text{ (evaluation: } i_l) \text{ with description: } i_k.$$

Since existing recommendation datasets lack explicit reasoning traces, we synthesize a subset of samples with explainable rationale by leveraging a powerful reasoning model QwQ<sup>2</sup>. Specifically, we sample a few thousand instances from the training data and construct textual prompts according to Appendix C.2. We then repeatedly query the reasoning model until a correct prediction is produced and record the latest explanation content as the reason for interaction as in Figure 2(b).

#### 3.2.2 Thinking activation

To equip the recommendation model with explicit reasoning capability, we introduce a thinking acti-

<sup>2</sup><https://qwenlm.github.io/blog/qwq-32b/>



variation mechanism that jointly trains the model on both user-item interactions and synthesized reasoning examples. During training, we perform mixed sampling from conventional recommendation data with binary labels, and reasoning data generated by prompting a reasoning model to produce natural language explanations for user preferences. For each sample, we concatenate the input prompt and its corresponding answer ( $i_i$  or  $r_{u,t}$ ) to form a language modeling instance as input for training:

$$\begin{aligned} \mathbf{E}^{qa} &= \text{Concat}(\mathbf{E}^q, \mathbf{E}^a), \\ pos &= -\text{Length}(\mathbf{E}^a), \end{aligned} \quad (4)$$

where  $\mathbf{E}^q, \mathbf{E}^a$  means embedding of question and answer constructing from  $\mathcal{S}$ , respectively. While  $pos$  marked the position of the first answer token. This mixed sampling setup encourages the model to learn both prediction accuracy and semantically aligned reasoning behavior, within a shared representation space as shown in Figure 2(c).

The overall objective of training optimization consists of two loss components. For thinking instances, the standard token-level cross-entropy loss over the input sequence is used:

$$\mathcal{L}_{think, logits} = \text{LLM}_\theta(\mathbf{E}^{qa}). \quad (5)$$

For recommendation instances, we use the standard binary cross-entropy (BCE) loss between the predicted relevance score  $\hat{l}$  and the observed ground-truth label  $l$  of  $y$ :

$$\begin{aligned} posid &= \text{TKZ}('Yes'), \\ \hat{l} &= \text{logits}[pos][posid], \\ \mathcal{L}_{rec} &= \text{BCE}(\hat{l}, l). \end{aligned} \quad (6)$$

To jointly optimize both objectives on a unified batch, we always use two losses, adjusting the weights of each loss according to instances:

$$\mathcal{L} = \begin{cases} \alpha\mathcal{L}_{rec} + \beta\mathcal{L}_{think}, & \text{thinking} \\ \eta\mathcal{L}_{rec} + \gamma\mathcal{L}_{think}, & \text{recommend} \end{cases}. \quad (7)$$

$\alpha, \beta, \eta, \gamma$  are weights for each loss. The model learns not only to rank relevant items but also to generate grounded reasons for its predictions, effectively activating the model's thinking ability.

### 3.3 Recommendation Experts Fusion

#### 3.3.1 Base Expert fine-tuning

To enable each expert to better capture user-specific preferences and behavioral patterns, we partition users into distinct groups based on latent features.

Specifically, we utilize user embeddings derived from various pretrained small collaborative models, each of which implicitly encodes the interaction semantics of users. These embeddings serve as the basis for grouping. We aggregate the representations across all users and perform unsupervised clustering to obtain  $N$  groups  $\mathcal{S}' = \{\mathcal{S}_{1:N}\}$ . The resulting clusters are then used to partition the training and validating data, allowing each expert to specialize in a subset of users with similar representations to simplify the preference modeling.

We first fine-tune a global expert using the full data under the framework described in Section 3.2.2, applying LoRA-based adaptation to enable generalizable recommendations with thinking. Building on this global expert  $\text{LoRA}_{global}$ , we further adapt selected LoRA layers (the last 8 layers) using grouped data, enabling the model to preserve general thinking ability while adapting to more nuanced user preferences. As a result, we obtain a set of base experts  $\{\text{LoRA}_{1:N}\}$ , which serve as candidates for dynamic expert fusion.

#### 3.3.2 Instance Wise Expert Fusion

To determine which expert is most suitable for a given user, we assign the mean of user features extracted within each group by the corresponding small model as representations of experts  $\mathcal{E} = \{\mathbf{e}_{1:N}^c\}$ . We then compute the match between user features and expert representations to estimate each base expert's involvement in modeling the user's preferences. The cosine similarity and softmax functions were used to obtain participation scores  $\mathbf{w}^u = \{w_{1:N}^u\}$  of user  $u$  based on experts:

$$\begin{aligned} \mathbf{e}_n^c &= \text{Mean}(\mathbf{e}_s^u), u \in \mathcal{S}_n, \\ \mathbf{w}^u &= \text{Softmax}(\text{Cosim}(\mathbf{e}_s^u, \mathcal{E})/\tau), \end{aligned} \quad (8)$$

where  $\tau$  is the temperature coefficient. We introduce a gating mechanism to filter users with highly averaged  $H(\mathbf{w}^u) > 0.95 \log N$  or concentrated  $\max(\mathbf{w}^u) > 0.5 + 0.6/N$  preference profiles, assigning them directly to a global or base expert; the remaining users are served by instance-wise fusion over multiple experts as shown in Figure 2(d). The threshold is calculated as follows:

$$H(\mathbf{w}^u) = - \sum_{n=1}^N w_n^u \log w_n^u, \quad (9)$$

$$\text{LoRA}^u = \begin{cases} \text{LoRA}_{global}, & \text{averaged} \\ \text{LoRA}_{argmax(\mathbf{w}^u)}, & \text{concentrated} \\ \sum_{n=1}^N w_n^u \text{LoRA}_n, & \text{otherwise} \end{cases}. \quad (10)$$

## 4 Experiments

We conduct experiments on real-world datasets to answer three main research questions: **RQ1:** How does ThinkRec perform in comparison to existing recommendation methods? **RQ2:** Why is the thinking activation method in ThinkRec essential? **RQ3:** How does the fusion of experts influence recommendation performance?

### 4.1 Experimental Setup

**Datasets.** We conduct experiments on three datasets: *ML1M* refers to MovieLens-1M<sup>3</sup>. *Yelp* refers to the Yelp Open dataset<sup>4</sup>. *Book* refers to the “book” subset of the Amazon Product Review dataset<sup>5</sup>. The data processing details are in the Appendix A.1, and the statistical information on the processed datasets is in Table 1.  $\sigma$ ,  $\mu$  refer to the standard deviation and mean of user interactions.

**Evaluation Metrics.** We employ four widely used recommendation metrics: *AUC*, *UAUC*, *Normalized Discounted Cumulative Gain (NDCG)*, and *Mean Average Precision (MAP)*. And we employ *METEOR* (Banerjee and Lavie, 2005) and *BLEURT* (Sellam et al., 2020) to measure the generated reasons. *METEOR* incorporates synonym matching and word order. *BLEURT* is fine-tuned on human judgments to directly predict text quality.

**Baselines.** To fully evaluate the proposed method ThinkRec, we compare it with traditional collaborative methods MF (Koren et al., 2009), LightGCN (He et al., 2020), SASRec (Kang and McAuley, 2018), and LLM-based methods Prompt4NR (Zhang and Wang, 2023), TALLRec (Bao et al., 2023), CoLLM (Zhang et al., 2025). We extend the above LLM-based method to LLM Llama3-8B<sup>6</sup> for a fair comparison and tune the LLM with LoRA to manage computational costs. The complete introduction is in Appendix A.3.

**Implementation Details** We implement all the compared methods using PyTorch 2.5. We adopt BCE when not otherwise specified. LLM-based methods are optimized using AdamW, while other models use Adam. For hyperparameters such as learning rate, embedding size, and weight decay of traditional methods, we conduct grid search over commonly used ranges, and follow the original papers for other baseline-specific settings. For LLM-

Datasets	#Train	#Valid	#Test	#User	#Item	$\sigma/\mu$
ML1M	33,891	10,401	7,331	5,945	3,687	1.48
Yelp	1,637,168	144,929	144,929	40,617	60,014	1.24
Book	650,865	56,262	56,262	22,686	47,059	1.07

Table 1: Statistics of Datasets.

based methods, we adopt CoLLM’s configuration for optimizer, sequence truncation and LoRA settings. We set the sampling ratio of reasoning data to 0.2, and full hyperparameter settings and search ranges are provided in the Appendix A.2.

### 4.2 Overall Performance (RQ1)

We compare our ThinkRec method with baseline methods on the accuracy of recommendations and the quality of generated reasons.

#### 4.2.1 Accuracy of Recommendation

The results of ThinkRec and SOTA recommendations are in Table 2. We observe some improvement in our method on all three datasets. Notably, ThinkRec improves the previous SOTA CoLLM by +.0582 (9.13%) on AUC in Yelp and +.0623 (8.72%) on AUC in ML1M, demonstrating substantial gains in global ranking accuracy. CoLLM, which integrates collaborative signals into LLMs using external traditional models, achieves the second-best AUC on ML1M (0.7141) and Book (0.7830), confirming its effectiveness in leveraging user-item interaction patterns. However, its performance in user-level ranking metrics is relatively less competitive, especially on ML1M and Yelp, where ThinkRec offers more personalized modeling through instance-wise expert fusion.

Other LLM-based methods, such as Prompt4NR and TALLRec, also demonstrate competitive performance, particularly on datasets with rich textual item metadata such as Book. For instance, TALLRec achieves the highest NDCG@5 (0.7683) on ML1M, and both TALLRec and Prompt4NR slightly outperform ThinkRec in MAP@5 on Book. These results suggest that instruction tuning and prompt engineering are beneficial when item descriptions provide substantial semantic context.

Turning to traditional recommendation models, SASRec clearly outperforms MF and LightGCN, especially on ML1M and Yelp. Its self-attention-based sequential modeling effectively captures temporal patterns in user behavior, yielding the best UAUC (0.6687) and MAP@5 (0.4747) among non-LLM baselines on ML1M. However, its performance drops significantly on Book, where user

<sup>3</sup><https://grouplens.org/datasets/movielens/1m/>

<sup>4</sup><https://business.yelp.com/data/resources/open-dataset/>

<sup>5</sup><https://nijianmo.github.io/amazon/index.html>

<sup>6</sup>[huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct](https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct)

Datasets	ML1M				Yelp				Book			
Methods	AUC	UAUC	NDCG@5	MAP@5	AUC	UAUC	NDCG@5	MAP@5	AUC	UAUC	NDCG@5	MAP@5
MF	0.6401	0.6079	0.7286	0.4520	0.5838	0.5389	0.8120	0.2552	0.6592	0.5527	0.6805	0.2887
LightGCN	0.6140	0.6230	0.7333	0.4600	0.5360	0.5179	0.8076	0.2520	0.5622	0.4985	0.6406	0.2598
SASRec	0.6956	<u>0.6687</u>	0.7663	<u>0.4747</u>	0.6184	<b>0.6096</b>	<u>0.8564</u>	<u>0.2785</u>	0.5411	0.5197	0.6550	0.2701
Prompt4NR	0.6936	0.6433	0.7665	0.4652	0.6272	0.6034	0.8348	0.2705	0.6764	<u>0.5699</u>	<b>0.7023</b>	<u>0.3048</u>
TALLRec	0.6872	0.6553	<u>0.7683</u>	0.4706	0.5334	0.5206	0.7988	0.2538	0.6632	0.5568	<u>0.7023</u>	<b>0.3049</b>
CoLLM	<u>0.7141</u>	0.6672	0.7585	0.4647	<u>0.6373</u>	0.5961	0.8420	0.2734	<u>0.7830</u>	0.5672	0.6917	0.2968
Ours	<b>0.7764</b>	<b>0.6775</b>	<b>0.7747</b>	<b>0.4774</b>	<b>0.6955</b>	<u>0.6065</u>	<b>0.8585</b>	<b>0.2826</b>	<b>0.8302</b>	<b>0.5705</b>	0.6858	0.2977

Table 2: Comparison of prediction performance between ThinkRec and the baselines across the three evaluation datasets. The best results are highlighted in **bold** and sub-optimal results are underlined.

interactions are more sparse and less sequential, revealing the limitations of sequence-only models.

Overall, ThinkRec delivers the most balanced and robust performance across all datasets and evaluation metrics. It achieves the best performance of almost all metrics on ML1M and Yelp, and also obtains the top AUC and UAUC on Book. Its consistent top-tier results confirm the effectiveness of combining thinking activation and expert fusion. These components jointly enhance both global and user-specific ranking quality, making ThinkRec a scalable and interpretable recommendation.

Datasets	ML1M		Yelp		Book	
Methods	M	B	M	B	M	B
Prompt4NR	0.0010	0.2013	0.0205	0.1675	0.0003	0.1957
TALLRec	<u>0.0275</u>	<u>0.2607</u>	<u>0.0379</u>	<u>0.2420</u>	<u>0.0301</u>	<u>0.1931</u>
CoLLM	0.0003	0.1626	0.0001	0.1785	0.0097	0.1636
Ours	<b>0.0333</b>	<b>0.3104</b>	<b>0.0616</b>	<b>0.2683</b>	<b>0.0546</b>	<b>0.2828</b>

Table 3: Quality evaluation of generated reasons. "M" refers to "METEOR" and "B" refers to "BLEURT".

#### 4.2.2 Quality of Generated Reasons

Table 3 summarizes the quality of the reasons generated by the LLM-based recommenders. Compared with the generated reasons from QwQ using the METEOR and BLEURT metrics, the reasons generated by our method significantly outperform those of the three LLM-based baselines. Our method achieves an average relative improvement of 56.54% on METEOR and 23.35% on BLEURT across all datasets, suggesting better fluency, coherence, and semantic relevance. These results validate the effectiveness of our thinking activation mechanism, which explicitly aligns recommendations with user-centric reasoning via joint training on reasoning-augmented samples. The improvement in both syntactic and learned semantic metrics confirms that ThinkRec not only provides accurate recommendations but also produces more coherent, grounded, and human-aligned explanations, a crucial step toward reasonable and trustworthy LLM-based recommendations.

Datasets	ML1M			Yelp			Book		
Methods	UAUC	N@5	M@5	UAUC	N@5	M@5	UAUC	N@5	M@5
w/o both	0.6658	0.7643	0.4674	0.5904	0.8429	0.2736	0.5017	0.6381	0.2613
w/o think	0.6599	0.7570	0.4623	0.5865	0.8402	0.2702	0.4692	0.6284	0.2548
w/o experts	0.6765	0.7740	0.4742	0.5999	0.8562	0.2791	0.5631	0.6801	0.2939
Ours	<b>0.6775</b>	<b>0.7747</b>	<b>0.4774</b>	<b>0.6065</b>	<b>0.8585</b>	<b>0.2826</b>	<b>0.5705</b>	<b>0.6858</b>	<b>0.2977</b>

Table 4: Ablation studies of key components in ThinkRec. "N" refers to "NDCG", "M" refers to "MAP".

### 4.3 In-depth Analysis

#### 4.3.1 Ablation Studies (RQ2)

To evaluate the importance of explicit reasoning in recommendation, we ablate the "thinking" component of ThinkRec (w/o think), which disables reasoning supervision. As shown in Table 4, this leads to significant performance degradation across all datasets. For example, UAUC on the Book dataset drops from 0.5705 to 0.4692. Interestingly, this is even lower than the variant where both thinking and expert mechanisms are removed (w/o both). Without thinking, the recommendation task is effectively reduced to binary classification, where enabling multi-expert modeling may overfit shallow interactions, thus compromising generalization.

We then assess the contribution of the expert personalization module, which removes the latent-feature-based user grouping and experts fusion mechanism (w/o experts). As shown in Table 4, this also leads to consistent performance drops—for instance, MAP@5 on Yelp falls from 0.2826 to 0.2791. Notably, only when thinking is enabled does multi-expert modeling begin to show substantial benefits. With reasoning supervision, group-specific LoRA modules can effectively specialize in distinct user groups, capturing fine-grained preference signals that would otherwise be blurred in the global model. These findings highlight the consistency and complementarity between thinking and multi-expert modeling, providing a semantically rich space that allows user grouping to generalize rather than overfit, enabling expert models to move beyond surface interaction patterns and

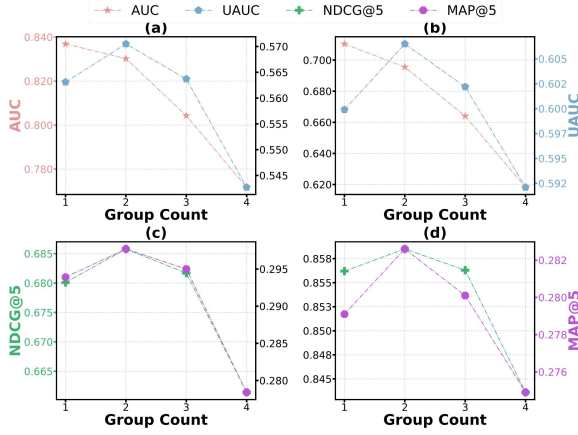


Figure 3: The influence of performance with the number of experts on Book (left panel) and Yelp (right panel).

capture deeper preference semantics.

### 4.3.2 Study on the Fusion of Experts (RQ3)

*Analysis of the number of experts.* As shown in Figure 3, the number of expert groups increases from 1 to 4, the model exhibits a characteristic 'rise-then-fall' performance trend, revealing the trade-off between personalization capacity and generalization. In the early stages, fine-tuning LoRA modules within user groups significantly enhances the model's ability to capture diverse preferences, resulting in notable gains in user-level and Top-N ranking metrics such as UAUC, NDCG@5, and MAP. However, with further partitioning, each subgroup receives fewer training samples, making the model prone to overfitting, thereby degrading ranking performance. Notably, the AUC metric consistently decreases with more experts, reflecting the deterioration of the consistency of global representation with expert specialization, validating the inherent tension between "global consistency" and "local specificity" in recommender systems. These results indicate that more experts do not necessarily equate to better performance; instead, the optimal group number should be dynamically adjusted according to user behavior diversity, and the frequency of interactions.

*Analysis of the grouping features.* Under a fixed two-group setting, we further investigate how the choice of user grouping features affects model performance. Specifically, we employ user embeddings generated by MF, LightGCN, and SASRec to construct different grouping strategies. As shown in Figure 4, with the increase of group difference (Cohen's d), the performance of ThinkRec consistently improves. This trend highlights the

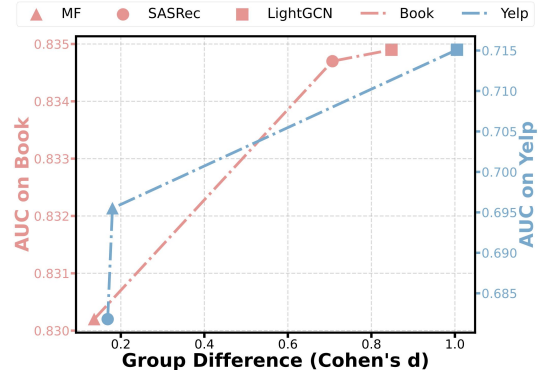


Figure 4: The influence of performance on the accuracy of grouping (Cohen's d of grouped datasets).

importance of semantic decoupling among expert groups. When user preferences across groups exhibit stronger heterogeneity, the LoRA modules assigned to each group can learn more complementary preference representations, thereby enhancing the system's modeling capacity and global discriminative power. In contrast, if the groups are semantically overlapping, expert models tend to learn redundant representations, and the fused output becomes increasingly similar to that of the global model, limiting performance gains. Therefore, leveraging high-quality user behavior modeling methods as the basis for grouping can amplify divergence across user groups, enabling multi-expert systems to achieve better personalized expressiveness while preserving a global perspective.

## 5 Conclusion

In this work, we explore how to activate the explicit reasoning capabilities of LLMs for recommendation tasks, aiming to improve both interpretability and accuracy. We highlight crucial challenges in existing LLM-based recommendation approaches, including the absence of reasoning supervision and insufficient modeling of fine-grained user preferences. To this end, we propose the ThinkRec, which incorporates a thinking activation framework containing item augmentation and mixed sampling training, an instance-wise expert fusion mechanism based on latent user features. The results in three real-world datasets confirm that ThinkRec performs better in both the accuracy of recommendations and the quality of explanations. We believe this work lays the foundation for a new paradigm that bridges reasoning and recommendation in LLMs. In the future, we plan to extend this line of research toward thinking-based item generation and review-informed user modeling.



## Limitations

ThinkRec has several limitations. First, the reasoning supervision relies on synthetic explanation data generated by a separate LLM, which may introduce noise or reinforce undesired patterns. Second, our expert modeling approach is currently based on offline clustering and may not capture evolving user preferences over time. In addition, our evaluation does not include real-world user engagement feedback. We leave these for future work to address.

## Ethical Considerations

We use only publicly available benchmark datasets (e.g., ML1M, Yelp, Book) that are widely adopted in the recommendation community and contain no personally identifiable information (PII) in our work. We take initial steps to mitigate oversimplified reasons and hallucinated content in the textual outputs via reasoning-aligned training.

## References

Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

Keqin Bao, Jizhi Zhang, Yang Zhang, Xinyue Huo, Chong Chen, and Fuli Feng. 2024. [Decoding matters: Addressing amplification bias and homogeneity issues in recommendations for large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10540–10552, Miami, Florida, USA. Association for Computational Linguistics.

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, RecSys '23*, page 1007–1014, New York, NY, USA. Association for Computing Machinery.

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, RecSys '22*, page 299–315, New York, NY, USA. Association for Computing Machinery.

Shijie Geng, Juntao Tan, Shuchang Liu, Zuohui Fu, and Yongfeng Zhang. 2023. [VIP5: Towards multi-modal foundation models for recommendation](#). In

*Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9606–9620, Singapore. Association for Computational Linguistics.

Maria Glenski, William I. Sealy, Kate Miller, and Dustin Arendt. 2021. [Improving synonym recommendation using sentence context](#). In *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*, pages 74–78, Virtual. Association for Computational Linguistics.

Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yong-Dong 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, SIGIR '20*, page 639–648, New York, NY, USA. Association for Computing Machinery.

Ngai Lam Ho, Roy Ka-Wei Lee, and Kwan Hui Lim. 2023. [Btrec: Bert-based trajectory recommendation for personalized tours](#). *Preprint*, arXiv:2310.19886.

Andreea Iana, Goran Glavaš, and Heiko Paulheim. 2024. [Train once, use flexibly: A modular framework for multi-aspect neural news recommendation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9555–9571, Miami, Florida, USA. Association for Computational Linguistics.

Jianchao Ji, Zelong Li, Shuyuan Xu, Wenyue Hua, Yingqiang Ge, Juntao Tan, and Yongfeng Zhang. 2023a. [Genrec: Large language model for generative recommendation](#). *Preprint*, arXiv:2307.00457.

Yitong Ji, Aixin Sun, Jie Zhang, and Chenliang Li. 2023b. [A critical study on data leakage in recommender system offline evaluation](#). *ACM Trans. Inf. Syst.*, 41(3).

Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

Wang-Cheng Kang and Julian McAuley. 2018. [Self-attentive sequential recommendation](#). In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 197–206.

Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. [Matrix factorization techniques for recommender systems](#). *Computer*, 42(8):30–37.

Lei Li, Yongfeng Zhang, and Li Chen. 2023. [Prompt distillation for efficient llm-based recommendation](#). In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM '23*, page 1348–1357, New York, NY, USA. Association for Computing Machinery.

Jianghao Lin, Rong Shan, Chenxu Zhu, Kounianhua Du, Bo Chen, Shigang Quan, Ruiming Tang, Yong Yu, and Weinan Zhang. 2024a. [Rella: Retrieval-enhanced large language models for lifelong sequential behavior comprehension in recommendation](#). In *Proceedings of the ACM Web Conference 2024*,

713	WWW '24, page 3497–3508, New York, NY, USA.	770
714	Association for Computing Machinery.	771
715	Xinyu Lin, Wenjie Wang, Yongqi Li, Fuli Feng, See-	772
716	Kiong Ng, and Tat-Seng Chua. 2024b. <a href="#">Bridging</a>	773
717	<a href="#">items and language: A transition paradigm for large</a>	774
718	<a href="#">language model-based recommendation</a> . <i>Preprint</i> ,	775
719	arXiv:2310.06491.	776
720	Han Liu, Xianfeng Tang, Tianlang Chen, Jiapeng Liu,	777
721	Indu Indu, Henry Peng Zou, Peng Dai, Roberto Fer-	778
722	nandez Galan, Michael D Porter, Dongmei Jia, Ning	
723	Zhang, and Lian Xiong. 2024. <a href="#">Sequential LLM</a>	779
724	<a href="#">framework for fashion recommendation</a> . In <i>Proceed-</i>	780
725	<i>ings of the 2024 Conference on Empirical Methods</i>	781
726	<i>in Natural Language Processing: Industry Track</i> ,	782
727	pages 1276–1285, Miami, Florida, US. Association	783
728	for Computational Linguistics.	784
729	Jiahao Liu, Xueshuo Yan, Dongsheng Li, Guangping	785
730	Zhang, Hansu Gu, Peng Zhang, Tun Lu, Li Shang,	786
731	and Ning Gu. 2025. <a href="#">Improving llm-powered recom-</a>	787
732	<a href="#">mendations with personalized information</a> . <i>Preprint</i> ,	788
733	arXiv:2502.13845.	789
734	Liangchen Luo, Yinxiao Liu, Rosanne Liu, Samrat	
735	Phatale, Meiqi Guo, Harsh Lara, Yunxuan Li, Lei	790
736	Shu, Yun Zhu, Lei Meng, Jiao Sun, and Abhinav	791
737	Rastogi. 2024. <a href="#">Improve mathematical reasoning in</a>	792
738	<a href="#">language models by automated process supervision</a> .	793
739	<i>Preprint</i> , arXiv:2406.06592.	794
740	Guangyuan Ma, Hongtao Liu, Xing W, Wanhui Qian,	795
741	Zhepeng Lv, Qing Yang, and Songlin Hu. 2023.	796
742	<a href="#">PUNR: Pre-training with user behavior modeling</a>	797
743	<a href="#">for news recommendation</a> . In <i>Findings of the As-</i>	798
744	<i>sociation for Computational Linguistics: EMNLP</i>	799
745	2023, pages 8338–8347, Singapore. Association for	800
746	Computational Linguistics.	
747	Joshua Park and Yongfeng Zhang. 2025. <a href="#">Agen-</a>	801
748	<a href="#">trec: Agent recommendation using sentence em-</a>	802
749	<a href="#">beddings aligned to human feedback</a> . <i>Preprint</i> ,	803
750	arXiv:2501.13333.	804
751	Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020.	
752	<a href="#">BLEURT: Learning robust metrics for text genera-</a>	805
753	<a href="#">tion</a> . In <i>Proceedings of the 58th Annual Meeting of</i>	806
754	<i>the Association for Computational Linguistics</i> , pages	807
755	7881–7892, Online. Association for Computational	808
756	Linguistics.	809
757	Yancheng Wang, Ziyang Jiang, Zheng Chen, Fan Yang,	810
758	Yingxue Zhou, Eunah Cho, Xing Fan, Yanbin Lu,	811
759	Xiaojiang Huang, and Yingzhen Yang. 2024a. <a href="#">Rec-</a>	
760	<a href="#">mind: Large language model powered agent for</a>	812
761	<a href="#">recommendation</a> . In <i>Findings of the Association</i>	813
762	<i>for Computational Linguistics: NAACL 2024</i> , pages	814
763	4351–4364, Mexico City, Mexico. Association for	815
764	Computational Linguistics.	816
765	Zhefan Wang, Yuanqing Yu, Wendi Zheng, Weizhi Ma,	
766	and Min Zhang. 2024b. <a href="#">Macrec: A multi-agent col-</a>	817
767	<a href="#">laboration framework for recommendation</a> . In <i>Pro-</i>	818
768	<i>ceedings of the 47th International ACM SIGIR Con-</i>	819
769	<i>ference on Research and Development in Information</i>	820
	<i>Retrieval</i> , SIGIR '24, page 2760–2764, New York,	821
	NY, USA. Association for Computing Machinery.	822
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	823
	Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le,	824
	and Denny Zhou. 2022. <a href="#">Chain-of-thought prompt-</a>	825
	<a href="#">ing elicits reasoning in large language models</a> . In	
	<i>Proceedings of the 36th International Conference on</i>	
	<i>Neural Information Processing Systems</i> , NIPS '22,	
	Red Hook, NY, USA. Curran Associates Inc.	
	Xiangpeng Wei, Haoran Wei, Huan Lin, Tianhao Li, Pei	
	Zhang, Xingzhang Ren, Mei Li, Yu Wan, Zhiwei Cao,	
	Binbin Xie, Tianxiang Hu, Shangjie Li, Binyuan Hui,	
	Bowen Yu, Dayiheng Liu, Baosong Yang, Fei Huang,	
	and Jun Xie. 2023. <a href="#">Polylm: An open source polyglot</a>	
	<a href="#">large language model</a> . <i>Preprint</i> , arXiv:2307.06018.	
	Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng,	
	Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin	
	Jiang. 2023. <a href="#">Wizardlm: Empowering large language</a>	
	<a href="#">models to follow complex instructions</a> . <i>Preprint</i> ,	
	arXiv:2304.12244.	
	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran,	
	Thomas L. Griffiths, Yuan Cao, and Karthik	
	Narasimhan. 2023. <a href="#">Tree of thoughts: deliberate prob-</a>	
	<a href="#">lem solving with large language models</a> . In <i>Pro-</i>	
	<i>ceedings of the 37th International Conference on</i>	
	<i>Neural Information Processing Systems</i> , NIPS '23,	
	Red Hook, NY, USA. Curran Associates Inc.	
	Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak	
	Shafran, Karthik Narasimhan, and Yuan Cao. 2022.	
	<a href="#">React: Synergizing reasoning and acting in language</a>	
	<a href="#">models</a> . <i>arXiv preprint arXiv:2210.03629</i> .	
	Dan Zhang, Sining Zhou, Ziniu Hu, Yisong Yue,	
	Yuxiao Dong, and Jie Tang. 2024a. <a href="#">Rest-mcts*: Llm</a>	
	<a href="#">self-training via process reward guided tree search</a> .	
	<i>arXiv preprint arXiv:2406.03816</i> .	
	Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun,	
	Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-	
	Rong Wen. 2024b. <a href="#">Agentcf: Collaborative learning</a>	
	<a href="#">with autonomous language agents for recommender</a>	
	<a href="#">systems</a> . In <i>Proceedings of the ACM Web Conference</i>	
	2024, WWW '24, page 3679–3689, New York, NY,	
	USA. Association for Computing Machinery.	
	Yang Zhang, Fuli Feng, Jizhi Zhang, Keqin Bao, Qifan	
	Wang, and Xiangnan He. 2025. <a href="#">Collm: Integrating</a>	
	<a href="#">collaborative embeddings into large language models</a>	
	<a href="#">for recommendation</a> . <i>IEEE Transactions on Knowl-</i>	
	<i>edge and Data Engineering</i> , 37(5):2329–2340.	
	Zijian Zhang, Shuchang Liu, Ziru Liu, Rui Zhong,	
	Qingpeng Cai, Xiangyu Zhao, Chunxu Zhang, Qi-	
	dong Liu, and Peng Jiang. 2024c. <a href="#">Llm-powered</a>	
	<a href="#">user simulator for recommender system</a> . <i>Preprint</i> ,	
	arXiv:2412.16984.	
	Zizhuo Zhang and Bang Wang. 2023. <a href="#">Prompt learn-</a>	
	<a href="#">ing for news recommendation</a> . In <i>Proceedings of</i>	
	<i>the 46th International ACM SIGIR Conference on</i>	
	<i>Research and Development in Information Retrieval</i> ,	

## A More Settings of Experiments

### A.1 Datasets Details

**ML1M** contains user ratings on movies, collected between 2000 and 2003, with ratings on a scale of 1 to 5. We convert these ratings into binary labels using a threshold of 3. **Yelp** includes user reviews, ratings for businesses such as restaurants and retail shops, as well as textual information about the businesses. We convert these ratings into binary labels using a threshold of 3. **Book** compiles user reviews of books from Amazon, collected between 1996 and 2018, with review scores ranging from 1 to 5. We transform these review scores into binary labels using a threshold of 4. Ratings greater than the threshold are labeled as “positive” ( $y = 1$ ), while the rest are labeled as “negative” ( $y = 0$ ).

To better simulate real-world recommendation scenarios and prevent data leakage (Ji et al., 2023b), we split the dataset into training, validation, and testing sets based on the interaction timestamp. Specifically, for Amazon-Book, given its large scale, we just preserve interactions from the year 2017, allocating the first 10 months for training, and the remaining 2 months for validation and testing. As for Yelp, we preserve interactions from 2010 to 2022, allocating the first 10 years for training, and the remaining 2 years for validation and testing. Given the sparse nature of Book and Yelp, we filtered out users and items with fewer than 20 interactions. As for ML1M, we preserve the interactions from the most recent twenty months, using the first 10 months for training, the remaining 10 months for validation and testing.

### A.2 Implementation Details

Regarding traditional models, we explore the learning rate within the range of  $[1e-2, 1e-3, 1e-4]$ , the embedding size within the range of  $[64, 128, 256]$ , and weight decay within the range of  $[1e-2, 1e-3, \dots, 1e-7]$ . As for LLM-based methods, we set the learning rate to  $1e-4$ , and weight decay to  $1e-3$  to align with CoLLM. For SASRec, we establish the maximum length of historical interaction sequences to 25. We adopt TALLRec’s practice of setting the maximum sequence length to 10 for all other methods. Regarding other specific parameters of the baseline models, we adhere to the configurations outlined in their

original papers. For the LoRA module, we follow the same configuration as CoLLM, setting  $r, \alpha, \text{dropout}, \text{target modules}$  to 8, 16, 0.05, and “[q proj, v proj]”, respectively. We set the sample rate of reasoning data and recommendation data to 0.2 and 0.8, and the weight of each loss function as follows:  $\alpha = 0.1, \beta = 0.9, \eta = 0.9, \gamma = 0.1$ . We set the temperature coefficient  $\tau$  to 0.1.

### A.3 Baselines Details

To fully evaluate the proposed method ThinkRec, we compare it with:

- **MF** (Koren et al., 2009) This refers to Matrix Factorization, a representative latent factor-based collaborative filtering method.
- **LightGCN** (He et al., 2020) A representative graph-based collaborative filtering method, which uses a simplified graph convolutional neural network to enhance the modeling of user interest.
- **SASRec** (Kang and McAuley, 2018) A representative sequential recommendation method, which uses self-attention to encode sequential patterns to model user interest.
- **Prompt4NR** (Zhang and Wang, 2023) It uses both fixed and soft prompts to utilize traditional Language Models for recommendation.
- **TALLRec** (Bao et al., 2023) This is a state-of-the-art LLMRec method that aligns LLM with recommendations through aligns LLM with recommendations through instruction tuning.
- **CoLLM** (Zhang et al., 2025) It effectively integrates collaborative information into LLMs by harnessing the capability of external traditional models to capture the information.

## B Pseudocode of ThinkRec

We briefly introduce ThinkRec in Algorithm 1 and 2. First, as shown in Algorithm 1, items are augmented and reasons are generated. Then global and base experts are trained on according datasets. Second, as shown in Algorithm 2, the participation of each expert is calculated according to their representations and the user’s latent feature. Then, user-specific experts are specified according to participation and thresholds.



---

**Algorithm 1: Thinking Activation**

---

**Module 1:** ▷ *Keywords Extraction*

**Target:** Metadata of each item  $i_d \mapsto$  Description of each item  $i_k$

**Input:** Metadata of each item  $i_d$

**Output:** Description of each item  $i_k$

**Module 2:** ▷ *Reasoning Data Generation*

**Target:** Interaction dataset  $\mathcal{S} \mapsto$  Reasoning dataset  $\mathcal{S}_r$

**Input:** Interaction dataset  $\mathcal{S}$

**Output:** Reasoning dataset  $\mathcal{S}_r$

**Module 3:** ▷ *Expert Training*

**Target:** Backbone model  $\text{LLM}(\cdot)$ , collaborative encoder  $f_\psi(\cdot; \mathcal{S})$ , dataset  $\mathcal{S}, \mathcal{S}_r \mapsto$  Global expert  $\text{LoRA}_{\text{global}}$ , base experts  $\{\text{LoRA}_{1:N}\}$

**Input:** Backbone model  $\text{LLM}(\cdot)$ , collaborative encoder  $f_\psi(\cdot; \mathcal{S})$ , dataset  $\mathcal{S}, \mathcal{S}_r$

**Output:** Global expert  $\text{LoRA}_{\text{global}}$ , base experts  $\{\text{LoRA}_{1:N}\}$

---

---

**Algorithm 2: Instance-wise Experts Fusion**

---

**Module 1:** ▷ *Participation Calculation*

**Target:** User embedding  $\mathbf{e}_s^u$ , representations of experts  $\{\mathcal{E}_{1:N}\} \mapsto$  Participation of experts  $\mathbf{w}^u$

**Input:** User embedding  $\mathbf{e}_s^u$ , representations of experts  $\{\mathcal{E}_{1:N}\}$

**Output:** Participation of experts  $\mathbf{w}^u$

**Module 2:** ▷ *Expert Fusion*

**Target:** Participation of experts  $\mathbf{w}^u$ , Global expert  $\text{LoRA}_{\text{global}}$ , base experts  $\{\text{LoRA}_{1:N}\} \mapsto$  Instance-wise Expert  $\text{LoRA}^u$

**Input:** Participation of experts  $\mathbf{w}^u$ , Global expert  $\text{LoRA}_{\text{global}}$ , base experts  $\{\text{LoRA}_{1:N}\}$

**Output:** Instance-wise Expert  $\text{LoRA}^u$

---

## C Prompt Templates

### C.1 Prompt for Summarizing Metadata

**Summarization**

Your task is to extract up to ten keywords from the given profile of the book above. Answers should contain only keywords, which should be separated by commas.

Keywords:

### C.2 Prompt for Reasoning Data Synthesis

**First Turn**

A user has given high ratings to the following books:  $\langle \text{HisItemList} \rangle$ . Using all available information, make a prediction about whether the user would enjoy the book titled  $\langle \text{TargetItemTitle} \rangle$ ?

**Reflect Turn**

▷ The correct response is  $\langle \text{answer} \rangle$ . Reflect on multiple aspects based on historical information and explain the reason for the oversight based on the previous analysis. Reanalyze to make a prediction about whether the user would enjoy the book titled  $\langle \text{TargetItemTitle} \rangle$ ?

▷ The accurate answer is  $\langle \text{answer} \rangle$ . Delve into various aspects considering historical data, elucidate the cause of the oversight according to the preceding analysis. Conduct a reanalysis to forecast whether the user will take pleasure in the book named  $\langle \text{TargetItemTitle} \rangle$ ?

▷ The right response is  $\langle \text{answer} \rangle$ . Reflect on a variety of aspects with reference to historical information, and account for the oversight based on the earlier analysis. Reanalyze to determine whether the user would appreciate the book titled  $\langle \text{TargetItemTitle} \rangle$ ?

### C.3 Prompt for Recommendation

**Recommendation**

#Question: A user has given ratings to the following books:  $\langle \text{HisItemList} \rangle$ . Additionally, we have information about the user's preferences encoded in the feature  $\langle \text{UserID} \rangle$ . Based on the descriptions and the user's enjoyment of each book in the historical sequence, construct a persona of the user's preferences and reevaluate whether the user would enjoy the book titled  $\langle \text{TargetItemTitle} \rangle$  with the feature  $\langle \text{TargetItemID} \rangle$ . Please begin your analysis with "Yes" or "No".

#Answer:



## D Case Study

Among existing LLM-based recommendations, CoLLM and Prompt4NR yield disorganized symbols. TALLRec frequently generates sentences with unrelated elements, such as code or hallucinated facts, failing to reflect coherent reasoning. In contrast, ThinkRec demonstrates structured, step-by-step reasoning aligned with user history and target item semantics, enabling it to produce accurate and interpretable recommendations.

### Prompt4NR

```
")) '""'; 3.}", ' '); }') , } "); }; "" "" " I think I understand the problem, but it's a bit unclear to me. ") }; } "" " );
```

### TALLRec

```
");" # 3. The user's response: "1"; (the number of the book to be analyzed) # - If a book has been rated, then it is included in this list. # * For each book that matches with any other book on the same list: # + Check if both books are novels and have similar genres # - In case they match these criteria, add them together as one single entry "" "" def predict_book_rating(book_id): for i in range(len(user_ratings)): if int(i) == book_id: return True result = [] for j in user_ratings: if len(result) > 0: flag = False for k in result[-1]: if j[book_name] == k[0][book_title]: if 'novel' in j['genres'] or 'fiction' or ": if ('romance', ") not in k: result.pop() break else: flag=True continue elif j.get('genres') != None and set(j['genre']) & set(k[2]): flag=False # #!/usr/bin/env python from flask import Flask, request, jsonify import json app = Flask(__name__) @app.route('/predict-book-rating') def predict(): data = {'user': ['The Duchess Deal: Girl Meets Duke'; 'The V Card';...]} print(data) if __name__ == '__main__': app.run(debug=True) "" Here you can use Python script. Please note that I am using 'flask' framework which allows us to create simple web applications. Here we will make an API endpoint 'predict-books-ratings'. You should send JSON object where key "books" contains array of book titles.
```

### CoLLM

```
"); #</#>'; ' ', ' '; '# '; "" " I have been trying to solve this problem for a long time, and I am still unable. Can you help me?
```

### ThinkRec

**Yes.** Okay, let's see here. The user has rated several books positively and I need to figure out if they'd like the new book *"If I Run (If i Run Series)"*. First, looking at their liked books: *They enjoy thrillers with elements of suspense, crime, mystery, family dynamics, abuse themes, and some historical or biographical aspects too*. For example, in "A Child Called It" there's childhood trauma, healing, which is a strong theme. Then series like Heist and Millennium have action, romance maybe? But mostly thriller/suspense. Even though it mentions romance in those, the main focus seems on the heists and cybercrime. The next one, "Saving Parker", includes American history and children's angles again. *So historical fiction blended with personal stories might be appealing*. Also, Jack Reacher novels are part of this list, so that suggests they like gritty detective work or military/authoritative figures. Michael Bennett and Alex Cross also indicate police procedural or investigative mysteries. Now the target book is an If I'm Found sequel, focusing on family, suspense/mystery, murder, love, marriage/divorce... If you run has plot twist, dialogue, compulsive reading, etc., which aligns with what users like from other books. Themes include family and suspense. Their existing likes include both family-focused (e.g., Saving Parker) along with intense plots (Heist, Stieg Larson). The presence of murder and investigation matches Haunted by Michael Bennet. Marriage/divorced themes aren't explicitly mentioned elsewhere, but relationships are present in others (Before We... has love story?). *So overall, given the overlap in key genre tags (family, thriller, suspension), plus positive reviews on related topics, I think the answer should be Yes.*

## E Notation

Symbol	Description
<b>Sets and Indices</b>	
$\mathcal{S}$	Sequential recommendation dataset
$\mathcal{S}' = \{\mathcal{S}_{1:N}\}$	Grouped Sequential recommendation dataset
$u, t$	User and time indices
$N$	Total number of users
$T_u$	Number of behaviors of user $u$
$\mathcal{E} = \{\mathbf{e}_{1:N}^c\}$	Representations of experts
<b>Variables and Hyperparameters</b>	
$x_{u,t}, y_{u,t}$	History and current behaviors of user $u$ at time $t$
$x, y$	Simplified notations for $(x_{u,t}, y_{u,t})$
$i_{id}, i_{txt}, i_t, i_d, i_k, i_l$	Item ID, textual information, title, description, keywords, yes/no label
$r_{u,t}$	Explanation or reason for recommendation at $(u, t)$
$\hat{l}, l$	Predicted and ground-truth labels (binary)
$\mathbf{e}_s^u, \mathbf{e}_s^i$	User/item embeddings from collaborative encoder
$\mathbf{E}_{txt}$	Token-level embedding of text
$\mathbf{emb}_{txt}^1$	First token embedding
$\mathbf{emb}_s^u, \mathbf{emb}_s^i$	User/item embeddings projected into language space
$\mathbf{E}^q, \mathbf{E}^a, \mathbf{E}^{qa}$	Embeddings of question, answer, and their concatenation
$\mathbf{w}^u = \{w_{1:N}^u\}$	Expert participation weights of user $u$
$d_1, d_2$	Dimensions of collaborative and language embeddings
$L$	Number of textual tokens
$pos, posid$	Position of answer token; token ID of “Yes”
$\tau$	Softmax temperature
$\alpha, \beta, \eta, \gamma$	Loss weighting coefficients
<b>Functions</b>	
$f_\psi(\cdot; \mathcal{S})$	Collaborative encoder function
$\text{TKZ}(\cdot)$	Tokenizer of LLM
$\text{WE}(\cdot)$	Word embedding lookup
$\text{proj}_\phi(\cdot)$	Projection into language space
$\text{Concat}(\cdot)$	Concatenation operation
$\text{Length}(\cdot)$	Sequence length
$\text{LLM}_\theta(\cdot)$	LLM with parameters $\theta$
$\text{BCE}(\cdot, \cdot)$	Binary cross-entropy loss
$\text{Softmax}(\cdot)$	Softmax function
$\text{Cosim}(\cdot, \cdot)$	Cosine similarity
$\text{Mean}(\cdot)$	Mean over a set of values
<b>Loss Functions</b>	
$\mathcal{L}_{think}$	Loss for the reasoning task
$\mathcal{L}_{rec}$	Loss for the recommendation task
$\mathcal{L}$	Final combined training loss

Table 5: Summary of Notations