

CRAB: Cross-environment Agent Benchmark for Multimodal Language Model Agents

Tianqi Xu^{1,2,3*}, Linyao Chen^{2,3,4*}, Dai-Jie Wu^{1*}, Yanjun Chen^{5*}, Zecheng Zhang, Xiang Yao^{2,3},
Zhiqiang Xie⁶, Yongchao Chen⁷, Shilong Liu⁸, Bochen Qian⁹, Anjie Yang^{2,3},
Zhaoxuan Jin^{2,3,11}, Jianbo Deng³, Philip Torr¹⁰, Bernard Ghanem^{1†}, Guohao Li^{2,10†}

¹KAUST ²Eigent.AI ³CAMEL-AI.org ⁴UTokyo ⁵CMU

⁶Stanford ⁷Harvard ⁸Tsinghua ⁹SUSTech ¹⁰Oxford ¹¹NU

*Equal contribution †Corresponding author

Correspondence: guohao.li@eigent.ai

Abstract

The development of autonomous agents increasingly relies on Multimodal Language Models (MLMs) to perform tasks described in natural language with GUI environments, such as websites, desktop computers, or mobile phones. Existing benchmarks for MLM agents in interactive environments are limited by their focus on a single environment, lack of detailed and generalized evaluation methods, and the complexities of constructing tasks and evaluators. To overcome these limitations, we introduce CRAB, the first cross-environment agent benchmark framework, incorporating a graph-based fine-grained evaluation method and an efficient task generation method. Our framework supports multiple devices and can be easily extended to any environment with a Python interface. Leveraging CRAB, we develop CRAB Benchmark-v0 comprising 120 tasks in computer desktop and mobile phone environments. We evaluated 6 advanced MLMs using different single and multi-agent system configurations on this benchmark. The experimental results demonstrate that the single agent with GPT-4o achieves the best completion ratio of 38.01%.

1 Introduction

In recent years, Multimodal Language Models (MLMs) based GUI agents have emerged as a significant area of research, aiming to develop intelligent systems capable of interacting with digital environments autonomously, including desktop OS (Zhang et al.), websites (Zhou et al.; Koh et al.), smartphones (Zhang et al., 2023; Xing et al., 2024), and games (Vinyals et al.; Wang et al.). A crucial aspect of advancing these agents is the evaluation process, which directly impacts their effectiveness and evolution. There are two distinct types of GUI agent benchmarks: one-step benchmarks (Sun et al.; Chen et al., 2024; Lu et al., 2024) and task-completion benchmarks (Xie et al.; Koh et al.). One-step benchmarks typically provide a screen-

shot, a current instruction, and an action space, along with a correct action as the label. In this setting, the agent is required to select a single correct action without access to prior actions or full environment information. These benchmarks are primarily used to evaluate UI grounding capabilities. In contrast, task-completion benchmarks involve one or more interactive environments and require the agent to complete a task based on a given instruction. Existing task completion benchmarks are typically evaluated on a single platform, either Web, Android, or Desktop OS (Shi et al., 2017; Xing et al., 2024; Xie et al.). However, real-world applications often involve tasks that span multiple platforms, such as smartphone-PC collaboration or distributed server interactions. Traditional benchmarks do not support such cross-platform tasks, limiting their applicability to practical scenarios.

To address these limitations, we introduce a novel framework CRAB, **C**Ross-environment **A**gent **B**enchmark framework. It leverages multiple environments, enabling agents to interact with real-world digital systems independently. Another major drawback of existing benchmarks lies in two key areas: task construction and evaluation methodology. Task construction often requires significant human labor to manually design and write tasks, limiting scalability. To overcome this, we propose a **graph-based task construction approach** capable of generating diverse and dynamic tasks efficiently. Additionally, traditional evaluation paradigms used in the one-step benchmarks, particularly those designed for original large language models (LLMs), which do not align with the interactive nature of agents. Many existing benchmarks either focus solely on final task completion, which overlooks intermediate progress, or enforce predefined solution paths, restricting the flexibility of different correct approaches. Our proposed evaluation framework seeks to overcome these limitations by capturing both final outcomes and intermediate milestones,

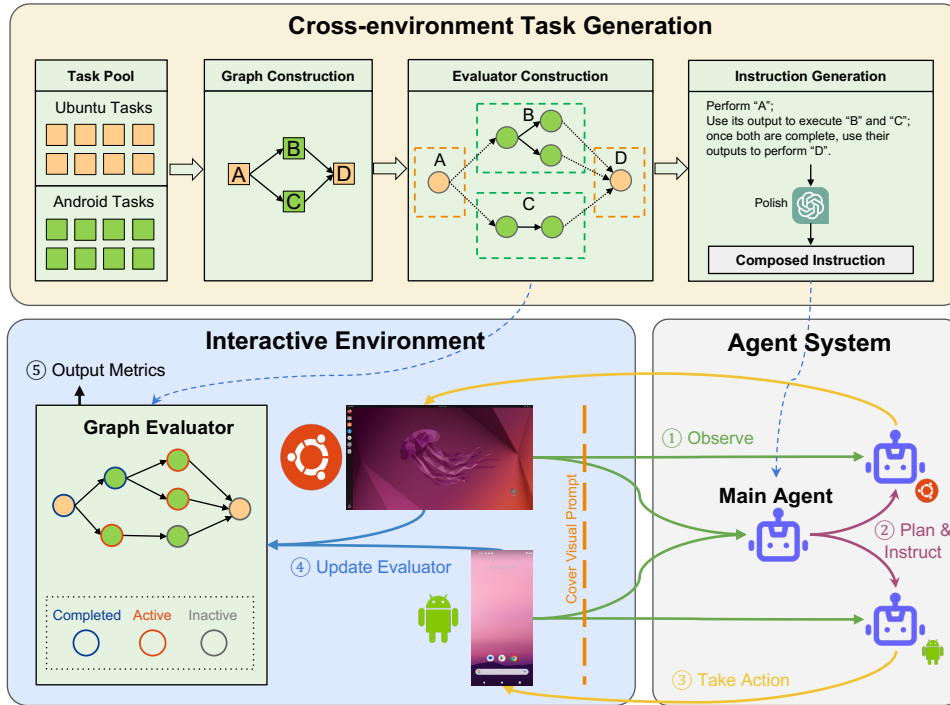


Figure 1: **Architecture of the CRAB demonstrating the task generation process and benchmarking workflow for a multi-agent system.** A task is generated by selecting subtasks from a task pool to form a graph, which is then used to construct the graph evaluator and generate instructions. The benchmarking workflow progresses through a cycle where the main agent observes, plans, and instructs the sub-agents, who then execute actions within their respective environments. The graph evaluator monitors the status of tasks within the environments, continuously updating and outputting the task completion metrics.

allowing for a more comprehensive assessment of an agent’s decision-making and adaptability.

Using the aforementioned CRAB framework, we propose a benchmark CRAB Benchmark-v0 with two cooperated environments that include an Android emulator and an Ubuntu desktop virtual machine. We have developed a total of 120 real-world tasks. These tasks address a wide array of common real-world applications and tools, including but not limited to calendars, email, maps, web browsers, and terminals, and facilitate common interactions between smartphones and desktops. Considerable time has been invested in verifying the accuracy and comprehensiveness of the instructions for sub-tasks, as well as the generalization and correctness of their evaluators. We test 6 popular MLMs, including GPT-4 Turbo, GPT-4o, Claude 3 Pro, Gemini 1.5 Pro, Pixtral-8B, and LLaVA-OneVision-72B across different structures of single-agent and multi-agent systems, totaling 12 different agent settings in our benchmarks. The experimental results show that the single agent structure with GPT-4o model achieves the best overall completion ratio of 38.01%, underscoring the necessity for ongoing development of more effective autonomous agents.

Our proposed metrics successfully distinguish between different methods better than previous metrics. We further analyze the different termination reasons that reflect the problems inherent in the communication within the multi-agent system.

2 Related Work

GUI agents With the proliferation of LLMs and MLMs, various studies (Sager et al., 2025; Zhang et al., 2025; Nguyen et al., 2024) have developed GUI agents across different platforms, including computer operating systems (Kim et al., 2023; Li et al., 2024; Zheng et al., 2024; Wu et al.; Zhang et al.; Li et al., 2024; Jia et al., 2024; Cheng et al.; Niu et al.; Baechler et al., 2024; Anthropic, 2024; Chen and Li, 2024; Chen et al., 2024; Agashe et al., 2024; Sun et al., 2024), web pages (He et al., 2024; Furuta et al., 2024; Deng et al., 2023; Deng et al.; Abuelsaad et al., 2024; Thil et al., 2024), and Android systems (Zhu et al., 2024; Wang et al., 2024; Zhang et al., 2024, 2023; Li et al., 2024; Nong et al., 2024). Although these studies have contributed to improving performance on their respective platforms through enhanced planning designs (He et al., 2024; Ma et al., 2024; Lai et al., 2024; Zhang and

Table 1: **Comparison of existing agent benchmark frameworks.** The columns details key features of each framework: *Interactive Environment* indicates the presence of either interactive environments or static datasets; *Multimodal Observation* specifies the availability of vision-based observations (e.g. screenshots); *Cross-platform* denotes support for multiple operating systems or platforms; *Evaluation* describes the evaluation metrics, categorized as *Goal-based* (checking environment state according solely on the final goal), *Trajectory-based* (comparing agent action trajectory with a gold actions sequence), *Multiple* (varied across tasks), *Intermediate-reward* (combines multiple signals with three strategies: Conjunctive Evaluation, Disjunctive Evaluation, and Order Constraint), *LLM-as-a-Judge* (Zheng et al., 2023), or *Graph-based* (a DAG with each node as an intermediate checkpoint); *Task Construction* shows the task construction method, including *Handmade* (handcrafted by human), *LLM-inspired* (using LLM to generate task drafts but still verified and annotated by human), *Template* (generated by filling in the blanks in task templates), or *Subtask Composition* (composing multiple sub’tasks to construct tasks and evaluators).

	Interactive Environment	Multimodal Observation	Cross-platform	Evaluation	Task Construction	# of apps or websites
MINIWOB++ (Shi et al., 2017)	Web	✓	✗	Goal-based	Handmade	1
WEBSHOP (Yao et al., 2022)	Web	✓	✗	Goal-based	Template	1
METAGUI (Sun et al.)	✗	✗	✗	Trajectory-based	Handmade	6
GAIA (Mialon et al.)	✗	✗	✗	Goal-based	Handmade	n/a
MIND2WEB (Deng et al., 2023)	✗	✗	✗	Goal-based	LLM-inspired	137
AGENTBENCH (Liu et al., 2024)	Multi-isolated	✗	✗	Multiple	Handmade	n/a
INTERCODE (Yang et al.)	Code	✗	✗	Goal-based	Handmade	n/a
WEBARENA (Zhou et al.)	Web	✓	✗	Goal-based	Template	6
OMNIACT (Kapoor et al.)	✗	✗	✗	Trajectory-based	Handmade	60+
VWEBARENA (Koh et al.)	Web	✓	✗	Goal-based	Template	4
ANDROIDARENA (Xing et al., 2024)	Android	✓	✗	Trajectory-based	LLM-inspired	9
OSWORLD (Xie et al.)	Linux / Windows	✓	✗	Goal-based	Template	9
MOBILE-ENV (Zhang et al., 2024)	Android	✓	✗	Intermediate-reward	Template	13
GUI-WORLD (Chen et al., 2024)	✗	✓	✗	LLM-as-a-Judge	LLM-inspired	not present
ANDROIDWORLD (Rawles et al., 2024)	Android	✓	✗	Goal-based	Template	20
WAA (Bonatti et al., 2024)	Windows	✓	✗	Goal-based	Handmade	6
GUI-ODYSSEY (Lu et al., 2024)	✗	✓	✗	Trajectory-based	LLM-inspired	201
CRAB	Linux & Android	✓	✓	Graph-based	Subtask Composition	25

Zhang, 2024; Wang et al., 2024), additional system components (You et al., 2024; Qin et al., 2025; Putta et al., 2024; Wang et al., 2024; Zhu et al., 2024; Shen et al., 2024), and novel training methods (Fereidouni et al., 2024; Meng et al., 2024; Wu et al., 2024; Lin et al., 2024; Fan et al., 2025; Liu et al., 2025; Sun et al., 2024), a robust and widely adopted benchmark for evaluation remains necessary. Recent GUI agent systems (Ge et al., 2025; Liu et al., 2025; Sun et al., 2024) have also begun to model GUI tasks across different platforms from a unified perspective, highlighting the need for benchmarks that focus on cross-platform capabilities.

Benchmarks for GUI agents Recently, various benchmarks have been applied to evaluate the capabilities of GUI agents across different platforms. Early works (Mialon et al.; Shi et al., 2017; Chen et al., 2021; Yao et al., 2022; Deng et al., 2023) focus on evaluating the automation capabilities of web-based agents, which were designed using the XML architecture of webpages. With the rise of multi-modal web agents, subsequent works (Koh et al.; Zhou et al.; Lù et al.; Zhang et al., 2024; Chen et al., 2023, 2024) have been built upon multi-modal inputs for web tasks. Researchers have also

expanded the focus from web-based platforms to mobile systems. Some studies (Deka et al., 2017; Sun et al.; Rawles et al., 2024; Xing et al., 2024; Zhang et al., 2024; Lee et al., 2024; Wang et al., 2024) rely solely on visual information, while others (Chai et al., 2024; Xu et al., 2024; Zhang et al., 2024; Lu et al., 2024; Fan et al., 2024; Pan et al., 2024) introduce additional inputs, such as XML data or control codes. In addition, efforts (Kim et al., 2023; Kapoor et al.; Xie et al.; Gao et al., 2024; Zheng et al., 2024) have been made to develop benchmarks for GUI control tasks in computer operating systems, covering various environments such as Ubuntu (Xie et al.), Windows (Bonatti et al., 2024), and other systems. These contributions have significantly advanced GUI automation. These benchmarks can be broadly categorized into two types based on evaluation methods: goal-based benchmarks (Mialon et al.; Xie et al.; Rawles et al., 2024) and trajectory-based benchmarks (Sun et al.; Xing et al., 2024). Mobile-env (Zhang et al., 2024) introduces a new perspective by incorporating immediate rewards into the evaluation process.

Existing benchmarks for GUI agents generally ignore the potential graph-like relationships between sequential actions, where some steps can

be reordered without affecting task success. This makes evaluations less realistic, as tasks in the real world don't always have a fixed order of steps. Although our benchmark does not include the largest number of apps, it uniquely emphasizes cross-environment evaluation, a significantly more complex and underexplored benchmarking scenario. In contrast to one-step benchmarks, which cannot accurately evaluate the agent performance in real world scenario, our approach involves fully interactive environments and introduces a graph-based evaluator. This requires greater implementation effort but provides increased flexibility and extensibility for future research.

3 The Crab Framework

3.1 Cross-environment Task

Compared to single-environment tasks, cross-environment tasks offer two main advantages for benchmarking agents. First, cross-environment tasks reflect real-world scenarios where humans use multiple devices simultaneously to accomplish tasks. Second, these tasks require sophisticated message processing and information transfer between environments. Such tasks demand that the agent plan actions, construct outputs for each environment, and remember what needs to be transferred, showcasing a high-level understanding of real-world and ability to solving complex tasks.

The following paragraph provides a formal definition of a cross-environment task and agent system. Consider an autonomous agent performing a task in a GUI environment. Such an environment typically has input devices (i.e., a mouse and keyboard) for human interaction and output devices (i.e., a screen) to allow human observation of its state. Formally, this environment is defined as a reward-free Partially Observable Markov Decision Process (POMDP), denoted by the tuple $M := (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{O})$, where \mathcal{S} represents the state space, \mathcal{A} the action space, $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ the transition function, and \mathcal{O} the observation space. Considering the collaborative nature of multiple devices in real-world scenarios, we can combine multiple environments into a set $\mathbf{M} = M_1, M_2, \dots, M_n$, where n is the number of environments and each environment $M_j = (\mathcal{S}_j, \mathcal{A}_j, \mathcal{T}_j, \mathcal{O}_j)$. We define a task that requires operations across multiple environments as a **cross-environment task**. This task is formalized as a tuple (\mathbf{M}, I, R) , in which \mathbf{M} is the environment set, I is the task objective

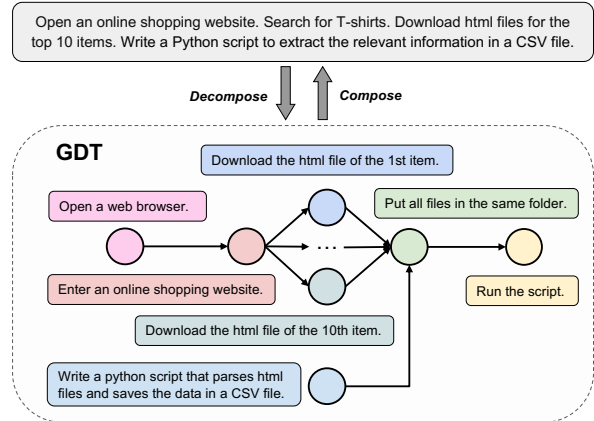


Figure 2: An example of Graph of Decomposed Tasks (GDT).

in the form of natural language instructions, and R is the reward function of the task. An **agent system**, designed to complete a task represented by an instruction I , can be modeled as a policy $\pi((m, a) \mid (I, H, o_1, \dots, o_n))$, which defines the probability of taking action a in environment m when receiving observation (o_1, \dots, o_n) from environment (M_1, \dots, M_n) with a history of actions H . An **agent** within the agent system operates with a fixed back-end MLM, a predefined system prompt, and retains its chat history. An agent system is composed of either a single agent responsible for all planning, reasoning, and action-taking or multiple agents connected through a communication workflow to collaborate.

CRAB uses a unified interface for agents to operate in all environments. We define an action by its name, the environment it belongs to, a concrete description of its functionality, and the parameters with descriptions. Through this approach, CRAB can adapt to any platform or modality, from devices to applications like browsers, by defining a few interactive functions. Implementation details are in the Appendix A.3.

3.2 Task Generation

Decomposing a complex task into several simpler sub-tasks has been proven to be an effective prompting method for LLMs (Khot et al., 2023). Some studies model sub-tasks using a graph structure. For instance, PLaG (Lin et al.) uses a graph-based structure to enhance plan reasoning within LLMs, while DyVal (Zhu et al., 2024) employs directed acyclic graphs (DAGs) to facilitate dynamic evaluation of LLMs, which decomposes a complex task into subtasks with both sequential and parallel dependencies. Based on this idea, we introduce the

Graph of Decomposed Tasks (GDT), which introduces a novel task decomposition method that represents subtasks within a DAG structure with clear input and output definition. In GDT, each node is a subtask, formalized as a tuple (m, i, r) , where m specifies the environment in which the sub-task is performed, i provides the natural language instruction, and r represents the reward function. r evaluates the state of m and outputs a boolean value to determine if the subtask is completed. During the decomposition process, we follow the principle that each subtask should **perform a single function within a distinct environment**, with clearly defined inputs and outputs that enable seamless integration with other tasks. For example, downloading a file from a URL to a specified file path constitutes a well-defined subtask: it accepts a URL as input and outputs the file’s contents. Edges within the GDT define sequential dependencies between subtasks. An example GDT is shown in Fig. 2.

The task generation process employs the reverse of decomposition process, which is subtask composition, to the realm of agent benchmark we ease task and evaluator creation. This is building GDTs by sub-tasks. We still need to address two main challenges: **(1) the need for manual creation of subtasks and (2) the complexity of modeling sequential and parallel relationships between them**. A template-based approach is commonly used to address the first issue by generating a large number of tasks efficiently. To tackle the second challenge, we use the well-defined input and output of subtasks. Specifically, if a subtask α produces an output that serves as an input for another subtask β , then α can be considered a legitimate prerequisite of β , allowing us to connect α and β with an directed edge in the GDT. To further refine our approach, we introduce a *sub-task template* structure. Each subtask is described using a natural language instruction template that includes several replaceable input attributes. The types of each input attribute and the task output should be defined carefully. To generate a GDT, input attributes can be filled with either a hand-crafted value corresponding to their type or linked to a task with the same output type as the input type.

Task descriptions are initially generated by GPT-4 from subtask prompts and refined by human reviewers. This approach, unlike naive templates, allows for a more detailed and scalable task composition.

3.3 Graph Evaluator

To assess the capabilities of MLM agents, most benchmarks (Shi et al., 2017; Deng et al., 2023; Koh et al.; Zhou et al.) evaluate performance solely based on the final state of the environment after the agent’s actions. They typically assess only whether the task was ultimately successful or not. However, this approach overlooks the agent’s incremental progress, which can be crucial for analyzing system shortcomings and leads to an incomplete evaluation of agent performance. For instance, consider two agents tasked with installing a new application on a computer: agent a successfully downloads the installer but fails during the installation process, whereas agent b does not even try to find the installer. Despite Agent a making more progress, both are deemed failures under the goal-based evaluation system, resulting in an incomplete assessment of their performance. An alternative method, *Trajectory-based Matching* (Xing et al., 2024; Kapoor et al.), abandons state-based evaluation and instead compares the agent’s actions against a predefined gold action sequence for each task, giving nuanced metrics. Nevertheless, this method faces challenges in real-world systems where tasks may have multiple valid execution paths. We propose a novel integrated approach, the *Graph Evaluator*, which provides fine-grained metrics and supports multiple valid paths.

To build a graph evaluator for a given subtask, we decompose it into a sequence of checkpoints, each representing a critical intermediate state required for subtask completion. We denote each intermediate state as s , representing a snapshot of the environment configuration at a specific step. Each checkpoint is associated with a **binary verification function** $b : \mathcal{S} \rightarrow \{0, 1\}$, which evaluates whether the current environment state s satisfies the desired condition. That is, $b(s) = 1$ if the state s matches the target specification of the checkpoint, and 0 otherwise. Formally, the graph evaluator is defined as a directed acyclic graph $\mathcal{G}_e = (\mathcal{V}, \mathcal{E})$, where each node $v \in \mathcal{V}$ corresponds to a checkpoint with an associated verification function b_v . An edge $(v_i, v_j) \in \mathcal{E}$ indicates that checkpoint v_j can only be evaluated after v_i has been successfully completed. During evaluation, a node $v \in \mathcal{V}$ becomes **active** if either it has no incoming edges: $\text{deg}^-(v) = 0$, or all its prerequisite checkpoints have been verified: $\forall (v', v) \in \mathcal{E}, b_{v'}(s) = 1$. After each agent action, the evaluator updates the

environment state s and checks all currently active nodes v using $b_v(s)$. If $b_v(s) = 1$, the node is marked completed, and its successors in the graph become active. This process repeats until no further activations occur, ensuring the evaluator’s progression is synchronized with the environment’s evolution.

Unlike trajectory-based methods, the Graph Evaluator focuses on **key states** rather than specific actions, allowing agents flexibility in execution. For instance, in a file-editing task, the evaluator checks if the file is edited and saved, regardless of whether a CLI or GUI editor is used. This ensures mandatory steps are completed while accommodating diverse execution paths.

From the task generation perspective, each sub-task template has an evaluator generator that uses the input attribute value to generate evaluator graphs. Once a GDT is constructed, the composed graph evaluator is created by interlinking evaluator graphs of subtasks in the GDT, see Fig. 1.

Given a graph evaluator synchronized with the environment state, it becomes possible to track agent progress through the current status of sub-task completions. Beyond the traditional **Success Rate (SR)**, which marks a task as *success* only when all subtasks are completed, we introduce three metrics aiming at assessing both performance and efficiency of agents, leveraging the detailed environment states provided by the graph evaluator. Specifically, the **Completion Ratio (CR)** measures the proportion of completed sub-task nodes relative to the total nodes in the graph, calculated as C / N , where C is the number of completed nodes and N is the total number of nodes. This metric offers a straightforward measure of an agent’s progress on a given task. The **Execution Efficiency (EE)**, calculated as CR / A , where A denotes the count of executed actions. It evaluates how efficiently actions are executed relative to the completion of nodes, reflecting the agent’s task execution efficiency. Lastly, the **Cost Efficiency (CE)**, calculated as CR / T , where T is the total number of model tokens used, evaluates the efficiency of resource consuming by the agent.

4 Experiments

4.1 The CRAB Benchmark

We build an agent benchmark CRAB Benchmark-v0 featuring with cross-environment, graph evaluator, and task generation through CRAB frame-

work. The environments consists of an Android smartphone emulator and a Ubuntu Linux desktop virtual machine. We establish both environments in a reproducible and standalone manner and utilize snapshots to ensure a consistent initial state for all environments. The observation space consists solely of the current system screen for both environments, captured in image format at each step of the agent’s interaction. We employ the Set-of-Marks visual prompt method (Yang et al.) to label each interactive element on the screen. Interactive elements are identified using the GroundingDINO (Liu et al.) with `icon.logo.text` prompt to locate all interactive icons. Additionally, Optical Character Recognition (OCR) is utilized through EasyOCR¹ to detect and label interactive text elements. The action spaces (Table 2) for Ubuntu and Android are distinct and designed to be close to the common interactions in the real devices. Detailed descriptions for the environment implementation are shown in Appendix A.2.

Table 2: **Action space of CRAB Benchmark-v0.** The actions at the top of the table apply to the Ubuntu environment, those in the middle to the Android environment, and those at the bottom are relevant across all environments.

Action Name (Parameters)	Description
<code>click(elem)</code>	Click on elem.
<code>right_click(elem)</code>	Right-click on elem.
<code>double_click(elem)</code>	Double-click on elem.
<code>write_text(text)</code>	Typing the specified text.
<code>press(key)</code>	Press a keyboard key.
<code>hotkey(keys)</code>	Press keyboard keys at the same time.
<code>scroll(direction)</code>	Scrolls page up or down.
<code>search_app(name)</code>	Search for application with name in the system.
<code>tap(elem)</code>	Tap on elem.
<code>long_tap(elem)</code>	Press and hold elem.
<code>swipe(elem, dire, dist)</code>	Swipe from elem in direction and distance.
<code>write_text(text)</code>	Typing the specified text.
<code>press(key)</code>	Press a key, can be <i>home</i> or <i>back</i> .
<code>show_all_drawer()</code>	Show the app drawer to list installed applications.
<code>submit(answer)</code>	Submit answer if needed.
<code>complete()</code>	State that a task is completed.
<code>wait()</code>	Wait the environment to process

We construct 17 subtask templates for Android and 19 for Ubuntu. The Ubuntu templates encompass a variety of tasks such as Command Line Interface (CLI) operations, file system management, search engine usage, desktop configurations, and map navigation. Conversely, the Android sub-task templates are primarily focused on the storage and transmission of messages via various applications. Each sub-task template is linked to a graph evaluator consisting of 1 to 4 nodes and verified by at least two related field experts. We make sure

¹<https://github.com/JaidedAI/EasyOCR>

that all tasks are reachable by human. We generate 104 tasks by sub-task composition and make 16 tasks by hand to include more complex scenarios that cannot easily be described by the sub-tasks. The dataset has 29 Android tasks, 73 Ubuntu tasks and 18 cross-platform tasks, totaling 120 tasks. Our tasks are intentionally designed to be more complex than those in other benchmarks, which naturally requires more time for design and experimentation. A single subtask in our benchmark might involve multiple operations across several applications. The format and the applications covered by the dataset are shown in Appendix A.4 and A.5, respectively. We have implemented a suite of evaluator functions. Specific techniques used in evaluators are demonstrated in Appendix A.2.

4.2 Baseline Agent System

At the core of MLM Agents are backend Multimodal Language Models that provide natural language and image understanding, basic device knowledge, task planning, and logical reasoning abilities. To run in CRAB Benchmark-v0, the backend model needs to support: (1) Accept multimodal mixed input, as the benchmark provides both screenshots and text instructions as prompts; (2) Handle multi-turn conversations; (3) Generate structured output through function calling.

We selected 4 commercial and 2 open source MLMs that meet these criteria for our experiments: GPT-4o (gpt-4o-2024-05-13) (OpenAI, 2024), GPT-4 Turbo (gpt-4-turbo-2024-04-09) (Achiam et al.), Gemini 1.5 Pro (May 2024 version) (Reid et al.), Claude 3 Opus (claude-3-opus-20240229) (Anthropic, Year), Pixtral-12B (Pixtral-12B-2409)², and LLaVA-OneVision-72B (llava-onevision-qwen-72b-ov-chat) (Li et al., 2024). These models serve as the backend models for our agents. Specifically, We use function calling feature in the four commercial models and JSON output in the two open source models that do not support function calling. Since the JSON output setting uses different prompts from the other, we employ a GPT-4o agent without function calling as the control group to the open source models.

Beyond the MLM backend, the structure of agent systems also influences overall performance. To examine how different multi-agent structures impact performance, we design three agent system structures: **single agent**, **multi-agent by function-**

ality, and **multi-agent by environment**. In the **single agent** structure, one agent manages all responsibilities, including observation analysis, planning, reasoning, and format the output action. The **multi-agent by functionality** structure splits tasks between a main agent, responsible for analysis and planning, and a tool agent that translates instructions into actions without accessing environmental observations. Meanwhile, in the **multi-agent by environment** setup, responsibilities are further distributed. A main agent processes all environmental observations for high-level planning, while each environment-specific sub-agent executes actions based on the main agent’s instructions, incorporating observations from their respective environments.

For all models, we utilized the default API parameters and retained two turns of historical messages. The max interaction turns are limited to 15. The agent can also terminate the task ahead if it thinks the task is completed. The screenshots are passed through PNG format with the highest quality that the APIs provide. Detailed agent and prompt designs are shown in Appendix B. In the experiment, we deployed four cloud machines cloned from the same disk image to ensure a consistent environment for all agents. Evaluation duration depends on the agent system, API response time, and task steps. Single-agent systems average 10 to 20 seconds per step, while multi-agent systems take 20 to 40 seconds. Running a single agent setting in the benchmark requires at least 30 hours to complete on one machine.

4.3 Result

The primary outcomes are detailed in Table 3. Aside from the *Success Rate*, *Completion Ratio*, *Execution Efficiency*, and *Cost Efficiency* mentioned above, we also present the reasons for agent termination to further investigate the factors preventing the agent system from completing the task.

Comparison of backend models. The GPT-4o and GPT-4 Turbo models, developed by OpenAI, achieved the highest average success rates and completion ratios (CR) among the tested models. Claude 3 outperforms Gemini 1.5 in terms of CR, but there remains a significant gap between the GPT-4 series and other models. Claude and Gemini have a higher Invalid Action Ratio, usually failing by clicking nonexistent elements on the screen or taking nonexistent actions. Regarding efficiency, the GPT-4 series also demonstrates strong perfor-

²<https://mistral.ai/news/pixtral-12b/>

Table 3: **Evaluation results on CRAB Benchmark-v0.** The *Model* column identifies the backend masked language models (MLMs) used. The *Structure* column describes the configuration of the agent system: *Single* means *single agent*; *By Func* is *multi-agent by functionality*; *By Env* indicates *multi-agent by environment*. We provide traditional metric of *Success Rate* (SR) alongside newly introduced metrics: *Completion Ratio* (CR), *Execution Efficiency* (EE), and *Cost Efficiency* (CE). Note that Gemini 1.5 Pro has an invalid CE because the Gemini API does not support retrieving token counts at the start time of experiments. The *Termination Reason* shows the ratio of reasons why the agent is terminated when the task is not success. *False Completion* (FC) indicates that the agent believes it has completed the task, but it actually has not; *Reach Step Limit* (RSL) means the agent has reached the step limit but has not completed the task; *Invalid Action* (IA) refers to the agent producing outputs that do not follow instructions, which may include invalid formats, nonexistent actions, or invalid action parameters.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4o	Single	14.17	38.01	4.15	5.29×10^{-4}	8.33	55.83	21.67
GPT-4o	By Func	15.00	34.00	3.93	5.31×10^{-4}	10.83	54.17	20.00
GPT-4o	By Env	14.17	33.34	3.84	2.74×10^{-4}	8.33	48.33	29.17
GPT-4 TURBO	Single	9.17	33.35	3.80	4.52×10^{-4}	8.33	65.00	17.50
GPT-4 TURBO	By Func	13.33	33.48	4.07	4.38×10^{-4}	10.83	40.00	35.83
GEMINI 1.5 PRO	Single	5.00	15.48	1.72	n/a	2.50	55.83	36.67
GEMINI 1.5 PRO	By Func	5.00	12.76	1.42	n/a	8.33	33.33	53.33
CLAUDE 3 OPUS	Single	3.33	19.60	1.95	1.85×10^{-4}	10.00	57.50	29.17
CLAUDE 3 OPUS	By Func	3.33	16.48	1.72	1.77×10^{-4}	28.33	34.17	34.17
GPT-4o w/o FC	Single	9.17	23.05	2.34	3.93×10^{-4}	5.00	42.50	43.33
PIXTRAL-12B	Single	0.83	9.50	0.75	0.87×10^{-4}	0.83	75.83	22.50
LLAVA-OV-72B	Single	0.83	6.64	0.52	1.02×10^{-4}	12.50	71.67	15.00
HUMAN	n/a	75.00	85.10	n/a	n/a	n/a	n/a	n/a

mance, with GPT-4o having a higher CE value compared to GPT-4 Turbo, highlighting its cost-effectiveness. GPT-4o’s performance drops after disabling tool calling feature, primarily due to its higher Invalid Action rate, showing the effectiveness of tool calling in generating structured output. In open source models, Pixtral-12B, with far fewer parameters, achieves a better CR compared to LLaVA-ov-72B, showcasing its efficiency. Although the open-source models generally understand screenshots and generate step-by-step plans correctly, they often fail to execute the correct actions according to the plan. Moreover, they do not effectively analyze task completion through observation. Once an incorrect action is performed, they tend to assume current step is success and proceed to the next step. Regarding platforms, we have three types of tasks: Ubuntu, Android, and cross-environment. The metrics for each type of task can reveal the model or structure preferences. We include further platform specific results in Appendix C.1.

Comparison of agent structures. The performance of multi-agent structures on all backend MLMs is slightly lower than that of single-agent structures, which is somewhat unconventional.

Based on the communication log, we find that multi-agent structures tend to experience information loss during inter-agent communication, leading to misunderstandings among downstream agents. This increases the likelihood of multi-agent structures taking invalid actions and incorrectly completing tasks. These experiments demonstrate that the design of the communication protocol and selecting the appropriate scenario are crucial for multi-agent systems. A detailed analysis is included in Appendix C.2. In terms of efficiency, multi-agent structures require more chat rounds, which can consume more tokens, resulting in a lower CE compared to single-agent settings.

Comparison of metrics. The completion ratio metric reveals a notable performance difference between models. For instance, even though GPT-4o with single agent structure and with multi-agent by environment structure have the same success rates, their completion ratios differ by up to 4.67%. This highlights the value of the completion ratio in assessing the effectiveness of different methods. For a more detailed analysis of each model and structure’s performance, we provide several case studies in the Appendix. C.3.

Key issues in solving cross-environment task.

The benchmark pipeline’s complexity makes it difficult to identify universal issues across tasks and models. However, the challenges in cross-platform tasks are similar to those in single-platform settings. Key issues include **action space discrepancies**, where diverse action spaces in cross-platform environments confuse single-agent architectures but can be mitigated by multi-agent setups tailored to each platform; **limited context length**, which prevents the ability to process entire history observations and becomes more severe for cross-platform scenarios with increasing screenshots; **coordinate grounding issues**, where advanced tools like GroundingDINO and OCR occasionally fail to detect all screen elements in too complicated GUI observation; and **icon recognition failures**, where the backend model correctly plans the next step but cannot accurately identify and interact with corresponding icons, even though the visual prompt detect them correctly.

5 Conclusion

We propose the CRAB framework, which introduces the cross-environment automatic task-performing problem, featuring advanced graph-based task generation and evaluation methods that reduce manual effort in task design and provide more dynamic and accurate agent assessments. Based on this framework, we present CRAB Benchmark-v0, a set of high-quality cross-environment tasks in smartphone and desktop environments. We tested various backend models and agent system structures on this dataset. The results reveal preferences for different agent settings.

Limitations

Despite our contribution to advancing cross-environment agent research, there are still some limitations. The sub-tasks are built upon the original apps in the Ubuntu and Android systems on Pixel devices, which limits the coverage of a wider range of applications. The current visual prompting methods do not fully recognize all interactive elements, hindering agent performance. Future work can focus on expanding the dataset and environments, testing more models, prompts, and multi-agent structures, as well as improving the use of visual prompting methods within the benchmark.

References

- Tamer Abuelsaad, Deepak Akkil, Prasenjit Dey, Ashish Jagmohan, Aditya Vempaty, and Ravi Kokku. 2024. [Agent-e: From autonomous web navigation to foundational design principles in agentic systems](#). *Preprint*, arXiv:2407.13032.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. [Gpt-4 technical report](#).
- Saaket Agashe, Jiuzhou Han, Shuyu Gan, Jiachen Yang, Ang Li, and Xin Eric Wang. 2024. [Agent s: An open agentic framework that uses computers like a human](#). *Preprint*, arXiv:2410.08164.
- Anthropic. 2024. Introducing computer use, a new claude 3.5 sonnet, and claude 3.5 haiku. <https://www.anthropic.com/news/3-5-models-and-computer-use>.
- Anthropic. Year. The claude 3 model family: Opus, sonnet, haiku. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.
- 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](#). *Preprint*, arXiv:2402.04615.
- Fabrice Bellard. 2005. Qemu, a fast and portable dynamic translator. In *USENIX annual technical conference, FREENIX Track*, volume 41, pages 10–5555. California, USA.
- Rogério Bonatti, Dan Zhao, Francesco Bonacci, Dillon Dupont, Sara Abdali, Yinheng Li, Yadong Lu, Justin Wagle, Kazuhito Koishida, Arthur Buckner, Lawrence Jang, and Zack Hui. 2024. [Windows agent arena: Evaluating multi-modal os agents at scale](#). *Preprint*, arXiv:2409.08264.
- 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](#). *Preprint*, arXiv:2407.17490.
- Dongping Chen, Yue Huang, Siyuan Wu, Jingyu Tang, Liuyi Chen, Yilin Bai, Zhigang He, Chenlong Wang, Huichi Zhou, Yiqiang Li, Tianshuo Zhou, Yue Yu, Chujie Gao, Qihui Zhang, Yi Gui, Zhen Li, Yao Wan, Pan Zhou, Jianfeng Gao, and Lichao Sun. 2024. [GUI-WORLD: A Dataset for GUI-oriented Multimodal LLM-based Agents](#).
- Qi Chen, Dileepa Pitawela, Chongyang Zhao, Gengze Zhou, Hsiang-Ting Chen, and Qi Wu. 2023. [WebVln: Vision-and-language navigation on websites](#). *Preprint*, arXiv:2312.15820.

- Wei Chen and Zhiyuan Li. 2024. [Octopus v2: On-device language model for super agent](#). *Preprint*, arXiv:2404.01744.
- Wei Chen, Zhiyuan Li, Zhen Guo, and Yikang Shen. 2024b. [Octo-planner: On-device language model for planner-action agents](#). *Preprint*, arXiv:2406.18082.
- Wentong Chen, Junbo Cui, Jinyi Hu, Yujia Qin, Junjie Fang, Yue Zhao, Chongyi Wang, Jun Liu, Guirong Chen, Yupeng Huo, Yuan Yao, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2024c. [Guicourse: From general vision language models to versatile gui agents](#). *Preprint*, arXiv:2406.11317.
- 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](#). *Preprint*, arXiv:2101.09465.
- Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Yantao Li, Jianbing Zhang, and Zhiyong Wu. [Seeclick: Harnessing gui grounding for advanced visual gui agents](#).
- Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hirschman, Daniel Afegan, 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.
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. [Mind2web: Towards a generalist agent for the web](#).
- Xiang Deng, Yu Gu, Boyuan Zheng, Shijie Chen, Samuel Stevens, Boshi Wang, Huan Sun, and Yu Su. 2023. [Mind2web: Towards a generalist agent for the web](#). *Preprint*, arXiv:2306.06070.
- Yue Fan, Lei Ding, Ching-Chen Kuo, Shan Jiang, Yang Zhao, Xinze Guan, Jie Yang, Yi Zhang, and Xin Eric Wang. 2024. [Read anywhere pointed: Layout-aware gui screen reading with tree-of-lens grounding](#). *Preprint*, arXiv:2406.19263.
- Yue Fan, Handong Zhao, Ruiyi Zhang, Yu Shen, Xin Eric Wang, and Gang Wu. 2025. [Gui-bee: Align gui action grounding to novel environments via autonomous exploration](#). *Preprint*, arXiv:2501.13896.
- Moghis Fereidouni, Adib Mosharraf, and A.b. Siddique. 2024. [Grounded language agent for product search via intelligent web interactions](#). In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, page 63–75. Association for Computational Linguistics.
- Hiroki Furuta, Kuang-Huei Lee, Ofir Nachum, Yutaka Matsuo, Aleksandra Faust, Shixiang Shane Gu, and Izzeddin Gur. 2024. [Multimodal web navigation with instruction-finetuned foundation models](#). In *The Twelfth International Conference on Learning Representations*.
- Difei Gao, Lei Ji, Zechen Bai, Mingyu Ouyang, Peiran Li, Dongxing Mao, Qinchen Wu, Weichen Zhang, Peiyi Wang, Xiangwu Guo, Hengxu Wang, Luwei Zhou, and Mike Zheng Shou. 2024. [Assistgui: Task-oriented desktop graphical user interface automation](#). *Preprint*, arXiv:2312.13108.
- Zhiqi Ge, Juncheng Li, Xinglei Pang, Minghe Gao, Kaihang Pan, Wang Lin, Hao Fei, Wenqiao Zhang, Siliang Tang, and Yueting Zhuang. 2025. [Iris: Breaking gui complexity with adaptive focus and self-refining](#). *Preprint*, arXiv:2412.10342.
- Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conference*, pages 11–15.
- 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](#). *Preprint*, arXiv:2401.13919.
- Chengyou Jia, Minnan Luo, Zhuohang Dang, Qiushi Sun, Fangzhi Xu, Junlin Hu, Tianbao Xie, and Zhiyong Wu. 2024. [Agentstore: Scalable integration of heterogeneous agents as specialized generalist computer assistant](#). *Preprint*, arXiv:2410.18603.
- Hanwen Jiang, Arjun Karpur, Bingyi Cao, Qixing Huang, and Andre Araujo. 2024. [Omniglu: Generalizable feature matching with foundation model guidance](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Raghav Kapoor, Yash Parag Butala, Melisa Russak, Jing Yu Koh, Kiran Kamble, Waseem Alshikh, and Ruslan Salakhutdinov. [Omniact: A dataset and benchmark for enabling multimodal generalist autonomous agents for desktop and web](#).
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2023. [Decomposed prompting: A modular approach for solving complex tasks](#). In *The Eleventh International Conference on Learning Representations*.
- Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2023. [Language models can solve computer tasks](#). *Preprint*, arXiv:2303.17491.
- Geunwoo Kim, Pierre Baldi, and Stephen McAleer. 2023b. [Language models can solve computer tasks](#). *Preprint*, arXiv:2303.17491.
- Avi Kivity, Yaniv Kamay, Dor Laor, Uri Lublin, and Anthony Liguori. 2007. [kvm: the linux virtual machine monitor](#). In *Proceedings of the Linux symposium*, volume 1, pages 225–230. Dttawa, Dntorio, Canada.
- Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and

- Daniel Fried. [Visualwebarena: Evaluating multi-modal agents on realistic visual web tasks.](#)
- Hanyu Lai, Xiao Liu, Iat Long Iong, Shuntian Yao, Yuxuan Chen, Pengbo Shen, Hao Yu, Hanchen Zhang, Xiaohan Zhang, Yuxiao Dong, and Jie Tang. 2024. [Autowebglm: A large language model-based web navigating agent.](#) *Preprint*, arXiv:2404.03648.
- Juyong Lee, Taywon Min, Minyong An, Dongyoon Hahm, Haeone Lee, Changyeon Kim, and Kimin Lee. 2024. [Benchmarking mobile device control agents across diverse configurations.](#) *Preprint*, arXiv:2404.16660.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. [Llava-onevision: Easy visual task transfer.](#) *arXiv preprint*.
- Hongxin Li, Jingran Su, Yuntao Chen, Qing Li, and ZHAO-XIANG ZHANG. 2024b. [Sheetcopilot: Bringing software productivity to the next level through large language models.](#) *Advances in Neural Information Processing Systems*, 36.
- Wei Li, William Bishop, Alice Li, Chris Rawles, Fola-iyoy Campbell-Ajala, Divya Tyamagundlu, and Oriana Riva. 2024c. [On the effects of data scale on ui control agents.](#) *Preprint*, arXiv:2406.03679.
- Yanda Li, Chi Zhang, Wanqi Yang, Bin Fu, Pei Cheng, Xin Chen, Ling Chen, and Yunchao Wei. 2024d. [Ap-pagent v2: Advanced agent for flexible mobile interactions.](#) *Preprint*, arXiv:2408.11824.
- Fangru Lin, Emanuele La Malfa, Valentin Hofmann, Elle Michelle Yang, Anthony Cohn, and Janet B. Pierrehumbert. [Graph-enhanced large language models in asynchronous plan reasoning.](#)
- Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Shiwei Wu, Zechen Bai, Weixian Lei, Lijuan Wang, and Mike Zheng Shou. 2024. [Showui: One vision-language-action model for gui visual agent.](#) *Preprint*, arXiv:2411.17465.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. [Grounding dino: Marrying dino with grounded pre-training for open-set object detection.](#)
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. 2024. [Agentbench: Evaluating LLMs as agents.](#) In *The Twelfth International Conference on Learning Representations*.
- Yuhang Liu, Pengxiang Li, Zishu Wei, Congkai Xie, Xueyu Hu, Xinchun Xu, Shengyu Zhang, Xiaotian Han, Hongxia Yang, and Fei Wu. 2025. [Infuguagent: A multimodal generalist gui agent with native reasoning and reflection.](#) *Preprint*, arXiv:2501.04575.
- 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.](#) *Preprint*, arXiv:2406.08451.
- Xing Han Lù, Zdeněk Kasner, and Siva Reddy. [Weblinx: Real-world website navigation with multi-turn dialogue.](#)
- Xinbei Ma, Zhuosheng Zhang, and Hai Zhao. 2024. [Coco-agent: A comprehensive cognitive mllm agent for smartphone gui automation.](#) *Preprint*, arXiv:2402.11941.
- Michael M. McKerns, Leif Strand, Tim Sullivan, Alta Fang, and Michael A. G. Aivazis. [Building a framework for predictive science.](#)
- Ziyang Meng, Yu Dai, Zezheng Gong, Shaoxiong Guo, Minglong Tang, and Tongquan Wei. 2024. [Vga: Vision gui assistant – minimizing hallucinations through image-centric fine-tuning.](#) *Preprint*, arXiv:2406.14056.
- Grégoire Mialon, Clémentine Fourrier, Craig Swift, Thomas Wolf, Yann LeCun, and Thomas Scialom. [Gaia: A benchmark for general ai assistants.](#)
- Dang Nguyen, Jian Chen, Yu Wang, Gang Wu, Namyong Park, Zhengmian Hu, Hanjia Lyu, Junda Wu, Ryan Aponte, Yu Xia, Xintong Li, Jing Shi, Hongjie Chen, Viet Dac Lai, Zhouhang Xie, Sungchul Kim, Ruiyi Zhang, Tong Yu, Mehrab Tanjim, Nesreen K. Ahmed, Puneet Mathur, Seunghyun Yoon, Lina Yao, Branislav Kveton, Thien Huu Nguyen, Trung Bui, Tianyi Zhou, Ryan A. Rossi, and Franck Dernoncourt. 2024. [Gui agents: A survey.](#) *Preprint*, arXiv:2412.13501.
- Runliang Niu, Jindong Li, Shiqi Wang, Yali Fu, Xiyu Hu, Xueyuan Leng, He Kong, Yi Chang, and Qi Wang. [Screenagent: A vision language model-driven computer control agent.](#)
- Songqin Nong, Jiali Zhu, Rui Wu, Jiongchao Jin, Shuo Shan, Xiutian Huang, and Wenhao Xu. 2024. [Mobileflow: A multimodal llm for mobile gui agent.](#) *Preprint*, arXiv:2407.04346.
- OpenAI. 2024. [Gpt-4 omni.](#) <https://openai.com/index/hello-gpt-4o/>.
- Jiayi Pan, Yichi Zhang, Nicholas Tomlin, Yifei Zhou, Sergey Levine, and Alane Suhr. 2024. [Autonomous evaluation and refinement of digital agents.](#) *Preprint*, arXiv:2404.06474.
- Guilherme Potje, Felipe Cadar, Andre Araujo, Renato Martins, and Erickson R Nascimento. 2024. [Xfeat: Accelerated features for lightweight image matching.](#) In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

- Pranav Putta, Edmund Mills, Naman Garg, Sumeet Motwani, Chelsea Finn, Divyansh Garg, and Rafael Rafailov. 2024. [Agent q: Advanced reasoning and learning for autonomous ai agents](#). *Preprint*, arXiv:2408.07199.
- Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, Wanjun Zhong, Kuanye Li, Jiale Yang, Yu Miao, Woyu Lin, Longxiang Liu, Xu Jiang, Qianli Ma, Jingyu Li, Xiaojun Xiao, Kai Cai, Chuang Li, Yaowei Zheng, Chaolin Jin, Chen Li, Xiao Zhou, Minchao Wang, Haoli Chen, Zhaojian Li, Haihua Yang, Haifeng Liu, Feng Lin, Tao Peng, Xin Liu, and Guang Shi. 2025. [Ui-tars: Pioneering automated gui interaction with native agents](#). *Preprint*, arXiv:2501.12326.
- Christopher Rawles, Sarah Clinckemahillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth Fair, Alice Li, William Bishop, Wei Li, Folawiyo Campbell-Ajala, Daniel Toyama, Robert Berry, Divya Tyamagundlu, Timothy Lillicrap, and Oriana Riva. 2024. [Androidworld: A dynamic benchmarking environment for autonomous agents](#). *Preprint*, arXiv:2405.14573.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#).
- Pascal J. Sager, Benjamin Meyer, Peng Yan, Rebekka von Wartburg-Kottler, Layan Etaiwi, Aref Enayati, Gabriel Nobel, Ahmed Abdulkadir, Benjamin F. Grewe, and Thilo Stadelmann. 2025. [Ai agents for computer use: A review of instruction-based computer control, gui automation, and operator assistants](#). *Preprint*, arXiv:2501.16150.
- Huawen Shen, Chang Liu, Gengluo Li, Xinlong Wang, Yu Zhou, Can Ma, and Xiangyang Ji. 2024. [Falconui: Understanding gui before following user instructions](#). *Preprint*, arXiv:2412.09362.
- Tianlin Shi, Andrej Karpathy, Linxi Fan, Jonathan Hernandez, and Percy Liang. 2017. [World of bits: An open-domain platform for web-based agents](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3135–3144. PMLR.
- Liangtai Sun, Xingyu Chen, Lu Chen, Tianle Dai, Zichen Zhu, and Kai Yu. [Meta-gui: Towards multimodal conversational agents on mobile gui](#).
- Qiushi Sun, Kanzhi Cheng, Zichen Ding, Chuanyang Jin, Yian Wang, Fangzhi Xu, Zhenyu Wu, Chengyou Jia, Liheng Chen, Zhoumianze Liu, Ben Kao, Guohao Li, Junxian He, Yu Qiao, and Zhiyong Wu. 2024. [Os-genesis: Automating gui agent trajectory construction via reverse task synthesis](#). *Preprint*, arXiv:2412.19723.
- Lucas-Andrei Thil, Mirela Popa, and Gerasimos Spanakis. 2024. [Navigating webai: Training agents to complete web tasks with large language models and reinforcement learning](#). In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, SAC '24, page 866–874. ACM.
- Oriol Vinyals, Igor Babuschkin, Wojciech M. Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H. Choi, Richard Powell, Timo Ewalds, Petko Georgiev, Junhyuk Oh, Dan Horgan, Manuel Kroiss, Ivo Danihelka, Aja Huang, Laurent Sifre, Trevor Cai, John P. Agapiou, Max Jaderberg, Alexander S. Vezhnevets, Rémi Leblond, Tobias Pohlen, Valentin Dalibard, David Budden, Yury Sulsky, James Molloy, Tom L. Paine, Caglar Gulcehre, Ziyu Wang, Tobias Pfaff, Yuhuai Wu, Roman Ring, Dani Yogatama, Dario Wünsch, Katrina McKinney, Oliver Smith, Tom Schaul, Timothy Lillicrap, Koray Kavukcuoglu, Demis Hassabis, Chris Apps, and David Silver. [Grandmaster level in starcraft ii using multi-agent reinforcement learning](#). 575(7782):350–354.
- Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. [Voyager: An open-ended embodied agent with large language models](#).
- Junyang Wang, Haiyang Xu, Haitao Jia, Xi Zhang, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. 2024. [Mobile-agent-v2: Mobile device operation assistant with effective navigation via multi-agent collaboration](#). *Preprint*, arXiv:2406.01014.
- Junyang Wang, Haiyang Xu, Jiabo Ye, Ming Yan, Weizhou Shen, Ji Zhang, Fei Huang, and Jitao Sang. [b. Mobile-agent: Autonomous multi-modal mobile device agent with visual perception](#).
- Luyuan Wang, Yongyu Deng, Yiwei Zha, Guodong Mao, Qinmin Wang, Tianchen Min, Wei Chen, and Shoufa Chen. 2024b. [Mobileagentbench: An efficient and user-friendly benchmark for mobile llm agents](#). *Preprint*, arXiv:2406.08184.
- Zora Zhiruo Wang, Jiayuan Mao, Daniel Fried, and Graham Neubig. 2024c. [Agent workflow memory](#). *Preprint*, arXiv:2409.07429.
- Zhiyong Wu, Chengcheng Han, Zichen Ding, Zhenmin Weng, Zhoumianze Liu, Shunyu Yao, Tao Yu, and Lingpeng Kong. [Os-copilot: Towards generalist computer agents with self-improvement](#).
- Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, and Yu Qiao. 2024. [Os-atlas: A foundation action model for generalist gui agents](#). *Preprint*, arXiv:2410.23218.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. [Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments](#).

- 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](#). ArXiv preprint arXiv:2402.06596.
- Yifan Xu, Xiao Liu, Xueqiao Sun, Siyi Cheng, Hao Yu, Hanyu Lai, Shudan Zhang, Dan Zhang, Jie Tang, and Yuxiao Dong. 2024. [Androidlab: Training and systematic benchmarking of android autonomous agents](#). *Preprint*, arXiv:2410.24024.
- Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. [Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v](#).
- John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. b. [Intercode: Standardizing and benchmarking interactive coding with execution feedback](#).
- Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. 2022. [Webshop: Towards scalable real-world web interaction with grounded language agents](#). *Advances in Neural Information Processing Systems*, 35:20744–20757.
- 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](#). *Preprint*, arXiv:2404.05719.
- Chaoyun Zhang, Shilin He, Jiaxu Qian, Bowen Li, Liqun Li, Si Qin, Yu Kang, Minghua Ma, Guyue Liu, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2025. [Large language model-brained gui agents: A survey](#). *Preprint*, arXiv:2411.18279.
- Chaoyun Zhang, Liqun Li, Shilin He, Xu Zhang, Bo Qiao, Si Qin, Minghua Ma, Yu Kang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. [Ufo: A ui-focused agent for windows os interaction](#).
- Chi Zhang, Zhao Yang, Jiakuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu. 2023. [Appagent: Multimodal agents as smartphone users](#).
- Danyang Zhang, Zhennan Shen, Rui Xie, Situo Zhang, Tianbao Xie, Zihan Zhao, Siyuan Chen, Lu Chen, Hongshen Xu, Ruisheng Cao, and Kai Yu. 2024. [Mobile-Env: Building Qualified Evaluation Benchmarks for LLM-GUI Interaction](#).
- Jiayi Zhang, Chuang Zhao, Yihan Zhao, Zhaoyang Yu, Ming He, and Jianping Fan. 2024b. [Mobileexperts: A dynamic tool-enabled agent team in mobile devices](#). *Preprint*, arXiv:2407.03913.
- Li Zhang, Shihe Wang, Xianqing Jia, Zhihan Zheng, Yunhe Yan, Longxi Gao, Yuanchun Li, and Mengwei Xu. 2024c. [Llamatouch: A faithful and scalable testbed for mobile ui task automation](#). *Preprint*, arXiv:2404.16054.
- Zhuosheng Zhang and Aston Zhang. 2024. [You only look at screens: Multimodal chain-of-action agents](#). *Preprint*, arXiv:2309.11436.
- Ziniu Zhang, Shulin Tian, Liangyu Chen, and Ziwei Liu. 2024d. [Mmina: Benchmarking multihop multimodal internet agents](#). *Preprint*, arXiv:2404.09992.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena](#).
- Longtao Zheng, Zhiyuan Huang, Zhenghai Xue, Xinrun Wang, Bo An, and Shuicheng Yan. 2024. [Agentstudio: A toolkit for building general virtual agents](#). *Preprint*, arXiv:2403.17918.
- Longtao Zheng, Rundong Wang, Xinrun Wang, and Bo An. 2024b. [Synapse: Trajectory-as-exemplar prompting with memory for computer control](#). In *The Twelfth International Conference on Learning Representations*.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. [Webarena: A realistic web environment for building autonomous agents](#).
- Kaijie Zhu, Jiaao Chen, Jindong Wang, Neil Zhenqiang Gong, Diyi Yang, and Xing Xie. 2024. [Dyval: Dynamic evaluation of large language models for reasoning tasks](#). In *The Twelfth International Conference on Learning Representations*.
- Zichen Zhu, Hao Tang, Yansi Li, Kunyao Lan, Yixuan Jiang, Hao Zhou, Yixiao Wang, Situo Zhang, Liangtai Sun, Lu Chen, and Kai Yu. 2024b. [Moba: A two-level agent system for efficient mobile task automation](#). *Preprint*, arXiv:2410.13757.

A Benchmark Detail

Section A.1 gives the dataset statistics of CRAB Benchmark-v0. Section A.2 introduces the implementation details and action space settings of the benchmark environments. Section A.3 describes the design logic and implementation of the CRAB framework. Section A.4 describes the our experiment settings in detail. Section A.5 describes the specific format defined in our framework that ease data extension and how to use them. We provides a detailed document to setup experiment environments and reproduce our results.³ Fig. 3 shows the structure of modules inside CRAB Benchmark-v0.

³<https://github.com/camel-ai/crab/blob/main/crab-benchmark-v0/README.md>

Crab Benchmark v0

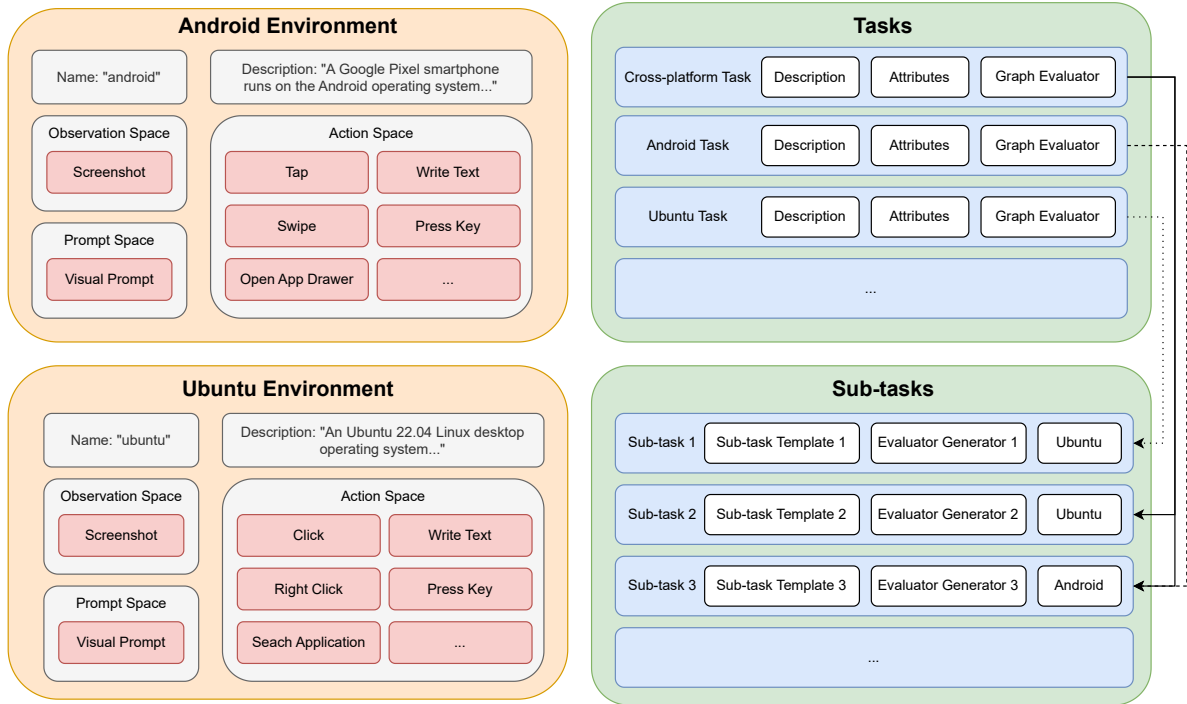


Figure 3: **Module Structure of CRAB Benchmark-v0.** The benchmark is divided into two primary sections: the left section, highlighted with warm hues, features two environments, while the right section, accentuated with cool hues, outlines various tasks. Each environment is defined by attributes including name, description, observation space, prompt method, and action space. Blocks marked in red denote actions. As for the tasks, they are composed of multiple sub-tasks and formulated by combine multiple evaluator sub-graphs derived from the sub-task evaluator generators. Arrows illustrate the compositional relationships between tasks and sub-tasks.

A.1 Dataset statistics

The applications in our task dataset along with the counts of tasks that utilize them is listed in Table 4 and 5. The task dataset covers a wide range of applications across two platforms, primarily focusing on daily life, programming, and office work scenarios. It is also worth noting that in our task settings, a single task often involves two or more applications. On average, each task contains 1.84 applications, according to our statistics.

The distribution of node counts of graph evaluators per task is provided in Table 6. Our task dataset includes graphs ranging from 1 to 11 nodes. It is important to note that the number of nodes depends on the complexity of the task, with more complex tasks involving larger graphs.

Our code and dataset will be open-sourced under Apache 2.0 License.

A.2 Environment Implementation Detail

The Ubuntu environment is launched on a QEMU/KVM (Bellard, 2005; Kivity et al., 2007) Virtual Machine, and the Android environment employs the Google Android Emulator⁴. Interaction with the Ubuntu environment is facilitated using PyAutoGUI⁵ and MSS⁶, which provide high-level commands for mouse and keyboard control and screen capture, respectively. For the Android environment, we use the Android Debug Bridge (ADB)⁷. The detailed action space is described in Table 2.

This paragraph shows the implementation detail of our evaluators. For sub-tasks on the Android platform, we incorporate XML-based evaluators (Xing et al., 2024). We dump UI layout as XML

⁴<https://developer.android.com/studio/run/emulator>

⁵<https://github.com/asweigart/pyautogui>

⁶<https://github.com/BoBoTiG/python-mss>

⁷<https://developer.android.com/tools/adb>

Table 4: Applications and their task counts in the Ubuntu environment.

App Name	Description	# Tasks
Terminal	GNOME terminal emulator with command line tools (e.g., cat, wget).	40
Firefox	Web browser with various web Apps (e.g., Google Docs and Search).	35
File Manager	GNOME official file manager.	25
GIMP	GNU Image Manipulation Program, open-source raster graphics editor.	13
System Setting	GNOME system setting GUI application.	11
VSCode	Code editor.	8
LibreOffice Writer	Word processor.	8
LibreOffice Impress	Presentation program.	7
LibreOffice Calc	Spreadsheet program.	6
Vim	CLI text editor.	6
Slack	Team communication platform.	1

path and verify whether the UI content matches the expected state. For the Ubuntu platform, we employ image matching techniques (Potje et al., 2024; Jiang et al., 2024) and OCR to handle scenarios where acquiring necessary state information through conventional APIs is challenging. Image matching offers fine-grained visual correspondences by comparing keypoint features between images, allowing us to assess spatial relationships among visual elements. Using OCR and image matching, we can accurately evaluate tasks such as verifying whether an agent has successfully created a slide with specified images, text content, and layouts—tasks for which trivial evaluation methods are lacking. We utilize EasyOCR¹ and XFeat⁸ as our primary tools for OCR and image matching. For tasks with real-time characteristics that may change over time, we implement crawler scripts to capture dynamic values at the moment of evaluation. These values are then compared with the results achieved by the agent upon task completion. We have a total of 59 evaluator functions with different types. Each task has 4.2 evaluators in average of the whole dataset.

A.3 Framework Design

CRAB offers a modular and extensible framework for evaluating agent performance in diverse tasks. At the heart of the framework lies the *action*, a unit operation representing the fundamental operation

within the benchmark. The *action* is essentially an executable Python function that can be defined with explicit typed parameters and a clear description. *actions* serve not only as building blocks but also as interfaces through which agents interact with the environment. The *evaluator* is a specialized *action* restricted to returning boolean values, signifying the success or failure of an agent’s task. It enhances the *actions* by analyzing the state of the environment and the sequence of *actions* executed by the agent, providing a decisive metric of task accomplishment. Additionally, multiple *evaluators* can be interconnected to form a graph evaluator for complex tasks (Sec. 3.3).

The *benchmark* is a key definition in the framework. A benchmark includes multiple *environments* and *cross-environment tasks*. The *environment* is formed by an action space and an observation space, which are both defined by a list of *actions*, and other essential parameters necessary for its configuration. This composite structure facilitates the execution and monitoring of *actions*, whether on local machines, remote servers, virtual machines, or physical devices networked together. A *task* encapsulates a natural language description and a graph evaluator.

CRAB utilizes Python functions to define all actions and evaluators, embodying a "code as configuration" philosophy. Each function’s docstring outlines its description and parameter definitions, which are then presented to the agent as structured prompts. Compared to traditional methods using

⁸https://github.com/verlab/accelerated_features

Table 5: Applications and their task counts in the Android environment.

App Name	Description	# Tasks
Google Map	Map application.	13
Google Calendar	Calendar application.	9
Gmail	Google mail service application.	7
Google Keep	Google note application.	6
Google Tasks	Google TO-DO list.	5
Messages	Android built-in message sending application.	5
Contacts	Android built-in contacts application.	5
Google Drive	Google Cloud Drive application.	4
Clock	Android built-in clock application.	2
Files	Android built-in file manager.	1
Settings	Android system setting.	1
Camera	Android built-in camera.	1
Google Docs	Google online word processor.	1
Phone	Android built-in phone calling application.	1

Table 6: Node count histogram.

# Nodes	1	2	3	4	5	6	7	8	9	10	11
# Tasks	5	16	29	26	14	18	7	1	0	3	1

data interchange formats like JSON or YAML, Python code configurations provide a more structured approach and fits in modern IDE.

By decoupling actions, environments, tasks, and evaluations, CRAB facilitates a plug-and-play architecture that can adapt to various scenarios. Such a system is scalable, maintainable and expandable, allowing researchers and developers to introduce new tasks and environments without restructuring the entire framework. Our implementation uses *networkx* (Hagberg et al.) for building graph and *dill* (McKerns et al.) for function serialization in our implementation.

A.4 Configuration by Modules

Building on the declarative and modular design of our framework, this section explains the configuration and potential extensibility of each module.

Environment The environments in CRAB are a combination of multiple different uses of actions with some environment metadata, such as name and natural language description. In CRAB Benchmark-v0, we use a computer desktop environment and a smartphone environment both based on virtual machine technology. The computer desktop environment, named *Ubuntu*, is installed from an ISO image of Ubuntu 22.04.4 LTS (Jammy Jel-

lyfish) downloaded from the Ubuntu Official website⁹. Necessary applications such as the LibreOffice suite (Writer, Calc, and Impress) and Slack are installed later via snap and apt, according to the task dataset requirements. The smartphone environment, named *Android*, is installed using pre-defined devices (Google Pixel 8 Pro with release name *R*) provided in Google Android Studio¹⁰. We install additional required applications such as *Keep Notes*, *Tasks*, and *Docs* from Google Play. The descriptions of the two environments in CRAB Benchmark-v0, which are inserted in the agent prompts, are as follows:

- **Ubuntu:** An Ubuntu 22.04 Linux desktop operating system. The interface displays a current screenshot at each step and primarily supports interaction via mouse and keyboard. You must use searching functionality to open any application in the system. This device includes system-related applications including Terminal, Files, Text Editor, Vim, and Settings. It also features Firefox as the web browser, and the LibreOffice suite—Writer, Calc, and Impress. For communication, Slack

⁹<https://releases.ubuntu.com/jammy/ubuntu-22.04.4-desktop-amd64.iso>

¹⁰<https://developer.android.com/studio>

is available. The Google account is pre-logged in on Firefox, synchronized with the same account used in the Android environment.

- **Android:** A Google Pixel smartphone runs on the Android operating system. The interface displays a current screenshot at each step and primarily supports interaction through tapping and typing. This device offers a suite of standard applications including Phone, Photos, Camera, Chrome, and Calendar, among others. Access the app drawer to view all installed applications on the device. The Google account is pre-logged in, synchronized with the same account used in the Ubuntu environment.

Action Action implementation in CRAB Benchmark-v0 utilize the dynamic feature of Python. It provides an intuitive method to define actions through Python function. Here is an example of action `search_application` in the Ubuntu environment:

```
@action
def search_application(name: str) ->
    None:
    """Search an application name.

    For exmaple, if you want to open an
    application named "slack",
    you can call search_application(name
    ="slack"). You MUST use this
    action to search for applications.

    Args:
        name: the application name.
    """
    pyautogui.hotkey("win", "a")
    time.sleep(0.5)
    pyautogui.write(name)
    time.sleep(0.5)
```

Listing 1: Define "search_application" action.

We extract key information from the function through the `@action` decorator as following:

- **Name:** The action name serves as the identifier for backend models. It should semantically match the action's behavior to improve the accuracy of the agent in executing the action. The function name is extracted as the action name. In this example, `search_application` is the assigned name.
- **Description:** The description provides a natural language explanation of the action to assist the agent in understanding how to use it. The main body of the function's docstring is

used as the description. For example, in this instance, the description outlines the basic usage of the action: *Search an application name*, along with an example of its usage.

- **Parameters:** The parameters are the arguments that the functions accept, offering flexibility for the agent to control the environment. Typically, a set of parameters is defined, each consisting of a name, type, and a natural language description. Parameters are extracted from the function's parameters along with their type annotations. Additionally, parameter descriptions are extracted from the `Args` section in the docstring. In this example, there is only one parameter named `name`, with a type of `str`, and its description is the application name.
- **Entry:** The entry represents the implementation of the function, defined within the function body to specify how the action is executed. When the agent invokes the function, the entry is executed with the provided parameters. In this example, we utilize the `pyautogui` package for keyboard control. Initially, it presses a hotkey to enter the application search panel in Ubuntu, then proceeds to type the application name provided by the parameters, finally displaying the search results.

Observation The observation space is represented by a set of actions. These observation actions are designed to be parameter-free and return an observation result. For instance, within the Ubuntu environment, the sole observation action available is the screenshot function, defined as follows:

```
@action
def screenshot() -> str:
    """Capture the current screen as a
    screenshot."""
    with mss() as sct:
        # Capture raw pixels from the screen
        sct_img = sct.grab(sct.monitors[1])
        # Convert to PNG format
        png = tools.to_png(sct_img.rgb,
            sct_img.size)
        # Encode to Base64 format for easier
        transmission
        base64_img = base64.b64encode(png).
            decode("utf-8")
        return base64_img
```

Listing 2: Define the "screenshot" observation action.

This action captures the screen's current view and encodes it in Base64 format. Additionally, visual

prompts are also defined by actions that utilize the output from an observation action as their input, further processing it to generate a visual prompt for the agent.

Evaluator The evaluator in CRAB Benchmark-v0 is crafted to assess the outcome of actions performed by the agent within the environment. The evaluator is defined as an action that outputs a boolean value. An example of an evaluator in the Ubuntu environment is the `check_text_in_current_window_name` function, outlined below:

```
@evaluator(env_name="ubuntu")
def check_text_in_current_window_name(
    text: str) -> bool:
    try:
        out = subprocess.check_output(
            ["xdotool", "getwindowfocus",
            "getwindowname"], text=True
        ).strip()
    except subprocess.CalledProcessError:
        return False
    return text in out
```

Listing 3: Define "check_text_in_current_window_name" evaluator.

The evaluator function is denoted with an `@evaluator` decorator and specifies its operating environment. The function's primary role is to execute a check within the system and return a boolean value indicating success or failure based on the condition being evaluated. Here, the function aims to verify whether a specified text appears in the title of the currently focused window. This is achieved through the use of the `subprocess` module to execute system commands that fetch the window's title, checking if the provided text parameter is contained within it.

Task Following a declarative programming paradigm, the task is defined as a data model. Here is an example of a cross-platform task in the dataset:

```
Task(
    id="a3476778-e512-40ca-b1c0-d7aab0c7f18b",
    description="Open \"Tasks\" app on Android, check the...",
    evaluator=path_graph(
        check_current_package_name("com.google.android.apps.tasks"),
        check_current_window_process("gnome-control-center"),
        check_color_scheme("prefer-dark")
    ),
)
```

```
)
```

Listing 4: Define a task.

In this model, each task is represented as an instance of the `Task` class, which is a subclass of `BaseModel` in *Pydantic*¹¹ package. Each task is uniquely identified by an ID and described by a detailed description. The evaluator component is structured as a graph evaluator, which integrates multiple evaluative functions into a directed graph using the *networkx*¹² package. Each evaluator within this graph must be appropriately parameterized to assess specific conditions relevant to the task. For example, the task demonstrated aims to open the "Tasks" app on Android and perform a series of verifications: it checks whether the correct Android app is opened, whether the current focused window's process name is `gnome-control-center`, and whether the color scheme is set to dark.

Sub-task The sub-task in CRAB is the unit component of in task construction. The following example is a sub-task template that we used to easily generate sub-tasks:

```
SubTask(
    id="0f589bf9-9b26-4581-8b78-2961b115ab49",
    description="Open \"{file_path}\" using vim in a terminal, write \"{content}\"", then save and exit vim.",
    attribute_dict={"file_path": "file_path", "content": "message"},
    output_type="file_path",
    evaluator_generator=lambda file_path, content: path_graph(
        check_current_window_process("gnome-terminal-server"),
        is_process_open("vim"),
        is_process_close("vim"),
        check_file_content(file_path, content),
    ),
),
```

Listing 5: Define a task.

In this sub-task model, each sub-task is defined using a similar approach to the main task. The attributes of the sub-task are outlined in an `attribute_dict`, which details the types and roles of each attribute used in the sub-task's operations. The `output_type` field specifies the expected type of output from the sub-task. The types reflected

¹¹<https://pydantic.dev/>

¹²<https://networkx.org/>

in `attribute_dict` and `output_type`, play a critical role in determining the compatibility and sequential logic of compose multiple sub-tasks. The evaluator for the sub-task is dynamically generated using a lambda function, which crafts an evaluator sub-graph based on the sub-task’s attributes.

A.5 Composed Task Format

We use a JSON format to save the composed tasks, which includes the task ID, overall task description, sub-tasks with their attribute values, and a graph structure represented in an adjacency list. The entire task dataset is defined by the sub-task pool in Python code and the task composition JSON files categorized by task platform.

```
{
  "description": "Combine Image 1 \"/\n/home/crab/Pictures/cat.png\" and\nImage 2 \"/\n/home/crab/assets/campus.\n.png\" using GIMP (GNU Image\nManipulation Program), placing Image\n1 on the left side of Image 2, and\nsave the combined image to \"/\n/home/\ncrab/Desktop/background.png\". Then,\nset this combined image as the\nscreen background of the system.",
  "tasks": [
    {
      "task": "4cf246ea-0a7f-43da\n-84b6-61d74a2699af",
      "attribute": {
        "image_path_1": "/home/\ncrab/Pictures/cat.png",
        "image_path_2": "/home/\ncrab/assets/campus.png",
        "output_path": "/home/\ncrab/Desktop/background.png"
      },
      "output": "/home/crab/\nDesktop/background.png"
    },
    {
      "task": "a207ef38-b3b2-4c6c-\na1e3-75c38162f5ba",
      "attribute": {
        "photo_path": "/home/\ncrab/Desktop/background.png"
      },
      "output": null
    }
  ],
  "adjlist": "0 1\n1",
  "id": "d3c917ff-406f-447a-87f5-\nb8d835cba750"
}
```

Listing 6: Define a composite task in JSON.

B Agent system

B.1 Agent Implementation

In this section, we outline the implementation of the agents used in our experiments, which lever-

age advanced multimodal language models from OpenAI, Anthropic, and Google. Each agent is designed to function in multi-environment setups, interacting with various action spaces defined by different environments.

General Framework All agents share a common architecture but are tailored to the specific APIs and capabilities of each language model provider.

Initialization Each agent is initialized with several key parameters, including a description, an action space, the model type, maximum tokens, history message length, and an optional environment description. The initialization process involves:

- **Action Space Conversion:** Actions defined for each environment are converted into a schema compatible with the respective API. This ensures that the actions can be correctly interpreted and executed by the language models.
- **System Message Setup:** Depending on whether the agent is configured for single or multiple environments, a system message is formatted to provide the model with context about the tasks and environments.

Interaction (Chat Method) The core functionality of each agent is encapsulated in its ability to interact with users through a chat method. This involves:

- **Content Parsing:** Input content is parsed and formatted to match the requirements of the respective API. This includes structuring user messages and any necessary contextual information.
- **Request Construction:** The request payload is constructed, incorporating the system message, chat history, and the newly parsed user input.
- **API Communication:** The constructed request is sent to the appropriate API, which generates a response. The agents handle API-specific constraints such as rate limits and response formats.
- **Response Handling:** The response from the API is processed to extract any tool calls suggested by the model. These are then appended to the chat history, maintaining a coherent conversation state.

Multi-Environment Support For agents configured to operate in multiple environments, additional logic ensures that actions are correctly associated with their respective environments. This involves modifying action names and descriptions to reflect their environmental context and handling responses accordingly.

Utilities and Shared Functions Several utility functions support the operation of these agents, facilitating tasks such as content parsing, action prompt generation, and schema conversion. These shared functions ensure consistency and reduce redundancy across the different agent implementations.

B.2 Inter-agent Communication Strategies

In this section we introduce the details of two multi-agent communications methods, which are introduced in 4.2.

Multi-agent Communication by Functionality

This setting involves two agents: a main agent prompted with the task description and a tool agent with the entire action space. The main agent generates the instruction for the next step and sends it to the tool agent. The tool agent chooses the proper action with parameters and a target environment, then feeds it back to the system.

Multi-agent Communication by Environment

This setting involves four agents in our benchmark setting: a main agent prompted with the task description and three tool agents, each corresponding to the environments of Android, Ubuntu, and Root, with the respective action spaces. The main agent generates the instruction for the next step and sends it to the tool agents. Each sub-environment agent receives the message containing the instruction and environment observation information. The environment agents process the message using their specialized models and action schemas, performing the required actions within their environments.

B.3 Agent Prompt

B.3.1 Single Agent

Prompt

You are a helpful assistant. Now you have to do a task as described below:
`**{task_description}**`.
You should never forget this task and always perform actions to achieve this task. And

this is the description of each given environment: `{env_description}`. A unit operation you can perform is called action in a given environment. For each environment, you are given a limited action space as function calls:

`{action_descriptions}`

You may receive a screenshot of the current system. You may receive a screenshot of a smartphone app. The interactive UI elements on the screenshot are labeled with numeric tags starting from 1.

In each step, You **MUST** explain what do you see from the current observation and the plan of the next action, then use a provided action in each step to achieve the task. You should state what action to take and what the parameters should be. Your answer **MUST** be a least one function call. You **SHOULD NEVER** ask me to do anything for you. Always do them by yourself using function calls.

Prompt

You are a helpful assistant. Now you have to do a task as described below:

`**{task_description}**`

You should never forget this task and always perform actions to achieve this task. And this is the description of each given environment: `{env_description}`. You will receive screenshots of the environments. The interactive UI elements on the screenshot are labeled with numeric tags starting from 1.

A unit operation you can perform is called Action. You have a limited action space as function calls: `{action_descriptions}`. You should generate JSON code blocks to execute the actions. Each code block **MUST** contain only one json object, i.e. one action. You can output multiple code blocks to execute multiple actions in a single step. You must follow the JSON format below to output the action.

```
{"name": "action_name",  
 "arguments": {"arg1": "value1",  
 "arg2": "value2"}}  
or if not arguments needed:
```



```
{"name": "action_name",  
  "arguments": {}}
```

You MUST use exactly the same "action_name" as I gave to you in the action space. You SHOULD NOT add any comments in the code blocks.

In each step, You MUST explain what you see from the current observation and the plan of the next action, then use a provided action in each step to achieve the task. You should state what action to take and what the parameters should be. Your answer MUST contain at least one code block. You SHOULD NEVER ask me to do anything for you. Always do them by yourself.

B.3.2 Multi-Agent by Functionality

Main Agent Prompt

You are a helpful assistant. Now you have to do a task as described below: {task_description}. And this is the description of each given environment: {env_description}. A unit operation you can perform is called action in a given environment. For each environment, you are given a limited action space as function calls:

{action_descriptions}

You may receive a screenshot of the current system. The interactive UI elements on the screenshot are labeled with numeric tags starting from 1. For each step, You must state what actions to take, what the parameters are, and you MUST provide in which environment to perform these actions.

Tool Agent Prompt

You are a helpful assistant in generating function calls. I will give you a detailed description of what actions to take next, you should translate it into function calls. please do not output any other information.

B.3.3 Multi-Agent by Environment

Main Agent Prompt

You are a main agent, and your goal is to plan and give instructions to sub-agents in each environment to complete the final task. Now you have to do a task as described below: {description}. The description of each given environment: {env_description}. For each step, you are required to provide high-level instructions detailing the next actions to be taken. Additionally, you must specify which sub-agent in the designated environment should execute these instructions. If a sub-agent is not needed for a particular step, you may instruct it to skip that step.

Root Environment Agent Prompt

You are a sub-agent responsible for the crab benchmark root environment. Your goal is to assist the main agent in completing the whole task: "{description}". You can only complete the task or submit the result when the main agent tells you the whole task has been completed. Otherwise, you can only call SKIP.

Sub-environment Agent Prompt

You are a sub-agent responsible for the {environment} environment. The description of the {environment} environment is: {env_description}. Your goal is to assist the main agent in completing the final task by performing actions in the {environment} environment according to the instructions from the main agent. The final task is described below: {task_description}. A unit operation you can perform is called action in a given environment. You can only execute action in the {environment} environment. For the {environment} environment, you are given a limited action space as function calls:

{action_descriptions}

The interactive UI elements on the screenshot are labeled with numeric tags starting from 1. For each step, You will receive an instruction telling you what you need to do

next. After analyzing the instruction you received and the current {environment} system, if you think you don't need to do anything in the current {environment} system, you should choose SKIP action. Otherwise, you must state what actions to take, what the parameters are, and you MUST provide in which environment to perform these actions. Your answer must be function calls. Please do not output any other information. You must make sure all function calls get their required parameters.

C Further Result Analysis

This section further discusses our experimental results in detail. Section C.1 categorizes the results into three types of tasks: Ubuntu, Android, and cross-platform, and provides further analysis. Section C.3 examines three specific tasks and analyzes the performance of different agent settings on each.

C.1 Result by Platforms

Table 7, 8 and 9 show the experiment results on Ubuntu Tasks, Android Tasks, and cross-platform Tasks, respectively.

We find that certain models demonstrate a distinct preference or better alignment with specific platforms. The GPT-4o, Gemini, and Claude models, for instance, show notably better outcomes on Android platforms. This suggests potential optimizations or intrinsic features within these models that cater effectively to the Android environment's requirements. Conversely, the GPT-4 Turbo model exhibits superior performance on Ubuntu tasks, hinting at possible architectural or training aspects that are better suited for that specific environment.

Cross-platform tasks necessitate functionality across different operating systems or platforms, demand a broader capability range and more sophisticated agent coordination. The importance of CR is especially critical in such environments, where it serves as a more reliable metric for distinguishing between agent models than SR. Given the presence of all Gemini, Claude, and open source model agents' SR is 0.0, indicating that Completion Ratio more effectively captures an agent model's capability, thereby better reflecting its robustness and adaptability to complex requirements. On cross-platform tasks, GPT-4 Turbo (Single) exhibits a

CR of 52.61%, which indicates that even though SR might be lower, the agent covers a significant portion of task objectives before termination.

Furthermore, analyzing the reasons for task termination offers additional insights into the operational challenges these models encounter. False Completion is notably prevalent in Android tasks. Reach Step Limit remains the most frequent cause of termination, particularly in cross-platform tasks. The Claude model exhibits a significantly high Invalid Action ratio in cross-platform tasks, indicating its difficulties in managing multi-environment scenarios effectively. The GPT-4o with JSON mode shows an extremely high IA ratio in Android tasks, proving the serious hallucination problem under this setting.

Overall, these findings underscore the necessity of selecting the appropriate agent model and configuration based on specific platform and task needs. The variability in model performance across different setups also highlights the ongoing need for development and refinement of multi-agent systems to enhance their versatility and efficacy in increasingly diverse and complex operational environments. These results comparing SR and CR also demonstrate the importance of our graph evaluator in agent evaluation.

C.2 Comparison between Single Agent and Multi-agent

The experimental results indicate that multi-agent structures perform slightly worse than single-agent systems, which is somewhat unusual. We analyse the possible reasons here.

First, comparing in False Completion Rate, we attribute the lower Success Rate (SR) of Multi-agent to a high False Completion Rate—where the agent incorrectly assumes that the task is complete. As observed in failure cases (e.g., the Cross-platform Task case study in Appendix C.3), Sub-agents often misinterpret the Main agent's instructions. Despite being required to perform a final action, the instructions lead Sub-agents to prematurely conclude that the task is complete, resulting in incorrect "complete" actions. While this issue also occurs in Multi-Env, it happens less frequently. By analysing the communication logs, we believe this is due to information loss during inter-agent communication. Sometimes, the main agent gives a correct instruction, but the sub-agent misunderstands it because it does not have the context. Natural language, while effective for aligning

with human understanding in LLM communication, is less suited for inter-agent communication, leading to information loss during compression and interpretation, which weakens the performance of multi-agent structures.

Next, comparing in Invalid Action Rate, we observe that in single-platform tasks, both Multi-Env and Multi-Func suffer from similar inter-agent communication issues, as indicated by their high Invalid Action rates. However, in cross-platform tasks (Table 9), the Single agent's Invalid Action rate is significantly higher than that of the Multi-agent by environment structures on GPT-4o model. Cross-platform tasks require frequent environment changes with varying action spaces, and if the model's performance output is inadequate, it often generates correct actions in the wrong environment, invalid actions in the correct environment, or correct actions in correct environment but in the wrong format. This phenomenon highlights the limitations of current general-purpose LLMs, where multi-agent structures can be advantageous. By assigning each agent a specific responsibility and a limited action space, multi-agent structures can mitigate these issues.

To improve multi-agent system performance, we suggest to follow two approaches: (1) Developing better multi-agent structures to minimize information loss during communication, and (2) Introducing a critical agent to correct hallucinations or information loss during communication. These improvements, however, come with a trade-off, namely an increase in token costs within the agent system. Within our benchmark framework, users can utilize the error log we provide to analyze the bottlenecks of their agents and refine their designs.

C.3 Case Study

To better understand how different agents perform the same task and exhibit varied properties, we present visual results along with detailed metrics and logs for three cases by platform. The screenshots illustrate the progress of agents executing tasks according to specific natural language instructions.

C.3.1 Cross-platform Task

Task: Open the "Tasks" app on an Android device, check the first incomplete task, and then execute it as described. The first task, found incomplete in the "Tasks" app, involves **switching the system to dark mode in Ubuntu via the "Set-**

tings" application.

This task exemplifies message passing across different environments, where the "incomplete task" serves as the critical information that the agent must relay and apply in the Ubuntu setting. These two phases—retrieving the task details via the phone and executing the task on a computer—are inseparably linked and cannot be treated as distinct tasks. The agent can only proceed to the second stage after successfully acquiring information from the first.

In this task, GPT-4o (single agent), GPT-4 Turbo (single agent), and GPT-4 Turbo (multi-agent by functionality) all successfully complete the task using the minimal steps necessary to locate and execute the task, demonstrating their efficiency in managing multiple environments simultaneously. On the other hand, both GPT-4o (multi-agent by functionality) and GPT-4o (multi-agent by environment) also perform commendably, completing the task up until the final step. However, after incorrectly performing the last step, they both erroneously conclude the task is completed and exit. This indicates a communication breakdown, where the sub-agents misinterpret the instructions from the main agent. The remaining four agents fail to complete the task. Agents equipped with the Gemini model do not even manage to open the "Tasks" app within the allocated step limit, whereas agents with the Claude model quickly open the "Tasks" app to complete the first step but fail at the task execution. The performance disparity between single-agent and multi-agent configurations in both the Gemini and Claude models highlights the variance in capability across different models and devices.

C.3.2 Ubuntu Task

Task: Create a new directory `"/home/crab/assets_copy"` and copy all files with the specified `"txt"` extension from `"/home/crab/assets"` to the directory `"/home/crab/assets_copy"`.

This task can be approached through multiple methods. An agent may opt for a straightforward strategy first using the `search_application` command to find the Terminal, then using Linux commands to create the directory and copy the necessary files. Alternatively, the agent could employ a GUI-based approach, manually creating the folder and selecting files through actions like `click` and `right_click`. We evaluate various agent systems in a single-agent setting for this task. As illustrated in Table 10–13, both GPT-4o and GPT-4 Turbo

Table 7: Evaluation results on Ubuntu tasks.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4o	Single	9.59	30.82	3.22	4.87×10^{-4}	6.85	58.90	24.66
GPT-4o	By Func	9.59	24.20	2.72	4.30×10^{-4}	5.48	63.01	21.92
GPT-4o	By Env	10.96	22.88	2.74	2.29×10^{-4}	5.48	43.84	39.73
GPT-4 TURBO	Single	10.96	31.09	4.08	5.57×10^{-4}	2.74	65.75	20.55
GPT-4 TURBO	By Func	12.33	28.95	3.70	4.18×10^{-4}	8.22	32.88	46.58
GEMINI 1.5 PRO	Single	1.37	7.76	0.63	n/a	0.00	47.95	50.68
GEMINI 1.5 PRO	By Func	1.37	3.31	0.33	n/a	0.00	20.55	78.08
CLAUDE 3 OPUS	Single	0.00	9.54	0.72	0.63×10^{-4}	8.22	58.90	32.88
CLAUDE 3 OPUS	By Func	0.00	4.93	0.46	0.47×10^{-4}	27.40	34.25	38.36
GPT-4o w/o FC	Single	10.96	22.58	2.30	4.49×10^{-4}	5.48	54.79	28.77
PIXTRAL-12B	Single	0.00	2.97	0.22	0.24×10^{-4}	1.37	80.82	17.81
LLAVA-OV-72B	Single	0.00	3.31	0.20	0.35×10^{-4}	17.81	64.38	17.81

Table 8: Evaluation results on Android tasks.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4o	Single	24.14	47.91	5.84	7.17×10^{-4}	13.79	58.62	3.45
GPT-4o	By Func	24.14	48.74	6.83	9.19×10^{-4}	24.14	37.93	13.79
GPT-4o	By Env	27.59	53.34	6.99	4.58×10^{-4}	13.79	44.83	13.79
GPT-4 TURBO	Single	6.90	27.08	2.60	2.87×10^{-4}	20.69	62.07	10.34
GPT-4 TURBO	By Func	20.69	37.01	5.00	5.92×10^{-4}	13.79	51.72	13.79
GEMINI 1.5 PRO	Single	17.24	34.52	4.82	n/a	10.34	65.52	6.90
GEMINI 1.5 PRO	By Func	17.24	35.99	4.31	n/a	31.03	37.93	13.79
CLAUDE 3 OPUS	Single	13.79	41.90	5.07	5.37×10^{-4}	20.69	55.17	10.34
CLAUDE 3 OPUS	By Func	13.79	44.02	4.75	5.35×10^{-4}	48.28	31.03	6.90
GPT-4o w/o FC	Single	10.34	14.29	1.72	2.94×10^{-4}	3.45	6.90	79.31
PIXTRAL-12B	Single	3.45	24.17	2.16	2.72×10^{-4}	0.00	65.52	31.03
LLAVA-OV-72B	Single	3.45	13.51	1.36	3.00×10^{-4}	3.45	93.10	0.00

Table 9: Evaluation results on cross-platform tasks.

Agent system		Metrics				Termination Reason		
Model	Structure	SR(%) ↑	CR(%) ↑	EE(%) ↑	CE(%) ↑	FC(%)	RSL(%)	IA(%)
GPT-4o	Single	16.67	51.24	5.21	3.98×10^{-4}	5.56	38.89	38.89
GPT-4o	By Func	22.22	50.00	4.15	3.13×10^{-4}	11.11	44.44	22.22
GPT-4o	By Env	5.56	43.54	3.22	1.60×10^{-4}	11.11	72.22	11.11
GPT-4 TURBO	Single	5.56	52.61	4.60	2.89×10^{-4}	11.11	66.67	16.67
GPT-4 TURBO	By Func	5.56	46.17	4.06	2.67×10^{-4}	16.67	50.00	27.78
GEMINI 1.5 PRO	Single	0.00	16.14	1.15	n/a	0.00	72.22	27.78
GEMINI 1.5 PRO	By Func	0.00	13.65	1.21	n/a	5.56	77.78	16.67
CLAUDE 3 OPUS	Single	0.00	24.50	1.93	1.24×10^{-4}	0.00	55.56	44.44
CLAUDE 3 OPUS	By Func	0.00	18.96	1.93	1.20×10^{-4}	0.00	38.89	61.11
GPT-4o w/o FC	Single	0.00	39.11	3.51	3.28×10^{-4}	5.56	50.00	44.44
PIXTRAL-12B	Single	0.00	12.35	0.62	0.44×10^{-4}	0.00	72.22	27.78
LLAVA-OV-72B	Single	0.00	9.07	0.48	0.53×10^{-4}	5.56	66.67	27.78

from OpenAI successfully interpret the task instructions and employ a simpler solution using Terminal commands. These agents also demonstrate superior capability in understanding the UI, selecting the correct commands, and accurately using the Terminal application to fulfill the task requirements.

Conversely, the Gemini and Claude agents, despite attempting to solve the task with Terminal, ultimately fail in different ways. Both agents struggle with precise clicking and selecting the correct icons for the intended actions, even though they share the same visual prompting mechanism as GPT-4o and GPT-4 Turbo. For instance, the Claude agent mistakenly opens the Ubuntu Desktop Guide instead of the Terminal and continues executing commands in the wrong application without realizing the error. The Gemini agent, on the other hand, unexpectedly opens the Firefox browser before correctly navigating to the Terminal but still interacts incorrectly with unrelated applications and icons. Unlike Claude, Gemini does not type in commands in the wrong applications but persists in exploring alternative methods using the Files application's UI. Despite taking significantly more steps than the GPT-4o and GPT-4 Turbo agents, neither the Claude nor the Gemini agents achieve the task's goal.

C.3.3 Android Task

Task: In Android, using the "Contacts" app, find the email of the contact named John Lauphin, then using the "Gmail" app, send an email to that contact with the subject "Hello John."

This task consists of sub-tasks across two different applications. Agents must sequentially open the two apps, retrieve the email address from the first app, and use it in the second app to send an email. This straightforward yet formal task can be completed using various methods. Agents may need to locate the contact in the Contacts app and then use the retrieved email address to send a message. We reports the performance of agents in a multi-agent setting for this challenging task. Following is the details of agents in operating the task.

GPT-4o multi-agent by functionality In steps 1-11, the agent tries to open the Contacts app but mistakenly opens Google Assistant multiple times. In steps 12-14, the agent successfully enters the Contacts app and finds the contact information. The agent then returns to the home page, and the pro-

cess is terminated due to the limitation of operation steps.

GPT-4 Turbo multi-agent by functionality In steps 1-2, the agent tries to open the Contacts app but mistakenly opens Google Messages. In steps 3-5, the agent opens the Contacts app and obtains the corresponding information. In steps 6-14, the agent repeatedly opens Google Chrome and Messages apps, failing to find the Gmail app as planned.

Gemini 1.5 Pro multi-agent by functionality In steps 1-2, the agent finds the Contacts app and enters it. However, the agent misunderstands the instruction, gets lost in creating a new contact with the given name, and cannot obtain the corresponding information.


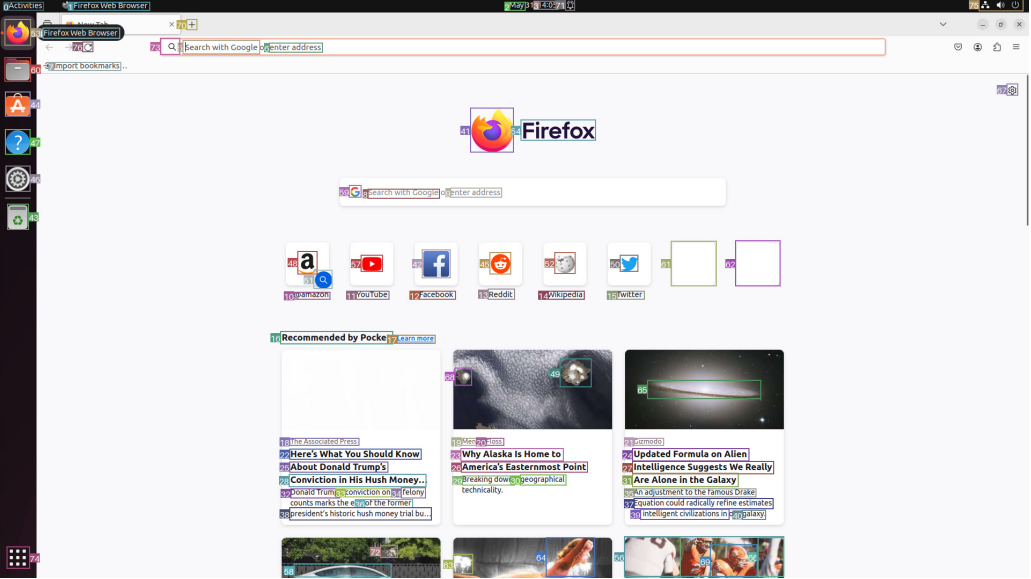
Claude 3 Opus multi-agent by functionality In steps 1-7, the agent tries to open the Contacts app but mistakenly opens Google Messages multiple times. In steps 7-11, the agent tries to open the Contacts app but mistakenly opens Google Assistant. In steps 12-14, the agent successfully enters the Contacts app and finds the contact information. The agent then returns to the home page, plans to open the Gmail app, and the process is terminated due to the limitation of operation steps.

GPT-4o multi-agent by environment In steps 1-7, the agent plans to open the Contacts app, but the operation fails due to an error in opening the app drawer, which prevents the agent from finding and tapping the Contacts app. In steps 8-11, the agent successfully enters the Contacts app and obtains the information. In steps 12-14, the agent opens the Gmail app, navigates to the sending page, and tries to input the retrieved email address as the recipient.

Analysis For the agents which are organized by functionality, Gemini 1.5 Pro struggles to complete the first operation. Although it recognizes and opens the Contacts app as instructed, it fails to proceed further. In contrast, Claude 3 Opus and GPT-4o successfully obtain the necessary information. In the initial phase, the multiple agents agree that opening the Contacts app is the first step. However, they often fail to find the correct position to tap, frequently opening incorrect apps such as Google Assistant and Messages. Once the agents do open the correct app, they usually find the email address of the contact quickly. Even when agents plan to go back home and open the Gmail app to send the message, due to the limitation of opera-

tions, the system ended. As shown in steps 3-5, GPT-4 Turbo quickly finishes the corresponding task after opening the correct app. However, similar to GPT-4o, GPT4-Turbo agents get stuck as they can not open the correct apps in the following steps. Besides, GPT-4o (multi-agent by environment) overcomes the issue encountered by GPT-4o (multi-agent by functionality). Even affected by not being able to access the app drawer, the system could still find and copy the corresponding information and change to the Gmail app for further operations.


Table 10: **Ubuntu task case with Gemini (Single):** Create a new directory `"/home/crab/assets_copy"` and copy all files with the specified `"txt"` extension from `"/home/crab/assets"` to the directory `"/home/crab/assets_copy"`.

Step	Agent Observation and Action
0	 <p data-bbox="794 936 906 967">click(5)</p>
1	 <p data-bbox="794 1576 906 1608">click(1)</p>

Continued on next page

Table 10 continued from previous page


2



click(4)

This screenshot shows the Firefox browser home page. A red rectangular box highlights the search bar, which contains the text "Search with Google" and "Enter address". The page features the Firefox logo at the top, a row of social media icons (Amazon, YouTube, Facebook, Reddit, Wikipedia, Twitter), and a "Recommended by Pocket" section with several article thumbnails and titles.

3



search_application("terminal")

This screenshot shows the same Firefox browser home page as in the previous image. A red rectangular box highlights the search bar. Additionally, a terminal window is open in the top right corner of the browser window, displaying the command "search_application('terminal')".

Continued on next page

Table 10 continued from previous page

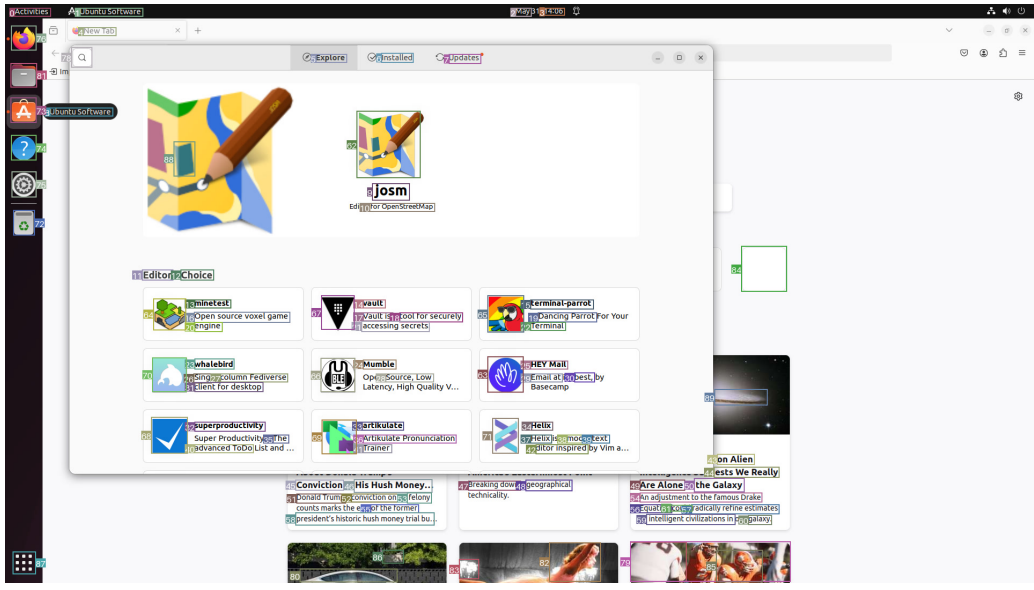
4



click(9)

Detailed description: This screenshot shows a Linux desktop environment with a dark theme. A search bar is positioned at the top center, displaying the text 'terminal'. Below the search bar, a terminal window is open, showing a prompt character '>' and the text '@terminal'. The desktop background is dark, and the left sidebar contains several application icons. The number '4' is located in the bottom-left corner of the image frame.

5



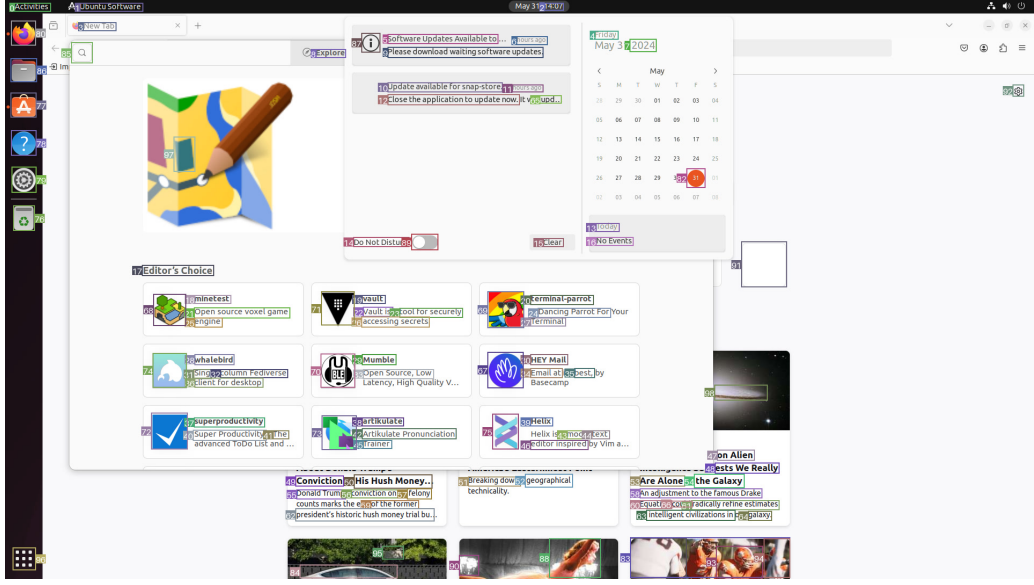
click(3)

Detailed description: This screenshot shows a Linux desktop environment with a light theme. A software center window is open, displaying a grid of application tiles. The top tile is for 'JOSM', an editor for OpenStreetMap. Below it, there is an 'Editor's Choice' section with several other application tiles, including 'minetest', 'Vault', 'Terminal-parrot', 'Whalebird', 'Humble', 'HEY Mail', 'Superproductivity', 'Artikulate', and 'Hella'. The desktop background is light, and the left sidebar contains several application icons. The number '5' is located in the bottom-left corner of the image frame.

Continued on next page

Table 10 continued from previous page

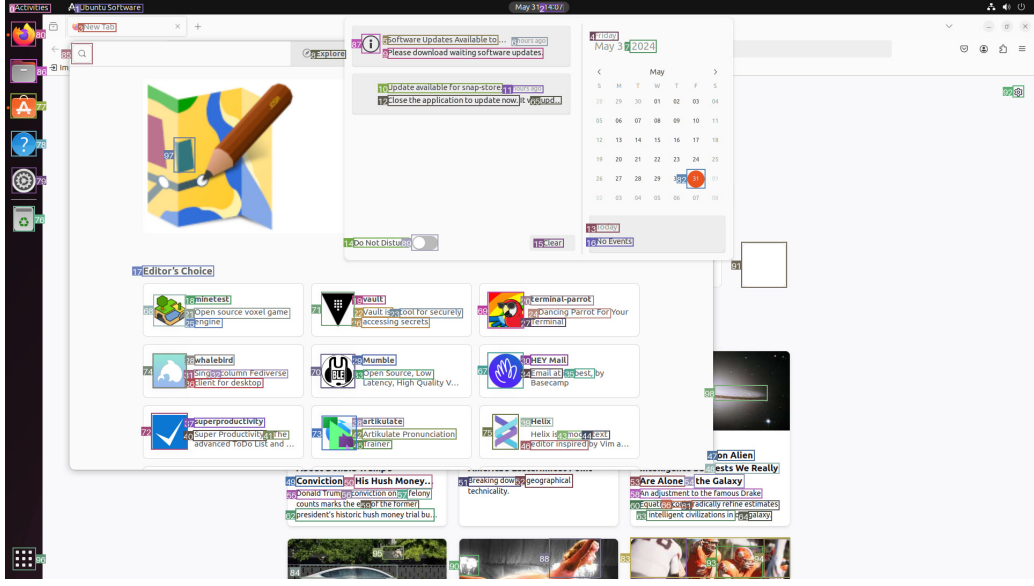
6



The screenshot shows the Ubuntu Software Center interface. At the top, there are several notification banners: 'Software Updates Available to download', 'Please download waiting software updates', 'Update available for snap-store', and 'Close the application to update now'. A calendar for May 2024 is visible in the top right. The main content area features an 'Editor's Choice' section with various application tiles, including 'minetest', 'vscodium', 'terminal-parrot', 'shateband', 'Aumble', 'SHEY Mail', 'superproductivity', 'Artikulate', and 'Helix'. Below this, there are news snippets such as 'Conviction of His Hush Money...', 'Breaking down geographical technicality.', and 'Are Alone the Galaxy'. A search bar at the bottom right contains the text 'terminal', and a red box highlights the search results for 'terminal-parrot'.

click(13)

7



The screenshot is identical to the one above, showing the Ubuntu Software Center interface with the search bar containing 'terminal'. A red box highlights the search results for 'terminal-parrot'.

search_application("terminal")

Continued on next page

Table 10 continued from previous page

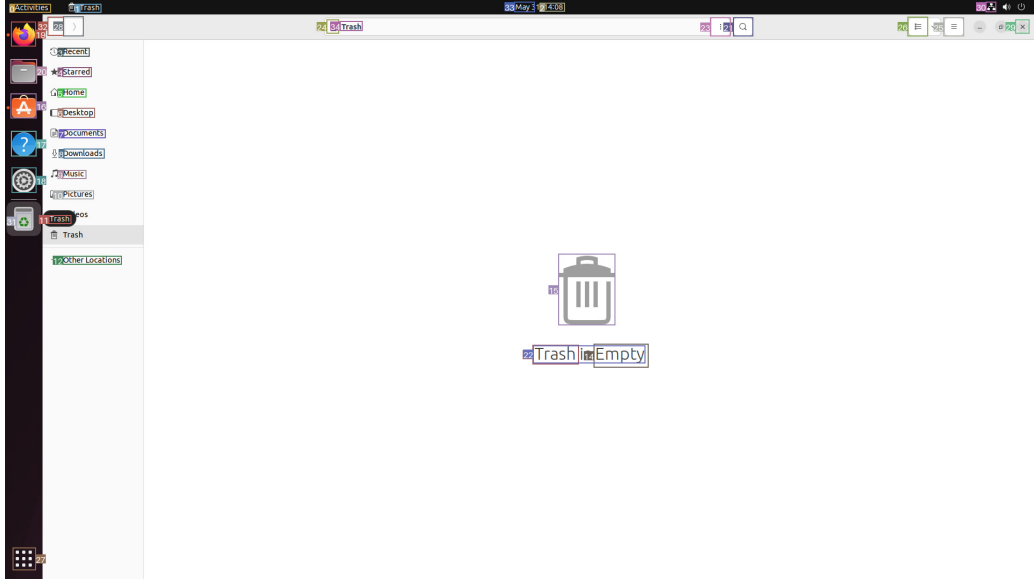
8



click(10)

Detailed description: This screenshot shows a Linux desktop environment. A terminal window is open in the center, displaying a 'gterminal' prompt. Above it, a search bar contains the text 'gterminal'. The desktop background is dark, and the left sidebar contains various application icons. The top panel shows the date 'May 3 10:13:00' and system status icons.

9



click(6)

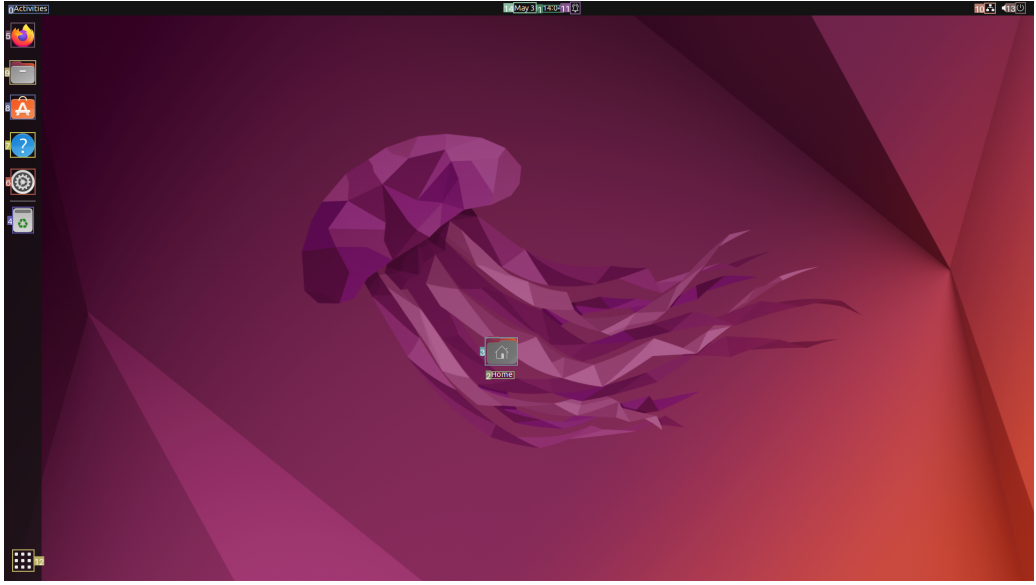
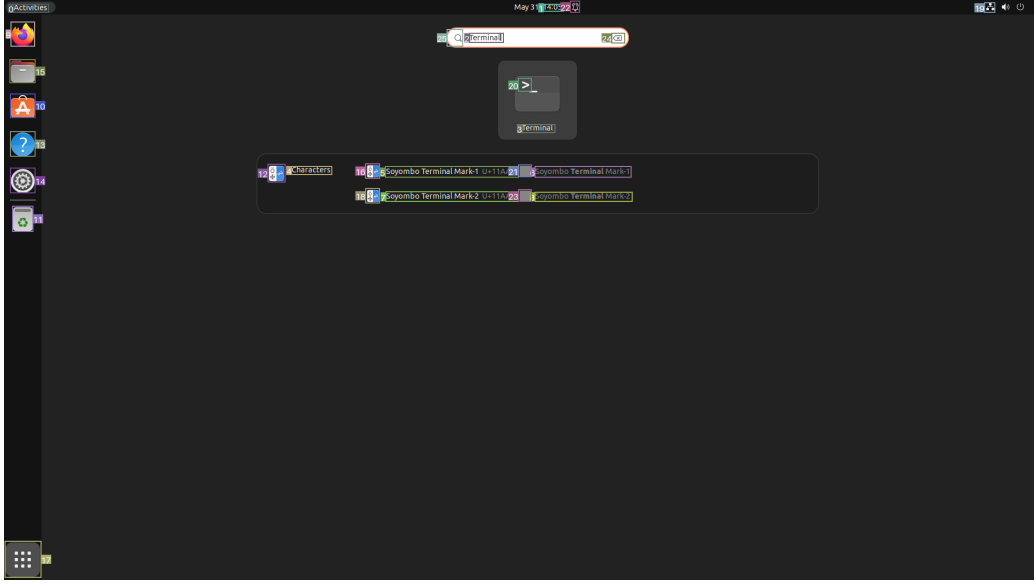
Detailed description: This screenshot shows a Linux desktop environment with a file manager window open. The window title is 'Trash' and it displays a trash can icon with the text 'Trash is Empty' below it. The left sidebar shows a list of locations including 'Recent', 'Starred', 'Home', 'Desktop', 'Documents', 'Downloads', 'Music', 'Pictures', 'Trash', and 'Other Locations'. The top panel shows the date 'May 3 10:13:00' and system status icons.

10-14

The agent is stuck at this stage and keeps clicking useless elements.

Completed Nodes / Total Nodes: 0 / 2 Termination Reason: Reach Step Limit

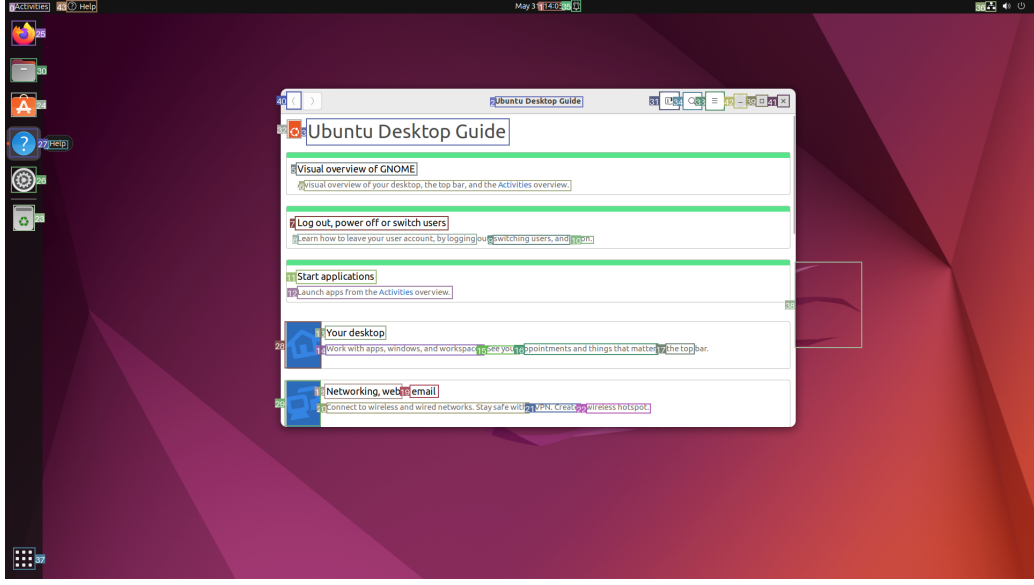
Table 11: **Ubuntu task case with Claude (Single):** Create a new directory `"/home/crab/assets_copy"` and copy all files with the specified `"txt"` extension from `"/home/crab/assets"` to the directory `"/home/crab/assets_copy"`.

Step	Agent Observation and Action
0	 <p data-bbox="619 936 1059 967">search_application("terminal")</p>
1	 <p data-bbox="772 1576 906 1608">click(13)</p>

Continued on next page

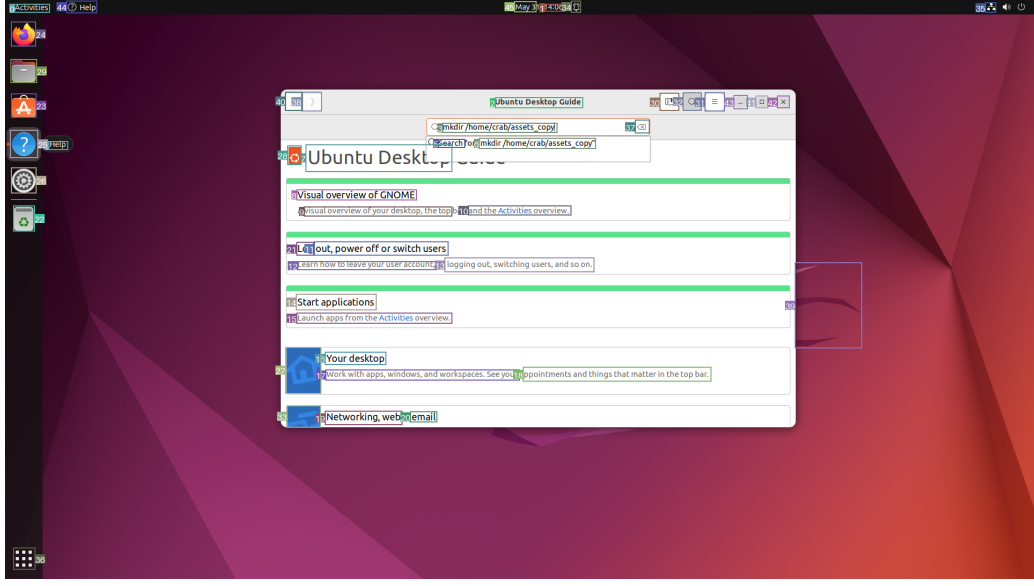
Table 11 continued from previous page

2



write_text("mkdir /home/crab/assets_copy")

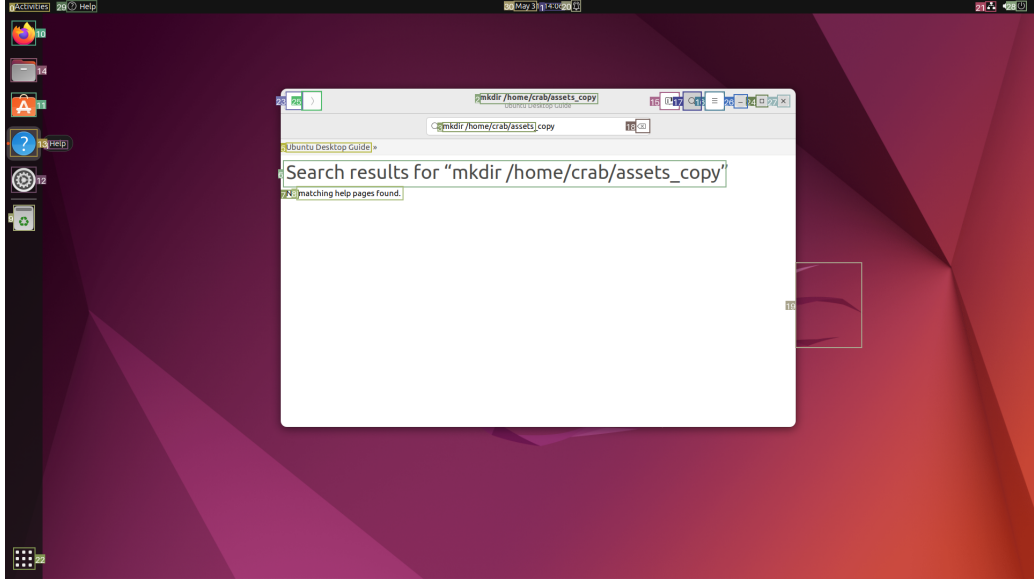
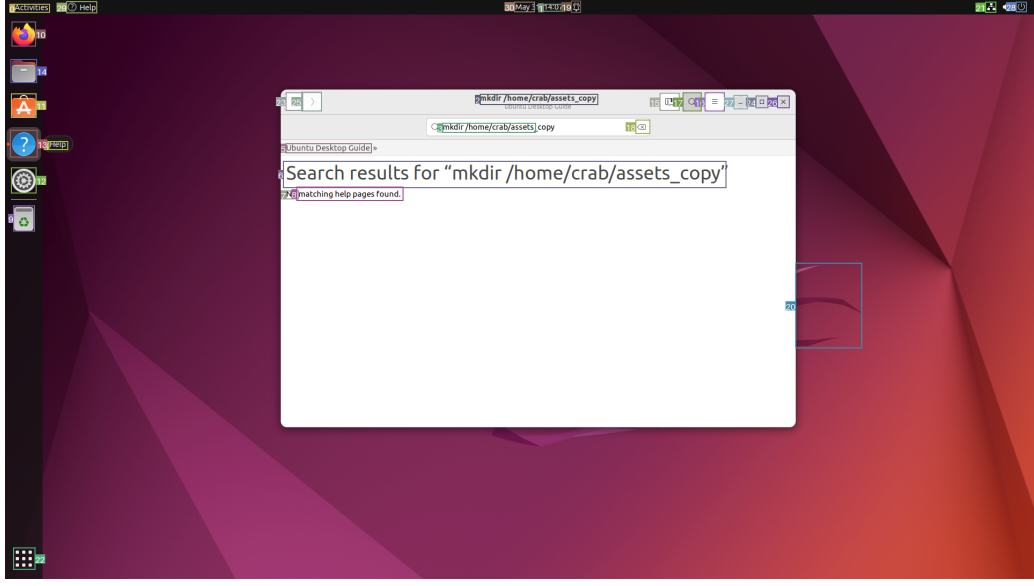
3



key_press("enter")

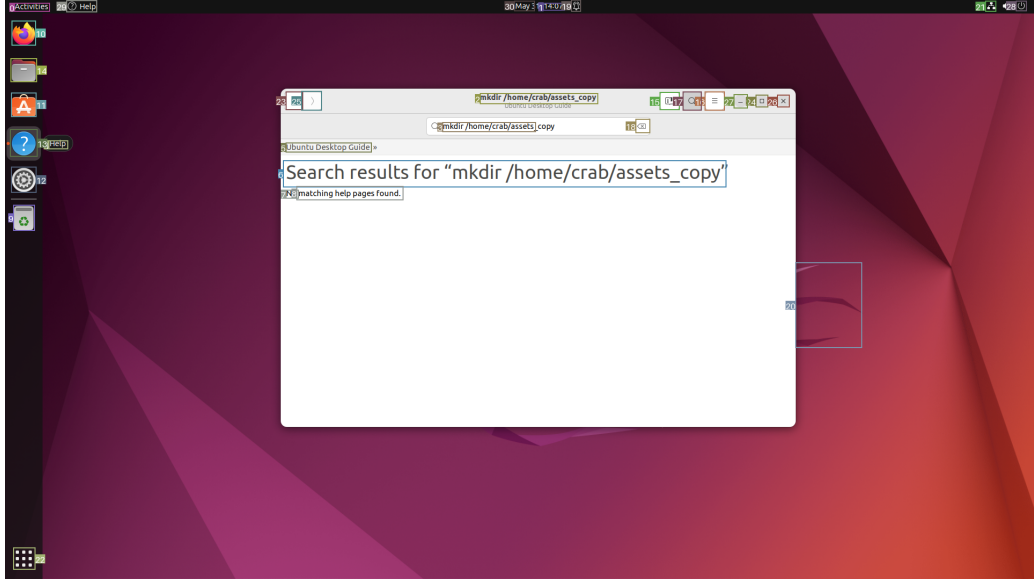
Continued on next page

Table 11 continued from previous page

4	 <p><code>write_text("cd /home/crab/assets")</code></p>
5	 <p><code>key_press("enter")</code></p>

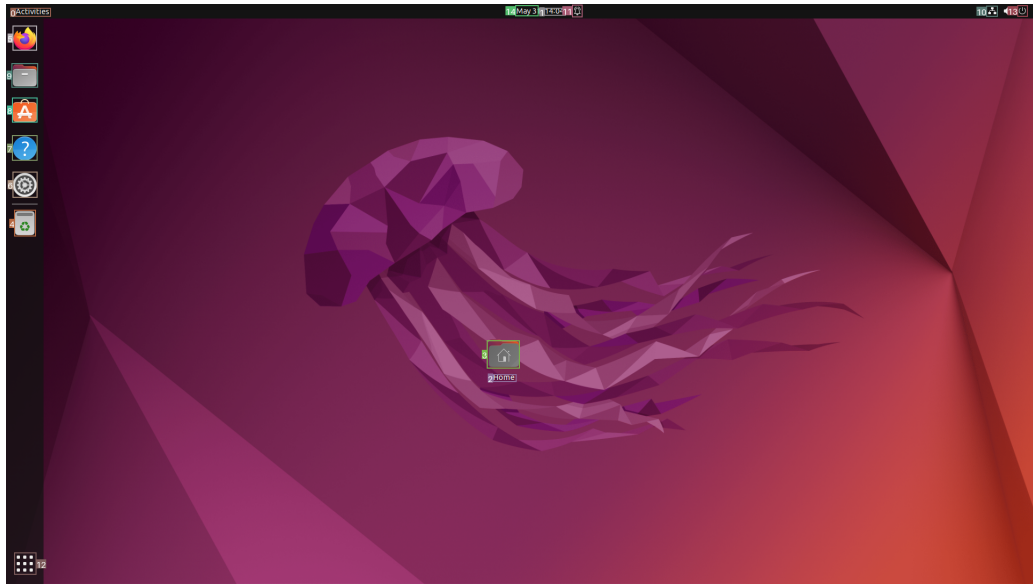
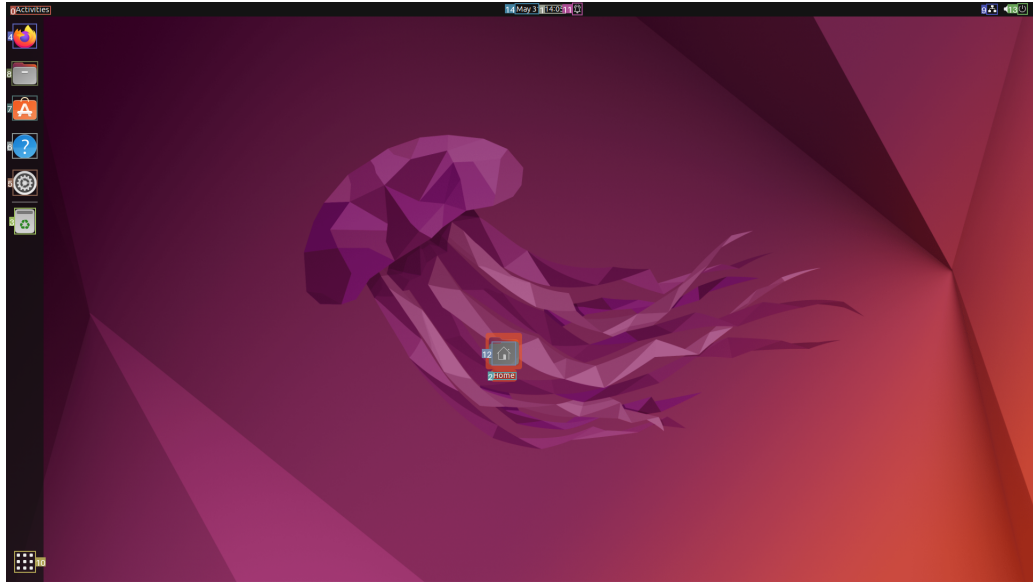
Continued on next page

Table 11 continued from previous page

6	 <p>write_text("mkdir assets_copy")</p>
7-14	The agent is stuck at this stage and keeps pressing keys.

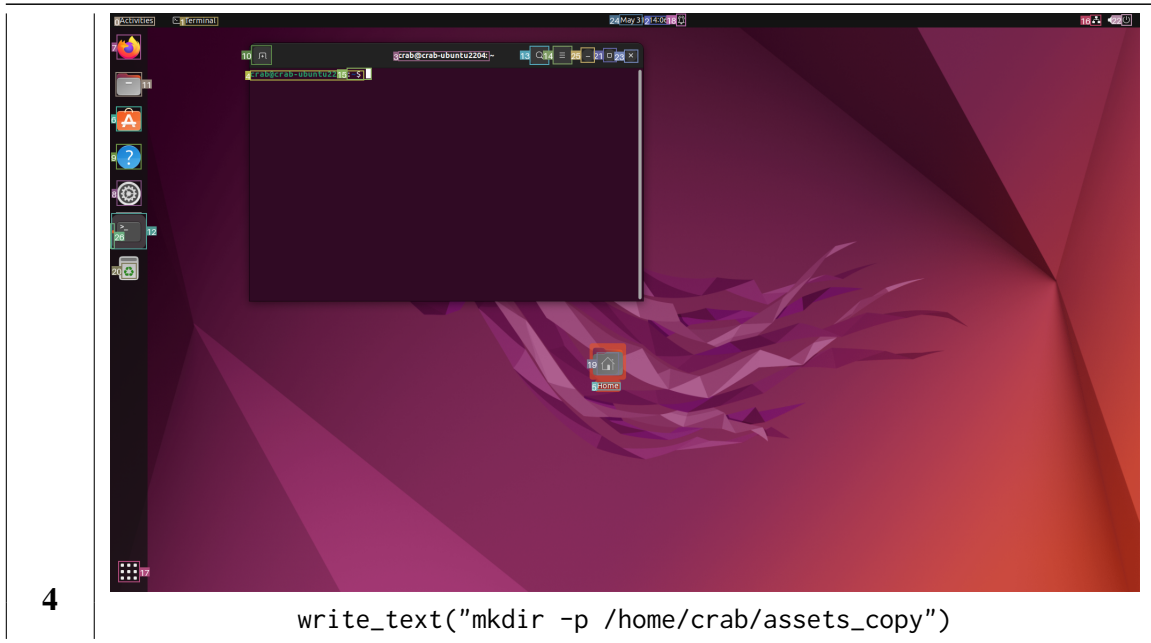
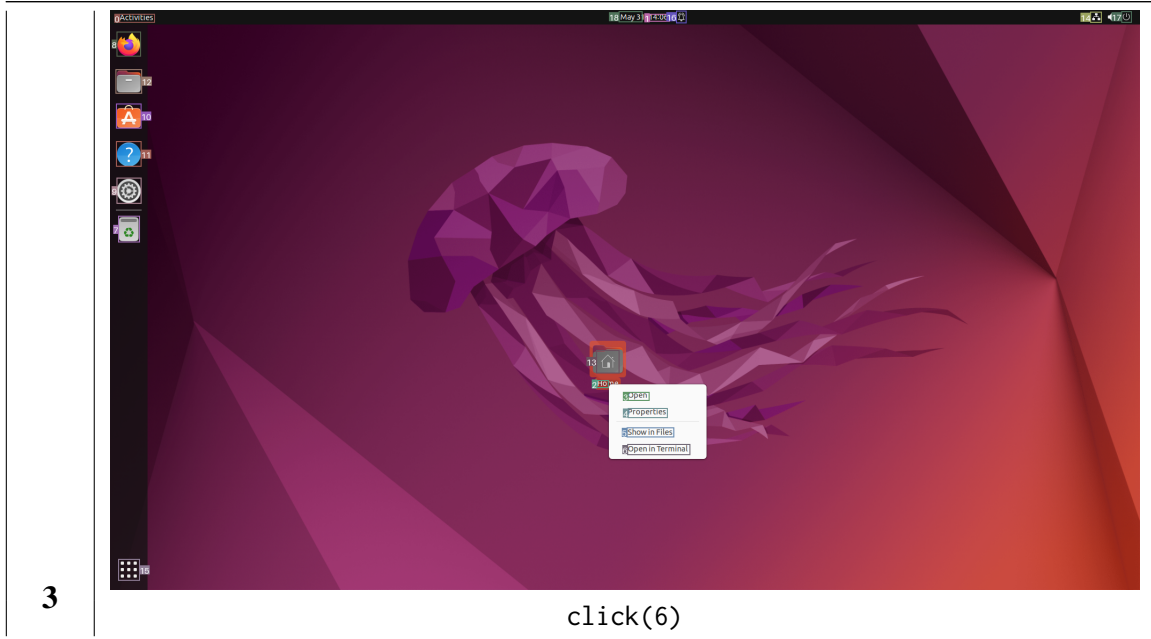
Completed Nodes / Total Nodes: 0 / 2 Termination Reason: Reach Step Limit

Table 12: **Ubuntu task case with GPT-4o (Single):** Create a new directory "/home/crab/assets_copy" and copy all files with the specified "txt" extension from "/home/crab/assets" to the directory "/home/crab/assets_copy".

Step	Agent Observation and Action
0	 <p>click(2)</p>
1-2	 <p>click(2) right_click(2)</p>

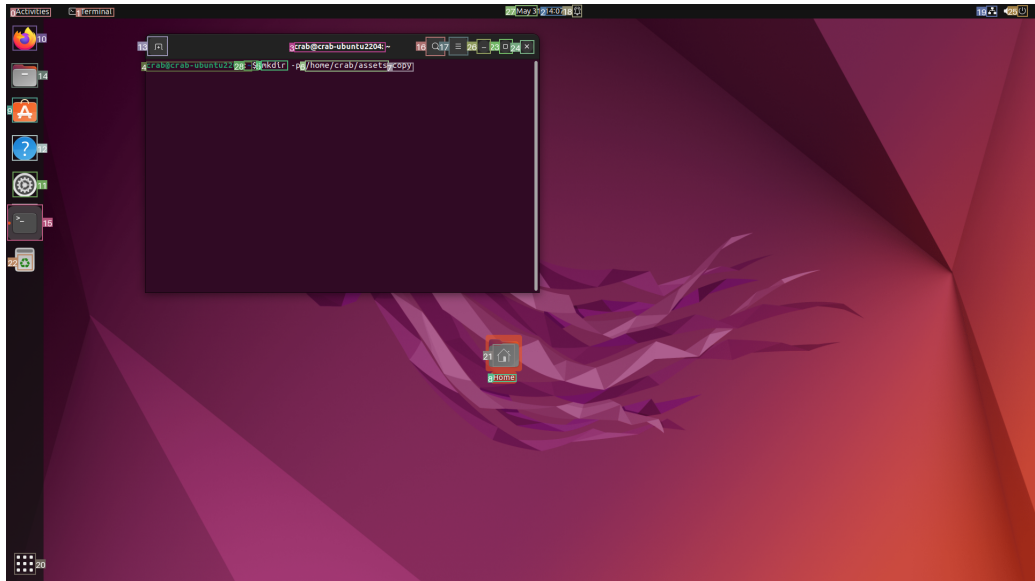
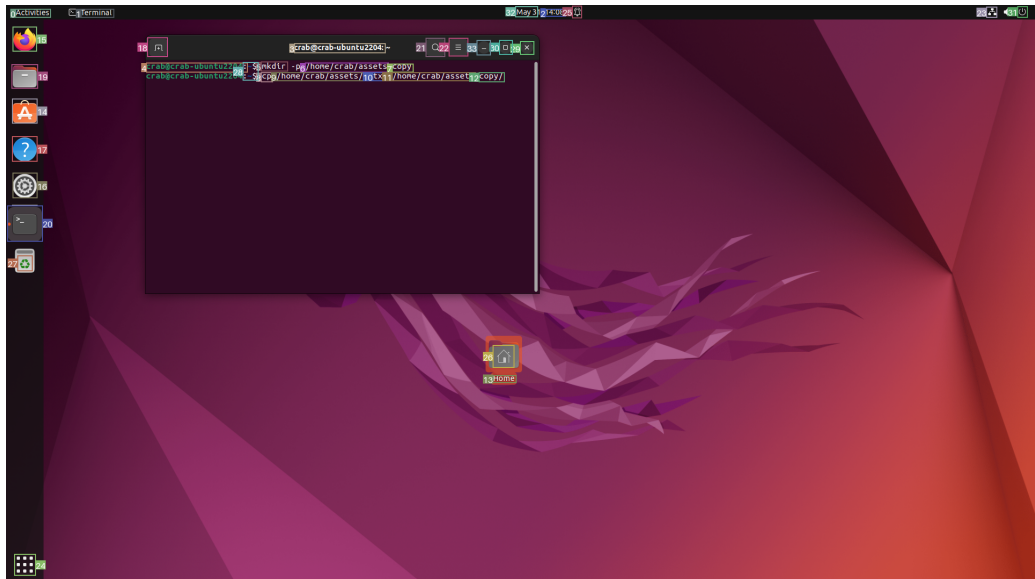
Continued on next page

Table 12 continued from previous page



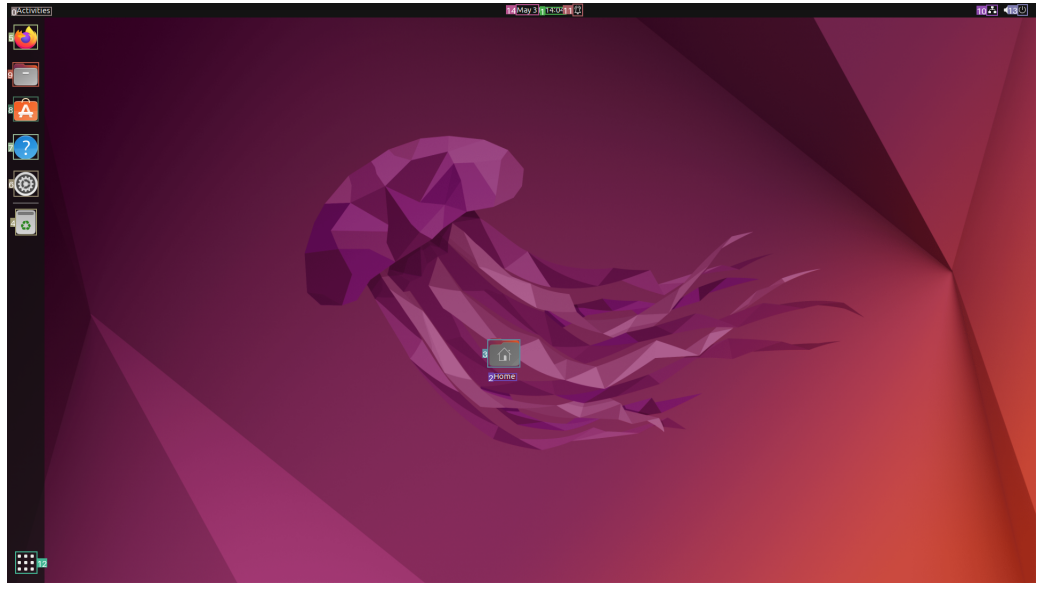
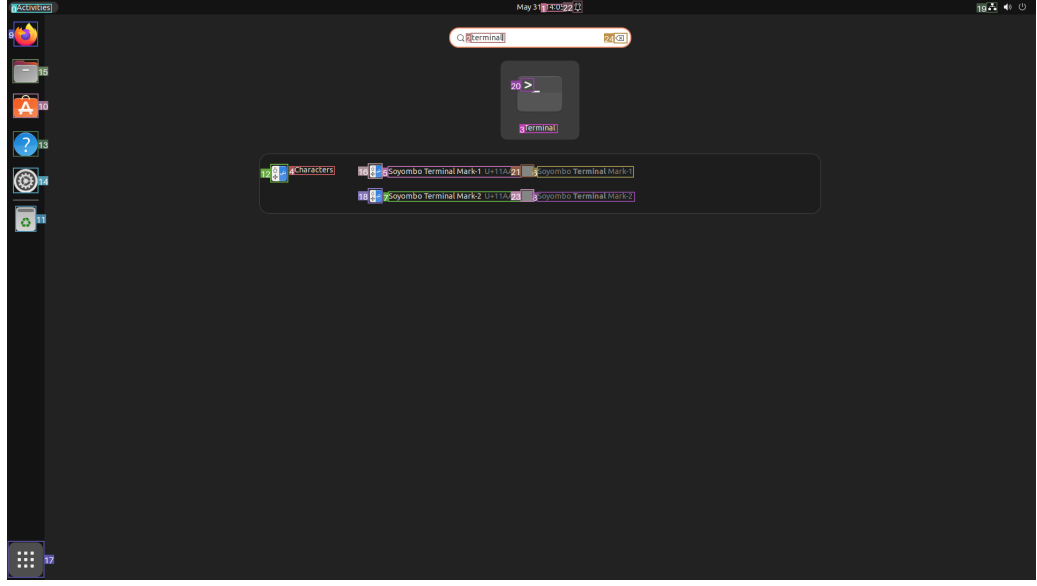
Continued on next page

Table 12 continued from previous page

5-6	
5-6	<pre>key_press("enter") write_text("cp /home/crab/assets/*.txt /home/crab/assets_copy/")</pre>
7	

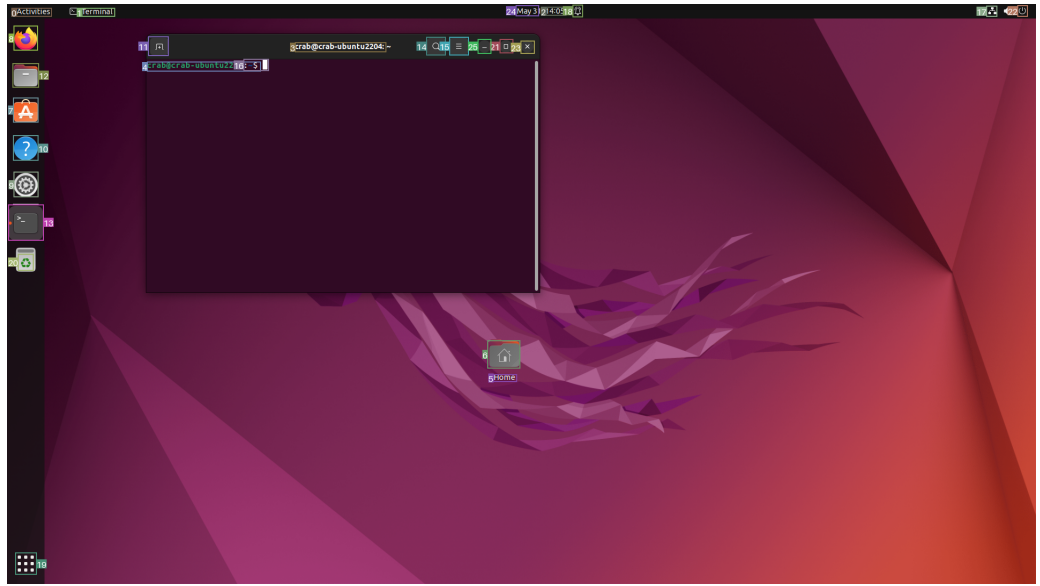
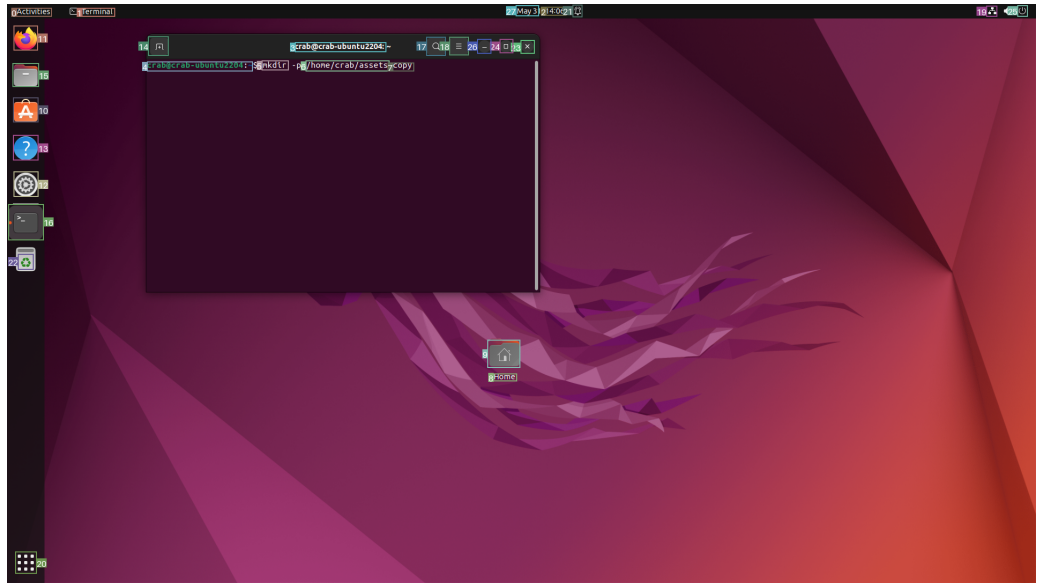
Completed Nodes / Total Nodes: 2 / 2 Termination Reason: Success

Table 13: **Ubuntu task case with GPT-4 Turbo (Single):** Create a new directory `"/home/crab/assets_copy"` and copy all files with the specified `"txt"` extension from `"/home/crab/assets"` to the directory `"/home/crab/assets_copy"`.

Step	Agent Observation and Action
0	 <p data-bbox="619 936 1056 967">search_application("terminal")</p>
1	 <p data-bbox="774 1574 906 1606">click(20)</p>

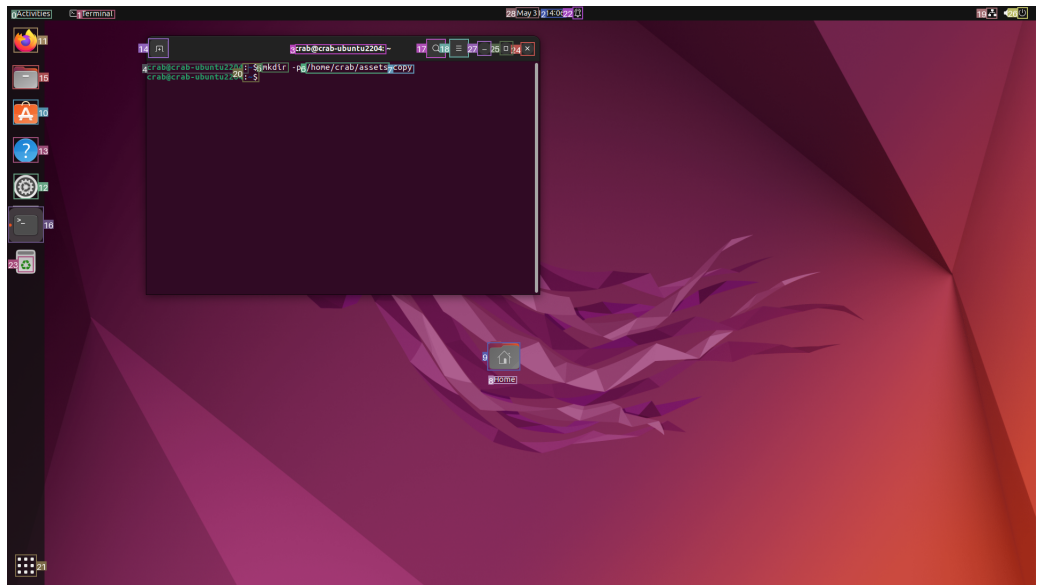
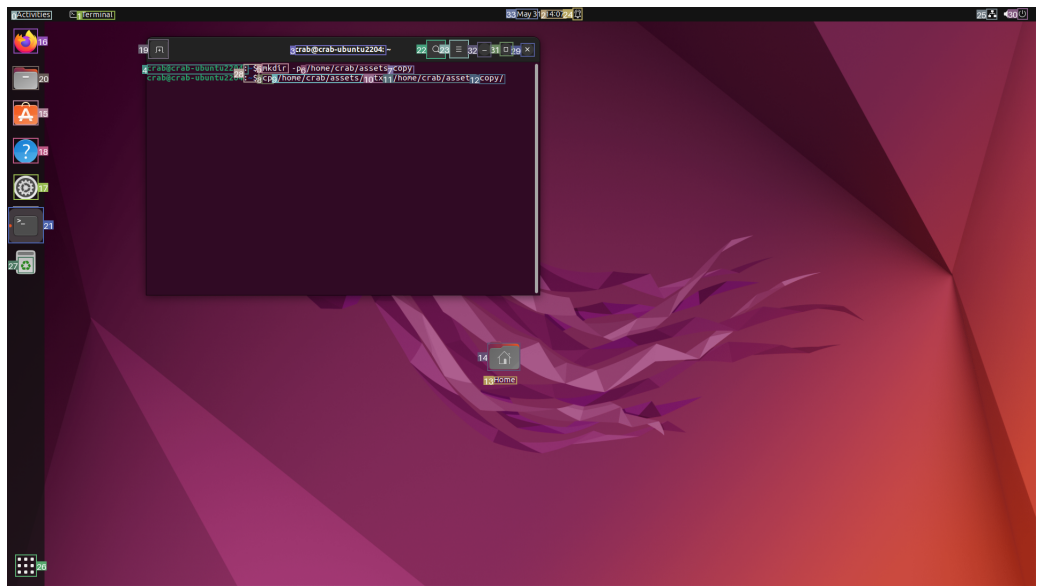
Continued on next page

Table 13 continued from previous page

2	 <p>A terminal window is open on an Ubuntu desktop. The prompt is <code>crab@crab-ubuntu2204:~\$</code>. The command <code>write_text("mkdir -p /home/crab/assets_copy")</code> has been entered and is highlighted in green.</p>
3	 <p>The terminal window shows the previous command has been executed. The prompt is now <code>crab@crab-ubuntu2204:~\$</code>. The command <code>key_press("enter")</code> has been entered and is highlighted in green.</p>

Continued on next page

Table 13 continued from previous page

4	 <p>A terminal window on an Ubuntu desktop. The prompt is <code>@crab@crab-ubuntu2204-</code>. The user has entered <code>cp /home/crab/assets/*.txt /home/crab/assets_copy</code>. The terminal output shows the command being executed. The desktop background is a purple and red geometric pattern. The terminal title bar says "Terminal".</p>
5	 <p>A terminal window on an Ubuntu desktop. The prompt is <code>@crab@crab-ubuntu2204-</code>. The user has entered <code>key_press("enter")</code>. The terminal output shows the command being executed. The desktop background is a purple and red geometric pattern. The terminal title bar says "Terminal".</p>

Completed Nodes / Total Nodes: 2 / 2 Termination Reason: Success