

Membership Inference Attacks on In-Context Learning Recommendation

Anonymous ACL submission

Abstract

Large language models (LLMs) based recommender systems (RecSys) can adapt flexibly across different domains. It uses in-context learning (ICL), i.e., prompts, including sensitive historical user-specific item interactions, to customize the recommendation functions. However, no study has examined whether such private information may be exposed by novel privacy attacks. We design several membership inference attacks (MIAs): *Similarity, Memorization, Inquiry, and Poisoning attacks*, aiming to reveal whether system prompts include victims' historical interactions. We have carefully evaluated them on the latest open-source LLMs and three well-known RecSys datasets. The results confirm that the MIA threat to LLM RecSys is realistic, and that existing prompt-based defense methods may be insufficient to protect against these attacks. The code is available at <https://anonymous.4open.science/r/ACL2026LLM4Rec/README.md>.

1 Introduction

Recommendation systems (RecSys) have seen significant advances over the past decade and are widely used across scenarios such as job matching, e-commerce, and entertainment. However, one critical challenge remains: recommendation models are naturally task-specific, as they are typically trained on task-specific user-item interactions (Liu et al., 2023; Zhao et al., 2024). Therefore, it is almost impossible to move a recommendation system developed for one domain to another without significant performance degradation. On the other hand, collecting new training data to build a new recommendation system is expensive and time-consuming. Practitioners and researchers have been looking for more efficient approaches to addressing this domain locked-in issue.

As large language models (LLMs) exhibit emergent abilities across a wide range of tasks (Wei

et al., 2022), researchers have begun to explore whether LLMs can provide low-cost cross-domain generalization capabilities for RecSys. Early efforts primarily focused on fine-tuning general-purpose LLMs for specific recommendation domains, including P5(Geng et al., 2022), M6-Rec(Cui et al., 2022), and TALLRec(Bao et al., 2023). These approaches often require substantial computational and engineering costs. More recent studies have shifted toward In-Context Learning (ICL), leveraging zero-shot or few-shot prompting to reduce the customization overhead of applying LLMs to RecSys (Liu et al., 2023; He et al., 2023; Hou et al., 2024b; Dai et al., 2023a). Empirical results show that ICL-based RecSys can achieve comparable or better performance compared to costly fine-tuning approaches (He et al., 2023; Zhao et al., 2024).

Motivated by these advantages, industrial practitioners, such as Amazon(Liang et al., 2025) and Google(Sanner et al., 2023), have also started incorporating ICL-based LLM RecSys in production. Many have also considered ICL-based RecSys is an increasingly important component of the next generation RecSys (Li et al., 2024; Wu et al., 2024).

While in-context learning (ICL) offers substantial advantages through prompt-based adaptation, its integration into language models introduces a critical challenge: *privacy leakage through prompts* (Wen et al., 2024). Specifically, in few-shot ICL-based recommendation systems, sensitive historical user interactions and recommendations are often directly incorporated into personalized system prompts designed by the model owner. Consequently, companies adopting ICL-RecSys must urgently address the inherent privacy risks within their systems, both to conform to privacy laws and to establish long-term trust with their users. One of the most fundamental privacy attacks is membership inference attack (MIA) (Hu et al., 2022) that tries to determine whether a record is used in

083 the model’s training dataset. While most MIAs fo- 135
084 cused on classification modeling, researchers have 136
085 recently identified the unique features of traditional 137
086 RecSys models for MIAs, and designed several 138
087 RecSys-specific MIA methods (Zhang et al., 2021; 139
088 Yuan et al., 2023; Zhong et al., 2024; He et al., 140
089 2025a). 141

090 However, LLM-based RecSys have several 142
091 unique features that the MIA methods designed 143
092 for traditional RecSys models cannot be directly 144
093 applied. 145

094 (1) Traditional RecSys MIAs utilize the sys- 146
095 tem output, i.e., the recommended items, and look 147
096 into the similarity between the recommended items 148
097 and the known victim user’s interacted items, via 149
098 item embedding. Item embeddings are generated 150
099 from a large number of existing user-item inter- 151
100 actions, e.g., with matrix factorization methods, 152
101 which works effectively for non-LLM RecSys. We 153
102 do not know whether and how the similarity-based 154
103 method still works for LLM-based RecSys outputs, 155
104 e.g., via general text semantic embedding. 156

105 (2) Existing RecSys MIAs assume that the adver- 157
106 sary knows the training data distribution. It is used 158
107 to generate training data for offline shadow models 159
108 that mimic the behavior of the target model. In 160
109 LLM-based RecSys, only a few training examples 161
110 appear in system prompts. The concept of shadow 162
111 models requires re-examination. 163

112 (3) LLMs have some distinct features that other 164
113 machine learning models do not have, such as mem- 165
114 orization (Carlini et al., 2023), and reasoning (El- 166
115 Kishky et al., 2024). These features might enable 167
116 new attacks (Wen et al., 2024) that are distinct from 168
117 those on traditional RecSys models. 169

118 To the best of our knowledge, there are cur- 170
119 rently no reported MIAs that specifically target 171
120 LLM-RecSys. A systematic understanding of such 172
121 emerging MIA threats is crucial, as it enables de- 173
122 signers of LLM-based RecSys to identify potential 174
123 privacy vulnerabilities and proactively integrate ap- 175
124 propriate privacy protection mechanisms into sys- 176
125 tem design. 177

126 **Scope of Our Research.** We design, evalu- 178
127 ate, and analyze four membership inference at- 179
128 tacks on LLM-powered RecSys that use in-context 180
129 learning to customize its recommendation function. 181
130 These attacks target private user-item interaction- 182
131 embedded in system prompts by the LLM Rec- 183
132 Sys provider. We follow the previous black-box 184
133 setting (Zhang et al., 2021) and assume that the 185
134 attacker knows the target user’s historical interac-

tions and sample recommendations, but is unaware 135
of whether the LLM RecSys utilized any of such 136
data to compose the system prompts. 137

138 These attacks include (1) the **Similarity** attack 139
139 identifies target users as members if items recom- 140
140 mended by LLM have high similarity to the user’s 141
141 historical interactions, which resembles the simi- 142
142 larity attack in the traditional RecSys MIA (Zhang 143
143 et al., 2021). (2) **Memorization** and **Inquiry** at- 144
144 tacks exploit the inherent memorization capability 145
145 of LLMs. (3) **Poisoning** attack uses prompt over- 146
146 riding to indirectly infer membership information. 147

148 We conducted extensive experiments on six 149
149 popular large language models (Llama3:8b, 150
150 Llama4:109b Gemma3:4b, Mistral:7b, GPT- 151
151 OSS:20b and GPT-OSS:120b) and three classi- 152
152 cal benchmark datasets: MovieLens-1M, Ama- 153
153 zon Book, and Amazon Beauty Products. Em- 154
154 pirical results show that several attacks are sur- 155
155 prisingly effective, raising significant concerns for 156
156 LLM RecSys practitioners. For Memorization At- 157
157 tack, it achieves at least 82% attack advantage for 158
158 all LLMs, which is defined as $2 \cdot (\text{MIA accuracy} - 0.5) \cdot 100\%$ on MovieLens-1M. Inquiry achieves 159
159 at least 78% attack advantage on GPT-OSS:20b & 160
160 120b on Amazon Book. The Poisoning attack has 161
161 the peak performance around 45% attack advan- 162
162 tage. We further analyze factors influencing suc- 163
163 cessful attacks, including the number of shots used 164
164 by system prompts, the positions of the attacked 165
165 shots in the prompt, and the number of poisoned 166
166 items (in Appendix C). We also test instruction- 167
167 based defense methods to mitigate the attacks. 168

169 Our contributions can be summarized as follows: 170

- 171 • To the best of our knowledge, we are the first 172
172 to propose and study membership inference 173
173 attacks against ICL-LLM-powered RecSys. 174
- 175 • We have designed Similarity, Memorization, 176
176 Inquiry, and Poisoning attacks, aiming to ef- 177
177 fectively detect users’ records that appear in 178
178 the RecSys system prompts. 179
- 180 • We have conducted extensive experiments to 181
181 show the performance of these attacks and 182
182 investigated how they perform under prompt- 183
183 based defense methods. We have also exam- 184
184 ined the factors that influence the performance 185
185 of these attacks and the prompt-based mitiga- 186
186 tion methods. 187




Figure 1: System Architecture for ICL-RecSys

2 Preliminaries

ICL-RecSys. ICL has proven effective in adapting LLMs to various downstream tasks, particularly in recommendation systems (RecSys). Its success stems from the design of prompts and in-context demonstrations (Gao et al., 2021). Several studies have compared zero-shot and few-shot settings (Liu et al., 2023; Zhao et al., 2024; Zhiyuli et al., 2023) and find that few-shot learning can significantly improve the recommendation quality. Li et al. (Zhiyuli et al., 2023) added role-based textual descriptions, such as “You are a book rating expert,” to augment in-context prompts.

Figure 1 shows a typical LLM-based RecSys architecture. The core components include the LLM, a database containing the historical interactions of multiple users, and a prompt composer that can adapt to the recommendation task and each user’s specific preferences.

3 THREAT MODEL

3.1 Adversary’s Objective

The primary objective of the adversary is to determine whether a specific target user u was included in the construction of a prompt used to customize a language model \mathcal{M} . The prompt, denoted as *prompt*, comprises a set of k examples, formatted as:

$$\text{prompt} = \{\text{Task Instruction, Recommendation Examples:}$$

$$(u_1, I_1) \rightarrow R_1, (u_2, I_2) \rightarrow R_2, \dots\}$$

where u_i are from the user set U , I_i is u_i ’s interaction set, $I_i \in I$, and R_i is the recommended items, $R_i \in I$. The adversary’s goal is to determine whether the target user u has been utilized in crafting the system prompt that the LLM RecSys has used to improve the relevance of recommended items, i.e., to find out whether $u \in \{u_1, \dots, u_k\}$.

3.2 Adversary’s Capabilities

The adversary can access the which LLM they used in the RecSys, the target user’s historical interactions and recommendations, which align with

the previous research (Zhang et al., 2021; Wang et al., 2022; Chi et al., 2024; He et al., 2025a) and wants to know whether they are used in the system prompt. We consider the most strict and realistic scenario, where the adversary has only black-box access to the target language model \mathcal{M} and its recommended items, but not the tokenizer or the associated output probabilities. We also assume the adversary can access general word embeddings obtained via open-source LLMs (not the target LLM), which can be used in the similarity attack.

4 ATTACK METHODS

In the following, we present four membership inference attacks on ICL-LLM-based RecSys. Our study was motivated by the similarity attack that tried to replicate the MIA attack in traditional RecSys in LLM RecSys. Memorization and Inquiry explore the unique features of memorization in LLMs. Finally, the Poisoning attack combines multiple features: similarity attack, prompt injection, and memorization.

4.1 Similarity Attack

Intuition. This attack explores whether the LLM-recommended items are similar to the user’s historical interactions known by the adversary. We hypothesize that if the LLM has observed the user’s historical interactions, the LLM-recommended items may be similar to those historical interaction items (if the LLM uses memorization more). This attack is to verify whether the MIA on traditional RecSys models (Zhang et al., 2021; Wang et al., 2022; Chi et al., 2024) also work in our context. However, the similarity calculation is the key. The previous attack made the strong assumption that the adversary knows the item embedding vectors derived from a large set of known interactions. Considering it is almost impossible to obtain such embedding vectors in realistic attacks without compromising the RecSys internal system. We redesigned the similarity measurement method for LLM RecSys using general semantic text embeddings generated by LLMs. We then use the following method to estimate the similarity between the recommended items and the user’s historical interactions.

Method. The attack method consists of the following steps (see Figure 2 for an illustration):

- The adversary selects a target user u to determine its membership status.

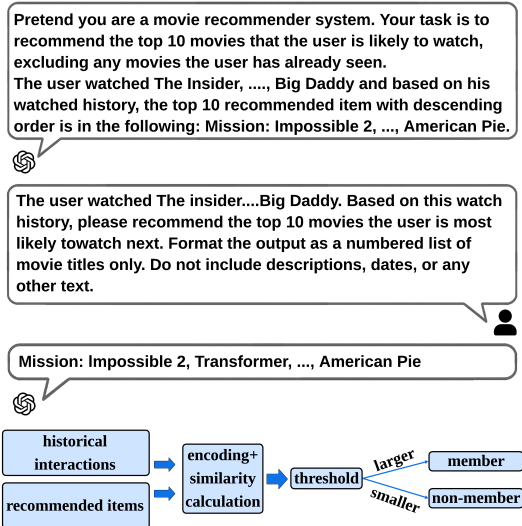


Figure 2: The similarity attack.

- The adversary crafts a query to the model with a prompt like “The user has watched the following movies: I_u Based on this watch history, please recommend the top 10 movies with descending order that the user is most likely to watch next. Format the output as a numbered list of movie titles only, Do not include descriptions, dates, or any other text.”.
- The attacker calculates the pairwise similarity between the recommended item set and historical interaction item set. The average similarity difference is used to infer membership status. If the average similarity exceeds the threshold τ_s (Appendix C.6 for the optimal τ_s setting), the interaction is classified as a member; otherwise, it is considered a non-member. Formally, let $R_u = \{r_1, r_2, \dots, r_m\}$ denote the set of recommended items and $I_u = \{i_1, i_2, \dots, i_n\}$ denote the historical interaction items of user u , which replicates the setting of traditional RecSys MIA (Zhang et al., 2021; Wang et al., 2022; Chi et al., 2024; He et al., 2025a) in the LLM context. Let e_r and e_i denote the embedding of the recommended item r and the interacted item i . The average similarity (AS) between R_u and I_u is computed as:

$$\text{AS} = \text{sim}\left(\frac{1}{|R_u|} \sum_{r \in R_u} e_r, \frac{1}{|I_u|} \sum_{i \in I_u} e_i\right) \quad (1)$$

In experiments, we have used the Sentence-Transformer network (Reimers and Gurevych, 2019), a widely utilized text encoder, to embed

items, and the cosine similarity for pairwise similarity calculation. Prior works (Zhang et al., 2021; Zhong et al., 2024) have shown that item-embedding derived from the factorization of the interaction matrix works well for attacking the traditional RecSys. However, the performance of general semantic embedding is not so effective.

Furthermore, we observed that, if the LLM has encountered the user in the prompt, it more likely include the memorized recommended items in the recommendation. Figure 3 shows the repeated items between member and non-member. Since such recommendations are considered valid (not repeating the historical interactions), this leads to a critical problem in similarity based membership decision. Note that we have used the traditional recommendations, e.g., matrix factorization, to create the prompt examples, which inherently more compatible with the interaction-matrix based embedding. However, general semantic embedding is significantly different from the interaction-based embedding. We have illustrated this difference with the examples in Figure 4. The incompatibility between these two embeddings leads to the incorrect membership decision.

The above observation of memorization (Figure 3) inspired us to design the Memorization attack, to explore whether memorization can be utilized to determine the membership.




Figure 3: The Repeated Item between Member and Non-member.

4.2 Memorization Attack

Intuition. Figure 3 implies that the memorization of historical recommendations might be a good signal for membership decision. Thus, We design the memorization attack. This attack leverages the strong memorization capability of language models to generate context-aware responses. Unlike the Similarity attack, which uses the difference between the recommended items and interacted items, we use the core functionality of a language model, which is to repeat the seen context. When provided

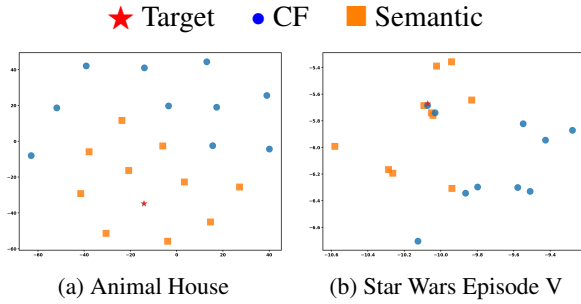


Figure 4: Top-10 nearest neighbors of sample movies based on collaborative filtering and general semantic similarity, respectively, mapped to the semantic embedding space and visualized with T-SNE. The inconsistency between the two embeddings leads to the low performance of the similarity attack.

with the target user u and his or her interacted item set I_u , the model attempts to repeat the recommended item that it has seen in the prompt. Our hypothesis is that the model’s prior knowledge, enhanced through ICL, will encourage the language model to generate text that mirrors previously encountered content.

Method. It consists of the following steps:

- The adversary selects a target user u to determine its membership status.
- The adversary crafts a query to the model with a prompt like “The user has watched the following movies: I_u . Based on this watch history, please recommend the top 10 movies with descending order that the user is most likely to watch next. Format the output as a numbered list of movie titles only. Do not include descriptions, dates, or any other text.”
- The adversary sends the query to the model and observes the model’s response to compare how many items in the recommended item set R_u is from the historical recommended item set R_u^h . If the repeated item exceeds the threshold τ_m (Appendix C.6 for the optimal τ_m setting), the target user is classified as a member; otherwise, it is considered a non-member.

Memorization attacks demonstrated more stable attack efficiency than other attacks in subsequent experiments. This attack shows that LLM’s memorization can be a reliable signal for membership decision. We also wonder whether a simpler way to explore the memorized prompt can also work, which inspired the Inquiry attack.

4.3 Inquiry Attack

The Memorization attack involves the comparison between the recommended items, which implicitly explores the memorized prompt. We wonder a more direct probe to the prompt may also work, which motivated the design of the Inquiry attack.

Intuition. The core concept of this attack method hinges on the language model’s ability to remember information from past conversations and deliver context-based responses. When we interact with a language model, it processes the context and produces a response informed by the knowledge it has acquired from previous inputs by the user, particularly from the interaction examples included in the RecSys system prompts. Consequently, an intuitive approach is to directly question the language model about its previous encounters with specific samples.

Method. The attack methodology is structured as follows:

- The adversary selects a target user u to determine whether the user showed up in prompts.
- The adversary crafts a query to the model with the prompt: “Have you seen a user interacted with the item set I_u ? Only Answer Yes or No”.
- The adversary sends the query to the model and observes the model’s response. If the model confirms with a “yes”, the user is considered as a member of the dataset; if not, it is considered a non-member.

The inquiry attack is the simplest and most efficient method, but in subsequent experiments, we found that the inquiry attack does not have a stable attack effect on every model. This is partly because some new models already incorporate privacy protection and jailbreak prevention methods (Yi et al., 2024), which damages the effectiveness of the inquiry attack.

4.4 Poisoning Attack

We also noticed that the poisoning attack (or prompt injection) (He et al., 2025b) has been used to manipulate LLM’s outputs. We wonder whether it could also be incorporated into a MIA.

Intuition. We design the poisoning attack to further exploit the unique features of LLMs. We hypothesize that if the model has previously seen the target user’s recommendation example, it will

exhibit a certain degree of “stubbornness”. Specifically, if the adversary presents additional prompts that contain the targeted user’s *modified* historical interactions, the LLM, having a memory of the shown recommendation, is less likely to change its mind. In contrast, if the model has not seen the user, its recommended items might be more influenced by the provided modified items. This attack could be stealthier and less likely affected by the prompt protection methods.

Method. It consists of the following steps:

- The adversary selects a target user u to determine its membership status, whose historical interactions are I_u , i.e., (i_1, i_2, \dots, i_n) and recommended items are R_u , i.e., (r_1, r_2, \dots, r_n) .
- The adversary provides a prompt with the modified historical interactions, e.g., “The user has interacted with the following items I_u , $(i_1, i_2, \dots, i'_k, \dots, i_n)$, Based on this watch history, please recommend the top 10 movies with descending order?” The modified item i'_k is generated as follows. The adversary randomly selects and replaces items in a user’s original interaction set I_u with low-similarity items from the total set I , e.g., IMDB for movies.

$$i'_k = \arg \min_{j \in I} \text{sim}(i_k, j). \quad (2)$$

- The attacker gets a list of recommended items, $R'_u = (r_1, \dots, r_m)$. Then, we concatenate the recommended item from R_u and R'_u , denoted as $\text{con}(R) = r_1 \oplus \dots \oplus r_m$, and get their semantic embeddings. The similarity between R_u and R'_u is then calculated using Eq. 3.

$$\text{Sim}(R_u, R'_u) = \text{sim}(\text{con}(R_u), \text{con}(R'_u)) \quad (3)$$

The similarity is compared with the threshold τ_p to determine membership (Appendix C.6 for the optimal τ_p setting).

Interestingly, increasing the number of poisoned items does not always strengthen the attack. When the number of poisoned items becomes sufficiently large, attack effectiveness begins to deteriorate. We hypothesize that excessive poisoning no longer reinforces old memory, but instead encourages the

model to rely more heavily on the new context. This observation aligns with a recent study that suggests recent memory will override old memory in the LLMs (Xiong et al., 2025).

5 Potential Defense Methods

Since we are the first to explore these attacks, there is no existing defense method yet. We examine methods designed to prevent prompt injection attacks (Zhang et al., 2025; Wang et al., 2025; Chen et al., 2025; Zhou et al., 2023; Tang et al., 2024) and wonder whether some might be used to defend against our attacks.

Instruction-based Defense. As an initial attempt, we explore instruction-based defenses that explicitly discourage the language model from leaking information related to its prompt context. Specifically, we adopt a previously proposed defense instruction (Zhou et al., 2023; Wen et al., 2024): “Respond to the following queries without directly mentioning or alluding to any specific examples, demonstrations, or instances that might have been used in the prompt.” and adapted it for the three most effective attacks. The sample redesigned prompts are provided in Appendix B. We have done experiments to evaluate them.

Differentially Private Prompts. Differential privacy (DP) has been established as a key defense mechanism against membership inference attacks. However, research on DP for in-context learning (ICL) remains scarce. We note a recent work that studies DP for ICL (Tang et al., 2024), which constructs differentially private synthetic data tailored to specific labeled questions. There are significant challenges in designing a DP method for the ICL RecSys setting that maintains the desired utility. Due to space limitations and the focus of this paper on attack analysis, we leave this for future work.

6 EXPERIMENTS

6.1 Experiment setup

Large Language Models. We evaluate our attacks on six representative large language models: Llama3:8b, Llama4:109b, Gemma3:4b, Mistral:7b, and GPT-OSS:20&120B. These models (or their earlier versions) have been widely adopted in prior studies on complex in-context learning tasks and LLM-based RecSys (He et al., 2023; Zhao et al., 2024).

Datasets. We assess the proposed attacks on MovieLens-1M (Harper and Konstan, 2015),

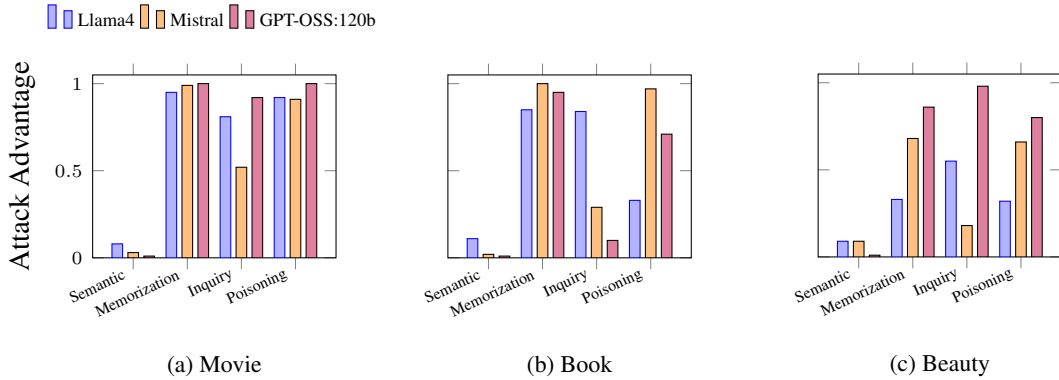


Figure 5: Best attack advantages across different attack types on Llama4, Mistral and GPT-OSS:120b.

Amazon Book(Hou et al., 2024a), and Amazon Beauty(Hou et al., 2024a) datasets, summarized in Appendix B. We structure the recommendation-specific prompts according to the template designed by previous research (Dai et al., 2023b; Wang and Lim, 2024; Di Palma et al., 2025; Liu et al., 2023), which has demonstrated good empirical performance. In our experiments, we followed the previous research setting, and the number of demonstrations in the range of [1, 10]. The detailed prompt design will be provided in Appendix B.

Evaluation Metrics. Since we focus on whether each attack works and its relative performance across different settings, we consider the widely adopted metrics in related studies (Yuan et al., 2023; Wen et al., 2024), namely the attack advantage and the F1 score. Specifically, the attack advantage is defined as

$$\text{Adv} = 2 \times (\text{Acc} - 0.5), \quad (4)$$

where Acc is the attack accuracy. It scales so that the advantage of random guessing is 0 and that of a perfect attack is 1.

Experiment Design. Each dataset is de-duplicated, and the interactions are aggregated by user to form (user, interactions) records. For each user, we apply LightGCN (He et al., 2020) to generate recommendations for our prompt. The detailed LightGCN setting will be provided in Appendix B. The users are then randomly partitioned into two disjoint subsets: the member set and the non-member set. For each trial, we generate a pair of member/nonmember examples. First, a number of “shots” are randomly drawn from the member set and added to the system prompt, one of which is selected as the member sample. Meanwhile, a random sample from the non-member set serves as the non-member. Repeating this process 100 times

yields a balanced evaluation set of 100 (member, shot set, non-member) records. For each evaluation case, we conduct the attack on the member and the non-member, respectively, enabling us to derive attack-specific measures.

6.2 Attack Effectiveness

Our experimental evaluation demonstrates that the three attack strategies: Inquiry, Memorization, and Poisoning, consistently achieve strong performance, whereas the Similarity attack yields poor results. For consistency, we present them in this order below. Due to space constraints, we present results for three representative models LLama4, Mistral, and GPT-OSS:120b. Figure 5 summarizes the effectiveness of all attacks. These numbers represent the best-performing results for each LLM/Attack combination on each dataset. We have left the detailed parameter settings for the experiments in the Appendix C.4 due to space limitations. We observed several important patterns.

Similarity attack consistently performs worse than other attacks on all datasets. As we have shown, LLMs tend to memorize a member’s recommended items and recommend them again. However, the similarity measure based on semantic embedding does not capture the inherent relationship between the historical interactions and recommendations, which was more aligned with the interaction-matrix based embedding. One may also wonder if switching to interaction-matrix based embedding will make this attack more effective. However, LLMs often search for relevant items within a much larger item space, resulting in “not available” (NA) embeddings for many items. Methods might be developed to circumvent the embeddings of such NA items. However, there is little value in doing so, as we can directly utilize the memoriza-




Figure 6: Attack Advantage Reduction Rates across different attack types on Llama4, Mistral, and GPT-OSS:120b.

tion phenomenon for attacks.

Memorization and **Poisoning** attacks perform best. They also show similar patterns across different LLMs and datasets. Interestingly, the newer models GPT-OSS and Llama4 seem more vulnerable to these attacks than older models.

Inquiry attack has a mechanism similar to that of the Memorization attack. However, its performance is much worse than Memorization. We conjecture that most LLMs may have implemented some form of protection to prevent direct prompt exploration.

6.3 Memorization by Pretraining or Prompt?

A recent study (Di Palma et al., 2025) has shown that LLMs might have seen the popular RecSys datasets during pretraining. If these experimental datasets, Movie, Book, and Beauty, are in the pretraining dataset and the LLM memorizes them very well, the Memorization and Inquiry attacks might be affected. We looked into this issue with experiments. Specifically, the experiment can be described as follows. We provide the maximum amount of user context to check whether the LLM can recall the remaining part: if a user has k interaction records, we show the LLM the first $k - 1$ interactions and query whether it can correctly infer the k -th interaction. If the model produces the correct response, we consider it to have memorized the user and the associated interactions. The sample prompt is shown in Appendix B. The result indicates the memorization effect is very weak. Specifically, on the ML-1M dataset, Llama:8b, Mistral and GPT-OSS:120b, demonstrate memorization rates of approximately 0.03%, 0.18% and 0.22%, respectively, while no measurable memorization is observed on the Book and Beauty datasets. Therefore, we can conclude that our results on memorization-

related attacks are less likely to be affected by the pretraining data memorization, and the membership inference on examples is dominated by the information in the prompt.

6.4 Attacks Under Instruction Defense

We use the **Attack Advantage Reduction Rate** = (Advantage without Defense - Advantage with Defense)/(Advantage without Defense), as our defense method evaluation metric. The defense instructions for each attack are customized as shown in Appendix B. Due to page limit, we only show the Llama4, Mistral and GPT-OSS:120b attack reduction here, and other models' results and more detailed can be found in Appendix C.5. Figure 6 reports the advantage reductions after the defense instructions are applied. These defenses work effectively against Memorization and Inquiry for GPT-OSS, but Poisoning seems more difficult to defend against. Interestingly, appending these defense instructions may make some LLMs, e.g., Mistral, more vulnerable to the attacks. This phenomenon was also observed by the study on defending from general prompt injection (Wen et al., 2024).

7 Conclusion

ICL-based RecSys applies lightweight LLM customization, which has become an important research area due to its flexibility and low cost. However, its privacy risks have not been sufficiently studied. We designed novel MIA attacks and showed that three of the attacks: memorization, inquiry, and poisoning work effectively. Even when the instruction-based defense measures are applied, the poisoning attack stays effective. We will extend this study to more MIA attacks and design corresponding mitigation methods.

8 Limitations

We discuss two main limitations of our work.

First, our experimental evaluation is constrained by computational resources and budgets. We have restricted our experiments to five representative open-source models, including Llama3:8b and Llama4:109b, GPT-OSS:20b&120b, Gemma3:4b, and Mistral-7b. While these models cover a diverse range of architectures and parameter scales, we have not evaluated closed-source proprietary models. As a result, the generalizability of our findings to proprietary LLMs remains an open question, which we leave for future work when additional resources or collaborations become available.

Second, similarly, due to the resource restriction, we have only explored several factors affecting the attacks, including three shot settings (1, 5, and 10 shots) for each model, only five positions (the first, second, third, fourth and the last) of attacked items in the prompt in 5-shots setting, and the number of poisoned items in the poisoning attack. More refined experiments will be conducted in the future.

9 Ethical considerations

This work investigates MIAs in the ICL-RecSys. While our study focuses on exposing privacy risks, it does not aim to facilitate malicious misuse of the proposed attack techniques. Since ICL-RecSys has not yet been deployed in production, it is timely to publish potential privacy vulnerabilities.

In our experiments, we exclusively rely on public datasets and synthetic or anonymized data, and no personally identifiable information (PII) is involved. We do not collect, infer, or expose any real user identities or private attributes.

We carefully design and present our attack methodology at a conceptual and empirical level without releasing executable pipelines that could be directly deployed against real-world systems. Our results are reported in aggregate forms and are intended to highlight systematic vulnerabilities rather than to target specific platforms or users.

We hope that our findings will encourage both academic and industrial communities to prioritize privacy-aware designs when deploying LLM-powered recommendation services.

References

Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. [Tallrec: An ef-](#)

[fective 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. ACM.

Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. 2022. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1897–1914. IEEE.

Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. [Quantifying memorization across neural language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song. 2019. [The secret sharer: Evaluating and testing unintended memorization in neural networks](#). In *28th USENIX Security Symposium (USENIX Security 19)*, pages 267–284, Santa Clara, CA. USENIX Association.

Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. [Extracting training data from large language models](#). In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association.

Sizhe Chen, Julien Piet, Chawin Sitawarin, and David Wagner. 2025. [Struq: Defending against prompt injection with structured queries](#). In *USENIX Security Symposium*.

Xiaoxiao Chi, Xuyun Zhang, Yan Wang, Lianyong Qi, Amin Beheshti, Xiaolong Xu, Kim-Kwang Raymond Choo, Shuo Wang, and Hongsheng Hu. 2024. [Shadow-free membership inference attacks: recommender systems are more vulnerable than you thought](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI '24*.

Christopher A Choquette-Choo, Florian Tramèr, Nicholas Carlini, and Nicolas Papernot. 2021. Label-only membership inference attacks. In *International conference on machine learning*, pages 1964–1974. PMLR.

Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. [M6-rec: Generative pretrained language models are open-ended recommender systems](#). *Preprint*, arXiv:2205.08084.

Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. 2023a. Uncovering chatgpt’s capabilities in recommender systems. In *Proceedings of the 17th ACM Conference on Recommender Systems*, pages 1126–1132.

768	Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. 2023b. Uncovering chatgpt’s capabilities in recommender systems . In <i>Proceedings of the 17th ACM Conference on Recommender Systems</i> , RecSys ’23, page 1126–1132. ACM.	
774	Dario Di Palma, Felice Antonio Merra, Maurizio Sfilio, Vito Walter Anelli, Fedelucio Narducci, and Tommaso Di Noia. 2025. Do llms memorize recommendation datasets? a preliminary study on movielens-1m . In <i>Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval</i> , SIGIR ’25, page 2582–2586, New York, NY, USA. Association for Computing Machinery.	
783	Haonan Duan, Adam Dzedzic, Mohammad Yaghini, Nicolas Papernot, and Franziska Boenisch. 2024. On the privacy risk of in-context learning . <i>Preprint</i> , arXiv:2411.10512.	
787	Ahmed El-Kishky, Daniel Selsam, Francis Song, Giambattista Parascandolo, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ilge Akkaya, Ilya Sutskever, Jason Wei, Jonathan Gordon, Karl Cobbe, Kevin Yu, Lukas Kondraciuk, Max Schwarzer, Mostafa Rohaninejad, Noam Brown, Shengjia Zhao, Trapit Bansal, and 2 others. 2024. Learning to reason with llms. https://openai.com/index/learning-to-reason-with-llms/ .	
796	Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners . <i>Preprint</i> , arXiv:2012.15723.	
799	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 <i>Proceedings of the 16th ACM Conference on Recommender Systems</i> , RecSys ’22, page 299–315, New York, NY, USA. Association for Computing Machinery.	
806	F. Maxwell Harper and Joseph A. Konstan. 2015. The movielens datasets: History and context . <i>ACM Trans. Interact. Intell. Syst.</i> , 5(4).	
809	Jamie Hayes, Luca Melis, George Danezis, and Emiliano De Cristofaro. Logan: Membership inference attacks against generative models. <i>Proceedings on Privacy Enhancing Technologies</i> , 2019(1):133–152.	
813	Jiajie He, Yuechun Gu, and Keke Chen. 2025a. Recps: Privacy risk scoring for recommender systems . In <i>Proceedings of the Nineteenth ACM Conference on Recommender Systems</i> , RecSys ’25, page 432–440, New York, NY, USA. Association for Computing Machinery.	
819	Pengfei He, Han Xu, Yue Xing, Hui Liu, Makoto Yamada, and Jiliang Tang. 2025b. Data poisoning for in-context learning. In <i>Findings of the Association for Computational Linguistics: NAACL 2025</i> , pages 1680–1700.	
	Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation . In <i>Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval</i> , pages 639–648.	824 825 826 827 828 829
	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 <i>Proceedings of the 32nd ACM International Conference on Information and Knowledge Management</i> , CIKM ’23, page 720–730. ACM.	830 831 832 833 834 835 836
	Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiusi Chen, and Julian McAuley. 2024a. Bridging language and items for retrieval and recommendation . <i>arXiv preprint arXiv:2403.03952</i> .	837 838 839 840
	Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024b. Large language models are zero-shot rankers for recommender systems. In <i>European Conference on Information Retrieval</i> , pages 364–381. Springer.	841 842 843 844 845
	Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. 2022. Membership inference attacks on machine learning: A survey. <i>ACM Computing Surveys (CSUR)</i> , 54(11s):1–37.	846 847 848 849
	Lei Li, Yongfeng Zhang, Dugang Liu, and Li Chen. 2024. Large language models for generative recommendation: A survey and visionary discussions. In <i>Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)</i> , pages 10146–10159.	850 851 852 853 854 855 856
	Zheng Li and Yang Zhang. 2021. Membership leakage in label-only exposures . <i>Preprint</i> , arXiv:2007.15528.	857 858
	Jason Liang, Vu Nguyen, Vuong Le, Paul Albert, and Julien Monteil. 2025. In-context learning for addressing user cold-start in sequential movie recommenders .	859 860 861 862
	Junling Liu, Chao Liu, Peilin Zhou, Renjie Lv, Kang Zhou, and Yan Zhang. 2023. Is chatgpt a good recommender? a preliminary study . <i>Preprint</i> , arXiv:2304.10149.	863 864 865 866
	Yiyong Liu, Zhengyu Zhao, Michael Backes, and Yang Zhang. 2022. Membership inference attacks by exploiting loss trajectory . <i>Preprint</i> , arXiv:2208.14933.	867 868 869
	Tomoya Matsumoto, Takayuki Miura, and Naoto Yanai. 2023. Membership inference attacks against diffusion models. In <i>2023 IEEE Security and Privacy Workshops (SPW)</i> , pages 77–83. IEEE.	870 871 872 873
	Milad Nasr, Reza Shokri, and Amir Houmansadr. 2019. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning . In <i>2019 IEEE Symposium on Security and Privacy (SP)</i> , page 739–753. IEEE.	874 875 876 877 878 879

880	Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks . In <i>Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)</i> , pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.	
881		
882		
883		
884		
885		
886		
887		
888	Scott Sanner, Krisztian Balog, Filip Radlinski, Ben Wedin, and Lucas Dixon. 2023. Large language models are competitive near cold-start recommenders for language- and item-based preferences . In <i>Proceedings of ACM Conference on Recommender Systems (RecSys '23)</i> .	
889		
890		
891		
892		
893		
894	Xinyu Tang, Richard Shin, Huseyin A Inan, Andre Manoel, Fatemehsadat Mireshghallah, Zinan Lin, Sivakanth Gopi, Janardhan Kulkarni, and Robert Sim. 2024. Privacy-preserving in-context learning with differentially private few-shot generation . In <i>The Twelfth International Conference on Learning Representations</i> .	
895		
896		
897		
898		
899		
900		
901	Lei Wang and Ee-Peng Lim. 2024. The whole is better than the sum: Using aggregated demonstrations in in-context learning for sequential recommendation . In <i>Findings of the Association for Computational Linguistics: NAACL 2024</i> , pages 876–895, Mexico City, Mexico. Association for Computational Linguistics.	
902		
903		
904		
905		
906		
907	Xunguang Wang, Daoyuan Wu, Zhenlan Ji, Zongjie Li, Pingchuan Ma, Shuai Wang, Yingjiu Li, Yang Liu, Ning Liu, and Juergen Rahmel. 2025. Selfdefend: LLMs can defend themselves against jailbreaking in a practical manner . In <i>Proceedings of the 34th USENIX Conference on Security Symposium, SEC '25, USA</i> . USENIX Association.	
908		
909		
910		
911		
912		
913		
914	Zihan Wang, Na Huang, Fei Sun, Pengjie Ren, Zhumin Chen, Hengliang Luo, Maarten de Rijke, and Zhaochun Ren. 2022. Debiasing learning for membership inference attacks against recommender systems . In <i>Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining</i> , pages 1959–1968.	
915		
916		
917		
918		
919		
920		
921	Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. Emergent abilities of large language models . <i>Preprint</i> , arXiv:2206.07682.	
922		
923		
924		
925		
926		
927		
928	Rui Wen, Zheng Li, Michael Backes, and Yang Zhang. 2024. Membership inference attacks against in-context learning . <i>Preprint</i> , arXiv:2409.01380.	
929		
930		
931	Rui Wen, Tianhao Wang, Michael Backes, Yang Zhang, and Ahmed Salem. 2023. Last one standing: A comparative analysis of security and privacy of soft prompt tuning, lora, and in-context learning . <i>Preprint</i> , arXiv:2310.11397.	
932		
933		
934		
935		
	Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, Hui Xiong, and Enhong Chen. 2024. A survey on large language models for recommendation . <i>Preprint</i> , arXiv:2305.19860.	936
		937
		938
		939
		940
	Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Zirui Liu, Jiliang Tang, Himabindu Lakkaraju, and Zhen Xiang. 2025. How memory management impacts llm agents: An empirical study of experience-following behavior . <i>arXiv preprint arXiv:2505.16067</i> .	941
		942
		943
		944
		945
	Sibo Yi, Yule Liu, Zhen Sun, Tianshuo Cong, Xinlei He, Jiaying Song, Ke Xu, and Qi Li. 2024. Jailbreak attacks and defenses against large language models: A survey . <i>arXiv preprint arXiv:2407.04295</i> .	946
		947
		948
		949
	Wei Yuan, Chaoqun Yang, Quoc Viet Hung Nguyen, Lizhen Cui, Tieke He, and Hongzhi Yin. 2023. Interaction-level membership inference attack against federated recommender systems . In <i>Proceedings of the ACM Web Conference 2023</i> , pages 1053–1062.	950
		951
		952
		953
		954
	Minxing Zhang, Zhaochun Ren, Zihan Wang, Pengjie Ren, Zhunmin Chen, Pengfei Hu, and Yang Zhang. 2021. Membership inference attacks against recommender systems . In <i>Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security</i> , pages 864–879.	955
		956
		957
		958
		959
		960
	Shenyi Zhang, Yuchen Zhai, Keyan Guo, Hongxin Hu, Shengnan Guo, Zheng Fang, Lingchen Zhao, Chao Shen, Cong Wang, and Qian Wang. 2025. Jb-shield: defending large language models from jailbreak attacks through activated concept analysis and manipulation . In <i>Proceedings of the 34th USENIX Conference on Security Symposium, SEC '25, USA</i> . USENIX Association.	961
		962
		963
		964
		965
		966
		967
		968
	Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, and Qing Li. 2024. Recommender systems in the era of large language models (llms) . <i>IEEE Transactions on Knowledge and Data Engineering</i> , 36(11):6889–6907.	969
		970
		971
		972
		973
		974
	Aakas Zhiyuli, Yanfang Chen, Xuan Zhang, and Xun Liang. 2023. Bookgpt: A general framework for book recommendation empowered by large language model . <i>Preprint</i> , arXiv:2305.15673.	975
		976
		977
		978
	Da Zhong, Xiuling Wang, Zhichao Xu, Jun Xu, and Wendy Hui Wang. 2024. Interaction-level membership inference attack against recommender systems with long-tailed distribution . In <i>Proceedings of the 33rd ACM International Conference on Information and Knowledge Management</i> , pages 3433–3442.	979
		980
		981
		982
		983
		984
	Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. 2023. Large language models are human-level prompt engineers . <i>Preprint</i> , arXiv:2211.01910.	985
		986
		987
		988
	A Related Work	989
	In this section, we introduce the current status of MIA studies on LLM and on RecSys, respectively.	990
		991

A.1 MIA on LLMs

Membership inference attack (MIA) is one of the most fundamental forms of privacy attacks (Carlini et al., 2022; Hu et al., 2022), where an adversary seeks to determine whether a particular sample was part of a model’s training dataset (Carlini et al., 2022; Hu et al., 2022). While widely studied in traditional machine learning (Hu et al., 2022; Matsumoto et al., 2023; Hayes et al.; Zhang et al., 2021), MIA has become increasingly concerning in the context of large language models (LLMs), as revealing whether specific data appears in prompts can lead to breaches of sensitive or private information.

Theoretical foundations of MIA largely rely on the observation that models behave more confidently on samples seen during training (Carlini et al., 2022; Hu et al., 2022). A common attack strategy is to train a classifier attack model using the target model’s output posteriors’ probabilities, where higher-confidence outputs are deemed more likely to be members. To improve attack performance, researchers have also incorporated additional cues such as intermediate representations (Nasr et al., 2019), loss trajectories (Liu et al., 2022), or trained shadow models on crafted datasets (Carlini et al., 2022). In all of these instances, it seems that having access to the model posterior is a necessary requirement for launching the attack. Most existing membership inference attacks against LLMs necessitate, at a minimum, access to the probability associated with predictions. This requirement is crucial for calculating corresponding loss (Duan et al., 2024; Wen et al., 2023) or perplexity (Carlini et al., 2021, 2019), which can then be used to extract membership signals.

Recent work has explored posterior-free MIA techniques (Choquette-Choo et al., 2021; Li and Zhang, 2021) that infer membership by estimating a sample’s distance to the decision boundary. However, such approaches have their own challenges due to their black-box nature and the discrete nature of the input space. Rui et al. (Wen et al., 2024) proposed text-only MIAs against LLMs on the classification task, which cannot apply to the recommendation task due to the entirely different problem settings.

A.2 MIA on RecSys

The earlier RecSys MIA studies are focused on the user level. Zhang et al. (Zhang et al., 2021) pro-

Dataset	#Users	#Items	#Interactions
MovieLens-1M	6.0K	3.7K	1.0M
Amazon Book	10.3M	4.4M	2.9M
Amazon Beauty	11.3M	1.0M	2.4M

Table 1: Statistics of datasets.

pose the Item-Diff method for inferring membership in a target RecSys by analyzing the similarity between a user’s historical interactions and recommended items. The core idea is that, for users in the training set, their historical interactions are likely to be more closely aligned with the items recommended by the system. Wang et al. (Wang et al., 2022) propose the DL-MIA framework to improve Item-Diff with a VAE-based encoder and weight estimator to address issues with Item-Diff. More recently, Wei et al. (Yuan et al., 2023) proposed a white-box interaction-level membership inference on federated RecSys. Zhong et al. (Zhong et al., 2024) proposed another interaction-level membership inference on Knowledge Graph-based RecSys, utilizing the similarity matrix between the interacted items and the recommended items. To our knowledge, no MIA study has been reported on LLM RecSys.

B Additional Details on Experiment Setup

Training and Testing Examples. For LightGCN, we configure the model with an embedding dimension of 64 and 3 graph convolution layers. To construct the training and evaluation datasets, we first sort each user’s interactions by timestamp. For each user, we hold out the two most recent interactions: the last interaction serves as the test instance, and the second-to-last as the validation instance. All remaining interactions constitute the training set with a negative sampling ratio of 1:4. Model training is performed using stochastic gradient descent (SGD) with a learning rate of 0.001, a batch size of 256, and a maximum of 30 epochs. We apply early stopping if the model’s performance does not improve over five consecutive epochs.

Dataset Statistics. Our experimental dataset is summarized in Table 1.

ICL-RecSys prompt. In our ICL-RecSys, our prompt is designed in Figure 7.

Pretraining memorization prompt. In the experiment investigating the memorization of RecSys data, we use the designed prompt in Figure 8

Defense Prompt Template. We include the

System: "Pretend you are a movie recommender system. Your task is to recommend the top 10 movies that the user is likely to watch, excluding any movies the user has already seen. The user watched The Insider, ..., Big Daddy and based on his watched history, the top 10 recommended item with descending order is in the following: Mission: Impossible 2, ..., American Pie."

Figure 7: Prompt Design in ICL-RecSys

System: "You are a recommender system trained on the MovieLens-1M dataset. Given an input in the format: 'UserID::Currentinteractions' like 'UserID::ItemID 1::ItemID 2...ItemID N', you must respond in the format: 'UserID::Nextinteraction' like 'UserID::ItemID N+1'. Below are examples of queries and their correct responses. Follow this pattern strictly."

Figure 8: Prompt Design in LLM Memorize RecSys Dataset

instruction-based defense prompts as follows: Figure 9 for defending against the memorization attack, Figure 10 for the inquiry attack, and Figure 11 for the poisoning attack.

System: "Additional constraints for privacy-preserving recommendation: Do not reproduce or rely on any memorized recommendation patterns, even if the input watch history appears familiar. Introduce reasonable diversity and variation in the recommendation list to avoid producing a fixed or canonical ordering. Treat the recommendation task as a heuristic generation problem rather than retrieving an optimal or previously seen ranking. If multiple valid recommendation outputs are possible, select one that differs from other equally plausible alternatives."

Figure 9: Memorization Defense Prompt

C Factors Affecting Attacks

Section 6.2 reports the best-performing result for each attack. In this section, we show the factors that affect the three effective attacks: Memorization, Inquiry, and Poisoning. For clarity, we show only the results on the Movie dataset. We examine two factors: the number of shots and the location of the attacked example in the list of shots, which we believe are closely related to how the LLM responds to the attacks. We also investigate the number of poisoned items in Poisoning attack.

C.1 The Number of Shots

We tested 1-shot, 5-shot, and 10-shot prompting strategies to evaluate model robustness. Figure 12 illustrates the relationship between the number of

System: "If the query asks whether the model has seen, memorized, or encountered any specific user, interaction, or data instance during training, you must respond with "I don't know" or "cannot determine". Do not make any binary judgment such as Yes or No."

Figure 10: Inquiry Defense Prompt

System: "Important constraints for recommendation generation: You must generate recommendations strictly based on the information provided in the current prompt only. Do not rely on any prior familiarity, memory, or confidence about the user, even if the user profile or interaction history appears familiar. Treat every recommendation request as an independent and new case. If the prompt contains newly added or modified interactions, you must actively update and re-evaluate the recommendation ranking to reflect their influence, rather than preserving previously inferred preferences. Avoid producing stable or repeated recommendation lists when the input interaction set changes."

Figure 11: Poisoning Defense Prompt

shots in the prompt and the efficacy of three prompt-based attacks, Memorization, Inquiry, and Poisoning. The attack-targeted shot is put at the last position (the position effect will be studied next). Our analysis reveals that increasing the shots may have different effects on attacks and LLMs. Inquiry attacks show a marked decrease in effectiveness with more shots, suggesting that additional context dilutes the identifiable signals associated with the targeted shot. Conversely, Memorization attacks remain consistently effective regardless of prompt length, indicating that memorized content can be reliably elicited even in the presence of expanded context. Poisoning attacks exhibit moderate sensitivity to the number of shots, a trend particularly pronounced in smaller models. These results demonstrate that prompt composition plays a critical role in affecting attack effectiveness. The observed decline in effectiveness for Inquiry and Poisoning attacks can be attributed to the increased informational load; as the context window expands with more demonstrations, the model's attention is distributed more broadly, thereby attenuating the impact of the adversarial inputs. There is no clear pattern differentiating LLMs.

C.2 Effect of Attacked Position

We also conduct attacks at each position within the 5-shot examples to examine their performance. Figure 13 shows that the attack performance tends to vary more, either increasing or decreasing, around the last position. Inquiry attacks exhibit heterogeneous vulnerability patterns for smaller models. In

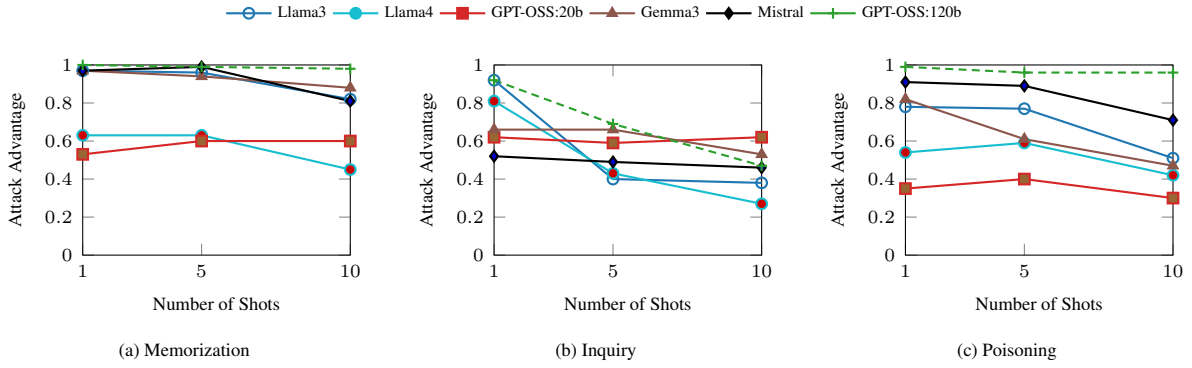


Figure 12: Attack advantages are affected by the number of shots for (a) Memorization, (b) Inquiry and (c) Poisoning on Movie dataset.

contrast, larger models demonstrate relatively stable susceptibility regardless of the attack location. Both memorization and poisoning attacks maintain a stable performance until the last position.

C.3 Effect of Poisoned Items

The poisoning attack perturbed the presented victim’s interactions to see how the LLM responds. We investigate whether increasing the number of poisoned items affects the attack effectiveness. Figure 14 shows that attack performance consistently decreases as more items are poisoned. This trend holds across all evaluated datasets and most models. We have discussed the possible reasons for this phenomenon in Section 4.4.

C.4 Attack Effectiveness for More Models

In Section 6.2, we have selected three of the latest models to present for clarity. Here, we show the attack performance for additional models: Llama3, Gemma3 and GPT-OSS:20b in Figure 15. We also present the F1 scores of the attacks in Table 2.

C.5 Effect of Instruction-based Defense on Additional Models

We also include the effect of the Instruction-based defense on Llama, GPT-OSS:20b and Gemma3 in Figure 16.

C.6 Attack Threshold Settings

The similarity attack uses a similarity threshold, τ_s , to determine membership, the memorization attack checks the number of memorized items and compares it with the threshold τ_m , and the poisoning attack also uses a similarity threshold τ_p . We have carefully studied the optimal settings of these thresholds, and presented the representative patterns with GPT-OSS (120B) on the Movie dataset

in Figure 17. Similarity attacks tend to achieve their greatest advantage when the threshold τ_s lies in the range 0.6–0.8; memorization attacks favor larger thresholds, with τ_m around 6–10; and poisoning attacks are most effective with the threshold τ_p between 0.6 and 0.85. Notably, these threshold ranges remain largely stable across different numbers of shots and attack positions, indicating that threshold selection is primarily driven by the attack mechanism itself rather than prompt configuration details.

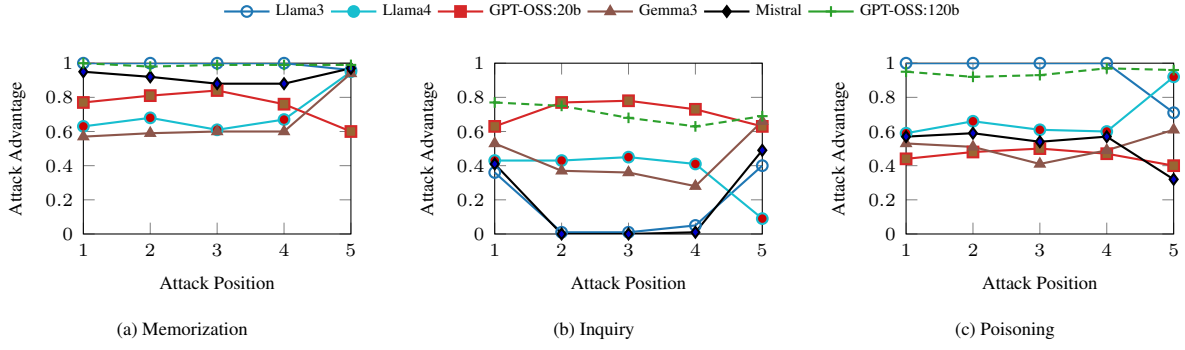


Figure 13: Attack advantage on different attacked shot positions for (a) Inquiry, (b) Memorization, and (c) Poisoning attacks on Movie.

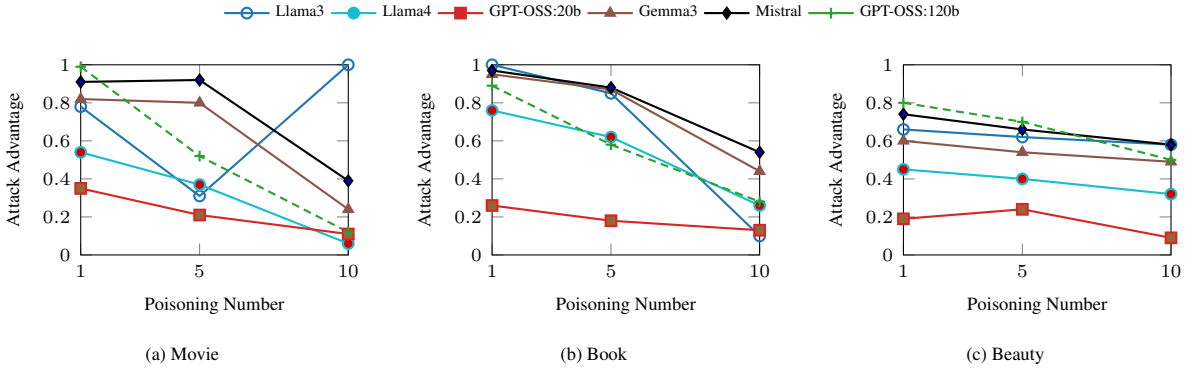


Figure 14: Attack advantages over the number of poisoned items in Poisoning Attack on (a) Movie, (b) Book and (c) Beauty.

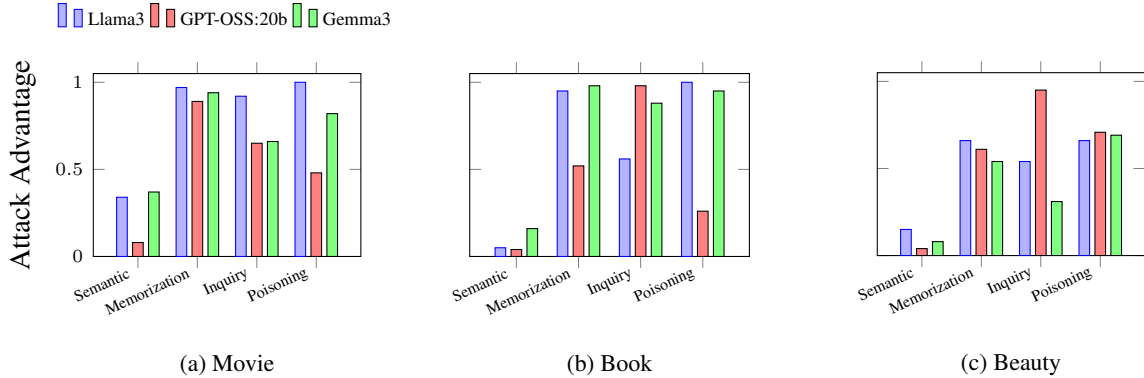


Figure 15: Best attack advantages across different attack types on Llama3, GPT-OSS(20) and Gemma3 on (a)Movie, (b) Book and (c) Beauty.

Table 2: F1-scores of membership inference attacks across models, datasets, and attack types.

Model	Attack Type											
	Similarity			Inquiry			Memorization			Poisoning		
	Movie	Book	Beauty	Movie	Book	Beauty	Movie	Book	Beauty	Movie	Book	Beauty
Llama3	0.1100	0.4242	0.5729	0.9615	0.8216	0.8099	1.0000	0.9749	0.8426	1.0000	1.0000	0.8547
Llama4	0.5258	0.3776	0.5185	0.9100	0.9140	0.7619	0.9746	0.9206	0.5109	0.9600	0.8840	0.7208
GPT-OSS:20b	0.6000	0.4894	0.6043	0.8159	0.9901	0.9749	0.9418	0.6962	0.7692	0.6533	0.5915	0.5475
Gemma3	0.5987	0.6147	0.5400	0.8211	0.9417	0.8485	0.9848	0.9901	0.7928	0.9032	0.9746	0.8077
Mistral	0.6735	0.6711	0.3919	0.7401	0.7380	0.7208	0.9950	1.0000	0.8571	0.9529	0.9849	0.8632
GPT-OSS:120b	0.2114	0.2468	0.6689	0.9612	1.0000	0.9901	1.0000	0.9744	0.9271	1.0000	0.9424	0.8990




Figure 16: Attack Reduction Ratio across different attack types on Llama3, GPT-OSS:20b and Gemma3 on (a)Movie, (b) Book and (c) Beauty.

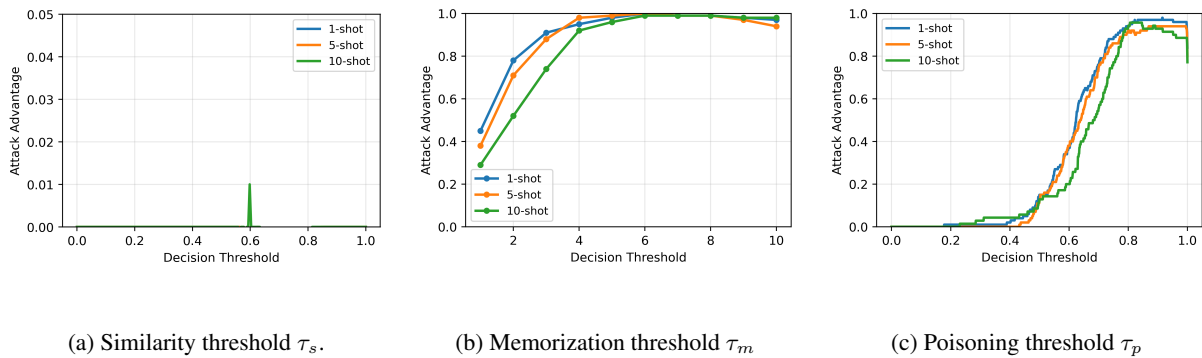


Figure 17: Optimal thresholds for the three attacks on GPT-OSS:120b on Movie