

EcoAgent: An Efficient Device-Cloud Collaborative Multi-Agent Framework for Mobile Automation

Biao Yi¹, Xueyu Hu¹, Yurun Chen¹, Shengyu Zhang^{1*}, Hongxia Yang², Fan Wu³

¹Zhejiang University

²The Hong Kong Polytechnic University

³Shanghai Jiao Tong University

{yi_biao, sy_zhang}@zju.edu.cn

Abstract

To tackle increasingly complex tasks, recent research on mobile agents has shifted towards multi-agent collaboration. Current mobile multi-agent systems are primarily deployed in the cloud, leading to high latency and operational costs. A straightforward idea is to deploy a device–cloud collaborative multi-agent system, which is nontrivial, as directly extending existing systems introduces new challenges: (1) reliance on cloud-side verification requires uploading mobile screenshots, compromising user privacy; and (2) open-loop cooperation lacking device-to-cloud feedback, underutilizing device resources and increasing latency. To overcome these limitations, we propose EcoAgent, a closed-loop device-cloud collaborative multi-agent framework designed for privacy-aware, efficient, and responsive mobile automation. EcoAgent integrates a novel reasoning approach, Dual-ReACT, into the cloud-based Planning Agent, fully exploiting cloud reasoning to compensate for limited on-device capacity, thereby enabling device-side verification and lightweight feedback. Furthermore, the device-based Observation Agent leverages a Pre-understanding Module to summarize screen content into concise textual descriptions, significantly reducing token usage and device-cloud communication overhead while preserving privacy. Experiments on Android-World demonstrate that EcoAgent matches the task success rates of fully cloud-based agents, while reducing resource consumption and response latency.

Code — <https://github.com/Yi-Biao/EcoAgent>

Extended version — <https://arxiv.org/abs/2505.05440>

Introduction

With the rapid advancement of (Multimodal) Large Language Models ((M)LLMs), mobile agents capable of autonomously interacting with smartphones have garnered increasing attention (Hu et al. 2024; Gao et al. 2024; Liu et al. 2025a; Wu et al. 2024a; Zhang et al. 2024a; Wang et al. 2024c; Chen et al. 2025c; Li et al. 2024a). A surge of mobile agents has been proposed based on (M)LLMs (Wang et al. 2024b; Wu et al. 2024b). In recent years, research on mobile agents (Wang et al. 2024a; Zhu et al. 2025; Wang

et al. 2025b) has further leveraged the reasoning capabilities of (M)LLMs by delegating planning, execution, and reflection tasks to distinct agents. This collaborative approach has significantly enhanced the ability of mobile agents to tackle complex tasks. However, due to the enormous parameter size of (M)LLMs, these agents predominantly operate in the cloud, resulting in high latency, high operational costs, and concerns about user privacy. Moreover, although some cloud-based agents (Qin et al. 2025; Team 2025) have been fine-tuned specifically for GUI tasks, the majority still rely on general-purpose (M)LLMs due to cost constraints. These cloud-based general agents cannot directly interact with UI elements and must depend on OCR and object detection modules for grounding, which further diminishes overall efficiency.

A natural solution is to deploy a device–cloud collaborative multi-agent system. Many recent works (Wu et al. 2024c; Cheng et al. 2024; Gou et al. 2025) have made preliminary attempts in this direction by employing cloud-based agents for planning and device-based agents for grounding, yielding improved performance. However, they face two major challenges: (1) Frequent uploading of mobile screenshots to the cloud for verification increases response latency and poses serious privacy risks. (2) Such collaboration is typically open-loop, characterized by one-way cloud-to-device instructions without device-to-cloud feedback, failing to fully exploit the potential of device–cloud cooperation to reduce operational costs and latency.

To address these challenges, we introduce **EcoAgent**, a closed-loop device–cloud collaborative multi-agent framework for privacy-aware, efficient, and responsive mobile automation. EcoAgent consists of a cloud-based **Planning Agent**, a device-based **Execution Agent**, and a device-based **Observation Agent**, working together in a feedback loop. As illustrated in Figure 1, EcoAgent achieves moderate latency and improved adaptability compared to existing agent architectures.

To support device-to-cloud feedback, the Planning Agent employs Dual-ReACT, a two-stage reasoning framework that first performs global ReACT-based planning over the instruction and initial screen, followed by local ReACT-based generation of step-by-step actions and corresponding expectations. On the device, the Execution Agent performs each step in the plan, while the lightweight Observation Agent

*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

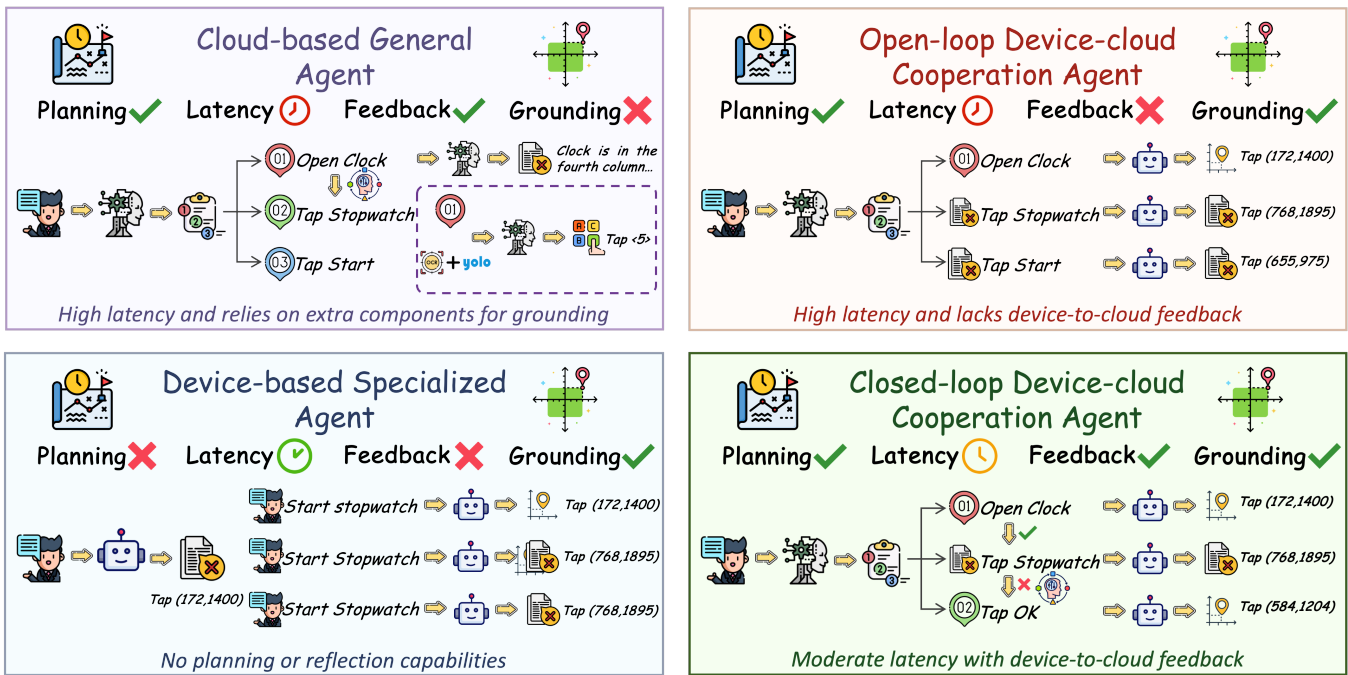


Figure 1: Comparison of four mobile agent architectures. Cloud-based General Agents exhibit strong planning abilities but suffer from high latency and rely on extra components for visual grounding. Device-based Specialized Agents achieve low latency and strong grounding but cannot plan or reflect. Open-loop Device-cloud Cooperation Agents enable grounding with precise control but lack device-to-cloud feedback. Closed-loop Device-cloud Cooperation Agents (**EcoAgent**) strike a balance by combining planning, feedback, and grounding, achieving moderate latency and improved adaptability.

compares the current screen with the expectation to verify execution status—without requiring strong reasoning. This enables efficient on-device verification and device-to-cloud feedback, with cloud assistance needed only when execution fails. To further reduce latency and protect user privacy, the Observation Agent incorporates a Pre-Understanding Module that compresses screenshots into concise textual summaries before transmission, significantly reducing token consumption and communication overhead. Finally, the Planning Agent integrates Memory and Reflection modules to support replanning in response to device-side feedback, enabling the system to adaptively recover from failure.

We evaluate EcoAgent on the widely-used benchmark AndroidWorld. Experimental results demonstrate that EcoAgent not only effectively handles complex mobile tasks, but also significantly reduces operational costs and response latency. Our contributions are summarized as follows:

- We propose **EcoAgent**, a closed-loop device–cloud collaborative framework that balances cloud-based reasoning with lightweight on-device verification, enabling privacy-aware and efficient mobile automation.
- We introduce **Dual-ReACT**, a novel two-stage planning mechanism that combines global and local ReACT reasoning to generate executable action plans with explicit expectations.
- We design a **Pre-Understanding Module** within the Observation Agent to compress visual inputs into textual

summaries, protecting user privacy and significantly reducing device–cloud communication overhead.

Related Work

Cloud-Based Mobile Agent

The rapid advancement of (M)LLMs has empowered mobile agents with strong capabilities in perception, reasoning, and decision-making. However, due to their massive model size and computational demands, (M)LLM-based mobile agents are typically deployed in the cloud, leading to significant latency and high operational costs, as illustrated in Figure 1. Moreover, as the vast majority of MLLMs are trained primarily on general natural images, cloud-based general agents (Zhang et al. 2023; Liu et al. 2024a; Wang, Li, and Li 2023; Wen et al. 2024c; Taeb et al. 2024; Lee et al. 2024; Huang et al. 2025; Zhang et al. 2024b; Liu et al. 2024b; Li et al. 2024b; Jiang et al. 2025; Yan et al. 2023; Song et al. 2024) struggle to accurately perceive and interact with mobile screens, often requiring extra components for UI element recognition and localization. For instance, Mobile-Agent (Liu et al. 2024a) integrates DINO (Liu et al. 2024a) and CLIP (Radford et al. 2021) recognition components, along with OCR-extracted text, to generate textual descriptions of screens, aiding MLLMs in screen understanding.

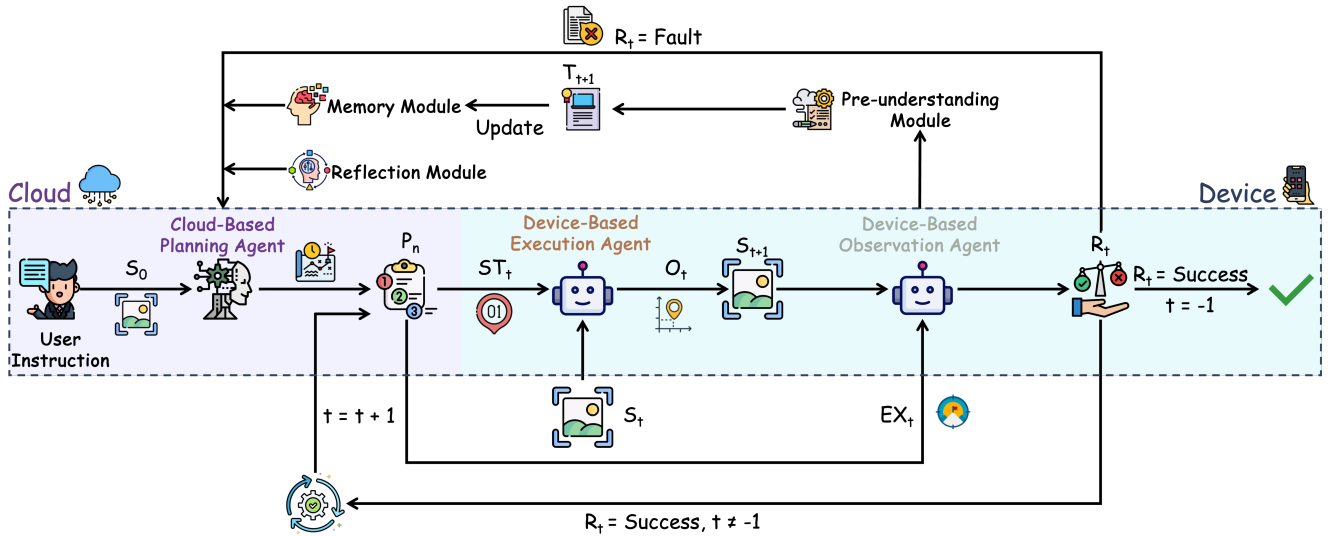


Figure 2: An overview of the EcoAgent framework. The Cloud-Based Planning Agent generates a task plan, which is executed and verified step-by-step by on-device agents. Feedback and pre-understood screen text from the Observation Agent enable dynamic replanning, forming a closed-loop system for efficient and adaptive task execution.

Device-Based Mobile Agent

To address the limitations of cloud-based general agents, recent studies (Hong et al. 2024; Lin et al. 2024; Liu et al. 2025b; Wu et al. 2024c; Wen et al. 2024b; Dai et al. 2025) have focused on fine-tuning multimodal Small Language Models ((M)SLMs) to enable fast and accurate mobile interactions. For example, ShowUI (Lin et al. 2024) proposes a vision-language-action model that integrates UI-guided visual token selection and interleaved vision-language-action streaming for efficient GUI manipulation. In addition, recent work (Wang et al. 2024d) has demonstrated the feasibility of deploying 7B MSLM directly on mobile devices, showing their potential for on-device automation. However, as illustrated in Figure 1, due to their limited model capacity, these device-based specialized agents often struggle with effective planning and reflection, making it difficult to handle complex mobile tasks.

Multi-Agent Mobile Agent

Due to the challenges single agents face in handling complex tasks, recent studies (Wang et al. 2024a; Zhu et al. 2025; Wang et al. 2025b; Chen et al. 2025b,a; Wang et al. 2025a) have explored enhancing agent capabilities through multi-agent collaboration. For example, M3A (Rawles et al. 2024) employs a dual-agent design, where one agent is responsible for decision-making and execution, while the other verifies execution status and performs reflection. However, cloud-based multi-agent systems like M3A further exacerbate latency and operational cost issues. A natural solution is device-cloud collaborative multi-agent systems. For instance, UGround (Gou et al. 2025) extends M3A by introducing an device-based grounding agent for execution, which improves task success rates. However, verification and planning still occur in the cloud, requiring frequent screenshot

uploads. As illustrated in Figure 1, such open-loop collaborations continue to suffer from high latency, high costs, and privacy concerns due to one-way cloud-to-device instruction flow without device-to-cloud feedback.

The EcoAgent Framework

In this section, we provide a detailed overview of the EcoAgent framework. EcoAgent adaptively coordinates cloud-based and device-based agents to execute user instructions, as illustrated in Figure 2. The framework consists of three specialized agents: a cloud-based Planning Agent and two device-based agents, the Execution Agent and the Observation Agent.

The overall workflow of EcoAgent is illustrated in Algorithm 1 and proceeds as follows. Firstly, the Planning Agent performs Dual-ReACT initial planning based on user instruction Ins and initial screen S_0 , generating an initial plan P_0 consisting of a sequence of steps ST_t and their corresponding expectations EX_t . The Execution Agent on the device side executes actions O_t based on the given step ST_t and the current screen S_t , obtains the next screen S_{t+1} . After each action, the Observation Agent evaluates whether the step was successfully executed by comparing the next screen S_{t+1} against the corresponding expectation EX_t to obtain the result R_t . To reduce token consumption and device-cloud communication overhead, the Observation Agent incorporates a Pre-Understanding Module, which converts high-token-cost images S_{t+1} into compact textual representations T_{t+1} and updates them to the memory module. If a step fails, the Planning Agent conducts Memory and Reflection empowered replanning, learning from the failure of the previous plan P_{n-1} to generate a new one P_n , and continues this process until the task is successfully completed.

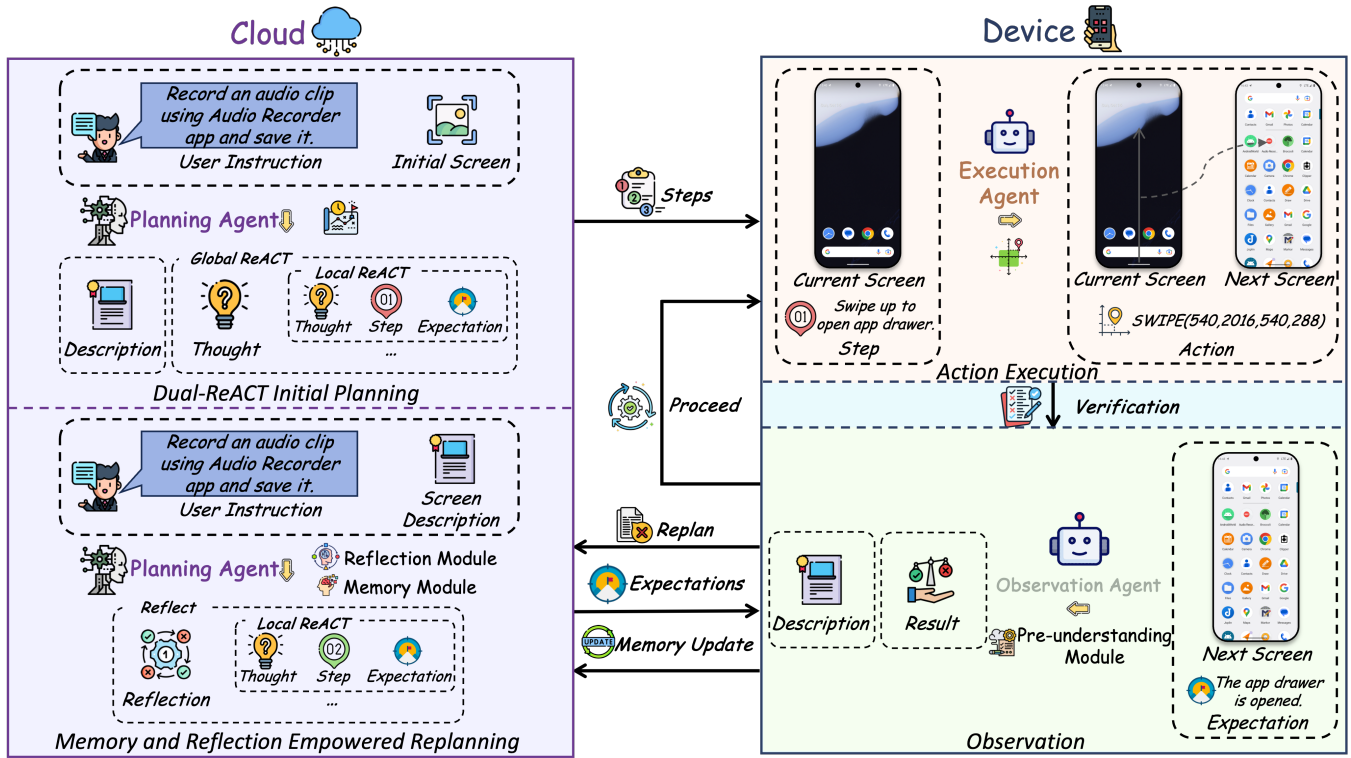


Figure 3: Illustration of the EcoAgent workflow. The cloud-based Planning Agent is responsible for Dual-ReACT initial planning as well as Memory and Reflection empowered replanning. The device-based Execution Agent and Observation Agent handle action execution and verification, respectively. These three agents collaborate to realize a closed-loop workflow.

Cloud-Based Planning Agent

Despite the difficulty MLLMs face in accurately identifying and grounding individual elements within mobile screens, their strong long-horizon planning and reflection capabilities make them well-suited for understanding screen context and adapting plans to dynamic environments. These strengths allow MLLMs to effectively decompose and solve complex mobile tasks—capabilities that MSLMs often lack. To leverage these strengths, we design a Cloud-based Planning Agent powered by MLLMs, which transforms high-level user instructions into adaptive and verifiable step-by-step plans. This agent operates through two primary mechanisms: Dual-ReACT Initial Planning and Memory and Reflection Empowered Replanning.

Dual-ReACT Initial Planning. Open-loop device–cloud cooperation agents often relies on cloud-side verification due to limited on-device capabilities, which makes it difficult to directly analyze execution outcomes from screenshots. However, general-purpose multimodal MSLMs are adept at describing screen content and judging whether it aligns with a given description. This motivates transforming complex task analysis into a simpler judgment by fully leveraging powerful cloud-side reasoning.

To this end, we propose a novel reasoning paradigm, Dual-ReACT, which extends the original ReACT framework (Yao et al. 2023) by applying reasoning and acting at both global and local levels.

As illustrated in Figure 3, given a user instruction Ins and an initial screen state S_0 , the Planning Agent first generates a global screen description to capture the overall context. Based on this, it conducts a Global ReACT process to derive a high-level execution plan by decomposing the task into intermediate subgoals. For each subgoal, a Local ReACT reasoning step determines how to achieve it, generating a concrete action step ST_t along with its expected outcome EX_t .

The resulting plan can be formulated as:

$$P_0 = \text{GIReACT}(Ins, S_0) = \{\text{LoReACT}(ST_1, EX_1), \dots, \text{LoReACT}(ST_t, EX_t)\}$$

This plan is then transmitted to the device, enabling effective device-to-cloud feedback and realizing a closed-loop.

Memory and Reflection Empowered Replanning. When execution fails, the Planning Agent initiates a memory-driven reflection process as shown in Figure 3. The Memory Module stores screen and action trajectories, which are then used as contextual input to the Reflection Module, enabling the system to analyze error trajectories and adaptively revise the task plan:

$$P_n = \text{Reflection}(Ins, P_{n-1}, \text{Memory})$$

This mechanism is inspired by prior work on LLM-based reflection for decision making (Shinn et al. 2023), which demonstrates the benefits of leveraging past experiences to iteratively refine agent behavior.

Algorithm 1: Workflow of the EcoAgent

Input: User instruction Ins , Initial screen state S_0

Output: Final system state S_f

```
1: Step 1: Initial Planning.
2:  $P_0 \leftarrow DualReACT(Ins, S_0)$ 
3: Step 2: Execution and Observation.
4: for each  $(ST_t, EX_t)$  in  $P_0$  do
5:   Execute current step:
6:    $O_t \leftarrow EA(S_t, ST_t)$ 
7:    $S_{t+1} \leftarrow O_t(S_t)$ 
8:   Verify result:
9:    $R_t \leftarrow OA(S_{t+1}, EX_t)$ 
10:  Compress screen:
11:   $T_{t+1} \leftarrow PreUnderstand(S_{t+1})$ 
12:  Update memory:
13:   $Memory \leftarrow Memory \cup \{T_{t+1}\}$ 
14:  if  $R_t == Fail$  then
15:    Step 3: Reflection and Replanning.
16:    Replan:
17:     $P_n \leftarrow Reflection(Ins, P_{n-1}, Memory)$ 
18:    Restart execution with updated plan  $P_n$ 
19:    break
20:  end if
21: end for
22: Final system state:  $S_f \leftarrow S_t$ 
23: Return  $S_f$ 
```

Device-Based Execution Agent

The Execution Agent, deployed on the device, is responsible for executing operation O_t based on the current screen state S_t and the corresponding step ST_t generated by the Planning Agent, which is denoted as:

$$O_t = EA(S_t, ST_t)$$

where EA represents the fine-tuned MSLM of the Execution Agent.

Leveraging fine-tuned MSLMs with strong grounding capabilities, the Execution Agent achieves high accuracy and low latency in input operations, which is often lacking in MLLM-based agents.

In addition to basic input execution, we design navigation operations that fully leverage the long-horizon planning and reflection capabilities of the cloud-based Planning Agent, allowing EcoAgent to promptly exit failure states and re-execute the correct actions. Below is a detailed description of the operation space:

• **Input Operations:**

- Tap(x, y): Tap the screen at coordinate (x, y) .
- Swipe($x_{start}, y_{start}, x_{end}, y_{end}$): Swipe from (x_{start}, y_{start}) to (x_{end}, y_{end}) .
- LongPress(x, y): Long press at coordinate (x, y) .
- InputText($text$): Input the string $text$ into the currently focused field.
- DeleteText(): Delete all content in the input field.

- OpenApp(app_name): Open specific application.

• **Navigation Operations:**

- PressBack(): Simulate the system back button.
- PressHome(): Simulate the system home button.

To improve the flexibility of input operations, we extend the action space with a new deletion action. We observe that existing mobile agents lack an explicit mechanism for text deletion, making it difficult to re-enter text once incorrectly input. To address this issue, we introduce the DeleteText() operation, which clears all content in the current input field located. This enhancement allows EcoAgent to effectively perform rename tasks.

Device-Based Observation Agent

Although the device-based Execution Agent can effectively recognize and locate screen elements, we find that due to the small parameter size of the fine-tuned MSLMs and the large amount of GUI data used for fine-tuning, their textual capabilities are significantly degraded, leading to challenges in assessing the success of each action. While cloud-based MLLMs could address this issue, they introduce significant latency and cost. To overcome this, we design the device-based Observation Agent, enabling accurate understanding of screen semantics and instruction-aligned expectations. Specifically, as shown in Figure 3 the Observation Agent is responsible for determining whether a step is successfully executed by comparing the post-execution screen state S_{t+1} with the expected outcome EX_t , denoted as:

$$R_t = OA(S_{t+1}, EX_t)$$

where R_t represents the result of execution verification, and OA represents the general MSLM of the Observation Agent.

Pre-Understanding Module. Screen history plays a crucial role in helping the Planning Agent identify error causes and replan accordingly. However, directly storing raw screenshots in the memory module presents two key challenges: (1) the original screen content must be transmitted to the cloud, potentially leading to user privacy leakage; and (2) image transmission introduces high device–cloud latency and significantly increases token consumption when processed by MLLMs. Moreover, effective replanning does not require access to all screen details, but only the screen transition trajectory—that is, the high-level semantic changes across screens that reflect execution progress or failure.

To address these issues, we introduce a Pre-Understanding Module within the Observation Agent. This module transforms the raw screen state S_{t+1} into a compact textual representation T_{t+1} :

$$T_{t+1} = PreUnderstand(S_{t+1})$$

Through the Pre-Understanding Module, screen images that typically consume over 1400 tokens per image are encoded into concise textual descriptions requiring only 50–150 tokens. This significantly reduces device–cloud communication overhead and MLLM token consumption. Moreover, by converting full-screen content into simplified textual descriptions rather than transmitting raw images, the risk of user privacy leakage is substantially mitigated.

Agent Architecture	Agent	Type	Foundation Model(s)	SR (%)
Device-based	ShowUI	Single	ShowUI (2B)	7.0
	InfiGUIAgent	Single	InfiGUIAgent (2B)	9.0
	OS-Atlas	Single	OS-Atlas-Pro (4B)	4.3
	V-Droid	Single	V-Droid (8B)	59.5
Cloud-based	AppAgent	Single	GPT-4o	11.2
	M3A	Multi	GPT-4o * 2	28.4
	Agent S2	Multi	GPT-4o * 4	54.3
Open-loop Device-Cloud	UGround-V1-2B	Multi	GPT-4o * 2, UGround-V1-2B	32.8
	UGround-V1-7B	Multi	GPT-4o * 2, UGround-V1-7B	44.0
Closed-loop Device-Cloud	EcoAgent (ShowUI)	Multi	GPT-4o, ShowUI (2B), Qwen2-VL-2B	25.6
	EcoAgent (OS-Atlas)	Multi	GPT-4o, OS-Atlas-Pro (4B), Qwen2-VL-2B	27.6

Table 1: Success Rate Comparison of Mobile Agents on AndroidWorld.

Architecture	Agent	MC	MT
Device-based	ShowUI	0	0
Cloud-based	AppAgent	6.46	15309
	M3A	13.39	87469
Open-loop	UGround-V1-2B	12.21	45192
Closed-loop	EcoAgent (ShowUI)	1.87	3545
	EcoAgent (OS-Atlas)	1.53	3240

Table 2: Operational Costs Comparison of Mobile Agents.

Architecture	Agent	Latency (s)
Device-based	ShowUI	1.2
	V-Droid	3.0
Cloud-based	APPAgent	7.1
	MobileAgent	15.9
	M3A	15.3
Open-loop	UGroud-V1-2B	18.2
	CogAgent	6.8
	AutoDroid	4.9
Closed-loop	EcoAgent (ShowUI)	3.9

Table 3: Latency Comparison of Mobile Agents.

Experiments

In this section, we compare EcoAgent with other mobile agents of different architectures in terms of performance, operational costs, and latency. In addition, we conduct an ablation study to validate the contribution of the proposed components within EcoAgent.

Benchmark

AndroidWorld (Rawles et al. 2024) is a dynamic benchmarking environment designed to evaluate autonomous agents on Android devices. It offers 116 programmatic tasks across 20 real-world Android applications. The benchmark runs on a live Android emulator, specifically configured with the Pixel 6 device model and Android 13 (API Level 33), ensuring a consistent and realistic testing environment. While other online benchmarks such as **MobileAgentBench** (Wang et al. 2024c) and **AndroidLab** (Xu et al. 2024) also cover a wide range of real-world tasks, they rely heavily on human evaluation to determine success or failure. This introduces challenges in reproducibility, scalability, and evaluation latency. In contrast, AndroidWorld offers fully programmatic evaluation, enabling consistent, automated, and reproducible measurements of agent performance without requiring manual intervention.

Baselines

To ensure a fair and comprehensive evaluation of EcoAgent, we compare it against several mobile agents of different architectures, including the device-based agents ShowUI (Lin et al. 2024), OS-Atlas (Wu et al. 2024c), InfiGUIAgent (Liu et al. 2025b) and V-Droid (Dai et al. 2025); the cloud-based single-agent AppAgent (Zhang et al. 2023) and MobileAgent (Wang et al. 2024b); the cloud-based multi-agent M3A (Rawles et al. 2024) and Agent S2 (Agashe et al. 2025); and the open-loop device–cloud collaborative multi-agent system UGround (Gou et al. 2025), CogAgent (Hong et al. 2024) and AutoDroid (Wen et al. 2024a).

For fairness, the cloud-based agents for all baselines are implemented using GPT-4o. For practical future deployment on mobile devices, EcoAgent adopts a lightweight design: the Observation Agent is built upon Qwen2-VL-2B, and the Execution Agent integrates the SOTA device-side models ShowUI (2B) and OS-Atlas-Pro (4B).

Metrics

We evaluate mobile agent performance along three dimensions: task success rate (SR), operational costs, and latency. The task success rate is defined as the percentage of tasks successfully completed by the agent. To assess operational

Executor	Planner	Observer	Ablation Setting			AndroidWorld		
			Executor	Planner	Observer	SR (%)	MC	MT
ShowUI	GPT-4o	Qwen2-VL-2B	✓			7.0	0	0
			✓	✓		15.5	1	2149
			✓	✓	✓	25.6	1.87	3545
OS-Atlas	GPT-4o	Qwen2-VL-2B	✓			4.3	0	0
			✓	✓		19.0	1	2181
			✓	✓	✓	27.6	1.53	3240

Table 4: Ablation results on AndroidWorld with different combinations of Execution Agent, Planning Agent, and Observation Agent.

costs, we consider two key metrics: Average MLLM Calls (MC) and Average MLLM Tokens (MT). MC denotes the average number of times the agent requests assistance from the cloud-side MLLM per task, while MT represents the average number of tokens consumed by the cloud model per task. Latency refers to the average time taken for the agent to complete each execution step.

Evaluation

Evaluation on Task Success Rate Table 1 presents the success rate (SR) comparison of various mobile agents on the AndroidWorld benchmark. Among all baselines, V-Droid achieves the highest SR at 59.5%, demonstrating its strong capability as a specialized device-based agent. Our proposed **EcoAgent** achieves SRs of 25.6% and 27.6% when using ShowUI and OS-Atlas as Execution Agents, respectively. This surpasses all single-agent baselines and even matches the performance of M3A (28.4%), a cloud-based dual-agent system using two GPT-4o models. This demonstrates that EcoAgent’s closed-loop design, which leverages lightweight on-device agents and cloud reflection via Dual-ReACT, can achieve a similar level of task success with significantly lower cloud reliance. Importantly, EcoAgent is orthogonal to V-Droid, and future work can explore integrating V-Droid as the Execution Agent, potentially combining its strong low-level execution ability with EcoAgent’s reasoning, reflection, and planning capabilities.

Evaluation on Operational Costs Table 2 presents the operational costs comparison of different mobile agents on the AndroidWorld benchmark. Compared to the cloud-based single-agent baseline AppAgent, EcoAgent (OS-Atlas) reduces MC by 76% and MT by 79%. When compared to the cloud-based dual-agent M3A, EcoAgent achieves even more significant savings: MC is reduced by 89% and MT by 96%. Although UGround-V1-2B introduces a device-based grounding agent to reduce cloud communication, it still requires $6.5\times$ more MC and $14\times$ more MT than EcoAgent (OS-Atlas). These results demonstrate that EcoAgent’s closed-loop device–cloud collaboration substantially reduces operational costs, making it a more practical and scalable solution for real-world mobile automation.

Evaluation on Latency Table 3 shows the latency comparison of mobile agents. Device-based agents, like ShowUI

and V-Droid, have the lowest latencies at 1.2 seconds and 3.0 seconds, respectively. Cloud-based agents like AppAgent, MobileAgent, and M3A exhibit much higher latencies, ranging from 7.1 seconds to 15.9 seconds, due to frequent cloud interactions. Open-loop device–cloud cooperation agents fall in between with latencies from 4.9 seconds to 18.2 seconds. EcoAgent (ShowUI) achieves a latency of 3.9 seconds, significantly lower than other cloud-based and open-loop device–cloud cooperation agents, demonstrating the efficiency of its closed-loop device–cloud collaboration.

Ablation Study

Table 4 reports the ablation results of EcoAgent on AndroidWorld under different agent combinations.

When only the Execution Agent is used, the success rate (SR) is very low (7.0% for ShowUI and 4.3% for OS-Atlas), demonstrating that relying solely on the Execution Agent leads to poor performance on complex tasks.

Introducing the Planning Agent with the Dual-ReACT framework substantially improves performance, increasing SR to 15.5% (ShowUI) and 19.0% (OS-Atlas). This validates the effectiveness of Dual-ReACT in providing structured task decomposition and reasoning.

Further adding the Observation Agent enables device-to-cloud feedback and integrates the Pre-Understanding Module. This enhances the SR to 25.6% (ShowUI) and 27.6% (OS-Atlas), demonstrating the crucial role of the Observation Agent. Notably, the increase in MC and MT is moderate. MC rises from 1 to approximately 1.5 to 1.9, while MT increases from around 2,100 to between 3,200 and 3,500. This indicates that the Pre-Understanding Module effectively reduces communication and computation costs while enabling closed-loop collaboration.

Conclusion

In this paper, we propose EcoAgent, a closed-loop device–cloud collaborative multi-agent framework for privacy-aware, efficient mobile automation. Experiments on AndroidWorld show EcoAgent reduces cloud reliance and performs robustly on complex tasks. While full on-device deployment of MSLMs remains challenging, we expect EcoAgent to become increasingly deployable as edge AI hardware and model efficiency continue to advance.

Acknowledgments

This work was supported by National Science and Technology Major Project (2022ZD0119100), National Natural Science Foundation of China (No. 62402429, U24A20326, 62441236), Key Research and Development Program of Zhejiang Province (No. 2025C01026), Ningbo Yongjiang Talent Introduction Programme (2023A-397-G) and Young Elite Scientists Sponsorship Program by CAST (2024QNR001). The author also gratefully acknowledges the support of the Zhejiang University Education Foundation Qizhen Scholar Foundation.

References

- Agashe, S.; Wong, K.; Tu, V.; Yang, J.; Li, A.; and Wang, X. E. 2025. Agent S2: A Compositional Generalist-Specialist Framework for Computer Use Agents. *arXiv:2504.00906*.
- Chen, Y.; Hu, X.; Liu, Y.; Wang, Z.; Liao, Z.; Chen, L.; Wei, F.; Qian, Y.; Zheng, B.; Yin, K.; and Zhang, S. 2025a. Graph2Eval: Automatic Multimodal Task Generation for Agents via Knowledge Graphs. *arXiv:2510.00507*.
- Chen, Y.; Hu, X.; Liu, Y.; Yin, K.; Li, J.; Zhang, Z.; and Zhang, S. 2025b. HarmonyGuard: Toward Safety and Utility in Web Agents via Adaptive Policy Enhancement and Dual-Objective Optimization. *arXiv:2508.04010*.
- Chen, Y.; Hu, X.; Yin, K.; Li, J.; and Zhang, S. 2025c. AEIA-MN: Evaluating the Robustness of Multimodal LLM-Powered Mobile Agents Against Active Environmental Injection Attacks. *arXiv:2502.13053*.
- Cheng, K.; Sun, Q.; Chu, Y.; Xu, F.; Li, Y.; Zhang, J.; and Wu, Z. 2024. SeeClick: Harnessing GUI Grounding for Advanced Visual GUI Agents. *arXiv:2401.10935*.
- Dai, G.; Jiang, S.; Cao, T.; Li, Y.; Yang, Y.; Tan, R.; Li, M.; and Qiu, L. 2025. Advancing Mobile GUI Agents: A Verifier-Driven Approach to Practical Deployment. *arXiv:2503.15937*.
- Gao, M.; Bu, W.; Miao, B.; Wu, Y.; Li, Y.; Li, J.; Tang, S.; Wu, Q.; Zhuang, Y.; and Wang, M. 2024. Generalist virtual agents: A survey on autonomous agents across digital platforms. *arXiv preprint arXiv:2411.10943*.
- Gou, B.; Wang, R.; Zheng, B.; Xie, Y.; Chang, C.; Shu, Y.; Sun, H.; and Su, Y. 2025. Navigating the Digital World as Humans Do: Universal Visual Grounding for GUI Agents. *arXiv:2410.05243*.
- Hong, W.; Wang, W.; Lv, Q.; Xu, J.; Yu, W.; Ji, J.; Wang, Y.; Wang, Z.; Zhang, Y.; Li, J.; Xu, B.; Dong, Y.; Ding, M.; and Tang, J. 2024. CogAgent: A Visual Language Model for GUI Agents. *arXiv:2312.08914*.
- Hu, X.; Xiong, T.; Yi, B.; Wei, Z.; Xiao, R.; Chen, Y.; Ye, J.; Tao, M.; Zhou, X.; Zhao, Z.; et al. 2024. Os agents: A survey on mllm-based agents for general computing devices use.
- Huang, T.; Yu, C.; Shi, W.; Peng, Z.; Yang, D.; Sun, W.; and Shi, Y. 2025. Prompt2Task: Automating UI Tasks on Smartphones from Textual Prompts. *ACM Transactions on Computer-Human Interaction*.
- Jiang, W.; Zhuang, Y.; Song, C.; Yang, X.; and Zhang, C. 2025. AppAgentX: Evolving GUI Agents as Proficient Smartphone Users. *arXiv:2503.02268*.
- Lee, S.; Choi, J.; Lee, J.; Wasi, M. H.; Choi, H.; Ko, S. Y.; Oh, S.; and Shin, I. 2024. Explore, Select, Derive, and Recall: Augmenting LLM with Human-like Memory for Mobile Task Automation. *arXiv:2312.03003*.
- Li, W.; Bishop, W.; Li, A.; Rawles, C.; Campbell-Ajala, F.; Tyamagundlu, D.; and Riva, O. 2024a. On the Effects of Data Scale on UI Control Agents. *arXiv:2406.03679*.
- Li, Y.; Zhang, C.; Yang, W.; Fu, B.; Cheng, P.; Chen, X.; Chen, L.; and Wei, Y. 2024b. AppAgent v2: Advanced Agent for Flexible Mobile Interactions. *arXiv:2408.11824*.
- Lin, K. Q.; Li, L.; Gao, D.; Yang, Z.; Wu, S.; Bai, Z.; Lei, W.; Wang, L.; and Shou, M. Z. 2024. ShowUI: One Vision-Language-Action Model for GUI Visual Agent. *arXiv:2411.17465*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; Zhu, J.; and Zhang, L. 2024a. Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection. *arXiv:2303.05499*.
- Liu, W.; Liu, L.; Guo, Y.; Xiao, H.; Lin, W.; Chai, Y.; Ren, S.; Liang, X.; Li, L.; Wang, W.; et al. 2025a. Llm-powered gui agents in phone automation: Surveying progress and prospects.
- Liu, Y.; Li, P.; Wei, Z.; Xie, C.; Hu, X.; Xu, X.; Zhang, S.; Han, X.; Yang, H.; and Wu, F. 2025b. InfiGUIAgent: A Multimodal Generalist GUI Agent with Native Reasoning and Reflection. *arXiv:2501.04575*.
- Liu, Z.; Li, C.; Chen, C.; Wang, J.; Chen, M.; Wu, B.; Wang, Y.; Hu, J.; and Wang, Q. 2024b. Seeing is Believing: Vision-driven Non-crash Functional Bug Detection for Mobile Apps. *arXiv:2407.03037*.
- Qin, Y.; Ye, Y.; Fang, J.; Wang, H.; Liang, S.; Tian, S.; Zhang, J.; Li, J.; Li, Y.; Huang, S.; et al. 2025. UI-TARS: Pioneering Automated GUI Interaction with Native Agents. *arXiv preprint arXiv:2501.12326*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rawles, C.; Clinckemaillie, S.; Chang, Y.; Waltz, J.; Lau, G.; Fair, M.; Li, A.; Bishop, W.; Li, W.; Campbell-Ajala, F.; et al. 2024. Androidworld: A dynamic benchmarking environment for autonomous agents. *arXiv preprint arXiv:2405.14573*.
- Shinn, N.; Cassano, F.; Gopinath, A.; Narasimhan, K.; and Yao, S. 2023. Reflexion: Language agents with verbal reinforcement learning. *Advances in Neural Information Processing Systems*, 36: 8634–8652.
- Song, Y.; Bian, Y.; Tang, Y.; Ma, G.; and Cai, Z. 2024. VisionTasker: Mobile Task Automation Using Vision Based UI Understanding and LLM Task Planning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, UIST '24, 1–17. ACM.

- Taeb, M.; Swearngin, A.; Schoop, E.; Cheng, R.; Jiang, Y.; and Nichols, J. 2024. Axnav: Replaying accessibility tests from natural language. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–16.
- Team, B. S. 2025. Seed1.5-VL Technical Report. *arXiv preprint arXiv:2505.07062*.
- Wang, B.; Li, G.; and Li, Y. 2023. Enabling conversational interaction with mobile ui using large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 1–17.
- Wang, J.; Xu, H.; Jia, H.; Zhang, X.; Yan, M.; Shen, W.; Zhang, J.; Huang, F.; and Sang, J. 2024a. Mobile-Agent-v2: Mobile Device Operation Assistant with Effective Navigation via Multi-Agent Collaboration. *arXiv:2406.01014*.
- Wang, J.; Xu, H.; Ye, J.; Yan, M.; Shen, W.; Zhang, J.; Huang, F.; and Sang, J. 2024b. Mobile-Agent: Autonomous Multi-Modal Mobile Device Agent with Visual Perception. *arXiv:2401.16158*.
- Wang, L.; Deng, Y.; Zha, Y.; Mao, G.; Wang, Q.; Min, T.; Chen, W.; and Chen, S. 2024c. MobileAgentBench: An Efficient and User-Friendly Benchmark for Mobile LLM Agents. *arXiv:2406.08184*.
- Wang, N.; Hu, X.; Liu, P.; Zhu, H.; Hou, Y.; Huang, H.; Zhang, S.; Yang, J.; Liu, J.; Zhang, G.; Zhang, C.; Wang, J.; Jiang, Y. E.; and Zhou, W. 2025a. Efficient Agents: Building Effective Agents While Reducing Cost. *arXiv:2508.02694*.
- Wang, Z.; Xu, H.; Wang, J.; Zhang, X.; Yan, M.; Zhang, J.; Huang, F.; and Ji, H. 2025b. Mobile-Agent-E: Self-Evolving Mobile Assistant for Complex Tasks. *arXiv:2501.11733*.
- Wang, Z.; Yang, J.; Qian, X.; Xing, S.; Jiang, X.; Lv, C.; and Zhang, S. 2024d. MNN-LLM: A Generic Inference Engine for Fast Large Language Model Deployment on Mobile Devices. In *Proceedings of the 6th ACM International Conference on Multimedia in Asia Workshops, MMAsia '24*, 1–7. ACM.
- Wen, H.; Li, Y.; Liu, G.; Zhao, S.; Yu, T.; Li, T. J.-J.; Jiang, S.; Liu, Y.; Zhang, Y.; and Liu, Y. 2024a. Autodroid: Llm-powered task automation in android. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 543–557.
- Wen, H.; Tian, S.; Pavlov, B.; Du, W.; Li, Y.; Chang, G.; Zhao, S.; Liu, J.; Liu, Y.; Zhang, Y.-Q.; and Li, Y. 2024b. AutoDroid-V2: Boosting SLM-based GUI Agents via Code Generation. *arXiv:2412.18116*.
- Wen, H.; Wang, H.; Liu, J.; and Li, Y. 2024c. DroidBot-GPT: GPT-powered UI Automation for Android. *arXiv:2304.07061*.
- Wu, B.; Li, Y.; Fang, M.; Song, Z.; Zhang, Z.; Wei, Y.; and Chen, L. 2024a. Foundations and Recent Trends in Multimodal Mobile Agents: A Survey. *arXiv preprint arXiv:2411.02006*.
- Wu, Q.; Gao, D.; Lin, K. Q.; Wu, Z.; Guo, X.; Li, P.; Zhang, W.; Wang, H.; and Shou, M. Z. 2024b. GUI Action Narrator: Where and When Did That Action Take Place? *arXiv:2406.13719*.
- Wu, Z.; Wu, Z.; Xu, F.; Wang, Y.; Sun, Q.; Jia, C.; Cheng, K.; Ding, Z.; Chen, L.; Liang, P. P.; and Qiao, Y. 2024c. OS-ATLAS: A Foundation Action Model for Generalist GUI Agents. *arXiv:2410.23218*.
- Xu, Y.; Liu, X.; Sun, X.; Cheng, S.; Yu, H.; Lai, H.; Zhang, S.; Zhang, D.; Tang, J.; and Dong, Y. 2024. Android-Lab: Training and Systematic Benchmarking of Android Autonomous Agents. *arXiv:2410.24024*.
- Yan, A.; Yang, Z.; Zhu, W.; Lin, K.; Li, L.; Wang, J.; Yang, J.; Zhong, Y.; McAuley, J.; Gao, J.; Liu, Z.; and Wang, L. 2023. GPT-4V in Wonderland: Large Multi-modal Models for Zero-Shot Smartphone GUI Navigation. *arXiv:2311.07562*.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. *arXiv:2210.03629*.
- Zhang, C.; He, S.; Qian, J.; Li, B.; Li, L.; Qin, S.; Kang, Y.; Ma, M.; Liu, G.; Lin, Q.; et al. 2024a. Large Language Model-Brained GUI Agents: A Survey. *arXiv preprint arXiv:2411.18279*.
- Zhang, C.; Yang, Z.; Liu, J.; Han, Y.; Chen, X.; Huang, Z.; Fu, B.; and Yu, G. 2023. AppAgent: Multimodal Agents as Smartphone Users. *arXiv:2312.13771*.
- Zhang, J.; Zhao, C.; Zhao, Y.; Yu, Z.; He, M.; and Fan, J. 2024b. MobileExperts: A Dynamic Tool-Enabled Agent Team in Mobile Devices. *arXiv:2407.03913*.
- Zhu, Z.; Tang, H.; Li, Y.; Liu, D.; Xu, H.; Lan, K.; Zhang, D.; Jiang, Y.; Zhou, H.; Wang, C.; Zhang, S.; Sun, L.; Wang, Y.; Sun, Y.; Chen, L.; and Yu, K. 2025. MobA: Multifaceted Memory-Enhanced Adaptive Planning for Efficient Mobile Task Automation. *arXiv:2410.13757*.