EDHUP: Evolving Dynamic Hierarchical User Profiling for Personalized Large Language Models

Anonymous ACL submission

Abstract

Personalized user profiling is increasingly critical for aligning large language models (LLMs) with complex user behaviors and enhancing user experience. However, existing profiling methods are often constrained by shallow labeling schemes or static key-value storage, limiting their ability to capture the rich, dynamically evolving structures and psychological motivations underlying complex user behaviors. To address these limitations, we propose a novel framework, EDHUP (Evolving Dynamic Hierarchical User Profiling), which integrates a dynamic update mechanism based on Bayesian inference with a hierarchical chain-of-thought reasoning process. This framework enables the generation of adaptive user profiles at multiple granularities. Experimental results demonstrate that EDHUP consistently outperforms baseline methods across diverse personalization tasks, substantially enhancing the effectiveness of downstream applications.

1 Introduction

The rapid advancement of large language models (LLMs) has demonstrated significant potential across diverse domains, including education (Wang et al., 2025c), healthcare (Chen et al., 2026), and intelligent assistants (Dam et al., 2024), exhibiting remarkable capabilities in understanding and generating human-like responses. As LLMs become increasingly integrated into everyday applications, users increasingly expect them to produce personalized responses that align with individual preferences, habits, and contextual scenarios, rather than providing generic, one-size-fits-all outputs. To address these demands, personalized LLM approaches (Purificato et al., 2024; Liu et al., 2025; Freitas et al., 2026) have been proposed to generate tailored responses, wherein user profiling plays a pivotal role. Accurate user profiles enable LLMs to capture individual characteristics, predict preferences, and produce responses that are well-aligned with each user's specific context. User profiling involves extracting distinctive user attributes from historical interaction data, including information such as age, occupation, hobbies, and preferences.

Traditional user profiling techniques (Godoy and Amandi, 2005; Al-Shamri, 2016) primarily rely on rule-based extraction or statistical machine learning methods and face significant limitations. These approaches require extensive manual feature engineering for each specific task and often struggle to capture the full spectrum of user profile characteristics. Recently, large language models (LLMs) have demonstrated substantial potential in user profiling (Wang et al., 2025a; Jiang et al., 2025; Zollo et al., 2025), enabling more accurate user modeling and profile generation across diverse tasks. Nevertheless, the application of LLMs for user modeling in real-world scenarios continues to face several challenges:

- Existing approaches generate tags with fixed formats and dimensions, lacking sufficient flexibility across different scenarios.
- Current methods often neglect the deep structural modeling of user profiles and fail to thoroughly mine user behaviors, limiting their ability to capture behavioral diversity and depth.
- User interests are dynamic, varying across contexts and over time. Consequently, user profiles generated by LLMs should adapt dynamically to temporal and contextual changes to enhance user experience.

To address these challenges, we propose EDHUP (Evolving Dynamic Hierarchical User Profiling), a dynamic hierarchical self-evolving user profiling framework. To accurately extract multi-level features from unstructured interaction data, we introduce the Hierarchical Chain-of-Thought

(H-CoT) technique for user profiling. H-CoT leverages information entropy theory within a hierarchical reasoning framework to guide the generation of user profiles at varying granularities. Importantly, to address the temporal fading and scenario-dependent shifts of user interests, we design a dynamic user profile updating mechanism based on Bayesian inference. By incorporating a confidence metric and the forgetting curve, the generated profiles can self-evolve and undergo a form of natural selection analogous to human memory. To evaluate the effectiveness of EDHUP, we compare it against several state-of-the-art user profile generation methods across tasks with evolving temporal and contextual scenarios. Experimental results demonstrate that EDHUP consistently outperforms baseline methods, providing interpretable, multi-granular, and dynamically adaptive user profile representations. Furthermore, to assess the practical utility of EDHUP-generated user profiles in downstream applications, we conduct experiments on two tasks: sequence recommendation and personalized conversational assistance. Results indicate that the profiles generated by EDHUP significantly enhance the personalization performance of these downstream tasks.

2 Method

Problem Setting. The objective of user profiling is to accurately construct user representations $\mathcal{Z}$ based on users' historical interaction data across diverse tasks, temporal periods, and contextual scenarios, denoted as $\mathcal{H}_u = x_1, x_2, \cdots, x_T$. We conceptualize the user representation $\mathcal{Z}$ as a set of latent variables underlying the observed interaction history $\mathcal{H}_u$. Accordingly, the goal of user profiling is to infer these latent variables by maximizing the posterior probability $P(\mathcal{Z} \mid \mathcal{H}_u)$. To address this problem, we propose a dynamic hierarchical self-evolving framework called EDHUP (Evolving Dynamic Hierarchical User Profiling) for user profiling. We first present the overall framework in Section 3.1, then introduce the Hierarchical Chain-of-Thought (H-CoT) technique for constructing multi-level user profiles in Section 3.2, and finally describe the Bayesian-driven profile updating mechanism in Section 3.3.

2.1 EDHUP Framework

The overall framework of our proposed EDHUP method is illustrated in Figure 1. To extract deep-level features of user profiles, and inspired by Maslow's Hierarchy of Needs (Huitt, 2007) and the Big Five Personality Traits (Roccas et al., 2002), we decompose the user profile $\mathcal{U}$ into three interconnected yet hierarchically distinct layers: the Basic Attribute layer $\mathcal{A}$, the Behavioral Characteristic layer $\mathcal{B}$, and the Psychological Motivation layer $\mathcal{P}$. This hierarchical decomposition reflects the observation that human cognition and behavior are interdependent and jointly shaped by underlying physiological attributes, intermediate behavioral patterns, and higher-order psychological motivations (Huitt, 2007). By constructing this multi-level architecture, EDHUP achieves a comprehensive understanding of users, providing a robust cognitive foundation for generating personalized and contextually appropriate responses. To accurately extract multi-level features from unstructured interaction data, we propose the Hierarchical Chain-of-Thought (H-CoT) reasoning technique for constructing multi-level user profiles. The reasoning process is organized into three logical hierarchies:

- **Basic Attribute Layer**: This foundational layer captures factual and relatively static information about users, such as occupation, age, and gender. Reasoning at this level relies on deductive inference and entity extraction.

- **Behavioral Characteristic Layer**: This layer represents dynamic aspects of user behavior, including interests, interaction preferences, and daily habits. To model diverse behavioral preferences across different granularities, we structure this layer as a recursive tree that allows infinite subdivision from coarse-grained interests to fine-grained sub-dimensions. Reasoning at this layer leverages inductive inference and hierarchical clustering.

- **Psychological Motivation Layer**: Positioned at the top of the hierarchy, this layer captures users' values, personality traits, and potential needs. Reasoning here involves abductive inference, which enables the model to infer the most likely psychological causes and underlying needs behind observed behaviors by utilizing commonsense knowledge and the causal reasoning capabilities of LLMs.

To address the challenges of user interests decaying over time and migrating across contexts, we design a dynamic evolution mechanism based on

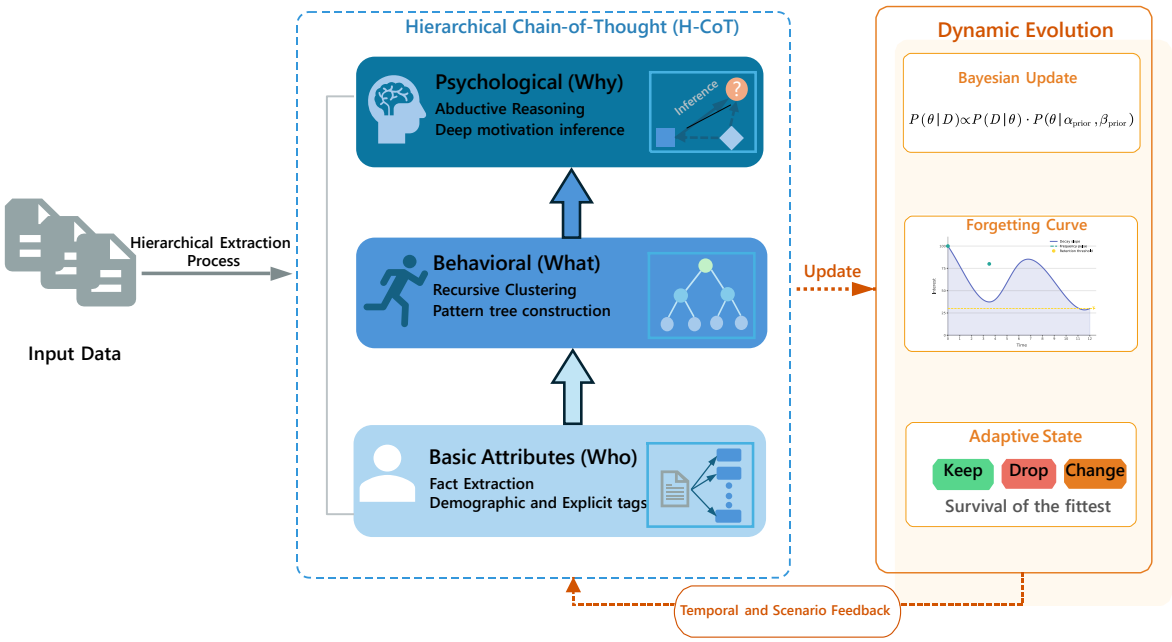

Figure 1: The overall framework of our proposed EDHUP framework.

Bayesian inference. By integrating Bayesian confidence metrics with forgetting curves, user profiles can self-evolve and undergo a natural selection process analogous to human memory. Through the combination of H-CoT reasoning and Bayesian-driven profile updating, EDHUP addresses two core challenges: the multi-level granularity mining of user profiles and their dynamic updating across contexts and temporal scales.

2.2 Hierarchical Chain-of-Thought for User Profiling

The Hierarchical Chain-of-Thought (H-CoT) formalizes the reasoning process as a search over a hierarchical tree, rather than the sequential generation used in traditional Chain-of-Thought (CoT) approaches. In H-CoT, each "thought" transcends its role as a mere intermediate computational step and instead represents a candidate node within the hierarchical structure of the user profile. Unlike conventional CoT, which focuses on a linear reasoning path, H-CoT simulates cognitive processes inspired by human cognition, progressing from surface to deeper layers and from concrete to abstract concepts, through a recursive hierarchical reasoning architecture.

H-CoT leverages the hierarchical reasoning capabilities of large language models (LLMs) to generate multi-level user profiles by guiding the model to reason progressively: from fact extraction to behavior induction and finally to motivation inference. The core idea of H-CoT is to model the user profile construction process as the decomposition of a joint probability distribution. For a complete user profile $\mathcal{Z} = Z^{(0)}, Z^{(1)}, Z^{(2)}$, the generation probability is formulated as:

$$P(\mathcal{Z}|\mathcal{H}_u) = P(Z^{(0)}|\mathcal{H}_u) \cdot P(Z^{(1)}|Z^{(0)}, \mathcal{H}_u) \cdot P(Z^{(2)}|Z^{(1)}, Z^{(0)}, \mathcal{H}_u), \quad (1)$$

where $Z^{(0)}$, $Z^{(1)}$, and $Z^{(2)}$ correspond to the basic attributes, behavioral characteristics, and psychological motivations of user profiles, respectively. This formulation explicitly captures the hierarchical dependencies within H-CoT: high-level mental inferences $Z^{(2)}$ are conditioned on middle-level behavioral features $Z^{(1)}$ and low-level basic attributes $Z^{(0)}$. As such, the overall user profile $\mathcal{Z}$ constitutes a structured, multi-layer object rather than a flat textual summary. This hierarchical generation mechanism mitigates the hallucination problem often encountered in direct user profile generation, as each reasoning layer is grounded in the outputs of the preceding layer. Furthermore, it enables LLMs to deeply infer latent user needs and personality traits from explicit historical interaction data, providing richer and more interpretable user representations.

Basic Attribute Reasoning. The Basic Attribute Layer forms the foundation of user profiles, encompassing demographic and static information. Information at this layer is typically explicit, factual, and

relatively stable. Within H-CoT, reasoning at this layer primarily involves information extraction and conflict resolution. This requires the LLM to scan user interaction histories and identify explicit entity information. We design a schema-constrained CoT prompt to extract underlying basic attribute information. Unlike open-ended generation, this prompt enforces reasoning within a predefined attribute schema, including fields such as age, gender, location, occupation, and device. The extracted basic attributes can be formally represented as:

$$Z^{(0)} = (\text{Att}_i, \text{Value}_i, \text{Conf}_i), , i = 1, \cdots, n, \quad (2)$$

where Att_i denotes a specific basic attribute, Value_i represents the corresponding feature value extracted by the LLM, and Conf_i reflects the model's confidence in the extracted feature, capturing uncertainty in attribute inference. Although basic attributes are generally stable, they remain updatable variables within the EDHUP framework. For instance, a user's place of residence may change due to job relocation, and such updates are captured through the Bayesian-driven profile updating mechanism discussed in Section 2.3.

Behavioral Characteristic Reasoning. To extract multi-level behavioral features, inspired by HERCULES (Petnehazi and Aradi, 2025) and TaxoAdapt (Kargupta et al., 2025) utilizing LLMs for hierarchical text clustering, we employ LLMs to recursively cluster and generalize behavioral features by integrating historical context with the foundational attributes extracted in the previous layer. In this layer, H-CoT performs inductive reasoning on user behavior patterns through hierarchical recursive clustering, generating multi-level behavioral profiles at varying granularities. Its primary objective is to map the user interaction history $\mathcal{H}_u$ into a hierarchical interest space $Z^{(1)}$ and perform fine-grained segmentation of behaviors. The behavioral characteristic layer is structured as a dynamically growing taxonomy tree rather than a flat list of tags. The root node corresponds to the user's historical interactions combined with the foundational profile from the basic attribute layer. Hierarchical recursive clustering is then applied based on behavioral feature patterns. At each level, we define a recursive operator:

$$\mathcal{R} = \begin{cases} \Phi_\theta\Big(\text{Pm}(z_{\text{par}}, \mathcal{H}_{\text{sub}}) \cup \bigcup_k \mathcal{R}(z^{(k)}_{\text{child}}, \mathcal{H}^{(k)}_{\text{sub}})\Big), \\ \emptyset, \qquad\qquad \text{If } \text{Stop}(z_{\text{par}}, \mathcal{H}_{\text{sub}}), \end{cases} \quad (3)$$

where Φ_θ denotes the LLM performing hierarchical clustering, $\text{Pm}(\cdot)$ is the dynamically constructed prompt template at the current level, and $\mathcal{H}^{(k)}_{\text{sub}}$ is the subset of the user's interaction history semantically relevant to the child node $z^{(k)}_{\text{child}}$. z_{par} represents the parent node. $\text{Stop}(z_{\text{par}}, \mathcal{H}_{\text{sub}})$ defines the stopping condition for recursive clustering. To determine when to terminate recursive subdivision, we introduce an adaptive stopping mechanism based on semantic entropy (Farquhar et al., 2024). For a candidate node z, the LLM generates N potential segmentation descriptions $d_1, \ldots, d_N$, which are clustered into K semantic clusters $C_1, \ldots, C_K$. The semantic entropy of the node is defined as:

$$H_{\text{sem}}(z) = -\sum_{k=1}^{K} \hat{p}(C_k) \log \hat{p}(C_k), \; \hat{p}(C_k) = \frac{|C_k|}{N}. \quad (4)$$

If $H_{\text{sem}}(z) > \varepsilon_{\text{high}}$, the feature exhibits high uncertainty or diversity, necessitating further recursive subdivision. Conversely, if $H_{\text{sem}}(z) < \varepsilon_{\text{low}}$, the feature has converged to a concrete, atomic-level concept, and recursion terminates.

Psychological Motivation Reasoning. The Psychological Motivation Layer captures implicit user traits that cannot be directly observed from interaction histories. At this layer, H-CoT employs psychological knowledge and causal inference to uncover latent motivations, including core needs, personality traits, and values, by leveraging both historical interaction behaviors and foundational attributes. The inputs to this layer include the historical interaction context $\mathcal{H}_u$, basic attributes $Z^{(0)}$, and behavioral characteristics $Z^{(1)}$. The goal is to infer latent psychological traits:

$$Z^{(2)} = (\text{Trait}_k, \text{Score}_k, \text{EvidenceChain}_k), \quad (5)$$

where Trait_k represents the inferred psychological trait, Score_k denotes the model's confidence in this trait, and EvidenceChain_k is a causal chain illustrating the inferred reasoning, e.g., Psych: High risk-taking $\xrightarrow{\text{causes}}$ Behavior: extreme skiing. This approach outputs both psychological labels and cross-level causal evidence chains, enabling interpretable and generalizable user profiles in novel scenarios.

Since the true posterior $P(\mathcal{Z} \mid \mathcal{H}_u)$ is intractable, H-CoT approximates it using a variational distribution $Q_\phi(\mathcal{Z} \mid \mathcal{H}_u)$. The decomposed form of Q_ϕ

is:

$$Q_\phi(\mathcal{Z} \mid \mathcal{H}_u) \approx q_{\phi_0}(Z^{(0)} \mid \mathcal{H}_u) \cdot q_{\phi_1}(Z^{(1)} \mid Z^{(0)}, \mathcal{H}_u) \cdot q_{\phi_2}(Z^{(2)} \mid Z^{(1)}, Z^{(0)}, \mathcal{H}_u), \quad (6)$$

where q_{ϕ_l} denotes the LLM's approximation to the posterior at level l. By generating the chain of thought hierarchically, H-CoT effectively maximizes the Evidence Lower Bound (ELBO):

$$\mathcal{L}_{\text{ELBO}}(\phi) = \mathbb{E}_{Q_\phi}\left[\log \frac{P(\mathcal{H}_u, \mathcal{Z})}{Q_\phi(\mathcal{Z} \mid \mathcal{H}_u)}\right]. \quad (7)$$

Theorem 2.1. *Let $\mathcal{M}$ be a large language model parameterized by ϕ. Given an observed user interaction history $\mathcal{H}u$ and a latent user profile variable set $\mathcal{Z} = Z^{(0)}, Z^{(1)}, Z^{(2)}$ structured as a Bayesian network, the H-CoT generation process is equivalent to maximizing the ELBO $\mathcal{L}_{ELBO}(\phi)$ of the marginal likelihood $P(\mathcal{H}_u)$.*

The proof is provided in the Appendix A.2. H-CoT thus constitutes a structured variational inference procedure rather than a heuristic for interpretability: it optimizes $\mathcal{L}_{\text{ELBO}}(\phi)$, minimizing $\mathbb{D}_{KL}(Q_\phi || P)$ and ensuring that the generated user profile $\mathcal{Z}$ closely approximates the true latent characteristics implied by the interaction history $\mathcal{H}_u$.

2.3 Bayesian-Driven Self-Evolving Update Mechanism for Dynamic Profiling

In real-world scenarios, user profiles are inherently dynamic. User interests may evolve over time and vary across contexts. For example, a user may frequently search for home furnishing information during a renovation period, but such interest should rapidly decay once the renovation is completed. If a recommendation system continues to suggest furniture-related items afterward, the user experience may be substantially degraded. To address this issue, we propose a dynamic evolution mechanism based on Bayesian inference, which incorporates confidence estimation and a forgetting mechanism to enable the self-updating of user profiles.

We model each feature z in a user profile as a random variable whose uncertainty is quantified by a confidence score C. Under the Bayesian framework, the confidence score is assumed to follow a Beta distribution, i.e., $C \sim \text{Beta}(\alpha, \beta)$, where α denotes the accumulated positive evidence, namely the number of times the feature has been confirmed or activated, and β denotes the accumulated negative evidence or uncertainty, namely the number of times the feature has been negated or remained inactive. Accordingly, the expected confidence of this feature is computed as:

$$\mathbb{E}[C] = \frac{\alpha}{\alpha + \beta}. \quad (8)$$

This expectation reflects the model's current belief regarding the reliability of the feature under the given temporal and contextual conditions.

Dynamic Forgetting Mechanism in the Temporal Dimension. To simulate the natural decay of human memory, we introduce a dynamic forgetting mechanism based on Ebbinghaus' forgetting curve (Murre and Dros, 2015) to model the temporal evolution of user profiles. Let t_{last} denote the last update time of a profile feature and t_{now} the current time, with the time interval defined as $\Delta t = t_{\text{now}} - t_{\text{last}}$. Prior to any update, the confidence of profile features decays as a function of this time interval. Within the Beta feature distribution framework, the hyperparameters are updated as:

$$\alpha_{\text{decayed}} = \alpha_{\text{old}} \cdot e^{-\lambda \cdot \Delta t}, \ \beta_{\text{decayed}} = \beta_{\text{old}} \cdot e^{-\lambda \cdot \Delta t},$$

where α_{old} and β_{old} denote the cumulative positive and negative evidence at the last update, respectively, α_{decayed} and β_{decayed} represent their decayed values at the current time, and $\lambda \in [0, 1]$ is the decay rate. Different feature types are assigned distinct decay rates. For example, the occupation basic attribute is relatively stable and thus has a small λ, while the emotional state psychological motivation tend to fluctuate and thus has a large λ. This decay mechanism ensures that if a profile feature remains inactive for an extended period, its α value gradually approaches zero, leading to a decrease in confidence $E[C]$. When the confidence falls below a predefined threshold $C_{\min}$, the feature is automatically pruned from the user profile.

Scenario-aware Bayesian Updating Strategy. User behaviors often vary significantly across different scenarios. For instance, a user's interaction style may be formal and rigorous in a work context, whereas it becomes relaxed and humorous in daily life. To dynamically adapt to such scenario-dependent behaviors, EDHUP incorporates a scenario-aware Bayesian updating mechanism, which adjusts feature confidence before and after scene changes and leverages the frequency of

profile feature occurrences within specific scenarios to strengthen or weaken confidence. Let S denote the user's current interactive environment (e.g., working mode, chatting mode, shopping mode). When a profile feature z is observed under a particular scenario S, this observation is treated as a Bayesian experiment. Denote by $\text{Occ}_S(z)$ the occurrence frequency of feature z in scenario S, with higher frequencies representing stronger evidence. The evidence strength L_z is defined as:

$$L_z = \log(1 + \text{Occ}_S(z)) \cdot W_{\text{scene}}, \tag{9}$$

where $W_{\text{scene}} \in [0, 1]$ is a scene-specific weight reflecting the importance of the scenario. To detect scene changes in real time, we construct a scene classifier using a specialized prompt for LLMs. If the current scene is consistent with the original scene associated with the profile feature, the positive evidence α is reinforced:

$$\alpha_{\text{new}} = \alpha_{\text{decayed}} + L_z, \ \beta_{\text{new}} = \beta_{\text{decayed}}.$$

Conversely, if the current scene conflicts with the profile feature (e.g., exhibiting casual behavior in a formal work setting), the uncertainty β is increased:

$$\alpha_{\text{new}} = \alpha_{\text{decayed}}, \ \beta_{\text{new}} = \beta_{\text{decayed}} + L_z.$$

This scenario-aware updating mechanism enables EDHUP to generate accurate and context-sensitive user profiles by dynamically adjusting the confidence of features according to both temporal and contextual evidence.

Confidence-based Decision Updating. After each interaction, we compute the average confidence change before and after updating a profile feature. The confidence change is defined as:

$$\Delta C = \mathbb{E}[C_{\text{new}}] - \mathbb{E}[C_{\text{old}}]. \tag{10}$$

If $\Delta C > \delta_{\text{pos}}$, where δ_{pos} denotes the positive reinforcement threshold, the corresponding profile feature is strengthened and retained as an active feature. If $\Delta C < \delta_{\text{neg}}$, where δ_{neg} denotes the negative attenuation threshold, the feature is weakened and treated as inactive. Furthermore, if $\mathbb{E}[C_{\text{new}}] < \tau_{\text{prune}}$, where τ_{prune} is the pruning threshold, the feature is removed from the user profile.

2.4 Experiments and Discussion

To demonstrate the effectiveness and superior self-evolution capability of the proposed EDHUP method, we compare it with various profile generation approaches across different scenarios and tasks. Specifically, we evaluate whether the profiles generated by EDHUP can benefit downstream applications in real-world settings. To this end, we select two representative downstream tasks: sequence recommendation, described in Section 2.4.2, and personalized conversational assistance, described in Section 2.4.3.

2.4.1 Experiments on the Wiki People Dataset

To evaluate user profile modeling techniques in dynamic settings, Prottasha et al. (2025) construct two high-quality benchmark datasets for user profile generating and updating using the Wiki People dataset. It comprises 42,786 user biographies and a wide range of profile attributes, encompassing personal details, life events, achievements, and other relevant information. For a comprehensive description of this benchmark, we refer readers to Prottasha et al. (2025). To simulate evolving user preferences, the benchmark introduces modifications to existing user profiles based on updated biographical information, reflecting real-world changes such as shifting interests or relocation. Following this protocol, we apply the proposed EDHUP method to construct and update user profiles on this benchmark and compare its performance with several popular LLMs, including GPT-5.2 (Singh et al., 2025), Qwen3.5 (Yang et al., 2025), and DeepSeek v3.2 (DeepSeek-AI, 2025). In addition, we conduct ablation studies to evaluate the contributions of different components of EDHUP: one variant using only the Chain-of-Thought reasoning (EDUP) and another employing only a static template (EUP). Implementation details and experimental settings for all methods are summarized in the Appendix A.4.1. Table 1 reports the performance of different user profile generation methods across various LLMs on the Wiki People benchmark. The results demonstrate that EDHUP generates context-aware and highly accurate user profiles while effectively adapting to evolving user behaviors. In contrast, static template-based methods such as EUP fail to flexibly accommodate such dynamic changes.

2.4.2 Experiments on the Amazon Reviews Dataset

Given a user's historical interaction sequence $h_1, \cdots, h_l$, the goal of sequential recommendation is to predict the item h_{l+1} that the user is likely to

Table 1: The performance of different profile generation methods based on LLMs in the Wiki People benchmark.

Method	Profile Construction						Profile Updating					
	GPT 5.2		Qwen 3.5		Deepseek 3.2		GPT 5.2		Qwen 3.5		Deepseek 3.2	
	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1
Nusrat	95.62	93.24	94.80	91.52	92.09	88.63	96.31	94.89	95.47	93.59	92.27	89.14
EUP	95.87	94.21	93.74	90.29	93.26	90.08	94.19	92.83	92.22	90.67	92.86	91.07
EDUP	97.31	96.59	97.35	96.64	96.91	94.52	97.63	95.59	98.12	96.05	97.23	94.89
EDHUP	**98.31**	**97.87**	**98.46**	**96.84**	**97.22**	**95.89**	**99.85**	**98.42**	**99.28**	**98.10**	**98.67**	**96.29**

be interested in during the next interaction round (Chen et al., 2023), where l denotes the length of the historical interaction sequence. In this work, we focus on product recommendation and adopt the large-scale Amazon Reviews 2023 dataset (Hou et al., 2024), which contains over 570 million reviews and 48 million items spanning 33 product categories. We select two representative domains, namely `Baby_Products` and `Movies_and_TV`. To demonstrate the effectiveness of the proposed EDHUP method, we compare it with several representative profile generation methods for sequential recommendation, including RLMRec (Ren et al., 2024), PALR (Zhao et al., 2024), and LettinGo (Wang et al., 2025a). To fairly evaluate the contribution of profile generation to recommendation performance, we also include KAR (Xi et al., 2024) as a basic baseline, which makes predictions solely based on users' historical interaction records. Following prior work, we adopt two widely used classification metrics, accuracy and weighted F1 score, to evaluate how effectively the generated user profiles support sequential prediction. The implementation details and experimental settings are summarized in Appendix A.4.2.

Figure 2 and Table 2 present the performance of different profile generation methods over the last five interaction iterations on the `Baby_Products` and `Movies_and_TV` domains of the Amazon Reviews 2023 dataset. The results show that EDHUP consistently outperforms existing profile generation methods across different domains. By generating self-evolving user profiles, EDHUP dynamically adapts to iterative changes in historical interaction data and continuously improves the prediction accuracy of downstream recommendation tasks. In contrast, other profile generation methods, such as PALR and RLMRec, exhibit noticeable fluctuations in prediction accuracy as the historical interaction sequence evolves, indicating their limited ability to adapt user profiles to changing historical contexts. Interestingly, KAR, which relies solely on historical interaction information, shows a consistent decline in accuracy as more historical context is incorporated. This may be because longer interaction histories introduce contextual conflicts and noisy preference signals, making the model more susceptible to confusion and reducing prediction accuracy. These findings further highlight the importance of dynamically adaptive user profiles for improving downstream sequential recommendation performance.

2.4.3 Experiments on the PERSONAMEM Benchmark

Recently, personalized AI assistants (Singh et al., 2025; Yang et al., 2025) have emerged as one of the most prominent applications of large language models (LLMs) worldwide. These systems deliver personalized responses and address diverse tasks by aligning with users' current contexts and preferences. Effectively doing so requires the system to internalize users' inherent traits and track the evolution of their preferences across time and scenarios. To evaluate the capability of personalized assistants in generating adaptive responses with dynamic user profiles, Jiang et al. (2025) propose the PERSONAMEM benchmark, which comprises 180 simulated user–LLM interaction histories. Since user characteristics change over time as different events occur within each task scenario, the length of historical interactions is critical for the accurate evolution of user profiles. Following the experimental setup of Jiang et al. (2025), we consider different lengths of historical interaction sequences (i.e., the most recent 30 and 50 interactions) to generate user profiles. We then combine the generated user profiles with seven types of user queries in the PERSONAMEM benchmark—recalling user-shared facts, suggesting new ideas, acknowledging latest preferences, tracking full preference evolution, revisiting reasons behind preference updates, providing preference-aligned recommendations, and general-

Table 2: The performance of different profile generation methods in the last five interaction iterations for downstream sequence recommendation tasks on the Amazon Reviews Dataset.

Method	Baby products										Movies and TV									
	S1		S2		S3		S4		S5		S1		S2		S3		S4		S5	
	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1	acc	F1
KAR	82.96	87.57	82.43	86.62	81.80	85.52	81.43	84.73	80.63	83.42	64.47	68.09	63.85	67.94	63.68	67.27	62.44	66.23	62.17	66.02
RLMRec	80.35	84.72	81.19	84.92	84.17	88.89	83.24	87.18	82.57	86.26	68.90	73.84	69.20	74.15	70.68	75.32	69.71	74.79	68.82	72.61
PALR	84.95	88.75	83.43	87.52	82.60	85.25	84.37	87.18	85.24	89.03	67.24	70.48	68.27	72.83	68.96	73.75	70.19	74.63	69.41	74.08
LettinGo	86.21	89.10	90.24	93.35	92.59	94.74	93.33	96.29	90.63	93.57	70.11	75.35	71.89	76.33	72.62	77.46	72.89	78.16	73.26	78.39
EDHUP	**90.67**	**95.18**	**92.89**	**96.14**	**93.10**	**97.63**	**95.42**	**98.07**	**96.56**	**98.95**	**72.49**	**78.84**	**73.16**	**79.92**	**74.43**	**80.62**	**75.89**	**82.98**	**76.14**	**81.17**

Table 3: The mean accuracy of different profile generation methods with various lengths of historical interactions for downstream personalized conversational assistant tasks on the PERSONAMEM benchmark, where 30H and 50H refers to the last 30 and 50 interactions respectively.

History	Method	Recall user shared facts	Suggest new ideas	Track full preference evolution	Revisit reasons behind preference updates	Acknowledge latest user preferences	Provide preference-aligned recommendations	Generalize to new scenarios
30H	Mem0	43.62	24.18	56.75	82.09	39.81	57.78	62.24
	A-Mem	65.38	39.70	57.03	86.17	45.65	71.03	64.80
	Memory OS	82.95	26.28	64.69	**88.57**	59.24	74.78	63.33
	O-Mem	78.92	35.31	73.90	80.72	62.22	78.29	76.16
	BMAM	82.25	**45.09**	69.57	92.88	58.89	71.43	**80.98**
	EDHUP	**85.47**	43.82	**76.12**	84.39	**69.62**	**82.36**	78.21
50H	Mem0	42.30	27.92	51.36	80.52	36.33	54.71	61.24
	A-Mem	61.48	43.22	58.09	85.92	46.50	70.67	62.41
	Memory OS	84.25	29.08	61.39	85.40	60.86	75.65	63.89
	O-Mem	75.13	45.52	76.30	82.13	67.72	79.50	81.64
	BMAM	86.22	47.63	70.28	93.05	58.40	70.67	**83.71**
	EDHUP	**89.36**	**48.91**	**81.29**	**95.68**	**70.95**	**83.60**	82.49

izing to new scenarios—to generate personalized responses across different times and contexts.

To demonstrate the superiority of the proposed EDHUP method, we compare it with several popular profile generation methods for personalized assistants, including Mem0 (Chhikara et al., 2025), A-Mem (Xu et al., 2025), Memory OS (Li et al., 2025), O-Mem (Wang et al., 2025b), and BMAM (Li et al., 2026). Implementation details and experimental settings for all methods are summarized in Appendix A.4.3. Table 3 reports the performance of different profile generation methods under varying historical interaction lengths on downstream personalized conversational assistant tasks within the PERSONAMEM benchmark. The results indicate that EDHUP consistently outperforms other profile generation methods across most tasks. Moreover, as the context length increases, EDHUP generates more adaptive profiles and produces more accurate responses in six out of the seven query types. These findings further confirm that EDHUP effectively constructs self-evolving user profiles capable of adapting to extended interaction histories and dynamically changing contexts.

3 Conclusion

In this paper, we propose EDHUP, an Evolving Dynamic Hierarchical User Profiling framework for personalized large language models. Compared to existing profile generation approaches, EDHUP generates adaptive user profiles in a structured manner through the Hierarchical Chain-of-Thought strategy. Furthermore, it produces dynamic user profiles that evolve over time via a Bayesian-driven self-updating mechanism. Experimental results demonstrate that EDHUP constructs self-evolving profiles that effectively adapt to diverse contexts and downstream tasks.

4 Limitations

our proposed EDHUP framework is primarily designed for personalized LLM applications and may not generalize seamlessly to emerging agent-based systems, such as OpenClaw (Steinberger and community, 2026) and Hermes (Nous Research, 2026) agents. In future work, we aim to extend EDHUP to a broader range of applications, including real-time personalized agents, and to further enhance its adaptability and robustness in long-horizon personalized interaction scenarios.

References

Mohammad Yahya H Al-Shamri. 2016. User profiling approaches for demographic recommender systems. *Knowledge-Based Systems*, 100:175–187.

Haonan Bian, Zhiyuan Yao, Sen Hu, Zishan Xu, Shaolei Zhang, Yifu Guo, Ziliang Yang, Xueran Han, Huacan Wang, and Ronghao Chen. 2026. RealMem: Benchmarking LLMs in real-world memory-driven interaction. *arXiv preprint arXiv:2601.06966*.

Sully F Chen, Anton Alyakin, Andreas Seas, Eunice Yang, Joanne J Choi, Jin Vivian Lee, Amelia L Chen, Pranav I Warman, Rochelle T Bitolas, Robert J Steele, and 1 others. 2026. Llm-assisted systematic review of large language models in clinical medicine. *Nature medicine*, pages 1–8.

Ting Chen, Wei-Li Han, Hai-Dong Wang, Yi-Xun Zhou, Bin Xu, and Bin-Yu Zang. 2007. Content recommendation system based on private dynamic user profile. In *International Conference on Machine Learning and Cybernetics*, volume 4, pages 2112–2118.

Xiaoqing Chen, Zhitao Li, Weike Pan, and Zhong Ming. 2023. A survey on multi-behavior sequential recommendation. *arXiv preprint arXiv:2308.15701*.

Prateek Chhikara, Dev Khant, Saket Aryan, Taranjeet Singh, and Deshraj Yadav. 2025. Mem0: Building production-ready ai agents with scalable long-term memory. *arXiv preprint arXiv:2504.19413*.

Sumit Kumar Dam, Choong Seon Hong, Yu Qiao, and Chaoning Zhang. 2024. A complete survey on llm-based ai chatbots. *arXiv preprint arXiv:2406.16937*.

DeepSeek-AI. 2025. Deepseek-v3.2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*.

Sebastian Farquhar, Jannik Kossen, Lorenz Kuhn, and Yarin Gal. 2024. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017):625–630.

Diogo Nuno Freitas, Katherin Varela, and Eduardo Fermé. 2026. User profiling and its dynamics: A narrative review. *The European Journal on Artificial Intelligence*, 0(0):30504554251407092.

Susan Gauch, Mirco Speretta, Aravind Chandramouli, and Alessandro Micarelli. 2007. User profiles for personalized information access. *The adaptive web: methods and strategies of web personalization*, pages 54–89.

Daniela Godoy and Analia Amandi. 2005. User profiling in personal information agents: a survey. *The Knowledge Engineering Review*, 20(4):329–361.

Jian Guan, Junfei Wu, Jia-Nan Li, Chuanqi Cheng, and Wei Wu. 2025. A survey on personalized alignment - the missing piece for large language models in real-world applications. In *ACL (Findings)*, pages 5313–5333.

Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiusi Chen, and Julian McAuley. 2024. Bridging language and items for retrieval and recommendation. *arXiv preprint arXiv:2403.03952*.

William Huitt. 2007. Maslow's hierarchy of needs. *Educational psychology interactive*, 23.

Gawesh Jawaheer, Peter Weller, and Patty Kostkova. 2014. Modeling user preferences in recommender systems: A classification framework for explicit and implicit user feedback. *ACM Transactions on Interactive Intelligent Systems*, 4(2):8:1–8:26.

Bowen Jiang, Zhuoqun Hao, Young-Min Cho, Bryan Li, Yuan Yuan, Sihao Chen, Lyle Ungar, Camillo J. Taylor, and Dan Roth. 2025. Know me, respond to me: Benchmarking llms for dynamic user profiling and personalized responses at scale. In *COLM*.

Priyanka Kargupta, Nan Zhang, Yunyi Zhang, Rui Zhang, Prasenjit Mitra, and Jiawei Han. 2025. Taxoadapt: Aligning llm-based multidimensional taxonomy construction to evolving research corpora. In *ACL*, pages 29834–29850.

Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A. Hale. 2024. The benefits, risks and bounds of personalizing the alignment of large language models to individuals. *Nature Machine Intelligence*, 6(4):383–392.

Yang Li, Jiaxiang Liu, Yusong Wang, Yujie Wu, and Mingkun Xu. 2026. Bmam: Brain-inspired multi-agent memory framework. *arXiv preprint arXiv:2601.20465*.

Zhiyu Li, Chenyang Xi, Chunyu Li, Ding Chen, Boyu Chen, Shichao Song, Simin Niu, Hanyu Wang, Jiawei Yang, Chen Tang, and 1 others. 2025. Memos: A memory os for ai system. *arXiv preprint arXiv:2507.03724*.

Jiahong Liu, Zexuan Qiu, Zhongyang Li, Quanyu Dai, Wenhao Yu, Jieming Zhu, Minda Hu, Menglin Yang, Tat-Seng Chua, and Irwin King. 2025. A survey of personalized large language models: Progress and future directions. *arXiv preprint arXiv:2502.11528*.

Jaap MJ Murre and Joeri Dros. 2015. Replication and analysis of ebbinghaus' forgetting curve. *PloS one*, 10(7):e0120644.

Nous Research. 2026. Hermes Agent. `https://github.com/nousresearch/hermes-agent`, commit `98d75dea5a86aec599b1e081f8bbe9170bd3f964`, 2026-04-27; release `v0.11.0`, 2026-04-23.

Gabor Petnehazi and Bernadett Aradi. 2025. Hercules: Hierarchical embedding-based recursive clustering using llms for efficient summarization. *arXiv preprint arXiv:2506.19992*.

Nusrat Jahan Prottasha, Md Kowsher, Hafijur Raman, Israt Jahan Anny, Prakash Bhat, Ivan Garibay, and Ozlem Garibay. 2025. User profile with large language models: Construction, updating, and benchmarking. *arXiv preprint arXiv:2502.10660*.

Erasmo Purificato, Ludovico Boratto, and Ernesto William De Luca. 2024. User modeling and user profiling: A comprehensive survey. *arXiv preprint arXiv:2402.09660*.

Xubin Ren, Wei Wei, Lianghao Xia, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. Representation learning with large language models for recommendation. In *WWW*, pages 3464–3475. ACM.

Sonia Roccas, Lilach Sagiv, Shalom H Schwartz, and Ariel Knafo. 2002. The big five personality factors and personal values. *Personality and social psychology bulletin*, 28(6):789–801.

Yunxiao Shi, Wujiang Xu, Zhang Zeqi, Xing Zi, Qiang Wu, and Min Xu. 2025. PersonaX: A recommendation agent-oriented user modeling framework for long behavior sequence. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 5764–5787, Vienna, Austria. Association for Computational Linguistics.

Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, Akshay Nathan, Alan Luo, Alec Helyar, Aleksander Madry, Aleksandr Efremov, Aleksandra Spyra, Alex Baker-Whitcomb, Alex Beutel, Alex Karpenko, and 465 others. 2025. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*.

Peter Steinberger and community. 2026. OpenClaw: Personal AI assistant runtime. https://github.com/openclaw/openclaw, commit 3120401f53e789caf565e60ba29cb9751829b1b6, 2026-04-27.

Hang Su, Yun Yang, Tianyang Liu, Xin Liu, Peng Pu, and Xuesong Lu. 2025. Personalized question answering with user profile generation and compression. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 4744–4763, Suzhou, China. Association for Computational Linguistics.

Lu Wang, Di Zhang, Fangkai Yang, Pu Zhao, Jianfeng Liu, Yuefeng Zhan, Hao Sun, Qingwei Lin, Weiwei Deng, Dongmei Zhang, Feng Sun, and Qi Zhang. 2025a. Lettingo: Explore user profile generation for recommendation system. In *KDD*, pages 2985–2995.

Piaohong Wang, Motong Tian, Jiaxian Li, Yuan Liang, Yuqing Wang, Qianben Chen, Tiannan Wang, Zhicong Lu, Jiawei Ma, Yuchen Eleanor Jiang, and Wangchunshu Zhou. 2025b. O-mem: Omni memory system for personalized, long horizon, self-evolving agents. *arXiv preprint arXiv:2511.13593*.

Shen Wang, Tianlong Xu, Hang Li, Chaoli Zhang, Joleen Liang, Jiliang Tang, Philip S. Yu, and Qingsong Wen. 2025c. Large language models for education: A survey and outlook. *IEEE Signal Processing Magazine*, 42(6):51–63.

Haoyang Wen, Zhenxin Xiao, Eduard H. Hovy, and Alexander G. Hauptmann. 2023. Towards open-domain twitter user profile inference. In *ACL (Findings)*, pages 3172–3188.

Yunjia Xi, Weiwen Liu, Jianghao Lin, Xiaoling Cai, Hong Zhu, Jieming Zhu, Bo Chen, Ruiming Tang, Weinan Zhang, and Yong Yu. 2024. Towards open-world recommendation with knowledge augmentation from large language models. In *RecSys*, pages 12–22. ACM.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. 2025. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*.

Yue Xu, Qian Chen, Zizhan Ma, Dongrui Liu, Wenxuan Wang, Xiting Wang, Li Xiong, and Wenjie Wang. 2026. Toward personalized LLM-powered agents: Foundations, evaluation, and future directions. *arXiv preprint arXiv:2602.22680*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Jiarui Zhang. 2024. Guided profile generation improves personalization with large language models. *arXiv preprint arXiv:2409.13093*.

Weizhi Zhang, Xiaokai Wei, Wei-Chieh Huang, Zheng Hui, Chen Wang, Michelle Gong, and Philip S. Yu. 2026. MemoryCD: Benchmarking long-context user memory of LLM agents for lifelong cross-domain personalization. In *Lifelong Agent Workshop at ICLR 2026*.

Zhehao Zhang, Ryan A. Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, Ruiyi Zhang, Jiuxiang Gu, Tyler Derr, Hongjie Chen, Junda Wu, Xiang Chen, Zichao Wang, Subrata Mitra, Nedim Lipka, and 2 others. 2025. Personalization of large language models: A survey. *Transactions on Machine Learning Research (TMLR)*, 2025.

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). *IEEE Transactions on Knowledge and Data Engineering*, 36(11):6889–6907.

Thomas P. Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, and Hongseok Namkoong. 2025. Personalllm: Tailoring llms to individual preferences. In *ICLR*.

A Appendix

A.1 Related Work

User profiling plays a crucial role in daily life, enabling applications such as personalized shopping experiences (Jawaheer et al., 2014; Al-Shamri, 2016), intelligent assistants (Wen et al., 2023; Zhang et al., 2025), and personalized medical services (Guan et al., 2025). Its primary objective is to distill heterogeneous user traces into compact representations of a user's attributes, interests, preferences, and needs (Gauch et al., 2007; Godoy and Amandi, 2005).

Traditional User Profiling Methods. Traditional user profiling techniques (Godoy and Amandi, 2005) can be broadly categorized into rule-based, content-based, collaborative, and supervised learning paradigms. Rule-based methods (Gauch et al., 2007; Chen et al., 2007) extract predefined attributes or interests from registration forms, search logs, browsing records, and textual content, whereas content-based methods (Jawaheer et al., 2014) represent users using keyword vectors, topic distributions, or aggregated item features. While these methods perform efficiently in well-defined domains, their representational capacity is constrained by manually designed schemas, sparse labels, and task-specific feature engineering. Consequently, they often produce flat and fragmented profiles that capture surface-level correlations but fail to reveal the underlying semantic motivations behind user behavior. Moreover, they struggle with open-ended natural language histories, long-tail interests, and cold-start users, where reliable attributes must be inferred from sparse or noisy evidence (Al-Shamri, 2016). These limitations make traditional profiles difficult to transfer across scenarios and insufficient for capturing the diversity and dynamics of user behavior in rich, complex contexts.

LLM-based User Profiling. With the advent of large language models (LLMs), user profiling has gained renewed significance. Applications such as personalized assistants, writing aids, educational question-answering systems, and recommendation engines require continuous, in-depth understanding of users to generate responses that are helpful, contextually appropriate, and aligned with individual preferences (Kirk et al., 2024; Xu et al., 2026). Recent studies have leveraged LLMs as semantic reasoners to construct user profiles from natural language histories and behavioral sequences. Compared with conventional feature engineering, LLM-based profiling can summarize sparse personal context, normalize heterogeneous signals, and provide interpretable descriptions of user preferences. For example, Guided Profile Generation (GPG) (Zhang, 2024) demonstrates that converting raw personal context into natural-language profiles enhances personalization in both preference prediction and response generation tasks. In recommendation systems, LLM-based representation learning and profile generation approaches enrich user modeling by incorporating semantic knowledge from item text and interaction histories (Ren et al., 2024; Zhao et al., 2024; Wang et al., 2025a). PersonaX (Shi et al., 2025) addresses long behavior sequences by clustering historical interactions and generating cached multi-persona profiles for recommendation agents. In conversational and assistant settings, PERSONAMEM (Jiang et al., 2025) evaluates LLMs' ability to internalize user traits, track profile evolution, and respond appropriately in novel scenarios, highlighting that direct prompting still struggles to capture dynamic preference shifts. Recent user profiling benchmarks (Jiang et al., 2025; Su et al., 2025; Bian et al., 2026; Zhang et al., 2026) further emphasize that effective personalization requires not only recognizing stable traits but also tracking evolving preferences across sessions. To address these challenges, we propose EDHUP (Evolving Dynamic Hierarchical User Profiling), which generates user profiles that are interpretable, multi-granular, and self-evolving, making them particularly suitable for personalized LLM applications.

A.2 Proof

Proof of Theorem 3.1. Given that the user profiling task can be viewed as inference processes within probabilistic graphical models, the goal is to compute posterior distribution of the latent user profile Z given observations H_u:

$$P(Z|H_u) = \frac{P(H_u, Z)}{P(H_u)} = \frac{P(H_u|Z)P(Z)}{\int P(H_u, Z)dZ}. \tag{11}$$

Due to high dimensionality of the latent space Z and complex dependencies within the interaction history H_u, the marginal likelihood $P(H_u)$ is computationally intractable. Consequently, exact inference of $P(Z|H_u)$ is impossible. To overcome this intractability, we employ variational inference

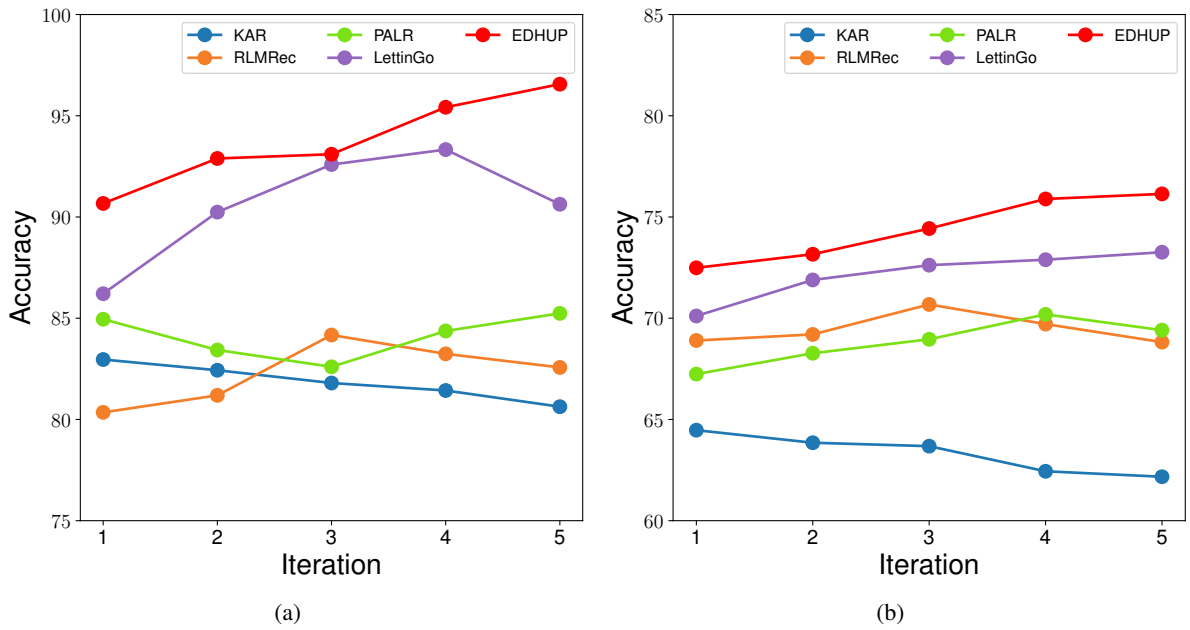

Figure 2: The mean accuracy changes over the last five interaction iterations of various different profile generation methods for downstream sequence recommendation tasks on the Amazon Reviews Dataset: (a) the baby products domain (b) the Movies and TV domain.

technique and introduce a variational distribution family $\mathcal{Q}$ and seek a distribution $Q_\phi(Z|H_u) \in \mathcal{Q}$ that best approximates the true posterior $P(Z|H_u)$. Its objective is to minimize the KL divergence between the approximation distribution $Q_\phi(Z|H_u)$ and the truth distribution $P(Z|H_u)$:

$$\phi^* = \arg\min_\phi \mathbb{D}_{KL}(Q_\phi(Z|H_u)||P(Z|H_u)), \tag{12}$$

where the KL divergence is defined as:

$$\mathbb{D}_{KL}(Q_\phi||P) = \mathbb{E}_{Q_\phi}[\log Q_\phi(Z|H_u)/P(Z|H_u)]. \tag{13}$$

Then substituting $P(Z|H_u) = P(H_u, Z)/P(H_u)$, we rearrange this term and obtain the log-marginal likelihood decomposition:

$$\log P(H_u) = \underbrace{\mathbb{E}_{Q_\phi}[\log(P(\mathcal{H}_u, \mathcal{Z})/Q_\phi(\mathcal{Z}|\mathcal{H}_u))]}_{\mathcal{L}(\phi)\text{ (ELBO)}} + \underbrace{\mathbb{D}_{KL}(Q_\phi||P)}_{\geq 0}. \tag{14}$$

Since $\log P(H_u)$ is constant with respect to ϕ, minimizing the KL divergence is mathematically equivalent to maximizing the ELBO $\mathcal{L}(\phi)$. The core innovation of H-CoT lies in the structural constraint it imposes on the variational distribution. Based on the d-separation properties of our defined Bayesian Network, we factorize the approximate posterior Q_ϕ autoregressively, mirroring the sequential inference capability of the LLM:

$$Q_\phi(Z|H_u) \approx \prod_{l=0}^{L} q_{\phi_l}(Z^{(l)}|Z^{(<l)}, H_u). \tag{15}$$

Then substituting the specific layers $Z^{(0)}, Z^{(1)}, Z^{(2)}$ defined in our model:

$$Q_\phi(Z|H_u) = q_{\phi_0}(Z^{(0)}|H_u) \cdot q_{\phi_1}(Z^{(1)}|Z^{(0)}, H_u) \cdot q_{\phi_2}(Z^{(2)}|Z^{(1)}, Z^{(0)}, H_u). \tag{16}$$

When the LLM performs H-CoT, it generates tokens sequentially following the prompt's structure. This process can be viewed as sampling from the factorized variational distribution. The model first samples $Z^{(0)} \sim q_{\phi_0}(\cdot|H_u)$. Conditioned on the output of Level 0, it samples $Z^{(1)} \sim q_{\phi_1}(\cdot|Z^{(0)}, H_u)$.Finally, it samples $Z^{(2)} \sim q_{\phi_2}(\cdot|Z^{(1)}, Z^{(0)}, H_u)$.

By fine-tuning or prompting the LLM to produce coherent reasoning chains at each level l, we are effectively maximizing the conditional log-probabilities $\log q_{\phi_l}$. Consequently, the H-CoT process maximizes the sum of these expectations, which constitutes the entropy term and the reconstruction term of ELBO. Thus, we have proven that H-CoT is not merely a heuristic for interpretability but a structured variational inference procedure.

It optimizes the ELBO $\mathcal{L}(\phi)$, thereby minimizing $\mathbb{D}_{KL}(Q_\phi || P)$ and ensuring that the generated user profile Z approximates the true latent characteristics implied by the interaction history H_u. □

A.3 Algorithm Details of EDHUP

To integrate Hierarchical Chain-of-Thought (H-CoT) reasoning and Bayesian-driven update mechanisms into our proposed EDHUP framework, we first initialize the user profile based on the historical interaction data $\mathcal{H}_u$. We extract the Basic Attribute Layer $Z^{(0)}$ using schema-constrained deductive reasoning. Building upon $Z^{(0)}$, we construct the Behavioral Characteristic Layer $Z^{(1)}$ through a recursive hierarchical clustering process, which dynamically subdivides behavioral features until the semantic entropy $H_{\text{sem}}(z)$ falls below a predefined threshold ε_{low}. Subsequently, we employ abductive reasoning to infer the Psychological Motivation Layer $Z^{(2)}$ from the lower-layer representations, finalizing the initial three-dimensional user profile $\mathcal{Z}$. All extracted features z are initialized with confidence parameters α and β following a Beta distribution.

During the continuous interaction phase, we develop a Bayesian-based self-evolving profile update mechanism. For each new interaction within a specific scene S at time t_{now}, we first design a time-decay mechanism, which is implemented by allowing the existing confidence parameters α and β to decay over time based on the elapsed time Δt. We then calculate the evidence strength L_z based on the frequency of occurrence of a feature in the current scene. If the user's behavior aligns with the existing feature's expected scenario, the positive evidence α is reinforced; if it conflicts, the uncertainty β is increased. Finally, we compute the expected confidence change ΔC to strengthen or weaken profile features. Any feature with an expected confidence $E[C_{\text{new}}]$ below the pruning threshold τ_{prune} will be removed from the current user's profile. The detailed algorithm of our proposed EDHUP framework is summarized in Algorithm 1.

Algorithm 1: EDHUP

Input: History $\mathcal{H}_u$, Current Scene S, Thresholds $\{\varepsilon_{\text{low}}, \lambda, \tau_{\text{prune}}, W_{\text{scene}}\}$
Output: Dynamic User Profile $\mathcal{Z} = \{Z^{(0)}, Z^{(1)}, Z^{(2)}\}$

```
// Phase 1: Profile Initialization via H-CoT
```
Extract $Z^{(0)}$ via schema-constrained CoT;
Construct $Z^{(1)}$ via recursive clustering $\mathcal{R}$ until $H_{\text{sem}}(z) < \varepsilon_{\text{low}}$;
Infer $Z^{(2)}$ via abductive reasoning on $\{\mathcal{H}_u, Z^{(0)}, Z^{(1)}\}$;
Initialize $\mathcal{Z} \leftarrow \{Z^{(0)}, Z^{(1)}, Z^{(2)}\}$ with parameters $\{\alpha, \beta\}$;
```
// Phase 2: Bayesian-Driven Self-Evolution
```
for *each interaction at* t_{now} **do**
  $\Delta t \leftarrow t_{\text{now}} - t_{\text{last}}$;
  for *each profile feature* $z \in \mathcal{Z}$ **do**
    $(\alpha, \beta) \leftarrow (\alpha, \beta) \cdot e^{-\lambda \Delta t}$; `// Temporal Forgetting`
    $L_z \leftarrow \log(1 + \text{Occ}_S(z)) \cdot W_{\text{scene}}$;
    if *S matches the expected attribute of z* **then**
      $\alpha \leftarrow \alpha + L_z$; `// Reinforce`
    else
      $\beta \leftarrow \beta + L_z$; `// Weaken`
    end
    $E[C] \leftarrow \alpha/(\alpha + \beta)$; `// Bayesian Update`
    if $E[C] < \tau_{prune}$ **then**
      $\mathcal{Z} \leftarrow \mathcal{Z} \setminus \{z\}$ **else** $t_{\text{last}} \leftarrow t_{\text{now}}$;
  end
end
return $\mathcal{Z}$.

A.4 Experimental settings

A.4.1 Experimental settings of Section 2.4.1

To evaluate the performance of various user profile generation methods the Wiki people benchmark, we follow the settings of Prottasha et al. (2025) and use the latest most cost-effective LLM models include GPT 5.2, Qwen3.5-35B-A3B, Deepseek v3.2 to extract user profile features. To ensure fairness

and strike a balance between the accuracy and diversity of profile generation, we set the generation temperature and top parameters for all LLMs to 0.6 and 0.9, respectively. Unlike the method of Nusrat (Prottasha et al., 2025), we use a similar generation prompt for LLMs during the profile construction and updating processes a similar prompt except for the input data. During the profile building phase, the input of LLMs consists only the Wiki People data. When updating the user profile, the input for LLMs includes the most recently updated Wiki People data and the previously builted user profile. The generation prompts of different methods are shown in Figures 3 to 5. we employ two classic predictive indicators accuracy and F1 to evaluate the effectiveness of various profile generation methods in terms of portrait generation and updating processes.

A.4.2 Experimental settings of Section 2.4.2

To evaluate the performance of various user profile generation methods on the downstream sequence recommendation task within the large E-commerce Amazon Reviews 2023 (Hou et al., 2024) dataset, we follow the settings of LettinGo (Wang et al., 2025a). Note that in practice, the product items in historical interaction sequences and the item to predict are usually from the same domain defined in the Amazon Reviews datasets (Hou et al., 2024). We leverage this rich information to construct user profiles and predict products that users are likely to be interested in next interaction round. The dataset provides rich information, including user reviews, such as ratings and review texts, item metadata, such as product descriptions and prices, and relational signals, such as user–item co-purchase graphs.

To ensure that meaningful user profiles can be generated from sufficiently rich historical interaction sequences, following Wang et al. (2025a), we retain users with more than 70 historical interactions. We then combine each user's historical interaction data with the corresponding generated profile to iteratively predict the items that the user may be interested in over the last five interaction rounds. Similar to GPG (Zhang, 2024) and LettinGo (Wang et al., 2025a), we directly employ an LLM as the sequential recommendation model to predict products that users are likely to prefer. The inputs to the recommendation system include historical interaction records, the generated user profile, and a candidate item. The system then predicts the user's preference toward the candidate item, categorizing it into one of three classes: *like*, *dislike*, or *neutral*. We employ the same Qwen3.5-35B-A3B LLM model for profile generation and product recommendation with different prompts and generation hpper-parameters, respectively. For the profile generation, the generation temperature and top parameter are set to 0.6 and 0.9 seperately, which makes a balance between the accuracy and diversity of profile generation. For the product recommendation, the model temperature is set to 0 to ensure deterministic outputs. The prompts of profile generation and product recommendation are shown in Figure 6 and Figure 7 seperately.

A.4.3 Experimental settings of Section 2.4.3

To evaluate the performance of various user profile generation methods on the downstream personalized AI assistant task on the PERSONAMEM benchmark, we follow the settings of Jiang et al. (2025) and use the same benchmark for a fair comparison. It contains up to 60 sessions of multi-turn conversations across 15 diverse real-world tasks, ranging from personal topics such as family, dating, health, and finance, therapy and legal advice to recommendations on books, music, movies, and food, and practical tasks like travel planning, online shopping, studying tips, and home decoration. In each task scenario, the LLM model selects the most appropriate response from four different choices according to the user query, historical interactions and current situation. The evaluation metic is the accuracy of multiple-choice anwers for different questions. The conversation sessions of different tasks are interleaved in the chronological order of their occurrence. To evaluate the ability of personalized assistants in generating personalized responses with dynamic user profiles, We employ the same Qwen3.5-35B-A3B LLM model for profile generation and product recommendation with different prompts and generation hpper-parameters, respectively. For the profile generation, the generation temperature and top parameter are set to 0.7 and 0.9 seperately, which makes a balance between the accuracy and diversity of profile generation. For the selection of multiple-choice anwers from different questions, the model temperature is set to 0 to ensure deterministic outputs. The prompts of profile generation and personalized response options are shown in Figure 8 and Figure 9 seperately.

Generating and updating user profiles with the fixed static template.

Given a user resume file: {profile_text} and {old user profiles } (if exists), a detailed, objective, and complete user profile is analyzed and mined. The core summary dimensions include basic characteristics, behavioral characteristics, and psychological motivation (all need to be covered, no relevant information is marked as "unknown ", and the expected confidence of inferring each profile is $E[C_{\text{new}}]$). Note that this is strictly based on content distillation. Do not add any guesswork not already mentioned. The user profiles include:
**Basic features: Basic attributes of users, mainly including Name, Age, Birth place, Gender, Education, Occupation, industry, profession, awards, etc.
**Behavior characteristics: user's interests and behavior habits, mainly including likes, dislikes, hobbies, interests, research interests, work experience, etc.
**Psychological motivation: user's potential needs, personality characteristics, personality traits, mainly including Responsibility, Memberships, influences, etc.
The updating rule for the generated profile involves recalculating the confidence score $E[C_{\text{new}}]$ based on the current newest context and the time interval. If the confidence score of the current profile falls below the threshold τ_{prune} , this feature is forgotten from the current generated user profiles.
At last, please return the profile results in json format, do not include a "'json"' field in the output, and do not return anything else.

Figure 3: Prompt for generating user profile given the persona description in the EUP method .

Generating and updating user profiles with the CoT template.

Given a user resume file: {profile_text} and {old user profiles } (if exists), a detailed, objective, and complete user profile is analyzed and mined. The core summary dimensions include basic characteristics, behavioral characteristics, and psychological motivation (all need to be covered, no relevant information is marked as "unknown ", and the expected confidence of inferring each profile is $E[C_{\text{new}}]$). Note that this is strictly based on content distillation. Do not add any guesswork not already mentioned. The user profiles include:
**Basic features: Basic attributes of users, mainly including Name, Age, Birth place, Gender, Education, Occupation, industry, profession, awards, etc.
**Behavior characteristics: user's interests and behavior habits, mainly including likes, dislikes, hobbies, interests, research interests, work experience, etc.
**Psychological motivation: user's potential needs, personality characteristics, personality traits, mainly including Responsibility, Memberships, influences, etc.
The reasoning process of EDUP first extracts basic features from the persona description, and then discover behavior characteristics using basic features and the original persona description information. At last, use the psychological knowledge and causality to infer the implicit user's psychological motivation by combing basic features, the discovered behavior characteristics and the persona description information.
The updating rule for the generated profile involves recalculating the confidence score $E[C_{\text{new}}]$ based on the current newest context and the time interval. If the confidence score of the current profile falls below the threshold τ_{prune} , this feature is forgotten from the current generated user profiles.
At last, please return the profile results in json format, do not include a "'json"' field in the output, and do not return anything else.

Figure 4: Prompt for generating user profile given the persona description in the EDUP method .

Generating and updating user profiles with the H-CoT template.

Given a user resume file: {profile_text} and {old user profiles } (if exists), a detailed, objective, and complete user profile is analyzed and mined. The core summary dimensions include basic characteristics, behavioral characteristics, and psychological motivation (all need to be covered, no relevant information is marked as "unknown ", and the expected confidence of inferring each profile is $E[C_{\text{new}}]$). Note that this is strictly based on content distillation. Do not add any guesswork not already mentioned. The user profiles include:
**Basic features: Basic attributes of users, mainly including Name, Age, Birth place, Gender, Education, Occupation, industry, profession, awards, etc.
**Behavior characteristics: user's interests and behavior habits, mainly including likes, dislikes, hobbies, interests, research interests, work experience, etc.
**Psychological motivation: user's potential needs, personality characteristics, personality traits, mainly including Responsibility, Memberships, influences, etc.
The reasoning process of EDUP first extracts basic features from the persona description, and then discover behavior characteristics using basic features and the original persona description information. After that, it dynamically subdivides behavioral features and discover similar behavior patterns until the semantic entropy $H_{\text{sem}}(z)$ falls below a predefined threshold ε_{low}. At last, use the psychological knowledge and causality to infer the implicit user's psychological motivation by combing basic features, the discovered behavior characteristics and the persona description information.
The updating rule for the generated profile involves recalculating the confidence score $E[C_{\text{new}}]$ based on the current newest context and the time interval. If the confidence score of the current profile falls below the threshold τ_{prune} , this feature is forgotten from the current generated user profiles.
At last, please return the profile results in json format, do not include a "'json"' field in the output, and do not return anything else.

Figure 5: Prompt for generating user profile given the persona description in the EDHUP method .

Generating and updating user profiles with the H-CoT template.

Given the provided purchase history summary {purchase_history} and {old user profiles } (if exists), a detailed, objective, and complete user profile is analyzed and mined. The core summary dimensions include basic characteristics, behavioral characteristics, and psychological motivation (all need to be covered, no relevant information is marked as "unknown ", and the expected confidence of inferring each profile is $E[C_{\text{new}}]$). Note that this is strictly based on content distillation. Do not add any guesswork not already mentioned. The user profiles include:
**Basic features: Basic attributes of users, mainly including likely_age_range, Gender, Education, Occupation, region, etc.
**Behavior characteristics: user's interests and behavior habits, mainly including likes, dislikes, hobbies, purchase_style, etc.
**Psychological motivation: user's potential needs, personality characteristics, personality traits, mainly including needs, Memberships, traits, etc.
The reasoning process of EDUP first extracts basic features from the persona description, and then discover behavior characteristics using basic features and the original persona description information. After that, it dynamically subdivides behavioral features and discover similar behavior patterns until the semantic entropy $H_{\text{sem}}(z)$ falls below a predefined threshold ε_{low}. At last, use the psychological knowledge and causality to infer the implicit user's psychological motivation by combing basic features, the discovered behavior characteristics and the persona description information.
The updating rule for the generated profile involves recalculating the confidence score $E[C_{\text{new}}]$ based on the current newest context and the time interval. If the confidence score of the current profile falls below the threshold τ_{prune} , this feature is forgotten from the current generated user profiles.
At last, please return the profile results in json format, do not include a "'json"' field in the output, and do not return anything else.

Figure 6: Prompt for generating user profile given user's purchase history summaries in the EDHUP method .

Predicting the next item the user might purchase.

You are an e-commerce recommender. Given the user's last {purchase history }, the {generated user profile}, and a candidate product, predict the user's preference. User attitudes fall into three categories: like, dislike, and neutral. Output exactly one label from the set like, dislike, neutral. Do not provide any explanation or extra text.
** Prediction criteria
Your prediction should be based on the semantic consistency between the candidate product and the user's historical and inferred preferences, including but not limited to: product category relevance, brand preference, functional match, style consistency, price-level consistency, inferred needs and interests from the user profile.
** Label assignment rules
Assign like when the candidate product is strongly aligned with the user's observed or inferred preferences. Assign dislike when the candidate product is clearly misaligned with or contradictory to the user's observed or inferred preferences. Assign neutral when the evidence is insufficient, ambiguous, or does not support a clear positive or negative judgment.

Figure 7: Prompt for Predicting the next item the user might purchase with the generated user profile.

Generating and updating user profiles with the H-CoT template.

Given the user's {conversation histories } and {old user profiles } (if exists), a detailed, objective, and complete user profile is analyzed and mined. The core summary dimensions include basic characteristics, behavioral characteristics, and psychological motivation (all need to be covered, no relevant information is marked as "unknown ", and the expected confidence of inferring each profile is $E[C_{\text{new}}]$). Note that this is strictly based on content distillation. Do not add any guesswork not already mentioned. The user profiles include:
**Basic features: Basic attributes of users, mainly including sex, likely_age_range, education, occupation, region, nickname, social_relations, etc.
**Behavior characteristics: user's interests and behavior habits, mainly including likes, dislikes, hobbies, purchase_style, work_behavior, life_behavior, social_behavior, interaction_behavior, etc.
**Psychological motivation: user's potential needs, personality characteristics, personality traits, mainly including potential_needs, personality_traits, emotional_state, etc.
The reasoning process of EDUP first extracts basic features from the persona description, and then discover behavior characteristics using basic features and the original persona description information. After that, it dynamically subdivides behavioral features and discover similar behavior patterns until the semantic entropy $H_{\text{sem}}(z)$ falls below a predefined threshold ε_{low}. At last, use the psychological knowledge and causality to infer the implicit user's psychological motivation by combing basic features, the discovered behavior characteristics and the persona description information.
The updating rule for the generated profile involves recalculating the confidence score $E[C_{\text{new}}]$ based on the current newest context and the time interval. If the confidence score of the current profile falls below the threshold τ_{prune} , this feature is forgotten from the current generated user profiles.
At last, please return the profile results in json format, do not include a "'json"' field in the output, and do not return anything else.

Figure 8: Prompt for generating user profile given user's personalized conversation histories in the EDHUP method .

Selecting the best appropriate response.

You are a personalized AI assistant. Given the user's { historical multi-turn conversations }, the {generated user profile}, and four candidate responses, selecting the most appropriate response from four different choices according to the user query, historical interactions and current situation. Output exactly one label from the four different choices. Do not provide any explanation or extra text.

** **selection criteria**

The selection should be made by comprehensively considering the user's historical interactions, inferred preferences and traits from the user profile, the current conversational context, and the suitability of each candidate response. When evaluating the candidate responses, consider the following aspects:

1. Contextual relevance. Determine whether the candidate response is directly relevant to the current user query and consistent with the ongoing conversational context.
2. Historical consistency. Examine whether the candidate response is coherent with the user's historical multi-turn conversations, including previously expressed preferences, intentions, habits, communication patterns, and recurring topics.
3. Profile alignment Assess whether the candidate response aligns with the generated user profile, such as the user's inferred interests, preferences, background characteristics, communication style, and likely needs.
4. Personalization quality Prefer responses that are more personalized and better tailored to the individual user, rather than generic or weakly related alternatives.
5. Situational appropriateness Select the response that best fits the current conversational situation, including the user's immediate intent, emotional tone, and practical needs.
6. Coherence and usefulness Prefer responses that are coherent, natural, informative, and helpful in advancing the conversation.

** **Decision rules**

Choose the candidate that is most appropriate overall under the above criteria. If multiple candidates appear plausible, prefer the one with stronger personalization, better contextual fit, and higher utility. Do not rely on unsupported assumptions that cannot be inferred from the provided conversation history, user profile, or current query. Base the decision only on the given information.

** **Label assignment rules**

Assign like when the candidate product is strongly aligned with the user's observed or inferred preferences. Assign dislike when the candidate product is clearly misaligned with or contradictory to the user's observed or inferred preferences. Assign neutral when the evidence is insufficient, ambiguous, or does not support a clear positive or negative judgment.

Figure 9: Prompt for selecting the best appropriate response of multi-turn conversations across different tasks with the generated user profile.