Foundations and Recent Trends in Multimodal Mobile Agents: A Survey

Anonymous ACL submission

Abstract

Mobile agents are essential for automating tasks in complex and dynamic mobile environments. As foundation models evolve, they offer increasingly powerful capabilities for understanding and generating natural language, enabling real-time adaptation and processing of multimodal data. This survey provides a comprehensive review of mobile agent technologies, with a focus on recent advancements in foundation models. Our analysis begins by exploring key representative works in mobile benchmarks and interactive environments, aiming to fully understand the research focuses and their limitations. We then introduce the core components and categorize these advancements into two main approaches: prompt-based methods, which utilize large language models (LLMs) for instruction-based task execution, and training-based methods, which fine-tune multimodal models for mobile-specific applications. By discussing key challenges and outlining future research directions, this survey offers valuable insights for advancing mobile agent technologies.

1 Introduction

Mobile agents have achieved notable success in handling complex mobile environments, enabling the automation of task execution across various applications with minimal human intervention (Zhang et al., 2023a; Li et al., 2024; Bai et al., 2024). These agents are designed to perceive, plan, and execute in dynamic environments, making them highly suitable for mobile platforms that demand real-time adaptability. Over the years, research on mobile agents has evolved from simple rule-based systems to more sophisticated models capable of handling multimodal data and complex decision-making processes (Shi et al., 2017; Rawles et al., 2023). At the same time, foundation models, such as large language models (LLMs) and multimodal models, have become pivotal in enabling mobile agents to better understand, generate, and adapt to their environments, thus expanding their capabilities in mobile systems (Ma et al., 2024; Bai et al., 2024).

The growing complexity of mobile environments and the increasing demand for automation highlight the importance of mobile agents (Deng et al., 2024a). They play a critical role in applications such as mobile interfaces, autonomous navigation, and intelligent assistance, allowing for more efficient and intelligent task execution (Yang et al., 2023). As mobile technologies advance, mobile agents are expected to operate in environments that require continuous adaptation to changing inputs, multimodal data processing, and interaction with various user interfaces (Zhang et al., 2023a). The ability to process and integrate diverse data sources in real-time makes mobile agents essential for enabling seamless user experiences and efficient operations in dynamic mobile platforms.

Despite their progress, mobile agents face several challenges. Traditional evaluation methods often fail to capture real-world mobile tasks' dynamic and interactive nature, limiting their assessment accuracy (Deng et al., 2024a). To address this, recent benchmarks such as AndroidEnv (Toyama et al., 2021) and Mobile-Env (Zhang et al., 2023a) have been developed to evaluate agents in more realistic, interactive mobile environments, focusing on adaptability and task performance. These benchmarks provide a more comprehensive assessment of mobile agents by measuring task completion and their ability to respond to changes in the environment.

Addressing complex tasks while ensuring mobile agents are multimodal, scalable, adaptable, and resource-efficient remains a significant challenge. Recent advancements in multimodal mobile agent research can be categorized into prompt-based and training-based methods. Prompt-based methods leverage large language models (LLMs), such as ChatGPT (OpenAI, 2023) and GPT-4 (OpenAI, 2023), to handle complex tasks by using instruction

prompting and chain-of-thought (CoT) reasoning (Zhang et al., 2024b). Notable works such as AppAgent (Yang et al., 2023) and AutoDroid (Wen et al., 2024) have demonstrated the potential of prompt-based systems in interactive mobile environments, although scalability and robustness remain ongoing challenges. On the other hand, training-based methods focus on fine-tuning multimodal models, such as LLaVA (Liu et al., 2023a) and Qwen-VL (Bai et al., 2023), specifically for mobile applications. These models can handle rich, multimodal data by integrating visual and textual inputs, improving their ability to perform tasks like interface navigation and task execution (Ma et al., 2024; Dorka et al., 2024).

This survey provides a comprehensive review of multimodal mobile agent technologies, focusing on recent advancements and ongoing challenges. First, we explore the benchmarks and evaluation methods used to assess mobile agent performance. Furthermore, we explore the core components of mobile agents, including perception, planning, action, and memory, which collectively enable agents to operate effectively in dynamic environments. Next, we categorize mobile agents into prompt-based and training-based approaches, discussing their strengths and limitations in improving agent adaptability, reasoning, and task execution. Finally, we discuss the future development directions of multimodal mobile agents. By providing a clear understanding of the current state of mobile agent research, this survey identifies key areas for future exploration and offers insights into the development of more adaptive, efficient, and capable mobile agents.

2 Mobile Datasets and Benchmarks

Benchmarks establish a standardized testing environment for evaluating and comparing the performance of mobile agents across both static and interactive settings, covering areas such as user interface automation, task completion, and real-world application scenarios.

Currently, as shown in Table 1, many benchmarks for GUI interaction rely on static datasets (Sun et al., 2022; Deng et al., 2024b; Niu et al., 2024; Roßner et al., 2020), which provide fixed ground-truth annotations and evaluate models by comparing their action sequences to predefined solutions. This method is problematic, as it penalizes alternative valid approaches, marking them as failures even if the task is successfully completed. Interactive benchmarks, such as AndroidArena (Xing et al., 2024), also use action sequence similarity as a primary evaluation metric, resulting in an inadequate assessment of agent performance. While recent studies on LLM-based GUI agents (Yang et al., 2023; Wang et al., 2024a; Zhang et al., 2024a) incorporate LLMs or human evaluations, these experiments are often conducted in uncontrolled open environments, leading to issues with reproducibility and comparability of results.

2.1 Static Datasets

Static datasets provide a controlled and predefined set of tasks with annotated ground-truth solutions, making them essential for evaluating the performance of mobile agents in fixed environments. These datasets are primarily used to assess task automation, where agents are required to follow predetermined actions or commands to complete specific tasks.

Early research links referring expressions to UI elements on a screen, with each instance containing a screen, a low-level command, and the corresponding UI element. For example, the RicoSCA dataset (Deka et al., 2017) uses synthetic commands, while MiniWoB++ (Liu et al., 2018) includes sequences of low-level commands for multi-step tasks.

Recent research has shifted towards task-oriented instructions, where each episode contains action-observation pairs, including screenshots and tree-structured representations like Android's View Hierarchy or the Document Object Model in web environments. For instance, the PixelHelp (Li et al., 2020a) dataset contains 187 high-level task goals with step-by-step instructions from Pixel Phone Help pages, while the UGIF (Venkatesh et al., 2022) dataset extends similar queries to multiple languages. Meanwhile, MoTIF(Burns et al., 2021), includes 4.7k task demonstrations, with an average of 6.5 steps per task and 276 unique task instructions. AITW (Rawles et al., 2024b) is much larger, featuring 715,142 episodes and 30,378 unique prompts, some inspired by other datasets.

2.2 Simulation Environments

Simulation environments provide dynamic platforms where agents engage with the environment in real time, receiving feedback and adjusting their

Dataset	Templates	Attach	Task	Reward	Platform
Static Dataset					
RICOSCA(Deka et al., 2017)	259k	-	Grounding	-	Android
ANDROIDHOWTO(Deka et al., 2017)	10k	-	Extraction	-	Android
PixelHelp(Li et al., 2020a)	187	-	Apps	-	Android
Screen2Words(Wang et al., 2021)	112k	XML	Summarization	-	Android
META-GUI(Lee et al., 2021)	1,125	-	Apps+Web	-	Android
MoTIF(Wang et al., 2022)	4,707	-	Apps	-	Android
UGIF(Venkatesh et al., 2022)	4184	XML	Grounding	-	Android
AitW(Rawles et al., 2024b)	30k	-	Apps+Web	-	Android
AitZ(Zhang et al., 2024b)	2504	-	Apps+Web	-	Android
AMEX(Chai et al., 2024)	3k	XML	Apps+Web	-	Android
Ferret-UI(You et al., 2024)	120k	-	Apps	-	IOS
GUI-World(Chen et al., 2024a)	12k	-	Apps+Web	-	Multi Platforms
Mobile3M(Chen et al., 2024a)	3M	-	Apps	-	Android
Odyssey(Lu et al., 2024)	7735	-	Apps+Web	-	Multi Platforms
Interactive Environment					
MiniWoB++(Liu et al., 2018)	114	-	Web (synthetic)	Sparse Rewards	-
AndroidEnv(Toyama et al., 2021)	100	-	Apps	Sparse Rewards	Android
AppBuddy(Shvo et al., 2021)	35	-	Apps	Sparse Rewards	Android
Mobile-Env(Zhang et al., 2023a)	224	XML	Apps+Web	Dense Rewards	Android
AndroidArena(Wang et al., 2024c)	221	XML	Apps+Web	Sparse Rewards	Android
AndroidWorld(Rawles et al., 2024a)	116	-	Apps+Web	Sparse Rewards	Android
DroidTask(Wen et al., 2024)	158	XML	Apps+Web	-	Android
B-MoCA(Lee et al., 2024)	60	XML	Apps+Web	-	Android
Mobile-Bench(Deng et al., 2024a)	832	XML	Apps+Web	-	Android

Table 1: Comparison of various platforms based on templates, attach information, tasks, rewards, and supported platforms. In particular, the reward mechanisms are categorized as Sparse Rewards and Dense Rewards. Sparse Rewards are given only when the agent reaches a specific goal or completes the task, making learning more difficult due to the lack of immediate feedback. On the other hand, Dense Rewards provides feedback after each step or action, helping the agent learn the correct strategy more quickly.

actions accordingly. Unlike static datasets, these environments allow for continuous, adaptive interactions, making them critical for evaluating agents in more complex, evolving scenarios.

Before the rise of LLM-based agents, research primarily focused on reinforcement learning (RL)-based agents. A prominent example is AndroidEnv (Toyama et al., 2021), which provided RL agents with an environment to interact with mobile applications via predefined actions and rewards. However, with advancements in LLMs, the focus has shifted towards agents that can use natural language understanding and generation to perform more flexible and adaptive tasks (Liu et al., 2024; Sun et al., 2024b,a).

Simulation Environments are a key focus in current research on LLM-based agents, particularly in their ability to explore decision paths autonomously through interactions with the environment (Liu et al., 2024; Sun et al., 2024b,a). In mobile settings, these agents are designed to handle complex, multi-step tasks and simulate human-like behaviors for app automation (Wen et al., 2023a,b; Liu et al., 2023c; Yao et al., 2022a; Shvo et al., 2021). A notable example is Mobile-Env (Zhang et al., 2023a), created to evaluate how well agents manage multi-step interactions in mobile environments. Ultimately, this research aims to improve the adaptability and flexibility of LLM-based agents, allowing them to function in dynamic, real-world environments with minimal reliance on predefined scripts or manual input.

2.3 Real-world Environments

Real-world Environments present a significant opportunity to address one of the main limitations of closed-reinforcement learning settings: their inability to fully capture the complexity and variability of real-world interactions. While controlled environments are useful for training and testing agents, they often miss the dynamic elements of real-world scenarios, where factors like changing content, unpredictable user behavior, and diverse device configurations are crucial. To overcome these challenges, researchers are increasingly exploring open, real-world environments for LLM-based GUI agents, enabling them to learn and adapt to the intricacies of live systems and evolving situations (Gao et al., 2023a; Wang et al., 2024b; Zhang et al., 2024a; Yang et al., 2023). However,

deploying agents in open-world settings introduces several risks. These include safety concerns, irreproducible results, and the potential for unfair comparisons. To mitigate these issues and ensure fair, reproducible evaluations, researchers advocate for strategies such as fixing dynamic online content and employing replay mechanisms during evaluation (Liu et al., 2018; Shi et al., 2017; Zhou et al., 2023). These methods help create a more controlled testing environment, even within the broader scope of open-world deployments.

2.4 Evaluation Methods

In evaluating agent performance, trajectory evaluation, and outcome evaluation are two main methods. Trajectory evaluation focuses on how well agent actions align with predefined paths. In contrast, outcome evaluation emphasizes whether the agent achieves its final goals, focusing on results rather than the specific process. The following sections will explore recent research advancements in these two areas, highlighting how more comprehensive evaluation strategies can enhance our understanding of agent performance in complex environments.

Trajectory Evaluation Recent improvements in GUI interaction benchmarks have focused on step-by-step assessments, comparing predicted actions to reference action trajectories to evaluate agent performance effectiveness (Rawles et al., 2024b; Zhang et al., 2021). While this approach is effective in many cases, task completion often has multiple valid solutions, and agents might explore different paths that do not necessarily follow the predefined trajectories. To improve the flexibility and robustness of these evaluations, Mobile Env introduced intermediate state evaluations, enabling reliable assessment across a wider range of tasks (Zhang et al., 2023a).

Outcome Evaluation An agent's success is determined by assessing whether it reaches the desired final state, treating task goals as subsets of hidden states, regardless of the path taken to achieve them. These final states can be identified through various system signals. Relying on a single signal type may not capture all relevant state transitions, as certain actions, such as form submissions, may only be visible in the GUI and not in system logs (Toyama et al., 2021) or databases (Rawles et al., 2024a). Shifting to outcome-based evaluation and using multiple signals can make GUI interaction benchmarks more reliable and adaptable, allowing agents to show their full abilities in various scenarios (Wang et al., 2024c; Rawles et al., 2024a).

3 The Components of Mobile Agents

As shown in Fig 1, This section outlines the four fundamental components of mobile agents: perception, planning, action, and memory. Together, these components enable agents to perceive, reason, and execute within dynamic mobile environments, adapting their behavior dynamically to improve task efficiency and robustness.

3.1 Perception

Perception is the process through which mobile agents gather and interpret multimodal information from their surroundings. In mobile agents, the perception component focuses on handling multimodal information from different environments, extracting relevant information to aid in planning and task execution.

Early research on mobile agents (Zhang et al., 2021; Sunkara et al., 2022; Song et al., 2023) primarily relied on simple models or tools to convert images or audio into text descriptions. However, these approaches often generate irrelevant and redundant information, hampering effective task planning and execution, especially in content-heavy interfaces. Additionally, the input length limitations of LLMs further amplified these challenges, making it difficult for agents to filter and prioritize information during task processing. Existing visual encoders, mostly pre-trained on general data, are not sensitive to interactive elements in mobile data. To address this, recent studies by Seeclick (Cheng et al., 2024) and CogAgent (Hong et al., 2024) have introduced mobile-specific datasets that enhance visual encoders' ability to detect and process key interactive elements, such as icons, within mobile environments.

In contexts where API calls are accessible, Mind2Web (Deng et al., 2024b) introduces a method for processing HTML-based information. This method ranks key elements of HTML data and filters crucial details to improve LLM perception of interactive components (Li et al., 2024). Meanwhile, Octopus v2 (Chen and Li, 2024) leverages specialized functional tokens to streamline function calls, significantly enhancing on-device language model efficiency and reducing computational overhead.

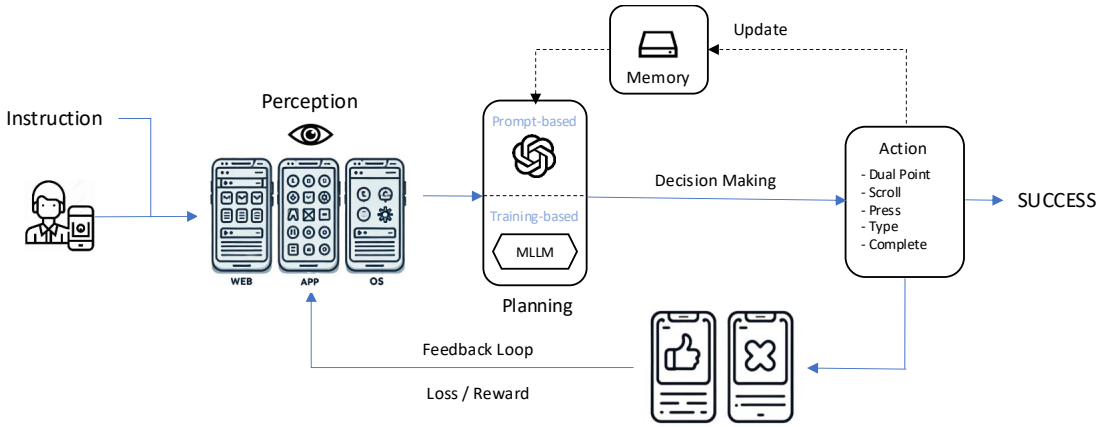

Figure 1: This pipeline shows the decision-making process of mobile agents: User instructions are processed through web, app, and OS interfaces, followed by planning with prompt-based or training-based methods. Actions are taken, and feedback is used to update memory, enabling continuous learning to achieve success.

3.2 Planning

Planning is central to mobile agents, enabling them to formulate action strategies based on task objectives and dynamic environments. Unlike agents in static settings, mobile agents must adapt to ever-changing inputs while processing multimodal information.

Planning in mobile agents can be done either programmatically or using natural language. Programmatic formats, like those in AiTW (Rawles et al., 2024b), are ideal for precise system execution. On the other hand, natural language formats, as seen in CoCo-Agent (Ma et al., 2024), bridge the gap between task instructions and the agent's existing conversational skills, making it easier for the agent to adapt and generalize to tasks in different domain.

Planning strategies can be categorized as dynamic or static. In dynamic planning, agents break down tasks into sub-goals but do not re-plan if errors occur (Zhang et al., 2024b). In contrast, static planning adjusts the plan based on real-time feedback, enabling agents to revert to earlier states and re-plan (Gao et al., 2023b; Wang et al., 2024a).

Recent advances in prompt engineering have further enhanced mobile agent planning. Omni-Act (Kapoor et al., 2024) employs prompt-based techniques to structure multimodal inputs and improve reasoning capabilities. This approach allows agents to integrate external tools and adjust output formats dynamically and efficiently.

3.3 Action

The action component demonstrates how agents execute tasks in a mobile environment by utilizing three key aspects: screen interactions, API calls, and agent interactions. Through screen interactions, agents tap, swipe, or type on GUIs, imitating human behavior to navigate apps. They also make API calls to access deeper system functions, such as issuing commands to automate tasks beyond the GUI (Chen and Li, 2024). Additionally, by collaborating with other agents, they enhance their ability to adapt to complex tasks, ensuring efficient task execution across diverse environments (Zhang et al., 2024b).

Screen Interactions In mobile environments, interactions often involve actions like tapping, swiping, or typing on virtual interfaces. Agents, such as those in AiTW, AITZ, and AMEX (Rawles et al., 2024b; Chai et al., 2024; Zhang et al., 2024b), perform GUI-based actions by mimicking human interactions, ensuring they work smoothly with native apps. These actions go beyond simple gestures, including complex multi-step processes requiring agents to dynamically adapt to changes or new inputs (Lee et al., 2021; Wang et al., 2022).

API Calls Mobile agents rely on various methods to interact with GUIs and perform tasks that require deep integration with mobile operating systems, with API calls serving as the foundation (Chen and Li, 2024; Kapoor et al., 2024). Building on API calls, mobile agents can further leverage HTML and XML data to access core functions, modify device settings, retrieve sensor data, and automate app navigation, extending their capabilities beyond GUI-based inputs (Chai et al., 2024; Chen and Li, 2024; Li et al., 2024). By integrating these approaches, agents can efficiently complete tasks

while gaining a more comprehensive understanding of their environment.

3.4 Memory

Memory mechanisms are crucial for mobile agents, allowing them to retain and use information across tasks. Current research maps in-context learning to short-term and long-term memory to external vector stores.

Short-term Memory Managing task continuity and adaptation effectively requires temporarily storing and reasoning about information, much like human working memory. Recent advancements have focused on improving the memory capabilities of mobile agents. For example, Auto-UI (Zhan and Zhang, 2023) incorporates historical text information to enhance decision-making by retaining past context, while UI-VLM (Dorka et al., 2024) uses image-based memory storage. Unlike single-modality agents, multimodal agents must handle short-term memory across various data types, including text, images, and interactions, ensuring that crucial information from different sources is retained.

Long-term Memory Handling long-term complex information requires a more efficient memory system. While external vector stores can retrieve past experiences, they function differently from human long-term memory, which is highly structured and interconnected. A combination of parametric memory and vector databases is currently used to mimic human-like long-term memory. Parametric memory stores implicit and semantic information, while vector databases hold more recent semantic and episodic memories. To make querying easier, some approaches convert multimodal inputs into a unified text format for storage, simplifying retrieval and integration during task execution (Yang et al., 2023; Wang et al., 2024b; Wen et al., 2024).

4 The Taxonomy of Mobile Agents

This section introduces a taxonomy of mobile agents, categorizing them into two primary types: prompt-based methods and training-based methods. As shown in Table 2, prompt-based agents leverage the advancements in LLMs to interpret and execute instructions through natural language processing, often focusing on tasks that require dynamic interaction with GUIs. On the other hand, training-based methods involve fine-tuning models or applying reinforcement learning to enhance agents' decision-making and adaptability over time.

4.1 Prompt-based Methods

Recent advancements in LLMs have demonstrated significant potential in developing autonomous GUI agents, particularly in tasks that require instruction following (Sanh et al., 2022; Taori et al., 2023; Chiang et al., 2023) and chain-of-thought (CoT) prompting (Nye et al., 2022; Wei et al., 2022). CoT prompting (Wei et al., 2022; Kojima et al., 2022; Zhang et al., 2023d), in particular, has proven effective in enabling LLMs to handle step-by-step processes, make decisions, and execute actions. These capabilities have shown to be highly beneficial in tasks involving GUI control (Rawles et al., 2023).

GUI Tools Enabling LLMs to interact with GUI is essential, as these models are primarily designed to process natural language rather than visual elements. GUI tools play a crucial role in bridging this gap by allowing LLMs to interpret and interact with visual elements through text-based commands, making it possible for the models to process and respond to graphical interface components. This multimodal integration significantly boosts the efficiency and flexibility of mobile agents in complex environments. Techniques like icon recognition and OCR (Zhang et al., 2021; Sunkara et al., 2022; Song et al., 2023) are used to parse GUI elements, which then converts the parsed elements into HTML layouts. However, this method relies heavily on external tools (Rawles et al., 2023; Wen et al., 2023a) and app-specific APIs (Zhou et al., 2023; Gur et al., 2023), often resulting in inefficiencies and errors during inference. Although some research has investigated multimodal architectures to process different types of inputs (Sun et al., 2022; Yan et al., 2023), these approaches still depend on detailed environment parsing for optimal performance. Given the importance of accurate GUI grounding, newer studies (Cheng et al., 2024; Hong et al., 2023) have begun exploring pre-training methods to improve agent performance in GUI tasks.

Memory Mechanism Effective task execution in prompt-based methods relies on a strong memory mechanism to retain and use relevant information. In agents like AppAgent (Yang et al., 2023), the agent employs an exploration phase for memory, allowing it to learn and adapt to new applications by

Method	Input Type	Model	Training	Memory	Multi-agents
Prompt-based Methods					
ResponsibleTA (Zhang et al., 2023c)	Image&Text	GPT-4	None	✓	✗
DroidGPT (Wen et al., 2023b)	Text	ChatGPT	None	✗	✗
AppAgent (Yang et al., 2023)	Image&Text	GPT-4	None	✓	✗
MobileAgent (Wang et al., 2024b)	Image&Text	GPT-4	None	✓	✗
MobileAgent v2 (Wang et al., 2024a)	Image&Text	GPT-4	None	✓	✓
AutoDroid (Wen et al., 2024)	Image&Text	GPT-4	None	✓	✗
AppAgent V2 (Li et al., 2024)	Image&Text	GPT-4	None	✓	✗
VLUI (Lee et al., 2024)	Image&Text	GPT4	None	✗	✗
Training-based Methods					
MiniWob (Liu et al., 2018)	Image	DOMNET	RL-based	✗	✗
MetaGUI (Sun et al., 2022)	Image&Text	VLM	Pre-trained	✗	✗
CogAgent (Hong et al., 2023)	Image&Text	CogVLM	Pre-trained	✗	✗
AutoGUI (Zhang and Zhang, 2023)	Image&Text	MMT5	Finetune	✓	✗
ResponsibleTA (Zhang et al., 2023c)	Image&Text	VLM	Finetune	✓	✗
UI-VLM (Dorka et al., 2024)	Image&Text	LLaMA	Finetune	✓	✗
Coco-Agent (Ma et al., 2024)	Image&Text	MMT5	Finetune	✓	✗
DigiRL (Bai et al., 2024)	Image&Text	MMT5	RL-based	✗	✗
SphAgent (Chai et al., 2024)	Image&Text	VLM	Finetune	✗	✗
Octopus v2 (Chen and Li, 2024)	Text	Gemma	Finetune	✗	✗
Octo-planner (Chen et al., 2024c)	Text	Gemma	Finetune	✗	✓
MobileVLM (Wu et al., 2024)	Image&Text	Qwen-VL	Finetune	✗	✗
OdysseyAgent (Lu et al., 2024)	Image&Text	Qwen-VL	Finetune	✓	✗

Table 2: Comparison of Mobile Agents: A Detailed Overview of Input Types, Models, Training Methods, Memory Capabilities, and Multi-agent Support.

storing interactions from prior explorations. This approach enables the agent to retain knowledge without needing additional training data. MobileAgent (Wang et al., 2024b,a) automates mobile app operations by analyzing screenshots with visual tools, avoiding reliance on system code.

Complex Reasoning In agent systems, complex reasoning refers to the ability of models to process, analyze, and integrate information from multiple sources to solve intricate tasks. This capability enhances decision-making, planning, and adaptability by enabling agents to draw connections between different data inputs, evaluate various outcomes, and execute informed actions in dynamic environments. CoAT (Zhang et al., 2024b) enhances GUI agent performance by integrating semantic information into action generation. It combines screen descriptions, action reasoning, next-action descriptions, and predicted outcomes to improve decision accuracy and consistency.

4.2 Training-based Methods

In contrast to prompt-based methods, training-based approaches involve explicit model optimization. These agents fine-tune large language models like LLama (Zhang et al., 2023b) or multimodal models such as LLaVA (Liu et al., 2023a) by collecting multimodal instruction-following data or accessing API to obtain instruction information. This enhancement allows these models to function as the core "brain" for reasoning and planning and to execute these plans.

Pre-trained VLMs In mobile environments, pre-trained VLMs have become powerful tools for decision-making and interaction. Models like LLaVA (Liu et al., 2023a) and Qwen-VL (Bai et al., 2023), pre-trained on large-scale general datasets, capture both visual and language information effectively. However, their applicability in mobile settings is limited by the lack of sensitivity to interactive elements specific to mobile data. To improve the responsiveness of pre-trained models to interactive elements in mobile data, CogAgent (Hong et al., 2023) collected a large-scale mobile dataset for pre-training representations. CogAgent (Hong et al., 2023) integrates visual and textual inputs for GUI agents, improving interaction with complex mobile UIs using VLMs. Spotlight (Li and Li, 2022) is a vision-language model for mobile UI tasks, relying solely on screenshots and specific regions, supporting multi-task and few-shot learning, trained on a large-scale dataset. VUT (Li et al., 2021) employs a dual-tower Transformer

for multi-task UI modeling, achieving competitive performance with fewer models and reduced computational costs.

Fine-Tuning The process of fine-tuning pre-trained VLMs with commonsense reasoning capabilities has been facilitated by large-scale mobile datasets, such as AitW (Rawles et al., 2024b), through the Visual Instruction Tuning approach. Mobile data is highly structured and information-rich, making it challenging to accurately identify the position of a specific element, especially in densely packed images. ScreenAI (Baechler et al., 2024) uses LLMs to generate synthetic data for screen annotation, identifying UI elements' types and locations to create large datasets for tasks like question answering and UI navigation. In contrast, AMEX (Chai et al., 2024) employs multi-level annotations, including GUI element grounding, functionality descriptions, and complex natural language instructions, offering more detailed training data for mobile AI agents. Both methods enhance model performance by fine-tuning with the use of constructed synthetic datasets.

Auto-GUI (Zhan and Zhang, 2023) introduces autonomous GUI control through direct interface interaction, using a chain-of-action technique for improved prediction. UI-VLM (Dorka et al., 2024) leverages multimodal data to generate image-text sequences for enhanced task performance. COCO-Agent (Ma et al., 2024) simplifies grounding tasks with modified instructions and element layouts. Octo-planner (Chen et al., 2024c) separates planning and execution, while AutoDroid (Wen et al., 2024) automates tasks by converting app exploration data into actionable knowledge, enhancing automation with fine-tuning and functionality matching.

Reinforcement Learning A dynamic approach to training mobile agents is offered by reinforcement learning, which enables them to learn from interactions with environments. This method is particularly effective in scenarios where the agent must adapt to sequential decision-making tasks or optimize its actions based on rewards. The WoB (Shi et al., 2017) platform enables reinforcement learning in real web environments by allowing agents to interact with websites using human-like actions. This work (Shi et al., 2017) converts web tasks into question-answering tasks through crowdsourcing, improving task generalization across different environments. MiniWoB++ (Liu et al., 2018) introduces workflow-guided exploration, which integrates expert workflows with task-specific actions, accelerating learning and improving task efficiency in web-based tasks. DigiRL (Bai et al., 2024) combines offline and online reinforcement learning to train device control agents. It scales online training using a VLM-based evaluator that supports real-time interaction with 64 Android emulators, enhancing the efficiency of RL-based agent training.

5 Future Work

This survey highlights key advancements in mobile agents, but important challenges remain. Future research should focus on three main areas:

Security and Privacy Strengthen security with advanced encryption and authentication methods. Develop privacy-preserving techniques to ensure data protection during agent interactions, even in open environments.

Adaptability Enhance real-time adaptability by enabling agents to adjust to changing conditions, such as resource availability and environmental factors, to maintain performance in dynamic settings.

Collaboration Improve multi-agent collaboration through better communication protocols and coordination strategies, allowing agents to work together more efficiently on complex tasks.

6 Conclusion

This survey provides a comprehensive overview of multimodal mobile agent technologies. Firstly, we reviewed advancements in mobile agents' benchmarks, which improve mobile agent assessments but still require more comprehensive methods to capture real-world dynamics. Next, we discussed the core components—perception, planning, action, and memory—that enable mobile agents to adapt to their environments, forming the foundation of their functionality. We then presented a taxonomy of mobile agents, differentiating between prompt-based and training-based methods, each with strengths and challenges in scalability and adaptability. Finally, we highlighted future research directions, focusing on security, adaptability, and multi-agent collaboration to advance mobile agent capabilities.

7 Limitations

This survey focuses on recent advancements in LLM-based mobile agents but provides limited coverage of traditional, non-LLM-based systems. The lack of discussion on older rule-based agents may limit the broader context of mobile agent technology development.

References

Gilles Baechler, Srinivas Sunkara, Maria Wang, Fedir Zubach, Hassan Mansoor, Vincent Etter, Victor Cărbune, Jason Lin, Jindong Chen, and Abhanshu Sharma. 2024. Screenai: A vision-language model for ui and infographics understanding. *arXiv preprint arXiv:2402.04615*.

Hao Bai, Yifei Zhou, Mert Cemri, Jiayi Pan, Alane Suhr, Sergey Levine, and Aviral Kumar. 2024. Digirl: Training in-the-wild device-control agents with autonomous reinforcement learning. *arXiv preprint arXiv:2406.11896*.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.

Andrea Burns, Deniz Arsan, Sanjna Agrawal, Ranjitha Kumar, Kate Saenko, and Bryan A Plummer. 2021. Mobile app tasks with iterative feedback (motif): Addressing task feasibility in interactive visual environments. *arXiv preprint arXiv:2104.08560*.

Yuxiang Chai, Siyuan Huang, Yazhe Niu, Han Xiao, Liang Liu, Dingyu Zhang, Peng Gao, Shuai Ren, and Hongsheng Li. 2024. Amex: Android multi-annotation expo dataset for mobile gui agents. *arXiv preprint arXiv:2407.17490*.

Dongping Chen, Yue Huang, Siyuan Wu, Jingyu Tang, Liuyi Chen, Yilin Bai, Zhigang He, Chenlong Wang, Huichi Zhou, Yiqiang Li, et al. 2024a. Gui-world: A dataset for gui-oriented multimodal llm-based agents. *arXiv preprint arXiv:2406.10819*.

Qi Chen, Dileepa Pitawela, Chongyang Zhao, Gengze Zhou, Hsiang-Ting Chen, and Qi Wu. 2024b. Webvln: Vision-and-language navigation on websites. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 1165–1173.

Wei Chen and Zhiyuan Li. 2024. Octopus v2: On-device language model for super agent. *arXiv preprint arXiv:2404.01744*.

Wei Chen, Zhiyuan Li, Zhen Guo, and Yikang Shen. 2024c. Octo-planner: On-device language model for planner-action agents. *arXiv preprint arXiv:2406.18082*.

Xingyu Chen, Zihan Zhao, Lu Chen, Danyang Zhang, Jiabao Ji, Ao Luo, Yuxuan Xiong, and Kai Yu. 2021. Websrc: A dataset for web-based structural reading comprehension. *arXiv preprint arXiv:2101.09465*.

Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Yantao Li, Jianbing Zhang, and Zhiyong Wu. 2024. Seeclick: Harnessing gui grounding for advanced visual gui agents. *ArXiv preprint*, abs/2401.10935.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. https://vicuna.lmsys.org.

Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibschman, Daniel Afergan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. 2017. Rico: A mobile app dataset for building data-driven design applications. In *Proceedings of the 30th annual ACM symposium on user interface software and technology*, pages 845–854.

Shihan Deng, Weikai Xu, Hongda Sun, Wei Liu, Tao Tan, Jianfeng Liu, Ang Li, Jian Luan, Bin Wang, Rui Yan, et al. 2024a. Mobile-bench: An evaluation benchmark for llm-based mobile agents. *arXiv preprint arXiv:2407.00993*.

Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Sam Stevens, Boshi Wang, Huan Sun, and Yu Su. 2024b. Mind2web: Towards a generalist agent for the web. *Advances in Neural Information Processing Systems*, 36.

Nicolai Dorka, Janusz Marecki, and Ammar Anwar. 2024. Training a vision language model as smartphone assistant. *arXiv preprint arXiv:2404.08755*.

Difei Gao, Lei Ji, Zechen Bai, Mingyu Ouyang, Peiran Li, Dongxing Mao, Qinchen Wu, Weichen Zhang, Peiyi Wang, Xiangwu Guo, et al. 2023a. Assistgui: Task-oriented desktop graphical user interface automation. *ArXiv preprint*, abs/2312.13108.

Difei Gao, Lei Ji, Luowei Zhou, Kevin Qinghong Lin, Joya Chen, Zihan Fan, and Mike Zheng Shou. 2023b. Assistgpt: A general multi-modal assistant that can plan, execute, inspect, and learn. *arXiv preprint arXiv:2306.08640*.

Yiduo Guo, Zekai Zhang, Yaobo Liang, Dongyan Zhao, and Duan Nan. 2023. Pptc benchmark: Evaluating large language models for powerpoint task completion. *arXiv preprint arXiv:2311.01767*.

Izzeddin Gur, Hiroki Furuta, Austin Huang, Mustafa Safdari, Yutaka Matsuo, Douglas Eck, and Aleksandra Faust. 2023. A real-world webagent with planning, long context understanding, and program synthesis. *ArXiv preprint*, abs/2307.12856.

Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. 2024. Webvoyager: Building an end-to-end web agent with large multimodal models. *arXiv preprint arXiv:2401.13919*.

Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. 2023. Cogagent: A visual language model for gui agents. *ArXiv preprint*, abs/2312.08914.

Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. 2024. Cogagent: A visual language model for gui agents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14281–14290.

Raghav Kapoor, Yash Parag Butala, Melisa Russak, Jing Yu Koh, Kiran Kamble, Waseem Alshikh, and Ruslan Salakhutdinov. 2024. Omniact: A dataset and benchmark for enabling multimodal generalist autonomous agents for desktop and web. *arXiv preprint arXiv:2402.17553*.

Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and Daniel Fried. 2024. Visualwebarena: Evaluating multimodal agents on realistic visual web tasks. *ArXiv preprint*, abs/2401.13649.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *ArXiv preprint*, abs/2205.11916.

Hung-Yi Lee, Mitra Mohtarami, Shang-Wen Li, Di Jin, Mandy Korpusik, Shuyan Dong, Ngoc Thang Vu, and Dilek Hakkani-Tur, editors. 2021. *Proceedings of the 1st Workshop on Meta Learning and Its Applications to Natural Language Processing*. Association for Computational Linguistics, Online.

Juyong Lee, Taywon Min, Minyong An, Changyeon Kim, and Kimin Lee. 2024. Benchmarking mobile device control agents across diverse configurations. *arXiv preprint arXiv:2404.16660*.

Gang Li and Yang Li. 2022. Spotlight: Mobile ui understanding using vision-language models with a focus. *arXiv preprint arXiv:2209.14927*.

Yanda Li, Chi Zhang, Wanqi Yang, Bin Fu, Pei Cheng, Xin Chen, Ling Chen, and Yunchao Wei. 2024. Appagent v2: Advanced agent for flexible mobile interactions. *arXiv preprint arXiv:2408.11824*.

Yang Li, Jiacong He, Xin Zhou, Yuan Zhang, and Jason Baldridge. 2020a. Mapping natural language instructions to mobile ui action sequences. *arXiv preprint arXiv:2005.03776*.

Yang Li, Gang Li, Luheng He, Jingjie Zheng, Hong Li, and Zhiwei Guan. 2020b. Widget captioning: Generating natural language description for mobile user interface elements. *arXiv preprint arXiv:2010.04295*.

Yang Li, Gang Li, Xin Zhou, Mostafa Dehghani, and Alexey Gritsenko. 2021. Vut: Versatile ui transformer for multi-modal multi-task user interface modeling. *arXiv preprint arXiv:2112.05692*.

Evan Zheran Liu, Kelvin Guu, Panupong Pasupat, Tianlin Shi, and Percy Liang. 2018. Reinforcement learning on web interfaces using workflow-guided exploration. *arXiv preprint arXiv:1802.08802*.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. *ArXiv preprint*, abs/2304.08485.

Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. 2023b. Agentbench: Evaluating llms as agents. *ArXiv preprint*, abs/2308.03688.

Yuhan Liu, Xiuying Chen, Xiaoqing Zhang, Xing Gao, Ji Zhang, and Rui Yan. 2024. From skepticism to acceptance: Simulating the attitude dynamics toward fake news. *arXiv preprint arXiv:2403.09498*.

Zhe Liu, Chunyang Chen, Junjie Wang, Mengzhuo Chen, Boyu Wu, Xing Che, Dandan Wang, and Qing Wang. 2023c. Chatting with gpt-3 for zero-shot human-like mobile automated gui testing. *arXiv preprint arXiv:2305.09434*.

Quanfeng Lu, Wenqi Shao, Zitao Liu, Fanqing Meng, Boxuan Li, Botong Chen, Siyuan Huang, Kaipeng Zhang, Yu Qiao, and Ping Luo. 2024. Gui odyssey: A comprehensive dataset for cross-app gui navigation on mobile devices. *arXiv preprint arXiv:2406.08451*.

Xinbei Ma, Zhuosheng Zhang, and Hai Zhao. 2024. Coco-agent: A comprehensive cognitive mllm agent for smartphone gui automation. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 9097–9110.

Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, et al. 2021. Webgpt: Browser-assisted question-answering with human feedback. *arXiv preprint arXiv:2112.09332*.

Runliang Niu, Jindong Li, Shiqi Wang, Yali Fu, Xiyu Hu, Xueyuan Leng, He Kong, Yi Chang, and Qi Wang. 2024. Screenagent: A vision language model-driven computer control agent. *arXiv preprint arXiv:2402.07945*.

Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. 2022. Show your work: Scratchpads for intermediate computation with language models. In *Deep Learning for Code Workshop*.

OpenAI. 2023. Chatgpt. https://openai.com/blog/chatgpt/. 1, 2.

OpenAI. 2023. Gpt-4 technical report.

Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, et al. 2023. Toolllm: Facilitating large language models to master 16000+ real-world apis. *arXiv preprint arXiv:2307.16789*.

Christopher Rawles, Sarah Clinckemaillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William Bishop, Wei Li, Folawiyo Campbell-Ajala, et al. 2024a. Androidworld: A dynamic benchmarking environment for autonomous agents. *arXiv preprint arXiv:2405.14573*.

Christopher Rawles, Alice Li, Daniel Rodriguez, Oriana Riva, and Timothy Lillicrap. 2024b. Androidinthewild: A large-scale dataset for android device control. *Advances in Neural Information Processing Systems*, 36.

Christopher Rawles, Alice Li, Daniel Rodriguez, Oriana Riva, and Timothy P Lillicrap. 2023. Androidinthewild: A large-scale dataset for android device control. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Daniel Roßner, Claus Atzenbeck, and Daniel Urban. 2020. Weblinks: Augmenting web browsers with enhanced link services. In *Proceedings of the 3rd Workshop on Human Factors in Hypertext*, pages 1–5.

Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. Multitask prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.

Tianlin Shi, Andrej Karpathy, Linxi Fan, Jonathan Hernandez, and Percy Liang. 2017. World of bits: An open-domain platform for web-based agents. In *International Conference on Machine Learning*, pages 3135–3144. PMLR.

Maayan Shvo, Zhiming Hu, Rodrigo Toro Icarte, Iqbal Mohomed, Allan D Jepson, and Sheila A McIlraith. 2021. Appbuddy: Learning to accomplish tasks in mobile apps via reinforcement learning. In *Canadian AI*.

Yunpeng Song, Yiheng Bian, Yongtao Tang, and Zhongmin Cai. 2023. Navigating interfaces with ai for enhanced user interaction. *ArXiv preprint*, abs/2312.11190.

Hongda Sun, Hongzhan Lin, Haiyu Yan, Chen Zhu, Yang Song, Xin Gao, Shuo Shang, and Rui Yan. 2024a. Facilitating multi-role and multi-behavior collaboration of large language models for online job seeking and recruiting. *arXiv preprint arXiv:2405.18113*.

Hongda Sun, Yuxuan Liu, Chengwei Wu, Haiyu Yan, Cheng Tai, Xin Gao, Shuo Shang, and Rui Yan. 2024b. Harnessing multi-role capabilities of large language models for open-domain question answering. In *Proceedings of the ACM on Web Conference 2024*, pages 4372–4382.

Liangtai Sun, Xingyu Chen, Lu Chen, Tianle Dai, Zichen Zhu, and Kai Yu. 2022. META-GUI: Towards multi-modal conversational agents on mobile GUI. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6699–6712, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Srinivas Sunkara, Maria Wang, Lijuan Liu, Gilles Baechler, Yu-Chung Hsiao, Jindong Chen, Abhanshu Sharma, and James W. W. Stout. 2022. Towards better semantic understanding of mobile interfaces. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5636–5650, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.

Daniel Toyama, Philippe Hamel, Anita Gergely, Gheorghe Comanici, Amelia Glaese, Zafarali Ahmed, Tyler Jackson, Shibl Mourad, and Doina Precup. 2021. Androidenv: A reinforcement learning platform for android. *arXiv preprint arXiv:2105.13231*.

Sagar Gubbi Venkatesh, Partha Talukdar, and Srini Narayanan. 2022. Ugif: Ui grounded instruction following. *arXiv preprint arXiv:2211.07615*.

Bryan Wang, Gang Li, Xin Zhou, Zhourong Chen, Tovi Grossman, and Yang Li. 2021. Screen2words: Automatic mobile ui summarization with multimodal learning. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pages 498–510.

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. *arXiv preprint arXiv:2406.01014*.

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*.

Luyuan Wang, Yongyu Deng, Yiwei Zha, Guodong Mao, Qinmin Wang, Tianchen Min, Wei Chen, and Shoufa Chen. 2024c. Mobileagentbench: An efficient and user-friendly benchmark for mobile llm agents. *arXiv preprint arXiv:2406.08184*.

Xintong Wang, Florian Schneider, Özge Alacam, Prateek Chaudhury, and Chris Biemann. 2022. MOTIF: Contextualized images for complex words to improve human reading. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2468–2477, Marseille, France. European Language Resources Association.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *ArXiv preprint*, abs/2201.11903.

Hao Wen, Yuanchun Li, Guohong Liu, Shanhui Zhao, Tao Yu, Toby Jia-Jun Li, Shiqi Jiang, Yunhao Liu, Yaqin Zhang, and Yunxin Liu. 2023a. Empowering llm to use smartphone for intelligent task automation. *ArXiv preprint*, abs/2308.15272.

Hao Wen, Yuanchun Li, Guohong Liu, Shanhui Zhao, Tao Yu, Toby Jia-Jun Li, Shiqi Jiang, Yunhao Liu, Yaqin Zhang, and Yunxin Liu. 2024. Autodroid: Llm-powered task automation in android. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, pages 543–557.

Hao Wen, Hongming Wang, Jiaxuan Liu, and Yuanchun Li. 2023b. Droidbot-gpt: Gpt-powered ui automation for android. *arXiv preprint arXiv:2304.07061*.

Jason Wu, Siyan Wang, Siman Shen, Yi-Hao Peng, Jeffrey Nichols, and Jeffrey P Bigham. 2023. Webui: A dataset for enhancing visual ui understanding with web semantics. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–14.

Qinzhuo Wu, Weikai Xu, Wei Liu, Tao Tan, Jianfeng Liu, Ang Li, Jian Luan, Bin Wang, and Shuo Shang. 2024. Mobilevlm: A vision-language model for better intra-and inter-ui understanding. *arXiv preprint arXiv:2409.14818*.

Mulong Xie, Sidong Feng, Zhenchang Xing, Jieshan Chen, and Chunyang Chen. 2020. Uied: a hybrid tool for gui element detection. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1655–1659.

Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, et al. 2024. Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments. *arXiv preprint arXiv:2404.07972*.

Mingzhe Xing, Rongkai Zhang, Hui Xue, Qi Chen, Fan Yang, and Zhen Xiao. 2024. Understanding the weakness of large language model agents within a complex android environment. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6061–6072.

An Yan, Zhengyuan Yang, Wanrong Zhu, Kevin Lin, Linjie Li, Jianfeng Wang, Jianwei Yang, Yiwu Zhong, Julian McAuley, Jianfeng Gao, et al. 2023. Gpt-4v in wonderland: Large multimodal models for zero-shot smartphone gui navigation. *ArXiv preprint*, abs/2311.07562.

Zhao Yang, Jiaxuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. Appagent: Multimodal agents as smartphone users. *ArXiv preprint*, abs/2312.13771.

Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022a. Webshop: Towards scalable real-world web interaction with grounded language agents. *Advances in Neural Information Processing Systems*, 35:20744–20757.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022b. React: Synergizing reasoning and acting in language models. *arXiv preprint arXiv:2210.03629*.

Keen You, Haotian Zhang, Eldon Schoop, Floris Weers, Amanda Swearngin, Jeffrey Nichols, Yinfei Yang, and Zhe Gan. 2024. Ferret-ui: Grounded mobile ui understanding with multimodal llms.

Zhuosheng Zhan and Aston Zhang. 2023. You only look at screens: Multimodal chain-of-action agents. *arXiv preprint arXiv:2309.11436*.

Chaoyun Zhang, Liqun Li, Shilin He, Xu Zhang, Bo Qiao, Si Qin, Minghua Ma, Yu Kang, Qingwei Lin, Saravan Rajmohan, et al. 2024a. Ufo: A ui-focused agent for windows os interaction. *arXiv preprint arXiv:2402.07939*.

Danyang Zhang, Lu Chen, and Kai Yu. 2023a. Mobile-env: A universal platform for training and evaluation of mobile interaction. *arXiv preprint arXiv:2305.08144*.

Jiwen Zhang, Jihao Wu, Yihua Teng, Minghui Liao, Nuo Xu, Xiao Xiao, Zhongyu Wei, and Duyu Tang. 2024b. Android in the zoo: Chain-of-action-thought for gui agents. *arXiv preprint arXiv:2403.02713*.

Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023b. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *ArXiv preprint*, abs/2303.16199.

Xiaoyi Zhang, Lilian de Greef, Amanda Swearngin, Samuel White, Kyle Murray, Lisa Yu, Qi Shan, Jeffrey Nichols, Jason Wu, Chris Fleizach, et al. 2021. Screen recognition: Creating accessibility metadata for mobile applications from pixels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15.

Zhizheng Zhang, Xiaoyi Zhang, Wenxuan Xie, and Yan Lu. 2023c. Responsible task automation: Empowering large language models as responsible task automators. *arXiv preprint arXiv:2306.01242*.

Zhuosheng Zhang and Aston Zhang. 2023. You only look at screens: Multimodal chain-of-action agents. *arXiv preprint arXiv:2309.11436*.

Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023d. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations*.

Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Yonatan Bisk, Daniel Fried, Uri Alon, et al. 2023. Webarena: A realistic web environment for building autonomous agents. *ArXiv preprint*, abs/2307.13854.

A Example Appendix

A.1 Complementary Technologies

Effective complementary technologies are vital for enhancing the performance and usability of mobile agents, in addition to key components like benchmarks, VLM models, fine-tuning methods, and advanced reasoning skills. These technologies facilitate seamless interactions with mobile environments, allowing agents to adapt, learn, and perform complex tasks efficiently.

UIED (Xie et al., 2020) detects and classifies GUI elements using computer vision and deep learning, supporting interactive editing. WebGPT (Nakano et al., 2021) fine-tunes GPT-3 for web-based question answering using imitation learning and human feedback. WebVLN (Chen et al., 2024b) trains AI agents to navigate websites with question-based instructions, incorporating HTML for deeper understanding.

A.2 Available related technologies

Additionally, OmniACT (Kapoor et al., 2024) offers a comprehensive platform for evaluating task automation across various desktop applications and natural language tasks. WebVoyager (He et al., 2024) introduces an automated evaluation protocol using GPT-4V, capturing screenshots during navigation and achieving an 85.3% agreement with human judgments. Furthermore, Widget Captioning (Li et al., 2020b) sets a benchmark for improving UI accessibility and interaction by providing 162,859 human-annotated phrases that describe UI elements from multimodal inputs, paving the way for advancements in natural language generation tasks. Above all, leveraging a diverse set of system signals provides a more comprehensive and accurate assessment of an agent's performance (Xie et al., 2024).

On desktop platforms, research has focused on evaluating how well LLM-based agents utilize APIs and software tools to complete tasks such as file management and presentations (Qin et al., 2023; Guo et al., 2023). AgentBench (Liu et al., 2023b) offers a flexible, scalable framework for evaluating agent tasks, while PPTC Benchmark (Guo et al., 2023) targets the evaluation of LLM-based agents' performance in PowerPoint-related tasks.

Dataset	Templates	Attach	Task	Reward	Platform
Static Dataset					
WebSRC(Chen et al., 2021)	400k	HTML	Web	-	Windows
WebUI(Wu et al., 2023)	400k	HTML	Web	-	Windows
Mind2Web(Deng et al., 2024b)	2,350	HTML	Web	-	Windows
OmniAct(Kapoor et al., 2024)	9802	Ocr/Seg	Web	-	Windows
WebLINX(Roßner et al., 2020)	2,337	HTML	Web	-	Windows
ScreenAgent(Niu et al., 2024)	3005	HTML	Web	-	Windows
Interactive Environment					
WebShop(Yao et al., 2022a)	12k	-	Web	Product Attrs Match	Windows
WebArena(Zhou et al., 2023)	241	HTML	Web	url/text-match	Windows
VisualWebArena(Koh et al., 2024)	314	HTML	Web	url/text/image-match	Windows
Ferret-UI(You et al., 2024)	314	HTML	Web	url/text/image-match	Windows
OSWorld(Xie et al., 2024)	369	-	Web	Device/Cloud state	Linux

Table 3: Comparison of various platforms based on parallelization, templates, tasks per template, rewards, and supported platforms.

Method	Input Type	Model	Training	Memory	Task	Multi-agents
Prompt-based Methods						
ReAct (Yao et al., 2022b)	Text	GPT-4	None	✓	Web	✗
MM-Navigator (Yan et al., 2023)	Image&Text	GPT-4	None	✗	Apps+Web	✗
MindAct (Deng et al., 2024b)	Text	GPT-4	None	✓	Apps+Web	✗
OmniAct (Kapoor et al., 2024)	Text	GPT-4	None	✗	Apps+Web	✗
Training-based Methods						
VUT (Li et al., 2021)	Image&Text	Encoder-Decoder	Pre-trained	✗	Web	✗
Spotlight (Li and Li, 2022)	Image&Text	Encoder-Decoder	Pre-trained	✗	Web	✗
ScreenAI (Baechler et al., 2024)	Image&Text	Encoder-Decoder	Pre-trained	✗	Web	✗
ScreenAgent (Niu et al., 2024)	Image&Text	CogAgent	Pre-trained	✓	Web	✗
SeeClick (Cheng et al., 2024)	Image&Text	Qwen	Pre-trained	✗	Web	✗

Table 4: Comparison of Mobile Agents: A Detailed Overview of Input Types, Models, Training Methods, Memory Capabilities, Tasks, and Multi-agent Support. Web* means synthesized web data.