

Me-Agent: A Personalized Mobile Agent with Two-Level User Habit Learning for Enhanced Interaction

Shuoxin Wang^{1*}, Chang Liu^{2*}, Gowen Loo^{3*}, Lifan Zheng⁴
Kaiwen Wei⁵, Huanqian Yan⁶, Yinyi Zeng⁷, Jingyuan Zhang⁸, Yu Tian^{9†}

¹Yunnan University ²Hong Kong Polytechnic University ³University of Electronic Science and Technology of China
⁴Southeast University ⁵Chongqing University ⁶Beihang University
⁷University of Chinese Academy of Sciences
⁸Kuaishou Technology ⁹Tsinghua University
shuoxin_wang@163.com, tianyu1810613@gmail.com

Abstract

Large Language Model (LLM)-based mobile agents have made significant performance advancements. However, these agents often follow explicit user instructions while overlooking personalized needs, leading to significant limitations for real users, particularly without personalized context: (1) inability to interpret ambiguous instructions, (2) lack of learning from user interaction history, and (3) failure to handle personalized instructions. To alleviate the above challenges, we propose Me-Agent, a learnable and memorable personalized mobile agent. Specifically, Me-Agent incorporates a two-level user habit learning approach. At the prompt level, we design a user preference learning strategy enhanced with a Personal Reward Model to improve personalization performance. At the memory level, we design a Hierarchical Preference Memory, which stores users' long-term memory and app-specific memory in different level memory. To validate the personalization capabilities of mobile agents, we introduce User FingerTip, a new benchmark featuring numerous ambiguous instructions for daily life. Extensive experiments on User FingerTip and general benchmarks demonstrate that Me-Agent achieves state-of-the-art performance in personalization while maintaining competitive instruction execution performance. Our code is publicly available at: <https://github.com/Gianthaha/personal-mobile-agent.git>.

1 Introduction

With the continuous advancement of large language models (LLMs) (Achiam et al., 2023; Anthropic, 2024), LLM-based mobile agent (Qin et al., 2025; Ye et al., 2025) have achieved significant progress in performance and generalization, enabling them to execute complex tasks across various application scenarios. However, their learning paradigm

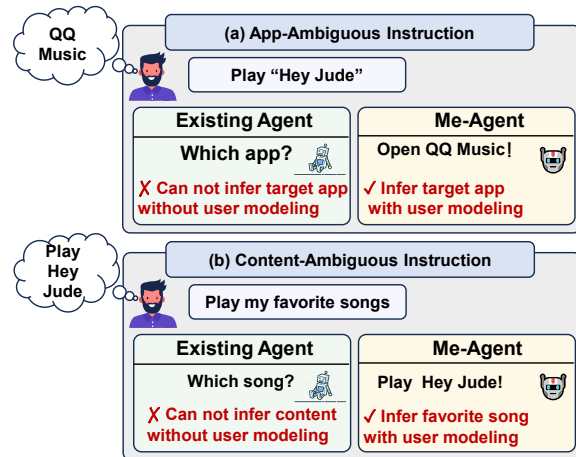


Figure 1: Two key challenges for mobile agents in ambiguous instruction.

still primarily relies on explicit user instructions, lacking systematic modeling and long-term adaptation to users' latent preferences and personalized needs. Particularly when user background information is absent or application scenarios are complex and dynamic, existing LLM-based mobile agents (Wang et al., 2024a, 2025b; Li et al., 2024c) exhibit three prominent limitations: 1) difficulty accurately interpreting natural language instructions with vague or implicit user intentions; 2) inability to continuously learn from and update user behavioral patterns across multi-turn interactions; and 3) limited capability in processing personalized instructions and preference configurations, hindering user-centric interactive experiences.

While prior research has made substantial progress in LLM personalization, applying these methods to mobile agents reveals significant limitations. First, existing methods (Li et al., 2024a; Zhang et al., 2025b) often rely on fine-tuning, which proves impractical for mobile agents due to limited computational resources and storage. Cloud-based training alternatives introduce addi-

*Equal contribution.

†Corresponding author

tional costs and raise privacy and security concerns. Second, many approaches (Yao et al., 2025; Kong et al., 2024) express user preferences and context through continuously expanding prompts, but without filtering mechanisms, this degrades inference efficiency and complicates storage and version management as interactions accumulate.

To address these limitations, we propose **Me-Agent**, a learnable and memorable personalized mobile agent that adapts to user preferences through two-level habit modeling without additional training. Me-Agent operates at two levels: 1) In prompt level, we introduce User Preference Learning strategy with a Personal Reward Model that ranks candidate responses by user preference alignment, guiding personalized decision-making. 2) In memory level, we construct a Hierarchical Preference Memory that stores and retrieves users' long-term memory and app usage patterns, mitigating prompt length issues while enabling continuous modeling of user trajectories and behavioral patterns. Me-Agent offers two key advantages: 1) it requires no model parameter modifications, making it suitable for API-based deployment; and 2) it employs a hierarchical optimization strategy that effectively mitigates memory decay.

To evaluate the personalization capabilities of mobile agents in executing ambiguous instructions, we construct User FingerTip to assess mobile agents' performance in inferring user intent, capturing personalized preferences, and adapting through multi-turn interactions. Extensive experiments demonstrate that Me-Agent achieves state-of-the-art performance on personalization metrics and surpasses baselines on performance metrics. We further evaluate Me-Agent on the E-dataset (Wang et al., 2025b), a general benchmark with complex, high-difficulty, multi-app instructions. Me-Agent achieves a task completion rate of 89.3%, a 39.7% improvement over the baseline (63.9%), demonstrating strong generalization to challenging mobile tasks. Our main contributions are as follows:

- We design Me-Agent, a learnable and memorable personalized mobile agent that adapts to user preferences through two-level habit modeling without additional training.
- We propose User Preference Learning, a parameter-free module for modeling user preferences by a Personal Reward Model.
- We design Hierarchical Preference Memory,

which stores users' long-term and app-specific memory at different levels.

- We develop a personalized benchmark featuring two types of ambiguous instructions to test mobile agents' implicit preference reasoning.

2 Related Work

2.1 LLM-based Mobile Agents

Early mobile agents use single-agent architectures (Wang et al., 2024b; Zhang et al., 2025a), which are prone to decision drift and error accumulation in long-horizon planning and cross-application tasks. Recent research on multi-agent architectures (Wang et al., 2024a, 2025b; Li et al., 2024c; Loo et al., 2025) focuses on improving instruction execution success rates through collaborative decision-making, hierarchical structures, self-evolution mechanisms, and structured knowledge bases for enhanced accuracy in complex scenarios. However, most existing mobile agents still rely on explicit user instructions and utilize limited user-related information, hindering their ability to predict user intent and model preferences. PerPilot (Wang et al., 2025a) introduces personalization by resolving user-specific entities via memory retrieval, but it primarily focuses on entity resolution and does not capture long-term user behavior or preferences. In contrast, Our work focuses on modeling user behavioral preferences in mobile agents to enable more personalized instruction execution.

2.2 Personalization in LLMs

Personalization of large language models is important research direction, which can be categorized into four main approaches (Zhang et al., 2024). Parameter-level methods (Hu et al., 2022; Dettmers et al., 2023) enable efficient personalization through fine-tuning techniques, while prompt and context-driven methods (Salemi et al., 2024; Zhuang et al., 2024) use user information as prompts to guide model outputs. Additionally, memory and retrieval-augmented methods (Fan et al., 2025; Shi et al., 2025; Maragheh et al., 2025) enhance personalization by maintaining long-term user memory, and preference alignment methods (Ouyang et al., 2022; Li et al., 2024b; Poddar et al., 2024) leverage reinforcement learning to model diverse user preferences effectively. While these methods effectively mitigate the personalization issues of LLMs, their application in the

Type	Form	Instruction Example	app^*	c^*
Original	$\langle app, a, c \rangle$	“Play Hey Jude on QQ Music”	–	–
Type I	$\langle \emptyset, a, c \rangle$	“Play Hey Jude”	QQ Music	–
Type II	$\langle \emptyset, a, r \rangle$	“Play my favorite song”	QQ Music	Hey Jude

Table 1: Examples of instruction construction. app^* and c^* represent the ground-truth application and content.

LLM-based agent encounters challenges like training difficulties and contextual redundancy. To overcome these challenges, we propose Me-Agent, a personalized mobile agent framework that models user habits at both the prompt and memory levels, enabling implicit intent inference from ambiguous instructions without training.

3 User FingerTip

Real-world users tend to use ambiguous expressions—omitting application names (e.g., "I want to listen to music") or using imprecise references (e.g., "play what I usually play"). Such instructions require the agent to infer application and content preferences from the user’s behavioral history. However, existing personalized datasets (Yang et al., 2025; Wang et al., 2025a) fail to assess this capability due to two major limitations: 1) personalized information is explicitly stated in the commands; 2) the small scale of the datasets limits comprehensive capability evaluation. To fill this gap, we extend **FingerTip** to build a new personalized benchmark, introducing two types of ambiguous instructions to test mobile agents’ implicit preference reasoning.

3.1 Data Construction

The basic instruction can be formalized as a triplet

$$I_{orig} = \langle app, a, c \rangle, \quad (1)$$

where app denotes the target application, a denotes the action, and c denotes the operation content. In real scenarios, users tend to use ambiguous expressions—omitting application names (e.g., "I want to listen to music") or replacing specific content with vague references (e.g., "order my usual"). As shown in Table 1, we design two types of ambiguous instructions to construct User FingerTip.

Type I: App-Ambiguous Instruction. Application preference reasoning is a fundamental capability for personalized mobile agents, as multiple applications often coexist within the same functional category (e.g., NetEase Music and QQ Music). However, existing datasets (Wang et al., 2025b;

Statistic	Value
Users	60
Functional Categories	12
Applications	33
Type I Instructions	300
Type II Instructions	267
Training Instructions per User	15
Test Instructions per User	5

Table 2: Dataset statistics.

Zhang et al., 2025a; Wang et al., 2024a) explicitly provide target applications in instructions, making it infeasible to evaluate whether mobile agents can infer user preferences from behavioral history. To address this, we construct Type I instructions by removing the application name:

$$I_{type\ I} = \langle \emptyset, a, c \rangle, \quad (2)$$

where a denotes the user action and c denotes the explicit content. The ground truth is annotated as the user’s most frequently used application within the corresponding category.

Type II: Content-Ambiguous Instruction. While Type I focuses on application preference reasoning, comprehensive personalization also requires understanding user content preferences. We design Type II instructions to evaluate joint reasoning: inferring the target application while retrieving preferred content from behavioral history. Specifically, we extend Type I by replacing explicit content c with implicit references r (e.g., “my favorite”, “the usual”), yielding:

$$I_{type\ II} = \langle \emptyset, a, r \rangle. \quad (3)$$

Each instruction is annotated with both the target application a^* and the user-preferred content c^* inferred from historical behavior.

3.2 Data Statistics

Table 2 summarizes the dataset statistics. User FingerTip includes 60 users, 33 applications across 12

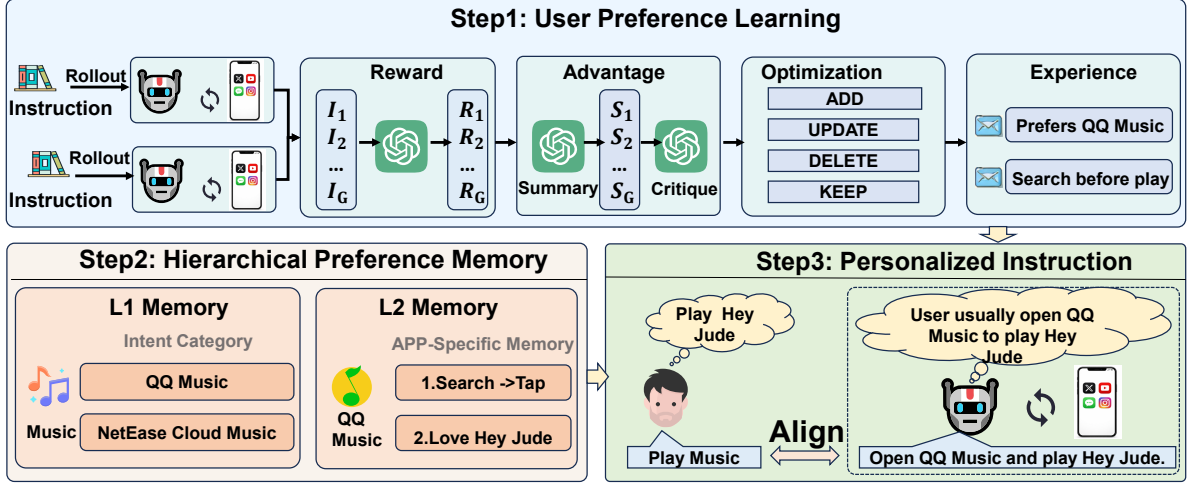


Figure 2: Overview of the proposed Me Agent framework.

categories, with 300 Type I and 267 Type II instructions in total. We partition the dataset at the user level: for each user, 15 instructions are randomly sampled for training and 5 for testing. The training set enables agents to learn user preferences, while the test set evaluates preference reasoning capabilities. Training and test instructions are non-overlapping to ensure evaluation validity.

4 Methodology

4.1 Problem Formulation

Given user instruction I , user preference profile U , historical action sequence $A_{\text{history}} = \{a_1, \dots, a_t\}$, current screenshot observation O_t , and available action space \mathcal{A}_t , the mobile agent predicts the next action as:

$$a_{t+1} = f(I, A_{\text{history}}, O_t, \mathcal{A}_t, U), \quad (4)$$

where $f(\cdot)$ denotes the agent. It iterates this process until instruction I is completed.

4.2 Me-Agent

Current LLM-based agents (Wang et al., 2025b; Maragheh et al., 2025) face the following challenges: 1) a lack of capability for modeling user preferences, making it difficult to understand users' vague intent in instructions; 2) limitations in mobile training, as the computational power of mobile devices is insufficient to support local training; and 3) difficulties in tracking the dynamic evolution of preferences, due to a lack of cross-iteration continuous learning mechanisms, which hampers the ability to update user behavior dynamically, resulting in a decline in personalization effectiveness over

time. To address the above issues, we propose Me-Agent, a training-free personalized mobile agent framework. As shown in the Figure 2, Me-Agent is composed of **User Preference Learning** and **Hierarchical Preference Memory**.

4.3 User Preference Learning

Inspired by Training-Free GRPO (Cai et al., 2025), we propose User Preference Learning (UPL), a parameter-free module for modeling user preferences. Specifically, UPL includes two targeted improvements: (1) a VLM-based reward model that evaluates task completion from visual observation, and (2) predefined experience categories for structured preference accumulation. Built upon these designs, UPL operates through four stages: Rollout, Reward, Advantage, and Optimization.

Rollout. For each user instruction I , we perform G independent rollouts using the mobile agent $f(\cdot)$ to obtain diverse execution trajectories that provide comparative samples for subsequent intra-group analysis:

$$\{\tau_1, \tau_2, \dots, \tau_G\} \sim f(I, \mathcal{E}), \quad (5)$$

where \mathcal{E} denotes the current experience pool. Each trajectory $\tau_g = \{(O_t^g, A_t^g)\}_{t=1}^{T_g}$ represents a complete instruction execution, comprising a sequence of screenshots and corresponding actions. A sampling temperature $T > 0$ is applied to introduce exploration diversity, ensuring coverage of different operational paths.

Reward. We design a VLM-based reward model that leverages a Vision-Language Model to assess instruction execution from screenshot sequences.

The reward model offers several advantages: 1) It is easily accessible, extracting information directly from screenshots, facilitating data collection across various application scenarios. 2) It provides enriched feedback, reflecting instruction execution completeness levels and catering to personalized information. It is defined as:

$$r_g = R_{\text{VLM}}(I, \{O_1^g, O_2^g, \dots, O_{T_g}^g\}), \quad (6)$$

where I denotes the instruction and $\{O_1^g, \dots, O_{T_g}^g\}$ denotes the screenshot sequence. The VLM evaluates four dimensions: goal achievement, step validity, result visibility, and error detection. The reward $r_g \in [0, 1]$ provides a continuous signal for fine-grained trajectory quality assessment.

Advantage. The personalized mobile agent needs not only to identify which trajectory is superior but also to understand user preference patterns. By having an LLM analyze successful and failed trajectories, we extract natural language experiences that are interpretable, editable, and can be directly injected into prompts to guide subsequent decisions. Given the trajectory set $\{(\tau_g, r_g)\}_{g=1}^G$ for instruction I , we extract experiences through two phases:

Trajectory Summarization. Each trajectory is summarized into structured experiences:

$$S_g = f(I, \tau_g, r_g). \quad (7)$$

Comparative Critique. All operational summaries for the same instruction are compared to extract generalizable experiences:

$$\mathcal{E}_{\text{new}}^I = f(I, \{S_g\}_{g=1}^G). \quad (8)$$

The critique process focuses on six experience categories: user preferences, UI navigation, action sequences, element identification, context awareness, and completion signals.

Optimization. As learning iterates, newly extracted experiences may overlap or conflict with existing ones. Direct appending leads to redundant pool expansion, while naive overwriting risks losing valuable historical memory. To address this, we design an integration mechanism that enables agents to compare new experiences against the existing pool and assigns each one an operation:

$$\{op_j\} = f(\mathcal{E}_{\text{old}}, \mathcal{E}_{\text{new}}), \quad (9)$$

where $op_j \in \{\text{ADD}, \text{UPDATE}, \text{Delete}, \text{KEEP}\}$ denotes inserting new information, refining existing experiences, deleting redundant ones, or remaining unchanged, respectively.

4.4 Hierarchical Preference Memory

The User Preference Learning module produces an experience pool \mathcal{E} that stores general user preferences and is directly injected into prompts. However, app-specific knowledge is extensive and only relevant when interacting with the corresponding app. Injecting all such memory into prompts would cause context overflow and attention dilution. To address this, we introduce Hierarchical Preference Memory (HPM), which stores application-specific memory externally and retrieves it as needed for each target application.

Personalized Memory. Mobile instructions exhibit a natural hierarchical structure: users first select an app category (e.g., music, navigation), then perform app-specific operations. We design a hierarchical memory accordingly:

Level-1 Memory \mathcal{M}^{L1} is organized by intent categories:

$$\mathcal{M}^{L1} = \{(c_i, \mathcal{A}_i)\}_{i=1}^C, \quad (10)$$

where c_i denotes the category name and \mathcal{A}_i is the set of apps under this category.

Level-2 Memory \mathcal{M}^{L2} stores app-specific memory for each app:

$$\mathcal{M}_{\text{app}}^{L2} = \{\text{workflows, preferences}\}, \quad (11)$$

including workflows (instruction execution patterns and UI element positions) and the content preferences (user-specific preferences within the app).

For successful trajectories with reward $r_g \geq \theta$, we use an LLM to extract workflows and content preferences from the execution trace. The memory supports dynamic updates: similar workflows are merged based on embedding similarity and success counts, while content preferences are adjusted using frequency statistics to reflect evolving user habits.

Personalized Inference. During inference, the mobile agent must resolve two types of ambiguity in user instructions: *application ambiguity* (which app to use) and *content ambiguity* (what specific content to operate on). We design a two-stage inference mechanism that leverages the learned experience pool \mathcal{E} and user memory \mathcal{M} .

Application Resolution. When instruction I explicitly specifies an application, the agent uses it directly. Otherwise, it infers the target application

Type	Method	ASA \uparrow	BERTScore \uparrow	PS \uparrow	HSS \uparrow	TCR \uparrow	TSR \uparrow	RF \uparrow	AF \uparrow	Step \downarrow
Type I	Mobile-Agent V2	0.600	—	—	—	0.369	0.200	0.464	0.698	—
	Mobile-Agent E	0.400	—	—	—	0.763	0.600	0.917	0.950	—
	Me-Agent	1.000	—	—	—	0.960	0.800	0.870	0.855	—
Type II	Mobile-Agent V2	0.200	0.526	0.265	0.428	0.778	0.300	0.838	0.898	10.80
	Mobile-Agent E	0.200	0.549	0.275	0.682	0.543	0.125	0.838	0.900	7.60
	Me-Agent	1.000	0.725	0.534	0.890	0.888	0.500	0.935	0.876	6.20

Table 3: Personalized evaluation results on the User FingerTip benchmark.

Method	TCR \uparrow	TSR \uparrow	AF \uparrow	RP \uparrow	Step \downarrow
Mobile-Agent-v2	0.414	0.200	0.467	0.570	20.20
Mobile-Agent-E	0.639	0.667	0.617	0.702	18.40
Me-Agent	0.893	0.600	0.861	0.978	12.53

Table 4: Performance evaluation on the E dataset.

based on user preferences:

$$a^* = \begin{cases} a, & \text{if } a \in I, \\ \arg \max_{a \in \mathcal{A}_{\text{cat}(I)}} P(a | \mathcal{E}), & \text{otherwise.} \end{cases} \quad (12)$$

where $\mathcal{A}_{\text{cat}(I)}$ denotes the set of applications within the functional category of instruction I , and $P(a | \mathcal{E})$ represents the user’s preference probability for application a derived from the experience pool.

Content Retrieval. When the instruction contains implicit references (e.g., “my favorite”), the mobile agent retrieves relevant content from HPM through a two-step process. First, we retrieve the top- k candidate contents $\mathcal{C}_{\text{cand}}$ based on semantic similarity:

$$\mathcal{C}_{\text{cand}} = \text{Top}_k(\text{sim}(\mathbf{e}_I, \mathbf{e}_c) | c \in \mathcal{M}_{a^*}), \quad (13)$$

where \mathcal{M}_{a^*} denotes the memory associated with the target application and $\text{sim}(\cdot, \cdot)$ measures the semantic similarity between embeddings. Then, the LLM selects the most appropriate content by reasoning over the candidates and user history:

$$c^* = f(I, \mathcal{C}_{\text{cand}}). \quad (14)$$

Prompt Injection. The resolved application a^* , retrieved content c^* , and relevant experiences from \mathcal{E} are injected into the prompt to enable personalized execution based on user preferences.

5 Experiments

5.1 Experimental Setup

Dataset & Metrics. We assess personalization capabilities on User FingerTip and general capabilities on E Datasets (Wang et al., 2025b). We

evaluate agents from two perspectives. **Personalization metrics** include: **App Selection Accuracy (ASA)**, measuring whether the agent selects user-preferred apps among functionally equivalent options; **BERTScore**, computing semantic similarity between agent outputs and ground-truth references (Type II only); **Preference Score (PS)**, evaluating reasoning alignment with user preferences via LLM scoring (Type II only); and **Human Satisfaction Score (HSS)**, a weighted composite of app selection correctness (0.4), content match correctness (0.4), and task completion (0.2). **Performance metrics** comprise: **Task Completion Ratio (TCR)**, measuring the proportion of achieved sub-goals; **Task Success Rate (TSR)**, a binary metric requiring all requirements to be fulfilled; **Action Fidelity (AF)**, assessing correctly executed atomic actions; **Reflection Precision (RP)**, quantifying the rate of correct diagnoses during error recovery; and **Step**, the average number of inference steps per instruction.

Baseline. We compare our approach against two baselines: Mobile-Agent-v2 (Wang et al., 2024a) and Mobile-Agent-E (Wang et al., 2025b). To ensure a fair comparison, we evaluate Me-Agent against the baselines using the same atomic action space and perception models. We use DBNet (Liao et al., 2020) and the ModelScope ConvNextViT-document for OCR detection and recognition. We utilize GroundingDINO (Liu et al., 2024) for icon grounding and GPT-4o (Hurst et al., 2024) for generating captions for each cropped icon.

Implementation Details. All experiments are conducted on standardized cloud infrastructure. We set the group size to 2 and the sampling temperature to 0.3. Following Training-Free GRPO (Cai et al., 2025), we perform 2 epochs with a batch size of 5 over the training queries. In each update, distilled experience is written into \mathcal{E} and \mathcal{M} via natural-language operations, while all LLM/VLM

Type	Backbone	ASA \uparrow	BERTScore \uparrow	PS \uparrow	HSS \uparrow	TCR \uparrow	TSR \uparrow	RF \uparrow	AF \uparrow
Type I	GPT-4o	1.00	—	—	—	0.976	0.833	0.933	1.000
	Claude-3.5-Sonnet	1.00	—	—	—	0.920	1.000	0.935	0.975
	Gemini-2.5-Pro	1.00	—	—	—	1.000	0.958	1.000	1.000
Type II	GPT-4o	1.00	0.554	0.800	0.428	0.900	0.600	0.825	0.825
	Claude-3.5-Sonnet	1.00	0.534	0.850	0.682	0.697	0.400	0.733	0.783
	Gemini-2.5-Pro	1.00	0.534	0.750	0.890	0.893	0.937	0.933	0.893

Table 5: Evaluation under different reasoning backbones.

backbones, the reward model, and perception modules remain frozen. To ensure reproducibility, we maintain identical software versions and hardware configurations across all experimental runs.

5.2 Personalized Evaluation

We conduct a comprehensive evaluation of Me-Agent on the User FingerTip. As shown in Table 3, It significantly outperforms Mobile-Agent-v2 and Mobile-Agent-E on personalization metrics, while also demonstrating superior performance on task execution metrics.

Personalization Capabilities. Equipped with UPL and HPM, Me-Agent learns preferences from user historical behavior, achieving an ASA of 1.000 on both Type I and Type II instructions. In contrast, Mobile-Agent-v2 and Mobile Agent E lack preference modeling capabilities, which likely contributes to their lower ASA when facing multiple applications within the same functional category (e.g., QQ Music vs. NetEase Music). Furthermore, Me-Agent achieves notably higher BERTScore and PS on Type II instructions, indicating its ability to infer user content preferences. Consistently, Me-Agent also attains the highest HSS (0.890 vs. 0.428 for Mobile-Agent-v2), and its ranking alignment with PS further validates that our preference modeling translates into objectively verifiable user satisfaction. These results confirm the learning and memory capabilities of the Me-Agent in user personalization: it captures application-level preferences through an experience pool, while hierarchical memory enables personalization at the content level.

Instruction Execution Performance. Me-Agent achieves the highest TCR and TSR across both instruction types. Notably, as instruction complexity increases from Type I to Type II, the baselines show substantial performance degradation, while Me-Agent maintains stable performance. We at-




Figure 3: Ablation Study on the User Fingertip.

tribute this robustness to the HPM module, which records UI element positions and action sequences from successful interactions, enabling the agent to reuse proven operational paths rather than exploring complex interfaces from scratch.

5.3 Performance Evaluation

We evaluate the performance of Me-Agent on the E dataset (Wang et al., 2025b). As shown in Table 4, Me-Agent achieves the best performance on Task Completion Rate (0.893), Action Accuracy (0.861), and Reflection Precision (0.978). For Task Success Rate, Me-Agent obtains 0.600, which is slightly lower than Mobile-Agent-E (0.667) but substantially higher than Mobile-Agent-v2 (0.200).

These improvements stem from Me-Agent’s experience learning mechanism. Baselines often fail due to inaccurate UI element recognition rather than incorrect planning. Me-Agent addresses this through UPL, which extracts structured knowledge—action sequences, UI navigation paths, and completion signals—by comparing successful and failed attempts, and HPM, which organizes this knowledge by application for on-demand retrieval during inference.

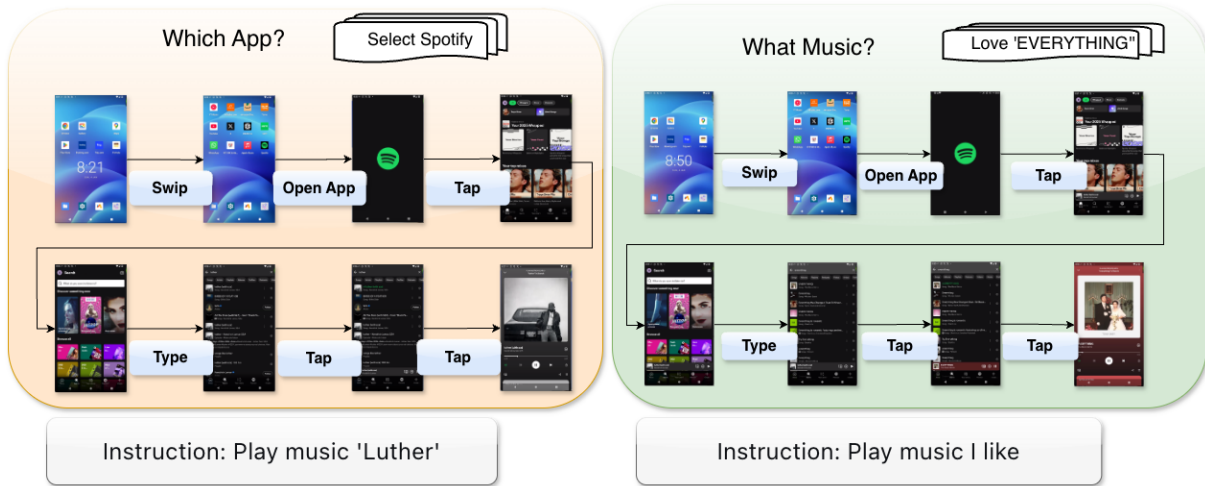


Figure 4: Example cases of Me-Agent executing personalized mobile instructions

5.4 Efficiency Evaluation

We further evaluate runtime efficiency using average step count, which directly reflects cumulative LLM/VLM calls since per-step cost is uniform across methods. As shown in Tables 3 and 4, Me-Agent requires fewer steps on both datasets (e.g., 12.53 vs. 18.40/20.20 on E Dataset), as the experience pool and HPM supply structural priors that guide near-optimal execution paths.

5.5 Ablation Study

We conduct ablation studies to analyze the contributions of the proposed **UPL** and **HPM** modules. The base method is vanilla Mobile-Agent-E. We then incrementally incorporate the **UPL** module (**BASE+UPL**), the **HPM** module (**BASE+HPM**), and their combination (**BASE+UPL+HPM**). As shown in Figure 3: 1) Mobile-Agent-E exhibits clear limitations in personalized settings, validating our motivation that current agents lack the ability to model user preferences. 2) UPL leads to a substantial improvement in task completion, demonstrating the effectiveness of shifting policy optimization from parameter space to context space. 3) HPM significantly enhances semantic quality and preference alignment, indicating that hierarchical memory enables more accurate resolution of ambiguous instructions through on-demand retrieval of application-specific knowledge. 4) The two modules are complementary: UPL captures general user preferences, while HPM provides fine-grained execution knowledge. Their combination achieves the best overall performance.

5.6 Impact of Backbone Models

We evaluate Me-Agent using three backbone models: **GPT-4o**, **Claude-3.5-Sonnet**, and **Gemini-2.5-Pro**. As shown in Table 5, our method achieves stable personalized evaluations and task performance across all models for both Type I and Type II tasks. This robustness stems from the modular design of Me-Agent: the UPL module encodes preferences as natural language experiences that are compatible with any LLM, while the HPM module leverages structured retrieval to reduce dependence on backbone-specific capabilities.

5.7 Case Study

Figure 4 illustrates Me-Agent’s personalization capabilities. For Play the song ‘Luther,’” Me-Agent directly selects the correct app and track from historical habits, unlike baselines that request clarification (Li et al., 2024c; Wang et al., 2025b). For vague instructions like Play the music I like,” it recommends preference-aligned content by synthesizing listening history, eliminating redundant confirmations.

6 Conclusions

We propose Me-Agent, a personalized mobile agent that incorporates User Preference Learning for parameter-free preference modeling in context space, and Hierarchical Preference Memory for on-demand retrieval of app-specific knowledge. We also introduce User FingerTip, a benchmark with ambiguous daily-life instructions. Experiments show that Me-Agent outperforms baselines in both personalization and task performance.

Acknowledgements

The work is supported by the National Natural Science Foundation of China (62506050), China Postdoctoral Science Foundation Funded Project (2024M763867). We are deeply grateful to our advisor Yu Tian for his extraordinary patience with students and his creative insights into research, both of which will continue to inspire us in our future work.

Ethical Considerations

User FingerTip is built by extending FingerTip (Yang et al., 2025): we inherit its 60 users, 33 applications, and 12 functional categories, and construct 300 Type I and 267 Type II ambiguous instructions to probe implicit preference reasoning—a capability not assessed by the original benchmark. All user behavior data is derived from the source benchmark under its original terms; no additional personal information is collected.

For human evaluation, we recruited PhD students in computer science as evaluators. Detailed evaluation guidelines, including scoring rubrics and illustrative examples, were provided to all evaluators prior to the evaluation. The evaluation task involved only assessing agent outputs on predefined metrics and did not involve any sensitive or personal information.

Limitations

UI Dynamics and Temporal Validity . We extract UI positions and action sequence from successful task executions. However, these learned experiences may become invalid when the application is updated or unexpected pop-ups appear. Such UI changes can cause execution failures because the stored positions no longer match the current interface.

Limited Contextual Factors in Personalization

. Our personalization approach mainly relies on users' historical preferences and behavior patterns. It does not consider dynamic factors such as the user's current location or emotional state. This may lead to poor recommendations when the user's current situation differs from their usual patterns.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. Claude 3.5 sonnet. <https://www.anthropic.com/news/3-5-models-and-computer-use>. Accessed: 2026-01-03.
- Yuzheng Cai, Siqi Cai, Yuchen Shi, Zihan Xu, Lichao Chen, Yulei Qin, Xiaoyu Tan, Gang Li, Zongyi Li, Haojia Lin, and 1 others. 2025. Training-free group relative policy optimization. *arXiv preprint arXiv:2510.08191*.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms, 2023. URL <https://arxiv.org/abs/2305.14314>, 2.
- Wenqi Fan, Pangjing Wu, Yujuan Ding, Liangbo Ning, Shijie Wang, and Qing Li. 2025. Towards retrieval-augmented large language models: Data management and system design. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, pages 4509–4512. IEEE.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Aobo Kong, Shiwan Zhao, Hao Chen, Qicheng Li, Yong Qin, Ruiqi Sun, Xin Zhou, Enzhi Wang, and Xiaohang Dong. 2024. Better zero-shot reasoning with role-play prompting. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4099–4113.
- Xinyu Li, Ruiyang Zhou, Zachary C Lipton, and Liu Leqi. 2024a. Personalized language modeling from personalized human feedback. *arXiv preprint arXiv:2402.05133*.
- Xinyu Li, Ruiyang Zhou, Zachary C Lipton, and Liu Leqi. 2024b. Personalized language modeling from personalized human feedback. *arXiv preprint arXiv:2402.05133*.
- Yanda Li, Chi Zhang, Wenjia Jiang, Wanqi Yang, Bin Fu, Pei Cheng, Xin Chen, Ling Chen, and Yunchao Wei. 2024c. Appagent v2: Advanced agent for flexible mobile interactions. *arXiv preprint arXiv:2408.11824*.

- Minghui Liao, Zhaoyi Wan, Cong Yao, Kai Chen, and Xiang Bai. 2020. Real-time scene text detection with differentiable binarization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11474–11481.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, and 1 others. 2024. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European conference on computer vision*, pages 38–55. Springer.
- Gowen Loo, Chang Liu, Qinghong Yin, Xiang Chen, Jiawei Chen, Jingyuan Zhang, and Yu Tian. 2025. Mobilerag: Enhancing mobile agent with retrieval-augmented generation. *arXiv preprint arXiv:2509.03891*.
- Reza Yousefi Maragheh, Pratheek Vadla, Priyank Gupta, Kai Zhao, Aysenur Inan, Kehui Yao, Jianpeng Xu, Praveen Kanumala, Jason Cho, and Sushant Kumar. 2025. Arag: Agentic retrieval augmented generation for personalized recommendation. *arXiv preprint arXiv:2506.21931*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Sriyash Poddar, Yanming Wan, Hamish Ivison, Abhishek Gupta, and Natasha Jaques. 2024. Personalizing reinforcement learning from human feedback with variational preference learning. *Advances in Neural Information Processing Systems*, 37:52516–52544.
- Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, and 1 others. 2025. Uitars: Pioneering automated gui interaction with native agents. *arXiv preprint arXiv:2501.12326*.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. Lamp: When large language models meet personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392.
- Teng Shi, Jun Xu, Xiao Zhang, Xiaoxue Zang, Kai Zheng, Yang Song, and Han Li. 2025. Retrieval augmented generation with collaborative filtering for personalized text generation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1294–1304.
- Junyang Wang, Haiyang Xu, Haitao Jia, Xi Zhang, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. 2024a. Mobile-agent-v2: Mobile device operation assistant with effective navigation via multi-agent collaboration. *Advances in Neural Information Processing Systems*, 37:2686–2710.
- Junyang Wang, Haiyang Xu, Jiabo Ye, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. 2024b. Mobile-agent: Autonomous multi-modal mobile device agent with visual perception. *arXiv preprint arXiv:2401.16158*.
- Xin Wang, Zhiyao Cui, Hao Li, Ya Zeng, Chenxu Wang, Ruiqi Song, Yihang Chen, Kun Shao, Qiaosheng Zhang, Jinzhuo Liu, and 1 others. 2025a. Perpillot: Personalizing vlm-based mobile agents via memory and exploration. *arXiv preprint arXiv:2508.18040*.
- Zhenhailong Wang, Haiyang Xu, Junyang Wang, Xi Zhang, Ming Yan, Ji Zhang, Fei Huang, and Heng Ji. 2025b. Mobile-agent-e: Self-evolving mobile assistant for complex tasks. *arXiv preprint arXiv:2501.11733*.
- Qinglong Yang, Haoming Li, Haotian Zhao, Xiaokai Yan, Jingtao Ding, Fengli Xu, and Yong Li. 2025. Fingertip 20k: A benchmark for proactive and personalized mobile llm agents. *arXiv preprint arXiv:2507.21071*.
- Bingsheng Yao, Bo Sun, Yuanzhe Dong, Yuxuan Lu, and Dakuo Wang. 2025. Dprf: A generalizable dynamic persona refinement framework for optimizing behavior alignment between personalized llm role-playing agents and humans. *arXiv preprint arXiv:2510.14205*.
- Jiabo Ye, Xi Zhang, Haiyang Xu, Haowei Liu, Junyang Wang, Zhaoqing Zhu, Ziwei Zheng, Feiyu Gao, Junjie Cao, Zhengxi Lu, and 1 others. 2025. Mobile-agent-v3: Fundamental agents for gui automation. *arXiv preprint arXiv:2508.15144*.
- Chi Zhang, Zhao Yang, Jiaxuan Liu, Yanda Li, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2025a. Appagent: Multimodal agents as smartphone users. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–20.
- Xiaotian Zhang, Ruizhe Chen, Yang Feng, and Zuozhu Liu. 2025b. Persona-judge: Personalized alignment of large language models via token-level self-judgment. *arXiv preprint arXiv:2504.12663*.
- Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, and 1 others. 2024. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*.
- Yuchen Zhuang, Haotian Sun, Yue Yu, Rushi Qiang, Qifan Wang, Chao Zhang, and Bo Dai. 2024. Hydra: Model factorization framework for black-box llm personalization. *Advances in Neural Information Processing Systems*, 37:100783–100815.

A Prompt

In this appendix, we provide the system and user prompts used in mobile UI automation tasks. Tables 6 and 7 present the prompts for trajectory analysis and single-task experience extraction. Tables 8 and 9 are used for experience integration and global experience consolidation. Table 10 works for guiding task execution based on learned experiences.

B Examples of Learned Experiences

In this appendix, we provide some examples in table 11, which are extracted from 49 learned experiences by Me-Agent.

C Examples of Dataset

In this appendix, we provide an example user from User FingerTip, including both training and testing sets. Table 12 shows the training set tasks for a randomly selected user, with each task specifying the instruction, intent category, and target app. Table 13 presents the corresponding testing set tasks, where each test task is divided into Type I (APP-Ambiguous Instruction) and Type II(Content-Ambiguous Instruction), along with the target app and the ground truth content.

Table 6: The Prompt for Rollout

System

You are an expert in mobile phone operations and UI automation. Your task is to analyze a mobile operation trajectory and summarize what worked and what did not, focusing on concise and actionable insights.

User

Task Info: App: {app}; Intent: {intent_category}; Instruction: {instruction}.

Mobile Operation Trajectory: A sequence of steps, each containing *Thought* (planned intent), *Action* (Tap, Swipe, Type, etc.), and *Summary* (on-screen result or observation).
{trajectory}

Your Task: Analyze the trajectory and provide a concise summary (**under 200 words**) covering:

- (1) **Successful Actions:** UI interactions that worked as intended;
- (2) **Failed Actions:** interactions that did not achieve the goal;
- (3) **UI Understanding:** correctness of UI element identification and interaction;
- (4) **Navigation Pattern:** whether the navigation sequence was reasonable;
- (5) **User Patterns:** recurring app preferences or frequent behaviors;
- (6) **Anti-Patterns:** actions leading to loops or failures, repeated ineffective taps, or cases where continued interaction was counterproductive.

Output Format: Step-by-step summary:

Step 1: [Action and outcome]

Step 2: [Action and outcome]

...

Overall: [Key insights explaining what worked or failed and why].

Table 7: The Prompt for Single Query Critique

System

You are an expert in mobile UI automation and user behavior analysis. Your task is to compare multiple attempts at the same mobile task and extract actionable lessons, with special focus on learning user preferences and habits.

CRITICAL: Output MUST be wrapped in <Experiences></Experiences> tags.

User

Task Info: App: {app}; Intent: {intent_category}; Instruction: {instruction}.

Multiple Attempts: Summaries of repeated executions: {attempts}.

Your Task: Compare attempts and extract **5–8 actionable experiences** for future tasks, focusing on:

(1) **User Preferences** (priority): preferred apps, frequent destinations, shopping/music habits, routines, inferred percentages;

(2) **UI Navigation:** locations of key UI elements and optimal paths;

(3) **Action Sequences:** correct operation order, timing, recovery;

(4) **Element Identification:** distinguishing similar or dynamic UI elements;

(5) **Context Awareness:** app-specific behaviors, success/failure cases;

(6) **Task Completion:** screens/states indicating completion and when to STOP.

Output Format (MANDATORY):

<Experiences>

1. [User Preference – {intent_category}] <specific preference (percentages if known)>

2. [UI Navigation – {app}] <key UI element locations>

3. [Action Sequence – {intent_category}] <operation order>

4. [Element Identification – {app}] <element identification guidance>

5. [Context] <when approaches work or fail>

6. [Task Completion – {intent_category}] <completion signals>

</Experiences>

Reminders: start with user preferences; one experience per line; be specific and actionable; include percentages when inferable.

Table 8: The Prompt for Group Experience Update

System

You are an expert in knowledge management for mobile automation. Your task is to process a batch of new experiences from a single task and decide how to integrate them with existing experiences. **CRITICAL:** Each line in the new experiences is ONE independent experience and must be processed separately. You must output valid JSON only, with no additional text or explanations.

User

Existing Experiences Database:

{existing_experiences}

New Experiences: Newly extracted experiences from recent task attempts. Each numbered line represents ONE independent experience.

{new_experiences}

Your Task: For EACH new experience, choose exactly one operation:

ADD: completely new information not covered by existing experiences;

UPDATE: refines, corrects, or expands an existing experience (specify the ID to replace);

DISCARD: redundant with existing experiences or not useful.

Important Rules:

- User Preference experiences: **ALWAYS ADD or UPDATE;**
- Similar experiences with new information (e.g., new items or updated percentages): use UPDATE;
- Completely redundant experiences: use DISCARD;
- Each operation must reference exactly ONE new experience line.

Output Format (JSON ONLY): Return ONLY valid JSON in the following structure:

```
[
  { "operation": "ADD", "content": "[User Preference - Music] User consistently uses QQ Music for music playback", "reason": "New preference" },
  { "operation": "UPDATE", "id": "G3", "content": "[UI Navigation - QQ Music] Search icon at top-right; tap to activate keyboard", "reason": "Refines existing experience" },
  { "operation": "DISCARD", "content": "[Context] Wait for page to load", "reason": "Redundant information" }
]
```

CRITICAL: Output must be valid JSON parsable by `json.loads()`, with no text before or after the array.

Table 9: The Prompt for Batch Experience Update

System

You are an expert in knowledge management for mobile automation. Your task is to create a comprehensive, non-redundant experience base that captures both operational knowledge and user preferences. You must output valid JSON only, with no additional text or explanations.

User

Current Experiences and Proposed Operations:

{experiences_and_operations}

Your Task: Review all proposed operations and produce a final revision plan that:

(1) resolves conflicts between operations; (2) removes redundancies by merging similar experiences; (3) ensures experiences are specific and actionable; (4) maintains a clean and organized experience base; (5) prioritizes user preference experiences, which should always be preserved and updated.

Guidelines for High-Quality Experiences:

- Keep each experience concise (1–2 sentences, max 50 words);
- Each experience should be independently useful;
- Prefer specific, actionable guidance over generic advice;
- Use consistent prefixes: [Category] or [Category – SubCategory];
- User Preference experiences MUST use format: [User Preference – IntentCategory].

Output Format (JSON ONLY): Return ONLY valid JSON in the following structure:

```
[
  { "operation": "ADD", "content": "New experience with proper [Category] prefix" },
  { "operation": "UPDATE", "id": "G2", "content": "Updated experience replacing G2"
},
  { "operation": "DELETE", "id": "G5" }
]
```

Operation Types:

ADD: create a new experience; **UPDATE:** replace an existing experience by ID; **DELETE:** remove an existing experience by ID.

CRITICAL RULES:

- Output must be valid JSON parsable by `json.loads()`;
 - Do not include any text before or after the JSON array;
 - All content fields must be complete experience strings with [Category] prefixes;
 - IDs must exactly match existing experience IDs (e.g., G2, G15);
 - Never delete user preference experiences unless absolutely necessary.
-

Table 10: The Prompt for Problem Solving with Learned Experiences

Learned Experiences

{experiences}

Your Task

{problem}

Important Instructions

Use the learned experiences above to guide your operations. They contain valuable insights from previous successful attempts and documented user preferences.

Pay special attention to User Preference experiences- they tell you which apps, locations, or items the user commonly prefers. When multiple options are available, ALWAYS choose based on documented user preferences. **Example:**

If experiences indicate “User prefers QQ Music (75%)”, you should use QQ Music for music-related tasks by default.

Table 11: Sample Experiences Extracted from User Interaction Trajectories (ID included)

Experience (ID + English Translation)
G0: [User Preference - Movies] User prefers accessing movie and TV series content through Douban’s “Books, Movies, and TV” section for ratings and rankings.
G1: [UI Navigation - Douban] Navigate to “Books, Movies, and TV”, then “Douban Rankings”, and switch to “Movie Rankings” or “TV Series Rankings” for rankings.
G3: [Error Avoidance] Use “Tap_Type_and_Enter” to efficiently input text and avoid repeated tapping; ensure the search bar is engaged before typing or using shortcuts.
G5: [UI Navigation - Xiaohongshu] The search icon is located at the top-right corner, serving as a reliable starting point for searches.
G8: [User Preference - Shopping] User prefers the JD app for precise product searches such as MacBook Air.
G10: [Action Sequence - JD] Open the JD app → Tap the search bar → Clear existing text → Type “latest iPad” → Execute the search.
G11: [Task Completion - JD] The task is complete when the search results display the latest iPad models.
G23: [User Preference - Music] User prefers QQ Music for music-related tasks, indicating a strong preference for this app.
G24: [Action Sequence - Music] Open QQ Music → Tap the search bar → Enter the song name or select from search history → Tap the play button next to the song.
G25: [Task Completion - Music] The task is complete when the song begins playing.
G34: [User Preference - Reading] User prefers using the WeChat Reading app for novels, indicating a strong preference for this application.
G42: [User Preference - Posts] User frequently searches for “weight loss tips” in Xiaohongshu, indicating a strong interest in this content.
G43: [Action Sequence - Posts] Open Xiaohongshu → Tap the search icon → Utilize search history by selecting “weight loss tips”.
G44: [Context] Utilizing search history is effective when the user has previously searched for the same term.
G45: [Task Completion - Posts] The task is complete when the search results for “weight loss tips” are displayed.

Table 12: Training set instructions for an example user

Task ID	Instruction	Intent Category	Target App
train_0	Open WeChat Reading to read The Three-Body Problem	Reading	WeChat Reading
train_1	Open WeChat Reading to read Xiangtu Zhongguo	Reading	WeChat Reading
train_2	Open Xiaohongshu to search posts about home bartending	Posts	Xiaohongshu
train_3	Open Xiaohongshu to search travel guides for Summer Palace	Posts	Xiaohongshu
train_4	Open Xiaohongshu to search fitness strategies and open the first one	Posts	Xiaohongshu
train_5	Open JD.com to check price of latest iPad	Shopping	JD.com
train_6	Open JD.com to search for price of monitors	Shopping	JD.com
train_7	Open JD.com to search for price of latest MacBook Air	Shopping	JD.com
train_8	Open Douban to check currently high-rated movies	Posts	Douban
train_9	Open QQ Music to play the song "My Future"	Music	QQ Music
train_10	Open QQ Music to play the song "Replicate Memories"	Music	QQ Music
train_11	Open QQ Music to play the song "The Best Day"	Music	QQ Music
train_12	Open Douban to check the currently highest-rated TV dramas and their reviews	Posts	Douban
train_13	Open Xiaohongshu to search for body fat rate related to fitness	Posts	Xiaohongshu
train_14	Open Xiaohongshu to search for fat loss tips	Posts	Xiaohongshu

Table 13: Testing set instructions for an example user.

Task ID	Intent Category	Type I Instruction	Type II Instruction	Target App	Target Content
test_1	Posts	Search for posts related to fitness strategies	Search for sports strategies that interest me	Xiaohongshu	Sports strategies the user likes: fitness-related activities
test_2	Shopping	Check the price of iPhone 16 Pro Max	Check the prices of digital products that interest me	JD.com	Digital products the user likes: iPad, monitor, Mac
test_3	Music	Play the song "Chun Ni (Spring Mud)"	Play music that I like	QQ Music	Music the user likes: My Future, Replicate Memories
test_4	Posts	Check the highest-rated movies of 2024 and their reviews	Check the movies I follow and their reviews	Douban	No restrictions, just open what the user follows
test_5	Reading	Read "Xiangtu Zhongguo"	Read books that I like	WeChat Reading	Books the user likes: The Three-Body Problem, Xiangtu Zhongguo