
Personas within Parameters: Fine-Tuning Small Language Models with Low-Rank Adapters to Mimic User Behaviors

Himanshu Thakur
JPMorgan Chase & Co
Palo Alto, CA 94304
himanshu.thakur@chase.com

Eshani Agrawal
JPMorgan Chase & Co
Palo Alto, CA 94304
eshani.agrawal@jpmchase.com

Smruthi Mukund
JPMorgan Chase & Co
Palo Alto, CA 94304
smruthi.mukund@chase.com

Abstract

A long-standing challenge in developing accurate recommendation models is simulating user behavior, mainly due to the complex and stochastic nature of user interactions. Towards this, one promising line of work has been the use of Large Language Models (LLMs) for simulating user behavior. However, aligning these general-purpose large pre-trained models with user preferences necessitates: (i) effectively and continuously parsing large-scale tabular user-item interaction data, (ii) overcoming pre-training-induced inductive biases to accurately learn user specific knowledge, and (iii) achieving the former two at scale for millions of users. While most previous works have focused on complex methods to prompt an LLM or fine-tune it on tabular interaction datasets, our approach shifts the focus to extracting robust textual user representations using a frozen LLM and simulating cost-effective, resource-efficient user agents powered by fine-tuned Small Language Models (SLMs). Further, we showcase a method for training multiple low-rank adapters for groups of users or *persona*, striking an optimal balance between scalability and performance of user behavior agents. Our experiments provide compelling empirical evidence of the efficacy of our methods, demonstrating that user agents developed using our approach have the potential to bridge the gap between offline metrics and real-world performance of recommender systems.

1 Introduction

In the rapidly evolving landscape of recommendation systems, user behavior simulation has become crucial for cost, safety, and efficiency benefits. A promising line of work towards this has been Agentic-LLM systems to simulate a user in a recommendation environment. While prior work has been able to effectively establish the value of such systems, representing a user faithfully and at scale has for a large part been unsolved. Thus, we posit that developing robust user simulation agents requires the ability to continuously capture a user’s static and dynamic behavior from large volumes of noisy interactions.

To address these challenges, we propose a method to build scalable and efficient user behavior agents. First, we demonstrate that a single LLM can transform large volumes of user interactions into meaningful textual representations. Next, we use low-rank adapters to fine-tune a SLM on this data,

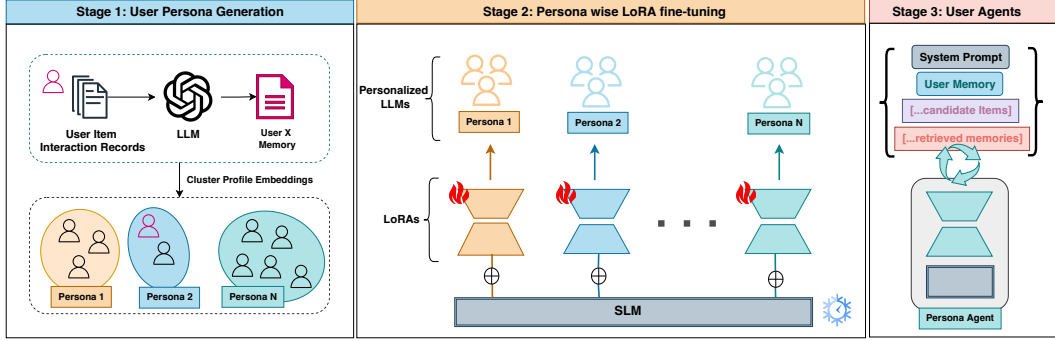


Figure 1: An overview of our proposed approach. First, we iterate over all user-item interaction records to generate user profile features and explanations for certain unique interactions. Then, we cluster users based on their profile embedding to generate N personas. In stage 2, we train a low-rank adapter on all user interaction data within the persona keeping the base SLM weights frozen. Finally, in stage 3, we utilize at persona-level SLMs to build user agents.

creating a user agent that closely mimics actual user actions. While previous works have employed Retrieval Augmented Generation (RAG) systems to build agents, we argue that fine-tuning an SLM is more effective and scalable in real-world applications. Additionally, we show that users can be grouped into personas, balancing the number of parameters required with the personalization quality of these user agents. Thus, our main contributions can be summarized as follows:

1. We demonstrate a hierarchical knowledge distillation process that utilizes an LLM and self-reflection to convert tabular user interactions into textual user profiles and ground their interactions in reason. We find that both of these components are key to improving the user agent’s performance.
2. We show that low-rank adaptation of SLMs can match or exceed the performance of frozen LLMs for building personalized agents. Furthermore, we show that Retrieval Augmented Fine-tuning (RAFT) helps in the effective utilization of short-term and long-term memory needed to build the user agent.
3. Finally, we show that clustering users based on their profiles generated during distillation helps strike a balance between personalization performance and the number of model parameters needed. Furthermore, we find that fine-tuning multiple Low-Rank Adapters (LoRA), one for each of these clusters, leads to better performance compared to using the same model for all users.

2 Related Works

In recent years, the field of LLMs has witnessed significant advancements, particularly in the domains of personalization, simulation of human behavior, and recommendation systems. This section reviews the most pertinent works, categorized by their thematic focus.

Simulating User Agents using LLMs: The simulation of human behavior through LLMs shows promise for interactive systems. Jiang et al. [4] and Park et al. [10] enhance virtual agents with personality traits and lifelike interactions. Wang et al. [15] and Yang et al. [16] improve recommendation accuracy and adaptability, highlighting LLMs’ potential in enhancing user interactions and recommendations. Our works aims to progress this research for SLMs.

Enriching Users and Items using LLMs: Sun et al. [13] propose a framework to enhance user representation by distilling interactions into a hierarchy of user descriptions. This method leverages persona-specific data to improve the contextual appropriateness of response predictions. Zhang et al. [17] explore the use of generative agents in recommendation systems, treating items as agents that engage in multi-turn interactions with users, thereby enriching each other’s metadata. This approach proves particularly effective in cold start scenarios and demonstrates robust generalization capabilities. We employ a combination of both approaches in this work to generate enriched user and item information.

Parameter-Efficient Tuning for Personalization: Tan et al. [14] use low-rank adapters to create instruction-tuned LLMs for each user, enhancing accuracy and generalization of user behavior prediction. Our approach builds on this by incorporating comprehensive user profiles and enriched interactions to better encode preferences. Also, our aim is to check the effectiveness of the approach on SLMs. The minimal updated parameters and plug-and-play nature of PEFT make it ideal for efficient SLM personalization. Our work embeds user history within user level LoRA parameters, providing each user agent with a robust and expressive long-term memory.

LLMs for Personalization: Current research on LLM personalization focuses on crafting prompts using historical user data to customize responses [1]. Strategies include: (1) Vanilla Personalized Prompts: Use in-context and few-shot learning to embed user behavior history (e.g., personal rating histories by Dai et al. [2] and Kang et al. [5], task-specific interaction histories by Liu et al. [7]). (2) Retrieval-Augmented Personalized Prompts: Retrieve relevant user history to enhance prompts (e.g., LaMP [12], AuthorPred [6], Pearl [9]). (3) Profile-Augmented Personalized Prompts: Summarize user preferences into natural language profiles (e.g., Richardson et al. [11], ONCE [8]). These methods aim to improve the personalization of LLM responses. While these advancements are notable, the current methods struggle to effectively leverage long-term user preferences especially in SLMs. Our approach overcomes this limitation by embedding long-term preferences directly into the model parameters using the personal PEFT module.

3 Methodology

In this section, we define the standard notation and methodologies used throughout the paper. We begin by introducing the fundamental concepts and terminologies that form the basis of our approach. This includes the representation of users, items, and their interactions within a traditional recommender system. We then describe our approach on construction and training of user agents to predict user preferences.

3.1 Preliminaries

In a traditional recommender system, we have a set of users $U = \{u\}$, a set of items $I = \{i\}$, and a set of their interaction records denoted by $D = \langle u, i \rangle$. In our system, a user agent is built using an open source SLM and has short-term and long-term memories, denoted by M_s and M_l , respectively. Thus, a user agent can be defined as:

$$y_i \leftarrow f_{\text{SLM}}(M_s, \text{item}_i, \{\text{item}_1, \dots, \text{item}_k\}),$$

where M_s is the user agent’s short-term memory, item_i is the description of the i -th item, and $\{\text{item}_1, \dots, \text{item}_k\} \subseteq M_l$ are items retrieved using item_i as a query by a similarity function $m(e(\text{item}_i), e(\text{item}_k)) < \delta$. We implement m as cosine similarity and e using a frozen embedding model, with δ being a threshold determined heuristically. Finally, y_i is the response of the SLM-based user agent and can have multiple tokens in its output.

3.2 Distilling User Preferences and Enriching Interactions

A key challenge of using SLMs for building user agents has been the appropriate representation of the user. Inspired by [13] and [17], we use an hierarchical self-reflection based algorithm that inputs a chronologically sorted user-item interaction dataset in batches D and outputs $D_{\text{enriched}} = \{\langle u_{\text{enriched}}, i_{\text{enriched}} \rangle\}$, where the user and item descriptions reflect the preferences of the user and the item, respectively. The process results in generation of these two key features per user:

User Profile (M_s): A description of general traits of the user inferred based on their interaction history. This act as the user’s short term memory M_s .

Enriched User Interaction (M_l): A textual description that explains the why the user liked or dislikes a given item. To create a persona agent with an SLM, we use M_s as a fixed system prompt during training and inference. We also use user profile embeddings to identify user clusters or personas. Additionally, M_l is embedded and stored in a database for use in RAG and RAFT. During training, we fetch the k nearest items using the current item embedding and use these as queries to retrieve their enriched versions.

3.3 Personalized Low-Rank Adapters

Central to our work is parameter efficiency, as we aim to simulate a very large number of users, which is typical in a recommender system. Previous research [14] has demonstrated that training multiple low-rank adapters per user can effectively capture user-specific knowledge and style. However, the number of parameters required for this approach would be enormous. Therefore, we propose that multiple users can be clustered into specific behavioral groups or personas, and training a low-rank adapter for each group could suffice.

We group users into K (note the difference in notation for k retrieval items) personas and fine-tune an SLM using K separate low-rank adapters, one per persona, on a downstream user-item interaction task. Once trained, we load the LoRA corresponding to the user’s persona during inference to predict the target(s). We study the effects of these two types of memory with two fine-tuning variants:

Using M_s only: For a user U , we train a low-rank adapter on the enriched user persona from the distillation step given as: $y \leftarrow F_{\theta_{u,k}}(\text{user_profile}, \text{movie_to_rate})$.

Using M_s and M_l : For a user U , we train a low-rank adapter on the enriched user profile and top- k relevant item interactions retrieved by a similarity function $m(e(U + I), e(I + I')) < \delta$ given as: $y \leftarrow F_{\theta_{u,k}}(\text{user_profile}, \text{movie_to_rate}, \text{top_k_enriched_interactions})$. We also perform experiments with training a single LoRA for all users, called θ_{single} , to compare with our method of training a separate LoRA for each persona, called θ_{persona} .

4 Experiments and Results

In this section, we detail the models, datasets, and methodologies used to evaluate the effectiveness of our proposed approach for generating user personas and simulating user interactions in a recommender system. For distillation and persona generation, we utilize GPT-4o to generate M_s and M_l , and text_ada_002 embeddings to embed M_s and M_l . We employ LLaMA-3-8B-Instruct (referred as LLaMA-3) as our LLM and Phi-3-Mini-4k-Instruct as our SLM (referred as Phi-3-Mini), which has 3.8 billion parameters and a context length of 4,092 tokens. To evaluate our approach, we use the MovieLens-1M dataset [3], containing 1 million ratings from 6,000 users on 4,000 movies. The dataset includes movie information (IDs, titles, genres, actors, release years) and user ratings (1 to 5). We use this dataset to (i) initialize user profiles, capturing their unique tastes and social traits, and (ii) train an SLM-powered user agent to accurately rate movies. We pre-process the data to include users with more than 100 interactions, ensuring long interaction histories.

We conducted experiments on 200 users with 100-200 interactions each. Using GPT-4, we distilled the historical interactions and clustered users into 4 personas with the KMeans++ algorithm. Then, using cosine similarity, we identified the top-1 nearest interactions for each prompt. We fine-tuned the SLM using low-rank adapters of rank 256 (details in Appendix A). For fair comparison, the dataset for θ_{single} was sub-sampled to match the largest dataset used for θ_{persona} . Moreover, given that each adapter in θ_{persona} has access to different amounts of data, we trained two versions of θ_{single} to enable fair comparisons: one using a dataset size equal to the smallest persona or $\theta_{\text{single (min)}}$ and using the largest $\theta_{\text{single (max)}}$. For each user, we merged their LoRA with the base SLM to predict movie ratings on a Likert scale of 1 to 5. We evaluated these agents using Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Unrelated Response Rate (URR), which is the percentage of times the model produces an unrelated response.

SLM vs LLM for Personalization: We begin by comparing vanilla prompting on both LLaMA-3 and Phi-3-Mini. We use both short-term and long-term memories as these models are not fine-tuned and not adding long-term memory could result in poor performance. In this particular instance, we observe in Table 4(No.1 and 4) that LLaMA-3 beats Phi-3-Mini in terms of both MAE and RMSE. However, when we fine-tuned the SLM (No. 7), we observe that the performance increases significantly and beats LLaMA-3 baseline. We also note that only using short term memory improves the frozen model’s performance (No. 0 vs 1 and 3 vs 4). To understand the utility of user persona’s we prompt LLaMA-3 with only the long-term memories $k = 1$ and find that it achieves the least performance (No.2 vs 0 and 1). However, LLaMA-3 does not produced an unrelated response at all while both non fine-tuned and fine-tuned versions of Phi-3-Mini produce unrelated responses across the runs.

No.	Base Model	Adapter(s)	Memories Used	RMSE (\downarrow)	MAE (\downarrow)	URR (\downarrow)
0	LLaMA-3	-	M_s	1.158	0.847	0%
1	LLaMA-3	-	$M_s + M_l$	1.221	0.883	0%
2	LLaMA-3	-	M_l	1.229	0.888	0%
3	Phi-3-Mini	-	M_s	1.203	0.863	4.23%
4	Phi-3-Mini	-	$M_s + M_l$	1.315	0.952	2.93%
5	Phi-3-Mini	$\theta_{\text{single (max)}}$	M_s	1.312	0.940	1.30%
6	Phi-3-Mini	$\theta_{\text{single (min)}}$	M_s	1.180	0.848	1.93%
7	Phi-3-Mini	$\theta_{\text{single (max)}}$	$M_s + M_l$	1.150	0.834	1.15%
8	Phi-3-Mini	$\theta_{\text{single (min)}}$	$M_s + M_l$	1.337	1.042	3.01%
9	Phi-3-Mini	θ_{persona}	M_s	1.292	0.937	2.29%
10	Phi-3-Mini	θ_{persona}	$M_s + M_l$	1.171	0.881	1.77%

Table 1: Performance comparison of user agents built using different model and dataset strategies. The absence of an adapter means we did not fine-tune and only prompted the base model. $\theta_{\text{single (min)}}$ signifies that a single LoRA was trained on ≈ 2100 samples, while $\theta_{\text{single (max)}}$ indicates it was trained on ≈ 5200 samples. The presence of M_s or M_l indicates the type of memories that were used to train and evaluate the models. The numbers in bold show the best scores per category, and underlines show the comparison between single LoRA and persona LoRA.

Does user clustering help? We find that clustering users into specific personas increases performance when the model is fine-tuned on $M_s + M_l$, as observed in Table 4. Specifically, in the comparison between No. 10 and Nos. 7 and 8, we see that No. 10 performs better, as the average RMSE of Nos. 7 and 8 is higher. Although No. 7 outperforms No. 10, it uses 2.5 times more training data, which accounts for its better performance. (As per Appendix B, Figure 3, the performance of θ_{persona} improves with an increase in training dataset size.) Similarly, No. 9 achieves a lower RMSE compared to No. 5, but higher than No. 6. This discrepancy can be attributed to the lesser data diversity in each θ_{persona} 's training set compared to θ_{single} , which could lead to performance drops among some persona clusters (Appendix B, Figure 2.)

Are long-term memories (M_l) useful? In 2 out of 3 fine-tuning settings, we find that using long-term memories in addition to short-term memories provides better results than training only on short-term memories. This can be observed in the comparison between No. 5 and No. 7, where RMSE reduces by 0.181 and MAE by 0.056 Table (4). Similarly, between No. 9 and No. 10, the latter achieves a 0.16 lower RMSE and 0.065 lower MAE. However, we observe an anti-pattern in the comparison between No. 6 and No. 8, where No. 6 achieves a lower RMSE, indicating that the model using long-term memory could not generalize well. We identify that our approach does not account for the quality of data retrieved during fine-tuning, and we leave this as a line of future work.

5 Conclusions

In conclusion, our work has demonstrated that distilling user interaction data into representative user knowledge, manifested as short-term and long-term memories, is highly effective in enabling user behavior agents. Furthermore, we have shown that fine-tuning SLMs with a low-rank adapter per persona significantly improves the model's ability to mimic user behavior, compared to training a single adapter for all users. In summary, our approach bridges the gap between offline evaluation metrics and real-world performance by utilizing scalable user simulation agents, especially in settings where users have extensive interaction histories. We are optimistic that our findings will pave the way for more advanced and personalized user interaction systems, showcasing the strength and achievement of our method.

6 Limitations and Future Work

Our results demonstrate the promising use of Small Language Models (SLMs) over Large Language Models (LLMs) in user simulation settings. However, few limitations serve as future research directions. The distillation process, reliant on LLM prompting for quality outputs, can be slow with large datasets. Fine-tuning is constrained by computationally intensive hyperparameter tuning, and exploring other parameter-efficient methods beyond low-rank adapters could yield better results. While clustering users into personas shows positive outcomes, optimizing persona generation by considering more easily acquired user features is needed. We encourage further exploration of SLMs in user simulation for recommender systems and suggest training smaller models to generate enriched user personas using our knowledge distillation technique from LLMs.

References

- [1] Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, and Xingmei Wang. When large language models meet personalization: Perspectives of challenges and opportunities. *arXiv preprint arXiv:2307.16376*, 2023.
- [2] Sunhao Dai, Ninglu Shao, Haiyuan Zhao, Weijie Yu, Zihua Si, Chen Xu, Zhongxiang Sun, Xiao Zhang, and Jun Xu. Uncovering chatgpt’s capabilities in recommender systems. *arXiv preprint arXiv:2305.02182*, 2023.
- [3] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4), dec 2015.
- [4] Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. Personallm: Investigating the ability of large language models to express personality traits. *arXiv preprint arXiv:2305.02547*, 2023.
- [5] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv:2305.06474*, 2023.
- [6] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. Teach llms to personalize—an approach inspired by writing education. *arXiv preprint arXiv:2308.07968*, 2023.
- [7] Junling Liu, Chao Liu, Renjie Lv, Kang Zhou, and Yan Zhang. Is chatgpt a good recommender? a preliminary study. *arXiv preprint arXiv:2304.10149*, 2023.
- [8] Qijiong Liu, Nuo Chen, Tetsuya Sakai, and Xiao-Ming Wu. Once: Boosting content-based recommendation with both open- and closed-source large language models. *arXiv:2305.06566*, 2023.
- [9] Sheshera Mysore, Zhuoran Lu, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, and Tara Safavi. Pearl: Personalizing large language model writing assistants with generation-calibrated retrievers. *arXiv preprint arXiv:2311.09180*, 2023.
- [10] Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22, 2023.
- [11] Chris Richardson, Yao Zhang, Kellen Gillespie, Sudipta Kar, Arshdeep Singh, Zeynab Raeesy, Omar Zia Khan, and Abhinav Sethy. Integrating summarization and retrieval for enhanced personalization via large language models. *arXiv preprint arXiv:2310.20081*, 2023.
- [12] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. Lamp: When large language models meet personalization. *arXiv preprint arXiv:2304.11406*, 2023.

- [13] Chenkai Sun, Ke Yang, Revanth Gangi Reddy, Yi R Fung, Hou Pong Chan, ChengXiang Zhai, and Heng Ji. Persona-db: Efficient large language model personalization for response prediction with collaborative data refinement. *arXiv preprint arXiv:2402.11060*, 2024.
- [14] Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. Democratizing large language models via personalized parameter-efficient fine-tuning. *arXiv preprint arXiv:2402.04401*, 2024.
- [15] Yancheng Wang, Ziyang Jiang, Zheng Chen, Fan Yang, Yingxue Zhou, Eunah Cho, Xing Fan, Xiaojiang Huang, Yanbin Lu, and Yingzhen Yang. Recmind: Large language model powered agent for recommendation. *arXiv preprint arXiv:2308.14296*, 2023.
- [16] Zhengyi Yang, Jiancan Wu, Zhicai Wang, Xiang Wang, Yancheng Yuan, and Xiangnan He. Generate what you prefer: Reshaping sequential recommendation via guided diffusion. *Advances in Neural Information Processing Systems*, 36, 2024.
- [17] Junjie Zhang, Yupeng Hou, Ruobing Xie, Wenqi Sun, Julian McAuley, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. Agentcf: Collaborative learning with autonomous language agents for recommender systems. In *Proceedings of the ACM on Web Conference 2024*, pages 3679–3689, 2024.

A Details on Models, Training and Dataset

We conduct experiments on 200 users with 100-200 interactions per user. For the llm baselines and single LoRA runs, our training set size was 5132 rows, creating by sampling a fixed percentage of rows at random per user. The test set in all runs was kept to be the same, where we have 7496 rows. For persona setup, the training dataset size varied from 2k-5k for each cluster. In the persona setup, we also split the aforementioned test file per cluster based on the user id. Thus, each persona LoRA was used to evaluate only the users within that cluster. Also, both train and test sets have the same 200 users but different interactions, where interactions were split temporally following a 60:40 train test split. For the distillation of MovieLens dataset, we use GPT-4o to generate M_s and M_l by passing batches of interactions (batch size was kept to be less than or equal to 10). We rejected generations that were longer than 2000 words to ensure we are able to use the smaller context length provided by Phi-3.

We cluster user profiles using KMeans++ to form $K = 4$ personas, determined by the Elbow method. To find the top- k nearest interactions, we use cosine similarity and append the enriched interaction text (from M_l) to the prompt and use $k = 1$ for all experiments. During fine-tuning, we train low-rank adapters for key and query matrices in the attention layer with a rank of 256, alpha of 32, and dropout of 0.1. We use a learning rate of 3×10^{-4} and train for 2 epochs using the AdamW optimizer. During inference, we prompt the model to generate 4 tokens, using a temperature of 0.3 and top-p set to 50 and look for the ground-truth using a regex search. To make comparison of θ_{single} and θ_{persona} fair, we subsample the dataset used to train θ_{single} (by interactions per user) so the dataset size resembles the size of the largest dataset used to train θ_{persona} . We conduct all our experiments on 8xA100 GPUs with 80 GB of GPU memory on each core. We use the torchtune library for fine-tuning and evaluation, OpenAI APIs for distillation and sklearn metrics.

B More on Persona

In Figure 2, we find that each persona cluster has a slightly different set of users, mainly by genre. For instance, persona 1 is all about method movies while persona 2 is about wars, battles and heroic stories. However, the token counts of these words are very low, indicating that these persona have overlaps as well.

In Figure 3, we observe that the performance for each persona group depends on the size of the training set used. Both RMSE and MAE on the test set decrease (improve) as the training size increases. This indicates that larger training datasets contribute to better model performance, reducing both RMSE and MAE values. Each number show the performance of the fine-tuned model of the persona cluster on their respective test sets.




Figure 2: Top-10 distinct words found in each persona. We compute this by obtaining all nouns and adjectives using the Spacy en_core_web_sm model, and finding one-vs-all set difference amongst the words in each cluster, retaining only top-10.

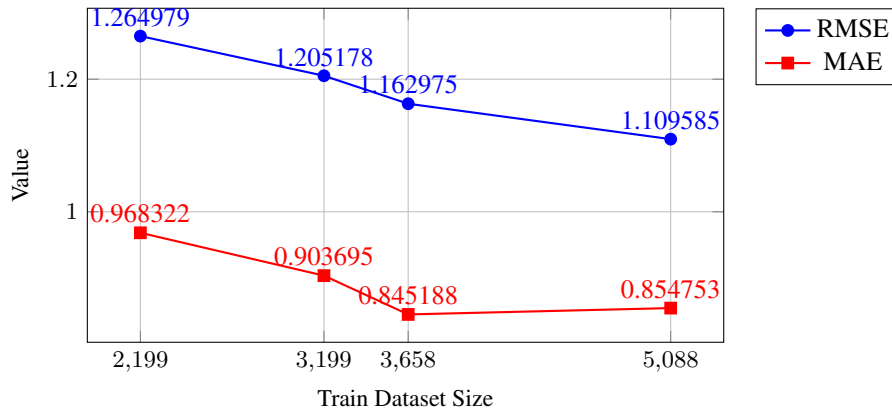


Figure 3: RMSE and MAE by Train Dataset Size

C Data Samples

Movie Rating prompt using Short Term Memory only

PERSONA: You excel at role-playing. Picture yourself as a user exploring a movie recommendation system. You have the following social traits:. Your activity trait is described as: An Occasional Viewer, seldom attracted by movie recommendations. Only curious about watching movies that strictly align the taste. The movie-watching habits are not very infrequent. And you tend to exit the recommender system if you have a few unsatisfied memories.. Your conformity trait is described as: A Balanced Evaluator who considers both historical ratings and personal preferences when giving ratings to movies. Sometimes give ratings that are different from historical rating.. Your diversity trait is described as: A Cinematic Trailblazer, a relentless seeker of the unique and the obscure in the world of movies . The movie choices are so diverse and avant-garde that they defy categorization.. Beyond that, your profile description expressing your preferences and dislikes is: I have a strong appreciation for movies that blend wit, depth, and compelling narratives. My highest ratings are often reserved for films that offer a mix of dark humor, psychological intrigue, and intense drama. I am particularly drawn to complex characters and intricate plots, as seen in my preference for thrillers and dramas

with strong performances by seasoned actors. I enjoy historical and action-packed films, especially those with a rebellious or heroic theme. My taste also includes a fondness for classic comedies and animated features that provide a nostalgic or heartwarming experience. I tend to favor movies with ensemble casts where each actor brings a unique element to the story. While I appreciate a variety of genres, I am less enthusiastic about films that lack depth or fail to engage me emotionally. Overall, my movie-watching habits reflect a desire for thought-provoking and emotionally resonant storytelling, with a particular interest in films that challenge conventional narratives and offer fresh perspectives. However, I have realized that I also appreciate visually stunning and action-packed crime thrillers, as well as heartwarming adventure films, even if they are in the children's genre. Additionally, I have a newfound appreciation for unique blends of genres, such as comedy and westerns, when they offer a fresh and engaging narrative. I also value films with strong character development and intricate plots, even if they are set in historical or war contexts, as long as they provide a compelling narrative and emotional depth. I have also come to appreciate older cinematic styles and classic films more than I initially thought, as long as they offer a compelling narrative and emotional depth. I also enjoy high-stakes, fast-paced narratives and unique storytelling approaches, as well as films with a strong visual and sci-fi element, provided they offer a compelling narrative..". The activity characteristic pertains to the frequency of your movie-watching habits. The conformity characteristic measures the degree to which your ratings are influenced by historical ratings. The diversity characteristic gauges your likelihood of watching movies that may not align with your usual taste.

TASK: Rate the movie given below on a likert scale from 1 to 5, where 1 means you hated it, 2 means you disliked it, 3 means you are neutral about it, 4 meaning you liked it and 5 means you absolutely loved it. Always respond with either one of 1, 2, 3, 4 or 5. Do not output anything else.

MOVIE: The movie Big Lebowski, The was released in the year 1998 is of Comedy, Crime, Mystery, Thriller genre with historical average rating of 3.74.

SUMMARY: A laid-back slacker gets caught up in a case of mistaken identity and embarks on a wild adventure involving bowling, kidnapping, and a rug that really tied the room together.. The cast of the movie is as follows: Jeff Bridges, John Goodman, Julianne Moore, Steve Buscemi, David Huddleston, Philip Seymour Hoffman, Tara Reid, Philip Moon, Mark Pellegrino, Peter Stormare, Flea, Torsten Voges, Jimmie Dale Gilmore, Jack Kehler, John Turturro

RATING:

Movie Rating prompt using both Short Term and Long Term Memories

PERSONA: You excel at role-playing. Picture yourself as a user exploring a movie recommendation system. You have the following social traits:. Your activity trait is described as: An Occasional Viewer, seldom attracted by movie recommendations. Only curious about watching movies that strictly align the taste. The movie-watching habits are not very infrequent. And you tend to exit the recommender system if you have a few unsatisfied memories.. Your conformity trait is described as: A Balanced Evaluator who considers both historical ratings and personal preferences when giving ratings to movies. Sometimes give ratings that are different from historical rating.. Your diversity trait is described as: A Cinematic Trailblazer, a relentless seeker of the unique and the obscure in the world of movies . The movie choices are so diverse and avant-garde that they defy categorization.. Beyond that, your profile description expressing your preferences and dislikes is: I have a strong appreciation for movies that blend wit, depth, and compelling narratives. My highest ratings are often reserved for films that offer a mix of dark humor, psychological intrigue, and intense drama. I am particularly drawn to complex characters and intricate plots, as seen in my preference for thrillers and dramas with strong performances by seasoned actors. I enjoy historical and action-packed films, especially those with a rebellious or heroic theme. My taste also includes a fondness for classic comedies and animated features that provide a nostalgic or heartwarming experience. I tend to favor movies with ensemble casts where each actor brings a unique element to the story. While I appreciate a variety of genres, I am less enthusiastic about films that lack depth or fail to engage me emotionally. Overall, my movie-watching habits reflect a desire for thought-provoking and emotionally resonant storytelling, with a particular interest in films that challenge conventional narratives and offer fresh perspectives. However, I have realized that I also appreciate visually stunning and action-packed crime thrillers, as well as heartwarming adventure films, even if they are in the children's genre. Additionally, I have a newfound appreciation for unique blends of genres, such as comedy and westerns, when they offer a fresh and engaging narrative. I also value films with strong character development and intricate plots, even if they are set in historical or war contexts, as long as they provide a compelling narrative and emotional depth. I have also come to appreciate older cinematic styles and classic films more than I initially thought, as long as they offer a compelling narrative and emotional depth. I also enjoy high-stakes, fast-paced narratives and unique storytelling approaches, as well as films with a strong visual and sci-fi element, provided they offer a compelling narrative..". The activity characteristic pertains to the frequency of your movie-watching habits. The conformity characteristic measures the degree to which your ratings are influenced by historical ratings. The diversity characteristic gauges your likelihood of watching movies that may not align with your usual taste.

Given your persona and memory of some of movies you watched in the past, think carefully and rate the movie given in the end. Always use your memory at your own discretion as not everything is helpful. Also, historical average ratings of the movie mean an average rating of how other people have rated it, not you.

YOUR MEMORIES

MOVIE: The movie Falling Down was released in the year 1993 is of Action, Drama genre with historical average rating of 3.45.

MY MEMORY: A man's descent into madness and violence as he navigates through the frustrations of modern life.. The cast of the movie is as follows: Michael Douglas, Robert Duvall, Barbara Hershey, Rachel Ticotin, Tuesday Weld, Frederic Forrest, Lois Smith, Joey Singer, Ebbe Roe Smith, Michael Paul Chan, Raymond J. Barry, D.W. Moffett, Steve Park, Kimberly Scott, James Keane. I rated movie Falling Down (1993) as 4 because it offers a compelling narrative of a man's descent into madness and violence, aligning with my appreciation for psychological intrigue and intense drama. The strong performances by Michael Douglas and Robert Duvall, along with the thought-provoking and emotionally resonant storytelling, make it a highly engaging film.

TASK: Rate the movie given below on a likert scale from 1 to 5, where 1 means you hated it, 2 means you disliked it, 3 means you are neutral about it, 4 meaning you liked it and 5 means you absolutely loved it. Always respond with either one of 1, 2, 3, 4 or 5. Do not output anything else.

MOVIE: The movie Big Lebowski, The was released in the year 1998 is of Comedy, Crime, Mystery, Thriller genre with historical average rating of 3.74.

SUMMARY: A laid-back slacker gets caught up in a case of mistaken identity and embarks on a wild adventure involving bowling, kidnapping, and a rug that really tied the room together.. The cast of the movie is as follows: Jeff Bridges, John Goodman, Julianne Moore, Steve Buscemi, David Huddleston, Philip Seymour Hoffman, Tara Reid, Philip Moon, Mark Pellegrino, Peter Stormare, Flea, Torsten Voges, Jimmie Dale Gilmore, Jack Kehler, John Turturro

RATING: